

Language Models Largely Exhibit Human-like Constituent Ordering Preferences

Ada Defne Tur^{††*}, Gaurav Kamath^{††*}, Siva Reddy^{††§}

[†]McGill University, Canada, [‡]Mila - Quebec AI Institute, Canada,

[§]Canada CIFAR AI Chair

*Equal Contribution

Correspondences to ada.tur@mila.quebec

Abstract

Though English sentences are typically inflexible vis-à-vis word order, constituents often show far more variability in ordering. One prominent theory presents the notion that constituent ordering is directly correlated with constituent weight: a measure of the constituent’s length or complexity. Such theories are interesting in the context of natural language processing (NLP), because while recent advances in NLP have led to significant gains in the performance of large language models (LLMs), much remains unclear about how these models process language, and how this compares to human language processing. In particular, the question remains whether LLMs display the same patterns with constituent movement, and may provide insights into existing theories on when and how the shift occurs in human language. We compare a variety of LLMs with diverse properties to evaluate broad LLM performance on four types of constituent movement: heavy NP shift, particle movement, dative alternation, and multiple PPs. Despite performing unexpectedly around particle movement, LLMs generally align with human preferences around constituent ordering.¹

1 Introduction

Despite the fact that word order in English is typically strict, constituents in post-verbal positions can be highly flexible in their ordering (Chomsky, 2002; Wasow and Arnold, 2003). A number of specific phenomena are prime examples of this movement; we show these in Figure 1.

The cause of this movement has been attributed to a variety of factors, including lexical bias (Baars et al., 1975; Hartsuiker et al., 2005; Dell and Reich, 1981), semantic connectedness (Kayne, 1983;

Behaghel, 1932), and information structure (Lambrecht, 1994; Chafe, 1976; Behaghel, 1932; Gundel, 1988). However, one of the most prominent factors, as prior research has suggested, is constituent weight and complexity (Quirk et al., 1975; Behaghel, 1909; Wasow, 1997a).

As the following examples show, we tend to prefer sentences where longer and more complex constituents are moved to the end of the sentence:

- (5) a. I met [the man]_{NP} [at the park]_{PP}.
- b. *I met [at the park]_{PP} [the man]_{NP}.
- c. I met [at the park]_{PP} [the tall man selling water to marathon runners]_{NP}.
- d. ?I met [the tall man selling water to marathon runners]_{NP} [at the park]_{PP}.

The typical constituent order shown in (5-a), for example, is not readily perturbed to the constituent order shown in (5-b). In sentences (5-c) and (5-d), however, where the NP is considerably longer and more complex than the PP, the reverse is true.²

Linguistic theory thus suggests that phrases and constituents are specifically ordered to be presented in increasing complexity, or *weight*; essentially, the larger the constituent, the further to the end of the sentence we expect it to appear (Quirk et al., 1975; Behaghel, 1909; Wasow, 1997a; Futrell et al., 2015). Consider the example in Figure 3:

Other orderings of this sentence, if greatly violating this principle of weight, would likely be considered undesirable. This relationship between the complexity of post-verbal constituents and their ordering raises several questions:

- *What are the exact effects of weight on con-*

²Note that, following common linguistics notation conventions, ‘*’ refers to a sentence judged to be highly unacceptable, while ‘?’ refers to a sentence whose acceptability is questionable.

¹All code and data is available at <https://github.com/McGill-NLP/Constituent-Movement>.

Heavy NP Shift (HNPS)

- (1) a. I met [the tall man selling water to marathon runners]_{NP} [at the park]_{PP}.
b. I met [at the park]_{PP} [the tall man selling water to marathon runners]_{NP}.

Particle Movement (PM)

- (2) a. She looked [up]_{particle} [her question]_{NP} on her computer.
b. She looked [her question]_{NP} [up]_{particle} on her computer.

Dative Alternation (DA)

- (3) a. He sent [her]_{IndirectObj} [a gift]_{DirectObj} for her birthday.
b. He sent [a gift]_{DirectObj} [to her]_{IndirectObj} for her birthday.

Multiple PP Shift (MPP)

- (4) a. I went [to the mall]_{PP} [with my sister]_{PP} on Sunday.
b. I went [with my sister]_{PP} [to the mall]_{PP} on Sunday.

Figure 1: Examples of constituent movement types: Heavy NP Shift (HNPS), Particle Movement (PM), Dative Alternation (DA) and Multiple PP Shift (MPP).

stituent ordering, in terms of gradient and ceiling effects, of increasing complexities on sentence acceptability?

- Which measures of ‘weight/complexity’ best explain the effects on constituent ordering?
- How exactly do LLM preferences around constituent shifting align with human constituent shifting preferences?

Psycholinguistic research has provided some insight into these questions for human language processing, but the same cannot yet be said about increasingly powerful non-human language processors (Medeiros et al., 2021; Wasow, 2002). Prior research, however, supports the abilities of modern language models in assigning relative linguistic plausibility scores aligning similar to human preferences (Linzen et al., 2016; Marvin and

Linzen, 2018). Additionally, we hypothesize that the human-feedback mechanism incorporated in instruction-tuned models will present even more similar judgements.

In this work, we study the behavior of LLMs with constituent movement in English. We model Heavy NP Shift (HNPS), Dative Alternation (DA), Particle Movement (PM), Multiple PP Shift (MPP)—see Figure 1—as a function of its weight. Weight corresponds to a number of selected measures: word length, syllable weight, token length, and modifier weight. We analyze these measures to determine which best explains constituent ordering effects. Figure 2 outlines our contributions, which are as follows:

- We evaluate the preferences of models in regards to constituent movement, using a novel constituent shift dataset containing both synthetic and naturally occurring data.
- We study the motivating factors for constituent ordering preferences in models, over a variety of candidates.
- We compare the behaviors of models with human judgements, analyzing correlative trends between the two.

2 Background

At a high level, English follows a subject-verb-object (SVO) ordering; beyond this basic structure, other objects, modifiers, constituents, and clauses can be added to form more complex sentences (Hengeveld, 1992). The organization and format of how and when each constituent in a sentence is delivered, or its ordering, can be highly flexible (Bakker, 1998; Namboodiripad, 2019, 2017).

Constituent shifting is the process of reordering the constituents of a sentence, such that the original meaning of the sentence is maintained, and all semantic truth conditions are unchanged. This work focuses on four specific types of shift, with a prime commonality: each shift involves the movement of constituents from a post-verbal position (i.e. appearing after the verb of a sentence) to another post-verbal position (Wasow and Arnold, 2003; Wasow, 2002).³ Table 1 demonstrates how we define shifted/unshifted sentences.

³Topicalization, for instance, is not considered, as it can involve movement from a post-verbal position to a pre-verbal position, and vice-versa.

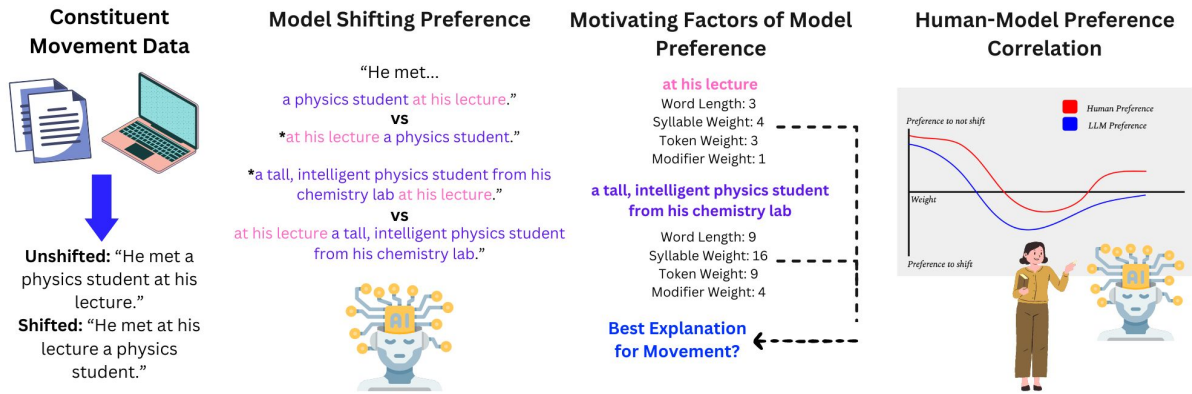


Figure 2: We categorize our work into three main experiments. Our first experiment evaluates model response to constituent movement, our second experiment analyzes what motivates LLM constituent ordering preferences, and our third experiment compares model preferences with human judgements.

	Unshifted Form	Shifted Form
HNPS	S + V + NP + PP <i>I met [the tall man selling water]_{NP} [at the park]_{PP}.</i>	S + V + PP + NP <i>I met [at the park]_{PP} [the tall man selling water]_{NP}.</i>
PM	S + V + NP + PRT <i>She looked [up]_{PRT} [her question]_{NP} on her computer.</i>	S + V + PRT + NP <i>She looked [her question]_{NP} [up]_{PRT} on her computer.</i>
DA	S + V + NP1 + NP2 <i>He sent [her]_{NP} [a gift]_{NP} for her birthday.</i>	S + V + NP1 + PP <i>He sent [a gift]_{NP} [to her]_{PP} for her birthday.</i>
MPP	S + V + PP1 + PP2 <i>I went [to the mall]_{PP} [with my sister]_{PP} on Sunday.</i>	S + V + PP2 + PP1 <i>I went [with my sister]_{PP} [to the mall]_{PP} on Sunday.</i>

Table 1: Unshifted versus shifted forms for each shift. Note that the unshifted and shifted form for MPP is ambiguous; the unshifted and shifted forms cannot be derived separately given an example.

