Cognitive Psychology 166 (2026) 101829

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Children use speaker knowledge to learn novel subordinate vs. basic-level nouns

June Choe *, Anna Papafragou

Department of Linguistics, University of Pennsylvania, 3401-C Walnut St., Suite 300C, Philadelphia, PA 19104, United States

ARTICLE INFO

Keywords:
Subordinates
Contrast
Pragmatics
Speaker knowledge
Informativeness
Word learning

ABSTRACT

Children typically prefer the so-called "basic"-level (e.g., dog) versus narrower subordinate (e.g., dalmatian) meanings in acquisition. Subordinate conjectures are facilitated by contrast; for example, when shown a dalmatian and a corgi each labelled "dax" and "wug," learners take "dax" to mean dalmatian, as opposed to dog. But what is the nature of this effect? On a pragmatic account, contrast is helpful because it highlights a communicative pressure on the speaker to be more informative than usual when describing the target referent, given relevant alternatives. Another possibility, however, is that contrastive forms create a simple lexical heuristic. We raise and test two predictions that follow from the pragmatic, but not the lexical, account: in learning subordinate nouns, children use contrastive information even when there are no competing lexical labels (Exp. 1); furthermore, they consider the speaker's knowledge of contrasting alternatives during word learning, as a fully Gricean rich-computation model expects (Exp. 2). Results support a pragmatic (over a purely lexical) account of subordinate conjectures. Thus subordinate noun learning offers an early demonstration of children's ability to integrate informativeness with the speaker's epistemic state during word learning.

1. Introduction

1.1. The challenge of subordinate noun learning

Word learning involves not only the task of mapping labels onto referents in the physical world, but also the task of identifying the intended meaning that a speaker encodes in their choice of a word (Grigoroglou & Papafragou, 2021). This challenge of discovering meaning is pervasive and permeates even the acquisition of concrete nouns for familiar object categories, especially for narrow and specific word meanings such as *dalmatian*, as opposed to *dog* (Callanan & Markman, 1982; Choe & Papafragou, 2023; Clark, 1987; Clark, 1997; Gelman et al., 1989; Markman, 1989; Quine, 1960; Waxman, 1990). In the literature, this distinction is characterized as a difference in specificity between subordinate-level and basic-level categories, respectively. Furthermore, the so-called "basic-level" categories are generally thought to be privileged in early word learning (e.g., Rosch et al., 1976; Cruse, 1977; Mervis & Crisafi, 1982; Blewitt, 1994; Callanan et al., 1994; Murphy, 2002, 2016; Choe & Papafragou, 2023). By comparison, subordinate nouns are more difficult to acquire under word learning conditions characterized by simple ostensive labelling (on the separate but similar problem posed by superordinate nouns such as *animal*, see the review in Choe & Papafragou, 2025).

Young children's difficulties with acquiring subordinate nouns have been highlighted by many prior word learning experiments.

* Corresponding author.
E-mail addresses: yjchoe@sas.upenn.edu (J. Choe), anna4@sas.upenn.edu (A. Papafragou).

https://doi.org/10.1016/j.cogpsych.2026.101829
Received 10 September 2025; Received in revised form 4 June 2026; Accepted 16 July 2026
Available online 23 July 2026

Much of this work uses versions of a paradigm in which children are introduced to a target exemplar (e.g., a dalmatian) with a label and then asked to extend the label to a generalization set containing subordinate matches (other dalmatians), basic-level matches (different dog breeds), and distractors (e.g., other animals or non-living things; see, e.g., Waxman, 1990 for a classic design). In several studies, when a novel noun was used to refer to an object, preschool-aged children performed worse at generalizing the word's meaning to the subordinate level than when no label was presented at all (e.g., Waxman, 1990; Waxman et al., 1991). This contrasts with the finding that nouns facilitate taxonomic categorization at the basic and even the superordinate level (e.g., Golinkoff et al., 1995; Waxman, 1990; Waxman & Gelman, 1986). A similar pattern has been reported in production: children overwhelmingly use basic-level nouns to name objects (e.g., responding "dog" to each and every dog in a diverse set of many dogs pointed at individually) and seemingly generate labels at the subordinate level only when explicitly prompted (e.g., "But if these are all dogs, how can we tell the difference?"; Waxman, 1990).

How then are subordinate meanings acquired? According to an influential proposal, learners extract these narrower categories from the physical world by attending to the perceptual characteristics of labelling events. Under such mechanistic accounts, subordinate categories primarily emerge from the physical circumstances of observation, influenced by the perceptually salient details of exemplars and the learner's ability to hold onto these perceptual representations in memory as they are stored and refined across labelling events. Whereas objects are typically construed at the basic level by their most salient characteristics (e.g., for a dog, having a snout, a lean body, and four legs), subordinate-level categories are defined by more fine-grained features that are elusive to the typical circumstances of observation (e.g., for a dalmatian, having short, white fur with black spots). In turn, the availability of these subordinate-level features is heavily influenced by the manner of exemplar presentation, such as whether the learner sees three exemplars of a dalmatian simultaneously, side-by-side, versus seeing the same exemplars sequentially, separated across space and time (Jenkins et al., 2015; Jenkins et al., 2021; Spencer et al., 2011). Here, the simultaneous presentation is thought to activate subordinate categories due to the immediately available perceptual contrasts between similar exemplars, which improves the discovery and retention of the more specific features that are shared (e.g., Samuelson et al., 2009; Sloutsky et al., 2007; Sloutsky & Fisher, 2004; Spencer et al., 2011). On this view, it is the physical world that offers objects at different levels of categorization by making salient the more specific or global kinds of features; the learner in turn perceives the dalmatian-ness or the dog-ness of the labelled exemplars, and through that process discovers the word's meaning in a largely bottom-up fashion.

On a prominent alternative position, pointers to subordinate meanings are primarily found in language. As this literature points out, mapping words to subordinate-level categories is greatly helped by cues to a narrower word meaning, including contrastive information that guides children's word meaning conjectures away from the often-preferred basic level (see especially Waxman, 1990; Waxman et al., 1991; Waxman et al., 1997). For example, the penalty for subordinate-level categorization is reversed when the novel noun is anchored to a known basic-level term (e.g., "This is a *dax*. It's a kind of a dog"), a linguistic device that children are sensitive to in both comprehension and production (Callanan, 1985, 1989; Markman et al., 1980; Waxman et al., 1991). Similarly, when taught a novel compound noun headed by a basic-level term (e.g., "dalmatian-dog"), children restrict their generalizations to subordinate-level matches rather than extending them to the entire basic-level category (Clark et al., 1985; Gelman et al., 1989). In a particularly clear demonstration, three-to-four-year-olds were asked questions about an object (e.g., a dog) that contained misleading labels at either the subordinate (e.g., "Is this a dandelion?") or the basic level (e.g., "Is this a tree?"; Waxman & Hatch, 1992): when children responded by producing the correct label, they matched the linguistic level of their response to that of the question, showing that they can flexibly produce subordinate labels when the linguistic (contrastive) context supports them.

