

# COMPS: Conceptual Minimal Pair Sentences for testing Robust Property Knowledge and its Inheritance in Pre-trained Language Models

**Kanishka Misra**  
Purdue University  
kmisra@purdue.edu

**Julia Rayz**  
Purdue University  
jtaylor1@purdue.edu

**Allyson Ettinger**  
University of Chicago  
aettinger@uchicago.edu

## Abstract

A characteristic feature of human semantic cognition is its ability to not only store and retrieve the properties of concepts observed through experience, but to also facilitate the inheritance of properties (*can breathe*) from superordinate concepts (ANIMAL) to their subordinates (DOG)—i.e. demonstrate *property inheritance*. In this paper, we present COMPS, a collection of English minimal pair sentences that jointly tests pre-trained language models (PLMs) on their ability to attribute properties to concepts and their ability to demonstrate property inheritance behavior. Analyses of 22 different PLMs on COMPS reveal that they can easily distinguish between concepts on the basis of a property when they are trivially different, but find it relatively difficult when concepts are related on the basis of nuanced knowledge representations. Furthermore, we find that PLMs can show behaviors suggesting successful property inheritance in simple contexts, but fail in the presence of distracting information, which decreases the performance of many models sometimes even below chance. This lack of robustness in demonstrating simple reasoning raises important questions about PLMs’ capacity to make correct inferences even when they appear to possess the prerequisite knowledge.

## 1 Introduction

The ability to learn, update and deploy one’s knowledge about concepts (ROBIN, CHAIR) and their properties (*can fly*, *can be sat on*), observed during everyday experience is fundamental to human semantic cognition (Murphy, 2002; Rogers and McClelland, 2004; Rips et al., 2012). Knowledge of a concept’s properties, combined with the ability to infer the ISA relation (Sloman, 1998; Murphy, 2003) leads to an important behavior known as *property inheritance* (Quillian, 1967; Smith and Estes, 1978; Murphy, 2002), where subordinates of a concept inherit its properties. For instance, one is likely to infer that an entity called *luna* can

meow, has a tail, is a mammal, etc., even if *all* they know is that it is a cat. The close connection between a word’s meaning and its conceptual representation makes these abilities crucial to language understanding (Murphy, 2002; Lake and Murphy, 2021), making it critical for computational models of language processing to also exhibit behavior consistent with these capacities. Indeed, modern pre-trained language models (PLMs; Devlin et al., 2019; Brown et al., 2020, etc.) have made impressive empirical strides in eliciting general knowledge about real world concepts and entities (Petroni et al., 2019; Weir et al., 2020, *i.a.*), as well as in demonstrating isomorphism with real world abstractions like direction and color (Abdou et al., 2021; Patel and Pavlick, 2022), often times without even having been explicitly trained to do so. At the same time, their ability to robustly demonstrate such capacities has recently been called to question, owing to failures due to reporting bias (Gordon and Van Durme, 2013; Shwartz and Choi, 2020), lack of consistency (Elazar et al., 2021; Ravichander et al., 2020), and sensitivity to lexical cues (Kassner and Schütze, 2020; Misra et al., 2020; Pandia and Ettinger, 2021).

In this work, we cast further light on PLMs’ ability to robustly demonstrate knowledge about concepts and their properties. To this end, we introduce Conceptual Minimal Pair Sentences (COMPS), a collection of English minimal pair sentences, where each pair attributes a property (*can fly*) to two noun concepts: one which actually possesses the property (ROBIN), and one which does not (PENGUIN). Following standard practice in the minimal pairs evaluation paradigm (Warstadt et al., 2020, etc.), we test whether PLMs prefer sentence stimuli expressing correct property knowledge over those expressing incorrect ones. COMPS can be decomposed into three subsets, each containing stimuli that progressively isolate deeper understanding of the task of attributing properties to concepts,

by adding controls for more superficial heuristics. Our first subset—COMPS-BASE—measures the extent to which PLMs attribute properties to the right concepts, while varying the similarity of the positive (ROBIN) and the negative concepts (PENGUIN [high] vs. TABLE [low]). This controls for the possibility that models are relying on coarse-grained concept distinctions. For instance, in this setup a model should prefer (1a) over both versions of (1b).

- (1) a. A **robin** can fly.
- b. \*A (**penguin/table**) can fly.

Next, drawing on the phenomenon of property inheritance, the COMPS-WUGS set introduces a novel concept, WUG, expressed as the subordinate of the positive and negative concepts from a subset of the COMPS-BASE set, and tests the extent to which PLMs successfully attribute it the given property when it is associated with the positive concept. This increases the complexity of the reasoning task, as well as the distance between the associated concept (ROBIN) and property (*can fly*). These manipulations help to control for memorization of the literal phrases being tested, forcing models to judge properties for a novel concept that inherits the property from a known concept. In this task, given that a model successfully prefers (1a) over (1b), it should also prefer (2a) over (2b):

- (2) a. A wug is a **robin**. Therefore, a wug can fly.
- b. \*A wug is a **penguin**. Therefore, a wug can fly.

The final subset—COMPS-WUGS-DIST, combines the aforementioned controls by using negative concepts as distracting content and inserting them into the COMPS-WUGS stimuli. Specifically, we transform the stimuli of COMPS-WUGS by creating two subordinates for every minimal pair; one for the positive concept (ROBIN, subordinate: WUG) and the other for the negative concept (PENGUIN, subordinate: DAX), which acts as a distractor. This way, we control for the possibility that models may be relying on simple word associations between content words—of which there are only two in the prior tests—by introducing additional, irrelevant but contentful words into the context. Here, we consider models to be correct if they prefer (3a) over (3b), given that they prefer (1a) over (1b):

- (3) a. A **wug** is a robin. A **dax** is a penguin. Therefore, a **wug** can fly.
- b. \*A **wug** is a robin. A **dax** is a penguin. Therefore, a **dax** can fly.

Together, the three sets of stimuli tease apart more superficial predictive behaviors, such as contextual word associations, from more robust reasoning behaviors based on understanding of concept properties. While we can expect superficial predictive strategies to be brittle in the face of shallow perturbations and irrelevant distractions, robust property knowledge and reasoning behaviors should not.

We use COMPS to analyze robust property knowledge and its inheritance in 22 different PLMs, ranging from small masked language models to billion-parameter autoregressive language models. In our experiments with COMPS-BASE, we find PLMs to demonstrate strong performance in attributing properties to the correct concepts in our minimal pairs. However, we observe this strong performance largely when the concepts in the minimal pairs are trivially different (e.g., LION and TEA for the property *is a mammal*). When the concept pairs are similar (on the basis of different knowledge representations), we find models' performance to degrade substantially, by as much as 25 points. We observe a similar trend in our analyses on COMPS-WUGS—models first appear to show desirable behavior, potentially indicating proficiency in the more complex property inheritance reasoning. However, their overall performance declines drastically when investigated in the presence of distractors (i.e., on COMPS-WUGS-DIST). This failure is particularly pronounced in larger autoregressive PLMs, whose performance in fact drops below chance in cases where distracting information is proximal to the queried property, indicating the presence of a proximity effect. Together, our findings highlight brittleness of PLMs with conceptual knowledge and reasoning, as evidenced by failures in the face of simple controls. We make our code and data available at: <https://github.com/kanishkamisra/comps>.

## 2 Conceptual Minimal Pair Sentences (COMPS)

### 2.1 Connections to prior work

Prior work in exploring property knowledge in PLMs has adopted two different paradigms: one which uses probing classifiers to test if the applicability of a property can be decoded from the representations of LMs (Forbes et al., 2019; Da and Kasai, 2019; Derby et al., 2021); and the other which uses cloze-testing, in which LMs are tasked to fill in the blank in prompts that describe specific prop-

erties/factual knowledge about the world (Petroni et al., 2019; Weir et al., 2020). We argue that both approaches—though insightful—have key limitations for evaluating property knowledge, and that minimal pair testing overcomes these limitations to a beneficial extent.

Apart from ongoing debates surrounding the validity of probing classifiers (see Hewitt and Liang, 2019; Ravichander et al., 2021; Belinkov, 2022), the probing setup does not allow the testing of property knowledge in a precise manner. Specifically, several properties are often perfectly correlated in datasets such as the one we use here (see §2.2). For example, the property of being an animal and being able to breathe and grow, etc., are all perfectly correlated with one another. Even if the model’s true knowledge of these properties is highly variable, probing its representations for them would yield the exact same result, leading to conclusions that overestimate the model’s capacity for some properties, while underestimating for others. Evaluation using minimal pair sentences overcomes this limitation by allowing us to explicitly represent the properties of interest in language form, thereby allowing precise testing of property knowledge.

Similarly, standard cloze-testing of PLMs (Petroni et al., 2019; Weir et al., 2020; Jiang et al., 2021) also faces multiple limitations. First, it does not allow for testing of multi-word expressions, as by definition, it involves prediction of a single word/token. Second, it does not yield faithful conclusions about one-to-many or many-to-many relations: e.g. the cloze prompts “Ravens can \_\_\_\_.” and “\_\_\_\_ can fly.” do not have a single correct answer. This makes our conclusions about models’ knowledge contingent on choice of one correct completion over the other. The minimal pair evaluation paradigm overcomes these issues by generalizing the cloze-testing method to multi-word expressions—by focusing on entire sentences—and at the same time, pairing every prompt with a negative instance. This allows for a straightforward way to assess correctness: the choice between multiple correct completions is transformed into one between correct and incorrect, at the cost of having several different instances (pairs) for testing knowledge of the same property. Additionally, the minimal pairs paradigm allows us also to shed light on how the nature of negative samples affects model behavior, which has been missing in approaches using probing and cloze-testing. The us-

age of minimal pairs is a well-established practice in the literature, having been widely used in works that analyze syntactic knowledge of LMs (Marvin and Linzen, 2018; Futrell et al., 2019; Warstadt et al., 2020). We complement this growing literature by introducing minimal-pair testing to the study of conceptual knowledge in PLMs.

