

Not quite Sherlock Holmes: Language model predictions do not reliably differentiate impossible from improbable events

James A. Michaelov^{1,2,3}, Reeka Estacio³, Zhien Zhang³, Benjamin K. Bergen³

¹ Department of Brain and Cognitive Sciences, MIT ² MIT Libraries CREOS

³ Department of Cognitive Science, UCSD

jamic@mit.edu, {rdestaci, zhz067, bkbergen}@ucsd.edu

Abstract

Can language models reliably predict that possible events are more likely than merely improbable ones? By teasing apart possibility, typicality, and contextual relatedness, we show that despite the results of previous work, language models' ability to do this is far from robust. In fact, under certain conditions, all models tested—including Llama 3, Gemma 2, and Mistral NeMo—perform at worse-than-chance level, assigning higher probabilities to impossible sentences such as ‘the car was given a parking ticket by the brake’ than to merely unlikely sentences such as ‘the car was given a parking ticket by the explorer’.

🌸 [Data, Code, and Analyses](#)

1 Introduction

“How often have I said to you that when you have eliminated the impossible, whatever remains, however improbable, must be the truth?”

Sherlock Holmes, *The Sign of the Four*
(Doyle, 1890)

Consider the following scenario:

- (1) Marissa forgot to bring her pillow on her camping trip. As a substitute for her pillow, she filled up an old sweater with...
 - (a) clothes
 - (b) leaves
 - (c) water

(Glenberg and Robertson, 2000)

Clothes is the obvious best answer here, and that most preferred by humans when presented with the options shown above (see Glenberg and Robertson, 2000). Nonetheless, *leaves* is also a possible continuation. In fact, Glenberg and Robertson (2000) find that experimental participants rate *leaves* to be

significantly more sensible than *water*, despite the fact that the sentence *as a substitute for her pillow, she filled up an old sweater with leaves* describes a highly unusual event.

This example (and indeed, all of Glenberg and Robertson, 2000) hinges specifically on the physical properties of objects and their ‘affordances’ (i.e., what they can be used to do; Gibson, 1966, 1979). But it exemplifies a broader behavior that humans exhibit largely unconsciously, namely, using knowledge of the world to distinguish between the impossible and the merely atypical. In this study, we focus on whether language models also have the capability to do so.

A growing body of work on world knowledge and event understanding investigates whether language models can select the most likely or plausible of a set of possibilities (e.g., Zellers et al., 2019; Sakaguchi et al., 2020; Bisk et al., 2020; see subsection 2.1). But this work often conflates different ways that the ‘incorrect’ options can be incorrect; some are impossible, but others are simply less typical. While finding that models are better or worse at identifying likely events is in itself important, it is only one component of robust event understanding. In fact, any use of language models in practice, whether as part of larger systems or directly as chat assistants, is likely to involve novel and unexpected situations. Thus being able to tell the difference between the impossible and the merely unlikely is crucial, especially in critical domains where the use of language models has been suggested, such as medicine (see, e.g. Van Veen et al., 2024; Liévin et al., 2024; Singhal et al., 2025).

For this reason, we argue that it is important to also assess the extent to which language models are able to differentiate between possible but improbable events and impossible ones. We call this the **Sherlock Holmes Task** after the famous epigraphic quote (Doyle, 1890).

In this study, we evaluate language models on a

specific version of the Sherlock Holmes task, asking whether they are able to reliably assign a higher probability to sentences describing possible (but in some cases, atypical or unlikely) events (e.g., *the car was given a parking ticket by the delinquent*; Vega-Mendoza et al., 2021) than their impossible equivalents (e.g., *the car was given a parking ticket by the stamp*). We draw on experimental stimuli from previous research investigating how humans process such sentences (Vega-Mendoza et al., 2021; Chow and Phillips, 2013). Prior work shows that language models assign impossible sentences with words semantically related to their context higher probabilities than equivalent sentences with unrelated words (Michaelov and Bergen, 2022). Thus, we also include a further adversarial component, specifically considering cases where impossible sentences include related words (e.g., *the car was given a parking ticket by the brake*) and possible sentences include unrelated words (e.g., *the car was given a parking ticket by the explorer*).

In contrast to previous work suggesting that language models have relatively good world knowledge and event understanding capabilities (e.g., Kauf et al., 2023), we find that both event atypicality and semantic relatedness lead to significant drops in performance. In the most highly adversarial case, where we compare the probabilities assigned to possible but atypical sentences with unrelated words to those assigned to impossible sentences with related words (e.g., *the car was given a parking ticket by the explorer* vs. *brake*), we find that all language models tested perform at or below chance—that is, they assign the impossible sentences a higher probability half or more of the time. We further find that this effect does not disappear with scale—in fact, smaller models often do better than their larger counterparts. Together, these results suggest that language models may rely on typicality and semantic relatedness cues when making predictions rather than actual world knowledge, and thus, that any previously-ascribed event understanding capabilities are far from robust.

2 Background

2.1 Event Typicality

Texts vary in the extent to which they describe a state of affairs that is likely and congruent with reality. In studies of human language comprehension, this has been operationalized in a variety of ways, including *plausibility* (e.g., Paczynski and Kuper-

berg, 2012; Vega-Mendoza et al., 2021), *sensibility* (e.g., Glenberg and Robertson, 2000), *typicality* (e.g., Urbach and Kutas, 2010), *combinability* (e.g., Chow and Phillips, 2013), and *canonicity* (e.g., Nieuwland and Van Berkum, 2006). For consistency, in this paper we use the term *event typicality* as a general term encompassing the general unifying idea behind these—the extent to which an event described is likely to occur.

One common way to evaluate language models’ sensitivity to event typicality is in commonsense and physical reasoning benchmarks (e.g., Levesque et al., 2012; Sakaguchi et al., 2020; Talmor et al., 2019; Bisk et al., 2020; Zellers et al., 2018, 2019), which often at least implicitly target this capability (Davis and Marcus, 2015; Storks et al., 2020). Take, for example, the following item from the widely-used HellaSwag dataset (Zellers et al., 2019):

- (2) A woman is outside with a bucket and a dog. The dog is running around trying to avoid a bath. She...
 - (a) rinses the bucket off with soap and blow dry [*sic*] the dog’s head.
 - (b) uses a hose to keep it from getting soapy.
 - (c) gets the dog wet, then it runs away again.
 - (d) gets into a bath tub with the dog.

In this example, continuation (c) is ‘correct’—it corresponds to the most typical sequence of events. And indeed, continuation (c) is the one that actually occurs in the ActivityNet Captions dataset (Krishna et al., 2017) from which the example is derived. However, the ‘incorrect’ continuations, while less plausible and appropriate, also all describe possible continuations of the event in question—we don’t see any truly impossible continuations such as [*she*] and *the dog are chased by the bucket*.

Several studies have explicitly investigated event possibility (Michaelov et al., 2023; Kauf et al., 2023; Hanna et al., 2023). These studies generally base their analyses on experimental stimuli derived from human psycholinguistic studies where impossible sentences always involve animacy violations such as *the peanut was in love* (Michaelov et al., 2023; Hanna et al., 2023; original stimulus from Nieuwland and Van Berkum, 2006), or *the laptop bought the teacher* (Kauf et al., 2023; original stimulus from Fedorenko et al., 2020).

These studies investigate whether language models can differentiate between possible and impos-

sible events. Yet in all cases, the best continuation is a typical one. This is true even in the case of Kauf et al. (2023), who investigate whether there is a difference between the performance of language models at distinguishing plausible from impossible events compared to distinguishing plausible from merely implausible (i.e., atypical) events. Thus, none of these studies are able to successfully tease apart the effects of possibility and typicality, which is important because it is precisely in atypical situations that it is important to distinguish between possible and impossible events.

2.2 Semantic Relatedness

A number of studies have shown that the word predictions of both humans (e.g., Ettinger et al., 2016; Uchida et al., 2021) and language models (e.g., Misra et al., 2020; Michaelov et al., 2024) correlate with the word’s semantic relatedness to its context. For example, Misra et al. (2020) find that adding words like *airplane* (either alone or in a sentence context) before sentences like *I wanted to become a...* increases the probability of BERT predicting words like *pilot* compared to when adding a control word such as *table*.

This tendency is in principle adaptive—it would be surprising if within an utterance or connected discourse, one were to encounter a sentence completely unrelated to anything mentioned previously (see, e.g., Grice, 1975, 1989; Sperber, 1986). Thus, contextual relatedness may often function as a reliable heuristic of what is to come next.

However, the evidence suggests that using contextual semantic relatedness as a partial basis for prediction may not always be beneficial. Consider, for example, the following text from a study by Metusalem et al. (2012):

- (3) We’re lucky to live in a town with such a great art museum. Last week I went to see a special exhibit. I finally got in after waiting in a long...
- (a) line
 - (b) painting
 - (c) toothbrush

Metusalem et al. (2012) find that while *painting* and *toothbrush* are both impossible sentence continuations, and are equally unlikely to be offered up as possible continuations by experimental participants, *painting* is more strongly predicted during

the process of language comprehension, which Metusalem et al. (2012) argue is due to its relation to the event discussed in context (i.e., visiting an art museum). Michaelov and Bergen (2022) replicate this finding in language models, finding that they also predict *painting* to be more likely. While in this context, such a result may not be particularly harmful, as it amounts to a difference between two impossible continuations, it does suggest that reliance on contextual semantic similarity could be problematic in other cases.

2.3 Possibility

We are only aware of two studies that directly test whether language models are able to successfully assign sentences describing possible but atypical events a higher probability than impossible ones. Kauf et al. (2023) look across a large set of sentences of comparable length and structure, finding that overall, impossible sentences are assigned lower probabilities than merely implausible ones.

Jones et al. (2022), on the other hand, carry out a study on the full set of stimuli from Glenberg and Robertson (2000), as exemplified in (1). Thus they directly investigate cases where the texts differ only by one word—a *critical word*—that determines whether they are likely (as in *she filled up an old sweater with clothes*; see (1)), possible but atypical (i.e., afforded; e.g., *she filled up an old sweater with leaves*), or impossible (i.e., nonafforded; e.g., *she filled up an old sweater with water*). Jones et al. (2022) find that BERT and RoBERTa are not able to consistently assign the possible but atypical events a higher probability than the impossible ones, but (on average) GPT-3 does; though not to the extent that it can fully account for human sensibility judgments.