Subordinate noun learning is additionally helped by the explicit mention of contrastive information about distinct subordinate-level categories with different names, such as the fact that *collies* help shepherds take care of sheep while *settlers* help hunters find birds (Waxman et al., 1991, 1997; see also Klibanoff & Waxman, 2000; Waxman & Klibanoff, 2000; Waxman et al., 2009, for somewhat different examples beyond the noun domain). In a more recent study using cross situational learning (Wang & Trueswell, 2019; cf. Wang & Trueswell, 2022), three-to-five-year-old children were more likely to generalize the meaning of a novel label at the subordinate level (e.g., taking "dax" to mean *dalmatian* after seeing three dalmatians) if a second novel label was simultaneously paired with other subordinate-level exemplars from the same basic-level category (e.g., three non-dalmatian dogs). In the same study, children generalized the meaning of a novel label at the basic level (e.g., taking "dax" to mean *dog* after seeing three dalmatians) when a second label was instead paired with exemplars from a different basic-level category (e.g., three birds). Finally, in a study with adult learners, contrast was interpreted as a cue to a subordinate meaning even when the input was presented in an artificial language, and effects of contrast extended beyond the precise moment of labelling, driven by inferences about the relevance of lexical alternatives presented at generalization; moreover, these effects disappeared when a visual alternative was merely present without being labelled (Choe & Papafragou, 2023).

To summarize, it has been extensively demonstrated that contrast helps constrain the space of possible word meanings, and that this is especially helpful for the task of discovering subordinate (versus basic-level) noun meanings. A remaining question is how exactly contrast drives the narrowing of a novel word's hypothesized meaning. Under mechanistic accounts, what matters is the immediate comparisons available in the referential world: subordinate meanings arise from circumstances of observation that highlight fine-grained featural distinctions. On this view, contrast is not necessarily mediated by language, though labels may offer additional cues to attention and encoding (Sloutsky, 2010; but see Wang & Trueswell, 2019; Choe & Papafragou, 2023). Alternatively, contrast may be a linguistic-communicative phenomenon, operating in at least two different ways. On one possibility, contrast operates as a surface-level heuristic: when the child encounters or already knows an alternative word, the availability of that competing label guides them toward a different (e.g., narrower) meaning for the target (e.g., Golinkoff et al., 1992; Golinkoff et al., 1994; Markman, 1987, 1989). This kind of lexical heuristic may be sufficient to explain contrast effects on subordinate meanings, since much of the prior literature involved children learning from one-to-one mappings between a single label and a referent. On a second

possibility, contrast reveals a pragmatic choice on the part of the speaker (Clark, 1987, 1990, 2017; Clark & Wong, 2002; Diesendruck, 2005). On this view, learners do not only track differences in linguistic forms but also consider why the speaker chose one form over another, drawing on inferences about the speaker's mental state and communicative goals (Au & Glusman, 1990; Diesendruck et al., 2004; Diesendruck et al., 2010). Currently, the rich prior demonstrations of the role of contrast for moving learners beyond the basic level have not conclusively settled on a single one of these theoretical options over the others.

1.2. Subordinate meanings, contrast, and pragmatics

In the present study, we build on the well-established finding that contrastive information facilitates subordinate-level hypotheses for word meanings (see previous section). Our goal is to understand the mechanism by which contrastive cues exert their effects, and embed this mechanism into a broader theory of how children acquire word meanings. Specifically, we present a pragmatic account of children's subordinate noun learning (including a pragmatic notion of contrast) and test its explanatory potential against alternative possibilities.

We pursue the hypothesis that subordinate meanings arise from the expectation that the speaker intends to be informative in their choice of a word, given the availability of relevant alternatives (see also Choe & Papafragou, 2023). On our pragmatic account, what constitutes informativeness depends on contextual factors including what information is shared between the speaker and listener, and what alternatives are salient in the discourse context (Grice, 1975; cf. Sperber & Wilson, 1986). For nouns denoting taxonomic categories, the subordinate-basic-superordinate distinction can be understood as forming a scale on the basis of informativeness (Horn, 1972), whereby the subordinate level is typically more informative than the basic level. In turn, the acquisition of subordinate nouns is helped by clear cues to the speaker's intent. This problem of lexical choice from a scale can be understood as audience design, involving the assessment of what level of specificity would be the most conversationally appropriate (Brennan & Clark, 1996). This framing also helps capture one interpretation of the basic-level bias, as the level of specificity that is informative enough to satisfy the needs of a "generic addressee" (Brown & Dell, 1987; Choe & Papafragou, 2023; Cruse, 1977; Degen et al., 2020; Grigoroglou & Papafragou, 2019; Lockridge & Brennan, 2002).

On this pragmatic proposal, contrast is helpful for learning subordinate nouns because it allows learners to interpret the speaker's choice of a word as an intentional act of narrowly identifying the target among alternatives which differ in a particular dimension. Thus when a speaker provides two distinct labels for a pair of contrasting exemplars at the subordinate level (e.g., "dalmatian" and "corgi"), learners infer that the speaker's intent is more specific in using those words, given the availability of the basic-level lexical alternative ("dog"; Choe & Papafragou, 2023). Note that contrast need not involve two competing linguistic forms; contrastive inferences may also be facilitated via communicative stimuli beyond words, such as drawings or pictures (Kampa et al., 2024; Kampa & Papafragou, 2023). For example, Kampa et al. (2024) showed that children as young as three years old can derive quantity-based inferences from pictures as well as words, selecting pragmatically appropriate referents based on information that was omitted from a pictorial or linguistic description. This suggests that the inferential machinery underlying contrastive interpretations is not restricted to lexical alternatives but is governed by general principles of communication that also apply to non-linguistic forms.