Our property inheritance analyses closely relate to the ‘Leap-of-Thought’ (LoT) framework of Talmor et al. (2020). In particular, LoT holds the taxonomic relations between concepts *implicit* and tests whether models can abstract over them to make property inferences—e.g., testing the extent to which models assign *Whales have bellybuttons* the ‘True’ label, given that *Mammals have bellybuttons* (with the implicit knowledge here being *Whales are mammals*). With COMPS-WUGS (and COMPS-WUGS-DIST), we instead explicitly provide the relevant taxonomic knowledge in the context and target whether PLMs can behave consistently with knowledge they have already demonstrated (in the base case, COMPS-BASE) and attribute the property in question to the correct subordinate concept. This also relates to recent work that measures consistency of PLMs’ word prediction capacities in eliciting factual knowledge (Elazar et al., 2021; Ravichander et al., 2020).

## 2.2 Ground-truth Property Knowledge data

For our ground-truth property knowledge resource, we use a subset of the CSLB property norms collected by Devereux et al. (2014), which was further extended by Misra et al. (2022). The original dataset was constructed by asking 123 human participants to generate properties for 638 everyday concepts. Contemporary work has used this dataset by taking as positive instances all concepts for which a property was generated, while taking the rest as negative instances (Lucy and Gauthier, 2017; Da and Kasai, 2019, etc.) for each property. While this dataset has been popularly used in related literature, Misra et al. (2022) recently discovered striking gaps in coverage among the properties included in the dataset.<sup>1</sup> For example, the property *can breathe* was only generated for 6 out of 152 animal concepts, despite being applicable for all of them—as a result, contemporary work can be expected to have wrongfully penalized models that attributed this property to animals that could

<sup>1</sup>See also Sommerauer and Fokkens (2018) and Sommerauer (2022), who also discuss this limitation.

indeed breathe, and similarly for other properties. To remedy this issue, Misra et al. (2022) manually extended CSLB’s coverage for 521 concepts and 3,645 properties. We refer to this extended CSLB dataset as XCSLB, and we use it as our source for ground-truth property knowledge.

### 2.3 Choosing negative samples

We rely on a diverse set of knowledge representation sources to construct negative samples for COMPS. Each source has a unique representational structure which gives rise to different pairwise similarity metrics, on the basis of which we pick out negative samples for each property:

**Taxonomy** We consider a hierarchical organization of our concepts, by taking a subset of WordNet (Miller, 1995) consisting of our 521 concepts. We use the wup similarity (Wu and Palmer, 1994) as our choice of taxonomic similarity.

**Property Norms** We use the XCSLB dataset and organize it as a matrix whose rows indicate concepts and columns indicate properties that are either present (indicated as 1) or absent (indicated as 0) for each concept. As our similarity measure, we consider the jaccard similarity between the row vectors of concepts. This reflects the overlap in properties between concepts, and is prevalent in studies utilizing conceptual similarity in cognitive science (Tversky, 1977; Sloman, 1993, etc.).

**Co-occurrence** We use the co-occurrence between concept words as an unstructured knowledge representation. For quantifying similarity, we use the cosine similarity of the GloVe vectors (Pennington et al., 2014) of our concept words.

**Sampling Strategy** Each property ( $p_i$ ) in our dataset splits the set of concepts into two: a set of concepts that possess the property ( $Q_{p_i}$ ), and a set of concepts that do not ( $\neg Q_{p_i}$ ). We sample  $\min(|Q_{p_i}|, 10)$ —i.e., at most 10—concepts from  $Q_{p_i}$  and take them to be our positive set. Then for each concept in the positive set, we sample from  $\neg Q_{p_i}$  the concept that is most similar (depending on the source) to the positive concept and take it as a negative concept for the property. We additionally include a negative concept that is randomly sampled from  $\neg Q_{p_i}$ , leaving out the concepts sampled on the basis of the three previously described knowledge sources. Examples of the four types of negative samples for the concept ZEBRA and the property *has striped patterns* are shown in Table 1.

Knowledge Rep.	Negative Concept	Similarity
Taxonomy	HORSE	0.88
Property Norms	DEER	0.63
Co-occurrence	GIRAFFE	0.75
Random	BAT	-

Table 1: Negatively sampled concepts selected on the basis of various knowledge representational mechanisms, where the property is *has striped patterns*, and the positive concept is ZEBRA.

### 2.4 Minimal Pair Construction

Following our negative sample generation process, we end up with total of 49,280 pairs of positive and negative concepts that span across 3,645 properties (14 pairs per property, on average). Every property is associated with a property phrase—a verb phrase which expresses the property in English, as provided in XCSLB. Using these materials, we construct our three datasets of minimal pair sentence stimuli, examples of which are shown in Figure 1.


**COMPS-BASE** The COMPS-BASE dataset contains minimal pair sentences that follow the template: “[DET] [CONCEPT] [property-phrase].” where [DET] is an optional determiner, and [CONCEPT] is the noun concept. Applying this template to our generated pairs results in 49,280 instances. See Figure 1a for an example.

**COMPS-WUGS** We test property inheritance in PLMs using only the animal kingdom subset of COMPS-BASE (152 concepts, 944 properties, and 13,888 pairs), keeping the same negative samples. We convert the original minimal pair sentences in COMPS-BASE, in which the positive concept is an animal, into pairs of two-sentence stimuli by first introducing a new concept (WUG) to be the subordinate of the concepts in the original minimal pair. We then express its property inheritance in a separate sentence. Our two sentence stimuli follow the template: “A wug is a [CONCEPT]. Therefore, a wug [property-phrase].” Although we use *wug* as our running example for the subordinate concept, we use four different nonsense words {*wug*, *dax*, *blicket*, *lep*} equal numbers of times, to avoid making spurious conclusions based on a single nonsense word.<sup>2</sup> Introducing an intervening novel concept allows us to robustly control for simple word-level associations between concepts and properties that models might have picked up during

<sup>2</sup>As we describe in §4, we also tried a different set of nonce words, to address concerns about possible impacts of using nonce words from existing literature (e.g., *wug*).

**Property:** *can fly*  
**Positive:** ROBIN  
**Negative:** PENGUIN  
**Subordinate:** WUG  
**COMPS-BASE:** A (robin/penguin) can fly.  
**COMPS-WUGS:** A wug is a (robin/penguin).  
Therefore, a wug can fly.

(a) Instances of COMPS-BASE and COMPS-WUGS.



(b) Distraction scheme for stimuli in COMPS-WUGS-DIST, where the distractor is inserted either **before** or **in-between** each COMPS-WUGS stimulus.

Figure 1: Examples of materials used in our experiments. In this example, ROBIN is the positive concept.

training. Figure 1a shows an example.

**COMPS-WUGS-DIST** To add distracting information, we follow Pandia and Ettinger (2021) and convert the COMPS-WUGS stimuli by associating a different subordinate concept (DAX) with the negative concept ([NEG-CONCEPT]), and inserting it **before** or **in-between** the sentence containing the positive concept and its subordinate, separately. This results in two subsets (**before** and **in-between**) of three-sentence minimal pair stimuli, which differ in the subordinate to which the property is attributed. We use the following template to create our stimuli: “A **wug** is a [CONCEPT]. A **dax** is a [NEG-CONCEPT]. Therefore, a (**wug/dax**) [property-phrase].” That is, we have stimuli that resemble COMPS-WUGS but instead deal with a pair of competing subordinate concepts in context.<sup>3</sup> See Figure 1b for an example.

### 3 Methodology

#### 3.1 Models Investigated

We investigate property knowledge and property inheritance capacities of 22 different PLMs, belonging to six different families. We evaluate four widely used masked language modeling (MLM) families: (1) ALBERT (Lan et al., 2020), (2) BERT (Devlin et al., 2019), (3) ELECTRA (Clark et al., 2020), and (4) RoBERTa (Liu et al., 2019); as well as two auto-regressive language modeling families: (1) GPT2 (Radford et al., 2019), and (2) the GPT-Neo (Black et al., 2021) and GPT-J models (Wang and Komatsuzaki, 2021) from EleutherAI. We also use distilled versions of BERT-base, RoBERTa-base, and GPT2, trained using the method described by Sanh et al. (2019). We list each model’s parameters, vocabulary size, and training corpora in Table 3 (Appendix A).

<sup>3</sup>We again choose from our list of four nonsense words (*wug*, *dax*, *blicket*, and *fep*), which amounts to 12 unique ordered pairs, after accounting for counterbalancing.

#### 3.2 Measuring Performance

To evaluate models on COMPS, we compare their log-probabilities for the property phrase—conditioned on contexts (to the left) containing the positive and negative noun concepts. That is, we hold the property phrase constant, and compare across minimally differing conditions to evaluate the probability with which a property is attributed to each concept. For example, we score stimuli in COMPS-BASE, e.g., “A *dog can bark.*” as:

$$\log p(\text{can bark.} \mid \text{A dog}),$$

its corresponding stimulus in COMPS-WUGS, “A *wug is a dog. Therefore, a wug can bark.*” as:

$$\log p(\text{can bark.} \mid \text{A wug is a dog. Therefore, a wug}),$$

and similarly—assuming CAT as the negative concept—the corresponding stimuli in our COMPS-WUGS-DIST subset, “A *wug is a dog. A dax is a cat. Therefore, a wug can bark.*” as:<sup>4</sup>

$$\log p(\text{can bark.} \mid \text{A wug is a dog. A dax is a cat. Therefore, a wug}).$$

This approach to eliciting conditional LM judgments is equivalent to the “scoring by premise” method (Holtzman et al., 2021), which has been shown to result in stable comparisons across items. Additionally, this also takes into account the potential noise due to frequency effects or tokenization differences (Misra et al., 2021). Estimating these conditional log-probabilities using auto-regressive PLMs can be directly computed in a left-to-right manner. For MLMs, we use their conditional pseudo-loglikelihoods (Salazar et al., 2020) as a proxy for conditional log-probabilities.

