

Characterizing the Role of Similarity in the Property Inferences of Language Models

Juan Diego Rodriguez[†] Aaron Mueller^{‡,§} Kanishka Misra^{★,χ}

[†]The University of Texas at Austin [‡]Northeastern University

[¶]Technion – Israel Institute of Technology ^χToyota Technological Institute at Chicago

juand-r@utexas.edu aa.mueller@northeastern.edu kanishka@ttic.edu

Abstract

Property inheritance—a phenomenon where novel properties are projected from higher level categories (e.g., birds) to lower level ones (e.g., sparrows)—provides a unique window into how humans organize and deploy conceptual knowledge. It is debated whether this ability arises due to explicitly stored taxonomic knowledge vs. simple computations of similarity between mental representations. How are these mechanistic hypotheses manifested in contemporary language models? In this work, we investigate how LMs perform property inheritance with behavioral and causal representational analysis experiments. We find that taxonomy and categorical similarities are not mutually exclusive in LMs’ property inheritance behavior. That is, LMs are more likely to project novel properties from one category to the other when they are taxonomically related and at the same time, highly similar. Our findings provide insight into the conceptual structure of language models and may suggest new psycholinguistic experiments for human subjects.¹

1 Introduction

Categories are fundamental to human semantic cognition. Our knowledge of categories allows us to draw everyday inferences; a prominent example of which is *property inheritance*, where properties are projected from a category to its members. For example, if we learn that *dogs* have the T9 hormone, we can reasonably assume that *corgis* also have the T9 hormone. In cognitive psychology, an obvious and previously popular explanation for property inheritance is that it is the natural consequence of our minds organizing categories hierarchically into taxonomies (Collins and Quillian, 1969; Glass and Holyoak, 1974; Murphy, 2002).

¹Data and code available at <https://github.com/aaronmueller/lm-property-inheritance>.

[★]Work partly done at UT-Austin before joining TTIC.

Given that **dogs** are *daxable*, is it true that **corgis** are *daxable*?

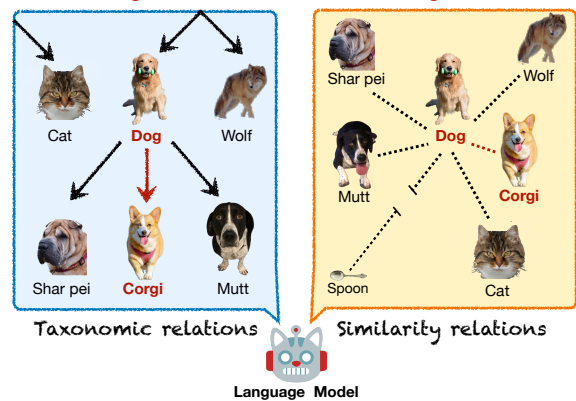


Figure 1: **Property inheritance** involves projecting properties from a category to its members. For example, if **dogs** are *daxable*, are **corgis** *daxable*? Here, *daxable* is a nonce word used to study property inheritance without any confounding effects from parametric knowledge. Language models may rely on taxonomic (left) and/or similarity (right) relations to perform property inheritance. We investigate the interplay between these two effects in LMs’ property inheritance judgments using both behavioral and mechanistic analyses.

At the same time, empirical evidence has called this assumption into question. Sloman (1998) showed that humans were highly sensitive to the similarity of the categories when performing property inheritance; for example, they were more likely to project novel properties from *birds* to *robins* (a more typical bird) than from *birds* to *penguins* (a less typical bird), despite them agreeing that both were members of the bird category. Sloman concluded that instead of explicitly using stored taxonomies, humans might simply be computing categorical similarities to demonstrate inheritance-compatible behavior. In such a case, a property is more likely to be shared between categories insofar as the similarity between them is high enough.

Debates about the mechanisms that underlie human property inheritance offer an exciting avenue to analyze conceptual organization and use in language models (LMs). First, due to the close connec-

tion of a word to its conceptual meaning (Murphy, 2002; Lupyán and Lewis, 2019; Lake and Murphy, 2021), property inheritance provides an opportunity to diagnose meaning-sensitivities in LMs (Piantadosi and Hill, 2022; Bender and Koller, 2020). Next, similarity effects like those found by Sloman (1998) are related to the notion of *content effects* shown by humans and LMs on logical reasoning tasks (Wason, 1968; Evans, 1989; Lampinen et al., 2024); the presence of similarity effects goes against the idea that humans rely on abstract taxonomic principles (where every category member would be equally likely to inherit the property). Finally, the principles of taxonomy vs. similarity offer two clear hypotheses about what mechanisms could govern property inheritance behavior in a system. This allows us to go beyond prior work investigating property inheritance in LMs (Misra et al., 2023), which has largely focused on their *behavior*, leaving open questions about the underlying mechanisms.

In this study, we empirically investigate the roles of taxonomic relations and categorical similarities on LMs’ property inheritance behavior, using both behavioral and causal interpretability methods on four LMs. First, we behaviorally analyze LMs’ sensitivity to taxonomic relations when performing property inheritance (§3.4). We demonstrate that noun similarity, known to correlate with human property inheritance behavior (Sloman, 1998), also correlates strongly with property inheritance judgments in language models, and explains many of the false positives and false negatives. Then, we causally localize property inheritance to specific activation subspaces (§4) using distributed alignment search (DAS; Geiger et al., 2024b). By training and evaluating DAS on controlled subsets of the data, we show that the subspaces responsible for property inheritance reflect both taxonomic and similarity features. Together, these results indicate that LMs are susceptible to non-trivial content effects rooted in noun similarities, and moreover that taxonomy and similarity are *fundamentally entangled* in model representations.

2 Conceptual Organization in LMs

To what extent is the similarity vs. taxonomy debate relevant to LMs? By definition, LMs learn representations from distributional statistics of tokens in context, so any reasoning they demonstrate is a result of distributional similarity. While technically

true, this viewpoint may be an oversimplification: two words can be distributionally similar if they participate in antonymy (*hot* vs. *cold*), metonymy (*car*, *wheel*), hypernymy (*robin*, *bird*), or even if they share thematic relations (*dog*, *bone*). For property inheritance, however, it is only the **hypernymy** relation that plays a critical role (Murphy, 2002). Is it possible that something structural like a taxonomy might arise through distributional statistics? It has been postulated that symbolic behavior can emerge as a result of sub-symbolic processes internal to a neural network (Smolensky, 1986, 1988; McClelland et al., 2010). Empirically, this has been shown in a number of recent works, e.g., Nanda et al. (2023) on LMs and modular arithmetic, and Feng et al. (2023) on LMs learning to simulate dynamic programming—both of which require operating over symbolic representations.

The above discussion raises multiple interrelated questions: when performing property inheritance, do LMs show sensitivity to taxonomic relations? If so, is their behavior *graded*, where properties are less likely to be projected to concepts that are taxonomically related to the higher-level category but low in similarity? Or, is it more *binary*, where there is no significant difference between how likely a property is projected from the higher-level category to its high- and low-similarity members? If their behavior is graded, then does it differ between non-category members that might share similarities with the higher level category (e.g., bird and giraffe) vs. those that do not (e.g., bird and chair)?

Related work By focusing on the mechanisms that underlie LMs’ property inheritance behavior, our work contributes to a rich body of work that investigates conceptual representations in language models (Bhatia and Richie, 2021; Misra et al., 2021; Patel and Pavlick, 2022; Abdou et al., 2021; Grand et al., 2022; Wu et al., 2023a; Park et al., 2024, *i.a.*). Complementing these, our focus in this work is to understand how LMs *use* their conceptual structure to project properties from one category to another—an important task in the broader space of inductive problems solved by humans (Kemp and Jern, 2014).

In this vein, recent works have investigated LMs’ property inheritance behavior—either by purely behavioral analyses (Misra et al., 2023; Wu et al., 2023a), or by manipulating LMs’ internal representations to edit their taxonomic knowledge (Powell et al., 2024; Cohen et al., 2024), and measuring

the implications of the edit. Importantly, these works assume that changes to taxonomic membership must be directly reflected in the LMs’ behavior, without necessarily considering the impact of categorical similarity. Grounding our investigation in empirical observations about humans (Sloman, 1998), we abstract away from this assumption, and explicitly consider the interplay between similarity and taxonomic relations by characterizing it using both behavioral and causal analysis methods.

3 Experimental Materials

3.1 Data

We used the THINGS dataset (Hebart et al., 2019, 2023) as our primary repository of noun categories. THINGS is a repository of human responses on odd-one-out tasks for 1,854 unique object categories—e.g., *dogs* and *chairs*. Each object concept is also annotated with its superordinate category (*dog–animal*). We performed a few modifications to this dataset. First, many objects belonged to more than one category—e.g., *cat* is an *animal* but it is also a *mammal*; we make these into separate entries. Second, we also paired each object with its WordNet sense, since one of the similarity methods we use involves ground-truth knowledge of word senses (we expand on this in §3.2). We discarded instances for which we could not find a valid sense in WordNet. These manipulations result in 2,016 pairs of taxonomically related object-category pairs. Our final dataset has 44 superordinate and 1,281 subordinate categories, respectively. Next we describe our method to sample pairs that are not taxonomically related.

