



HAL
open science

Conceptual Hierarchy in Child-Directed Speech: Implicit Cues are More Reliable

Kyra Wilson, Michael Frank, Abdellah Fourtassi

► **To cite this version:**

Kyra Wilson, Michael Frank, Abdellah Fourtassi. Conceptual Hierarchy in Child-Directed Speech: Implicit Cues are More Reliable. *Journal of Cognition and Development*, 2023, 24 (4), pp.563-580. 10.1080/15248372.2023.2178436 . hal-04521280

HAL Id: hal-04521280

<https://hal.science/hal-04521280>

Submitted on 26 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Conceptual Hierarchy in Child-Directed Speech: Implicit Cues are More Reliable

Kyra WILSON¹, Michael C. FRANK¹, and & Abdellah FOURTASSI²

¹ Department of Psychology

Stanford University

² Department of Computer Science

Aix-Marseille University

Author Note

All data and analytic code are available at <https://github.com/afourtassi/concepts>

A preliminary version of this work was presented in the cognitive science society conference and appeared in the non-archival proceedings as:

Fourtassi, A., Wilson, K., & Frank, M., C. (2020). Discovering Conceptual Hierarchy Through Explicit and Implicit Cues in Child-Directed Speech. In Proceedings of the 42nd Annual Meeting of the Cognitive Science Society

Correspondence concerning this article should be addressed to Abdellah FOURTASSI, Postal address. E-mail: abellah.fourtassi@univ-amu.fr

Abstract

In order for children to understand and reason about the world in an adult-like fashion, they need to learn that conceptual categories are organized in a hierarchical fashion (e.g., a dog is also an animal). While children learn from their first-hand observation of the world, social knowledge transmission via language can also play an important role in this learning.

Previous studies have documented several cues in parental talk that can help children learn about conceptual hierarchy. However, these studies have used different datasets and methods that have made it difficult to compare the relative usefulness of various linguistic cues to conceptual knowledge and to test whether they scale up to naturalistic speech. Here, we study a large-scale corpus of English child-directed speech – collected in North America and the UK – and used a unified classification-based evaluation method which allowed us to investigate and compare cues that vary in terms of how explicit the information they offer is. We found the more explicit cues to be too sparse or too noisy in child-directed speech, making them unlikely to support robust learning. In contrast, the implicit cues offered a more reliable source of information. Further, we investigated developmental changes from 3 to 6 years of age, and we found no differences in the availability of these cues in the input. Our work confirms the utility of caregiver talk for conveying conceptual information and supporting the development of early taxonomic knowledge. It provides a first step toward a cognitive model that would combine perceptual- and language-based mechanisms, leading to a more comprehensive account of children’s conceptual development.

Keywords: Conceptual learning, child-directed speech, language and cognition

Conceptual Hierarchy in Child-Directed Speech: Implicit Cues are More Reliable

Introduction

A hallmark of human conceptual knowledge is its hierarchical organization. For example, the same entity (e.g., a dog) can be considered in a nested structure as, e.g., a husky, a dog, a mammal, and an animal. Such hierarchical organization is fundamental to human cognition as it allows, among other things, the generalization of knowledge through inference. For example, upon children’s learning that female mammals produce milk to feed their young, they can infer that female giraffes also produce milk because the category “giraffe” is *included* in the category “mammal” (e.g., Inhelder & Piaget, 2013; Markman, 1989; Murphy, 2004; Sloutsky, 2010). The current study asks how this hierarchical knowledge develops. More specifically, we ask whether children can learn conceptual hierarchy from the *language* they hear around them.

Although taxonomic knowledge takes time to develop into an adult-like form (Blaye et al., 2006; Lucariello et al., 1992; Unger et al., 2020), researchers have noted that children as young as 3 years old already show signs of hierarchical knowledge in their early lexicon. For example, observational studies have found that children use superordinate words like “food” and “animal” (e.g., Fenson et al., 1994; Frank et al., 2021), and they use different labels to refer to the *same* object, shifting their conceptual perspective from one level of abstractness to another (e.g., using the labels “dog” and “animal” to refer to a dog) (e.g., Clark, 1997). In more controlled in-lab experiments, preschool children are able, depending on the situation, to interpret the meaning of a novel word (e.g., “Dax”) either at the basic or at the superordinate level (Callanan, 1989; Markman & Hutchinson, 1984). Critically, children do not simply extend the meanings of words based only on perceptual similarity. Under some circumstances, e.g., when using the label to name multiple objects from the target superordinate category, they do prefer taxonomic relations even when a perceptually similar — but taxonomically unrelated — alternative is available (Gentner & Namy, 1999;

Liu et al., 2001).

How do children begin acquiring the hierarchical structure of semantic knowledge? Early accounts considered this learning to be the consequence of the emergence of a *domain-general* logic of class inclusion — that one category can be part of a larger one — which develops through middle childhood (Inhelder & Piaget, 2013; Winer, 1980). The logic of classes may represent a mature, adult-like structuring principle of concepts, enabling reasoning on both familiar and unfamiliar domains. Nevertheless, the above-reviewed evidence suggests that children begin showing, much earlier, the ability for hierarchical reasoning in *specific domains* they are familiar with or interested in (e.g., the domain of food or dinosaurs). This fact suggests that the input they receive in specific domains (whether through perceptual or linguistic means) plays an important role in shaping this organization (See also Carey, 1987; Chi et al., 1989; Inagaki & Hatano, 2002).

There are two types of input that may serve to help children learn this type of hierarchical conceptual structure: perceptual/functional input and language input. According to perceptual/functional accounts, children rely on their first-hand observation of the world, allowing them to form abstract concepts using primarily their ability to pick up on similarities between entities in terms of their perceptual features or functions (Madole & Oakes, 1999; McClelland & Rogers, 2003; Quinn & Eimas, 2000; Sloutsky, 2015, 2010; Smith & Heise, 1992). Using their inductive reasoning skills, they synthesize these concepts into a nested tree structure via performing some sort of hierarchical clustering (Kemp et al., 2007; Saxe et al., 2019; Tenenbaum et al., 2006). Finally, they learn to attribute multiple labels to the same entity (e.g., dog, mammal, animal), with more abstract/broad labels mapping onto larger regions in the nested tree structure (Xu & Tenenbaum, 2007).

Direct perceptual/functional input is an essential source of information for children's early conceptual development. However, the current study focuses on an additional source of information – knowledge transmitted from other people via language input. Indeed,

superordinate concepts (more than basic-level ones) do not necessarily share similar — sensory accessible — features. For example, the categories “animal” and “plants” are organized for Groote Eylandt Australian aborigines rather into three categories: “biological”, “food”, and “totemic” (Waddy, 1982). Thus, learning conceptual hierarchy requires additional, culture-specific input, primarily via language (Csibra & Gergely, 2009; Gelman, 2009; Harris, 2012).