Based on this simple method of eliciting relative acceptability measures from PLMs, we evaluate

<sup>4</sup>Here we show an example where the distractor is added **in-between** the context specifying the positive concept, and the queried property knowledge.




Figure 2: Accuracies of PLMs on COMPS-BASE under various negative sampling schemes. Chance performance for all rows is 50%, except for ‘Overall,’ where it is 6.25%. Refer to Table 3 for unabbreviated model names.

a model’s accuracy on all COMPS stimuli as the percentage of times its log-probability for a property is greater when conditioned on the context that attributes the property to the positive—as opposed to the negative—concept. Since all cases are forced-choice tasks between two instances, chance performance is set to 50%. Table 4 (Appendix B) shows examples of all COMPS stimuli and GPT-J’s conditional log-probabilities for them.

## 4 Experiments and Analyses

### 4.1 Base property knowledge of PLMs and their sensitivity to similarity effects

We begin by evaluating the 22 PLMs on COMPS-BASE. Here we focus on the extent to which models robustly associate properties to the correct concepts across stimuli with varying kinds of similarity between the positive and negative concepts. We report accuracies of the 22 PLMs on COMPS-BASE across the four different negative sampling schemes that we specified in §2.3. We additionally report a more stringent accuracy measure that we refer to as ‘Overall accuracy,’ which is calculated for every property and its positive concept, as the percentage of times a model correctly attributes the property to the positive concept in **all four types of negative sampling schemes**. Chance performance for only the ‘Overall’ case is then 6.25% ( $0.5^4 \times 100$ ). Figure 2 shows these results.

From Figure 2, we see that models strongly distinguish between positive and negative concepts in cases where they are dramatically different—i.e., where negative concepts were sampled randomly (e.g., BEAR [positive] vs BOTTLE [negative] for the property *can breathe*). However performance drops substantially when there are subtler

differences between the two concepts—e.g., the concepts WALRUS (positive) and SHARK (negative) for the property *is a mammal*. For instance, the best performing model in any similarity-based negative sampling scheme (GPT-J, 76%, ‘Co-oc’) only slightly outperforms the worst model in the random negative sampling scheme (Neo-125M, 71%). The performance of PLMs is not substantially different across the three similarity-based negative sampling schemes, suggesting that the dynamics of model sensitivity in attributing properties to concepts are largely harmonized across various types of similarities. As a result of models’ insensitivity in presence of similar negative concepts, the overall accuracies are very modest in value, with the overall accuracy of the best performing model (GPT-J) being only 53%. This overall performance is, however, significantly above chance (6.25%). We discuss additional findings, such as performance by property type and model size, in Appendix C, since they are incidental to the main conclusions of this analysis.

### 4.2 Property inheritance in PLMs

Having established the base property knowledge of PLMs, we now investigate the extent to which they can show behavior that is consistent with reasoning required to handle property inheritance. We first investigate their performance on COMPS-WUGS, created using the subset of COMPS-BASE containing only animal concepts (see §2.4 for stimulus construction). Table 2 shows average accuracies obtained by PLMs on our property inheritance stimuli, and compares them to average accuracies on COMPS-BASE—aggregating across all negative sampling schemes. Recall that the stimuli in COMPS-WUGS present a more challenging property attribution task than in COMPS-BASE, by

COMPS subset	Size	Acc.
BASE	49.3K	68.4 <sub>1.7</sub>
BASE ( <i>animal kingdom only</i> )	13.8K	67.1 <sub>2.0</sub>
WUGS	13.8K	68.9 <sub>2.3</sub>
WUGS-DIST ( <b>before</b> )	13.8K	59.2 <sub>3.9</sub>
WUGS-DIST ( <b>in-between</b> )	13.8K	47.2 <sub>4.5</sub>

Table 2: Average accuracy (and standard error of the mean) of PLMs ( $N = 22$ ) on each of our COMPS subsets. Chance performance is 50% throughout.

not only controlling for coarse-grained similarity effects, but also introducing an intervening novel concept that is expected to inherit the properties of the positive concept. By measuring attribution of properties more indirectly, these stimuli increase the complexity of the reasoning and control for memorization of the literal phrase initially tested with COMPS-BASE.

Table 2 shows the average accuracy of the PLMs on each subset of COMPS. Despite the increase in complexity, we see that PLMs actually show slightly stronger performance on COMPS-WUGS (68.9%) than on COMPS-BASE (67.1%). This means that there are instances in which models prefer the property in the positive context over the negative context (4a > 4b), but show the opposite behavior in COMPS-BASE (4d > 4c).

- (4) a. A wug is a **robin**. Therefore, a wug can fly.  
 b. A wug is a **penguin**. Therefore, a wug can fly.  
 c. A **robin** can fly.  
 d. A **penguin** can fly.

This pattern of performance could lead to spurious conclusions that models are successfully executing property inheritance, when in fact they show a lack of the pre-requisite property knowledge based on their failure on COMPS-BASE. We will discuss these inconsistencies in more detail below. Overall, however, the relatively strong performance on COMPS-WUGS suggests that models are largely unaffected when we control for simple memorization of tested phrases—e.g., *robin can fly*—by linking known concepts to properties through an intervening subordinate concept (*wug*). This suggests that models are not relying on simple memorization, but does not control for the possibility of simple association between content words (*robin* and *fly*)—for this we turn to COMPS-WUGS-DIST.

The COMPS-WUGS-DIST test assesses whether models retain strong property attribution performance when content words in the context are not all relevant for the property prediction. The stimuli thus include irrelevant distractor concepts and their




Figure 3: Distribution of model performance on COMPS-WUGS and COMPS-WUGS-DIST (both subsets) across possible outcomes (correct = ✓, incorrect = ✗) of the models on corresponding minimal pairs in COMPS-BASE. Error bars indicate 95% CI, while dashed line indicates chance performance (50%).

subordinates—which, in a robust model, should not affect attribution of the property to the correct concept (see §2.4 for stimulus construction).

From Table 2, the average accuracies of PLMs on both subsets of COMPS-WUGS-DIST (**before** and **in-between**) indicate that overall, models now show clear degradation in property inheritance performance as a result of the distracting information. Specifically, the PLMs’ performance drops by 9.7 points on instances when the distracting information is added **before** the relevant context and queried property, and by 21.7 points on instances where it is added **in-between** the two, relative to the undistracted property inheritance stimuli (COMPS-WUGS). Notably, the latter drop in performance brings models level with chance accuracy (we fail to reject the null hypothesis that avg. accuracy of models is 50%;  $p = .62$ , Wilcoxon signed rank exact test), highlighting a pronounced lack of robustness in PLMs’ capacity to attribute properties to the correct concepts in their input context.

**Accounting for spurious performance** The COMPS-WUGS results above raise the concern that models are often showing spurious performance: accurately demonstrating property inheritance behavior without actually possessing the right property knowledge. To shed more light on this potential issue, we plot the distribution of model accuracies on our property inheritance stimuli (COMPS-WUGS and COMPS-WUGS-DIST) divided based on their outcomes on the corresponding stimuli in COMPS-BASE. Figure 3 shows these distributions. In COMPS-WUGS and both subsets of COMPS-WUGS-DIST, models show this spurious correct behavior on 41.3%, 55.6%, and 42.8% of instances in which they produce incorrect judgments on the corresponding COMPS-BASE stimuli (**yellow** bars




Figure 4: Accuracies of individual models (grouped by family, in increasing order based on number of parameters) on COMPS-WUGS and COMPS-WUGS-DIST. Black dashed line indicates chance performance (50%). Refer to Table 3 for unabbreviated model names. Error bands indicate 95% Bootstrap CIs.

in Figure 3). This non-trivial proportion of cases with spurious performance further reinforces the idea that PLMs’ successful predictions on these tests are likely relying on heuristics rather than robust inferences about property knowledge. We can remove the effects of these spurious instances by filtering to items in which models give the correct answer on COMPS-BASE (blue bars in Figure 3)—though we see that the overall conclusions remain the same after this filtering.

**On the pronounced effect of proximity in autoregressive PLMs** Our previous discussion summarized the aggregate property inheritance behavior of the 22 PLMs we considered—we now zoom in for a model-wise analysis. Figure 4 shows models’ relative accuracies on COMPS-WUGS and COMPS-WUGS-DIST, filtering to items with correct COMPS-BASE performance, as in the blue bars of Figure 3. Consistent with our overall findings, we observe distracting content to substantially degrade model performance across the board.<sup>5</sup> A particularly noteworthy pattern is that the degradation in autoregressive PLM families—GPT2 and EleutherAI—shows a stark sensitivity to *proximity effects*. While these classes of model seem to suffer less when distracting content is added **before** the context containing the positive concept (thus placing the distraction farther from the queried property), they show substantially worse performance when the opposite is the case (i.e., when distraction is added **in-between**, and is therefore closer to the queried property). This degradation due to proximity of the distracting content becomes **catastrophically worse as models grow larger in the number of**

**pre-trained parameters**—in fact bringing their performance down to as much as 26.2 points **below chance** (in GPT-J, which has 6B parameters). While MLMs also show similar levels of degraded performance in presence of distraction, they do not seem to show any systematic sensitivity to proximity effects, likely due to their bidirectional nature.