Negative Sampling Because we are explicitly analyzing the effect of *similarity* in LMs’ property inheritance behavior, we include similarity measures as a central component for deriving pairs of related categories that are *not* taxonomically related. We perform this sampling as follows: for each superordinate category C , we construct the negative sample space \mathcal{N} as the set of items that are *outside* that category—e.g., for *bird*, this could mean non-birds such as *zebra*, *sofa*, etc. We then compute the similarity of each item in \mathcal{N} to the superordinate category using a given similarity measure. Finally, if k is the number of category members of C , then we sample the top $k/2$ and bottom $k/2$ concepts from \mathcal{N} , based on the similarity values. This way, for each superordinate category, we have k category members, and k non-category members,

with an even split between high and low-similarity non-taxonomic items.

Stimuli design We follow precedent from research in the psychology of concepts (Osherson et al., 1990; Sloman, 1998), and create stimuli in a premise-conclusion format, commonly used to analyze category-based inference in humans. In each instance, the premise expresses a statement where a category is said to have some property (*birds have the T9 hormone*), and the conclusion asks if a different category also has that property (*robins have the T9 hormone*). Following the same precedent, we use properties that are expressed by nonce words (e.g., *are daxable*). Since LMs are likely to have less—if any—prior knowledge about these nonce properties, this design decision allows us to isolate LMs’ inference behavior to the relations between categories (i.e., taxonomy and similarity), devoid from any interference from knowledge of real properties (e.g., *can fly*) which may already be embedded in their learned weights. We experimented with multiple different surface form realizations of the stimuli, and found the following template to work best on average:²

Answer the question. Given that **A** is/are *daxable*, is it true that **B** is/are *daxable*?
Answer with Yes/No. The answer is:

where **A** and **B** are the premise and conclusion categories, respectively. Most nouns were pluralized, except for those which were more naturally expressed in singular form (mass nouns, e.g., *honey*). A model that is sensitive to taxonomic relations when performing property inheritance should be more likely to generate *Yes* than *No* for a taxonomic pair such as (*toy*, *doll*), while the reverse should be true for non-taxonomic pairs such as (*fruit*, *puppy*).

3.2 Similarity scores

It is non-trivial to predict the nature of the ‘similarity’ that might have an effect on a system’s category-based inference behavior. This is true even for humans—most experiments used human derived similarity ratings to construct their stimuli (Osherson et al., 1990; Sloman, 1998), but what these ratings actually captured was largely unknown. To this end, we used two different types of similarity metrics, one sensitive to global distributional contexts of word-senses, and the other sensitive to the visual and conceptual properties of

²Some models demonstrated better performance when we omitted “The answer is:” from this prompt. The entire set of templates is made available in Appendix C.

Similarity Type	Category	Taxonomic Neighbors	Non-taxonomic Neighbors
Word-Sense	vehicle	<i>car</i> (0.85), <i>sled</i> (0.70), <i>hot-air balloon</i> (0.59)	<i>headlight</i> (0.82), <i>spinach</i> (0.11), <i>calzone</i> (0.09)
SPOSE	bird	<i>eagle</i> (0.95), <i>pelican</i> (0.90), <i>penguin</i> (0.82)	<i>bee</i> (0.92), <i>gazelle</i> (0.87), <i>cabinet</i> (0.08)

Table 1: Examples of taxonomically related and non-taxonomically related neighbors of the *vehicle* and *bird* categories determined by the two different types of similarities. Values in parentheses indicate similarities.

objects (*animacy*, *color*, etc.). Both metrics were derived from vector-space embeddings, and similarity was calculated using cosines between vectors of the input concepts.³

Word-sense Similarity We used the LMMS-ALBERT-xxl model (Loureiro et al., 2022) which was constructed by aggregating contextualized embeddings from the ALBERT-xxl LM (Lan et al., 2020) for lexical items with tagged WordNet senses. This model can be considered as a distributional semantic model like Glove (Pennington et al., 2014) that isolates the distributional information of the particular sense of a word—i.e., the embedding for *bat* when used as a mammal is different from that used as sporting equipment.

SPoSE Similarity We used SPoSE (Zheng et al., 2019), a sparse, non-negative vector space model fitted to human behavioral judgments on odd-one-out tasks, where participants were tasked to choose a concept from a triple that was least similar to the other two (e.g., {*crow*, *sparrow*, *ship*}). We used the 49-dimensional embedding released by (Hebart et al., 2023), which achieved a correlation of up to 0.9 with human judgments of perceived similarity (Kaniuth et al., 2024). Since the embeddings in SPoSE are for subordinate categories, we augmented this dataset with superordinate category vectors by taking the mean of the vector representations of the category members for each category—e.g., for *bird*, we took the average of the embeddings of {*crow*, *sparrow*, ...}.

Table 1 shows examples of taxonomically and non-taxonomically related neighbors of a few categories according to both types of similarity.⁴ For each similarity type, we bin the category pairs in our stimuli into ‘high’ or ‘low’ similarity by calculating the median similarity for premise category and categorizing instances where the similarity between the premise and the conclusion is above the median to be ‘high’, and those below to be ‘low’.

³We discuss other possible metrics we could have used, and our reasons for not using them in the limitations.

⁴Additional examples are given in Appendix E.

This way, every premise category in our stimuli is associated with equal numbers of high and low similarity conclusion categories. In total, we have 4032 category pairs, for each type of similarity.

3.3 Models

We analyze four instruction-tuned LMs: Mistral 7B Instruct v0.2 (Jiang et al., 2023), the 2B and the 9B parameter variants of the Gemma 2 family (Rivière et al., 2024), and the 8B parameter variant of Llama 3 Instruct (Dubey et al., 2024).⁵ We found varying degrees of success in using the chat-templates for these LMs, and therefore used it only when it resulted in a non-trivial improvement over the standard method of passing inputs to the LMs. All models were accessed using the transformers library (Wolf et al., 2020).

3.4 Behavioral Characterization of Property Inheritance

We begin by characterizing the role of taxonomic relations and similarity in LMs’ property inheritance behavior. To this end, we collect LMs’ responses to our stimuli, and analyze the conditions under which they are more likely to extend the property from the premise category to the conclusion category. We do this by comparing their relative probabilities for ‘Yes’ vs. ‘No’ on our stimuli, which we compute as follows:⁶

$$P_{\text{rel}}(x) = \frac{p_{\theta}(x \mid \text{prefix})}{\sum_{x \in \{\text{Yes}, \text{No}\}} p_{\theta}(x \mid \text{prefix})},$$

where ‘prefix’ is the stimulus, formatted as discussed in §3, and $p_{\theta}(\cdot)$ denotes the LM’s next

⁵Initial experiments showed these LMs’ non instruction-tuned counterparts to perform worse at property inheritance, which corroborates recent evidence about the advantage of instruction-tuning for property inheritance (Misra et al., 2024).

⁶The LMs that we use have two different tokens to represent ‘Yes’—one with and one without space, and similarly for ‘No’. To address this potential issue, we take the maximum of the probabilities over the two options for both cases. That is, $p_{\theta}(\text{Yes}) = \max(p_{\theta}(\text{Yes}), p_{\theta}(_ \text{Yes}))$, and similarly for ‘No’. We note that these tokens appeared as the top predicted tokens for all models evaluated except Gemma-2-2B-IT.

Model	Word-Sense				SPoSE			
	TS	PS	MS	ρ	TS	PS	MS	ρ
Gemma-2-2B-IT	0.84	0.85	0.89	0.26	0.79	0.79	0.88	0.59
Mistral-7B-Instruct-v0.2	0.84	0.85	0.84	0.42	0.80	0.81	0.83	0.62
Llama-3-8B-Instruct	0.86	0.86	0.98	0.37	0.79	0.79	0.98	0.68
Gemma-2-9B-IT	0.90	0.89	0.85	0.34	0.85	0.84	0.84	0.65

Table 2: Behavioral sensitivities of LMs across both similarity types (Word-Sense and SPoSE). TS: Taxonomic Sensitivity; PS: Property Sensitivity; MS: Mismatch Sensitivity; ρ : Spearman correlation of model scores ($P_{\text{rel}}(\text{Yes})$) with similarity ($p < .001$ throughout).

token distribution which we computed using the minicons library (Misra, 2022). Using these relative probabilities, we compute the following metrics to analyze LMs’ property inheritance behavior:

Taxonomic Sensitivity (TS) We measure the extent to which LMs are sensitive to the taxonomic relations between the premise and conclusion category. We do so by computing the proportion of examples where $P_{\text{rel}}(\text{Yes}) > 0.5$ when the premise and conclusion are taxonomically related, and $P_{\text{rel}}(\text{Yes}) < 0.5$ when they are not.

Property Sensitivity (PS) The TS metric is computed for stimuli where the only nonce property is *is/are taxable*. To what extent does the specific surface form of the property play a role in LMs’ property inheritance behavior? To this end, we generate a different set of stimuli where we randomly replace both instances of *is/are taxable* with *has/have feps* for half of the examples. We maintain the even balance between taxonomically and non-taxonomically related premise-conclusion pairs. Given this data, we then compute the taxonomic sensitivity in this setting as before. This metric penalizes bias for a particular property.