Regarding how children may learn conceptual hierarchy from language, observational studies have noted that when parents introduce words at the superordinate level, they typically also provide the basic level term (Callanan, 1985; Shipley et al., 1983). For example, parents rarely point to an object and say “this is an animal!” Instead, they usually *anchor* the superordinate word “animal” at the basic level by saying something along the lines of “This is a duck; a duck is a kind of animal.” Such an anchoring strategy provides children with a categorization of the same object at different levels, which may help them understand the hierarchical organization.

Also within the language-based account, more recent research, prompted by advances in machine learning (e.g., Landauer & Dumais, 1997; Mikolov et al., 2013), found that the statistical distribution of basic-level terms in parental speech can lead to coherent structures at the superordinate level (Fourtassi et al., 2019; Frermann & Lapata, 2016; Huebner & Willits, 2018). For example, though “fish” and “bird” do not look very similar, people talk about them in similar linguistic contexts, typically leading to similar distributions in speech. Having a shared pattern of co-occurrence cues taxonomic relations and, in fact, several studies have suggested that it does help to acquire such relations in development (Brown & Berko, 1960; Ervin, 1961; McNeill, 1963; Sloutsky et al., 2017).

These linguistic cues can be thought of as ends in a continuum that varies from explicit to implicit. The “is-a-kind-of” cue is the most explicit cue since both the labels (i.e., the basic and superordinate labels) and their hierarchical relationship are explicitly stated.

The “distributional cue” is the most implicit cue since, on the one hand, the superordinate term is not required and, on the other hand, the hierarchical relationship (that is, the fact that co-occurring basic-level terms are part of a higher-level category) can only be induced.

While the above-reviewed studies have focused on these extremes, one can think of other cues that have an intermediate status on this continuum. We introduce and test two such cues. The first, which we name the “pragmatic cue,” is when parents hint at the hierarchical relationship between two concepts pragmatically without using an explicit inclusion expression. For example, instead of saying “a cow is a kind of animal,” parents can say the following: “Do you want a cow or do you want another animal?” (see Table 1 for more examples). In addition, we also study an “Affordance-based cue” whereby linguistic labels for action affordances provide another implicit linguistic cue for category membership. For example, food items could be identified as members of a single superordinate category by virtue of their affordance of “eating,” which translates linguistically to food labels serving as objects of the verb “eat.” In other words, one can infer that “apple” and “banana” are related taxonomically to a higher-level category of items that afford eating (i.e., food) by virtue of their occurrence in sentences like “Dad eats an apple!” and “would you like to eat a banana?” (see Figure 1).

The current study

Previous studies examining individual linguistic cues to conceptual hierarchy have varied in terms of both the datasets and methods they used, which has made comparison difficult, thus hindering cumulative scientific progress on this question. On the one hand, implicit cues have generally been studied using large-scale data and have been evaluated based on their ability to provide an accurate similarity space for words. On the other hand, explicit cues have been studied mainly in the context of small-scale experiments and have been tested mainly by counting the frequency of a given linguistic expression (e.g., “X is a kind of Y”).

In this study, we make a systematic comparison of explicit and implicit cues using similar methods. In particular, we quantify the information provided by each of these cues in Figure 1 using spontaneous – rather than scripted – child-directed speech, using language data from naturalistic child-caregiver interactions in North America and the UK (MacWhinney, 2000). We test information contained in this input data across the entire early lexicon as measured by the MacArthur-Bates Communicative Inventory (CDI) (Fenson et al., 1994). We focus on the hierarchical relations between basic-level terms in this early lexicon in relation to six superordinate categories: Animals, Furniture, Clothes, Food, Toys, and vehicles. This set of superordinate concepts was chosen because it combines all the concepts that had been studied previously and for which CDI data were available. Further, a data-driven approach using CDI words and input from CHILDES (Fourtassi et al., 2019) led to a set of categories made of “animal”, “food,” “clothes,” and a broad category of “artifacts” which included instances of toys, vehicles, and furniture. Here, we kept the last three categories differentiated as in previous experimental work.

Our research question is to investigate the relative usefulness of each linguistic cue in providing information that children can utilize for learning hierarchical conceptual relationships. To address this question, we proceed as follows. We begin by introducing the child-directed data and the derived linguistic cues for conceptual hierarchy. We then explain the task and evaluation methods that we use to study and compare these cues. Next, we present the results quantifying the relative informativeness of each of the four cues by 3 three years of age as well as their change in development up to 6 years old. Finally, we discuss our findings in light of the literature on early linguistic and conceptual development.

Methods

Data

We constructed a large-scale corpus by aggregating all English-language transcripts (collected in North America and the UK) from the version CHILDES 2018.1, downloaded

from chldes-db (MacWhinney, 2000; Sanchez et al., 2019) We selected all utterances in which the speaker was not tagged as “Child” or “Target_Child.” In other words, we only kept input from adults, excluding language produced by children. In the main analyses, we limit the input to the speech addressed to children up to three years of age as this is the age where early signs of conceptual hierarchy appear according to the above-reviewed developmental literature. That said, we also include results from speech addressed to children up to six years to investigate potential developmental changes in the input. The final corpus contained around 2 million utterances and 8.3 million words from a collection of 4,939 transcripts across $N = 660$ children (mean age = 26.5 months, $SD = 5.9$, median = 27.5). We decided to study the six following superordinate categories: “animal”, “furniture”, “clothes”, “food”, “toys” and “vehicles.” For each of these categories, we used the corresponding basic-level terms available in the English-language MacArthur-Bates Communicative Inventory (CDI) (Fenson et al., 1994; Frank et al., 2017), a parent-report instrument that provides a partial listing and categorization of words produced by children 18–30 months.

Cues to Conceptual Hierarchy and their Feature Vectors

As shown in Figure 1, we consider four linguistic cues to conceptual hierarchy, ranging from most explicit to most implicit: “is-a-kind-of”, pragmatic, affordance-based, and distributional cues. To facilitate comparison, we represented all these cues as vectors and we tested how these vectors allow us to classify basic-level terms into superordinate categories. Below we describe in detail how these vectors were constructed. To summarize, for the “is-a-kind-of” and the pragmatic cues where the hierarchical relation relies on the explicit mention of the superordinate words, the vectors were based on the six superordinate words, and each cell in the vector corresponded to the frequency with which the hierarchical relation was found between the basic-level word at hand and the superordinate word in the current cell (see an illustration in Figure 2, left). For the affordance-based cue, the vector was constructed in a similar fashion, except that the superordinate terms were replaced with their corresponding affordances (e.g., “eatable” for food and “wearable” for clothes) and each of

the six cells in the vector corresponded to the frequency with which the basic-level word has been used with the verb that marks the affordance. Finally, the implicit distributional cue does not rely on an explicit category marker. Instead, the feature vector was an “embedding” in a high dimensional space derived based on the words’ shared pattern of co-occurrence.

Is-a-kind-of

This cue tests the extent to which parents use explicit expressions of class inclusion (Callanan, 1985). For each word at the basic level, X (e.g., cow), we construct a feature vector of length 6, where each cell corresponds to one of the 6 superordinate categories under study, Y_i (e.g., animal or food). The value in cell i corresponds to the frequency with which X appears with Y_i in one of the following expressions: “ X is a/an Y_i ” or “ X is a kind/type/sort of Y_i ”.