**Results on GPT-3** In addition to our main experiments, we also evaluate GPT-3 (Brown et al., 2020) models on a small subset of COMPS stimuli (denoted as miniCOMPS). Results from this analysis (shown in Appendix C.1) are largely aligned with our main conclusions, with all GPT-3 models—including the largest one (175B parameters)—performing worse than chance on the **in-between** subset of miniCOMPS-WUGS-DIST, while performing substantially better on miniCOMPS-WUGS and miniCOMPS-WUGS-DIST (**before**). Together with our main results, this indicates that scaling alone may be insufficient to elicit robust inferences about concepts and their properties.

**Choice of nonce words** Nonce words constitute an important design decision for our stimuli—we followed precedents in language acquisition research (Berko, 1958; Gopnik and Sobel, 2000, *i.a.*) and used previously existing nonce words (such as *wug* and *blicket*) to represent novel concepts in context. While these are expected to be novel for humans, they may appear in pre-training corpora on which PLMs are usually trained.<sup>6</sup> This raises a potential concern that PLMs could already be biased toward certain properties for these words (e.g., *wug* is commonly depicted as a bird), and

<sup>5</sup>See also Pandia and Ettinger (2021) for a similar degradation of performance on cloze-tasks involving factual retrieval.

<sup>6</sup>e.g., *wug* appears in wikipedia: [https://en.wikipedia.org/wiki/Jean\\_Berko\\_Gleason](https://en.wikipedia.org/wiki/Jean_Berko_Gleason) (accessed on Jan 23)



may struggle to associate them with different properties.<sup>7</sup> To explore this empirically, we conducted experiments with alternative nonce words (generated synthetically, similar to Kim et al. (2022); see Appendix C.2). Figure 7 (Appendix C.2) shows results on COMPS-WUGS and COMPS-WUGS-DIST with randomly sampled nonce words. We see that the new results are comparable to those in Figure 4, with models showing the same preference on both stimulus versions 80% of the time on average. This suggests that the choice of nonce words is not producing any noteworthy bias.

**Framing of novel taxonomic information** Another relevant stimulus design decision is the phrasing for introducing novel concepts in context. While we used “A wug is a [CONCEPT]” for our main experiments, we additionally tested with an alternate framing: “A wug is a type of [CONCEPT].” From Figure 9 (Appendix C.3), we again see that the overall patterns of results are comparable to the original results, with models showing the same preference across both versions of the stimuli on COMPS-WUGS and COMPS-WUGS-DIST 90% of the time, on average.

## 5 General Discussion and Conclusion

The overall goal of COMPS is to shed light on the extent to which PLMs can robustly (1) attribute to real world concepts (e.g., HORSE, WHALE) their correct properties (e.g., *is a mammal*); and (2) demonstrate behavior consistent with *property inheritance*: a reasoning process in which concepts are endowed with the properties of their superordinates (Smith and Estes, 1978; Sloman, 1998; Murphy, 2002). Testing PLMs for these abilities allows us to ask key questions about how they encode and transfer knowledge. To target these capabilities more precisely, and mitigate potential inflation of performance by superficial heuristics such as coarse-grained similarity and word association, we propose incrementally increasing levels of controls in constructing our minimal pair stimuli, progressively making the task of attributing properties to concepts more challenging.

Findings from our initial experiment on COMPS-BASE established that the basic capacity of models to attribute properties to everyday concepts is largely coarse grained. PLMs were more successful in making correct property attributions when

the candidate concepts were radically different, and struggled when the concepts shared semantic relations or had high co-occurrence. On testing for ‘property inheritance’ behavior (via COMPS-WUGS), PLMs initially appeared to demonstrate reasonable success, but they also showed spurious behavior in achieving correct performance on a non-trivial number of instances for which they did not succeed in the prerequisite base condition. Furthermore, this performance declined substantially in the presence of distracting information (COMPS-WUGS-DIST), providing further evidence that what property knowledge and reasoning we appear to see in these PLMs is more reliant on superficial heuristics than on ideal reasoning behavior. Of particular note is our finding of *catastrophic distraction* in large autoregressive PLMs, whose sensitivity to proximity effects brings their overall performance well under chance, especially when scaled up to billions of parameters.

Contemporary work has highlighted the promise of PLMs on high-level tasks requiring—among other things—access to proper relational knowledge between concepts (see Petroni et al., 2019; Safavi and Koutra, 2021; Piantadosi and Hill, 2022). By drawing on the concept of property inheritance, our experiments target reasoning ability based on perhaps the most well-established of relations—the taxonomic or the *IsA* relation (Murphy, 2003). Recent work has also alluded to the proficiency of PLMs in capturing taxonomic information about everyday objects and entities (Weir et al., 2020; Chen et al., 2021, though see Ravichander et al. (2020)). Findings from our controlled experiments suggest that PLMs’ approximation of the consequences of the taxonomic relation is at best noisy, in light of clear failures especially in presence of similarity-governed competition. We conclude from our analyses that instead of robustly extracting relational information and reasoning about properties of concepts, it is likely that the PLMs tested here are optimized to prefer superficial cues in making word predictions, leading to mistakes and inaccuracies in presence of irrelevant and distracting information. Since robust natural language understanding will be critically reliant on understanding of property knowledge and implications of property transfer, we hope that these findings will motivate adoption of rigorous assessment methods as well as work toward more robust property knowledge and reasoning in PLMs.

<sup>7</sup>We thank Reviewers 1 and 3, Najoung Kim and Kyle Mahowald for raising this concern.

## Limitations

**Zero-shot setup** Using a zero-shot setup to test PLMs for human-like capacities such as property inheritance (as we have done in this work) has recently come under scrutiny. In particular, Lampinen (2022) argues that such a setup could be problematic because PLMs are trained to imitate the language produced by countless individuals with different beliefs, cultures, and behaviors. As a result, PLMs are likely to be handicapped in assigning sufficient probability mass to the desired family of continuations, given minimal prompts without any particular task-specific context. Instead, Lampinen (2022) suggests the need for PLMs

“[...] to be guided into an experiment-appropriate behavioral context, analogously to the way cognitive researchers place humans in an experimental context, and orient them toward the task with instructions and examples.”

This criticism is valid, and it is possible that models could overcome their lack of robustness to distraction effects by observing examples of our stimuli in context, though this has largely been shown in PLMs that are significantly larger than the ones we have tested in this work (Brown et al., 2020; Chowdhery et al., 2022; Wei et al., 2022a).<sup>8</sup> Indeed, recent work has demonstrated these larger PLMs to achieve strong performance on other types of reasoning—such as those required for solving math problems, reversing sequences, etc.—by priming models to produce additional textual content that represents intermediate reasoning steps and explanations (Nye et al., 2021; Wei et al., 2022b; Lampinen et al., 2022), in a few-shot setting.<sup>9</sup> At the same time, a few-shot version of COMPS stimuli could expose models to the possibility of leveraging heuristics that are naturally absent in the zero-shot setup, and therefore such a setup would critically require the design of additional controls, which we leave for future work.

<sup>8</sup>though see recent work by Shi et al. (2023), who show distraction effects in such large PLMs in solving arithmetic reasoning problems, even after using sophisticated in-context prompting methods such as Chain-of-Thought (Wei et al., 2022b), Least-to-Most (Zhou et al., 2022), and Self-Consistency (Wang et al., 2022).

<sup>9</sup>See also Sinha et al. (2022), who analyze PLMs comparable in size to those studied in this work in a few-shot minimal-pair setting.

**Ideal reasoning behavior** Another limitation of our work is that it takes ideal and robust property inheritance behavior as the monolithic gold-standard for human cognition, something that recent work has cautioned against (Pavlick and Kwiatkowski, 2019; Dasgupta et al., 2022; Webson et al., 2023). Although we relied on a database of concept-property pairs that were largely generated by human participants, whether or not humans will be robust to the types of distraction that were observed in PLMs is an open question and requires further investigation. However, notably we are not making direct comparisons between models and humans here—we argue that our primary contribution of controlled stimuli that tease apart shallow processing from robust conceptual reasoning in PLMs bears substantial merit that is independent from any comparisons between humans and computational systems. Furthermore, we emphasize that we are setting a reasonable—and to a certain extent, human-independent—desideratum in this work, which is that models should robustly capture ground-truth knowledge about everyday concepts and their properties and reflect this knowledge in their inferences about newly introduced concepts.

**Behavioral evaluation** This work tests and analyses PLMs on property knowledge and property inheritance only from a behavioral perspective, which at its core is a correlational endeavor. Potential future work could complement our results by providing evidence from representational analyses, or by devising causal interventions, similar to those recently explored in the realm of syntactic agreement (Finlayson et al., 2021), or in testing of negation and hypernymy in NLI models (Geiger et al., 2020), among others. Importantly, this would require the development of new methods that shed light on how new information—such as the ones we use in COMPS-WUGS and COMPS-WUGS-DIST—is integrated into the model (see Misra et al. (2022) for an example of such an analysis for novel properties).

**Targeted language** Finally, COMPS only consists of sentences in English, thereby biasing our results only for PLMs trained in that language.

## Acknowledgments

For helpful comments we thank Najoung Kim, Tal Linzen, Brenden Lake, Kyle Mahowald, Andrew Lampinen, the anonymous reviewers, and audiences at the Computation and Psycholinguistics

Lab (NYU), the Human and Machine Learning Lab (NYU), the UChicago CompLing Lab, UT Austin Linguistics Grad Student Seminar, and the MIT Department of Brain and Cognitive Sciences. Any errors are our own. Our experiments were conducted with resources provided by the Rosen Center for Advanced Computing at Purdue University (McCartney et al., 2014). We are also grateful to Sam Huang for helping out with experiments conducted on GPT-3/3.5 models.