Crucially, on a pragmatic account, learners' hypotheses about word meaning are not only shaped by what children themselves know, but also by what they can assume about the speaker's knowledge and perspective. This view is informed by demonstrations of effects of the interlocutor's mental state in how informative children are during reference production and comprehension. For example, Nadig and Sedivy (2002) found that 4-to-5-year-olds modulate the specificity of their reference choices by tracking information in the common ground, suppressing their production of modifiers to refer to a target (e.g., referring to a taller cup as "a cup" versus "a tall cup") when the relevant visual alternative (e.g., a shorter cup) was not visible to their interlocutor. Complementing this evidence, Kampa and Papafragou (2020) showed that children can modulate their informativeness judgments depending on the speaker's knowledge state, successfully matching more-informative descriptions (e.g., "I see a box and a spoon") to contexts where the speaker had full visual access to the contents of a box being described, and less-informative statements (e.g., "I see a spoon") to contexts where the speaker only had partial knowledge. These findings suggest that children understand pragmatic constraints on informativeness involving inferences about relevant alternatives and the speaker's perspective.

The pragmatic account makes two predictions about the role of contrast in the acquisition of subordinate meanings. First, it predicts that children should use contrastive information to facilitate subordinate generalizations even when there are no competing lexical forms within the learning episode. Second, it predicts that children, in generating specific hypotheses for subordinate word meanings, should integrate information about speaker knowledge, such as whether contrastive lexical alternatives for the target referent are also available to the speaker. Together, these two predictions differ from alternative accounts of the role of contrast in subordinate noun learning reviewed earlier. Mechanistic accounts predict that contrast effects emerge primarily from the learner's observation of perceptual contrasts in the world; on these accounts, both labelling contrasts and the speaker's mental access to contrastive alternatives should be irrelevant for learning purposes. Lexical heuristic accounts predict that contrast should exert its effects only when two or more contrasting linguistic forms are explicitly present; these accounts cannot accommodate more flexible implementations of contrast (including cases where only one of the contrastive alternatives is fully labelled) and – because they conceive of contrast as a simple heuristic - are silent on the role of the speaker's mental access to contrastive alternatives.

In what follows, we test these predictions of the pragmatic account in two word learning experiments with 4- and 5-year-old children who have been extensively studied in the literature on the acquisition and use of subordinate nouns (e.g., Emberson et al., 2019; Jenkins et al., 2015; Wang & Trueswell, 2019; Waxman et al., 1997; Xu & Tenenbaum, 2007a) and the development of informativeness and perspective-taking (e.g., Nadig & Sedivy, 2002; Barner et al., 2011; Katsos & Bishop, 2011; Skordos & Papafragou, 2016; Papafragou et al., 2017; Papafragou et al., 2018; Kampa & Papafragou, 2020; Wilson et al., 2022; Choe & Papafragou, 2025). In

Experiment 1, we ask whether children's subordinate-level hypotheses for word meaning are driven by the presence of alternative exemplars that contrast with the target referent at the subordinate level. We predict that children should generalize the meaning of novel labels to the subordinate level when other subordinate-level exemplars are present (even if only one of these is labelled with a unique label) compared to when only the target referent is present and labelled. In Experiment 2, we ask whether the effect of contrast on subordinate word meanings is modulated by the speaker's knowledge of the potential alternatives to the target. We predict that children should show more subordinate-level generalizations when the speaker has full visual access to the alternatives compared to when those alternatives are hidden from the speaker's view.

2. Experiment 1

In Experiment 1, we tested whether inferences about the specificity of word meaning are driven by the presence of alternatives that contrast along a particular taxonomic dimension. As we have discussed earlier, our hypothesis for the effect of contrast is that it is guided by the expectation that the speaker intends to be more specific than usual in their choice of label. Thus, we predicted that – even within a single learning opportunity - learners would be more likely to generalize novel labels to the subordinate level when other subordinate-level exemplars were present and linguistically negated by the speaker during labeling, compared to when only the target referent was present and labelled. We tested this prediction by presenting learners with a novel label for a target object (e.g., a dalmatian) either in isolation or alongside two other subordinate-level alternatives from the same basic-level category (e.g., other dog breeds).[1]

2.1. Participants

Forty English-speaking 4-to-5-year-olds (M = 5;0, range = 4;0–5;11; 9 female, 19 male and 12 without reported gender) were recruited from local daycares in Philadelphia, and forty English-speaking adults were recruited from the undergraduate subject pool at the University of Pennsylvania and given course credit.

2.2. Procedure

Experiment 1 was administered via the PCIbex platform for web-based experiments (Zehr & Schwarz, 2018). Adult participants were tested virtually and asynchronously; child participants were tested in person by an experimenter who guided their interactions with the experiment on a laptop. Children received all instructions from the experimenter; adults read the same instructions on the screen. In the beginning of the experiment, participants were introduced to a cartoon character named Blue:

"We're going to play a game with a friend. Her name is Blue and she speaks a different language. Her language sounds like English but has some words that are different from English. Blue is looking for something from a pile of cards. Once she finds what she's looking for, it's your turn, so pay attention! Are you ready?"

The experiment then proceeded to the word learning task, consisting of eight trials. Each trial involved a learning phase where participants were exposed to a labelling event, and then a testing phase which probed whether participants generalized the meaning of the label at the subordinate or the basic level. In the learning phase, Blue appeared on the screen alongside a row of three cards, all facing down (Fig. 1). Blue began by stating what she was looking for (e.g., "I'm looking for a gavai!"). Then, she moved next to each card and began flipping them over one-by-one, stopping when the target was revealed.

In a between-participants design, we manipulated when the target was found during the card-flipping sequence. In the *No Contrast* condition, the first card that was revealed happened to be the target card (e.g., a dalmatian; Fig. 1A). When the target was revealed, it was marked by a red outline and labelled by Blue (e.g., "Look, I found what I was looking for! It's a gavai!"). Then Blue added, "I don't see another gavai", since only one entity (the target) was visible in this condition. Because Blue always stopped flipping cards upon finding the target, this meant that what was on the other two cards remained hidden at the moment of labelling (and would remain unknown to both Blue and the participant for that trial).

In the *Contrast* condition, Blue discovered the target (the dalmatian) on the third and last card flip (Fig. 1B). She had previously flipped over the other two cards that depicted two subordinate-level alternatives (a corgi and a terrier), and had made it clear that each was not the target by saying, "This is not a gavai". Again, when the target was found, it was marked by a red outline and labelled (e.g., "Look, I found what I was looking for! It's a gavai!"). In this condition, the two non-target cards were already known to both Blue and the participant prior to the labelling of the target, and remained visible on the screen once revealed.