Mismatch Sensitivity (MS) We measure the extent to which the LMs are sensitive to whether the properties mentioned in the premise and conclusion match. If they are mismatched, then regardless of the taxonomic relationship between the concepts, the LMs should always prefer ‘No’ over ‘Yes’. To test this, we generate stimuli where we uniformly substitute either the premise or the conclusion with a different property (e.g., *birds are taxable vs. robins have feps*), and compute the proportion of time $P_{\text{rel}}(\text{Yes}) < 0.5$.

Spearman correlation with similarity (ρ) Finally, we measure the Spearman correlation between $P_{\text{rel}}(\text{Yes})$ and the similarity between the premise and conclusion categories, to quantify the

role of similarity on LMs’ inference behavior.

We compute these metrics separately for both types of similarity. For the three sensitivity metrics (TS, PS, MS), chance sensitivity is 0.5, as they all involve pairwise comparisons. We also measure the average $P_{\text{rel}}(\text{Yes})$ across four slices of our data based on when the premise and conclusion categories were (i) non-taxonomically related, low similarity (-Tax, -Sim), (ii) non-taxonomically related, high similarity (-Tax, +Sim), (iii) taxonomically related, low similarity (+Tax, -Sim), and (iv) taxonomically related, high similarity (+Tax, +Sim).

3.5 Results and Analysis

Table 2 shows the aforementioned metrics for all four LMs, across both similarity types, while Figure 2 shows their average $P_{\text{rel}}(\text{Yes})$ values across the four different slices of our data. For both similarity types, all four LMs demonstrate substantially high TS and PS values—i.e., they are more likely to extend the novel property when the premise category is a superordinate of the conclusion category. This suggests a non-trivial role of taxonomic category membership in the models’ inference behavior. At the same time, the LMs also show positive correlation with categorical similarity, suggesting that their tendency to extend properties from the premise to the conclusion is also sensitive to the similarity between the two. More specifically, they show greater correspondence with similarities derived from SPoSE ($\rho \in [0.59, 0.68]$) than from the sense embeddings ($\rho \in [0.26, 0.42]$).

While LMs demonstrate sensitivity to similarity, their tendency to produce ‘Yes’ over ‘No’ seems to correspond better with the presence of taxonomic relations. In Figure 2, LMs on average prefer to extend the property from premise to conclusion for high-similarity pairs than for low-similarity ones. However, we see the average $P_{\text{rel}}(\text{Yes})$ to be largely greater than 0.5 for the slices of our data where the premise is a taxonomic superordinate of the conclusion, regardless of the similar-




Figure 2: LMs’ Average Relative Probability of ‘Yes’ for different conclusion categories and different types of similarities (Word-Sense vs SPoSE). LMs show a clear sensitivity to taxonomic relations, but also show an effect of similarity, where they are more likely to extend the property to a conclusion category when the premise and the conclusion categories are highly similar. Chance behavior is 0.50, as indicated by the dashed line.

ity bin (high vs. low), for both types of similarity. For non-taxonomically related premise conclusion pairs, while the LMs show greater $P_{rel}(\text{Yes})$ value when the premise and conclusion are highly similar, it is seldom above 0.5 (with the exception of Llama-3-8B-Instruct, for the SPoSE similarity). Finally, all LMs show robustness to cases where the property being projected between the premise and conclusions was different (largely high MS values). That is, whatever taxonomic/similarity sensitivities LMs demonstrate are largely for the well-formed instances of the task (i.e., when the properties between premise and conclusion match). These results suggest that **LMs’ property inheritance behavior can be explained jointly by taxonomic relations as well as categorical similarity.**

4 Causal Investigation of Taxonomy and Similarity

Given that LMs perform property inheritance with high taxonomic sensitivity, we now use causal interpretability tools (Mueller et al., 2024a; Geiger et al., 2024a) to localize this behavior. We then investigate (i) whether the subspaces that are causally responsible for inheritance perform inductive generalization in a manner that is better explained by the similarity of the noun concepts, rather than their taxonomic relationship; and (ii) whether taxonomic information and noun similarity are fundamentally entangled in the model representations. We investigate both questions using **distributed alignment search** (DAS; Geiger et al., 2024b), a causal interpretability method which has been used to localize and mechanistically explain syntax-sensitive behaviors (Arora et al., 2024) and arithmetic abilities (Wu et al., 2023b) in LMs. To avoid needing to search over all possible subspaces of an activation vector manually, we employ the Boundless DAS

variant (Wu et al., 2023b) implemented in pyvene (Wu et al., 2024) to learn the appropriate subspace.

We use DAS rather than other causal interpretability techniques⁷ because it allows us to investigate to what extent a specific hypothesized *causal model* is being implemented by the neural network. Because DAS isolates a specific subspace corresponding to a given hypothesized causal variable, the subspace can be evaluated in novel contexts post hoc to better characterize its functional role. We leverage this to evaluate whether subspaces that are sensitive to *taxonomic* relationships are also sensitive to *similarity-based but non-taxonomic* relationships. This type of experimental design involving ambiguous training sets with separate test and generalization sets has been used to characterize the inductive biases of LMs (McCoy et al., 2019; Si et al., 2023; Mueller et al., 2024b), though this has typically been evaluated behaviorally rather than at the level of internal mechanisms. There is a limited but growing literature validating how well discovered mechanisms (or explanations thereof) generalize to novel types of contexts (Geiger et al., 2020; Huang et al., 2023, 2024, *i.a.*).

Given an LM M , a hypothesized causal graph L , and counterfactual input (source, base) pairs (s, b) that target a specific intermediate variable V in the causal graph, DAS learns a rotation R for the activations of a given submodule, and then performs counterfactual interventions in this rotated space. This is done as follows: $L_V(s)$, the value of variable V in L when applied to source s , is patched into the corresponding variable of the base b ,

$$L_V(b) \leftarrow L_V(s),$$

⁷For example, those based on sparse autoencoders (Bricken et al., 2023; Marks et al., 2024; Huben et al., 2024) or neurons (Vig et al., 2020; Finlayson et al., 2021).

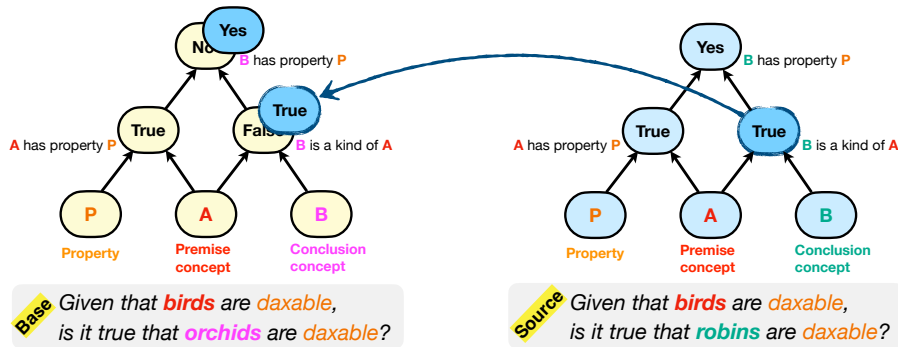


Figure 3: We hypothesize that language models rely on this causal graph to perform the categorical inference task. We are interested specifically in the node responsible for taxonomic judgments, so we use DAS (§4) to isolate the subspace encoding this causal variable. We causally verify its sensitivity to taxonomy by setting its value to what it would have been on counterfactual *source* inputs, and observing whether model behavior changes appropriately. Here, activations at the isolated subspace for the *base* are replaced with activations from the *source*. IIA measures to what extent this intervention results in the expected behavior given the hypothesized causal graph across inputs.

and then the final output of L is computed. Translating this to operations on the low-level computation graph, a rotation R is learned for activations in the network M from a position i to position j . We patch $R(M(s)_i)$ into $M(b)_j$:

$$M(b)_j \leftarrow R(M(s)_i),$$

and then compute the output of M .

After learning the rotation, we compute the interchange intervention accuracy (IIA) on a held-out test set. This scores how frequently the learned subspace performs as expected given the hypothesized causal graph. Intuitively, IIA is defined as the proportion of (s, b) pairs where patching in the rotated source activation into the base model M causes M 's output to match the causal model L 's output *after* the intervention. It thus measures the degree of alignment between the neural network and the causal graph. The intuition for learning a rotation is that the features of interest may not necessarily be aligned to the bases of the original activation space (Hinton et al., 1986; Smolensky, 1986; Elhage et al., 2022; Mueller et al., 2024a), rendering neuron-based approaches insufficient for capturing the target variable. We refer readers to Geiger et al. (2024b) for details.

In our experiments, we hypothesize that a given subspace will be sensitive to taxonomic relationships—if this is true, then passing an activation into this subspace from taxonomically unrelated inputs should yield a variable output of *False*. However, if we intervene on this subspace passing in activations from an alternate input where the nouns *are* taxonomically related (while holding all other activations in the model constant), then the

variable's output should flip to *True*. We illustrate the intuition and the causal graph that we hypothesize LMs implement in Figure 3. By only changing the premise and conclusion nouns across inputs (e.g., *birds*, *robins*, *orchids*), and keeping the rest of input constant, we effectively target the node “**B** is a kind of **A**” in Figure 3.⁸

We train DAS on a subset of 3,000 stimuli and evaluate on the rest. Implementation and training details are given in Appendix D. Next, we describe variations of the training and test sets used in our experiments.