However, because this study was aimed at modeling human responses, it does not report the proportion of sentences where the possible continuation was assigned a higher probability than the impossible one (i.e., model accuracy). Thus, it is unclear how consistent this pattern is across individual items. Additionally, the original stimuli, as constructed by Glenberg and Robertson (2000) explicitly balance the semantic relatedness of the afforded and unafforded options in order to ensure that it is not used as a heuristic—that is, to avoid cases where the afforded continuation has a higher semantic relatedness than the unafforded one. However, by the same token, this also removes the possibility of the reverse situation where

Lang.	Sentence	Possibility	Typicality	Relatedness	
Eng.	The cure for the disease was discovered by the <u>doctor</u> .	✓ Possible	✓ Typical	✓ Related	(PTR)
	The cure for the disease was discovered by the <u>patient</u> .	✓ Possible	✗ Atypical	✓ Related	(PAR)
	The cure for the disease was discovered by the <u>guest</u> .	✓ Possible	✗ Atypical	✗ Unrelated	(PAU)
	The cure for the disease was discovered by the <u>medication</u> .	✗ Impossible	✗ Atypical	✓ Related	(IAR)
	The cure for the disease was discovered by the <u>stamp</u> .	✗ Impossible	✗ Atypical	✗ Unrelated	(IAU)
Mand.	高材生把数学题解答了 'The student solved the math problem'	✓ Possible	✓ Typical	✓ Related	(PTR)
	高材生把数学题挂起了 'The student hung the math problem'	✓ Possible	✗ Atypical	✗ Unrelated	(PAU)
	高材生把数学题难倒了 'The student baffled the math problem'	✗ Impossible	✗ Atypical	✓ Related	(IAR)
	高材生把数学题困住了 'The student restrained the math problem'	✗ Impossible	✗ Atypical	✗ Unrelated	(IAU)

Table 1: Examples of the types of sentences used in the present study. All English sentences are drawn from Vega-Mendoza et al. (2021), and all Mandarin sentences are drawn from Chow and Phillips (2013).

reliance on a semantic relatedness heuristic may lead to mistakenly over-estimating the probability of semantically related words relative to unrelated words. Thus, the results may also over-estimate the extent to which language models can reliably predict possible events to be more likely than impossible ones.

3 The Present Study

As has been discussed, the effects of possibility, typicality, and contextual semantic relatedness on language model probability have each been investigated individually. However, these are all intertwined—all else being equal, we should expect that typical events are also possible, and that words describing typical and possible events are more likely to be related to their semantic contexts than words describing atypical and impossible events. In this study, therefore, we aim to disentangle the effect of each by investigating the effects on prediction of cases where these cues conflict.

To do this, our version of the Sherlock Holmes task uses a minimal pairs paradigm (see, e.g., Linzen et al., 2016; Marvin and Linzen, 2018; Warstadt et al., 2019, 2020) where the language model is presented with two sentences that differ by a critical word that makes the sentence either possible or impossible. In order to allow any language model to carry out our task (not just large or finetuned models, see Hu and Frank, 2024), we follow previous work such as BLiMP (Warstadt et al., 2020) and directly compare the probability that a language model assigns to each sentence to determine the more probable one.

As shown in Table 1, with each sentence, the

critical word (underlined) is manipulated such that the event it describes is either Possible or Impossible and either Typical or Atypical. The critical word itself is either semantically Related or Unrelated to its preceding context. Thus, a minimal pair might pit a Possible-Typical-Related (PTR) sentence (e.g., *the cure for the disease was discovered by the **doctor***) against its corresponding Impossible-Atypical-Unrelated sentence (e.g., *the cure for the disease was discovered by the **stamp***). We consider a language model to be correct if it assigns a higher probability to the former (possible) sentence than the latter (impossible) sentence.

In this study, as in previous work comparing possible and impossible sentences in language models (Kauf et al., 2023), we look at impossibility arising from animacy violations. An animacy violation occurs when an inanimate entity has an event role that requires animacy. For example, in *the cure for the disease was discovered by the...*, the discoverer needs to be animate, and since stamps are inanimate, completing the sentence with *stamp* is an animacy violation. Thus, when a sentence with an animacy violation is interpreted literally, it refers to an impossible state of affairs.

All English stimuli were drawn from a human study carried out by Vega-Mendoza et al. (2021; based on Paczynski and Kuperberg, 2012), and all Mandarin stimuli from a study carried out by Chow and Phillips (2013). English minimal pairs are exactly as shown in Table 1; while due to their nature, the Mandarin stimuli are embedded in larger sentences as they appear in Chow and Phillips (2013), which are identical across experimental conditions. All English tasks were made up of 154 sentence

pairs, and all Mandarin tasks of 57 sentence pairs.

4 Experiment 1: Typical vs. Atypical

4.1 Introduction

In our first experiment, we first assess how well language models tell apart sentences denoting typical and atypical (both possible and impossible) events. Like Kauf et al. (2023), we also test whether there is a difference in performance based on the type of atypical sentence compared (impossible vs. possible but atypical). We go beyond previous work both by looking at how performance is impacted by semantic relatedness, and by looking at a language other than English, namely, Mandarin.

4.2 Method

We construct minimal pairs based on the conditions described in Table 1. For both languages, we test whether language models assign a higher probability to Possible-Typical-Related events compared to impossible critical words (see Table 1). For English, we are also able to test whether language models assign a higher probability to Possible-Typical-Related sentences than to sentences with possible but atypical critical words.

We run our analyses using the Language Model Evaluation Harness (Gao et al., 2021) on 35 language models of the BLOOM (BigScience Workshop et al., 2023), Gemma (Gemma Team et al., 2024a,b), Llama (Llama Team, 2024), mGPT (Shlitzhko et al., 2024), Mistral (Jiang et al., 2023), OLMo (Groeneveld et al., 2024), Qwen (Yang et al., 2024; Qwen et al., 2025), SmoLLM (Allal et al., 2024), XGLM (Lin et al., 2022), and Yi (01.AI et al., 2025) model families (see Appendix A for full list). Because we are interested in language models’ world knowledge and the extent to which it is used in prediction—rather than their ability to answer questions about it—we limit our analysis to pretrained-only (or ‘base’) models.

4.3 Results

The results are shown in Figure 1 (the individual accuracy scores for all models is provided in Appendix A). In addition to comparing the accuracies of different tasks numerically, we also run pairwise statistical tests. Specifically, to test whether there is a difference between language model performance at two different tasks, we use logistic mixed-effects regressions to predict whether a given answer is correct or not, and predict this based on which task

it is, as well as including the maximal random effects structure that would converge with no singular fits, namely, random intercepts for each language model and sentence context. The specific statistical test was a likelihood ratio test including or excluding task as a predictor.

First, we see that on the typical vs. impossible tasks, accuracy is lower when the impossible critical word is related compared to when it is unrelated in both English ($\chi^2(1) = 616.36, p < 0.0001$) and Mandarin ($\chi^2(1) = 163.06, p < 0.0001$). The same is true with the typical vs. atypical comparison on English stimuli, where performance is lower when the implausible sentence is related rather than unrelated ($\chi^2(1) = 1417.80, p < 0.0001$). Finally, we replicate the result found by Kauf et al. (2023) that language models tend to be worse at the typical vs. atypical than the typical vs. impossible tasks. We find that this is the case both when the atypical or impossible critical word is related ($\chi^2(1) = 587.32, p < 0.0001$), and when it is unrelated ($\chi^2(1) = 63.09, p < 0.0001$).

4.4 Discussion

In line with previous work (Kauf et al., 2023), our results show that language models can differentiate between possible and impossible events as well as typical and atypical events, and are better at the former comparison.

In addition, we show for the first time that performance at each of these is affected by contextual semantic relatedness. If the atypical critical word (either possible or impossible) is semantically related to the event described in its context, it leads to decreased performance relative to when the atypical word is unrelated. This suggests that in cases where the atypical word is semantically related to the event in the context, there is an increased risk of it being calculated as more probable than the actually typical word than when the atypical word is semantically unrelated, including in cases where the atypical word renders the sentence impossible.

5 Experiment 2: Atypical vs. Impossible

5.1 Introduction

In Experiment 1, we showed that while language models can tell apart sentences denoting typical and atypical events, the extent to which this is the case is impacted by semantic relatedness, and specifically, the extent to which the critical word in a sentence is semantically related to its context. In

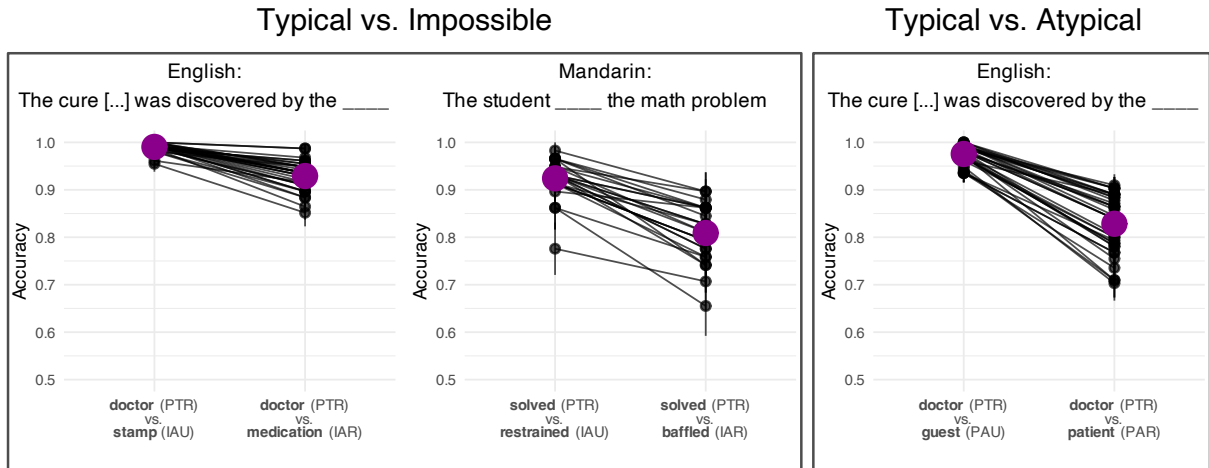


Figure 1: Language model scores for all tasks comparing sentences describing typical events (Possible-Typical-Related) to those describing impossible (Impossible-Atypical-Related or Impossible-Atypical-Unrelated) or merely atypical (Possible-Atypical-Related or Possible-Atypical-Unrelated) events. Individual model scores are shown in black (standard error calculated as in Biderman et al., 2024), while the mean over all models is shown in purple (standard error calculated over model means).

Experiment 2, we instead focus on the more consequential question of how well language models can tell apart impossible and merely atypical events, again looking at how this is impacted by the semantic relatedness of the critical word to the context. However, in this experiment we evaluate models on the more difficult case where possible events are also atypical (Possible-Atypical-Unrelated or Possible-Atypical-Related).

5.2 Method

The method is the same as in Experiment 1, as are the impossible stimuli. However, for the possible stimuli, we instead draw on the atypical but possible sentence types: Possible-Atypical-Unrelated (English and Mandarin) and Possible-Atypical-Related (English only). We also limit our analysis to the subset of the language models used in Experiment 1 that were trained on Mandarin.

5.3 Results

The results are shown in Figure 2 (individual accuracy scores for all models provided in Appendix A). First, we see that in all cases, an atypical but possible critical word leads to worse performance than a typical critical word.

There is a drastic drop in performance when the impossible critical word is related compared to when it is unrelated, both when the possible but atypical critical word is related (English only: $\chi^2(1) = 1879.50, p < 0.0001$) and when it is unrelated ($\chi^2(1) = 3792.85, p < 0.0001$;

Mandarin¹: $z = 13.18, p < 0.0001$). There is also a drop in performance on the English stimuli when the possible but atypical critical word is related compared to when it is unrelated (vs. Impossible-Atypical-Unrelated: $\chi^2(1) = 1277.45, p < 0.0001$; vs. Impossible-Atypical-Related: $\chi^2(1) = 3570.05, p < 0.0001$). When we compare performance to Experiment 1, we also see that for tasks where the possible critical word is related, the models perform significantly worse when the word is atypical rather than typical (vs. Impossible-Atypical-Unrelated: $\chi^2(1) = 273.69, p < 0.0001$; vs. Impossible-Atypical-Related: $\chi^2(1) = 1240.70, p < 0.0001$).