At the end of the learning phase, the experimenter said, "Now it's your turn to find gavais. Are you ready?" and the trial moved into the test phase after getting confirmation from the participants. Participants' word meaning generalizations were tested using the immediate generalization paradigm, in which participants are shown a set of pictures, from which they are instructed to select all instances of the target word. This paradigm has been proven successful in probing the specificity of learners' word meaning hypotheses (Caplan, 2022; Choe & Papafragou, 2023; Emberson et al., 2019; Jenkins et al., 2015; Lewis & Frank, 2018; Spencer et al., 2011; Wang & Trueswell, 2019, 2022; Xu & Tenenbaum, 2007a). There were six items in the generalization set: two subordinate-level matches (e.

[1] In a prior study, Gelman et al. (1989) taught 3- and 5-year-olds subordinate nouns via negation using a visual alternative (e.g., "This [pointing to one dog] is a dog; it's a dingo. This [pointing to the other dog] is a dog; it's not a dingo"). However, the effect was conflated with anchoring to the basic level. Therefore, it remains open whether linguistic negation can facilitate subordinate meanings on its own.

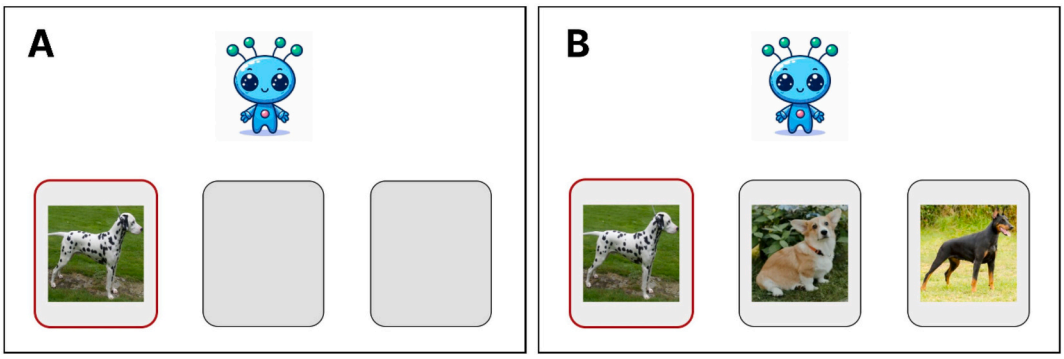

Fig. 1. Display at the moment of labelling the target (dalmatian) in the *No Contrast* (A) and *Contrast* (B) condition in a sample trial of Experiment 1.

Fig. 2. Images from the generalization set shown during the test phase of Experiment 1. Each image was presented individually.

g., two different dalmatians), two basic-level matches (e.g., two non-dalmatian dogs), and two distractors (e.g., a bird and a dresser; see Fig. 2). The items in the generalization set were shown one at a time, and for each item participants were asked whether they thought it was an instance of the target label (e.g., "Is this a gavai?"). The fact that participants did not get to see the full set of possible items up front increased the difficulty of the task but was done to minimize potential interference from the test phase in highlighting the relevance of subordinate-level categories (see Choe & Papafragou, 2023).

Blue's dialogue during the learning phase was played from pre-recorded audio from a female native speaker of English. The position of the target card in the row (leftmost or rightmost) was randomized between trials. The eight basic-subordinate category pairs used in critical trials alongside the novel target names were: dog/dalmatian (*gavai*), bus/school bus (*lorpim*), pepper/green pepper (*tuffee*), apple/red apple (*mipen*), dinosaur/T-rex (*tantol*), ball/basketball (*kapsin*), phone/smartphone (*spadil*), and chair/wooden chair (*kalmick*).

Participants' responses were coded using a strict criterion: a response was coded as Subordinate if only the two subordinate-level matches were selected, Basic if both the two basic-level matches were additionally selected, and Other in all other cases (Choe & Papafragou, 2023; Wang & Trueswell, 2019, 2022).[2]

2.3. Results

Starting with the data from adults, Fig. 3 shows the distribution of word meaning generalizations (coded as Subordinate, Basic, and Other) at test. As predicted, adults showed more subordinate-level generalizations of the target label in the Contrast condition (82.9%) compared to the No Contrast condition (58.8%). To statistically test for this difference, a logistic mixed-effects regression model was fitted to the proportion of Subordinate responses (Table 1), using the MixedModels library (Alday & Bates, 2025) in the Julia programming language (Bezanson et al., 2017). The model included a fixed effect of the Contrast manipulation (sum coded with No Contrast = -1 and Contrast = 1) with a maximum random effects structure up to by-participant and by-item random intercepts and slopes with correlations.[3] The model confirmed that Contrast was a significant predictor in the expected direction ($\beta = 2.160$, $SE = 0.673$, $p = .001$), consistent with the hypothesis that the presence of referential alternatives constrained word meaning generalizations at a narrower level.

Turning to the data from children, Fig. 4 shows the distribution of word meaning generalizations at test. Similar to adults, children showed more subordinate-level generalizations of the target label in the Contrast condition (64.4%) compared to the No Contrast condition (35.8%). A logistic mixed-effects regression model was fitted to the proportion of Subordinate responses (Table 2), with identical model specifications except for the addition of a predictor for Age (sum coded with four-year-olds = -1, five-year-olds = 1)

[2] Under an alternative coding scheme, incomplete generalizations receive partial scores (Lewis & Frank, 2018; Xu & Tenenbaum, 2007a). Even if we were to assign partial scores, the results would not change.

[3] For all reported models, any computational issues due to the overparameterization of the random effects structure in the initial maximal model were addressed through the iterative model reduction process suggested in Bates et al. (2018), whereby the correlation between random effects was dropped first and individual random effects were removed in the order of lowest variance until the model converged, in order to keep the random effects structure as informative as possible.

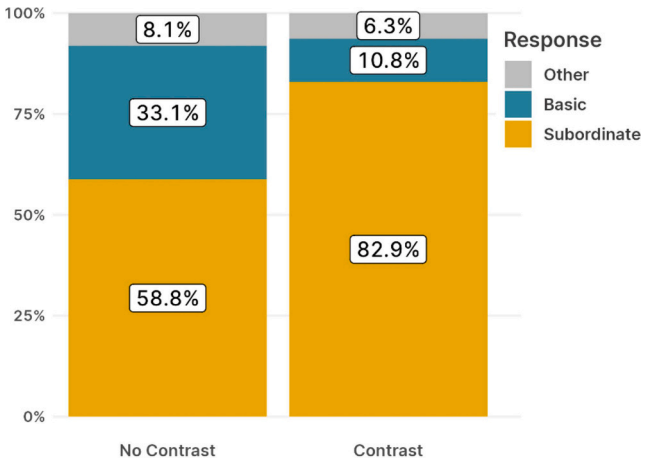

Fig. 3. Experiment 1 results (adults).