Pragmatic

Parents can express conceptual hierarchy between X and Y_i without necessarily using an explicit “is-a-kind-of”-like expression. In many cases, parents can hint at this hierarchy using a wide diversity of linguistic expressions (see Table 1 for examples from our dataset). The precise and exhaustive characterization of these expressions at scale is an open computational challenge. That said, as a first attempt to capture the diversity with which parents hint at conceptual hierarchy in a pragmatic fashion, we relaxed grammatical constraints between X to Y_i , and we kept only the requirement that X and Y_i should co-occur. More concretely, we represent each basic-level term, X , with a feature vector where each cell contains the frequency with which X co-occurs with the corresponding superordinate term Y_i (see Figure 2, left). This co-occurrence is determined using a fixed window of k utterances. Values of $k > 1$ allow us to capture the case where a relationship between X and Y_i is established across more than one utterance. For example:

- Mother: What kind of animal is this?

- Mother: It’s a giraffe!

Affordance-based

The superordinate label is not the only category marker that can cue conceptual hierarchy for a basic level term, especially when this category can also be characterized by its affordance. For example, “food” can be characterized as the category of things we eat and “clothes” as things we wear. Thus, children can learn that some concepts (e.g., “apple” and “bread”) are parts of a higher-level category (“things we eat”) by observing how these concepts co-vary with a cue of their common affordance (i.e., the verb “eat”).

We computed the feature vectors for the affordance-based cue as follows. In a first step, we tried to find a verb that could be used as an affordance marker for an entire category. We used “eat” for food, “wear” for clothes, “play” for toys, and “ride” for vehicles. The category “furniture” has no such obvious function verb. We decided to use the verb “use” because if there were a verb that could fit every member of the category of furniture, it would be that (even though it can also fit things that are not members of the category). For the animal category, we could find no verb that could categorize the instances.¹ In addition to these verbs, we also identified synonymous verbs for each category that could also signify an affordance for the category. However, these verbs were either not found in the corpus (words like “devour” and “utilize” are not used in child-directed speech) or they occurred in too many contexts to be useful for categorization (“have” can be used synonymously with “eat”, but it has additional meanings that make this not a useable cue). Because of this, only one verb was used for computing the feature vectors for each of the categories.

We detected the affordance-based relationships, syntactically, based on occurrence in a verb-complement structure.² For example, in the utterance “the bird eats the berries,” the

¹ This does not mean that animal terms cannot be used with verbs from other categories such as “ride” (vehicles) or “eat” (food). Thus, even if there was no specific verb for the category “animal,” we still produced word vectors for animal terms based on their use with such verbs.

² There are more complex structures that could, in principle, be used by parents. Here we used the simplest.

word “berries” was categorized as “eat”-able. For each basic-level term X , we computed a 6-cell-long feature vector where the value in each cell corresponded to the frequency with which X_i occurs in a verb-complement relationship with the verb/affordance Y_i at hand.

Distributional cue

Unlike the first three cues, the pure distributional cue is not based on an explicit category marker at the superordinate level. It is based, instead, only on the way basic-level terms are distributed together in speech. According to the distributional hypothesis (Harris, 1957), words that share similar patterns of co-occurrence (or distribution) in speech tend to be semantically close. For example, a child can posit that the words “apple” and “banana” must refer to objects that share semantic properties (and likely belong to the same superordinate category) because people around them tend to talk about both words in similar context using similar co-occurring words such as “eat,” “kitchen,” “dessert,” and “tree.”

Following previous research (Fourtassi et al., 2019), we quantified this cue using a model called Word2Vec, a prediction-based instantiation of the distributional hypothesis (Mikolov et al., 2013). In brief, Word2Vec is a neural network model that maximizes the likelihood of predicting the linguistic context given a word (or predicting a word given the context). For each prediction made, the model derives an error signal obtained by comparing the predicted vs. observed context. The error signal is then backpropagated through the neural network, improving the ability of future predictions. The trained model outputs a high-dimensional semantic space where words are represented as continuous vectors (or “embeddings”). Words that predict (or are predicted by) similar linguistic contexts will end up having vectors that are close in this semantic space. This distributional model and other variants have been applied fruitfully to study language development from word forms to the lexical-semantic organization through word meanings (e.g., Andrews et al., 2009; Fourtassi, 2020; Fourtassi et al., 2014; Frermann & Lapata, 2016; Hills et al., 2010; Huebner & Willits, 2018; Stella et al., 2017). Here we use Word2Vec to represent basic-level words as vectors in

a high-dimensional space, representing the distribution of these words in a latent semantic structure.

Task and Evaluation

Above, we explained how we characterized all cues in a unified formal framework. This framework allows us to directly compare the cues thanks to the similarity measures that we can derive from the vector-based representations. The similarity measure is commonly defined as the cosine of the angle formed by two vectors. Using this similarity, we test the ability of each cue to predict which pairs of basic-level words belong to the same superordinate category (e.g., “apple” and “bread”) and which pairs of words may belong to different categories (e.g., “apple”, “horse”) (see Figure 2, right).

More precisely, we listed all pairs of basic-level words in the CDI dataset and their four cosine similarities (from each of the four cues). Then, we evaluated the ability of each similarity measure to accurately predict whether the pairs belonged to “same” or “different” categories. We quantified performance in the task using a standard signal detection measure called the Area Under the ROC curve (hereafter AUC). The AUC score can be interpreted as the probability that, given two pairs of words (e.g., “apple”/“bread” and “apple”/“horse”), of which one is from the same category (i.e., “apple”/“bread”), the pairs are correctly classified. This probabilistic interpretation means that AUC values range between 0.5 (i.e., the cue is performing at chance) and 1 (i.e., the cue is perfect). We derived both a global AUC score across all six categories and a category-specific AUC score where we evaluated only the subset of pairs of words that contained at least an instance of a target category.³

³ A similar task and evaluation method have been used in previous work to evaluate cues to phonological categories in early development (Fourtassi et al., 2014; Fourtassi & Dupoux, 2013).

Results

Individual Cue Results

Instances of the most explicit cue type, the “*is-a-kind-of*” cue, were so rare that we could not even build feature vectors for basic-level words. In total, we found only four instances, all of them characterizing the “animal” category. Thus, we did not have meaningful results to report for this cue. Figure 3 shows the global AUC scores for the other cues across categories, as well as the AUC scores specific to each category.

The accuracy of the *pragmatic cue* was generally low. For this cue, we only report the results with $k = 1$, which captures relations between basic and superordinate words within a single utterance. Increasing the value of k led to worse, noisier performance. Regarding the *affordance-based cue*, the accuracy was relatively high for some categories, i.e., “food”, “clothes”, “vehicles,” and “toys” and low for others, i.e., “furniture” and “animal.” Finally, the *distributional cue* led to the best overall results across most superordinate categories.

Developmental change?