5.4 Discussion

These results show that language models are worse at differentiating between implausible and impossible events than they are between plausible and impossible events. While this may not be a surprising result given previous work (e.g., Jones et al., 2022; Kauf et al., 2023), as far as we know, this is the first study to provide direct evidence of this.

In addition, this experiment highlights, perhaps even more strikingly than Experiment 1, the role of semantic relatedness in prediction. We see that performance is impacted not only by the extent to which the critical word of impossible sentences is related to the context, but also the extent to which

¹For this comparison, the null regression to be used in the likelihood ratio test did not converge, so we instead used the asymptotic Wald test to calculate the effect of task.

Atypical (Possible) vs. Impossible

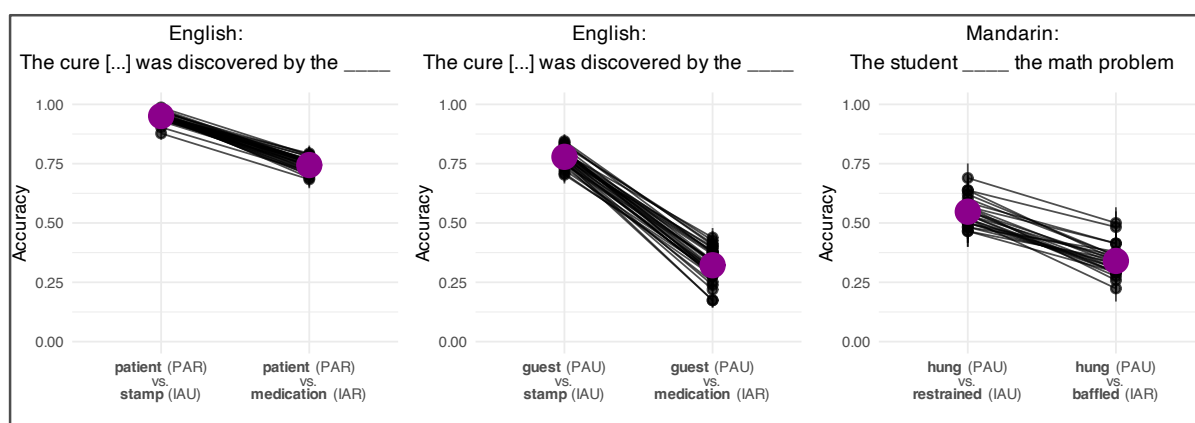


Figure 2: Language model scores for all tasks comparing sentences describing possible but atypical (Possible-Atypical-Related or Possible-Atypical-Unrelated) events to those describing impossible (Impossible-Atypical-Related or Impossible-Atypical-Unrelated) events. Individual model scores are shown in black, and overall mean scores in purple.

the critical word of the possible sentences is related. In all cases, sentences with semantically related critical words are assigned higher probabilities.

Crucially, in the specific case where there is an event-unrelated possible critical word and event-related impossible word, we see that performance drops below chance in both languages—that is, the language models calculate the event-related impossible words to be more likely than the event-unrelated possible words more than half of the time.

6 Experiment 3: Statistical Reliability

6.1 Introduction

In addition to replicating Kauf et al. (2023), we also make several novel findings in Experiments 1–2. Performance drops across the board when the possible critical word is atypical rather than typical, as well as when it is semantically unrelated rather than related. Additionally, performance drops when the impossible critical word is related rather than unrelated. While it may be unsurprising that language models appear to predict typical words to be more probable than atypical ones and semantically related words to be more probable than semantically unrelated words (especially given, e.g., Michaelov and Bergen, 2022; Michaelov et al., 2024), we show for the first time that this has a significant impact on their ability to assign higher probabilities to possible sentences than impossible ones.

6.2 Method

To investigate how robust this is and account for possible confounds, we carry out an analysis testing the extent to which the relatedness and typicality of each critical word impact language model performance. Because Vega-Mendoza et al. (2021) provide both the numeric values of relatedness (calculated using Latent Semantic Analysis; Dumais et al., 1988) and typicality (operationalized as human plausibility ratings), we use these, and thus only carry out our analyses on the English-language stimuli. We construct a logistic mixed-effects regression predicting whether a given language model correctly assigns the possible sentence a higher probability than the impossible one. As predictors, we include the semantic relatedness of the possible and impossible critical words, the typicality of the possible and impossible critical words, and the frequency of the possible and impossible critical words (a possible confound; see, e.g., McCoy et al., 2024). We also include random intercepts for each language model and sentence context, as well as random uncorrelated slopes of each predictor for each of these. Thus, we test the extent to which each predictor explains language model performance while accounting for the effects of the other predictors. We use likelihood ratio tests to compare the fit of regressions including or excluding the predictor of interest.

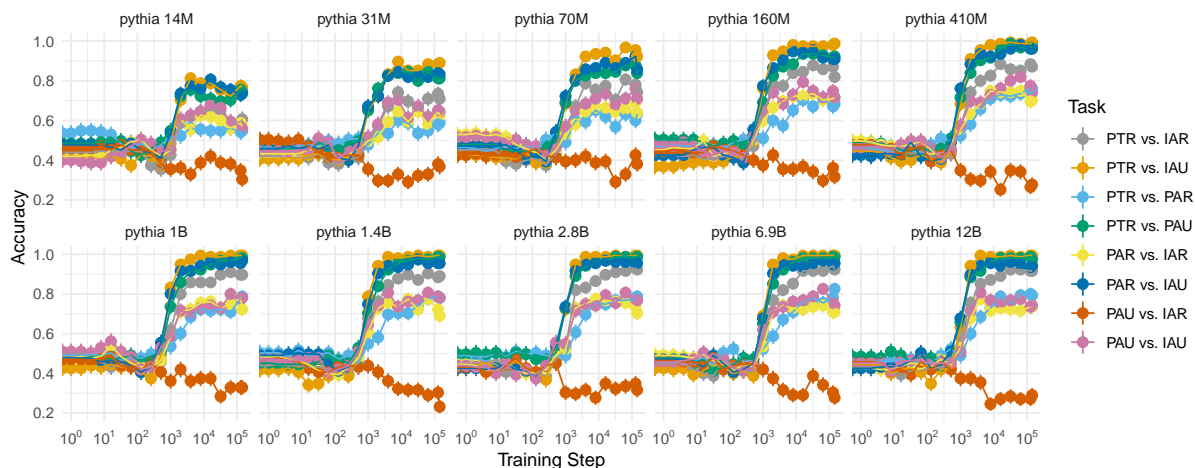


Figure 3: Pythia language model scores at all English tasks (from Experiments 1–2) over the course of training.

6.3 Results and Discussion

The results align with the results of Experiments 1–2. Specifically, we see that, for a given pair, the possible critical word being more semantically related improves performance ($\chi^2(1) = 176.17, p < 0.0001$), while the impossible critical word being more semantically related degrades performance ($\chi^2(1) = 197.58, p < 0.0001$). Similarly, we see that the possible critical word being more typical improves performance ($\chi^2(1) = 128.76, p < 0.0001$) and the impossible word being more typical degrades performance ($\chi^2(1) = 127.39, p < 0.0001$). Thus, we see that plausibility and semantic relatedness do increase the probability assigned to a sentence by a language model to the extent that it significantly alters performance at the task. We also see that higher possible critical word frequency improves performance ($\chi^2(1) = 394.32, p < 0.0001$) while higher impossible critical word frequency degrades performance ($\chi^2(1) = 302.54, p < 0.0001$), suggesting that this is indeed a possible confound, and one that should be studied in future work.

7 Experiment 4: Scaling Effects

7.1 Introduction

Experiments 1 and 2 demonstrated fragility in language models’ ability to differentiate between possible and impossible events, resulting from whether or not critical words in the sentences are related or unrelated to the context. However, there is a lot of variation between models. Given the general trend that larger models trained on more data tend to generally perform better (Kaplan et al., 2020;

Rae et al., 2022; Wei et al., 2022; Hoffmann et al., 2022), one might expect that the best-performing models on our tasks are the larger, more powerful models, and thus, that ultimately we may expect the relatedness effect on event understanding to disappear eventually with scale.

However, a preliminary examination of individual model performance at the Possible-Atypical-Unrelated vs. Impossible-Atypical-Related event comparison (see Table 2) suggests that this is not the case. In fact, for some model families we see the opposite pattern: BLOOM 560M, 1B, and 1.7B are all better-performing than the 3B and 7B models, XGLM 564M out-performs other XGLM models, mGPT 1.3B out-performs mGPT 13B, and Gemma 1 models out-perform Gemma 2 models.

In this experiment, therefore, we take a systematic approach to investigating whether the performance of language models at the English tasks in Experiments 1 and 2 is correlated with their scale (i.e., number of parameters and number of training tokens). To do this, we turn to the Pythia models (Biderman et al., 2023), a suite of language models of different sizes (in terms of number of parameters) that are provided at various checkpoints over the course of training. We investigate the patterns in model performance based on different numbers of parameters and over the course of training.

7.2 Method

The datasets are the same as those in Experiments 1 and 2. We use the 14M, 31M, 70M, 160M, 410M, 1B, 1.4B, 2.8B, 6.9B, and 12B parameter Pythia models (Biderman et al., 2023), at training steps 0, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1000, 2000,

4000, 8000, 16000, 32000, 64000, 128000, and 143000 (the fully-trained model), where each training step comprises of ~ 2 M tokens.

7.3 Results

Figure 3 shows a striking difference between the Possible-Atypical-Unrelated vs. Impossible-Atypical-Related (*the cure for the disease was discovered by the guest vs. medication*) comparison and all other tasks. With the other tasks, there is a general pattern of improvement over the course of training, with the larger models showing numerically the best final scores. However, for the critical Possible-Atypical-Unrelated vs. Impossible-Atypical-Related comparison—where in Experiment 2, we saw at best chance accuracy across all models in both languages—performance never increases above chance.

7.4 Discussion

The effect of relatedness is not one that language models can fully ‘grow out of’. On the Possible-Atypical-Unrelated vs. Impossible-Atypical-Related task, there is little difference between larger and smaller models, and if anything, the performance of larger models deteriorates over the course of training.

8 General Discussion

Previous work has shown that the predictions of language models are highly correlated with event typicality (Kauf et al., 2023; Michaelov et al., 2024). In our study, we similarly find that language models are sensitive to large differences in typicality, for example when distinguishing between impossible events and typical events. However, we also find that they perform significantly worse when distinguishing between impossible events and less-typical but still possible events. Thus, our results suggest that there is ample room for potential improvement at this end of the scale.

A larger and perhaps more striking result is the effect of semantic relatedness. A language model will become worse at distinguishing between a possible and impossible event when the possible event involves a critical word that is semantically unrelated to its preceding context, when the impossible critical word is semantically related, and especially when both of these are the case.

Combined with previous work showing that all else being equal, language models will often tend

to predict that words that are semantically related to the context are more likely than unrelated ones (Misra et al., 2020; Michaelov and Bergen, 2022), these results suggest that language models at least in part rely on this relatedness (in addition to typicality) as a cue for which words are more probable.