Table 1
Experiment 1: logistic mixed-effects model for Subordinate responses in adults.

	Estimate	*SE*	*t*	*p*
(Intercept)	1.976	0.521	3.792	<.001
Contrast	2.160	0.673	3.209	.001

and the interaction term. The model confirmed the predicted main effect of Contrast ($\beta = 1.190$, $SE = 0.237$, $p < 0.001$). Additionally, there was a main effect of Age, such that five-year-olds produced overall more Subordinate responses than four-year-olds ($\beta = 0.880$, $SE = 0.257$, $p = .010$). While not a core prediction from our hypothesis, this finding is consistent with the expectation that older children would be more knowledgeable about category names and concepts; five-year-olds may have been more capable of accessing the subordinate-level categories for the kinds of familiar objects that we used in the experiment. Lastly, the interaction effect was not significant ($\beta = 0.386$, $SE = 0.233$, $p = .097$).

One striking pattern in the child data is the large proportion of Other responses, totaling approximately a third of all responses. This was not entirely unexpected given the strict coding scheme and the more difficult sequential design of the test phase (Choe & Papafragou, 2023) but merited further analysis. Looking within the Other responses, we distinguished and coded for Incomplete Subordinate and Incomplete Basic responses (e.g., selection of just one of two dalmatians or just three of four dogs, respectively), as well as No-match responses which failed to generalize at all (i.e., no selection of any dalmatians). Additionally, because one of the distractor items in the generalization set was always an instance of the same global kind as the target (e.g., a bird, which is a kind of an animal), we reasoned that Superordinate generalizations were also a valid response for children to make (e.g., selecting all 5 animals, excluding the dresser), though those were not observed in adults. The overall distribution of these granular categories within the Other

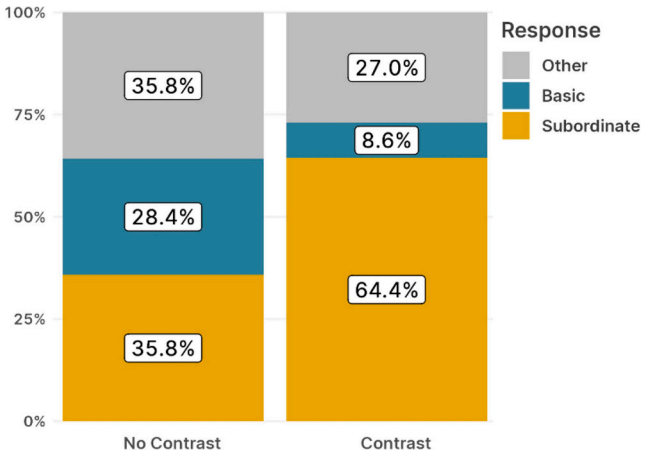

Fig. 4. Experiment 1 results (children).

Table 2
Experiment 1: logistic mixed-effects model for Subordinate responses in children.

	Estimate	*SE*	*t*	*p*
(Intercept)	0.112	0.535	0.209	.835
Contrast	1.190	0.237	5.011	<.001
Age	0.880	0.343	2.569	.010
Contrast * Age	0.386	0.233	1.659	.097

responses was: Incomplete Basic (47.1%), Incomplete Subordinate (25.5%), No-match (13.7%), and Superordinate (13.7%). Moreover, the shape of this distribution remained largely similar whether splitting by condition or age group. The large proportion of incomplete responses suggests that children may have, in fact, generated the subordinate and basic-level categories as the intended hypotheses for word meaning but simply had difficulty selecting all the necessary items from the generalization set at test. Corroborating this interpretation, a separate post-hoc model fitted to Other responses found only a significant main effect of Age ($\beta = -0.794$, $SE = 0.287$, $p = .006$): there were significantly fewer Other responses in five-year-olds (14.8%) than four-year-olds (39.3%), perhaps due to a greater availability of attentional resources in older children. We revisit the options within the generalization set in Experiment 2.

2.4. Discussion

In Experiment 1, we investigated whether children's inferences about word meaning are influenced by the presence of contrasting alternatives at the subordinate level. We found that 4-5-year-olds, like adults, were significantly more likely to interpret a novel label as referring to a subordinate-level category (e.g., dalmatian) when they observed contrasting alternatives from the same basic-level category (e.g., other dog breeds) compared to cases where no such alternatives were visible. This effect was robust across age groups, with five-year-olds showing an overall stronger tendency toward subordinate-level generalizations than four-year-olds. Importantly, linguistically negating that these visual alternatives shared a label with the target ("This is not a gavai") helped generate the contrastive inference promoting subordinate-level meanings for the target label *even though the visual alternatives were not themselves given unique labels*.

These findings are consistent with our hypothesis that inferences about word meaning are guided by pragmatic reasoning about the speaker's intent and the expectation for the speaker to be informative. Our findings complement prior theorizing of contrast as a pragmatic phenomenon (Choe & Papafragou, 2023; Clark, 1987, 1990; Diesendruck, 2005) and expand on the well-documented facilitative role of contrastive information in subordinate noun learning (Clark et al., 1985; Clark, 1987; Waxman, 1990; Waxman et al., 1991, 1997; and especially Waxman & Hatch, 1992) by showing that such effects can emerge through negation. However, our findings are hard to reconcile with proposals that view contrast only as a simple surface heuristic operating on two co-present competing words with mutually exclusive one-to-one mappings (e.g., Golinkoff et al., 1992, 1994; Markman, 1987, 1989). Relatedly, our results cannot be fully explained by perceptual accounts (e.g., Jenkins et al., 2015, 2021; Spencer et al., 2011) or other accounts of subordinate noun learning which do consider speaker intent as a filter but nevertheless rely on probabilistic inferences over positive evidence accumulated across multiple labelling events (e.g., Lewis & Frank, 2016; Xu & Tenenbaum, 2007b). While mechanistic and lexical heuristic accounts successfully capture effects of perceptual similarity and competing linguistic forms, respectively, neither predicts sensitivity to the speaker's epistemic state of the kind observed here.

In our design, participants in both conditions saw identical target exemplars and learned the target label from a single labelling event, with the critical difference being whether the alternative exemplars were also visible. A robust effect of contrast from this design suggests that learners were not simply tracking the perceptual or distributional properties of the target referent and the label but also taking into account the broader referential context to form expectations about the intent behind the choice of a word (Choe & Papafragou, 2023; cf. also Wang & Trueswell, 2019).