The results above are based on child-directed speech up to three years old. As we mentioned earlier, this choice was based on research showing that this is the age when early signs of conceptual hierarchy emerge in children’s lexicon. That said, several studies have documented changes in child-directed speech beyond 3 years (e.g., Huttenlocher et al., 2010; Jiang et al., 2022; Kunert et al., 2011), and here we ask whether and how our linguistic cues undergo changes as children grow older and their linguistic input gets richer. While previous work does not make direct predictions about what could change regarding our specific cues, one can expect, e.g., that the distributional cue would become even stronger because the accumulation of more data in an increasing number of contexts would make more accurate the detection of shared patterns of co-occurrence. As for the pragmatic and affordance-based cues, we did not have a priori expectations about whether their role would become more or less important with development.

We followed the same approach as above but with data addressed to children up to 4, 5, and 6 years old, respectively. We stopped at 6 years old as this is an age after which data become very sparse in the dataset CHILDES) (Table 2). This cumulative way to study change (e.g., data up to 4 years old include data up to 3 years old) parallels the cumulative nature of children’s linguistic experience. The results are summarized in Figure 4: They show that the performance of all cues remained stable across development (including for the distributional cue), suggesting, at first glance, that linguistic information in parental talk regarding conceptual hierarchy is already at ceiling by 3 years of age. Nevertheless, this lack of change can also be an artifact of data imbalance across children’s age in CHILDES database: By 6 years of age, the majority of data still comes from utterances directed to children aged 3 years or younger (around 62% of total input, see Table 2). The smaller amount of data in the older ages makes it very hard to find such changes, should they exist. For example, if the patterns of relative contributions are strong up to age 3, even if they do change with age, increasingly small amounts of data would lead to small changes in the results and an (erroneous) conclusion that parents are not changing their behavior.⁴

Cross-cue results

The analyses above explored how the cues fare individually. Here we investigate the extent to which they provide complementary vs. redundant information when combined. To this end, we fit six logistic regressions, one for each superordinate category. In all the regressions, the dependent variable was the binary classification of pairs of basic-level words belonging to same or different superordinate categories. The independent variables were the four similarity measures derived from the four cues. The results of the regressions, summarized in Table 3, indicate that, overall, each cue remains highly significant when

⁴ Another method of investigating developmental change would be to compare cues in each age range separately (rather than cumulatively). We tried this method as well. However, data imbalance across ages also made it difficult to draw meaningful conclusions: Later ages, with much smaller data sizes, were much noisier than earlier ones.

controlling for the other cues. Thus, overall, the cues provided *non-redundant* information to category membership.

Discussion

A crucial question in the study of both language and cognitive development is understanding how children acquire the complex hierarchical relationships that characterize mature human semantic knowledge. A particularly challenging task is learning how basic-level terms are related to abstract superordinate concepts. One difficulty of this task stems from the fact that first-hand perceptual experience does not always provide direct evidence for learning. For example, superordinate category, in particular, often requires additional social input that indicates how the child’s specific linguistic community carves up the world into abstract concepts (e.g., Csibra & Gergely, 2009; Gelman, 2009; Harris, 2012; Waddy, 1982). The current work is a first step to investigating the extent to which *language* that English-learning children (in North America and the UK) hear around them can provide explicit and implicit cues about conceptual hierarchy. The main novelty of our work is that we used a unified computational framework: It allowed us to directly compare different linguistic cues present in child-directed speech in terms of their ability to help learn the taxonomic relations linking basic-level words (present in children’s emerging lexicon) to six common superordinate categories.

The relative usefulness of cues to conceptual hierarchy in child-directed speech

The most explicit cue – where caregivers state the relationship between a basic-level term and its superordinate category label (e.g., “a dog is a kind of animal”) – did not scale up well to the naturalistic dataset we used. This finding contrasts with previous work that found this cue in parental speech (Callanan, 1985; Shipley et al., 1983). This contrast can be explained by the fact that these previous studies were done in the context of rather controlled settings and parents were aware of the task (e.g., teaching words at the superordinate level), whereas here we tested a corpus of largely *spontaneous* speech.

Caregivers can hint at conceptual hierarchy without necessarily stating it explicitly, however (see Table 1). To capture this pragmatic cue, we quantified the co-occurrence between basic-level and superordinate terms within utterances. This simple operationalization was meant to capture all possible ways the hierarchical relationship between two concepts can be expressed linguistically. However, it also made the representation susceptible to errors, mainly by increasing the rate of false alarms: A basic level term (e.g., “juice”) can also co-occur with a superordinate label of which it is not an instance (e.g., “Don’t pour the juice on your clothes!”). The rate of such false alarms was quite high, which explains the overall low – though not random – scores of this cue.

From a developmental point of view, this finding highlights the limitation (at scale) of a simple co-occurrence-based strategy: A deeper understanding of the meaning of the utterance is necessary if children are to learn conceptual hierarchy from pragmatic cues while avoiding false alarms. Future work should aim at providing a more refined computational implementation of these cues. While the ability of state-of-the-art language models in making general, reliable pragmatic inference is still poor, modeling work in Natural Language Processing that has focused on the specific task of inferring taxonomic relations from textual data has been more fruitful (e.g., Chami et al., 2022; Le et al., 2019; Zhang et al., 2022). While this literature has not yet managed to reliably extract taxonomic relations from pairs of terms in a sentence when the relation is not linguistically explicit (i.e., the way the pragmatic cue is defined), the methods developed – when relevant and cognitively plausible – can be a source of inspiration for researchers in child development to improve the study of children’s linguistic input in terms of inferring conceptual structure more generally.

Another linguistic cue we tested is based on affordances (e.g., basic-level terms for food are “eat”-able and can be detected as such if they occur as the grammatical object of the verb “eat”). The accuracy of the affordance-based cue was relatively high for the categories which had an obvious verb to cue its affordance, i.e., “food” (eat), “clothes”

(wear), “vehicles” (ride), and “toys” (play). The accuracy was low in the case of “furniture” since the verb “use” is not exclusive to this category and can also be used with instances of other categories, leading to false alarms. The accuracy for “animal” was also low as it was not characterized by any particular verb.⁵

Finally, the pure distributional cue was the most implicit one since it did not rely on a label or any other linguistic marker for the abstract category. The score for this cue was generally high, including for “animal” and “furniture,” two categories that were not reliably captured with any of the previous cues. This finding suggests that children can learn that certain basic-level terms share common abstract properties by realizing that they have a similar distribution in speech, i.e., that they are used in similar linguistic contexts. This strategy could be even more useful for high-level categories that do not have an explicit label, or for which the label could not be available to the young learners (e.g., “animate” vs. “inanimate”).