"In my laboratory we use it as an easily studied instance of mental grammar, allowing us to document [in great detail] [the psychology of linguistic rules] [from infancy to old age] [in both normal and neurologically impaired people] [in much the same way that biologists focus on the fruit fly *Drosophila* to study the machinery of the genes]."

Figure 3: Example from (Pinker, 2007).

In our experiments, we consider the following measures of weight: **word length**, which corresponds to the number of words in a constituent; **syllable weight**, which corresponds to the number of syllables in a constituent; **token length**, which refers to the number of tokens in a constituent, determined by respective model tokenizers; **modifier weight**, which refers to the number of adjective phrase (AdjP) and prepositional phrase (PP) modifying the constituent itself (plus 1 to include the weight of the base constituent itself as normalization).

Absolute weight itself, however, is not the most effective metric for observing motivations for con-

stituent movement; instead, there is reason to believe that it is the *relative* weight of constituents that determines their ordering (Wasow, 1997b). For instance, for phrases "with her grandmother" and "around the garden", the ratio of word lengths would be 1 (3:3); the ratio of syllables would be 1 (5:5), etc. For phrases "with her grandmother" and "around the decorated entryway garden with the large fountain", the ratio of word length would be $\frac{1}{3}$ (3:9); the ratio of syllables would be $\frac{5}{17}$ (5:17), etc. As the ratio of any metric increases beyond 1, we know that the weight of the first constituent is larger than the second, and thus, we predict that the motivation to shift will be greater.

3 Related Work

Constituent movement has been the subject of considerable linguistic literature; we categorize relevant contributions by our three research questions.

3.1 What are the exact effects of weight on constituent ordering?

Significant prior work has focused on investigating the effects of weight on constituent ordering (Arnold et al., 2000; Wasow and Arnold, 2003; Wasow, 1997a; Arnold et al., 2004; Hawkins, 1995; Behaghel, 1909; Hawkins, 2004); many contributions find gradient effects by which the shift becomes more frequent in examined language corpora as the relative weight of the relative constituent changes, suggesting that weight is the predominant factor in triggering the shift, even cross-linguistically (Wasow, 1997b; Wasow and Arnold, 2003; Faghiri and Thuilier, 2018; Wang and Liu, 2014; Hawkins, 1999; Quirk et al., 1975; Manetta, 2012). Furthermore, studies with human participants find similar results, where weight presents a primary role in the shift (Medeiros et al., 2021). The study, however, also suggests that this movement is constrained by ceiling effects, by which the efficacy of additional weight and complexity plateaus. More closely relevant to this work, Futrell and Levy (2018) conduct a similar analysis on post-verbal constituent movement using weight as a binary feature (i.e. ‘long’ vs ‘short’), and find similar trends with LSTMs.

3.2 What measure of weight best explains effects of constituent ordering?

Weight is a measure of a constituent’s complexity or size, but how best to measure it is less straightforward (Chomsky, 2008; Haegeman, 1991; Wasow, 1997a). Related research categorizes and analyzes the effects of three primary measures of weight on constituent movement, particularly with HNPS: the word length of the NP, the number of nodes in the NP’s syntactic structure, and the number of modifiers applied to the NP (Wasow and Arnold, 2003; Medeiros et al., 2021; Wasow, 1997b). The analyses found that, although the word length was statistically the strongest predictor for HNPS, “no single factor can account for observed constituent order alternation” (Medeiros et al., 2021, pg.6). Similarly, Wasow and Arnold (2003) find in a corpus study that for HNPS and DA, constituent movement was best accounted for when considering both word length and modifier weight together, as opposed to either on its own.

3.3 How exactly do LLM preferences around constituent shifting align with human constituent shifting preferences?

Research concerning the behaviors of computational models has also shown that models exhibit human-like preferences (Fujihara et al., 2022; Linzen et al., 2016; Marvin and Linzen, 2018; Kamath et al., 2024). Notably, prior work shows models learn syntactic alternations (Wilcox et al., 2019; Lau et al., 2017); more directly relevant to us, Futrell and Levy (2018) finds that behaviors of LSTMs appear to correlate closely with observed judgements of humans on corresponding data, suggesting that constituent movement is motivated similarly in both humans and models. This work, however, predates preference-aligned models (Ouyang et al., 2022); we hypothesize that the behaviors of such models will align even more closely with human preferences around constituent ordering.

4 Models and Data

4.1 Data

We both generate synthetic data using a template modified for the various shifts and structures and mine natural data from the Penn Treebank-2 corpus (Marcus et al., 1995). Each shift we consider has a standardized form that we can utilize for both processes, noted in Table 1. We also annotate for the aforementioned weight measures; syllable weight is computed using Syllapy;⁴ token weight is retrieved simultaneously with the model scoring process; modifier weight is counted when constructing modifier chunks on constituents.

4.1.1 Synthetic Data

We synthetically generate data using the process shown in Figure 4 in order to accumulate large amounts of iteratively more complex data for model evaluations. Using an overall frame for the sentence, we alternate subjects, verbs, moving constituents, and their modifiers; we do this for all variables on a variety of constituents.

4.1.2 Mined Data

Our synthetic data, however, contains limited syntactic variation, and may not represent naturally occurring data; thus, we mine from existing data to ground our results. We use the Penn Treebank-2 to retrieve sentences following the structure of each

⁴<https://github.com/mholtzscher/syllapy>

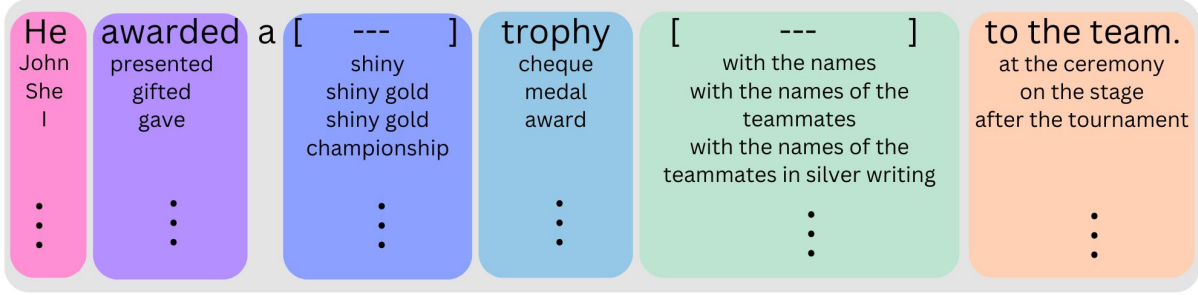


Figure 4: The outline for creating synthetic data, using varying modifier weights.

	Synthetic Data	Mined Data
HNPS	3,888	314
PM	4,136	131
DA	210,304	123
MPP	180,224	130
Total Size	398,552	698

Table 2: Dataset sizes by sentence count. Synthetic datasets for DA and MPP are larger due to modification of both constituents, rather than one, leading to more constituent weight ratios.

shift (Marcus et al., 1995). We collect over 1,000 examples for HNPS and DA, and approximately 400 for PM and MPP, selecting a random sample of 500 and 400 sentences, respectively. We manually inspect and exclude low-quality datapoints.⁵ As weight measures, we include constituent word length, syllable weight, and token length, but exclude modifier weight due to the complexity of accurately extracting this from such data.

4.2 Models

We select a range of open autoregressive models with diverse attributes: the entire GPT-2 model family (Radford et al., 2019) was analyzed to study behaviors over scaled model sizes; Llama-3 8B, Llama-3 8B Instruct (AI@Meta, 2024), Mistral v0.3 7B, Mistral v0.3 7B Instruct (Jiang et al., 2023), OLMo 7B, and OLMo 7B Instruct (Groeneveld et al., 2024) were used to compare standard and instruction-tuned models; BabyFlamingo and BabyOPT (Warstadt et al., 2023) were used to study LLMs trained on BabyLM data, a child-directed dataset to simulate language stimulus during the early human language acquisition period. We do not inspect closed-source models such as

⁵Low-quality data includes nonsensical or malformed sentences, unnecessary repetitions, or incorrect modifier structures.

GPT-4 (OpenAI, 2024), due to inaccessibility of their underlying logits.

5 Shifting Preference of Models

5.1 Approach

To observe the effects of constituent weight on ordering preferences, we compute the difference in log probabilities assigned by models to shifted and unshifted sentences at a range of constituent weight ratios (see Section 2). We begin by extracting log probabilities assigned to each sentence by the aforementioned models, using the `minicons` library (Misra, 2022). This score corresponds to the model’s judgement of the sentence’s linguistic plausibility, computed with the following equation:

$$M_{score}(\mathbf{w}) = \sum_{t=1}^T \log P_M(w_t | w_1, w_2, \dots, w_{t-1}; \theta).$$

M_{score} is the log probability score of the sequence $\mathbf{w} = [w_1, w_2, \dots, w_t, \dots, w_T]^T$, where w_t is the token at position t . The term P_M is the conditional probability from the model M of token w_t given the preceding tokens, while θ are the model parameters. The output of this formula is the sum of the probabilities of all tokens in the sequence, given previous tokens and model parameters, which equates to the overall sequence probability. The closer M_{score} is to 0, the more strongly the model judges the sequence to likely occur in human language.