## References

- Mostafa Abdou, Artur Kulmizev, Daniel Hershcovich, Stella Frank, Ellie Pavlick, and Anders Søgaard. 2021. [Can language models encode perceptual structure without grounding? a case study in color](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 109–132, Online. Association for Computational Linguistics.
- Yonatan Belinkov. 2022. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219.
- Jean Berko. 1958. The child’s learning of english morphology. *Word*, 14(2-3):150–177.
- Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#).
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Jamie Callan, Mark Hoy, Changkuk Yoo, and Le Zhao. 2009. [Clueweb09 data set](#).
- Catherine Chen, Kevin Lin, and Dan Klein. 2021. [Constructing taxonomies from pretrained language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4687–4700, Online. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. [Palm: Scaling language modeling with pathways](#). *arXiv preprint arXiv:2204.02311*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators](#). In *International Conference on Learning Representations*.
- Jeff Da and Jungo Kasai. 2019. [Cracking the contextual commonsense code: Understanding commonsense reasoning aptitude of deep contextual representations](#). In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 1–12, Hong Kong, China. Association for Computational Linguistics.
- Ishita Dasgupta, Andrew K Lampinen, Stephanie CY Chan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2022. [Language models show human-like content effects on reasoning](#). *arXiv preprint arXiv:2207.07051*.
- Steven Derby, Paul Miller, and Barry Devereux. 2021. [Representation and pre-activation of lexical-semantic knowledge in neural language models](#). In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 211–221, Online. Association for Computational Linguistics.
- Barry J Devereux, Lorraine K Tyler, Jeroen Geertzen, and Billi Randall. 2014. The Centre for Speech, Language and the Brain (CSLB) concept property norms. *Behavior Research Methods*, 46(4):1119–1127.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhisha 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.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. 2021. [Causal analysis of syntactic agreement mechanisms in neural language models](#). 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 1828–1843, Online. Association for Computational Linguistics.
- Maxwell Forbes, Ari Holtzman, and Yejin Choi. 2019. [Do Neural Language Representations Learn Physical Commonsense?](#) In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. 2019. [Neural language models as psycholinguistic subjects: Representations of syntactic state](#). In *Proceedings of NAACL-HLT 2019*, pages 32–42, Minneapolis, Minnesota. Association for Computational Linguistics.

- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Atticus Geiger, Kyle Richardson, and Christopher Potts. 2020. **Neural natural language inference models partially embed theories of lexical entailment and negation**. In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 163–173, Online. Association for Computational Linguistics.
- Aaron Gokaslan and Vanya Cohen. 2019. Openwebtext corpus. <http://Skylion007.github.io/OpenWebTextCorpus>.
- Alison Gopnik and David M Sobel. 2000. Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child development*, 71(5):1205–1222.
- Jonathan Gordon and Benjamin Van Durme. 2013. **Reporting bias and knowledge acquisition**. In *Proceedings of the 2013 workshop on Automated knowledge base construction*, pages 25–30.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.
- John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. 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 2733–2743.
- Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. **Surface form competition: Why the highest probability answer isn't always right**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7038–7051, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Nora Kassner and Hinrich Schütze. 2020. **Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.
- Najoung Kim, Tal Linzen, and Paul Smolensky. 2022. Uncontrolled lexical exposure leads to overestimation of compositional generalization in pretrained models. *arXiv preprint arXiv:2212.10769*.
- Brenden M Lake and Gregory L Murphy. 2021. Word meaning in minds and machines. *Psychological Review*.
- Andrew K Lampinen, Ishita Dasgupta, Stephanie CY Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L McClelland, Jane X Wang, and Felix Hill. 2022. Can language models learn from explanations in context? *arXiv preprint arXiv:2204.02329*.
- Andrew Kyle Lampinen. 2022. Can language models handle recursively nested grammatical structures? a case study on comparing models and humans. *arXiv preprint arXiv:2210.15303*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *International Conference on Learning Representations*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Li Lucy and Jon Gauthier. 2017. **Are distributional representations ready for the real world? evaluating word vectors for grounded perceptual meaning**. In *Proceedings of the First Workshop on Language Grounding for Robotics*, pages 76–85, Vancouver, Canada. 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.
- Gerry McCartney, Thomas Hacker, and Baijian Yang. 2014. **Empowering Faculty: A Campus Cyberinfrastructure Strategy for Research Communities**. *Education Review*.
- George A Miller. 1995. WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Kanishka Misra. 2022. minicons: Enabling flexible behavioral and representational analyses of transformer language models. *arXiv preprint arXiv:2203.13112*.
- 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.
- Kanishka Misra, Allyson Ettinger, and Julia Rayz. 2021. Do language models learn typicality judgments from text? In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*.

- Kanishka Misra, Julia Taylor Rayz, and Allyson Ettinger. 2022. A property induction framework for neural language models. In *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.
- Gregory L Murphy. 2002. *The Big Book of Concepts*. MIT press.
- M Lynne Murphy. 2003. *Semantic relations and the lexicon: Antonymy, synonymy and other paradigms*. Cambridge University Press.
- Sebastian Nagel. 2016. [CC-News](#).
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. 2021. Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Lalchand Pandia and Allyson Ettinger. 2021. [Sorting through the noise: Testing robustness of information processing in pre-trained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1583–1596, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Roma Patel and Ellie Pavlick. 2022. [Mapping language models to grounded conceptual spaces](#). In *International Conference on Learning Representations*.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP 2014*, pages 1532–1543.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? 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 2463–2473.
- Steven T Piantadosi and Felix Hill. 2022. Meaning without reference in large language models. *arXiv preprint arXiv:2208.02957*.
- M Ross Quilliam. 1967. Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral science*, 12(5):410–430.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI*.
- Abhilasha Ravichander, Yonatan Belinkov, and Eduard Hovy. 2021. [Probing the probing paradigm: Does probing accuracy entail task relevance?](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3363–3377, Online. Association for Computational Linguistics.
- Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. [On the systematicity of probing contextualized word representations: The case of hypernymy in BERT](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 88–102, Barcelona, Spain (Online). Association for Computational Linguistics.
- Lance J Rips, Edward E Smith, and Douglas L Medin. 2012. *Concepts and categories: Memory, meaning, and metaphysics*. Oxford University Press.
- Timothy T Rogers and James L McClelland. 2004. *Semantic cognition: A parallel distributed processing approach*. MIT press.
- Dana Rubinstein, Effi Levi, Roy Schwartz, and Ari Rapoport. 2015. [How well do distributional models capture different types of semantic knowledge?](#) In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 726–730, Beijing, China. Association for Computational Linguistics.
- Tara Safavi and Danai Koutra. 2021. [Relational World Knowledge Representation in Contextual Language Models: A Review](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1053–1067, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Thibault Sellam, Steve Yadlowsky, Jason Wei, Naomi Saphra, Alexander D’Amour, Tal Linzen, Jasmijn Bastings, Iulia Turc, Jacob Eisenstein, Dipanjan Das, et al. 2021. The multiberts: Bert reproductions for robustness analysis. *arXiv preprint arXiv:2106.16163*.

- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Chi, Nathanael Schärli, and Denny Zhou. 2023. [Large language models can be easily distracted by irrelevant context](#). *arXiv preprint arXiv:2302.00093*.
- Vered Shwartz and Yejin Choi. 2020. [Do neural language models overcome reporting bias?](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6863–6870, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Koustuv Sinha, Jon Gauthier, Aaron Mueller, Kanishka Misra, Keren Fuentes, Roger Levy, and Adina Williams. 2022. Language model acceptability judgements are not always robust to context. *arXiv preprint arXiv:2212.08979*.
- Steven A Sloman. 1993. Feature-based induction. *Cognitive psychology*, 25(2):231–280.
- Steven A Sloman. 1998. Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology*, 35(1):1–33.
- Edward E Smith and William K Estes. 1978. Theories of semantic memory. *Handbook of learning and cognitive processes*, 6:1–56.
- Pia Sommerauer and Antske Fokkens. 2018. [Firearms and tigers are dangerous, kitchen knives and zebras are not: Testing whether word embeddings can tell](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Brussels, Belgium. Association for Computational Linguistics.
- Pia Johanna Maria Sommerauer. 2022. Diagnosing semantic properties in distributional representations of word meaning.
- Alon Talmor, Oyvind Taffjord, Peter Clark, Yoav Goldberg, and Jonathan Berant. 2020. Leap-of-thought: Teaching pre-trained models to systematically reason over implicit knowledge. *Advances in Neural Information Processing Systems*, 33:20227–20237.
- Trieu H Trinh and Quoc V Le. 2018. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*.
- Amos Tversky. 1977. Features of similarity. *Psychological review*, 84(4):327.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Albert Webson, Alyssa Marie Loo, Qinan Yu, and Ellie Pavlick. 2023. Are language models worse than humans at following prompts? it’s complicated. *arXiv preprint arXiv:2301.07085*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022a. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022b. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Nathaniel Weir, Adam Poliak, and Benjamin Van Durme. 2020. Probing neural language models for human tacit assumptions. In *CogSci 2020*, pages 377–383. Cognitive Science Society.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of EMNLP 2020: Demos*, pages 38–45, Online. Association for Computational Linguistics.
- Zhibiao Wu and Martha Palmer. 1994. [Verb Semantics and Lexical Selection](#). In *32nd Annual Meeting of the Association for Computational Linguistics*, pages 133–138, Las Cruces, New Mexico, USA. Association for Computational Linguistics.
- Yian Zhang, Alex Warstadt, Xiaocheng Li, and Samuel R. Bowman. 2021. [When do you need billions of words of pretraining data?](#) 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 1112–1125, Online. Association for Computational Linguistics.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies

and reading books. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 19–27. IEEE.