In addition to studying each cue individually, we also tested how the cues interacted with one another when they were all combined to predict learning in each superordinate category. We found that the cues were largely non-redundant across most categories, suggesting that if children are able to use a combination of cues, they would have complementary information about category structure. Further, the cues did not fare similarly across categories, suggesting that children can rely more on different cues to learn different categories, e.g., they may use the affordance-based cue more to learn about “food”, “clothes” or “vehicle” and the distributional cue more to learn about “animals” and “furniture.” In other words, early on, children can learn that some words tend to occur as objects to the verbs “eat,” “wear,” or “ride,” enhancing the learning of these conceptual links. In parallel,

⁵ At the same time, the performance of the cue on this category was not totally random because animal instances tend to co-occur more consistently with some verbs from other categories (e.g., "ride a horse", "play with the dog", and "eat the chicken").

and as they accumulate more linguistic experience, they develop increasingly more accurate statistical patterns about which pairs of words share similar distribution in speech, leading to increasingly more refined knowledge about the categories “animals” and “furniture.”

Possible learning mechanisms

For all the cues discussed above, our goal was to test their ability to provide a reliable source of information in child-directed speech. However, our instantiation of the cues abstracted away from the children’s cognitive and information processing limitations, providing only an “ideal observer” point of view (Marr, 1982). That said, several experimental studies have provided evidence for the cognitive plausibility of the learning mechanisms that underlie the use of each of these cues.

For instance, preschool children ably use the “is-a-kind-of” cue to interpret the meaning of a novel word at the superordinate level (Callanan, 1989) (though the current study shows this cue to be highly impoverished in natural input). Both the pragmatic and affordance-based cues rely on the ability to track co-occurrence between pairs of words: the basic- vs superordinate-level labels in the former and the basic-level label vs. the verb cueing affordance in the latter. Extensive research in the last couple of decades has shown that even infants are capable of tracking the co-occurrence of various linguistic units (Saffran et al., 1996). In particular, Bannard and Matthews (2008) have shown toddlers encode together in memory collocational words such as “sit” and “chair” and Wojcik and Saffran (2015) have shown that they encode relationships between novel words co-occurring in a sentence. The distributional cue, on the other hand, requires not only sensitivity to word co-occurrence (as with the pragmatic and affordance-based cues) but also sensitivity to the words’ *shared* patterns of co-occurrence. For example, learners should be sensitive to the fact that “raven” and “salmon” co-occur with similar words and phrases such as “lay egg”, “live,” and “reproduce”, although the pair of words “raven” and “salmon” may not themselves co-occur with each other. There is evidence that even infants are sensitive to shared patterns of

co-occurrence (e.g., Lany & Saffran, 2011).

Note, however, that while the detection of reliably co-occurring pair of labels can create a mental association (via simple associative mechanisms), it is not clear how this association would develop into becoming a taxonomic relation (as opposed to, say, a thematic relation). In other words, detecting co-occurrence is a necessary but not a sufficient condition for the learning of conceptual *hierarchy*. In the case of the affordance-based cue, the co-occurrence information would need to be supplemented with an understanding of the syntactic relations (i.e., verb-complement structure). In the case of the pragmatic cue, it would need to be supplemented by inferences about the relationship that the speaker intends to convey between the basic and superordinate terms given the context. Indeed, the fact that our simple implementation of the pragmatic cue based on mere co-occurrence has led to poor performance is a strong testimony to the need for this additional, higher-level inference.

In contrast, for the pure distributional cue on its own, it is not clear that there is any additional information (syntactic or extra-linguistic) that would help children induce the conceptual hierarchy beyond information about word co-occurrence. On this account, the accumulation of direct co-occurrence in children’s memory (e.g., “fruits”-“hungry” and “vegetables”-“hungry”) leads to the realization (perhaps triggered by cognitive maturation, e.g., Bauer & San Souci, 2010; Savic, et al., 2023) that some words share similar pattern of co-occurrence (e.g., that “fruits” and “vegetables” co-occur in similar linguistic contexts; in this example, they both co-occur with “hungry”), which then foster the creation of a taxonomic – rather than a thematic – link (i.e., “fruits” and “vegetables” are instances of a single higher-level conceptual category, e.g., “food”) (Brown & Berko, 1960; Ervin, 1961; McNeill, 1963; Sloutsky et al., 2017). It is still not entirely clear what precise mechanistic process can explain how sensitivity to shared patterns of co-occurrence between two labels may give rise to a taxonomic link relating these labels to a single superordinate category, supporting conceptual inference and knowledge generalization. We refer the reader to Unger

and Fisher (2021) and Savic et al. (2023) for a review and discussion of some hypothesized mechanisms based on behavioral, computational, and/or neuroscientific evidence.

Conclusions

This work studied cues in the input that might support children’s acquisition of conceptual hierarchies. While most previous work focused on the role of first-hand observation, here we investigated another (complementary) source of input, i.e., the language children hear around them. We found implicit cues (the affordance-based and distributional cues) to be much more reliable than explicit cues (“is-a-kind-of” and pragmatic cues). While studies have shown that children are cognitively equipped to learn conceptual hierarchy from both implicit and explicit cues, our input analysis suggests that implicit cues could be the more important source of this knowledge.

That said, the current work is only a first step and more work is needed to refine the cues, especially the ones for which we only provided a simple approximation such as the pragmatic cue. In addition, while we emphasized the complementary aspect of linguistic cues compared to perceptual cues, future work should allow direct comparison between these two accounts in order to better estimate the relative role of each source of information in development (e.g., Andrews et al., 2014; Bruni et al., 2014; Fourtassi & Frank, 2020; Fourtassi et al., 2014). Ultimately, the goal would be to integrate both sources into a *cognitive* model that uses these cues in a principled way. Such a model should make precise developmental predictions about how the cues interact and how their role in learning changes across development.

Another direction for future work is to compare the findings we obtained with speech directed to English-learning children in North America and the UK to other languages. Understanding variation in children’s experiences allows us to determine which aspects of development are universal and which are culture-specific (Rogoff et al., 2018). The learning of conceptual hierarchy is an excellent case study; we know for example that superordinate

concepts have both similarities and differences across cultures (e.g. Waddy, 1982). In addition, the cues may vary in terms of their reliability from one language to the other, depending both on cultural practices/parenting styles and on the specificities of the language. Finally, it is crucial for future work to investigate whether and how variation in caregivers' language induces variations in children's conceptualization of the world.