Upon computing this score for each minimal pair of shifted-unshifted sentences (denoted as U and S), we calculate the difference of the M_{score} for each sentence:

$$M_{preference} = M_{score}(U) - M_{score}(S).$$

This metric aligns closely with the surprisal-based measure used by Futrell and Levy (2018),




Figure 5: OLMo 7B $M_{preference}$ scores with respect to different measures of weight.

and is in line with other surprisal-based metrics commonly used in work in psycholinguistics and computational linguistics (Linzen et al., 2016; Futrell et al., 2018; Wilcox et al., 2018; Schuster and Linzen, 2022; Baroni, 2022).

Intuitively, the value of $M_{preference}$ captures the model’s relative preference for the unshifted version of the sentence. If this value is >0 , the model has a stronger preference for the unshifted sentence, and if it is <0 , the model has a stronger preference for the shifted version of the sentence; values approaching 0 suggest no clear preference between the two.

Though the usage of instruction-tuned models attracts the idea of directly prompting such LMs for their linguistic preferences, Hu and Levy (2023) and Kamath et al. (2024) show results that suggests meta-linguistic prompting underestimates linguistic capacities; thus, we only consider raw probability scores in our experimentation.

5.2 Results: Are models motivated by weight to shift?

We analyze model $M_{preference}$ scores by weight ratios (see Section 2) considering various metric types in the analysis. Figure 5 shows OLMo 7B performance on HNPS; see Section A.2.1 for remaining models and variables. We observe similar trends overall across models.

For HNPS, we find that model $M_{preference}$ scores initially start positive, i.e., models prefer unshifted sentence, but converge above 0 as constituent weight ratios increase, indicate some continued preference for unshifted constituent orderings beyond a given point. In the case of DA and MPP, we see a similar effect of weight $M_{preference}$ scores, but with a plateau *below* zero, indicating

that models eventually converge upon a relative preference for the shifted version of a sentence. In the case of PM, however, we see scores initially drop sharply from around 0, but later somewhat rise.

6 Motivating Factors of Model Preference

Literature regarding human language finds varying levels of importance correlated with different measures of constituent weight, particularly in corpus studies and human judgement tasks (Wasow, 1997a; Wasow and Arnold, 2003; Medeiros et al., 2021). Whether or not this same trend is seen with language models is unknown. We conduct a regression analysis on the data collected for the first experiment, using a generalized additive mixed model (Wood, 2017; Sóskuthy, 2017), with the goal of measuring how significantly each weight measure both impacts the models’ judgements and serves to fit regression lines on the data.

6.1 Approach

To compare how well different measures of weight explain shifting preferences, we fit Generalized Additive Mixed Models, or GAMMs (Wood, 2017; Sóskuthy, 2017), on our $M_{preference}$ scores (see Section 5.1) as a function of various weight measures: word length, token length, syllable weight, and modifier weight. GAMMs allow for the fitting of highly non-linear relationships as the sum of multiple predictor-wise smooth functions: basis functions that allow for an arbitrary degree of smoothness. Crucially, aside from providing interpretable measures of goodness of fit, GAMMs also allow for grouping structures in the data to be captured as random effects (Wood, 2017, ch.6).

Bearing in mind that multiple measures of weight may jointly determine the accessibility of a shift (Wasow, 1997b; Wasow and Arnold, 2003), we analyze the relative importance of each weight measure in the following manner. For each model, first, we fit a GAMM on the model’s $M_{preference}$ scores as a function of all weight predictors, with verb-wise random intercepts and slopes. We then iteratively ablate each weight predictor while retaining all others, and compare the quality of fit yielded by the full model with that of the ablated model. Intuitively, this provides an indication of how important the dropped predictor is for the LLM: it captures how much less of the LLM’s be-

Model	Var	Full	Token _{ab}	Word _{ab}	Syll _{ab}	Mods _{ab}
GPT-2 Medium		.654	.629	.616	.540	.635
Llama-3	HNPS	.542	.524	.519	.485	.538
Llama-3 Instruct		.452	.438	.408	<u>.359</u>	.451
BabyLlama		.527	.514	.466	.415	.509
GPT-2 Medium		.605	.602	.581	<u>.534</u>	.580
Llama-3	PM	.608	.603	.586	.564	.590
Llama-3 Instruct		.627	.619	.599	.555	.602
BabyLlama		.719	.702	.663	.651	.701
GPT-2 Medium		.571	.561	.552	.568	.568
Llama-3	DA	.554	.544	.532	.538	.549
Llama-3 Instruct		.503	.493	.490	<u>.486</u>	.496
BabyLlama		.630	.616	.603	.622	.623
GPT-2 Medium		.368	.309	.302	.300	.356
Llama-3	MPP	.358	.316	.321	.271	.351
Llama-3 Instruct		.320	.284	.291	<u>.208</u>	.313
BabyLlama		.310	.297	.298	.281	.306

Table 3: R-squared scores for a subset of models; see full table in Appendix A.1. *Full* denotes R-squared values from GAMMs with all weight measure predictors included, while $[Metric]_{ab}$ denotes the R-squared score when excluding the metric; a larger difference compared to the original R-squared score denotes higher significance in contributing to the overall R-squared. Numbers bolded denote high R-sq.; numbers underlined denote low R-sq.

haviour is explained when information about that given measure of weight is ignored.

6.2 Results: Which measure of weight best explains shifting?

Table 3 presents a subset of the results of our regression analysis (full results in Appendix Table 5). Crucially, we find that **syllable weight** is often the most important predictor of LLM behavior around constituent ordering preference, since the drop in R-squared scores is highest when syllable weight is not used. For DA, word length seems to be the best predictor.

We find that GPT-2 Medium achieves the highest R-squared overall on both HNPS and MPP. Further, contrary to our hypotheses, across almost all shift types, instruction-tuned models consistently achieve lower R-squared scores than their base model counterparts. Table 3, for example, shows these results for Llama-3 and Llama-3 Instruct; see Appendix A.2.1 for remaining plots on synthetic data. The BabyLM models also present high R-squared scores on both PM and DA, with BabyLlama yielding the highest R-squared values (see Table 3 for BabyLlama results and Appendix A.2.1 for BabyOPT results). Finally, despite the high performance of GPT-2 Medium, we do not observe consistent improvements in R-squared values as model sizes scale.




Figure 6: Scatterplot of human judgment scores and OLMo 7B $M_{preference}$ scores on Heavy NP Shift data.

7 Human-Model Preference Correlation

7.1 Approach

To adequately compare the behaviors of LLMs with those of humans, a direct study of preferences on identical data points is necessary. We collect human judgements on a subset of data presented to models; though human judgements and model scores are not identical metrics, they can act as proxies when comparing relative trends.

We conduct a crowdsourced study through Prolific, collecting judgments from 126 native English speakers residing in Anglophone countries, on 500 sentence pairs. Each participant is presented 25 sentence pairs and asked to judge how natural they sound in relation to each other, assigning a score between 1 and 7; 1 corresponds to the first sentence presented appearing far more natural than the second, and 7 the reverse. We exclude datapoints with low inter-annotator agreement to minimize noise.⁶

7.2 Results: Do LLM and human preferences correlate?

To analyze how this human judgment data compares with model results, we compute the Spearman correlation between the average human score for each data point and the model’s $M_{preference}$ score. We present these scores in Table 4. We also plot model scores against human judgment data to observe the correlation visually; these plots are presented in Figure 6 and Section A.3.

Comparing preferences of humans and models on a statistical level introduces interesting find-

⁶Datapoints with standard deviation of participation responses >1.5 , as well as those from participants who consistently failed attention checks, were excluded.

	GPT-2	GPT-2 Med	GPT-2 Large	GPT-2 XL	BabyOPT	BabyLlama	Llama-3	Llama-3 I	Mistral v0.3	Mistral v0.3 I	OLMo	OLMo I
HNPS	.410	.440	.342	.390	.261	.420	.428	.386	.414	.344	.509	.431
PM	.212	.125	.213	.295	.196	.261	.305	.293	.315	.347	.430	.431
DA	.651	.565	.505	.541	.256	.391	.600	.511	.524	.449	.494	.478
MPP	.371	.395	.361	.579	.233	.357	.402	.222	.487	.412	.513	.263

Table 4: Spearman correlation values (absolute) between LLMs and humans on sentence judgement tasks.

ings. Notably, the GPT-2 and OLMo 7B classes of models appear to align most closely with human behaviors, achieving the highest correlation scores across all shift types.

Remarkably, most of the more highly correlated models are the base versions of their respective model category—the instruction-tuned models quite often performed in manners less correlated with human preferences than their base-model counterparts. This contrast is most stark in the case of MPP. Interestingly, no clear trend in alignment can be seen in the GPT-2 family as model sizes scale on any of the judgement tasks. We present remaining correlation figures in Section A.3. Crucially, correlations between model scores and human judgments are particularly low on the PM data.

8 Discussion

8.1 What are the exact effects of weight on constituent ordering?

In Figure 5 and Appendix A.2.1, we observe converging effects as weight increases, suggesting that prior theories defining weight as a prime factor in motivating the shift hold with computational models as well. Further, we observe, specifically with PM, that weight, beyond a certain threshold, begins to detriment motivation for shifting.

We observe similar effects on the scraped corpus data, in Appendix A.2.2, though with more noise. Given the specificity of the data itself, being rooted in financial reports and news articles, some noise and outliers were expected, and in some cases, observable trends remain. Similar to the synthetic data plots, we see convergence on HNPS, and some on MPP, as well as an initial drop followed by a rise in $M_{preference}$ scores for PM.