Balanced For each *base* stimulus in the set of 3,000 stimuli, we pick a *source* counterfactual stimulus by sampling from this set of stimuli without replacement. This ensures a balanced coverage of counterfactual stimuli in terms of similarity and whether or not the premise and conclusion are taxonomically related. This experimental setting is used to localize property inheritance in the network.

Control DAS requires training, raising concerns that its expressivity may lead to discovered causal effects where none exist (Arora et al., 2024), similar to concerns raised in the probing literature (Hewitt and Liang, 2019). To mitigate this, we follow Arora et al. (2024) and map the labels *Yes* and *No* to the semantically irrelevant words *chart* and *view*, respectively. This setting evaluates the extent to which DAS memorizes an arbitrary token mapping versus learning task-specific structure.

⁸We target this node since our focus is on taxonomic relations rather than property binding—where the goal is to check if LMs have appropriately bound the property with the noun concept—which has been investigated in other work (Feng and Steinhardt, 2024).

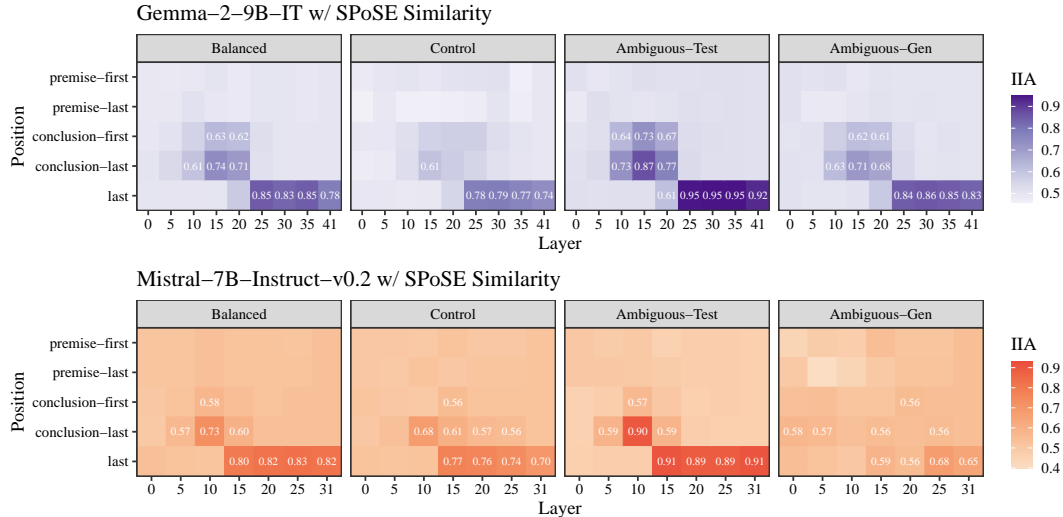


Figure 4: Interchange Intervention Accuracies (IIA) for the causal graph in Figure 3 for **Gemma-2-9B-IT** (Top) and **Mistral-7B-Instruct-0.2** (Bottom), when intervening at various layers and token positions, with negative samples derived using SPoSE similarities. Since the premise and conclusion nouns are often multiple tokens, we show IIA when intervening at the first and last positions of each. **Note:** Both models have different numbers of layers.

Ambiguous In this setting, we filter the training set such that learned rotations can be consistent with *either* similarity or taxonomic features. By only training on taxonomic examples with high similarity and negative examples with low similarity, we learn an intervention which is ambiguous with respect to what features it is using. To disambiguate these, we evaluate on two different test sets: an in-domain (**Ambiguous-Test**) test set composed in the same way as the train set (taxonomic examples with high similarity and negative examples with low similarity), and an out-of-domain generalization test set (**Ambiguous-Gen**), composed of taxonomic examples with low similarity and non-taxonomic examples with high similarity. A drop in IIA values between the Test and Generalization settings would show that learned interventions rely more on similarity than taxonomic features.

Unambiguous Finally, we test whether the subspace sensitive to taxonomic relationships is also sensitive to similarity-based relationships—i.e., whether taxonomic information and noun similarity are entangled in the model. To do this, we take the intervention trained under the **Balanced** setting and evaluate it on (**Ambiguous-Gen**), defined above. Here the training signal for DAS is unambiguous—it is only compatible with capturing taxonomic information. A drop in scores would then indicate that taxonomic and similarity features are entangled.

We use $P_{\text{rel}}(\text{Yes})$ rather than the top predicted token as done in (Wu et al., 2023b) when com-

puting IIA across experiments, in order to have a reasonable comparison against the control setting.⁹

4.1 Results and Analysis

Figure 4 shows the IIA results for Gemma-2-9B-IT and Mistral-7B-Instruct-v0.2 for the first three settings, and Table 3 shows results on the **Unambiguous** setting. Full results for the other models are given in Appendix B. Higher IIA values indicate a better causal alignment between the intervention at a particular location in the network and the causal graph. The activations at the premise token positions have poor alignment with the causal graph, while activations around the last conclusion token at layer 10 (for Mistral) and layers 10-20 (for Gemma) have a much bigger causal effect on property inheritance.

IIA values for the control experiment are always lower than for the balanced setting; thus, despite its expressivity, DAS can yield task-dependent results (i.e., comparatively lower IIA values in the control setting establishes that DAS is telling us something about the LM’s property inheritance abilities). High values for IIA occur in roughly the same positions for Ambiguous-Test and Ambiguous-Gen. There is a substantial drop from Ambiguous-Test to Ambiguous-Gen for all models and for every in-

⁹The interventions are not strong enough to force the model to output either of the control tokens chart, or view as their top-predicted token, resulting in 0% accuracy; on the other hand the difference between using the top predicted token and $P_{\text{rel}}(\text{Yes})$ was negligible in most other settings (<0.02 across experiments, except for Gemma-2-2B-IT).

Model	Word-Sense		SPoSE	
	Bal	Gen	Bal	Gen
Gemma-2-2B-IT	0.86	0.81	0.80	0.63
Mistral-7B-Instruct-v0.2	0.88	0.85	0.83	0.65
Llama-3-8B-Instruct	0.87	0.82	0.80	0.63
Gemma-2-9B-IT	0.90	0.85	0.85	0.67

Table 3: Maximum IIA values for the Unambiguous experiment. Interventions are trained under the Balanced train set, and evaluated on both the Balanced test set (Bal) as well as the Ambiguous-Gen set (Gen). The drop on the generalization test set suggests that the learned interventions rely non-trivially on similarity.

intervention position, showing that, **when given ambiguous data, the learned rotations tend to rely on similarity rather than taxonomic relations.** Similar observations hold for the Unambiguous experiment, with a drop in scores showing that, even when trained on data that is unambiguously consistent with taxonomic relations rather than similarity, the causal model is sensitive to similarity effects. In other words, **similarity and taxonomic features appear to be fundamentally entangled.**

5 Discussion and Conclusions

Our goal in this paper was to characterize category-based inferences in LMs. We derived inspiration from psychological experiments that have sought to understand how humans transfer properties from a higher level category like *mammals* to lower level ones such *dogs* and *cats* (Collins and Quillian, 1969; Glass and Holyoak, 1974; Murphy, 2002)—i.e., show *property inheritance*. Findings from these experiments show that instead of relying purely on abstract category-membership principles (where all members of a category are equally likely to inherit properties), humans show graded behavior proportional to the similarity between the categories (Sloman, 1998).

Through behavioral and representational analyses methods across four LMs, we found that **language models’ property inheritance behavior is sensitive to both taxonomic relations and similarity.** Behaviorally, LMs were more likely to extend the property from the premise to the conclusion when they were taxonomically related, and at the same time, showed positive correlations with the similarity between them. This finding was further reinforced in our representational investigations, where **we found taxonomy and similarity to be largely entangled in the LM subspaces causally responsible for property inheritance.** Our results

cast doubt on previous works’ assumption of pure taxonomic principles with respect to property inheritance in LMs, without considering similarity (Misra et al., 2023; Powell et al., 2024; Wang et al., 2024). In the broader context of LMs’ reasoning behaviors, our results showcase yet another instance where LMs were susceptible to human-like content effects (Lampinen et al., 2024).

Our findings also suggest interesting hypotheses for future human analyses. Importantly, Sloman (1998)’s experiments never tested humans on cases where the categories in a property inheritance setting did *not* share taxonomic relations (e.g., birds and bats). This makes it unclear if the principles of taxonomy and similarity are *also* not mutually exclusive for humans (Murphy et al., 2012)—e.g., it could be possible that both mechanisms are at play, where taxonomic relations demarcate the inference (i.e., whether the property will be inherited) and similarity modulates its strength. Future experiments could test this hypothesis and shed further light on the extent to which LMs and humans show similar sensitivities and biases during reasoning. Overall, our work shows how tools from modern interpretability research, primarily used to isolate sub-structures in neural networks, may also be used to perform hypothesis testing of the abstractions encoded within them.