References

- Andrews, M., Frank, S., & Vigliocco, G. (2014). Reconciling embodied and distributional accounts of meaning in language. *Topics in Cognitive Science*, *6*(3), 359–370. <https://doi.org/10.1111/tops.12096>
- Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*(3), 463. <https://doi.org/10.1037/a0016261>
- Bannard, C., & Matthews, D. (2008). Stored word sequences in language learning: The effect of familiarity on children's repetition of four-word combinations. *Psychological Science*, *19*(3), 241–248. <https://doi.org/10.1111/j.1467-9280.2008.02075.x>
- Bauer, P. J., & San Souci, P. (2010). Going beyond the facts: Young children extend knowledge by integrating episodes. *Journal of Experimental Child Psychology*, *107*(4), 452–465. <https://doi.org/10.1016/j.jecp.2010.05.012>
- Blaye, A., Bernard-Peyron, V., Paour, J.-L., & Bonthoux, F. (2006). Categorical flexibility in children: Distinguishing response flexibility from conceptual flexibility; the protracted development of taxonomic representations. *European Journal of Developmental Psychology*, *3*(2), 163–188. <https://doi.org/10.1080/17405620500412267>
- Brown, R., & Berko, J. (1960). Word association and the acquisition of grammar. *Child Development*, *31*(1), 1–14. <https://doi.org/10.2307/1126377>

- Bruni, E., Tran, N.-K., & Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, *49*, 1–47. <https://doi.org/10.1613/jair.4135>
- Callanan, M. A. (1985). How parents label objects for young children: The role of input in the acquisition of category hierarchies. *Child Development*, *56*(2), 508–523. <https://doi.org/10.2307/1129738>
- Callanan, M. A. (1989). Development of object categories and inclusion relations: Preschoolers' hypotheses about word meanings. *Developmental Psychology*, *25*(2). <https://doi.org/10.1037/0012-1649.25.2.207>
- Carey, S. (1987). *Conceptual change in childhood*. MIT Press.
- Chami, I., Abu-El-Haija, S., Perozzi, B., Ré, C., & Murphy, K. (2022). Machine learning on graphs: A model and comprehensive taxonomy. *Journal of Machine Learning Research*, *23*(89), 1–64. <https://doi.org/10.48550/arXiv.2005.03675>
- Chi, M. T., Hutchinson, J. E., & Robin, A. F. (1989). How inferences about novel domain-related concepts can be constrained by structured knowledge. *Merrill-Palmer Quarterly*, *35*(1), 27–62. <http://www.jstor.org/stable/23086424>
- Clark, E. V. (1997). Conceptual perspective and lexical choice in acquisition. *Cognition*, *64*(1), 1–37. [https://doi.org/10.1016/s0010-0277\(97\)00010-3](https://doi.org/10.1016/s0010-0277(97)00010-3)
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *13*(4), 148–153. <https://doi.org/10.1016/j.tics.2009.01.005>
- Ervin, S. M. (1961). Changes with age in the verbal determinants of word-association. *The American Journal of Psychology*, *74*(3), 361–372. <https://doi.org/10.2307/1419742>
- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., Tomasello, M., Mervis, C. B., & Stiles, J. (1994). Variability in early communicative development.

Monographs of the Society for Research in Child Development, 59(5).

<https://doi.org/10.1111/j.1540-5834.1994.tb00169.x>

Fourtassi, A. (2020). Word co-occurrence in child-directed speech predicts children's free word associations. In *Proceedings of the workshop on cognitive modeling and computational linguistics* (pp. 49–53). <https://doi.org/10.18653/v1/2020.cmcl-1.6>

Fourtassi, A., Dunbar, E., & Dupoux, E. (2014). Self-consistency as an inductive bias in early language acquisition. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 36).

Fourtassi, A., & Dupoux, E. (2013). A corpus-based evaluation method for distributional semantic models. In *51st annual meeting of the association for computational linguistics proceedings of the student research workshop* (pp. 165–171).

Fourtassi, A., & Frank, M. C. (2020). How optimal is word recognition under multimodal uncertainty? *Cognition*, 199, 104092. <https://doi.org/10.1016/j.cognition.2019.104092>

Fourtassi, A., Schatz, T., Varadarajan, B., & Dupoux, E. (2014). Exploring the relative role of bottom-up and top-down information in phoneme learning. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* (pp. 1-6). <https://doi.org/10.3115/v1/P14-2001>

Fourtassi, A., Scheinfeld, I., & Frank, M. C. (2019). The development of abstract concepts in children's early lexical networks. In *Proceedings of the workshop on cognitive modeling and computational linguistics* (pp. 129–133). <https://doi.org/10.18653/v1/W19-2914>

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44(3), 677–694. <https://doi.org/10.1017/S0305000916000209>

- Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2021). *Variability and consistency in early language learning: The wordbank project*. MIT Press.
- Frermann, L., & Lapata, M. (2016). Incremental bayesian category learning from natural language. *Cognitive Science*, *40*(6), 1333–1381. <https://doi.org/10.1111/cogs.12304>
- Gelman, S. A. (2009). Learning from others: Children’s construction of concepts. *Annual Review of Psychology*, *60*, 115–140.
<https://doi.org/10.1146/annurev.psych.59.103006.093659>
- Gentner, D., & Namy, L. L. (1999). Comparison in the development of categories. *Cognitive Development*, *14*(4), 487–513. [https://doi.org/10.1016/S0885-2014\(99\)00016-7](https://doi.org/10.1016/S0885-2014(99)00016-7)
- Harris, P. L. (2012). *Trusting what you’re told*. Harvard University Press.
- Harris, Z. S. (1957). Co-occurrence and transformation in linguistic structure. *Language*, *33*(3), 283–340. <https://doi.org/10.2307/411155>
- Hills, T. T., Maouene, J., Riordan, B., & Smith, L. B. (2010). The associative structure of language: Contextual diversity in early word learning. *Journal of Memory and Language*, *63*(3), 259–273. <https://doi.org/10.1016/j.jml.2010.06.002>
- Huebner, P. A., & Willits, J. A. (2018). Structured semantic knowledge can emerge automatically from predicting word sequences in child-directed speech. *Frontiers in Psychology*, *9*, 133. <https://doi.org/10.3389/fpsyg.2018.00133>
- Huttenlocher, J., Waterfall, H., Vasilyeva, M., Vevea, J., & Hedges, L. V. (2010). Sources of variability in children’s language growth. *Cognitive Psychology*, *61*(4), 343–365.
<https://doi.org/10.1016/j.cogpsych.2010.08.002>
- Inagaki, K., & Hatano, Giyoo. (2002). *Young children’s naive thinking about the biological world*. Psychology Press.

- Inhelder, B., & Piaget, J. (2013). *The early growth of logic in the child: Classification and seriation*. Routledge.
- Jiang, H., Frank, M. C., Kulkarni, V., & Fourtassi, A. (2022). Exploring patterns of stability and change in caregivers' word usage across early childhood. *Cognitive Science*, 46(7), e13177. <https://doi.org/10.1111/cogs.13177>
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical bayesian models. *Developmental Science*, 10(3), 307–321. <https://doi.org/10.1111/j.1467-7687.2007.00585.x>
- Kunert, R., Fernández, R., & Zuidema, W. (2011). Adaptation in child directed speech: Evidence from corpora. In *Proceedings of SemDial* (pp. 112–119).
- Landauer, T. K., & Dumais, S. T. (1997). A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211. <https://doi.org/10.1037/0033-295X.104.2.211>
- Lany, J., & Saffran, J. R. (2011). Interactions between statistical and semantic information in infant language development. *Developmental Science*, 14(5), 1207–1219. <https://doi.org/10.1111/j.1467-7687.2011.01073.x>
- Le, M., Roller, S., Papaxanthos, L., Kiela, D., & Nickel, M. (2019). Inferring concept hierarchies from text corpora via hyperbolic embeddings. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 3231–3241). <https://doi.org/10.18653/v1/P19-1313>
- Liu, J., Golinkoff, R. M., & Sak, K. (2001). One cow does not an animal make: Young children can extend novel words at the superordinate level. *Child Development*, 72(6), 1674–1694. <https://doi.org/10.1111/1467-8624.00372>

Lucariello, J., Kyratzis, A., & Nelson, K. (1992). Taxonomic knowledge: What kind and when? *Child Development, 63*(4), 978–998. <https://doi.org/10.2307/1131248>

MacWhinney, B. (2000). *The childe's project: Tools for analyzing talk. Transcription format and programs* (Vol. 1). Psychology Press.