The presence of such a heuristic would be in line with other work showing that the strong performance of language models and other artificial intelligence systems can often be explained by them learning simpler ‘shortcuts’ or other heuristics that correlate (often but not always) with the task at hand (see, e.g., Gururangan et al., 2018; McCoy et al., 2019, 2024; Abdou et al., 2020; Geirhos et al., 2020; Schramowski et al., 2020; Shah et al., 2020; Zhang et al., 2020; Du et al., 2021, 2022; Elazar et al., 2021; Kavumba et al., 2021; Ye and Kovashka, 2021; Stefanik, 2022).

As previously noted, such a surface-level heuristic is likely adaptive. Generally, words that are semantically related to their context are likely to be more appropriate sentence continuations than unrelated words—indeed, all the critical words in the typical sentences are semantically related to their context. This may at least in part explain why such a heuristic exists, not just in language models, as our study suggests, but also in humans (see, e.g., Chwilla et al., 2007; Parviz et al., 2011; Ettinger et al., 2016; Frank and Willems, 2017; Broderick et al., 2018; Uchida et al., 2021).

Crucially, this effect of contextual semantic relatedness does not disappear with increased model scale—in fact, while models trained on more data appear to become better at detecting larger differences in typicality, they also appear to increasingly rely on semantic relatedness to the context as a cue. Thus, this issue is not likely to disappear as language models continue to be trained on ever more data—in fact, it may get worse.

9 Conclusions

While contemporary language models can tell apart possible and impossible events (Jones et al., 2022; Kauf et al., 2023), we find that this breaks down when the possible events are not highly typical, and when semantic relatedness is not a cue to event possibility. This suggests that much of language models’ reported performance at reflecting event probability in previous work may be explained by language models’ sensitivity to differences in typicality and contextual semantic relatedness.

Limitations

Our study is limited in that it only looks at one type of impossibility—namely, animacy violations—but there are also many other ways in which a sentence could refer to an impossible event. Additionally, one possible issue with animacy-violating sentences, as Vega-Mendoza et al. (2021) note, is that they can be read as the beginning of a plausible sentence or figuratively as a plausible sentence. The former issue may be somewhat addressed by the fact that we add a period or full stop character at the end of each sentence (i.e., ‘.’ for English and ‘。’ for Mandarin). Additionally, Vega-Mendoza et al. (2021) explicitly address these issues by constructing their stimuli to avoid such interpretations, and further verify this by carrying out a norming study where experimental participants were asked to rate the plausibility of their sentences. They find that the impossible (i.e., animacy-violating sentences) were rated as significantly less plausible than possible sentences. Thus, we do not believe that this issue is likely to be of concern for the English stimuli. While Chow and Phillips (2013) do not report such a norming process for the Mandarin stimuli, the fact that we see almost indetical patterns in both the English and Mandarin stimuli suggests that this is unlikely to be a confound in their stimuli either.

A second limitation is that while we do expand our study beyond just English, the other language we include is Mandarin, which is another extremely widely-spoken high-resource language for which there are many high-quality language models. However, given that our study tests the limitations of such models, this is less of an issue—results of the kind that we find are more concerning for language models that otherwise show good performance and thus might be more likely to be erroneously trusted (for discussion, see, e.g., Bender et al., 2021; Raji et al., 2022).

Finally, we note that our dataset sizes are small. Each English-language task only includes 154 pairs, and each Mandarin task only includes 57 pairs. In principle, this means that the results may not be as robust as with a larger dataset. One way in which we address this is by including error bars in our figures, following the recommendation of the Association for Computational Linguistics Rolling Review Responsible Natural Language Processing Research checklist (i.e., the ARR Responsible NLP Research checklist; Carpuat et al., 2024; based

on Dodge et al., 2019; Beygelzimer et al., 2021; Rogers et al., 2021), which give an indication of the level of uncertainty in the results (for further discussion, see also, e.g., Reimers and Gurevych, 2017; Henderson et al., 2018; Forde and Paganini, 2019; Gundersen and Kjensmo, 2018; Marie et al., 2021; Gundersen et al., 2023; Kapoor et al., 2024; Biderman et al., 2024). Another is that the analyses reported in Experiment 3 on the English tasks are carried out at the item rather than task level, and thus also account for the number of experimental items.

Ethical Considerations

We do not believe that our study raises any ethical concerns. In fact, we hope that in demonstrating the limitations of contemporary language models, we will increase the extent to which readers are careful in their use of such technologies.

From an environmental perspective, our study is of minimal impact. All experiments were carried out on a computing cluster on NVIDIA A100 GPUs in under 2 GPU hours.

Our use of all the language models for research purposes falls within their terms of use and license agreements, as does our use of the Language Model Evaluation Harness. The experimental stimuli from Chow and Phillips (2013) and Vega-Mendoza et al. (2021) are scientific research materials, and as such, we believe that their use for scientific research falls under the category of fair use.

We provide all of our data, code and analyses at: https://osf.io/r6xns/?view_only=0567164a44f64530bde24c3bc5f1ddb

Acknowledgments

We would like to thank Wing-Yee Chow, Colin Phillips, Mariana Vega-Mendoza, Martin J. Pickering, Mante S. Nieuwland, Martin Paczynski, and Gina R. Kuperberg for making their experimental stimuli accessible. We would also like to thank the members of the Computational Psycholinguistics Laboratory at MIT and the Language and Cognition Laboratory at UCSD for their advice and discussion, in particular Cedegao E. Zhang, Roger Levy, and Catherine Arnett. James Michaelov was supported by a grant from the Andrew W. Mellon foundation (#2210-13947) during the writing of this paper.

References

- 01.AI, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin Wang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yanpeng Li, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. 2025. [Yi: Open Foundation Models by 01.AI](#). Preprint, arXiv:2403.04652.
- Mostafa Abdou, Vinit Ravishankar, Maria Barrett, Yonatan Belinkov, Desmond Elliott, and Anders Søgaard. 2020. [The Sensitivity of Language Models and Humans to Winograd Schema Perturbations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7590–7604, Online. Association for Computational Linguistics.
- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Leandro von Werra, and Thomas Wolf. 2024. [SmolLM - blazingly fast and remarkably powerful](#).
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?](#) 🐦. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pages 610–623, New York, NY, USA. Association for Computing Machinery.
- Alina Beygelzimer, Yann Dauphin, Percy Liang, and Jennifer Wortman Vaughan. 2021. [Introducing the NeurIPS 2021 Paper Checklist](#).
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, Usven Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. 2023. [Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling](#). In *Proceedings of the 40th International Conference on Machine Learning*, pages 2397–2430. PMLR.
- Stella Biderman, Hailey Schoelkopf, Lintang Sutawika, Leo Gao, Jonathan Tow, Baber Abbasi, Alham Fikri Aji, Pawan Sasanka Ammanamanchi, Sidney Black, Jordan Clive, Anthony DiPofi, Julien Etzaniz, Benjamin Fattori, Jessica Zosa Forde, Charles Foster, Jeffrey Hsu, Mimansa Jaiswal, Wilson Y. Lee, Haonan Li, Charles Lovering, Niklas Muennighoff, Ellie Pavlick, Jason Phang, Aviya Skowron, Samson Tan, Xiangru Tang, Kevin A. Wang, Genta Indra Winata, François Yvon, and Andy Zou. 2024. [Lessons from the Trenches on Reproducible Evaluation of Language Models](#). Preprint, arXiv:2405.14782.
- BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucicioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Vilanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Frohberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Al-mubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulqaila Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M. Sai-ful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwu, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat,

- Arjun Subramonian, Aurélie Névéol, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Na-joung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Uldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Daniel McDuff, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Onon-iwu, Habib Rezanejad, Hessie Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourier, Daniel León Perrián, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A. Castillo, Marianna Nezhurina, Mario Sängler, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S. Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sincee Sang-aaroonsiri, Sriшти Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2023. **BLOOM: A 176B-Parameter Open-Access Multilingual Language Model**. *Preprint*, arXiv:2211.05100.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. **PIQA: Reasoning about Physical Commonsense in Natural Language**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7432–7439.
- Michael P. Broderick, Andrew J. Anderson, Giovanni M. Di Liberto, Michael J. Crosse, and Edmund C. Lalor. 2018. **Electrophysiological Correlates of Semantic Dissimilarity Reflect the Comprehension of Natural, Narrative Speech**. *Current Biology*, 28(5):803–809.e3.
- Marine Carpuat, Marie-Catherine de Marneffe, Ivan Vladimir Meza Ruiz, Jesse Dodge, Margot Mieskes, ARR Editors-in-Chief, and ACL Ethics Committee. 2024. The ARR Responsible NLP Research checklist. <http://aclrollingreview.org/responsibleNLPresearch/>.
- Wing-Yee Chow and Colin Phillips. 2013. **No semantic illusions in the “Semantic P600” phenomenon: ERP evidence from Mandarin Chinese**. *Brain Research*, 1506:76–93.
- Dorothee J. Chwilla, Herman H. J. Kolk, and Constance T. W. M. Vissers. 2007. **Immediate integration of novel meanings: N400 support for an embodied view of language comprehension**. *Brain Research*, 1183:109–123.
- Ernest Davis and Gary Marcus. 2015. **Commonsense reasoning and commonsense knowledge in artificial intelligence**. *Communications of the ACM*, 58(9):92–103.
- Jesse Dodge, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A. Smith. 2019. **Show Your Work: Improved Reporting of Experimental Results**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2185–2194, Hong Kong, China. Association for Computational Linguistics.
- Arthur Conan Doyle. 1890. *The Sign Of Four*. Spencer Blackett, London.
- Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. 2022. **Shortcut Learning of Large Language Models in Natural Language Understanding: A Survey**. *Preprint*, arXiv:2208.11857.
- Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun, and Xia Hu. 2021. **Towards Interpreting and Mitigating Shortcut Learning Behavior of NLU models**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 915–929, Online. Association for Computational Linguistics.