## A Model Metadata

Table 3 shows the different models used in our experiments, along with their abbreviation, tokenization scheme, total parameters, vocabulary size, number of tokens encountered during training, and corpora on which they are pre-trained. All models were accessed using minicons (Misra, 2022),<sup>10</sup> a python library that serves as a wrapper around Huggingface’s transformers (Wolf et al., 2020), and provides a unified mechanism for eliciting log-probabilities in batch-wise manner for any autoregressive or masked LM that is accessible through the huggingface hub, or is trained using the transformers library. Experiments were performed using an NVIDIA V100 GPU (32 GB RAM) and took about 6 hours to run, discounting the time it took to download the models from the Huggingface Hub.<sup>11</sup>

## B Preview of COMPS stimuli

We show examples of stimuli from our COMPS-BASE, COMPS-WUGS, and COMPS-WUGS-DIST datasets in Listing 1 and Listing 2, respectively. Stimuli with distraction—i.e., in COMPS-WUGS-DIST—are similar to that in Listing 1, but with the `distraction_type` value set to either ‘before’ or ‘in-between’.

```
{
  "id": 12706,
  "property": "can fly",
  "acceptable_concept": "owl",
  "unacceptable_concept": "squirrel",
  "prefix_acceptable": "an owl",
  "property_phrase": "can fly.",
  "prefix_unacceptable": "a squirrel",
  "condition": "co-occurrence",
  "similarity": 0.62
}
```

Listing 1: An instance of COMPS-BASE. “condition” represents the negative sampling scheme, and “similarity” represents the similarity between the acceptable concept and the unacceptable concept on the basis of the condition (either Taxonomic, Property Norm, Co-occurrence, or Random).

<sup>10</sup><https://github.com/kanishkamisra/minicons>

<sup>11</sup><https://huggingface.co/models>.

```
{
  "item": 8343,
  "comps_id": 28798,
  "property": "has hooves",
  "acceptable_concept": "horse",
  "unacceptable_concept": "dog",
  "prefix_acceptable": "A dax is a
  ↪ horse. Therefore, a dax",
  "prefix_unacceptable": "A dax is a
  ↪ dog. Therefore, a dax",
  "property_phrase": "has hooves.",
  "negative_sample_type":
  ↪ "co-occurrence",
  "similarity": 0.62,
  "distraction_type": "undistracted"
}
```

Listing 2: An instance of COMPS-WUGS. “condition” and “similarity” are the same as in Listing 1. “distraction\_type” denotes the type of distraction used (undistracted, before, in-between).

Table 4 shows examples from each subset of COMPS, and the conditional log-probability scores as computed by GPT-J (Wang and Komatsuzaki, 2021), the largest LM tested on the full set of stimuli.

## C Additional findings and analyses

### C.1 Testing GPT-3/3.5

Recent work in scaling PLMs to hundred billion parameters has led to models such as GPT-3 (Brown et al., 2020), which are significantly larger than the largest model tested in the results discussed above (i.e., GPT-J, with 6B parameters). Testing them on the entire set of COMPS stimuli (49K + 3 × 13.8K pairs of sentences) is prohibitively expensive since they are only accessible through paid APIs. Nonetheless, we sampled a small set of COMPS stimuli—which we term as miniCOMPS—in order to get a glimpse of how well substantially larger PLMs elicit property knowledge and demonstrate reasoning behavior compatible with property inheritance. Specifically, we created miniCOMPS by sampling 1200 minimal pairs from each of our original COMPS subsets (matched in terms of real world concepts and properties across the subsets), such that all pairs of nonce words in the resulting miniCOMPS-WUGS-DIST end up being sampled equal number of times (100 times each).

Family	Model (Abbrev.)	Parameters	Vocab Size	Tokenization	Corpora	Tokens
ALBERT	albert-base-v2 (A-b)	11M	30,000	SentencePiece	WIKI and BC	3.3B
	albert-large-v2 (A-l)	17M				
	albert-xlarge-v2 (A-xl)	59M				
	albert-xxlarge-v2 (A-xxl)	206M				
BERT	distilbertbase-uncased (dB-b)	67M	30,522	WordPiece	WIKI and BC	3.3B
	bert-base-uncased (B-b)	110M				
	bert-large-uncased (B-l)	345M				
ELECTRA	electra-small (E-s)	13M	30,522	WordPiece	WIKI and BC	3.3B
	electra-base (E-b)	34M				
	electra-large (E-l)	51M				
RoBERTa	distilroberta-base (dR-b)	82M	50,265	Byte-pair encoding	OWTC BC, CC-NEWS, OWTC, and STORIES	2B 30B
	roberta-base (R-b)	124M				
	roberta-large (R-l)	355M				
GPT2	distilgpt2 (dGPT2)	82M	50,257	Byte-pair encoding	OWTC WEBTEXT	2B 8B*
	gpt2 (GPT2)	124M				
	gpt2-medium (GPT2-m)	355M				
	gpt2-large (GPT2-l)	774M				
	gpt2-xl (GPT2-xl)	1.5B				
EleutherAI	gpt-neo-125M (Neo-125M)	125M	50,257	Byte-pair encoding	PILE	300B 380B 420B 402B
	gpt-neo-1.3B (Neo-1.3B)	1.3B				
	gpt-neo-2.7B (Neo-2.7B)	2.7B				
	gpt-j-6B (GPT-J)	6B				

Table 3: Summary of the 22 models that we evaluate in this paper. **Legend for Corpora:** WIKI: Wikipedia; BC: BookCorpus (Zhu et al., 2015); CW: ClueWeb (Callan et al., 2009); CC: CommonCrawl GIGA: Gigaword (Graff et al., 2003); OWTC: OpenWebTextCorpus (Gokaslan and Cohen, 2019); CC-NEWS: CommonCrawl News (Nagel, 2016); STORIES: Stories corpus (Trinh and Le, 2018); WEBTEXT: WebText corpus (Radford et al., 2019); PILE: The Pile (Gao et al., 2020).

\*As estimated by Warstadt et al. (2020).

COMPS subset	Stimulus	Score
BASE	A <b>horse</b> has hooves.	-3.829
	A <b>dog</b> has hooves.	-4.963
WUGS	A fep is a <b>horse</b> . Therefore, a fep has hooves.	-2.153
	A fep is a <b>dog</b> . Therefore, a fep has hooves.	-3.392
WUGS-DIST ( <b>before</b> )	A wug is a dog. A fep is a horse. Therefore, a <b>fep</b> has hooves.	-2.919
	A wug is a dog. A fep is a horse. Therefore, a <b>wug</b> has hooves.	-2.895
WUGS-DIST ( <b>in-between</b> )	A fep is a horse. A wug is a dog. Therefore, a <b>fep</b> has hooves.	-3.616
	A fep is a horse. A wug is a dog. Therefore, a <b>wug</b> has hooves.	-3.092

Table 4: An example of matched stimuli across different COMPS subsets, as well as conditional log-probabilities elicited by GPT-J. Here, the property of interest is *has hooves*, the positive concept is HORSE, and the negative concept is DOG. The negative concept in this case was sampled using the co-occurrence knowledge representation method (see §2.3). Emboldened words indicate items that are different in the minimal pair. Refer to §3.2 for discussion on how ‘Score’ is computed.

**Models** As test subjects, we chose four GPT-3 models (Brown et al., 2020): ada, babbage, curie, davinci, with the last one being the largest (at 175B parameters), and an additional fifth davinci-based model called text-davinci-001,

which fine-tunes davinci on human-written demonstrations. We also test the recently proposed GPT-3.5 models, text-davinci-002 and text-davinci-003, which improve over davinci by additionally fine-tuning it on code *and* human-



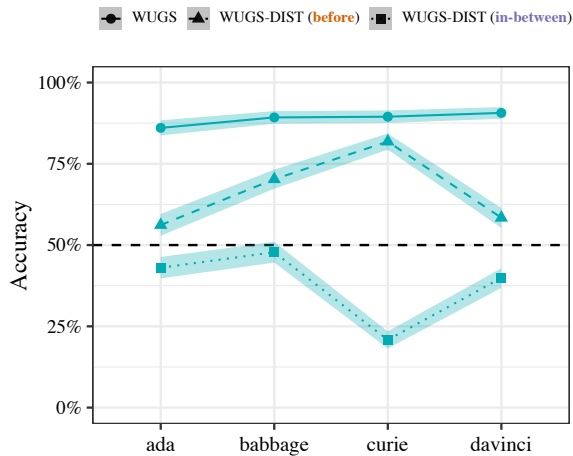


Figure 5: Accuracies of GPT-3 models (arranged in increasing order of the number of trained parameters) on miniCOMPS-WUGS and miniCOMPS-WUGS-DIST. Black dashed line indicates chance performance (50%). Error bands indicate 95% Bootstrap CIs.

written demonstrations (Ouyang et al., 2022).<sup>12</sup> All these models are autoregressive in nature, so we use the same scoring and evaluation method as described in §3.2. Since the four original GPT-3 models (ada, babbage, curie, davinci) are trained using the same LM objective on the same corpora, we analyze them separately from text-davinci-001, text-davinci-002, and text-davinci-003, which we only compare to davinci. We do this to remain consistent with the way we displayed results in §4—ordering models based on their number of trained parameters—and also because models in the text-davinci-XXX series use the same underlying davinci model augmented with additional training mechanisms (e.g., reinforcement learning and fine-tuning on human-feedback) and data (e.g., code) instead of increasing its size, to our knowledge.

**Results** Figure 5 shows the performance of the four GPT-3 models on miniCOMPS-WUGS and miniCOMPS-WUGS-DIST, while Figure 6 compares GPT-3 davinci to its code and human-feedback adapted counterparts. From Figure 5, we see robustness issues to persist even for GPT-3 models, similar to our main results. Models perform remarkably well in the absence of distraction (i.e., on miniCOMPS-WUGS), but struggle in its presence, especially when it is closer to the queried prop-

<sup>12</sup>These models are also known as InstructGPT, as discussed in <https://platform.openai.com/docs/model-index-for-researchers>.