Madole, K. L., & Oakes, L. M. (1999). Making sense of infant categorization: Stable processes and changing representations. *Developmental Review, 19*(2), 263–296. <https://doi.org/10.1006/drev.1998.0481>

Markman, E. M. (1989). *Categorization and naming in children: Problems of induction*. MIT Press.

Markman, E. M., & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology, 16*(1), 1–27. [https://doi.org/10.1016/0010-0285\(84\)90002-1](https://doi.org/10.1016/0010-0285(84)90002-1)

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. MIT Press.

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience, 4*(4), 310–322. <https://doi.org/10.1038/nrn1076>

McNeill, D. (1963). The origin of associations within the same grammatical class. *Journal of Verbal Learning and Verbal Behavior, 2*(3), 250–262. [https://doi.org/10.1016/S0022-5371\(63\)80091-2](https://doi.org/10.1016/S0022-5371(63)80091-2)

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems* (pp. 3111–3119).

- Murphy, G. (2004). *The big book of concepts*. MIT press.
- Quinn, P. C., & Eimas, P. D. (2000). The emergence of category representations during infancy: Are separate perceptual and conceptual processes required? *Journal of Cognition and Development, 1*(1), 55–61.
https://doi.org/10.1207/S15327647JCD0101N_6
- Rogoff, B., Dahl, A., & Callanan, M. (2018). The importance of understanding children's lived experience. *Developmental Review, 50*, 5–15.
<https://doi.org/10.1016/j.dr.2018.05.006>
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*(5294), 1926–1928.
<https://doi.org/10.1126/science.274.5294.1926>
- Sanchez, A., Meylan, S. C., Braginsky, M., MacDonald, K. E., Yurovsky, D., & Frank, M. C. (2019). Childes-db: A flexible and reproducible interface to the child language data exchange system. *Behavior Research Methods, 51*(4), 1928–1941.
<https://doi.org/10.3758/s13428-018-1176-7>
- Savic, O., Unger, L., & Sloutsky, V. M. (2023). Experience and maturation: The contribution of co-occurrence regularities in language to the development of semantic organization. *Child Development, 94*(1), 142–158.
<https://doi.org/10.1111/cdev.13844>
- Saxe, A. M., McClelland, J. L., & Ganguli, S. (2019). A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences, 116*(23), 11537–11546. <https://doi.org/10.1073/pnas.1820226116>
- Shipley, E. F., Kuhn, I. F., & Madden, E. C. (1983). Mothers' use of superordinate category terms. *Journal of Child Language, 10*(3).

<https://doi.org/10.1017/S0305000900005377>

Sloutsky, V. (2015). Conceptual development. In L. S. Liben, U. Müller, & R. M. Lerner (Eds.), *Handbook of child psychology and developmental science: Cognitive processes* (pp. 469–518). John Wiley & Sons, Inc.

<https://doi.org/10.1002/9781118963418.childpsy212>

Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science*, 34(7), 1244–1286. <https://doi.org/10.1111/j.1551-6709.2010.01129.x>

Sloutsky, V. M., Yim, H., Yao, X., & Dennis, S. (2017). An associative account of the development of word learning. *Cognitive Psychology*, 97, 1–30.

<https://doi.org/10.1016/j.cogpsych.2017.06.001>

Smith, L. B., & Heise, D. (1992). Perceptual similarity and conceptual structure. In *Advances in psychology*, 93, 233–272). [https://doi.org/10.1016/S0166-4115\(08\)61009-2](https://doi.org/10.1016/S0166-4115(08)61009-2)

Stella, M., Beckage, N. M., & Brede, M. (2017). Multiplex lexical networks reveal patterns in early word acquisition in children. *Scientific Reports*, 7.

<https://doi.org/10.1038/srep46730>

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10(7), 309–318.

<https://doi.org/10.1016/j.tics.2006.05.009>

Unger, L., & Fisher, A. V. (2021). The emergence of richly organized semantic knowledge from simple statistics: A synthetic review. *Developmental Review*, 60, 100949.

<https://doi.org/10.1016/j.dr.2021.100949>

Unger, L., Savic, O., & Sloutsky, V. M. (2020). Statistical regularities shape semantic organization throughout development. *Cognition*, 198, 104190.

<https://doi.org/10.1016/j.cognition.2020.104190>

Waddy, J. A. (1982). Biological classification from a groote eylandt aborigine's point of view.

Journal of Ethnobiology, 2(1), 63–77.

Winer, G. A. (1980). Class-inclusion reasoning in children: A review of the empirical

literature. *Child Development*, 51(2), 309–328. <https://doi.org/10.2307/1129264>

Wojcik, E. H., & Saffran, J. R. (2015). Toddlers encode similarities among novel words from meaningful sentences. *Cognition*, 138, 10–20.

<https://doi.org/10.1016/j.cognition.2015.01.015>

Xu, F., & Tenenbaum, J. B. (2007). Word learning as bayesian inference. *Psychological*

Review, 114(2), 245. <https://doi.org/10.1037/0033-295X.114.2.245>

Zhang, H., Liu, X., Pan, H., Ke, H., Ou, J., Fang, T., & Song, Y. (2022). Aser: Towards large-scale commonsense knowledge acquisition via higher-order selectional preference over eventualities. *Artificial Intelligence*, 309, 103740.

<https://doi.org/10.48550/arXiv.2104.02137>

Table 1

Examples of utterances from CHILDES where parents hint at hierarchical relations between basic- and superordinate-level terms.

Category	Utterance	Interlocutors	Corpus
Animals	Do you want a cow or do you want a different animal?	Mother to Max, 30 months	EllisWeismer
Furniture	Furniture means sofas and chairs and...	Mother to Naima, 23 months	Providence
Clothes	This is another clothes. See it's just like this shirt.	Investigator to Shem, 30 months	Clark
Food	She asked Lily what her favorite food was. If Lily says chocolate I am in trouble.	Mother about Lily, 24 months	Providence
Toys	You close the book and we'll get a different toy cause I think you're tired of this.	Mother to child, 13 months	Ambrose
Vehicles	The only vehicle you cut out so far is the train.	Mother to Warren, 30 months	Manchester

Table 2

Information about the corpora used in the analysis of developmental change.