Limitations

Alternative similarities There are other similarities that we could have used in our experiments. Examples include similarities computed over WordNet (e.g., Wu-Palmer similarity; Wu and Palmer, 1994), co-occurrence based similarities, or similarities computed from LM token embeddings or activations (Chronis and Erk, 2020). While it is tempting to use these, there are various shortcomings to each of them. WordNet similarities do not distinguish between co-hyponyms: concepts that are at the same level will have the same similarity with a given high level concept. Next, the sense based similarity used in this work *is* a form of distributional similarity as it is formed by aggregating distributional semantic embeddings of words occurring in free-text corpora, and has the added advantage of being informed by the sense of the word. Finally, using LM-specific similarities makes it difficult to use the same stimuli across models, since similarities from different models would have resulted in different negative samples. However,

future work could explore possible correlations between LM-derived similarities and their property inheritance behavior.

Abstract and Ad-hoc Concepts We have exclusively focused on concrete, object noun concepts in this work. However, these do not account for the entire range of concepts that humans acquire in their life times—i.e., it is unclear how property inheritance applies to abstract concepts like *love* or ad-hoc concepts (Barsalou, 1983) like *things you would pack for a camping trip*, which may not have taxonomic structure. We leave these investigations to future work.

Knowledge editing Previous work has studied whether language models are able to carry out inferences related to entities (Cohen et al., 2024) or concepts (Wang et al., 2024; Powell et al., 2024) which have been modified through knowledge editing. While we do not engage directly with knowledge edits, it would be interesting to measure the impact of knowledge editing on the LMs’ rotated subspaces during property inheritance.

Broader space of inductive problems While property inheritance is an important consequence of category structure and organization, it is only one type of inductive generalization that humans exhibit (Kemp and Jern, 2014). For instance, humans may also learn about novel concepts and ascribe known features to them (Smith and Estes, 1978; Murphy, 2002; Kemp, 2011), or they might learn a new higher level category and judge category memberships of known concepts (Xu and Tenenbaum, 2007), or they might generalize in a completely non-deductive manner (e.g., from robins to birds, or to its co-hyponyms, see Rips, 1975; Osherson et al., 1990).

Contextual similarities While we experimented with different types of similarities, human inductive inferences might not be explained by the same similarity every time—i.e., similarity in inductive problems can be context-sensitive (Heit and Rubinsten, 1994; Rogers and McClelland, 2004; Kemp and Tenenbaum, 2009). For instance, generalization of biological-sounding properties (*has an ulnar artery*) are often sensitive to taxonomic similarities. Bears and whales are likely to share these properties since they are mammals. On the other hand, behavioral properties (*studies its food before attacking*) might be influenced by similarities derived from other, non-taxonomic conceptual structures.

The aforementioned property could be shared by predatory animals such as tigers and eagles. It is, however, important to note that these caveats apply only when the properties are not *nonce*, like in our work—e.g., it is not clear if *is daxable* or *has feeps* are biological or otherwise. Future studies could explore these deeper nuances of similarity and category-based inferences in LMs using the methods we employ in this work.

Acknowledgments

A.M. is supported by a postdoctoral fellowship under the Zuckerman STEM Leadership program. K.M. acknowledges postdoctoral funding supported by NSF Grant 2139005 awarded to Kyle Mahowald. We also thank Martin Hebart for help on the THINGS dataset and the SPoSE embeddings, Atticus Geiger and Zhengxuan Wu for assistance with pyvene, and Katrin Erk and the anonymous reviewers for helpful discussion and feedback.

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.
- Aryaman Arora, Dan Jurafsky, and Christopher Potts. 2024. [CausalGym: Benchmarking causal interpretability methods on linguistic tasks](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14638–14663, Bangkok, Thailand. Association for Computational Linguistics.
- Lawrence W. Barsalou. 1983. Ad hoc categories. *Memory & cognition*, 11(3):211–227.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Sudeep Bhatia. 2023. Inductive reasoning in minds and machines. *Psychological Review*.
- Sudeep Bhatia and Russell Richie. 2021. Transformer networks of human concept knowledge. *Psychological Review*.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermy, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell,

- Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. 2023. [Towards monosemanticity: Decomposing language models with dictionary learning](#). *Transformer Circuits Thread*.
- Gabriella Chronis and Katrin Erk. 2020. [When is a bishop not like a rook? when it's like a rabbi! multi-prototype BERT embeddings for estimating semantic relationships](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 227–244, Online. Association for Computational Linguistics.
- Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. 2024. [Evaluating the ripple effects of knowledge editing in language models](#). *Transactions of the Association for Computational Linguistics*, 12:283–298.
- Allan M Collins and M Ross Quillian. 1969. Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior*, 8(2):240–247.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. 2022. [Toy models of superposition](#). *Preprint*, arXiv:2209.10652.
- Jonathan St BT Evans. 1989. *Bias in human reasoning: Causes and consequences*. Lawrence Erlbaum Associates, Inc.
- Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2023. [Towards revealing the mystery behind chain of thought: A theoretical perspective](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Jiahai Feng and Jacob Steinhardt. 2024. [How do language models bind entities in context?](#) In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- 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.
- Atticus Geiger, Duligur Ibeling, Amir Zur, Maheep Chaudhary, Sonakshi Chauhan, Jing Huang, Aryaman Arora, Zhengxuan Wu, Noah Goodman, Christopher Potts, and Thomas Icard. 2024a. [Causal abstraction: A theoretical foundation for mechanistic interpretability](#). *Preprint*, arXiv:2301.04709.
- 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.
- Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah D. Goodman. 2024b. [Finding alignments between interpretable causal variables and distributed neural representations](#). In *Causal Learning and Reasoning, 1-3 April 2024, Los Angeles, California, USA*, volume 236 of *Proceedings of Machine Learning Research*, pages 160–187. PMLR.
- Arnold L Glass and Keith J Holyoak. 1974. Alternative conceptions of semantic theory. *Cognition*, 3(4):313–339.
- Gabriel Grand, Idan Asher Blank, Francisco Pereira, and Evelina Fedorenko. 2022. [Semantic projection recovers rich human knowledge of multiple object features from word embeddings](#). *Nature Human Behaviour*, pages 1–13.
- Simon Jerome Han, Keith Ransom, Andrew Perfors, and Charles Kemp. 2022. [Human-like property induction is a challenge for large language models](#). In *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.
- Simon Jerome Han, Keith J Ransom, Andrew Perfors, and Charles Kemp. 2024. [Inductive reasoning in humans and large language models](#). *Cognitive Systems Research*, 83:101155.
- Brett K Hayes and Evan Heit. 2018. [Inductive reasoning 2.0](#). *Wiley Interdisciplinary Reviews: Cognitive Science*, 9(3):e1459.
- Martin N Hebart, Oliver Contier, Lina Teichmann, Adam H Rockter, Charles Y Zheng, Alexis Kidder, Anna Corriveau, Maryam Vaziri-Pashkam, and Chris I Baker. 2023. [Things-data, a multimodal collection of large-scale datasets for investigating object representations in human brain and behavior](#). *eLife*, 12:e82580.
- Martin N Hebart, Adam H Dickter, Alexis Kidder, Wan Y Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I Baker. 2019. [Things: A database of 1,854 object concepts and more than 26,000 naturalistic object images](#). *PloS one*, 14(10):e0223792.
- Evan Heit and Joshua Rubinstein. 1994. [Similarity and property effects in inductive reasoning](#). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(2):411.