Figure 6: Accuracies of davinci models (GPT-3 and GPT-3.5) on miniCOMPS-WUGS and miniCOMPS-WUGS-DIST. Black dashed line indicates chance performance (50%). Error bars indicate 95% Bootstrap CIs. davinci and text-davinci-001 are GPT-3 (Brown et al., 2020) models, while text-davinci-002 and text-davinci-003 are GPT-3.5 models.

erty. In particular, performance on miniCOMPS-WUGS-DIST (before) increases with an increase in parameters until the largest model (davinci), where the performance drops closer to chance. On miniCOMPS-WUGS-DIST (in-between), all models perform catastrophically worse than chance. This noteworthy pattern of proximity-based degradation in performance mimics the results shown in Figure 4, though we do not see a systematic decline in performance with an increase in parameters as observed in the GPT2 and EleutherAI models—with the 175B parameter model (davinci) demonstrating an increase in performance over the relatively smaller curie model.

While the above results demonstrate that simply scaling autoregressive PLMs is unlikely to overcome the lack of robustness against distracting content, we now test whether augmenting these large PLMs by additionally training on code (GPT-3.5 models) and aligning them with human-provided demonstrations (text-davinci-001 and both GPT-3.5 models) could lead to any improvements. For instance, training on code could provide training signals to PLMs that encourage entity tracking, which could potentially enable them, in our case, to resolve which subordinate concept (e.g., *wug* vs. *dax*) the target property is more likely to be associated with. Similarly, aligning with human-written demonstrations could potentially improve their truthfulness, which in our case, could lead to them to prefer correct property assignments. However, from Figure 6, we see no noteworthy improvements demonstrated by these augmented models. All augmented models achieved similar ac-

curacies on COMPS-WUGS as the davinci model (within 90.5% and 91%), suggesting that their augmentations preserved the general associations between the lexical items that denote everyday concepts and properties. On stimuli containing distraction (i.e., both subsets of COMPS-WUGS-DIST), either the models performed systematically worse as compared to davinci (with `text-davinci-002` showing below-chance performance on both subsets), or they showed mixed results, where an improvement on COMPS-WUGS-DIST (**before**) was accompanied by a decline on COMPS-WUGS-DIST (**in-between**).

Together, these results suggest that neither an increase in scale nor additional training methods such as alignment with human instructions/feedback or training on code prevents models from being distracted in associating properties to novel subordinate concepts introduced in the input context. In fact, the catastrophic effects of proximity-based distraction persists even for the most recent state of the art GPT-3/3.5 models.

## C.2 Results with alternate nonce words

Here we report results on COMPS-WUGS and COMPS-WUGS-DIST using an alternate set of nonce words, which we constructed by sampling (with replacement) from 26 lower-case ASCII alphabet characters. Specifically, we constructed novel character sequences—each assigned as a replacement for our original four nonce words—of lengths ranging from 4-8 by sampling in an alternate fashion from consonants (odd positions) and vowels (even positions).<sup>13</sup> A replication of Figure 4 using the stimuli with these newly sampled nonce words is shown in Figure 7. On comparing figures 4 and 7, we observe largely similar patterns of results on stimuli containing nonce words constructed using randomly sampled characters. That is, models generally performed well on COMPS-WUGS, while they struggled on COMPS-WUGS-DIST. There were some exceptions: (1) GPT-Neo 1.3B and 2.7B showed improvements (relative to the original stimuli) in cases where distraction is added closer to the queried property (i.e., **in-between**), though they still hover around chance performance, and additionally the performance of GPT-J, like in the original results is still substantially below chance; and (2) there were non-trivial improvements demon-

<sup>13</sup>the resulting set of words is: {*ruhisin, kifosa, rosibif, lepuvu*}, still amounting to 12 unique ordered pairs in the COMPS-WUGS-DIST stimuli.

Stimuli	Avg. Agreement
COMPS-WUGS	93.1 <sub>0.8</sub>
COMPS-WUGS-DIST ( <b>before</b> )	73.9 <sub>4.6</sub>
COMPS-WUGS-DIST ( <b>in-between</b> )	73.0 <sub>4.6</sub>
Overall	80.0 <sub>2.9</sub>

Table 5: Average agreement ( $\times 100$ ) in PLMs’ preference on stimuli containing original and synthetically constructed nonce words.

strated by ALBERT models (large and xl) on the **before** subset of COMPS-WUGS-DIST, and BERT-large on the **in-between** subset of COMPS-WUGS-DIST.

To precisely quantify the difference between the two sets of results, we measured the agreement between the predictions of the PLMs for both sets of stimuli, taken as the proportion of minimal pairs in which the models’ relative preference agree.

Figure 8 shows individual model agreement on COMPS-WUGS and COMPS-WUGS-DIST, while Table 5 shows agreement percentages averaged across all models. From these results we observe models to show greater robustness to the variability introduced by the choice of nonce words in stimuli with one novel concept (COMPS-WUGS) than in stimuli with multiple novel concepts (COMPS-WUGS-DIST). Despite this discrepancy, there is generally a high average agreement (80%) between a given model’s set of decisions on stimuli with original and alternative nonce words.

## C.3 Results with alternate templates

Here we report results on an alternate phrasing of our stimuli, where instead of using the original template for introducing novel concepts in context (*a wug is a [CONCEPT]*), we use: *A wug is a type of [CONCEPT]*, where *wug* indicates the novel concept. In all cases, we simply alter the template, keeping everything else constant, including the choice of nonce words.

Figure 9 shows accuracies of the models on stimuli with this alternate phrasing, while Figure 10 and Table 6 show individual and averaged overall agreement between models’ preference on original and the alternatively-phrased stimuli, respectively. The agreement percentages between models’ preferences are quite high (average agreement being 90%)—in fact even greater than the agreement observed as a result of altering the nonce words (Table 5), further cementing the robustness of our




Figure 7: Accuracies of individual models (grouped by family, in increasing order based on number of parameters) on COMPS-WUGS and COMPS-WUGS-DIST with **synthetically constructed nonce words**. Black dashed line indicates chance performance (50%). Refer to Table 3 for unabbreviated model names.

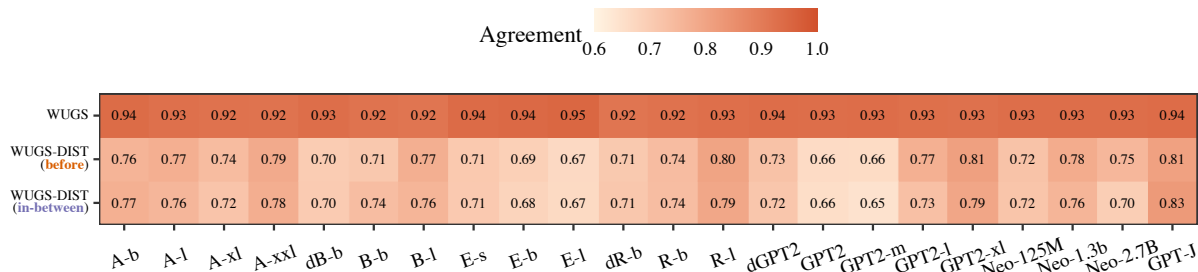


Figure 8: Proportion of cases (in COMPS-WUGS and both subsets of COMPS-WUGS-DIST) where each listed model's preference on the original stimuli matches that in stimuli with synthetically constructed nonce words, measured as the 'Agreement'. An agreement of 1.0 suggests that a given model's preferences are perfectly matched across both sets of stimuli.

Stimuli	Avg. Agreement
COMPS-WUGS	93.4 <sub>1.2</sub>
COMPS-WUGS-DIST (before)	88.5 <sub>3.1</sub>
COMPS-WUGS-DIST (in-between)	88.0 <sub>3.6</sub>
Overall	90.0 <sub>2.6</sub>

Table 6: Average agreement ( $\times 100$ ) in PLMs' preference on stimuli containing original (*A wug is a [CONCEPT].*) and alternate framing of novel taxonomic information (*A wug is a type of [CONCEPT].*).

results.

#### C.4 How does performance on COMPS-BASE vary by property type?

Devereux et al. (2014) have categorized the properties that we use in our experiments to lie in 5 different categories: (1) **Taxonomic**, e.g., *is a mammal*, *is a vehicle*, etc.; (2) **Functional**, e.g., *can keep the body warm*, *is used to hit nails*, etc.; (3) **Encyclopedic**, e.g., *uses electricity*, *is warm blooded*,

etc.; (4) **Visual Perceptual**, e.g., *has webbed feet*, *has thick fur*, etc.; and (5) **Other Perceptual**, e.g., *makes grunting sounds* and *is sharp*, etc. We report results of the 22 PLMs on the COMPS-BASE stimuli across the five different property types, in Figure 11.

From Figure 11, we observe that PLMs are substantially stronger in eliciting taxonomic properties of concepts as compared to other types, with highest overall accuracy being 70%, as compared to 48% on encyclopedic properties, 50% on visual perceptual properties, 57% on functional properties, and 43% on non-visual perceptual properties. Recall that chance accuracy for the 'Overall' scenario is just 6.25%, so these scores are fairly high. This corroborates evidence from previous work in analyzing property knowledge of distributional semantic models as well as LM representations to lack perceptual knowledge (Lucy and Gauthier, 2017; Da and Kasai, 2019; Rubinstein et al., 2015; Weir et al., 2020), likely due to reporting bias (Gordon and Van Durme, 2013; Shwartz and Choi, 2020).




Figure 9: Accuracies of individual models (grouped by family, in increasing order based on number of parameters) on COMPS-WUGS and COMPS-WUGS-DIST with **alternate framing of novel taxonomic information**. Black dashed line indicates chance performance (50%). Refer to Table 3 for unabbreviated model names.