Age	Mean (months)	SD	Median	Children	Transcripts	Utterances	Words
<3	26.5	5.9	27.5	660	4,939	1.996 M	8.262 M
<4	30.1	8.2	30	843	6,750	2.657 M	11.270 M
<5	33.1	11	31.3	971	7,889	3.093 M	13.393 M
<6	35	13	32.1	1,046	8,654	3.221 M	13.991 M

Table 3

Logistic regressions predicting the binary classification of pairs of basic-level words as belonging to same or different superordinate categories. The predictors are the pairs' similarity measures derived from each cue. We fit a different regression for each superordinate category. The predictors were centered and scaled for comparability.

	(Intercept)	Distributional	Affordance	Pragmatic
Animals	−2.741*** (0.085)	2.285*** (0.074)	0.022 (0.057)	0.179*** (0.050)
Furniture	−3.195*** (0.138)	2.040*** (0.127)	0.547*** (0.094)	−0.104 (0.080)
Toys	−3.244*** (0.155)	1.178*** (0.136)	0.620*** (0.113)	0.722*** (0.120)
Food	−2.616*** (0.112)	0.905*** (0.060)	2.112*** (0.092)	0.325*** (0.059)
Clothing	−3.101*** (0.183)	1.644*** (0.171)	1.535*** (0.153)	0.359* (0.146)
Vehicles	−4.663*** (0.348)	1.249*** (0.193)	2.211*** (0.245)	0.159 (0.138)

Note:

*p<0.05; **p<0.01; ***p<0.001

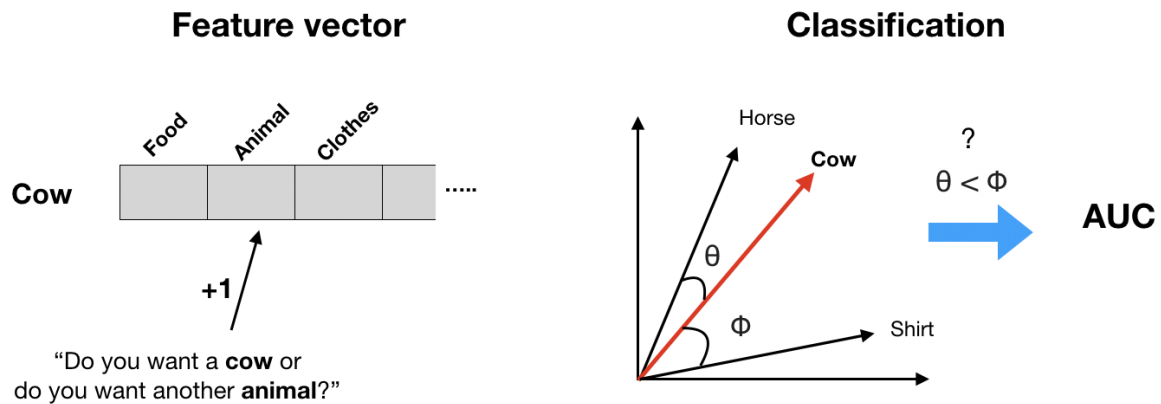


Figure 2

A schematic description of the task. For each basic-level word (here, ‘cow’) a feature vector is derived from child-directed speech based on how the cue is defined. Here, the vector cells correspond to the superordinate categories. The entry in a given cell (e.g., animal) is incremented when the word ‘cow’ co-occurs with the corresponding category label. The cue is evaluated based on its ability to classify pairs of words into ‘same’ or ‘different’ superordinate categories. Here, the pair ‘cow’-‘horse’ belongs to the same category. The corresponding vectors should be closer to each other than the vectors of a pair that belongs to different categories (e.g., ‘cow’-‘shirt’). This evaluation is quantified by a standard measure in signal detection theory called the Area Under the ROC Curve (AUC).

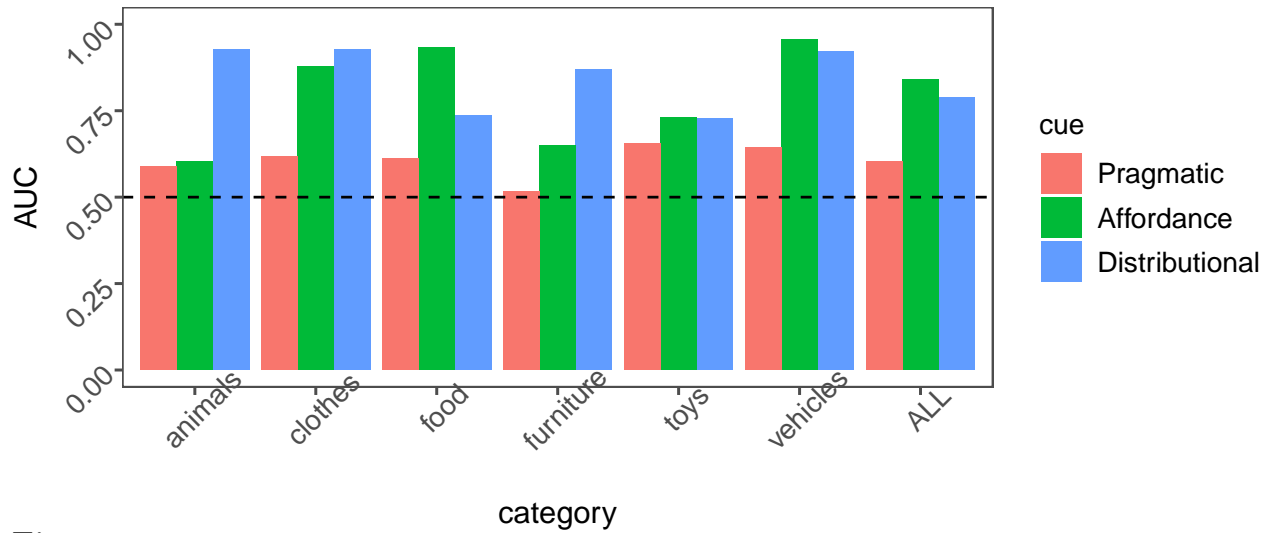


Figure 3

The Area Under the ROC Curve (AUC) scores of the cues for each category and across all categories ('ALL'). A value of 0.5 represents pure chance, and a value of 1 represents perfect performance.

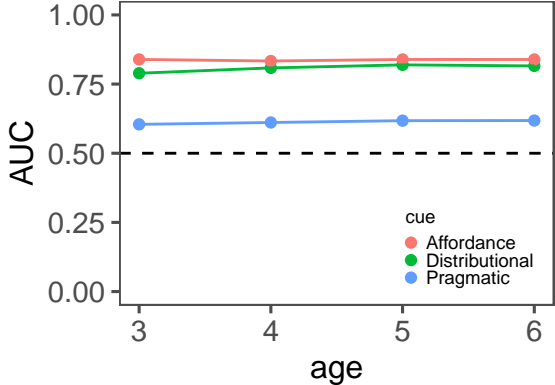


Figure 4

The Area Under the ROC Curve (AUC) scores for each cue (across all categories) using speech heard by children up to a particular age. A value of 0.5 represents pure chance, and a value of 1 represents perfect performance.