- 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, Hong Kong, China. Association for Computational Linguistics.
- G. E. Hinton, J. L. McClelland, and D. E. Rumelhart. 1986. *Distributed representations*, page 77–109. MIT Press, Cambridge, MA, USA.
- Jing Huang, Atticus Geiger, Karel D’Oosterlinck, Zhengxuan Wu, and Christopher Potts. 2023. **Rigorously assessing natural language explanations of neurons**. In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 317–331, Singapore. Association for Computational Linguistics.
- Jing Huang, Zhengxuan Wu, Christopher Potts, Mor Geva, and Atticus Geiger. 2024. **RAVEL: Evaluating interpretability methods on disentangling language model representations**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8669–8687, Bangkok, Thailand. Association for Computational Linguistics.
- Robert Huben, Hoagy Cunningham, Logan Riggs, Aidan Ewart, and Lee Sharkey. 2024. **Sparse autoencoders find highly interpretable features in language models**. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv:2310.06825*.
- Philipp Kaniuth, Florian P Mahner, Jonas Perkuhn, and Martin N Hebart. 2024. A high-throughput approach for the efficient prediction of perceived similarity of natural objects. *bioRxiv*, pages 2024–06.
- Charles Kemp. 2011. Inductive reasoning about chimeric creatures. *Advances in Neural Information Processing Systems*, 24:316–324.
- Charles Kemp and Alan Jern. 2014. A Taxonomy of Inductive Problems. *Psychonomic Bulletin & Review*, 21(1):23–46.
- Charles Kemp and Joshua B Tenenbaum. 2009. Structured Statistical Models of Inductive Reasoning. *Psychological Review*, 116(1):20.
- Brenden M Lake and Gregory L Murphy. 2021. Word meaning in minds and machines. *Psychological Review*.
- Andrew K Lampinen, Ishita Dasgupta, Stephanie CY Chan, Hannah R Sheahan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2024. Language models, like humans, show content effects on reasoning tasks. *PNAS nexus*, 3(7).
- 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*.
- Daniel Loureiro, Alípio Mário Jorge, and Jose Camacho-Collados. 2022. Lmms reloaded: Transformer-based sense embeddings for disambiguation and beyond. *Artificial Intelligence*, 305:103661.
- Gary Lupyan and Molly Lewis. 2019. From words-as-mappings to words-as-cues: The role of language in semantic knowledge. *Language, Cognition and Neuroscience*, 34(10):1319–1337.
- Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. 2024. **Sparse feature circuits: Discovering and editing interpretable causal graphs in language models**. *CoRR*, abs/2403.19647.
- James L McClelland, Matthew M Botvinick, David C Noelle, David C Plaut, Timothy T Rogers, Mark S Seidenberg, and Linda B Smith. 2010. Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, 14(8):348–356.
- 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.
- Kanishka Misra. 2022. **minicons: Enabling flexible behavioral and representational analyses of transformer language models**. *arXiv:2203.13112*.
- Kanishka Misra, Allyson Ettinger, and Kyle Mahowald. 2024. **Experimental Contexts Can Facilitate Robust Semantic Property Inference in Language Models, but Inconsistently**. *EMNLP 2024*.
- 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 Rayz, and Allyson Ettinger. 2023. **COMPS: Conceptual minimal pair sentences for testing robust property knowledge and its inheritance in pre-trained language models**. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2928–2949, Dubrovnik, Croatia. Association for Computational Linguistics.
- 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*.

- Aaron Mueller, Jannik Brinkmann, Millicent L. Li, Samuel Marks, Koyena Pal, Nikhil Prakash, Can Rager, Aruna Sankaranarayanan, Arnab Sen Sharma, Jiuding Sun, Eric Todd, David Bau, and Yonatan Belinkov. 2024a. [The quest for the right mediator: A history, survey, and theoretical grounding of causal interpretability](#). *CoRR*, abs/2408.01416.
- Aaron Mueller, Albert Webson, Jackson Petty, and Tal Linzen. 2024b. [In-context learning generalizes, but not always robustly: The case of syntax](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4761–4779, Mexico City, Mexico. Association for Computational Linguistics.
- Gregory L Murphy. 2002. *The Big Book of Concepts*. MIT press.
- Gregory L Murphy, James A Hampton, and Goran S Milovanovic. 2012. Semantic memory redux: An experimental test of hierarchical category representation. *Journal of Memory and Language*, 67(4):521–539.
- Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. 2023. [Progress measures for grokking via mechanistic interpretability](#). In *The Eleventh International Conference on Learning Representations*.
- Daniel N Osherson, Edward E Smith, Ormond Wilkie, Alejandro Lopez, and Eldar Shafir. 1990. Category-based Induction. *Psychological Review*, 97(2):185.
- Kiho Park, Yo Joong Choe, Yibo Jiang, and Victor Veitch. 2024. [The geometry of categorical and hierarchical concepts in large language models](#). *CoRR*, abs/2406.01506.
- Roma Patel and Ellie Pavlick. 2022. [Mapping language models to grounded conceptual spaces](#). In *International Conference on Learning Representations*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Steven T Piantadosi and Felix Hill. 2022. Meaning without reference in large language models. *arXiv:2208.02957*.
- Derek Powell, Walter Gerych, and Thomas Hartvigsen. 2024. [TAXI: Evaluating categorical knowledge editing for language models](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 15343–15352, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Lance J Rips. 1975. Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, 14(6):665–681.
- Morgane Rivi re, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, L onard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ram e, 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 Rogozinska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucinska, 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 McDonnell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sj sund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, and Lilly McNealus. 2024. [Gemma 2: Improving open language models at a practical size](#). *CoRR*, abs/2408.00118.
- Timothy T Rogers and James L McClelland. 2004. *Semantic Cognition: A Parallel Distributed Processing Approach*. MIT press.
- Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. 2023. [Measuring inductive biases of in-context learning with underspecified demonstrations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11289–11310, Toronto, Canada. Association for Computational Linguistics.
- 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.
- P. Smolensky. 1986. *Neural and conceptual interpretation of PDP models*, page 390–431. MIT Press, Cambridge, MA, USA.
- Paul Smolensky. 1988. On the proper treatment of connectionism. *Behavioral and brain sciences*, 11(1):1–23.

- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart M. Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Xiaohan Wang, Shengyu Mao, Ningyu Zhang, Shumin Deng, Yunzhi Yao, Yue Shen, Lei Liang, Jinjie Gu, and Huajun Chen. 2024. [Editing conceptual knowledge for large language models](#). *CoRR*, abs/2403.06259.
- Peter C Wason. 1968. Reasoning about a rule. *Quarterly journal of experimental psychology*, 20(3):273–281.
- 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 the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Weiqi Wu, Chengyue Jiang, Yong Jiang, Pengjun Xie, and Kewei Tu. 2023a. [Do PLMs know and understand ontological knowledge?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3080–3101, Toronto, Canada. Association for Computational Linguistics.
- Zhengxuan Wu, Atticus Geiger, Aryaman Arora, Jing Huang, Zheng Wang, Noah Goodman, Christopher Manning, and Christopher Potts. 2024. [pyvene: A library for understanding and improving PyTorch models via interventions](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: System Demonstrations)*, pages 158–165, Mexico City, Mexico. Association for Computational Linguistics.
- Zhengxuan Wu, Atticus Geiger, Thomas Icard, Christopher Potts, and Noah D. Goodman. 2023b. [Interpretability at scale: Identifying causal mechanisms in alpaca](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- 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.
- Fei Xu and Joshua B Tenenbaum. 2007. Word learning as Bayesian inference. *Psychological Review*, 114(2):245.
- Charles Y. Zheng, Francisco Pereira, Chris I. Baker, and Martin N. Hebart. 2019. [Revealing interpretable object representations from human behavior](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

A Sensitivity to directionality

Our previous analyses suggest that LMs demonstrate non-trivial taxonomic sensitivity in their property inheritance behavior. Are they also sensitive to the fact that the taxonomic relation is, in its strictest sense, asymmetrical (Murphy, 2002)? For instance, while *robins* are *birds*, one cannot say that *birds* are *robins*. Is this fact captured in LMs’ property projections? That is, if *robins* are *daxable*, are *birds* *daxable*? A finding of insensitivity might not necessarily be a negative result in the context of LMs’ conceptual representations. After all, generalizing from subordinates (*robins*) to superordinates (*birds*) is directly connected to contemporary work targeting property induction in LMs (Misra et al., 2022; Han et al., 2022; Bhatia, 2023; Han et al., 2024), where similarity has more of an uncontroversial role (Osherson et al., 1990; Sloman, 1993), and where humans show a myriad of interesting phenomena (Hayes and Heit, 2018).¹⁰ Instead, it might be informative to further characterize the LMs’ subspaces responsible for property inheritance—our previous findings suggest that they are sensitive to similarity, are they also sensitive to directionality? Or are they simply encoding the presence of taxonomic (and similarity) relations?

We behaviorally evaluate the sensitivity of LMs to the direction of the property projection as follows. We measure the **Directional Sensitivity** (DS) as the fraction of times an LM flips its relative preference from Yes to No when the order of the premise and conclusion nouns in our stimuli is swapped, *only* over all taxonomically related items. Additionally, we also measure the spearman correlation of DS with that of TS—i.e., the taxonomic sensitivity metric computed for the non-reversed property inheritance stimuli, computed for the tax-

¹⁰We refer to this analysis as the *directionality* of property projection given that *induction* is often used more generally to refer to many kinds of reasoning which are *ampliative*, and not determined by deductive logic (Kemp and Jern, 2014).

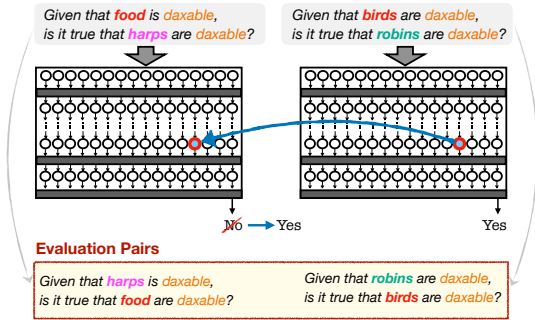


Figure 5: DAS interventions on pairs where the direction of the inference is flipped. Low IIA values for this experiment reveal whether the learned subspaces for property inheritance are sensitive to direction.

onomically related premise and conclusion. We refer to this as ρ_{DS}

Next we consider whether the learned DAS interventions from §4 generalize when reversing the direction of the premises and conclusions. We consider interventions trained under the **Balanced** setting with high IIA values (i.e., which localized property projection behavior in the regular property inheritance direction and showed taxonomic sensitivity), and evaluate them on source and base pairs with flipped premise and conclusion nouns. In order to measure the effect of direction, we only evaluate on examples where the intervention successfully flipped the label from No to Yes in the non-reversed direction, as illustrated in Figure 5. The evaluation set is constructed from source and base pairs where the premise and conclusions are taxonomic and non-taxonomic, respectively. We then filtered these so that (i) the LM predicts the right labels in the non-reversed case (i.e., shows taxonomic sensitivity) and (ii) the DAS intervention correctly flips the label from No to Yes. The order of the nouns are then reversed to establish whether the intervention causes the same behavior when the direction of inheritance is flipped. This is illustrated in Figure 5.