3. Experiment 2

Experiment 1 offered strong evidence that contrast via a combination of referential alternatives and linguistic negation serves as a cue to subordinate noun meaning. However, because this information about contrast was always shared by both the learner and the speaker, it was less clear whether (as our pragmatic account predicts) the effect of contrast involved rich pragmatic computations over *the specific speaker's* perspective and intent. Experiment 2 presents a stronger test of our pragmatic proposal by asking whether information about contrast is further sensitive to the speaker's knowledge of the visual alternatives that participate in the contrast. In a slightly modified design, we held the learner's knowledge of the alternatives constant while manipulating the speaker's knowledge of them. If, as we have proposed, the effect of contrast is driven by inferences about what the speaker intends to communicate, then learners should be more likely to generalize novel word meanings at the subordinate level when the alternatives are known to the speaker. In that context, the speaker is expected to be more specific in their label for the target object given the need to contrast it with the alternatives. On the other hand, if the visual contrasts are not available to the speaker, children may be more likely to interpret the speaker's label as encoding the broader basic-level category, since that is informative enough for reference from the perspective of a speaker with partial knowledge.

3.1. Participants

Forty English-speaking 4-to-5-year-olds (M = 5;3, range = 4;0–5;11; 19 female, 21 male) were recruited online from across the United States through social media posts. Their families were given a $10 gift card for participating. Forty English-speaking adults were also recruited from the undergraduate subject pool at the University of Pennsylvania. They were given course credit for participating.

3.2. Procedure

Experiment 2 was administered via the PCIbex platform for web-based experiments. Adult participants were again tested virtually and asynchronously; child participants were tested virtually by an experimenter who guided their interactions with the experimental materials over a zoom call. In the beginning of the experiment, participants were introduced to two characters, a boy named Bob and a girl named Sally, whose voices played from pre-recorded audio. For children, Sally and Bob were introduced by the experimenter who also communicated all subsequent instructions. For adults, instructions were identical but delivered in written form:

> "Let me introduce you to two friends: this is Sally ("Hi I'm Sally!"), and this is her friend Bob ("Hi I'm Bob!"). They speak a language that sounds like English but has some words that are different from English. Let's go see what they're doing! Are you ready?"

Participants were then shown two short videos designed to familiarize them with the display used in the task. The purpose of the videos was to ensure that participants could establish the link between a speaker's presence in the scene and their visual access to the objects that were also present in the scene. In the first video, Sally stood next to a table with a toy on it, turned to look at the toy, and then turned back to face the camera. The experimenter then asked, "Does Bob know what's on the table?" and recorded children's responses; adults entered their answer into a text box. All adults and all but two children responded with the correct (negative) answer to the question. For the remaining two children, the experimenter reminded them that "Bob doesn't know what's on the table because he is not in the room looking at what's on the table." The second video was identical, but with Bob now present, standing at the opposite side of the table from Sally. Both characters turned to look at the toy on the table, then turned back to face the camera. The experimenter then asked, "How about now – does Bob know what's on the table?" All participants responded with the correct (affirmative) answer to this question.

After participants were familiarized with the characters and the visual scene, they received the following instructions for the main task: "Now we're going to see them play with cards which have pictures on them. And sometimes, Sally will play a trick on Bob. So pay attention to what happens, okay?" As in Experiment 1, each trial involved a learning phase followed by a test phase. The learning phase presented a short scenario played over the course of three frames, with pre-recorded audio dialogue accompanying the frames (Fig. 5). In the first frame, Sally stood next to three cards which were placed on the table. The images on the cards consisted of three exemplars from the same basic-level category (e.g., dog) that contrasted at the subordinate level (e.g., different dog breeds; a dalmatian, a corgi, and a terrier). One of the three cards served as the target card (e.g., the dalmatian) that would later be labelled by Bob. The position of the target card was randomized between trials (leftmost vs. rightmost). The number of trials and the items used in the learning phase were identical to Experiment 1. During the first frame, participants heard Sally say, "Look, do you see the cards? (Pause.) Oh, I hear someone coming." Speaker Knowledge was then manipulated between participants. In the *Contrast Visible* condition, Sally said nothing further. In the *Contrast Hidden* condition, Sally added, "I'm gonna hide them."

In the second frame, Bob had entered the scene and was looking at what was on the table. In the Contrast Visible condition, Bob saw all three cards and said, "Oh, nice cards!". In the Contrast Hidden condition, Bob saw only the target card (the others had been hidden) and said, "Oh, a nice card!". In the final frame, Bob picked up the target card, showed it to the participants and said, "See? This is a gavai. Do you see the gavai? Help me find more gavais!" To wrap up the learning phase, the experimenter said, "Let's find more gavais."

At the test phase, participants were presented with a horizontal array of 6 images. For children, an animated hand moved sequentially over each image, and the experimenter said, "Do you see the moving hand? Tell me when the hand points to a gavai, okay?" The children verbally responded when the hand pointed to an image that they believed was also an instance of the target label. If they selected zero or just one image, the experimenter followed up with "Are there any (more) gavais?" to encourage complete responses (cf. Xu & Tenenbaum, 2007a). Adults freely clicked on their choices within the generalization set without the help of the moving hand.

The 6 items consisted of two subordinate matches (e.g., two different dalmatians), two basic-level matches (two different, non-dalmatian dogs), and two distractors (two different pieces of furniture; Fig. 6). These items were identical to the generalization set of Experiment 1 except that an earlier distractor from the same superordinate level as the target (e.g., bird) was replaced for simplicity so that both distractors were from a different semantic domain from the target (recall that Experiment 1 did not find a strong tendency for children to generalize at the superordinate level). Unlike Experiment 1, the six items in the generalization set were presented simultaneously to alleviate cognitive load for children. The order of images in the set was randomized between trials. Participants' final choices were recorded and coded as Subordinate, Basic, and Other responses for analysis.