8.2 What measure of weight best explains effects of constituent ordering?

The syllable weight was the best measure of weight for explaining motivations to shift in LLMs. This raises an obvious question—why is syllable weight a more effective predictor of

model behavior than token weight, which would intuitively be most aligned with a model’s processing of weight and complexity? This finding acts as initial evidence that models may induce linguistic information not just at the token level, but also implicitly at the level of syllables.⁷

8.3 How exactly do LLM preferences around constituent shifting align with human constituent shifting preferences?

Broadly speaking, a clear trend is maintained across humans and models, following what was presented by Futrell and Levy (2018) in their analysis. Where human language sees motivation for movement with increasing weight, model behavior follows closely. Our experiment, which includes graded data beyond binary weight categorizations, and a wider range of models, yields relatively high correlation effects between preferences of models and humans, as presented in Table 4—suggesting noticeably similar behaviors between the two on particular linguistic tendencies, with the notable exception of particle movement.

Interestingly, we observe an unexpected trend where instruction-tuned models, which consistently correlate less with human data than their corresponding base model, as well as, quite often, yield lower R-squared scores. This runs against our initial hypothesis around instruction-tuned models, and suggests inadequacy in providing consistent and explainable trends compared to base models.

8.4 Future Work

Our findings suggest that even though newer models are equipped with more parameters, training data, and the human-feedback mechanism, they fail to align better with human linguistic preferences than their earlier counterparts, raising ques-

⁷Admittedly, however, this evidence is not conclusive proof of such ‘knowledge’; it is also possible that syllabic information fits the data best simply because of how much it varies—syllabic information is more likely to be a signal for specific words and phrases, due to the fact that while any two tokens could have the same ratio of tokens between them, they may vary in their syllable weights.

tions for future study. Equally, it invites further research into how models generate such sentences in standard conversational usage. We also wish to study further the motivations of LLMs in constituent movement, primarily regarding analysis of theories suggesting disparities between ordering preference motivations of listeners and speakers (Wasow, 2002).

9 Conclusions

In this work, we present a thorough analysis of the behaviors of LLMs in response to constituent movement, using both a novel set of nearly 400K minimal pairs of variably ordered sentences, as well as naturally occurring data. We collect human judgements and model preference scores and observe comparable trends between the behaviors of humans and models. Such comparisons indicate that humans and LLMs largely hold similar linguistic preferences around constituent ordering, with the exception of particle movement. Our findings—and in particular, the surprising gap we find between instruction-tuned models and their vanilla counterparts—invite further research into when and how linguistic preferences of models and humans align.

10 Limitations

We only focus on constituent movement in English, even though this phenomena is known to manifest cross-linguistically (Faghiri and Thuilier, 2018; Wang and Liu, 2014; Hawkins, 1999; Quirk et al., 1975; Manetta, 2012; Fujihara et al., 2022).

11 Ethics Statement

Our experimentation poses no risks or harms for any participants involved. Human participants from our data study through Prolific were compensated on average US\$12 per hour. Anonymous participants were informed of the purpose of the study and how their responses would be used, as well as their rights regarding submitted data.

12 Acknowledgements

This work was partly funded by IVADO R3AI NLP Regroupement and a Doctoral Training Award from the *Fonds de Recherche du Québec—Société et Culture*. We also thank Morgan Sonderegger for his invaluable feedback and guidance.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Jennifer E. Arnold, Anthony Losongco, Thomas Wasow, and Ryan Ginstrom. 2000. [Heaviness vs. newness: The effects of structural complexity and discourse status on constituent ordering](#). *Language*, 76(1):28–55.
- Jennifer E. Arnold, Thomas Wasow, Ash Asudeh, and Peter Alrenga. 2004. [Avoiding attachment ambiguities: The role of constituent ordering](#). *Journal of Memory and Language*, 51(1):55–70.
- Bernard J Baars, Michael T Motley, and Donald G MacKay. 1975. Output editing for lexical status in artificially elicited slips of the tongue. *Journal of verbal learning and verbal behavior*, 14(4):382–391.
- Dik Bakker. 1998. Flexibility and consistency in word order patterns in the languages of europe. *Empirical Approaches to Language Typology*, 20:383–420.
- Marco Baroni. 2022. [On the proper role of linguistically-oriented deep net analysis in linguistic theorizing](#). *Preprint*, arXiv:2106.08694.
- Otto Behaghel. 1909. [Beziehungen zwischen umfang und reihenfolge von satzgliedern](#). *Indogermanische Forschungen (1909)*, 25(1909):110–142.
- Otto Behaghel. 1932. *Deutsche Syntax: Eine geschichtliche Darstellung. Band IV: Wortstellung, Periodenbau*. Carl Winters Universitätsbuchhandlung, Heidelberg.
- Wallace L. Chafe. 1976. Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In Charles N. Li, editor, *Subject and topic: A new typology of language*, pages 25–55. Academic Press, New York, NY.
- Noam Chomsky. 2002. *Syntactic Structures*. A Mouton classic. Mouton de Gruyter.
- Noam Chomsky. 2008. [The logical structure of linguistic theory](#). *Language*, 84:795–814.
- Gary S Dell and Peter A Reich. 1981. Stages in sentence production: An analysis of speech error data. *Journal of verbal learning and verbal behavior*, 20(6):611–629.
- Pegah Faghiri and Juliette Thuilier. 2018. [Relative weight and givenness in constituent ordering of typologically different languages: Evidence from French and Persian](#). The 31st Annual CUNY Conference on Human Sentence Processing. Poster.
- Riki Fujihara, Tatsuki Kuribayashi, Kaori Abe, Ryoko Tokuhisa, and Kentaro Inui. 2022. [Topicalization in language models: A case study on japanese](#). In *International Conference on Computational Linguistics*.

- Richard Futrell and Roger P. Levy. 2018. [Do rnns learn human-like abstract word order preferences?](#) *Preprint*, arXiv:1811.01866.
- Richard Futrell, Kyle Mahowald, and Edward Gibson. 2015. [Large-scale evidence of dependency length minimization in 37 languages.](#) *Proceedings of the National Academy of Sciences*, 112(33):10336–10341.
- Richard Futrell, Ethan Wilcox, Takashi Morita, and Roger Levy. 2018. [Rnns as psycholinguistic subjects: Syntactic state and grammatical dependency.](#) *Preprint*, arXiv:1809.01329.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Olmo: Accelerating the science of language models.](#) *Preprint*, arXiv:2402.00838.
- Jeannette K. Gundel. 1988. Universals of topic-comment structure. In Michael Hammond, Edith A. Moravcsik, and Jessica R. Wirth, editors, *Studies in Syntactic Typology*, pages 209–239. Benjamins, Amsterdam.
- Liliane Haegeman. 1991. *Introduction to Government and Binding Theory*. Blackwell, Oxford.
- Robert J Hartsuiker, Martin Corley, and Heike Martensen. 2005. The lexical bias effect is modulated by context, but the standard monitoring account doesn't fly: Related reply to baars et al.(1975). *Journal of Memory and Language*, 52(1):58–70.
- John A. Hawkins. 1995. *A Performance Theory of Order and Constituency*. Cambridge Studies in Linguistics. Cambridge University Press.
- John A. Hawkins. 1999. [The relative order of prepositional phrases in english: Going beyond manner–place–time.](#) *Language Variation and Change*, 11(3):231–266.
- John A. Hawkins. 2004. *Efficiency and Complexity in Grammars*. Oxford University Press.
- Kees Hengeveld. 1992. [Non-verbal predication: Theory, typology, diachrony.](#)
- Jennifer Hu and Roger Levy. 2023. [Prompting is not a substitute for probability measurements in large language models.](#) *Preprint*, arXiv:2305.13264.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b.](#) *Preprint*, arXiv:2310.06825.
- Gaurav Kamath, Sebastian Schuster, Sowmya Vajjala, and Siva Reddy. 2024. [Scope Ambiguities in Large Language Models.](#) *Transactions of the Association for Computational Linguistics*, 12:738–754.
- Richard S. Kayne. 1983. [Connectedness.](#) *Linguistic Inquiry*, 14(2):223–249.
- Knud Lambrecht. 1994. *Information Structure and Sentence Form: Topic, Focus, and the Mental Representations of Discourse Referents*. Cambridge Studies in Linguistics. Cambridge University Press.
- J. H. Lau, A. Clark, and S. Lappin. 2017. [Grammaticality, acceptability, and probability: A probabilistic view of linguistic knowledge.](#) *Cognitive Science*, 41(5):1202–1241.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of lstms to learn syntax-sensitive dependencies.](#) *Preprint*, arXiv:1611.01368.
- Emily Walker Manetta. 2012. [Reconsidering rightward scrambling: Postverbal constituents in hindi-urdu.](#) *Linguistic Inquiry*, 43:43–74.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1995. [Penn treebank ii 2 - ldc95t7.](#)
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- David J. Medeiros, Paul Mains, and Kevin B. McGowan. 2021. [Ceiling effects on weight in heavy np shift.](#) *Linguistic Inquiry*, 52(2):426–440.
- Kanishka Misra. 2022. [minicons: Enabling flexible behavioral and representational analyses of transformer language models.](#) *arXiv preprint arXiv:2203.13112*.
- Savithry Nambodiripad. 2017. *An experimental approach to variation and variability in constituent order*. University of California, San Diego.
- Savithry Nambodiripad. 2019. [A gradient approach to flexible constituent order.](#)
- OpenAI. 2024. [Gpt-4 technical report.](#) *Preprint*, arXiv:2303.08774.

- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Steven Pinker. 2007. *The Language Instinct*. Harper-Collins.
- Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. 1975. [Review of: A grammar of contemporary english, by randolph quirk, sidney greenbaum, geoffrey leech and jan svartvik](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Sebastian Schuster and Tal Linzen. 2022. When a sentence does not introduce a discourse entity, transformer-based models still sometimes refer to it. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 969–982.
- Márton Sóskuthy. 2017. [Generalised additive mixed models for dynamic analysis in linguistics: a practical introduction](#). *Preprint*, arXiv:1703.05339.
- Hua Wang and Haitao Liu. 2014. [The effects of length and complexity on constituent ordering in written english](#). *Poznan Studies in Contemporary Linguistics*, 50(4):477–494.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, and Ryan Cotterell. 2023. [Findings of the BabyLM challenge: Sample-efficient pretraining on developmentally plausible corpora](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 1–34, Singapore. Association for Computational Linguistics.
- Thomas Wasow. 1997a. [End-weight from the speaker’s perspective](#). *Journal of Psycholinguistic Research*, 26(3):347–361.
- Thomas Wasow. 1997b. [Remarks on grammatical weight](#). *Language Variation and Change*, 9(1):81–105.
- Thomas Wasow. 2002. [Postverbal behavior](#). *Language*, 80(2):327–331.
- Thomas Wasow and Jennifer Arnold. 2003. [Postverbal constituent ordering in English](#), page 119–154. De Gruyter Mouton, Berlin, New York.
- Ethan Wilcox, Roger Levy, and Richard Futrell. 2019. [What syntactic structures block dependencies in rnn language models?](#) *Preprint*, arXiv:1905.10431.
- Ethan Wilcox, Roger Levy, Takashi Morita, and Richard Futrell. 2018. [What do RNN language models learn about filler–gap dependencies?](#) In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 211–221, Brussels, Belgium. Association for Computational Linguistics.
- Simon N. Wood. 2017. *Generalized Additive Models: An Introduction with R, Second Edition*, 2 edition. Chapman and Hall/CRC.

A Appendix

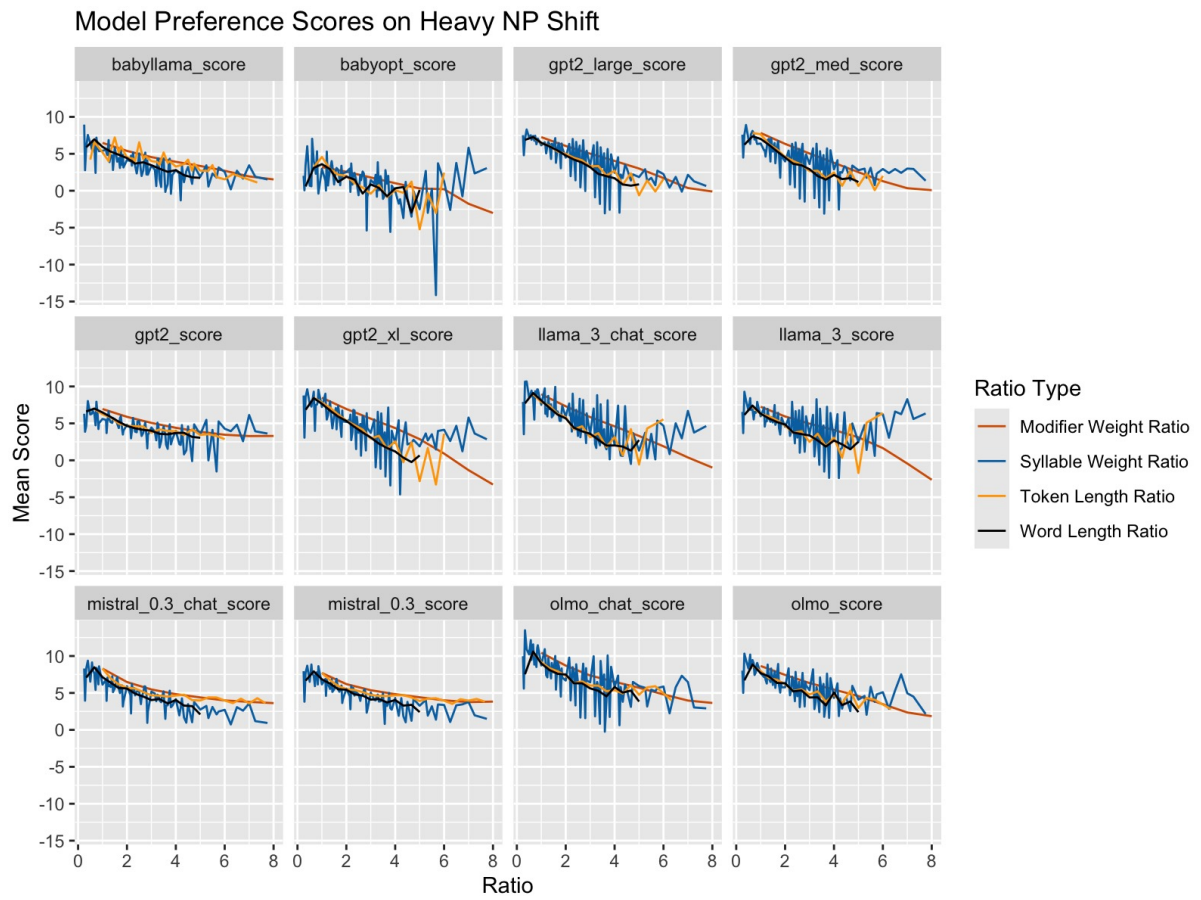
A.1 Regression Table

Model	Var	Full	Token _{ab}	Word _{ab}	Syll _{ab}	Mods _{ab}
GPT-2	PM	0.543	0.533	0.515	0.457	0.517
GPT-2 Med	PM	0.605	0.602	0.581	0.534	0.580
GPT-2 Large	PM	0.611	0.605	0.579	0.556	0.585
GPT-2 XL	PM	0.613	0.610	0.583	0.561	0.590
Llama-3	PM	0.608	0.603	0.586	0.564	0.590
Llama-3 I	PM	0.627	0.619	0.599	0.555	0.602
BabyOPT	PM	0.562	0.550	0.534	0.507	0.542
BabyLlama	PM	0.719	0.702	0.663	0.651	0.701
Mistral v0.3	PM	0.606	0.588	0.563	0.565	0.588
Mistral v0.3 I	PM	0.630	0.618	0.585	0.601	0.618
OLMo	PM	0.655	0.639	0.623	0.612	0.637
OLMo I	PM	0.511	0.505	0.494	0.469	0.481
GPT-2	MPP	0.379	0.325	0.337	0.264	0.370
GPT-2 Med	MPP	0.368	0.309	0.302	0.300	0.356
GPT-2 Large	MPP	0.311	0.265	0.273	0.213	0.302
GPT-2 XL	MPP	0.323	0.249	0.260	0.244	0.310
Llama-3	MPP	0.358	0.316	0.321	0.271	0.351
Llama-3 I	MPP	0.320	0.284	0.291	0.208	0.313
BabyOPT	MPP	0.441	0.433	0.435	0.368	0.432
BabyLlama	MPP	0.310	0.297	0.298	0.281	0.306
Mistral v0.3	MPP	0.325	0.302	0.315	0.239	0.314
Mistral v0.3 I	MPP	0.269	0.238	0.258	0.186	0.256
OLMo	MPP	0.306	0.303	0.296	0.222	0.294
OLMo I	MPP	0.263	0.257	0.253	0.237	0.250
GPT-2	HNPS	0.356	0.343	0.332	0.278	0.349
GPT-2 Med	HNPS	0.654	0.629	0.616	0.540	0.635
GPT-2 Large	HNPS	0.589	0.570	0.554	0.430	0.585
GPT-2 XL	HNPS	0.592	0.587	0.557	0.527	0.589
Llama-3	HNPS	0.542	0.524	0.519	0.485	0.538
Llama-3 I	HNPS	0.452	0.438	0.408	0.359	0.451
BabyOPT	HNPS	0.491	0.457	0.472	0.421	0.450
BabyLlama	HNPS	0.527	0.514	0.466	0.415	0.509
Mistral v0.3	HNPS	0.488	0.476	0.461	0.451	0.458
Mistral v0.3 I	HNPS	0.467	0.441	0.436	0.438	0.434
OLMo	HNPS	0.518	0.500	0.468	0.443	0.513
OLMo I	HNPS	0.411	0.391	0.377	0.290	0.395
GPT-2	DA	0.562	0.554	0.542	0.545	0.549
GPT-2 Med	DA	0.571	0.561	0.552	0.568	0.568
GPT-2 Large	DA	0.541	0.528	0.506	0.533	0.533
GPT-2 XL	DA	0.561	0.553	0.540	0.552	0.555
Llama-3	DA	0.554	0.544	0.532	0.538	0.549
Llama-3 I	DA	0.503	0.493	0.490	0.486	0.496
BabyOPT	DA	0.252	0.242	0.238	0.241	0.237
BabyLlama	DA	0.630	0.616	0.603	0.622	0.623
Mistral v0.3	DA	0.540	0.529	0.526	0.528	0.534
Mistral v0.3 I	DA	0.478	0.469	0.461	0.467	0.467
OLMo	DA	0.467	0.460	0.458	0.449	0.459
OLMo I	DA	0.392	0.384	0.388	0.374	0.385