We refer to the IIA values for this evaluation as **Subspace Directional Insensitivity (SDI)**, since higher scores mean the subspace is *less* sensitive to direction.

Results The behavioral DS results across models are given in Table 4. Gemma-2-9B-IT is the most directionally sensitive model: the order of the premise and conclusion concepts has an impact on property projection roughly half of the time. The other three models are far less sensitive to direction. This is also reflected in higher Spearman correla-

Model	DS	ρ_{DS}
Gemma-2-2B-IT	0.24	0.54
Mistral-7B-Instruct-v0.2	0.37	0.37
Llama-3-8B-Instruct	0.38	0.40
Gemma-2-9B-IT	0.49	0.31

Table 4: Directional sensitivities of LMs. DS: Directional Sensitivity; ρ : Spearman correlation of $P_{rel}(\text{Yes})$ between the original and reverse directions. DS and ρ are computed only over the taxonomic items.

Model	Word-Sense		SPoSE	
	Con- ℓ	Last	Con- ℓ	Last
Gemma-2-2B-IT	0.71	0.77	0.77	0.83
Mistral-7B-Instruct-v0.2	0.40	0.31	0.29	0.38
Llama-3-8B-Instruct	0.06	0.52	0.16	0.45
Gemma-2-9B-IT	0.08	0.57	0.07	0.39

Table 5: SDI results with interventions at conclusion-last token position (Con- ℓ) and final token position (Last). Lower numbers indicate that the subspace involved in property inheritance in the deductive case is sensitive to the directionality of the inference.

tions for those models.

Subspace Directional Insensitivity results are shown in Table 5 for interventions at different network locations. Since IIA values were high at both the final token position (Last) and the last token of the conclusion noun (Con- ℓ), we evaluated with interventions at both, in each case selecting the layer with highest IIA across each token position.¹¹

SDI is nearly always lower for Con- ℓ than for Last positions: it is harder to use the taxonomically sensitive property inheritance intervention to make the LM behave in the same way with reversed examples when intervening earlier in the network. Similar to the behavioral results, the subspace at conclusion-last token position and layer 15 for Gemma-2-9B-IT is the most sensitive to the directionality of the inference, while Gemma-2-2B-IT is the least sensitive. We note, however, that this only holds for the prompts used in our experiments—for example, using *Prompt 2* with chat template (Appendix C) makes Gemma-2-2B-IT more directionally sensitive than it is with *Prompt 1* without chat template.¹²

¹¹For Gemma-2-9B-IT, this was layer 15 for Con- ℓ , and layer 35 for Last; for the other models this was layer 10 for Con- ℓ and layer 15 for Last.

¹²Specifically, DS increases from 0.24 to 0.55, and SDI decreases (e.g., for Con- ℓ , to 0.22 and 0.27 for Word-Sense and SPoSE, respectively). Future work could systematically explore direction sensitivity for different prompts.

B Additional DAS Results

Additional results for the DAS experiments for Gemma-2-2B-IT and Llama-3-8B-Instruct are given in Figure 6. This figure also includes the full set of results from using the sense-based similarity measure to derive our negative samples. The full set of results from the **Unambiguous** DAS experiments are shown in Figure 7.

C Prompts

The formats for the prompts used in our experiments are given below.

Prompt 1: Answer the question. Given that **A** is/are *daxable*, is it true that **B** is/are *daxable*? Answer with Yes/No.\n

Prompt 2: Answer the question. Given that **A** is/are *daxable*, is it true that **B** is/are *daxable*? Answer with Yes/No. The answer is:

Prompt 3: Answer the question. Given that **A** is/are *daxable*, is it true that **B** is/are *daxable*?<0x0A>Answer with Yes/No.<0x0A>

Prompt 4: Given that **A** is/are *daxable*, is it true that **B** is/are *daxable*? Answer with Yes/No:

Table 6 indicates which the prompts used for each language model and whether the chat template¹³ was used. These selections were made on the basis of the models’ behavioral taxonomic sensitivities.

Model	Prompt	Chat Template
Gemma-2-2B-IT	Prompt 1	No
Mistral-7B-Instruct-v0.2	Prompt 3	No
Llama-3-8B-Instruct	Prompt 2	No
Gemma-2-9B-IT	Prompt 2	Yes

Table 6: Prompt formats used in our experiments.

D DAS Implementation Details

Boundless DAS interventions were trained for 2 epochs, with a batch size of 16 on NVIDIA RTX A6000 and NVIDIA A40 GPUs. We used the Adam optimizer with a learning rate of 1e-3, including gradient accumulation and a temperature schedule.¹⁴ All DAS experiments were performed

¹³https://huggingface.co/docs/transformers/main/en/chat_templating

¹⁴As used in https://github.com/stanfordnlp/pyvene/blob/main/tutorials/advanced_tutorials/Boundless_DAS.ipynb

using the pyvene library, version 0.1.1. We loaded all models using the default huggingface configuration, except for Gemma-2-9B-IT, which we loaded with `torch_dtype=torch.bfloat16` due to memory constraints. For the **Balanced**, **Unambiguous** and **Control** experiments, we trained on 3,000 stimuli and evaluated on 1,018 stimuli. For the **Unambiguous** setting we trained on 1,527 stimuli for the Word-Sense dataset and 2,017 stimuli for the SPoSE dataset. The **Ambiguous-Test** set contained 536 stimuli for Word-Sense and 671 for SPoSE, while the **Ambiguous-Gen** test set contained 482 stimuli for Word-Sense and 347 for SPoSE.

E Premise and Conclusion Concepts

The list of all premise categories used in our experiments is shown in Table 9. Tables 7 and 8 show examples from different slices of our data using both SPoSE and sense-based embeddings.

Slice	Pair	Similarity	$P_{rel}(\text{Yes})$
Taxonomically related, High similarity	<i>garden tool</i> → <i>shovel</i>	0.83	0.99
Taxonomically related, Low similarity	<i>garden tool</i> → <i>hose</i>	0.72	0.76
Non taxonomically related, High similarity	<i>garden tool</i> → <i>spear</i>	0.81	0.55
Non taxonomically related, Low similarity	<i>garden tool</i> → <i>ham</i>	0.13	0.01

Table 7: Examples of the four slices in our data, where similarity is calculated using SPoSE (Zheng et al., 2019; Hebart et al., 2023), and the model responses are from **Gemma-2-9B-IT**.

Slice	Pair	Similarity	$P_{rel}(\text{Yes})$
Taxonomically related, High similarity	<i>sea animal</i> → <i>whale</i>	0.79	0.99
Taxonomically related, Low similarity	<i>sea animal</i> → <i>coral</i>	0.70	0.74
Non taxonomically related, High similarity	<i>sea animal</i> → <i>aquarium</i>	0.75	0.42
Non taxonomically related, Low similarity	<i>sea animal</i> → <i>rack</i>	0.45	0.05

Table 8: Examples of the four slices in our data, where similarity is calculated using the sense-based embeddings (Loureiro et al., 2022), and the model responses are from **Gemma-2-9B-IT**.