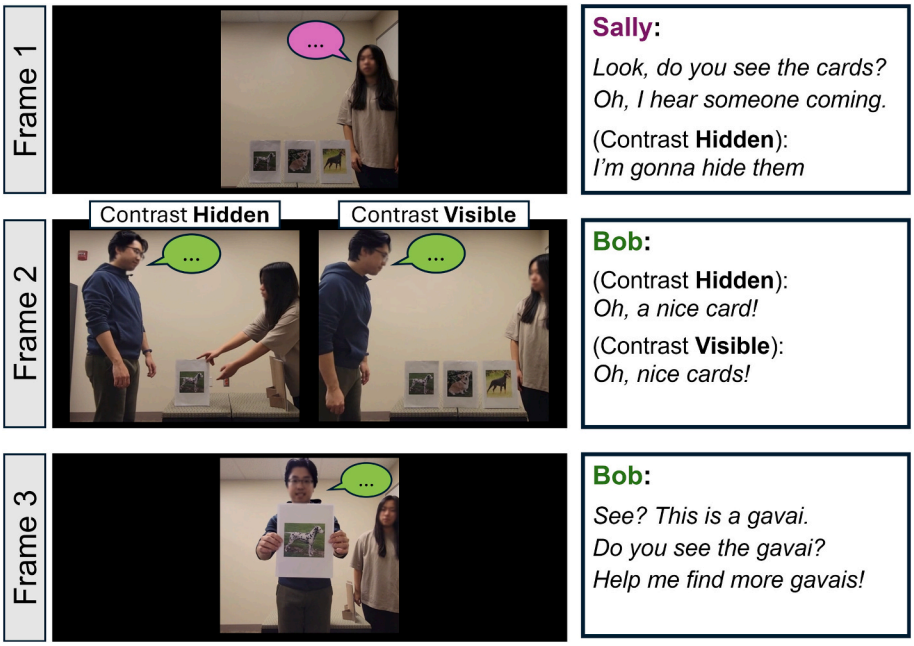

Fig. 5. A schematic of the learning phase from a critical trial in Experiment 2.

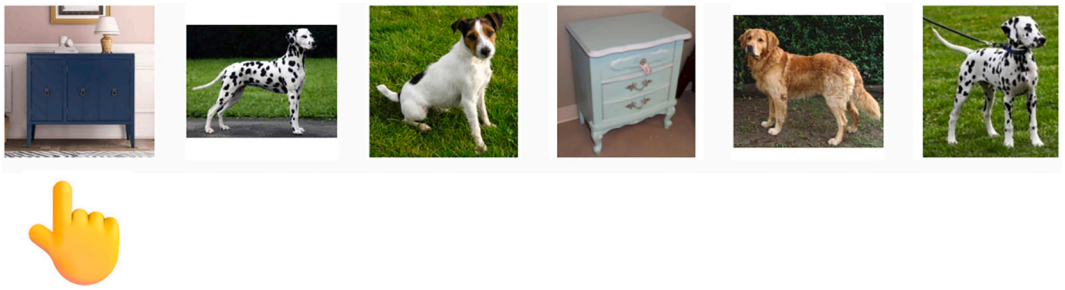

Fig. 6. Array of images from the generalization set shown in the test phase of Experiment 2. The animated hand was only present for children.

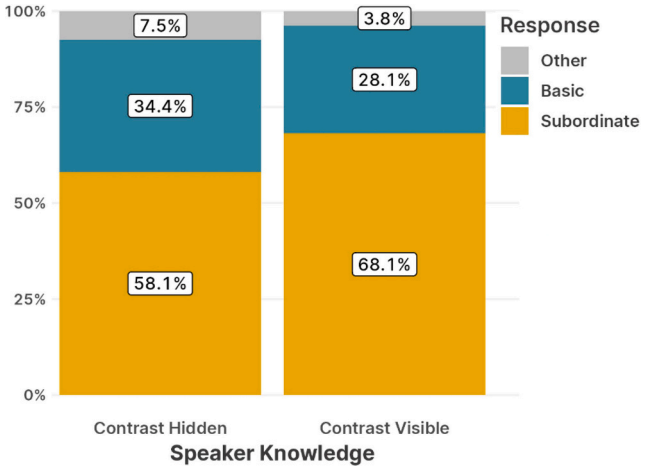

Fig. 7. Experiment 2 results (adults).

3.3. Results

Starting with the data from adults, Fig. 7 shows the distribution of word meaning generalizations at test across conditions. As expected, adults were more likely to generalize the novel label at the narrower, subordinate-level category when the speaker had access to the subordinate-level alternatives (68.1%, Contrast Visible condition) compared to when the speaker did not (58.1%, Contrast Hidden condition). A logistic mixed-effects regression model was fitted to the rate of Subordinate responses, with Speaker Knowledge as a fixed effect (sum coded: Contrast Hidden = -1, Contrast Visible = 1) and a maximal random effects structure up to by-participant and by-item intercepts and slopes, including correlations (Table 3). The model found a significant main effect of Speaker Knowledge in the expected direction ($\beta = 0.966$, $SE = 0.343$, $p = .005$).

Turning to the data from children, Fig. 8 shows the distribution of word meaning generalizations at test. Like adults, children showed more subordinate-level generalizations of the target label in the Contrast Visible condition (45.0%) compared to in the Contrast Hidden condition (31.9%). A logistic mixed-effects regression model was fitted to the proportion of Subordinate responses (Table 4), with identical model specifications except for the addition of a predictor for Age (sum coded with 4-year-olds = -1, 5-year-olds = 1) and the interaction term. The model confirmed the predicted main effect of Speaker Knowledge ($\beta = 0.515$, $SE = 0.189$, $p = .006$). The main effect of Age was not significant ($\beta = 0.189$, $SE = 0.421$, $p = .654$), though the direction of effect was the same as in Experiment 1. The interaction effect was also not significant ($\beta = -0.071$, $SE = 0.190$, $p = .707$). Notice that children made fewer Other responses overall compared to Experiment 1 (i.e., they were more likely to make complete generalizations to the subordinate and basic-level categories), most likely due to the simplifications to the test phase made in Experiment 2.

3.4. Discussion

In Experiment 2, we found that both children and adults were more likely to interpret a novel label as referring to a subordinate-level category when the speaker had knowledge about the subordinate-level alternatives compared to when those alternatives were unknown to the speaker. Because learners themselves always saw the visual alternatives in both conditions, this effect was tied to what the learner could infer about the speaker's mental state and, more specifically, what the speaker must have meant by their choice of the target label, given the Gricean injunction for the speaker to be informative and relevant. When the speaker had visual access to subordinate-level alternatives, participants inferred that the speaker had to be more informative and thus were more likely to generalize the meaning of the novel label at the subordinate level. However, when the speaker could only see the target object, such a demand on lexical specificity could not be expected of the speaker. As a result, learners were more likely to ignore the visual alternatives that they themselves (unlike the speaker) had seen and generalize the meaning of the novel label at the basic level.