- S. T. Dumais, G. W. Furnas, T. K. Landauer, S. Deerwester, and R. Harshman. 1988. [Using latent semantic analysis to improve access to textual information](#). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '88*, pages 281–285, Washington, D.C., United States. ACM Press.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhishava Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and Improving Consistency in Pretrained Language Models](#). *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Allyson Ettinger, Naomi Feldman, Philip Resnik, and Colin Phillips. 2016. Modeling N400 amplitude using vector space models of word representation. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, Philadelphia, USA.
- Evelina Fedorenko, Idan Asher Blank, Matthew Siegelman, and Zachary Mineroff. 2020. [Lack of selectivity for syntax relative to word meanings throughout the language network](#). *Cognition*, 203:104348.
- Jessica Zosa Forde and Michela Paganini. 2019. The Scientific Method in the Science of Machine Learning. In *Debugging Machine Learning Models Workshop at ICLR*.
- Stefan L. Frank and Roel M. Willems. 2017. [Word predictability and semantic similarity show distinct patterns of brain activity during language comprehension](#). *Language, Cognition and Neuroscience*, 32(9):1192–1203.
- Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2021. [A framework for few-shot language model evaluation](#). Zenodo.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut learning in deep neural networks](#). *Nature Machine Intelligence*, 2(11):665–673.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L. Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu-hui Chen, Zafarali Ahmed, Zhitaogong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024a. [Gemma: Open Models Based on Gemini Research and Technology](#). *Preprint*, arXiv:2403.08295.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshiev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju-yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay

- Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024b. [Gemma 2: Improving Open Language Models at a Practical Size](#). *Preprint*, arXiv:2408.00118.
- James J. Gibson. 1966. *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston.
- James J. Gibson. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston.
- Arthur M Glenberg and David A Robertson. 2000. [Symbol Grounding and Meaning: A Comparison of High-Dimensional and Embodied Theories of Meaning](#). *Journal of Memory and Language*, 43(3):379–401.
- H. P. Grice. 1989. *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass.
- H. Paul Grice. 1975. Logic and conversation. In Peter Cole and Jerry Morgan, editors, *Syntax and Semantics 3: Speech Acts*, pages 41–58.
- Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah Smith, and Hannaneh Hajishirzi. 2024. [OLMo: Accelerating the Science of Language Models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand. Association for Computational Linguistics.
- Odd Erik Gundersen, Kevin Coakley, Christine Kirkpatrick, and Yolanda Gil. 2023. [Sources of Irreproducibility in Machine Learning: A Review](#). *Preprint*, arXiv:2204.07610.
- Odd Erik Gundersen and Sigbjørn Kjensmo. 2018. [State of the Art: Reproducibility in Artificial Intelligence](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation Artifacts in Natural Language Inference Data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Michael Hanna, Yonatan Belinkov, and Sandro Pezzelle. 2023. [When Language Models Fall in Love: Animacy Processing in Transformer Language Models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12120–12135, Singapore. Association for Computational Linguistics.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. 2018. [Deep Reinforcement Learning That Matters](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katherine Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack William Rae, and Laurent Sifre. 2022. [Training Compute-Optimal Large Language Models](#). In *Advances in Neural Information Processing Systems*.
- Jennifer Hu and Michael Frank. 2024. [Auxiliary task demands mask the capabilities of smaller language models](#). In *First Conference on Language Modeling*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#). *Preprint*, arXiv:2310.06825.
- Cameron R. Jones, Tyler A. Chang, Seana Coulson, James A. Michaelov, Sean Trott, and Benjamin K. Bergen. 2022. [Distributional Semantics Still Can’t Account for Affordances](#). *Proceedings of the Annual Meeting of the Cognitive Science Society*, 44(44).
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray,

- Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling Laws for Neural Language Models](#). *Preprint*, arXiv:2001.08361.
- Sayash Kapoor, Emily M. Cantrell, Kenny Peng, Thanh Hien Pham, Christopher A. Bail, Odd Erik Gundersen, Jake M. Hofman, Jessica Hullman, Michael A. Lones, Momin M. Malik, Priyanka Nanayakkara, Russell A. Poldrack, Inioluwa Deborah Raji, Michael Roberts, Matthew J. Salganik, Marta Serra-Garcia, Brandon M. Stewart, Gilles Vandewiele, and Arvind Narayanan. 2024. [RE-FORMS: Consensus-based Recommendations for Machine-learning-based Science](#). *Science Advances*, 10(18):eadk3452.
- Carina Kauf, Anna A. Ivanova, Giulia Rambelli, Emanuele Chersoni, Jingyuan Selena She, Zawad Chowdhury, Evelina Fedorenko, and Alessandro Lenci. 2023. [Event Knowledge in Large Language Models: The Gap Between the Impossible and the Unlikely](#). *Cognitive Science*, 47(11):e13386.
- Pride Kavumba, Benjamin Heinzerling, Ana Brassard, and Kentaro Inui. 2021. [Learning to Learn to be Right for the Right Reasons](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3890–3898, Online. Association for Computational Linguistics.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. [Dense-Captioning Events in Videos](#). In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 706–715.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The Winograd Schema Challenge](#). In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*.
- Valentin Liévin, Christoffer Egeberg Hother, Andreas Geert Motzfeldt, and Ole Winther. 2024. [Can large language models reason about medical questions?](#) *Patterns*, 5(3).
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Nanam Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. 2022. [Few-shot Learning with Multilingual Generative Language Models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9019–9052, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the Ability of LSTMs to Learn Syntax-Sensitive Dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Llama Team. 2024. [The Llama 3 Herd of Models](#).
- Benjamin Marie, Atsushi Fujita, and Raphael Rubino. 2021. [Scientific Credibility of Machine Translation Research: A Meta-Evaluation of 769 Papers](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7297–7306, Online. Association for Computational Linguistics.
- 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.
- R. Thomas McCoy, Shunyu Yao, Dan Friedman, Mathew D. Hardy, and Thomas L. Griffiths. 2024. [Embers of autoregression show how large language models are shaped by the problem they are trained to solve](#). *Proceedings of the National Academy of Sciences*, 121(41):e2322420121.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Ross Metusalem, Marta Kutas, Thomas P. Urbach, Mary Hare, Ken McRae, and Jeffrey L. Elman. 2012. [Generalized event knowledge activation during online sentence comprehension](#). *Journal of Memory and Language*, 66(4):545–567.
- James A. Michaelov, Megan D. Bardolph, Cyma K. Van Petten, Benjamin K. Bergen, and Seana Coulson. 2024. [Strong Prediction: Language Model Surprisal Explains Multiple N400 Effects](#). *Neurobiology of Language*, 5(1):107–135.
- James A. Michaelov and Benjamin K. Bergen. 2022. [Collateral facilitation in humans and language models](#). In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 13–26, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- James A. Michaelov, Seana Coulson, and Benjamin Bergen. 2023. [Can Peanuts Fall in Love with Distributional Semantics?](#) In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 45.
- Kanishka Misra, Allyson Ettinger, and Julia Rayz. 2020. [Exploring BERT’s Sensitivity to Lexical Cues using Tests from Semantic Priming](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4625–4635, Online. Association for Computational Linguistics.
- Mante S. Nieuwland and Jos JA Van Berkum. 2006. [When peanuts fall in love: N400 evidence for the power of discourse](#). *Journal of cognitive neuroscience*, 18(7):1098–1111.

- Martin Paczynski and Gina R. Kuperberg. 2012. [Multiple influences of semantic memory on sentence processing: Distinct effects of semantic relatedness on violations of real-world event/state knowledge and animacy selection restrictions](#). *Journal of Memory and Language*, 67(4):426–448.
- Mehdi Parviz, Mark Johnson, Blake Johnson, and Jon Brock. 2011. [Using Language Models and Latent Semantic Analysis to Characterise the N400m Neural Response](#). In *Proceedings of the Australasian Language Technology Association Workshop 2011*, pages 38–46, Canberra, Australia.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 Technical Report](#). *Preprint*, arXiv:2412.15115.
- Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Maribeth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu, Erich Elsen, Siddhant Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsimppoukelli, Nikolai Grigorev, Doug Fritz, Thibault Sottiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d’Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake Hechtman, Laura Weidinger, Iason Gabriel, William Isaac, Ed Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu, and Geoffrey Irving. 2022. [Scaling Language Models: Methods, Analysis & Insights from Training Gopher](#). *Preprint*, arXiv:2112.11446.
- Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. [The Fallacy of AI Functionality](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’22, pages 959–972, New York, NY, USA. Association for Computing Machinery.
- Nils Reimers and Iryna Gurevych. 2017. [Reporting Score Distributions Makes a Difference: Performance Study of LSTM-networks for Sequence Tagging](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 338–348, Copenhagen, Denmark. Association for Computational Linguistics.
- Anna Rogers, Timothy Baldwin, and Kobi Leins. 2021. [‘Just What do You Think You’re Doing, Dave?’ A Checklist for Responsible Data Use in NLP](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4821–4833, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [WinoGrande: An Adversarial Winograd Schema Challenge at Scale](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8732–8740.
- Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger, Franziska Herbert, Xiaoting Shao, Hans-Georg Luigs, Anne-Katrin Mahlein, and Kristian Kersting. 2020. [Making deep neural networks right for the right scientific reasons by interacting with their explanations](#). *Nature Machine Intelligence*, 2(8):476–486.
- Harshay Shah, Kaustav Tamuly, Aditi Raghunathan, Prateek Jain, and Praneeth Netrapalli. 2020. [The Pitfalls of Simplicity Bias in Neural Networks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9573–9585. Curran Associates, Inc.
- Oleh Shliakhko, Alena Fenogenova, Maria Tikhonova, Anastasia Kozlova, Vladislav Mikhailov, and Tatiana Shavrina. 2024. [mGPT: Few-Shot Learners Go Multilingual](#). *Transactions of the Association for Computational Linguistics*, 12:58–79.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou, Kevin Clark, Stephen R. Pfohl, Heather Cole-Lewis, Darlene Neal, Qazi Mamunur Rashid, Mike Schaeckermann, Amy Wang, Dev Dash, Jonathan H. Chen, Nigam H. Shah, Sami Lachgar, Philip Andrew Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agüera y Arcas, Nenad Tomašev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle K. Barral, Dale R. Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. 2025. [Toward expert-level medical question answering with large language models](#). *Nature Medicine*, pages 1–8.
- Dan Sperber. 1986. *Relevance: Communication and Cognition*. Language and Thought Series. Harvard University Press, Cambridge, Massachusetts.
- Michal Stefanik. 2022. [Methods for Estimating and Improving Robustness of Language Models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational*

- Linguistics: Human Language Technologies: Student Research Workshop*, pages 44–51, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Shane Storks, Qiaozi Gao, and Joyce Y. Chai. 2020. [Recent Advances in Natural Language Inference: A Survey of Benchmarks, Resources, and Approaches](#). *Preprint*, arXiv:1904.01172.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Takahisa Uchida, Nicolas Lair, Hiroshi Ishiguro, and Peter Ford Dominey. 2021. [A Model of Online Temporal-Spatial Integration for Immediacy and Overrule in Discourse Comprehension](#). *Neurobiology of Language*, 2(1):83–105.
- Thomas P. Urbach and Marta Kutas. 2010. [Quantifiers more or less quantify on-line: ERP evidence for partial incremental interpretation](#). *Journal of Memory and Language*, 63(2):158–179.
- Dave Van Veen, Cara Van Uden, Louis Blanke-meier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerová, Nidhi Rohatgi, Poonam Hosamani, William Collins, Neera Ahuja, Curtis P. Langlotz, Jason Hom, Sergios Gatidis, John Pauly, and Akshay S. Chaudhari. 2024. [Adapted large language models can outperform medical experts in clinical text summarization](#). *Nature Medicine*, 30(4):1134–1142.
- Mariana Vega-Mendoza, Martin J. Pickering, and Mante S. Nieuwland. 2021. [Concurrent use of animacy and event-knowledge during comprehension: Evidence from event-related potentials](#). *Neuropsychologia*, 152:107724.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananeey, 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.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural Network Acceptability Judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent Abilities of Large Language Models](#). *Transactions on Machine Learning Research*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 Technical Report](#). *Preprint*, arXiv:2407.10671.
- Keren Ye and Adriana Kovashka. 2021. [A Case Study of the Shortcut Effects in Visual Commonsense Reasoning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(4):3181–3189.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A Large-Scale Adversarial Dataset for Grounded Commonsense Inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a Machine Really Finish Your Sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Hongming Zhang, Xinran Zhao, and Yangqiu Song. 2020. [WinoWhy: A Deep Diagnosis of Essential Commonsense Knowledge for Answering Winograd Schema Challenge](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5736–5745, Online. Association for Computational Linguistics.