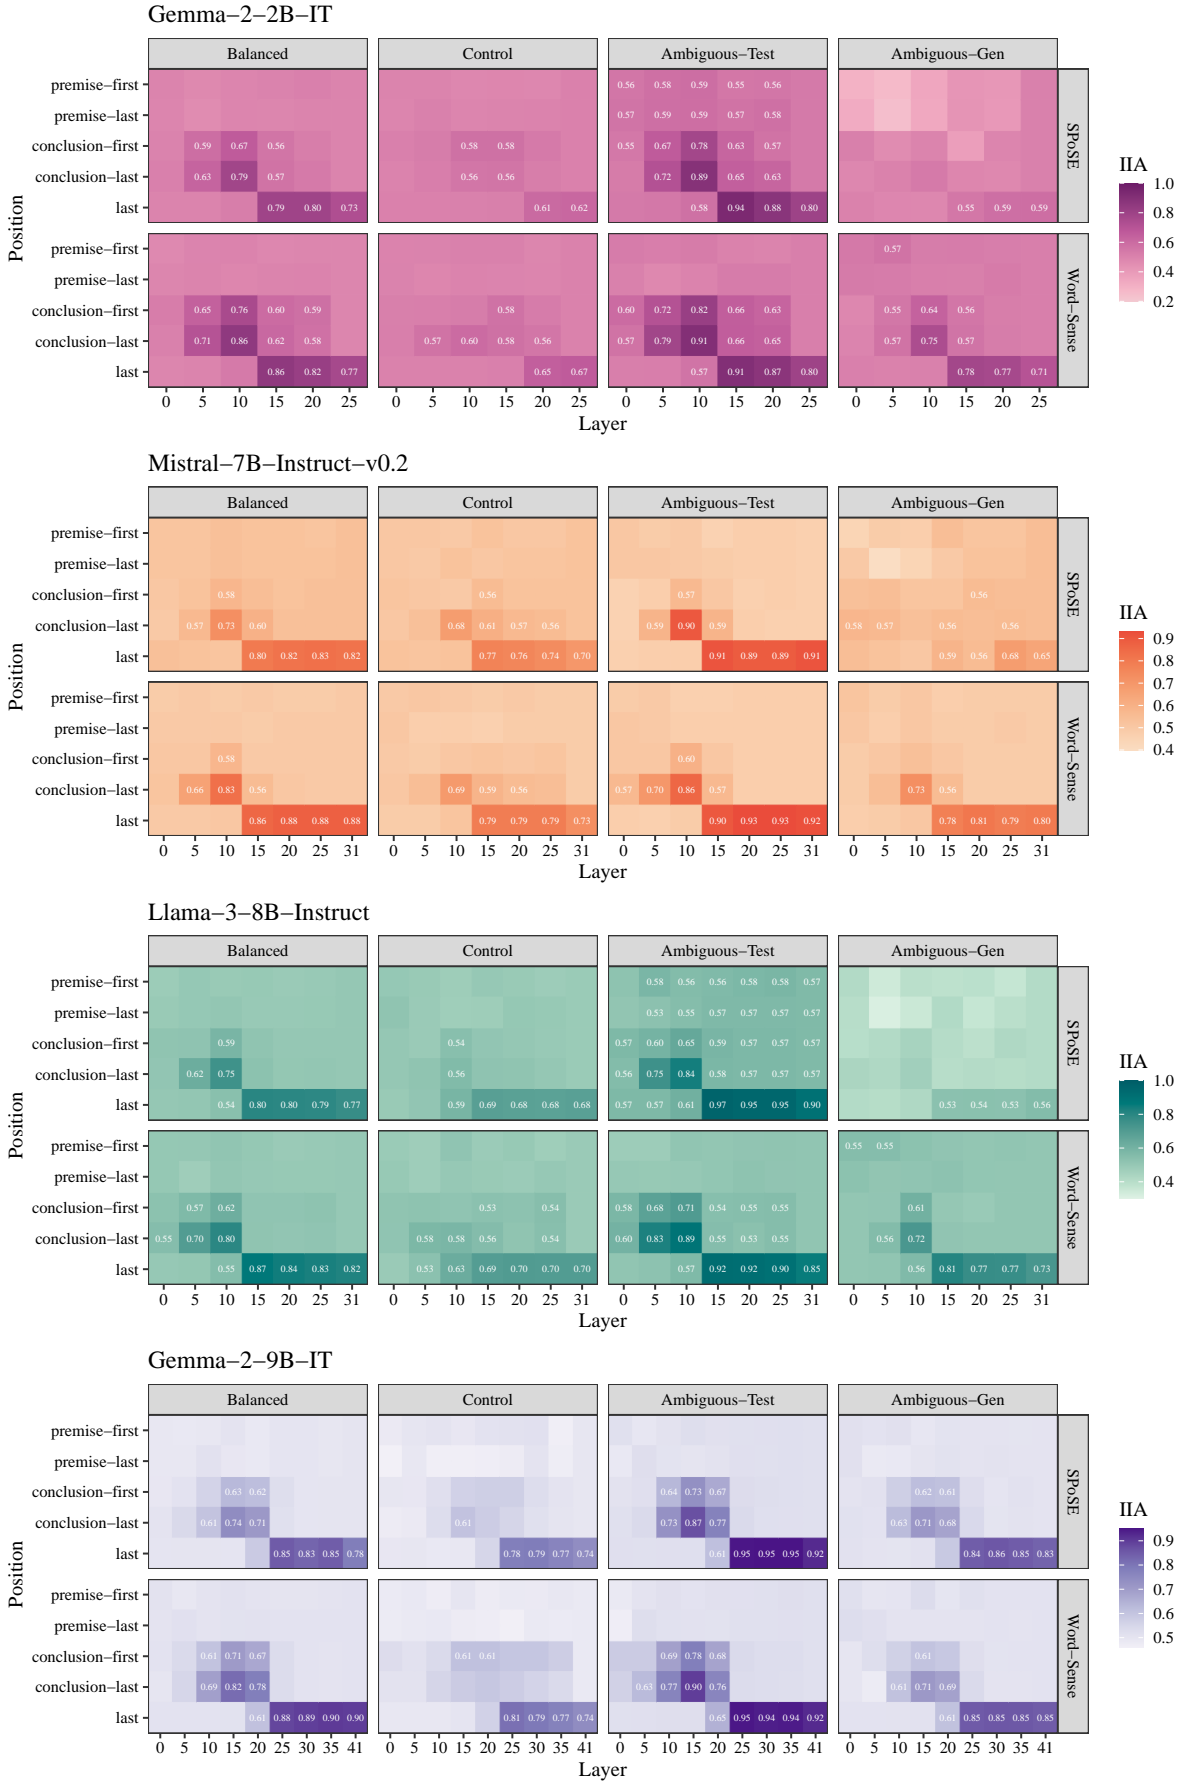


Figure 6: Interchange Intervention Accuracy (IIA) for the causal graph in Figure 3, for all LMs.

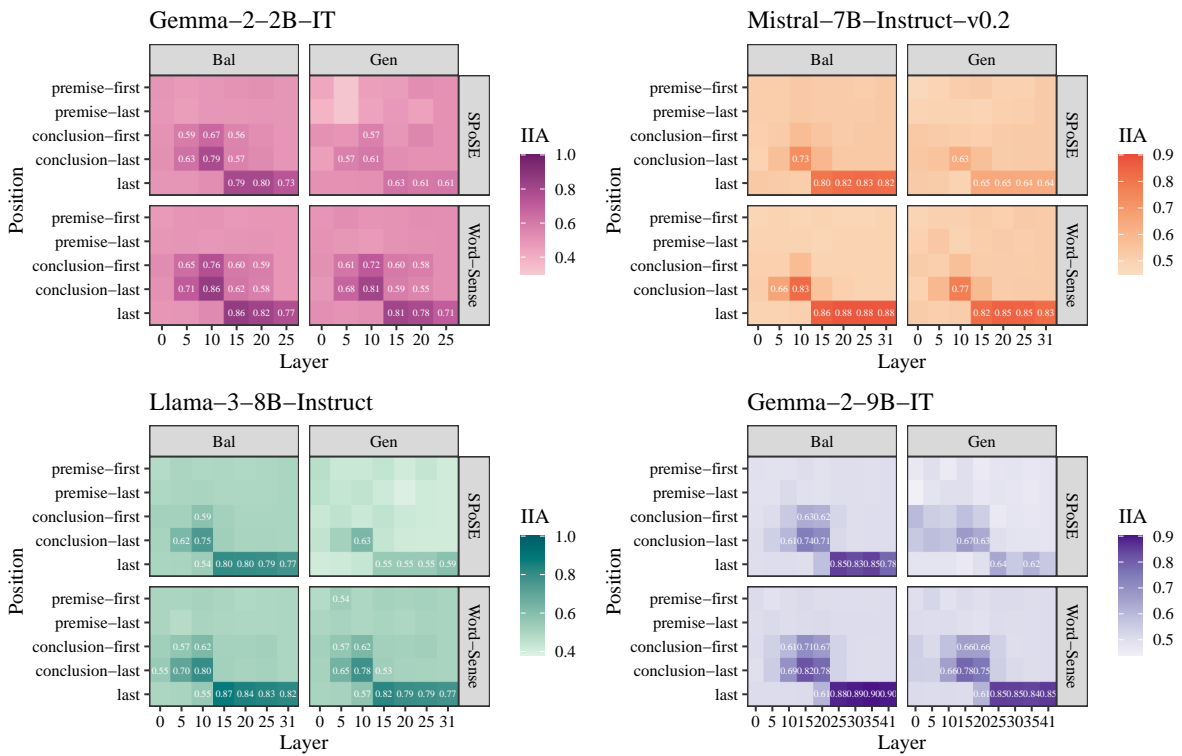


Figure 7: Interchange Intervention Accuracies (IIA) obtained on the Balanced and Generalization sets in the Unambiguous experiments, for all LMs.

premise category	<i>N</i>
food	291
animal	177
tool	142
clothing	107
container	105
mammal	88
electronic device	74
vehicle	70
weapon	48
plant	47
home decor	45
vegetable	42
accessory	37
dessert	36
furniture	36
breakfast	35
fruit	34
musical instrument	33
toy	33
fastener	31
toiletry	31
auto part	30
sea animal	30
bird	28
kitchen tool	27
medical equipment	26
school supply	26
office supply	24
seafood	24
kitchen equipment	20
drink	19
game	19
headwear	19
water vehicle	19
women's clothing	19
livestock	18
garden tool	17
insect	17
outerwear	16
protective clothing	16
candy	15
condiment	15
footwear	15
jewelry	15

Table 9: List of premise categories and their sizes (in terms of number of members), sorted by size. Many of these categories overlap—e.g., mammals are also animals, etc.