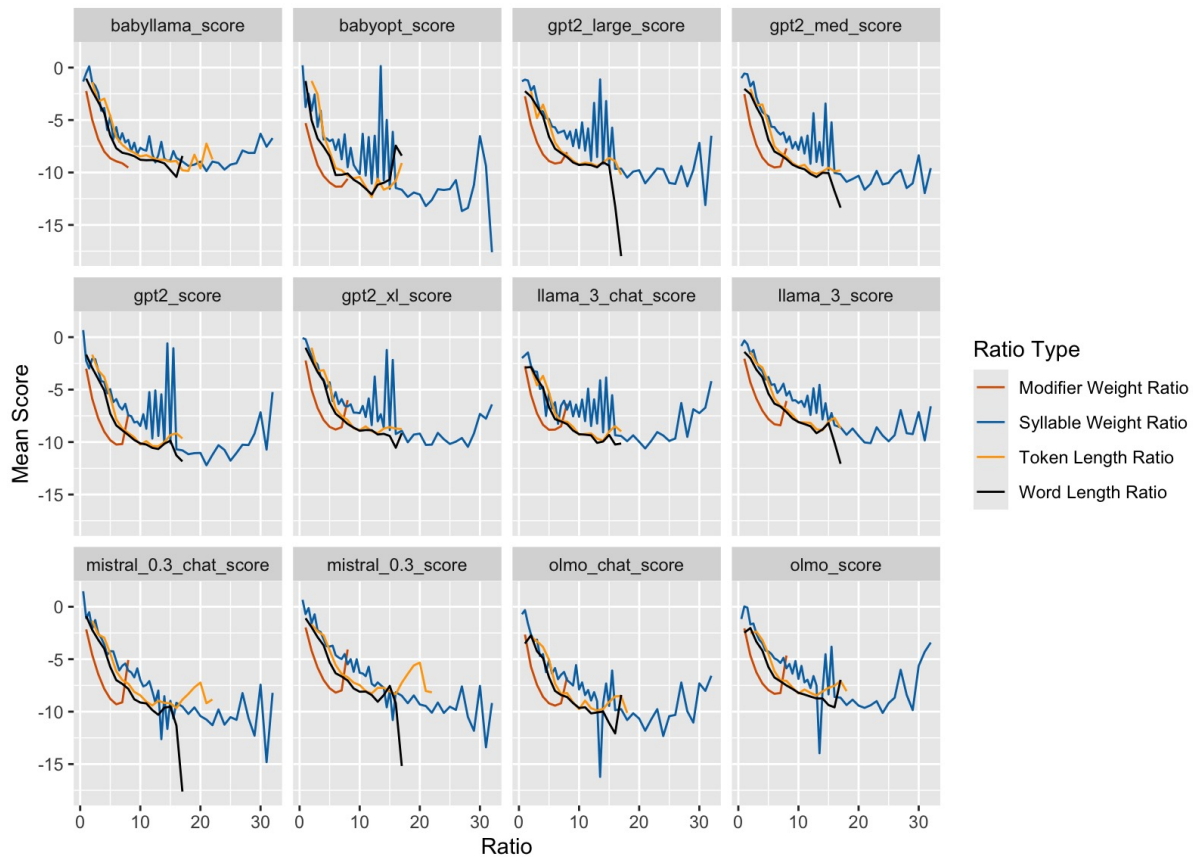
Table 5: R-sq. scores for each model on each variable. $[Metric]_{ab}$ denotes the R-sq. score when excluding the metric; a larger difference from the original score denotes higher significance in contributing to the overall score.