A Model Scores

We include the scores for all models. [Table 2](#) shows the scores on all English-language tasks from Experiments 1 and 2 for all language models used in the analysis. [Table 3](#) shows the scores on all Mandarin tasks from Experiments 1 and 2 for all language models used in the analysis. The remaining tables ([Table 4–Table 13](#)) show the scores on all English-language tasks from Experiment 3 on all checkpoints (0, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, and 143000) of all the Pythia models (14M, 31M, 70M, 160M, 410M, 1B, 1.4B, 2.8B, 6.9B, and 12B) tested.

Model	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
	vs.	vs.	vs.	vs.	vs.	vs.	vs.	vs.
	IAR	IAU	PAR	PAU	IAR	IAU	IAR	IAU
01-ai/Yi-1.5-6B	0.858	0.955	0.800	0.942	0.690	0.871	0.277	0.684
01-ai/Yi-1.5-9B	0.961	0.994	0.877	0.987	0.710	0.948	0.226	0.761
01-ai/Yi-6B	0.961	0.994	0.852	0.994	0.742	0.935	0.265	0.748
01-ai/Yi-9B	0.948	0.994	0.871	0.987	0.729	0.961	0.310	0.723
HuggingFaceTB/SmolLM-1.7B	0.884	0.987	0.832	0.981	0.710	0.987	0.277	0.703
HuggingFaceTB/SmolLM-135M	0.903	0.987	0.748	0.948	0.787	0.974	0.355	0.774
HuggingFaceTB/SmolLM-360M	0.923	0.974	0.819	0.987	0.768	0.974	0.252	0.729
Qwen/Qwen2-0.5B	0.935	1.000	0.839	1.000	0.761	0.981	0.335	0.800
Qwen/Qwen2-1.5B	0.942	1.000	0.832	0.994	0.761	0.961	0.348	0.813
Qwen/Qwen2-7B	0.942	1.000	0.877	1.000	0.723	0.942	0.361	0.787
Qwen/Qwen2.5-0.5B	0.935	0.994	0.800	0.981	0.781	0.987	0.381	0.839
Qwen/Qwen2.5-1.5B	0.955	0.987	0.852	0.994	0.755	0.968	0.342	0.787
Qwen/Qwen2.5-14B	0.974	1.000	0.903	0.994	0.794	0.948	0.387	0.774
Qwen/Qwen2.5-3B	0.916	0.968	0.781	0.935	0.729	0.910	0.342	0.768
Qwen/Qwen2.5-7B	0.955	1.000	0.890	0.994	0.761	0.955	0.400	0.781
ai-forever/mGPT	0.865	0.987	0.710	0.935	0.755	0.942	0.368	0.774
ai-forever/mGPT-13B	0.923	0.994	0.800	0.968	0.716	0.955	0.297	0.703
allenai/OLMo-2-1124-13B	0.961	1.000	0.903	0.981	0.768	0.948	0.394	0.800
allenai/OLMo-2-1124-7B	0.955	1.000	0.890	0.981	0.723	0.935	0.348	0.806
bigscience/bloom-1b1	0.929	1.000	0.710	0.961	0.774	0.955	0.419	0.826
bigscience/bloom-1b7	0.910	0.994	0.768	0.942	0.781	0.955	0.426	0.794
bigscience/bloom-3b	0.903	0.987	0.800	0.974	0.755	0.961	0.355	0.748
bigscience/bloom-560m	0.890	0.987	0.735	0.935	0.781	0.942	0.413	0.845
bigscience/bloom-7b1	0.942	0.994	0.800	0.981	0.768	0.955	0.381	0.794
facebook/xglm-1.7B	0.916	0.987	0.755	0.974	0.761	0.974	0.297	0.748
facebook/xglm-2.9B	0.910	0.994	0.774	0.987	0.742	0.981	0.277	0.768
facebook/xglm-4.5B	0.903	0.994	0.781	0.981	0.729	0.955	0.303	0.787
facebook/xglm-564M	0.884	0.994	0.703	0.968	0.761	0.974	0.329	0.794
facebook/xglm-7.5B	0.903	0.994	0.794	0.974	0.716	0.955	0.245	0.761
google/gemma-2-2b	0.929	1.000	0.852	0.994	0.697	0.961	0.181	0.761
google/gemma-2-9b	0.935	1.000	0.865	1.000	0.729	0.974	0.168	0.748
google/gemma-2b	0.942	1.000	0.826	0.981	0.742	0.948	0.400	0.794
google/gemma-7b	0.923	0.987	0.884	0.961	0.742	0.948	0.413	0.826
meta-llama/Llama-3.1-8B	0.981	1.000	0.903	1.000	0.794	0.974	0.310	0.794
meta-llama/Llama-3.2-1B	0.942	1.000	0.826	0.994	0.710	0.955	0.277	0.781
meta-llama/Llama-3.2-3B	0.942	1.000	0.871	0.994	0.755	0.955	0.297	0.813
meta-llama/Meta-Llama-3-8B	0.981	1.000	0.884	1.000	0.761	0.955	0.310	0.794
mistralai/Mistral-7B-v0.1	0.955	0.994	0.903	0.994	0.742	0.948	0.284	0.781
mistralai/Mistral-7B-v0.3	0.948	1.000	0.890	0.994	0.742	0.942	0.290	0.774
mistralai/Mistral-Nemo-Base-2407	0.955	1.000	0.897	0.987	0.755	0.942	0.310	0.787

Table 2: Scores on the English tasks for all models used in Experiments 1 and 2.

Model	PTR vs. IAR	PTR vs. IAU	PAU vs. IAR	PAU vs. IAU
01-ai/Yi-1.5-6B	0.759	0.914	0.500	0.690
01-ai/Yi-1.5-9B	0.897	0.983	0.362	0.621
01-ai/Yi-6B	0.845	0.966	0.414	0.586
01-ai/Yi-9B	0.879	0.966	0.362	0.603
Qwen/Qwen2-0.5B	0.776	0.931	0.414	0.569
Qwen/Qwen2-1.5B	0.862	0.931	0.310	0.500
Qwen/Qwen2-7B	0.897	0.948	0.345	0.552
Qwen/Qwen2.5-0.5B	0.793	0.914	0.328	0.569
Qwen/Qwen2.5-1.5B	0.862	0.897	0.345	0.638
Qwen/Qwen2.5-14B	0.828	0.931	0.362	0.500
Qwen/Qwen2.5-3B	0.707	0.776	0.483	0.638
Qwen/Qwen2.5-7B	0.862	0.931	0.293	0.517
ai-forever/mGPT	0.655	0.862	0.293	0.534
ai-forever/mGPT-13B	0.741	0.966	0.328	0.500
bigscience/bloom-1b1	0.828	0.931	0.293	0.552
bigscience/bloom-1b7	0.810	0.931	0.224	0.517
bigscience/bloom-3b	0.828	0.914	0.259	0.569
bigscience/bloom-560m	0.793	0.914	0.345	0.534
bigscience/bloom-7b1	0.862	0.966	0.310	0.517
facebook/xglm-1.7B	0.810	0.948	0.276	0.534
facebook/xglm-2.9B	0.776	0.931	0.310	0.552
facebook/xglm-4.5B	0.793	0.914	0.345	0.466
facebook/xglm-564M	0.759	0.862	0.328	0.483
facebook/xglm-7.5B	0.741	0.931	0.293	0.466
mistralai/Mistral-Nemo-Base-2407	0.862	0.948	0.379	0.483

Table 3: Scores on the Mandarin tasks for all models used in Experiments 1 and 2. Note that we only ran these analyses on language models that were reported to be trained on Mandarin.

Model	Step	PTR vs. IAR	PTR vs. IAU	PTR vs. PAR	PTR vs. PAU	PAR vs. IAR	PAR vs. IAU	PAU vs. IAR	PAU vs. IAU
EleutherAI/pythia-14m	0	0.452	0.419	0.535	0.484	0.419	0.477	0.445	0.406
EleutherAI/pythia-14m	1	0.452	0.419	0.535	0.484	0.419	0.477	0.445	0.406
EleutherAI/pythia-14m	2	0.452	0.419	0.548	0.484	0.419	0.477	0.445	0.400
EleutherAI/pythia-14m	4	0.452	0.419	0.548	0.484	0.426	0.477	0.445	0.394
EleutherAI/pythia-14m	8	0.452	0.413	0.542	0.490	0.413	0.471	0.458	0.394
EleutherAI/pythia-14m	16	0.445	0.419	0.529	0.490	0.452	0.465	0.452	0.426
EleutherAI/pythia-14m	32	0.445	0.406	0.484	0.497	0.465	0.452	0.477	0.439
EleutherAI/pythia-14m	64	0.439	0.374	0.458	0.452	0.484	0.484	0.452	0.471
EleutherAI/pythia-14m	128	0.439	0.432	0.484	0.426	0.503	0.452	0.452	0.503
EleutherAI/pythia-14m	256	0.381	0.426	0.452	0.458	0.452	0.439	0.445	0.484
EleutherAI/pythia-14m	512	0.355	0.374	0.477	0.490	0.413	0.419	0.381	0.445
EleutherAI/pythia-14m	1000	0.426	0.490	0.510	0.548	0.452	0.535	0.355	0.439
EleutherAI/pythia-14m	2000	0.581	0.703	0.523	0.723	0.606	0.742	0.361	0.606
EleutherAI/pythia-14m	4000	0.613	0.806	0.548	0.755	0.555	0.781	0.329	0.613
EleutherAI/pythia-14m	8000	0.632	0.787	0.542	0.723	0.619	0.755	0.387	0.645
EleutherAI/pythia-14m	16000	0.658	0.787	0.542	0.716	0.632	0.806	0.413	0.677
EleutherAI/pythia-14m	32000	0.652	0.755	0.542	0.697	0.606	0.768	0.387	0.671
EleutherAI/pythia-14m	64000	0.600	0.723	0.555	0.703	0.574	0.755	0.374	0.594
EleutherAI/pythia-14m	128000	0.581	0.774	0.535	0.755	0.587	0.735	0.342	0.568
EleutherAI/pythia-14m	143000	0.606	0.755	0.548	0.735	0.561	0.742	0.297	0.561

Table 4: Pythia 14M scores on the English-language tasks.

Model	Step	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
		vs. IAR	vs. IAU	vs. PAR	vs. PAU	vs. IAR	vs. IAU	vs. IAR	vs. IAU
EleutherAI/pythia-31m	0	0.419	0.406	0.465	0.465	0.426	0.452	0.503	0.445
EleutherAI/pythia-31m	1	0.419	0.406	0.465	0.465	0.426	0.452	0.503	0.445
EleutherAI/pythia-31m	2	0.413	0.406	0.458	0.458	0.426	0.452	0.497	0.445
EleutherAI/pythia-31m	4	0.413	0.413	0.458	0.465	0.426	0.465	0.490	0.445
EleutherAI/pythia-31m	8	0.413	0.406	0.471	0.465	0.426	0.452	0.503	0.452
EleutherAI/pythia-31m	16	0.419	0.426	0.477	0.452	0.452	0.471	0.471	0.452
EleutherAI/pythia-31m	32	0.432	0.432	0.471	0.484	0.465	0.458	0.439	0.510
EleutherAI/pythia-31m	64	0.387	0.413	0.497	0.465	0.477	0.477	0.413	0.490
EleutherAI/pythia-31m	128	0.387	0.400	0.484	0.452	0.406	0.413	0.413	0.406
EleutherAI/pythia-31m	256	0.406	0.413	0.503	0.477	0.426	0.400	0.426	0.445
EleutherAI/pythia-31m	512	0.413	0.432	0.503	0.497	0.452	0.471	0.426	0.452
EleutherAI/pythia-31m	1000	0.484	0.652	0.484	0.652	0.490	0.671	0.355	0.516
EleutherAI/pythia-31m	2000	0.574	0.723	0.535	0.761	0.516	0.723	0.303	0.587
EleutherAI/pythia-31m	4000	0.671	0.832	0.529	0.832	0.600	0.832	0.290	0.690
EleutherAI/pythia-31m	8000	0.742	0.897	0.600	0.832	0.645	0.845	0.329	0.697
EleutherAI/pythia-31m	16000	0.697	0.852	0.587	0.839	0.600	0.813	0.297	0.600
EleutherAI/pythia-31m	32000	0.645	0.852	0.535	0.800	0.652	0.832	0.323	0.671
EleutherAI/pythia-31m	64000	0.703	0.890	0.548	0.845	0.581	0.819	0.329	0.619
EleutherAI/pythia-31m	128000	0.742	0.884	0.574	0.806	0.632	0.839	0.374	0.639
EleutherAI/pythia-31m	143000	0.716	0.890	0.581	0.806	0.619	0.832	0.381	0.639

Table 5: Pythia 31M scores on the English-language tasks.