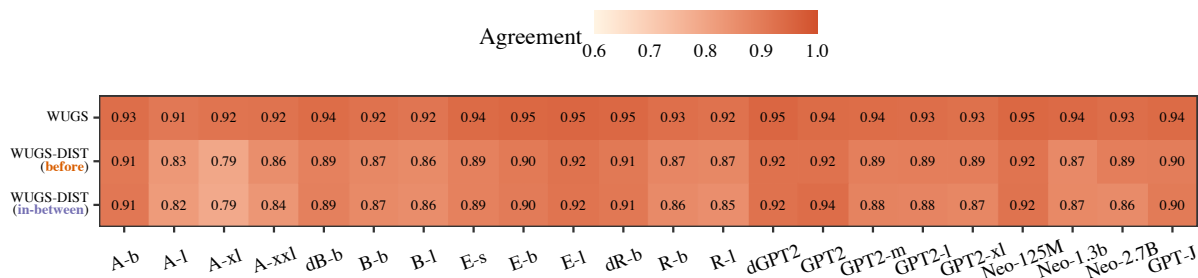


Figure 10: Proportion of cases (in COMPS-WUGS and both subsets of COMPS-WUGS-DIST) **where each listed model’s preference on the original stimuli matches that in stimuli with alternate framing of novel taxonomic information**, measured as the ‘Agreement’. An agreement of 1.0 suggests that a given model’s preferences are perfectly matched across both sets of stimuli.

However, different to most of these works, the gap between performance on perceptual properties and non-perceptual properties is small. We conjecture that this could be primarily due to the extension of the CSLB by Misra et al. (2022), which lead to an increase in coverage of property knowledge for several properties. For instance, the property *has teeth* was mentioned only for 45 out of 67 potential concepts, having been left out for concepts such as CALF,<sup>14</sup> BUFFALO, KANGAROO, etc. So it could be the case that previous research has underestimated the extent to which property knowledge is encoded by PLMs and other distributional semantic models of language.

### C.5 Does performance on COMPS-BASE depend on scale?

We plot the accuracies of PLMs on COMPS-BASE per model family (in order to control for differences in training corpora and tokenization) in Figure 12.

<sup>14</sup>the young one of a cow, and not the muscles in the vertebrate body

In all families except BERT, we see that accuracy increases with the model size, following standard scaling laws. We notice that distilBERT-base (Sanh et al., 2019) is able to outperform even BERT-large on stimuli with ‘Random’ negative samples, suggesting that pruning BERT might sometimes unintentionally improve the model’s ability to associate properties and concepts. We do however caution against interpreting these results as robust conclusion for scaling laws on COMPS-BASE. Such an endeavor would require comparing performance of models across multiple checkpoints with varying number of parameters, paired with rigorous statistical inference (Sellam et al., 2021; Zhang et al., 2021).

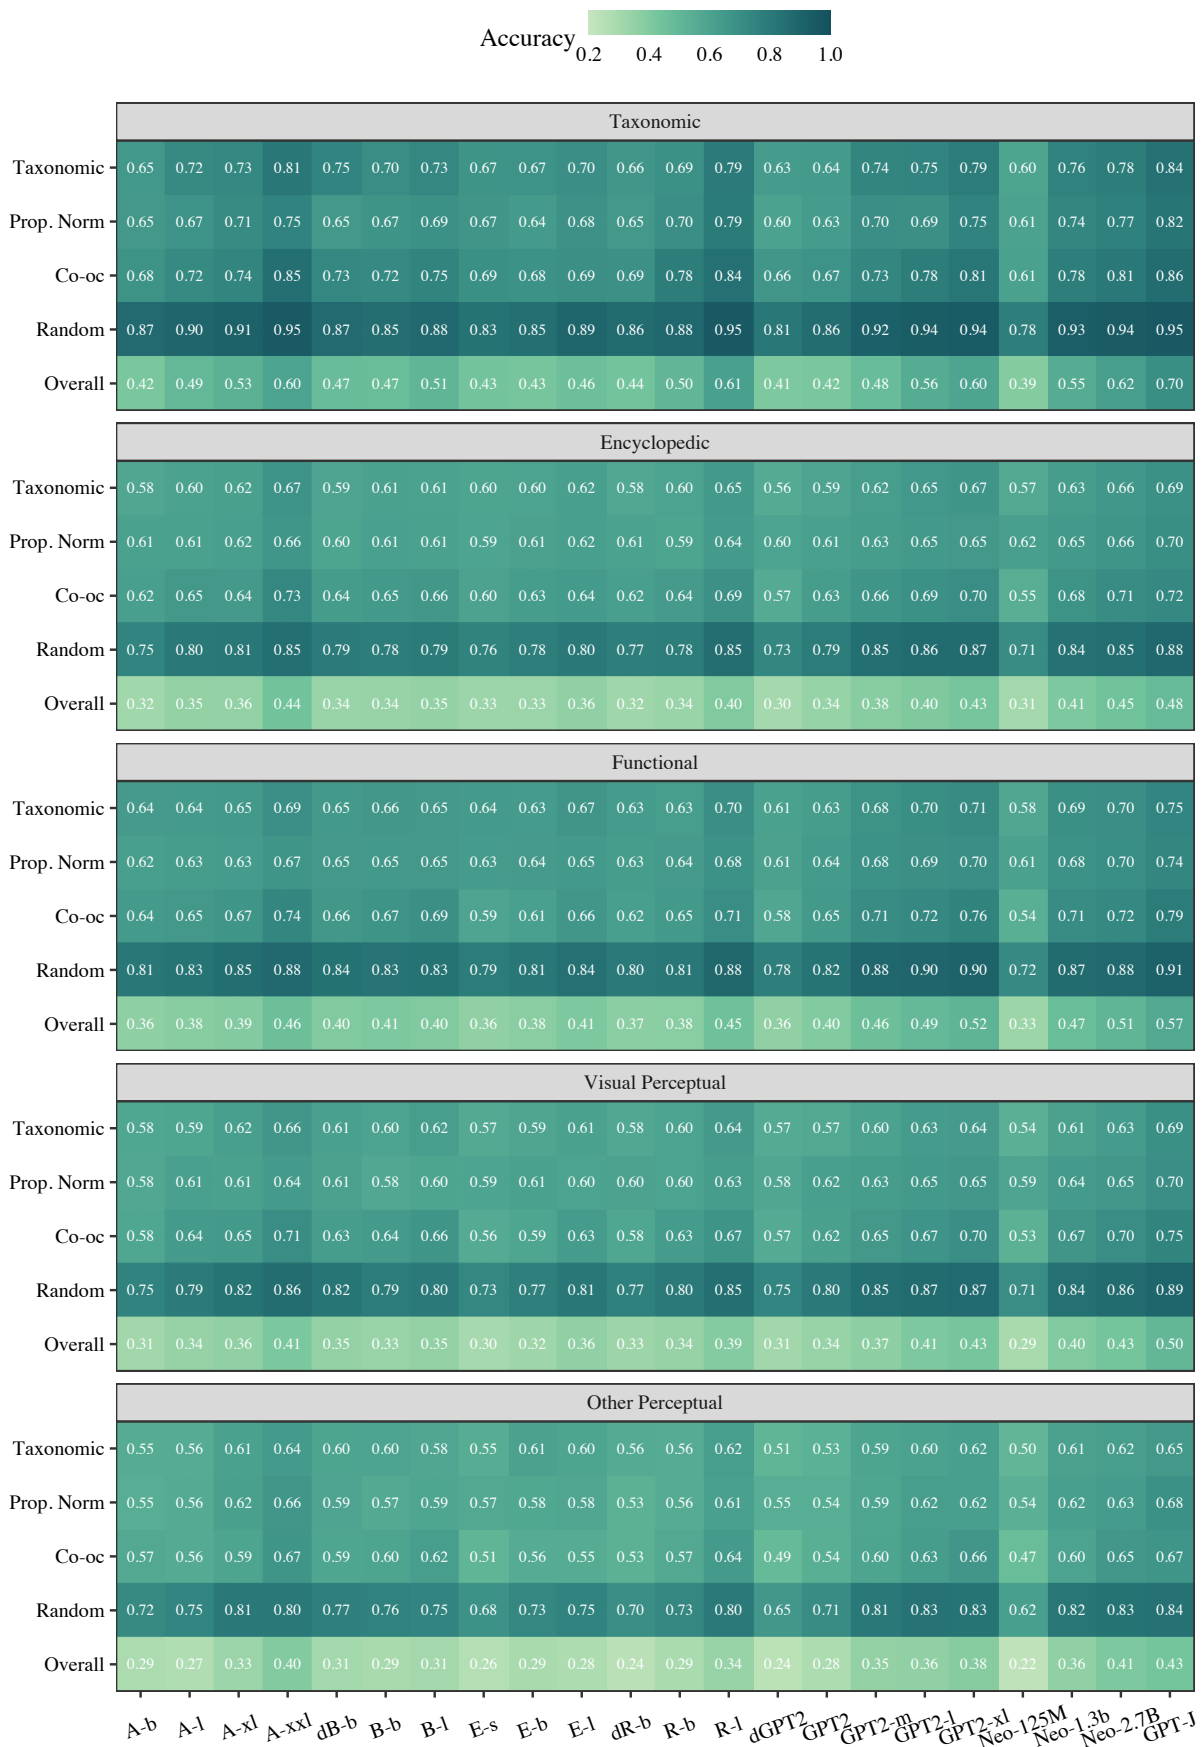


Figure 11: COMPS-BASE performance across five property types annotated in CSLB (Devereux et al., 2014).




Figure 12: Accuracy vs. parameters across various negative sampling strategies. Models are shaded based on number of parameters.