A.2 Model Behavior Plots

A.2.1 Synthetic Data

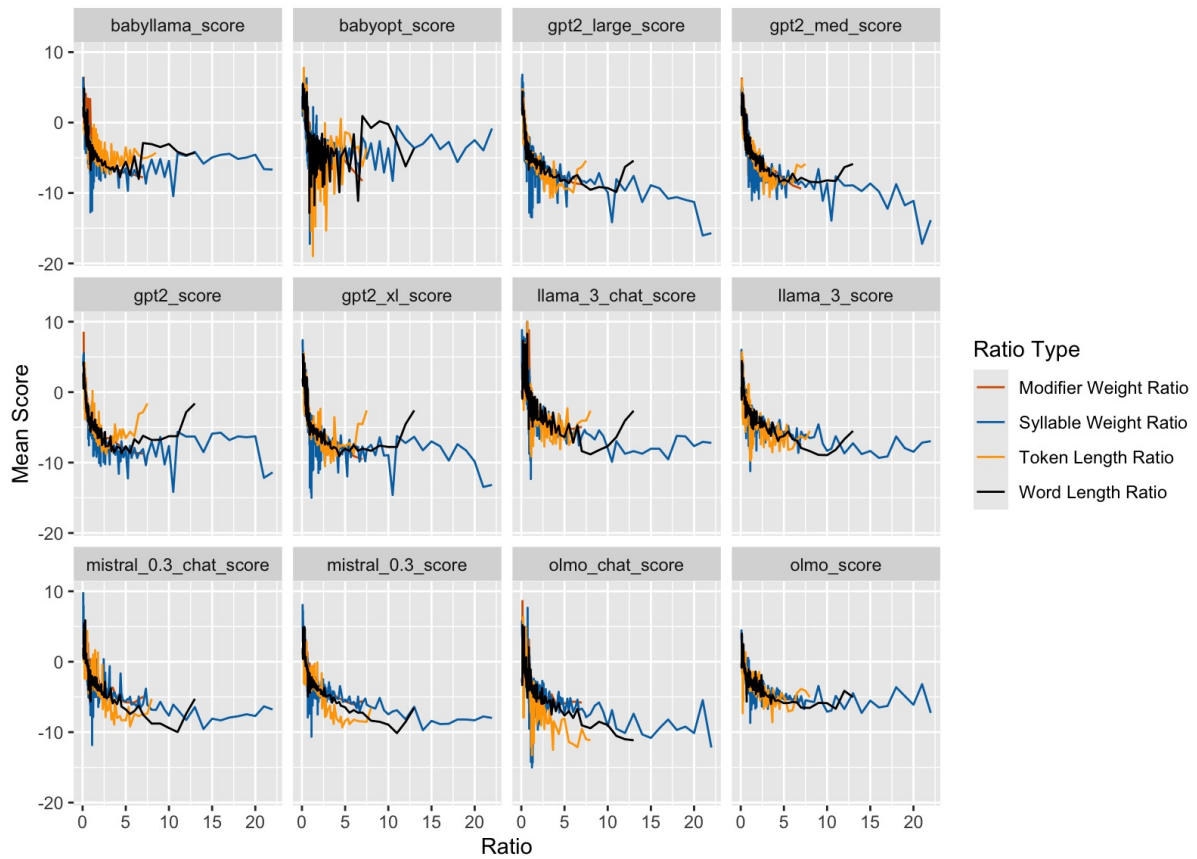


Model Preference Scores on Particle Movement

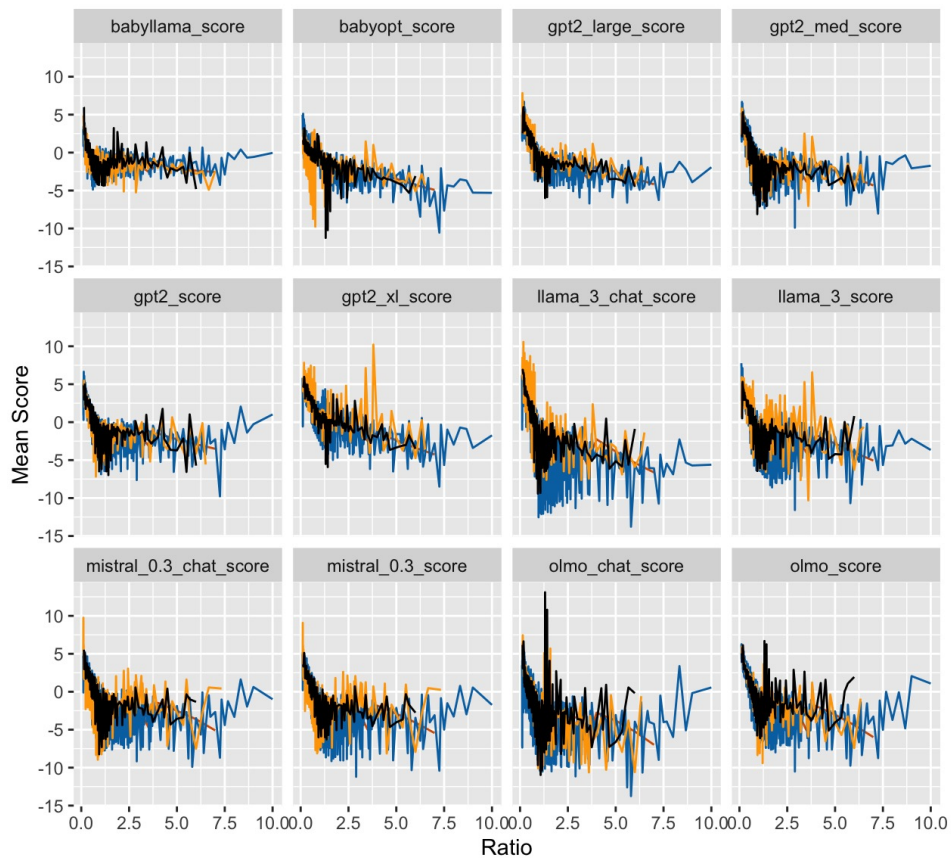


A.2.2 Mined Data

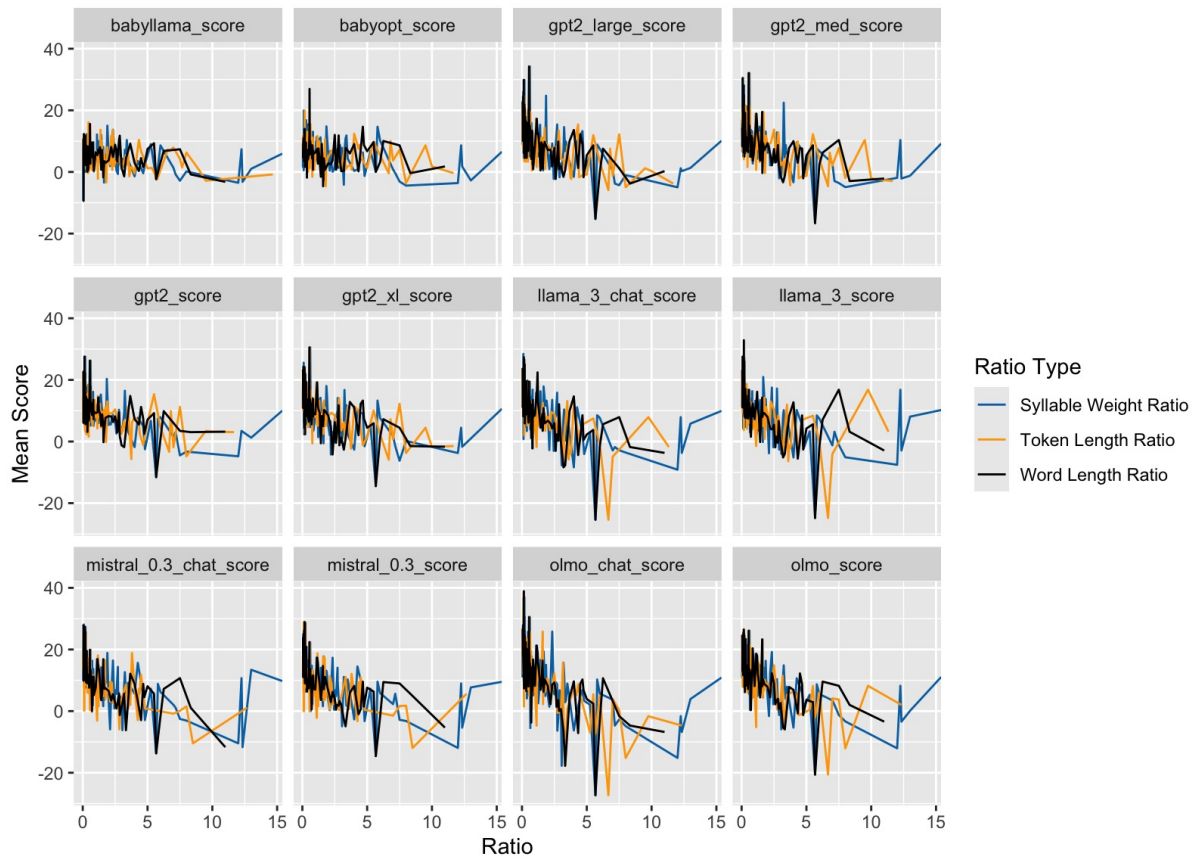
Model Preference Scores on Dative Alternation



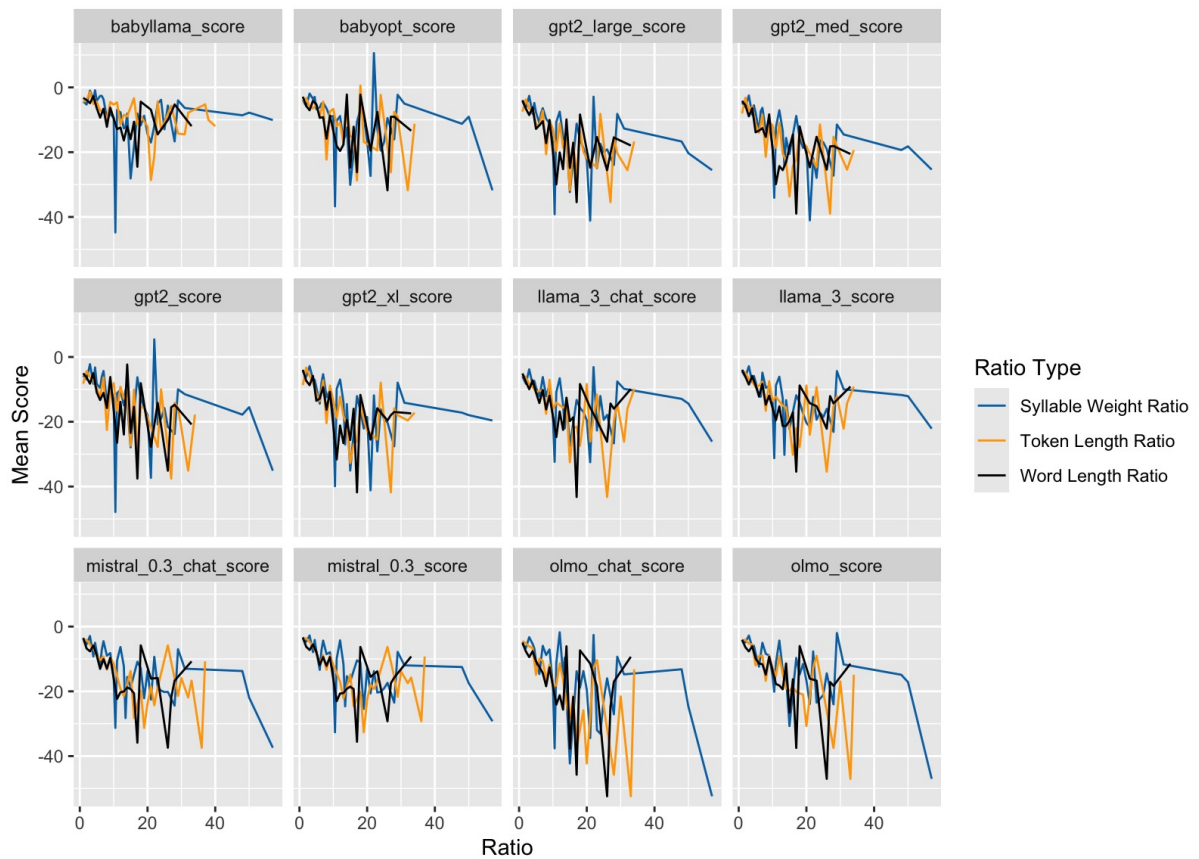
Model Preference Scores on Multiple PPs



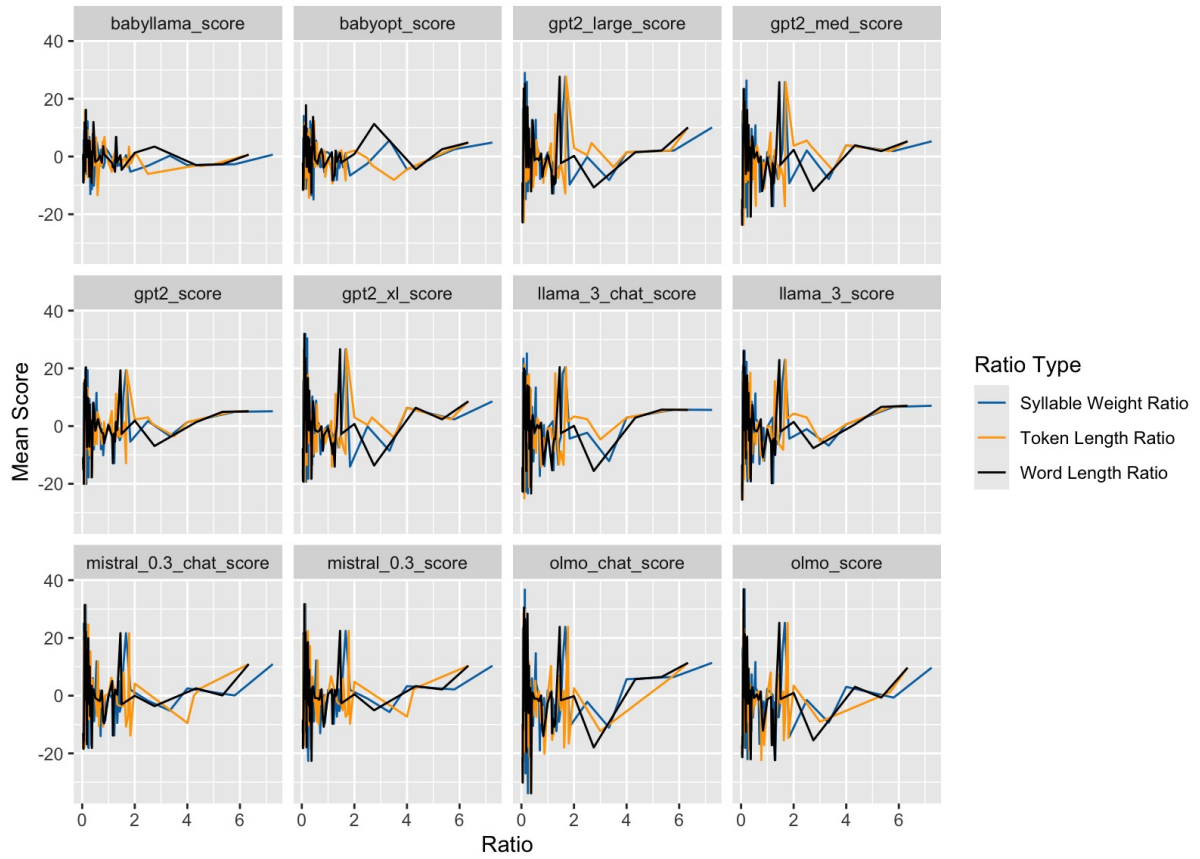
Model Preference Scores on Mined Heavy NP Shift