Model	Step	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
		vs. IAR	vs. IAU	vs. PAR	vs. PAU	vs. IAR	vs. IAU	vs. IAR	vs. IAU
EleutherAI/pythia-70m	0	0.445	0.426	0.477	0.477	0.529	0.471	0.432	0.516
EleutherAI/pythia-70m	1	0.445	0.426	0.477	0.477	0.529	0.471	0.432	0.516
EleutherAI/pythia-70m	2	0.445	0.426	0.465	0.477	0.535	0.471	0.439	0.529
EleutherAI/pythia-70m	4	0.452	0.419	0.452	0.477	0.529	0.471	0.426	0.516
EleutherAI/pythia-70m	8	0.445	0.413	0.465	0.477	0.523	0.471	0.432	0.516
EleutherAI/pythia-70m	16	0.394	0.413	0.477	0.452	0.510	0.477	0.432	0.490
EleutherAI/pythia-70m	32	0.381	0.413	0.465	0.445	0.477	0.452	0.419	0.503
EleutherAI/pythia-70m	64	0.400	0.439	0.477	0.445	0.452	0.406	0.439	0.484
EleutherAI/pythia-70m	128	0.406	0.394	0.458	0.445	0.439	0.439	0.458	0.471
EleutherAI/pythia-70m	256	0.374	0.387	0.503	0.471	0.400	0.394	0.426	0.419
EleutherAI/pythia-70m	512	0.432	0.477	0.529	0.497	0.452	0.542	0.413	0.458
EleutherAI/pythia-70m	1000	0.568	0.684	0.497	0.690	0.535	0.723	0.394	0.574
EleutherAI/pythia-70m	2000	0.671	0.832	0.600	0.800	0.581	0.806	0.413	0.645
EleutherAI/pythia-70m	4000	0.781	0.923	0.594	0.826	0.639	0.858	0.419	0.684
EleutherAI/pythia-70m	8000	0.755	0.935	0.652	0.832	0.671	0.865	0.394	0.710
EleutherAI/pythia-70m	16000	0.774	0.942	0.619	0.839	0.677	0.890	0.394	0.735
EleutherAI/pythia-70m	32000	0.703	0.903	0.581	0.845	0.639	0.865	0.290	0.677
EleutherAI/pythia-70m	64000	0.806	0.968	0.632	0.884	0.671	0.897	0.329	0.710
EleutherAI/pythia-70m	128000	0.781	0.955	0.594	0.871	0.690	0.923	0.439	0.761
EleutherAI/pythia-70m	143000	0.729	0.923	0.632	0.832	0.645	0.858	0.381	0.710

Table 6: Pythia 70M scores on the English-language tasks.

Model	Step	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
		vs. IAR	vs. IAU	vs. PAR	vs. PAU	vs. IAR	vs. IAU	vs. IAR	vs. IAU
EleutherAI/pythia-160m	0	0.426	0.368	0.471	0.503	0.445	0.413	0.445	0.471
EleutherAI/pythia-160m	1	0.426	0.368	0.471	0.503	0.445	0.413	0.445	0.471
EleutherAI/pythia-160m	2	0.426	0.368	0.477	0.503	0.458	0.413	0.452	0.465
EleutherAI/pythia-160m	4	0.426	0.381	0.471	0.503	0.452	0.406	0.452	0.484
EleutherAI/pythia-160m	8	0.426	0.387	0.471	0.503	0.452	0.445	0.439	0.471
EleutherAI/pythia-160m	16	0.406	0.394	0.484	0.490	0.490	0.465	0.439	0.445
EleutherAI/pythia-160m	32	0.406	0.374	0.477	0.471	0.419	0.465	0.471	0.426
EleutherAI/pythia-160m	64	0.413	0.394	0.458	0.490	0.439	0.426	0.458	0.439
EleutherAI/pythia-160m	128	0.413	0.419	0.490	0.497	0.445	0.413	0.445	0.439
EleutherAI/pythia-160m	256	0.406	0.413	0.503	0.484	0.406	0.387	0.426	0.413
EleutherAI/pythia-160m	512	0.471	0.497	0.523	0.535	0.503	0.555	0.400	0.510
EleutherAI/pythia-160m	1000	0.619	0.761	0.516	0.690	0.606	0.774	0.374	0.645
EleutherAI/pythia-160m	2000	0.781	0.916	0.600	0.826	0.671	0.865	0.432	0.742
EleutherAI/pythia-160m	4000	0.774	0.955	0.568	0.871	0.703	0.903	0.355	0.690
EleutherAI/pythia-160m	8000	0.839	0.981	0.645	0.890	0.729	0.942	0.348	0.729
EleutherAI/pythia-160m	16000	0.845	0.974	0.697	0.935	0.723	0.942	0.368	0.794
EleutherAI/pythia-160m	32000	0.877	0.974	0.665	0.935	0.742	0.955	0.348	0.755
EleutherAI/pythia-160m	64000	0.865	0.974	0.697	0.935	0.723	0.916	0.297	0.710
EleutherAI/pythia-160m	128000	0.871	0.981	0.658	0.923	0.729	0.910	0.361	0.748
EleutherAI/pythia-160m	143000	0.819	0.987	0.684	0.923	0.723	0.910	0.329	0.716

Table 7: Pythia 160M scores on the English-language tasks.

Model	Step	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
		vs. IAR	vs. IAU	vs. PAR	vs. PAU	vs. IAR	vs. IAU	vs. IAR	vs. IAU
EleutherAI/pythia-410m	0	0.471	0.445	0.471	0.490	0.490	0.426	0.471	0.465
EleutherAI/pythia-410m	1	0.471	0.445	0.471	0.490	0.490	0.426	0.471	0.465
EleutherAI/pythia-410m	2	0.465	0.445	0.465	0.497	0.484	0.419	0.471	0.471
EleutherAI/pythia-410m	4	0.445	0.439	0.458	0.484	0.477	0.432	0.465	0.465
EleutherAI/pythia-410m	8	0.452	0.413	0.445	0.477	0.471	0.432	0.458	0.452
EleutherAI/pythia-410m	16	0.445	0.394	0.445	0.452	0.471	0.432	0.432	0.477
EleutherAI/pythia-410m	32	0.432	0.400	0.445	0.445	0.477	0.490	0.452	0.497
EleutherAI/pythia-410m	64	0.452	0.432	0.445	0.465	0.490	0.426	0.471	0.497
EleutherAI/pythia-410m	128	0.419	0.406	0.484	0.445	0.452	0.394	0.394	0.465
EleutherAI/pythia-410m	256	0.406	0.387	0.503	0.458	0.419	0.406	0.426	0.413
EleutherAI/pythia-410m	512	0.432	0.452	0.497	0.510	0.477	0.490	0.452	0.452
EleutherAI/pythia-410m	1000	0.561	0.684	0.523	0.716	0.555	0.761	0.348	0.581
EleutherAI/pythia-410m	2000	0.735	0.910	0.581	0.839	0.632	0.884	0.303	0.690
EleutherAI/pythia-410m	4000	0.800	0.955	0.658	0.897	0.684	0.935	0.297	0.690
EleutherAI/pythia-410m	8000	0.826	0.981	0.716	0.923	0.742	0.923	0.342	0.735
EleutherAI/pythia-410m	16000	0.890	0.994	0.729	0.968	0.742	0.935	0.252	0.761
EleutherAI/pythia-410m	32000	0.858	1.000	0.723	0.994	0.742	0.961	0.335	0.800
EleutherAI/pythia-410m	64000	0.865	0.994	0.742	0.974	0.748	0.981	0.342	0.819
EleutherAI/pythia-410m	128000	0.890	0.994	0.774	0.968	0.697	0.961	0.271	0.774
EleutherAI/pythia-410m	143000	0.877	1.000	0.774	0.981	0.703	0.961	0.284	0.742

Table 8: Pythia 410M scores on the English-language tasks.

Model	Step	PTR vs. IAR	PTR vs. IAU	PTR vs. PAR	PTR vs. PAU	PAR vs. IAR	PAR vs. IAU	PAU vs. IAR	PAU vs. IAU
EleutherAI/pythia-1b	0	0.503	0.426	0.490	0.477	0.484	0.477	0.471	0.516
EleutherAI/pythia-1b	1	0.503	0.426	0.490	0.477	0.484	0.477	0.471	0.516
EleutherAI/pythia-1b	2	0.510	0.426	0.490	0.477	0.490	0.477	0.471	0.516
EleutherAI/pythia-1b	4	0.503	0.439	0.497	0.484	0.484	0.477	0.465	0.510
EleutherAI/pythia-1b	8	0.465	0.432	0.490	0.477	0.484	0.458	0.477	0.523
EleutherAI/pythia-1b	16	0.439	0.452	0.510	0.477	0.535	0.465	0.484	0.561
EleutherAI/pythia-1b	32	0.445	0.432	0.471	0.477	0.484	0.452	0.439	0.510
EleutherAI/pythia-1b	64	0.432	0.439	0.477	0.490	0.452	0.439	0.432	0.497
EleutherAI/pythia-1b	128	0.406	0.406	0.477	0.452	0.432	0.406	0.419	0.419
EleutherAI/pythia-1b	256	0.439	0.374	0.497	0.471	0.432	0.432	0.452	0.432
EleutherAI/pythia-1b	512	0.458	0.490	0.516	0.523	0.458	0.561	0.406	0.529
EleutherAI/pythia-1b	1000	0.619	0.826	0.535	0.735	0.587	0.800	0.368	0.581
EleutherAI/pythia-1b	2000	0.794	0.948	0.600	0.858	0.710	0.916	0.406	0.723
EleutherAI/pythia-1b	4000	0.865	0.974	0.684	0.916	0.716	0.916	0.361	0.710
EleutherAI/pythia-1b	8000	0.858	1.000	0.723	0.929	0.755	0.942	0.374	0.729
EleutherAI/pythia-1b	16000	0.871	0.994	0.716	0.955	0.755	0.961	0.361	0.748
EleutherAI/pythia-1b	32000	0.897	0.987	0.723	0.968	0.729	0.948	0.284	0.742
EleutherAI/pythia-1b	64000	0.916	1.000	0.710	0.968	0.748	0.968	0.323	0.800
EleutherAI/pythia-1b	128000	0.903	1.000	0.774	0.981	0.781	0.961	0.335	0.774
EleutherAI/pythia-1b	143000	0.903	1.000	0.794	0.981	0.723	0.974	0.335	0.781

Table 9: Pythia 1B scores on the English-language tasks.

Model	Step	PTR vs. IAR	PTR vs. IAU	PTR vs. PAR	PTR vs. PAU	PAR vs. IAR	PAR vs. IAU	PAU vs. IAR	PAU vs. IAU
EleutherAI/pythia-1.4b	0	0.432	0.419	0.503	0.452	0.484	0.497	0.458	0.465
EleutherAI/pythia-1.4b	1	0.432	0.419	0.503	0.452	0.484	0.497	0.458	0.465
EleutherAI/pythia-1.4b	2	0.439	0.419	0.503	0.445	0.484	0.497	0.471	0.465
EleutherAI/pythia-1.4b	4	0.432	0.426	0.484	0.445	0.471	0.503	0.458	0.452
EleutherAI/pythia-1.4b	8	0.439	0.406	0.510	0.445	0.452	0.510	0.452	0.458
EleutherAI/pythia-1.4b	16	0.432	0.342	0.484	0.439	0.452	0.497	0.426	0.477
EleutherAI/pythia-1.4b	32	0.432	0.355	0.497	0.458	0.413	0.497	0.413	0.471
EleutherAI/pythia-1.4b	64	0.413	0.394	0.503	0.490	0.400	0.445	0.406	0.400
EleutherAI/pythia-1.4b	128	0.400	0.400	0.471	0.465	0.394	0.413	0.426	0.419
EleutherAI/pythia-1.4b	256	0.387	0.387	0.490	0.452	0.432	0.432	0.445	0.439
EleutherAI/pythia-1.4b	512	0.452	0.484	0.452	0.529	0.477	0.516	0.426	0.503
EleutherAI/pythia-1.4b	1000	0.665	0.781	0.548	0.716	0.606	0.742	0.426	0.652
EleutherAI/pythia-1.4b	2000	0.761	0.942	0.581	0.839	0.735	0.903	0.406	0.742
EleutherAI/pythia-1.4b	4000	0.845	0.981	0.690	0.923	0.761	0.942	0.368	0.806
EleutherAI/pythia-1.4b	8000	0.877	0.987	0.710	0.961	0.735	0.935	0.316	0.729
EleutherAI/pythia-1.4b	16000	0.903	0.994	0.710	0.968	0.768	0.955	0.323	0.768
EleutherAI/pythia-1.4b	32000	0.884	1.000	0.742	0.981	0.723	0.974	0.323	0.768
EleutherAI/pythia-1.4b	64000	0.910	0.994	0.781	0.981	0.768	0.961	0.303	0.800
EleutherAI/pythia-1.4b	128000	0.890	1.000	0.787	0.994	0.723	0.961	0.323	0.781
EleutherAI/pythia-1.4b	143000	0.897	1.000	0.781	0.994	0.690	0.961	0.252	0.774

Table 10: Pythia 1.4B scores on the English-language tasks.

Model	Step	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
		vs. IAR	vs. IAU	vs. PAR	vs. PAU	vs. IAR	vs. IAU	vs. IAR	vs. IAU
EleutherAI/pythia-2.8b	0	0.439	0.419	0.490	0.497	0.439	0.426	0.426	0.445
EleutherAI/pythia-2.8b	1	0.439	0.419	0.490	0.497	0.439	0.426	0.426	0.445
EleutherAI/pythia-2.8b	2	0.439	0.432	0.484	0.503	0.439	0.432	0.426	0.439
EleutherAI/pythia-2.8b	4	0.452	0.452	0.484	0.497	0.419	0.419	0.406	0.439
EleutherAI/pythia-2.8b	8	0.432	0.426	0.490	0.503	0.406	0.426	0.406	0.400
EleutherAI/pythia-2.8b	16	0.394	0.439	0.477	0.477	0.452	0.445	0.439	0.452
EleutherAI/pythia-2.8b	32	0.484	0.458	0.510	0.503	0.419	0.413	0.477	0.394
EleutherAI/pythia-2.8b	64	0.419	0.426	0.497	0.497	0.400	0.413	0.426	0.413
EleutherAI/pythia-2.8b	128	0.400	0.406	0.503	0.471	0.406	0.381	0.406	0.374
EleutherAI/pythia-2.8b	256	0.387	0.368	0.490	0.471	0.413	0.439	0.439	0.439
EleutherAI/pythia-2.8b	512	0.477	0.523	0.484	0.516	0.484	0.542	0.445	0.555
EleutherAI/pythia-2.8b	1000	0.529	0.716	0.490	0.735	0.529	0.729	0.303	0.548
EleutherAI/pythia-2.8b	2000	0.819	0.948	0.600	0.884	0.729	0.916	0.297	0.716
EleutherAI/pythia-2.8b	4000	0.832	0.987	0.690	0.935	0.748	0.955	0.316	0.761
EleutherAI/pythia-2.8b	8000	0.871	0.994	0.761	0.968	0.781	0.955	0.277	0.742
EleutherAI/pythia-2.8b	16000	0.903	0.994	0.748	0.974	0.768	0.968	0.355	0.800
EleutherAI/pythia-2.8b	32000	0.910	1.000	0.774	0.994	0.748	0.961	0.323	0.761
EleutherAI/pythia-2.8b	64000	0.923	0.994	0.768	0.994	0.755	0.968	0.348	0.794
EleutherAI/pythia-2.8b	128000	0.935	0.994	0.781	0.994	0.735	0.955	0.355	0.774
EleutherAI/pythia-2.8b	143000	0.929	0.994	0.794	0.994	0.703	0.955	0.323	0.768

Table 11: Pythia 2.8B scores on the English-language tasks.

Model	Step	PTR	PTR	PTR	PTR	PAR	PAR	PAU	PAU
		vs. IAR	vs. IAU	vs. PAR	vs. PAU	vs. IAR	vs. IAU	vs. IAR	vs. IAU
EleutherAI/pythia-6.9b	0	0.445	0.419	0.477	0.477	0.490	0.477	0.465	0.465
EleutherAI/pythia-6.9b	1	0.445	0.419	0.477	0.477	0.490	0.477	0.465	0.465
EleutherAI/pythia-6.9b	2	0.445	0.413	0.477	0.477	0.490	0.477	0.465	0.465
EleutherAI/pythia-6.9b	4	0.445	0.419	0.484	0.477	0.484	0.477	0.458	0.471
EleutherAI/pythia-6.9b	8	0.413	0.394	0.465	0.452	0.477	0.471	0.458	0.477
EleutherAI/pythia-6.9b	16	0.400	0.406	0.484	0.465	0.497	0.432	0.445	0.497
EleutherAI/pythia-6.9b	32	0.387	0.445	0.445	0.413	0.471	0.484	0.477	0.477
EleutherAI/pythia-6.9b	64	0.458	0.439	0.503	0.432	0.458	0.452	0.426	0.471
EleutherAI/pythia-6.9b	128	0.445	0.400	0.523	0.445	0.419	0.406	0.439	0.426
EleutherAI/pythia-6.9b	256	0.419	0.413	0.465	0.477	0.445	0.419	0.445	0.413
EleutherAI/pythia-6.9b	512	0.400	0.387	0.465	0.471	0.432	0.452	0.445	0.432
EleutherAI/pythia-6.9b	1000	0.574	0.690	0.510	0.652	0.542	0.690	0.419	0.619
EleutherAI/pythia-6.9b	2000	0.761	0.948	0.581	0.852	0.703	0.903	0.374	0.723
EleutherAI/pythia-6.9b	4000	0.826	0.981	0.671	0.923	0.742	0.948	0.316	0.742
EleutherAI/pythia-6.9b	8000	0.858	0.981	0.710	0.968	0.735	0.942	0.290	0.755
EleutherAI/pythia-6.9b	16000	0.916	1.000	0.768	0.981	0.723	0.942	0.290	0.768
EleutherAI/pythia-6.9b	32000	0.923	1.000	0.787	0.987	0.735	0.961	0.381	0.781
EleutherAI/pythia-6.9b	64000	0.923	1.000	0.748	0.987	0.742	0.968	0.335	0.819
EleutherAI/pythia-6.9b	128000	0.942	0.994	0.768	0.994	0.716	0.968	0.310	0.755
EleutherAI/pythia-6.9b	143000	0.929	0.994	0.813	0.994	0.710	0.961	0.284	0.748

Table 12: Pythia 6.9B scores on the English-language tasks.

Model	Step	PTR vs. IAR	PTR vs. IAU	PTR vs. PAR	PTR vs. PAU	PAR vs. IAR	PAR vs. IAU	PAU vs. IAR	PAU vs. IAU
EleutherAI/pythia-12b	0	0.426	0.439	0.471	0.490	0.419	0.426	0.458	0.458
EleutherAI/pythia-12b	1	0.426	0.439	0.471	0.490	0.419	0.426	0.458	0.458
EleutherAI/pythia-12b	2	0.419	0.439	0.471	0.490	0.426	0.426	0.458	0.458
EleutherAI/pythia-12b	4	0.426	0.432	0.477	0.484	0.419	0.432	0.458	0.458
EleutherAI/pythia-12b	8	0.419	0.400	0.452	0.510	0.419	0.432	0.413	0.419
EleutherAI/pythia-12b	16	0.400	0.419	0.471	0.490	0.445	0.439	0.419	0.452
EleutherAI/pythia-12b	32	0.413	0.426	0.471	0.458	0.452	0.465	0.394	0.465
EleutherAI/pythia-12b	64	0.413	0.406	0.445	0.484	0.439	0.503	0.458	0.471
EleutherAI/pythia-12b	128	0.432	0.348	0.452	0.477	0.426	0.426	0.419	0.445
EleutherAI/pythia-12b	256	0.426	0.419	0.490	0.439	0.471	0.452	0.458	0.452
EleutherAI/pythia-12b	512	0.458	0.426	0.516	0.510	0.445	0.510	0.413	0.445
EleutherAI/pythia-12b	1000	0.574	0.632	0.484	0.665	0.529	0.671	0.368	0.510
EleutherAI/pythia-12b	2000	0.761	0.942	0.606	0.884	0.677	0.897	0.368	0.716
EleutherAI/pythia-12b	4000	0.858	0.994	0.684	0.916	0.774	0.942	0.355	0.806
EleutherAI/pythia-12b	8000	0.877	0.994	0.748	0.961	0.716	0.948	0.252	0.768
EleutherAI/pythia-12b	16000	0.897	0.994	0.787	0.987	0.729	0.948	0.277	0.761
EleutherAI/pythia-12b	32000	0.929	0.994	0.755	0.987	0.723	0.968	0.290	0.761
EleutherAI/pythia-12b	64000	0.923	0.987	0.781	0.987	0.716	0.955	0.265	0.774
EleutherAI/pythia-12b	128000	0.923	1.000	0.787	0.994	0.742	0.942	0.265	0.729
EleutherAI/pythia-12b	143000	0.923	1.000	0.774	0.981	0.735	0.948	0.290	0.735

Table 13: Pythia 12B scores on the English-language tasks.