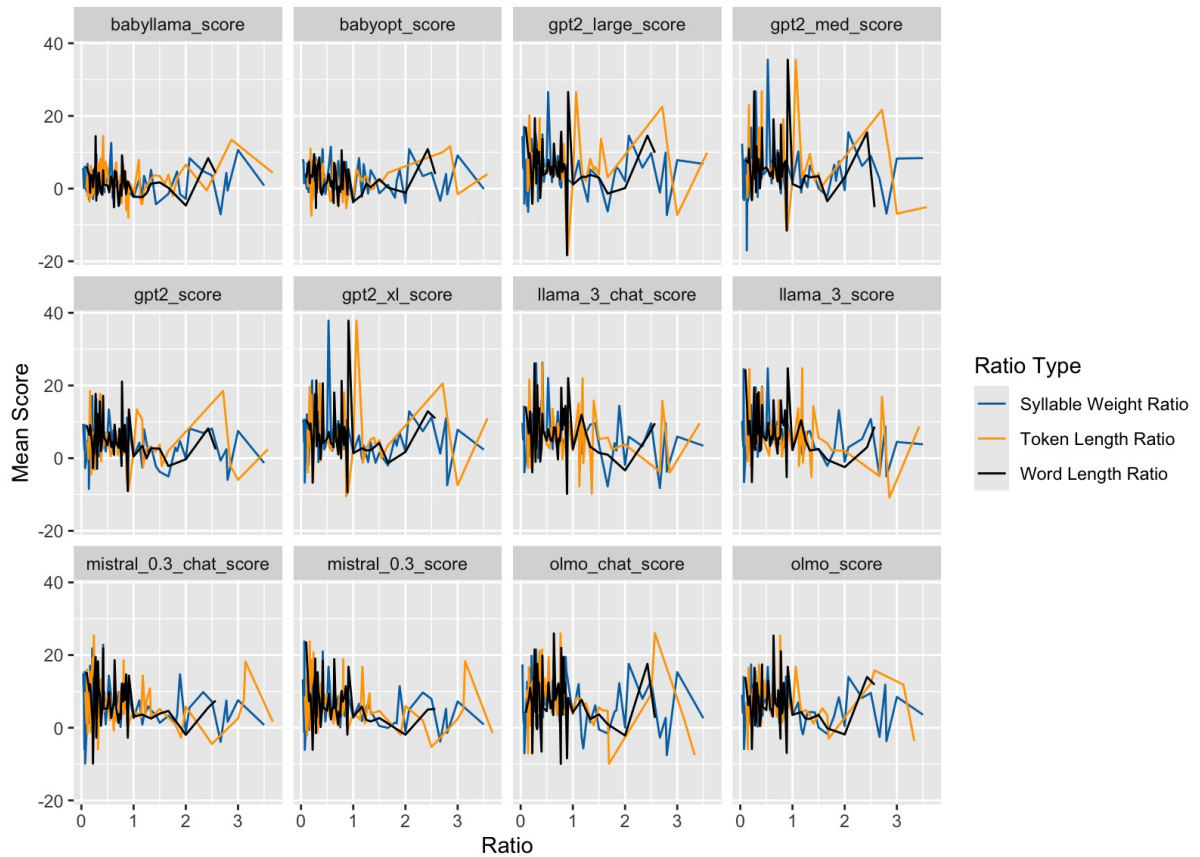
Model Preference Scores on Mined Particle Movement



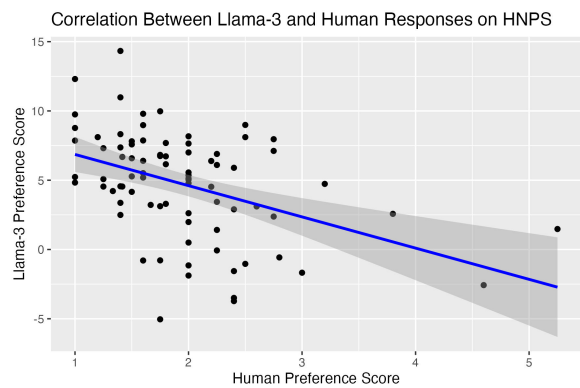
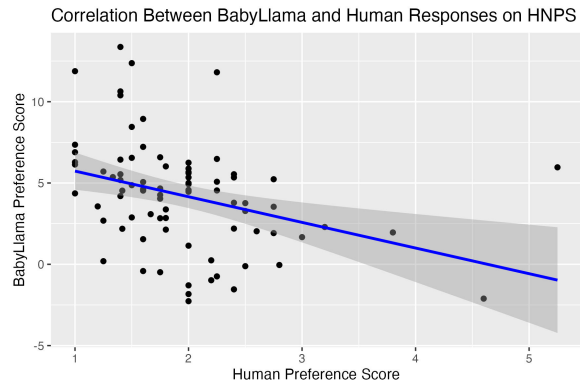

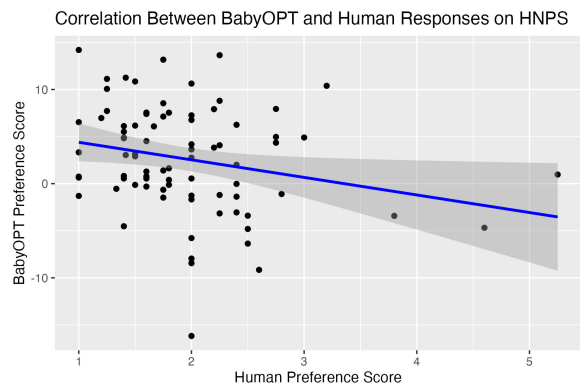

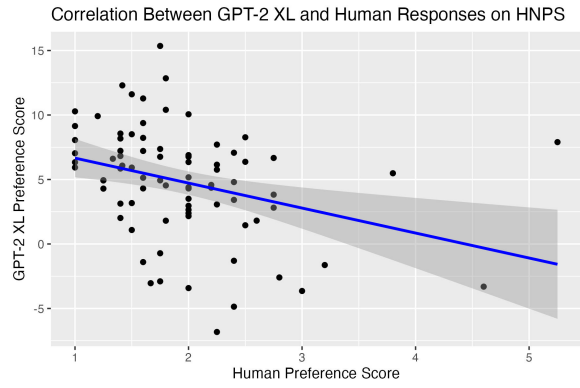

Model Preference Scores on Mined Dative Alternation




Model Preference Scores on Mined Multiple PPs




A.3 Human-Model Agreement Plots




Correlation Between Llama-3 Instruct and Human Responses on HN



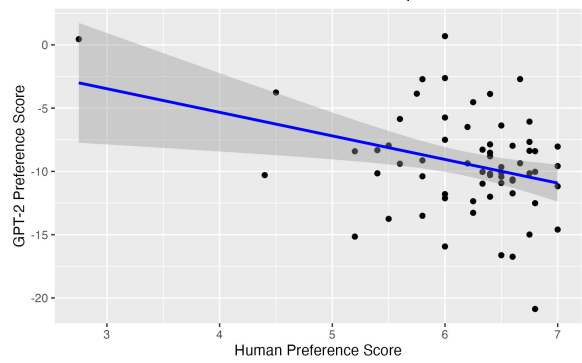
Correlation Between OLMo Instruct and Human Responses on HNP




Correlation Between Mistral v0.3 and Human Responses on HNPS




Correlation Between GPT-2 and Human Responses on PM




Correlation Between Mistral v0.3 Instruct and Human Responses on




Correlation Between GPT-2 Medium and Human Responses on PM



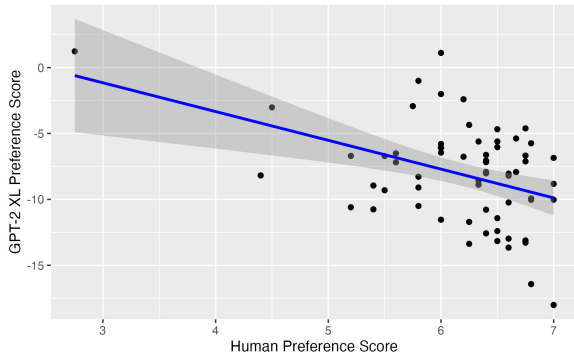
Correlation Between OLMo 7B and Human Responses on HNPS




Correlation Between GPT-2 Large and Human Responses on PM




Correlation Between GPT-2 XL and Human Responses on PM




Correlation Between Llama-3 Instruct and Human Responses on PM




Correlation Between BabyOPT and Human Responses on PM




Correlation Between Mistral v0.3 and Human Responses on PM




Correlation Between BabyLlama and Human Responses on PM



Correlation Between Mistral v0.3 Instruct and Human Responses on PM



Correlation Between Llama-3 and Human Responses on PM



Correlation Between OLMo and Human Responses on PM