In interpreting the results of Experiment 2, two patterns in the data merit discussion. First, looking across experiments, we observe generally lower Subordinate responses in the Contrast Visible condition of Experiment 2 (adults: 68.1%, children: 45.0%) than in the Contrast condition of Experiment 1 (adults: 82.9%, children: 64.4%). This is expected, since the design of Experiment 2 did not semantically exclude the basic level as a possible word meaning conjecture: Experiment 1 used explicit linguistic negation (e.g., "This is not a gavai") whereas Experiment 2 relied on the subtler cue of the speaker's choice to pick up and talk about just one of three items on the table. Indeed, a basic-level interpretation remained logically available in Experiment 2, since the speaker could have continued with, for example, "This is a gavai… those two are gavais too but I want to show you this one." Thus, the basic-level meaning always remained available as a logically licensed option, even if, as we predicted and confirmed in the data, that was a pragmatically dispreferred interpretation given the speaker's selective reference when all three exemplars were available in the common ground.

Second, within Experiment 2, one may wonder whether the seemingly even split between Subordinate and Basic responses in the Contrast Visible condition, particularly in the child data (45% vs. 51.9% respectively), is an indication of participants responding at random. However, two considerations argue against this interpretation. First, recall that each trial presented six test items including two unrelated distractors. If children were responding haphazardly to individual items (presumably the simpler version of responding at random), we would expect a high rate of Other responses. However, the extremely low rate of such responses in Experiment 2 suggests that children were, at minimum, engaging the task with enough intent and attention to consistently select coherent category-level generalizations. Second, an even split may not be the appropriate baseline for chance, given the well-documented basic-level bias that is notoriously difficult to overcome. If children were defaulting to prior expectations or disregarding the speaker knowledge manipulation entirely, we would have expected higher Basic responses as well as a similar distribution across both conditions, since the perceptual information at the moment of labelling was identical. Instead, we observe a systematic shift toward subordinate-level responses when the speaker had visual access to the alternatives.[4]

Taken together, the results of Experiment 2 provide strong evidence that word meaning inferences are guided by reasoning about speaker knowledge and intent, above and beyond information that is privately observed and known by the learner. While mechanistic and lexical heuristic accounts capture effects of perceptual similarity and competing linguistic forms, respectively, neither predicts sensitivity to the speaker's epistemic state of the kind observed here. As Experiment 1 and prior literature have shown, overt forms of

[4] Even when we set chance levels to 0.50, a two-tailed one-sample t-test comparing by-participant proportions of adults' Subordinate responses to chance shows no significant difference in the Contrast Hidden condition ($p = .187$), with adults displaying no bias for or against subordinate meanings, but a significant difference in the Contrast Visible condition ($p = .007$), where adults displayed a bias for subordinates. In children, the same test shows a significant difference in the Contrast Hidden condition ($p = .002$), with children displaying a bias against subordinate meanings (and in favor of basic-level ones; see Figure 8), but not in the Contrast Visible condition ($p = .421$), where no such bias was found. Overall, for both groups, the Contrast Visible moves learners away from the basic level, as we have argued for (but, for children specifically, it does not install a preference for subordinate meanings.)

Table 3
Experiment 2: logistic mixed-effects model for Subordinate responses in adults.

	Estimate	SE	*t*	*p*
(Intercept)	1.633	0.662	2.467	.014
Speaker Knowledge	0.966	0.343	2.817	.005

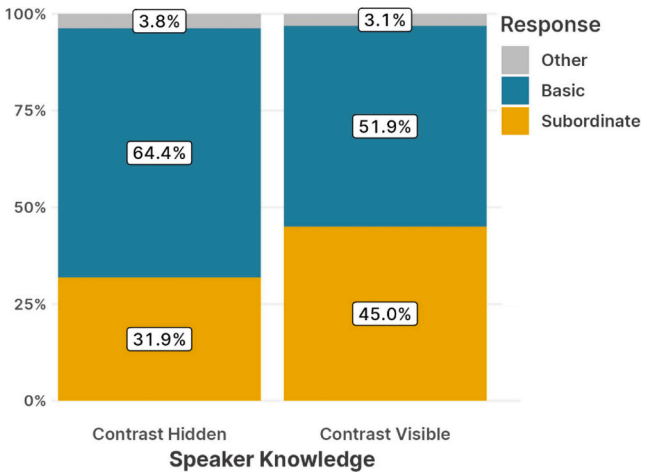

Fig. 8. Experiment 2 results (children).

Table 4
Experiment 2: logistic mixed-effects model for Subordinate responses in children.

	Estimate	*SE*	*t*	*p*
(Intercept)	−0.825	0.435	−1.895	.058
Speaker Knowledge	0.515	0.189	2.727	.006
Age	0.189	0.421	0.448	.654
Speaker Knowledge * Age	−0.071	0.190	−0.376	.707

contrast offer a strong cue to subordinate-level inferences. But here these inferences emerge from more general reasoning about why a speaker would choose to be specific in their reference given the information that is mutually available in the common ground and relevant in the discourse context.

4. General discussion

4.1. Acquiring subordinate nouns

In this study, we proposed that the acquisition of subordinate nouns poses a pragmatic puzzle involving inferences about the speaker's intent. Under broadly Gricean and post-Gricean accounts, listeners engage in a rich computational process that considers the speaker's communicative intentions given their goals and knowledge state (e.g., Grice, 1975; Horn, 1972; Sperber & Wilson, 1986). Applied to the domain of word learning, these approaches predict that learners filter their interpretation of the available evidence for word meaning through the speaker's perspective (see review in Grigoroglou & Papafragou, 2021). We have argued that information from contrast is one such piece of evidence: contrast is a helpful cue to word meaning precisely because it highlights a pragmatic pressure on the speaker, who strives to be informative, to be more specific than usual in their choice of a word. Building on prior findings showing that contrastive cues can reliably facilitate children's subordinate-level categorization and labeling (e.g., Choe & Papafragou, 2023; Clark, 1987; Clark et al., 1985; Klibanoff & Waxman, 2000; Wang & Trueswell, 2019; Waxman, 1990; Waxman et al., 1991, 1997; Waxman & Hatch, 1992; Waxman & Klibanoff, 2000), we hypothesized that young children additionally recruit epistemic reasoning about the speaker's knowledge and intent in learning subordinate nouns.

We tested and confirmed predictions from this proposal across two experiments. In Experiment 1, both children and adults were more likely to generalize novel labels at the subordinate level when contrastive alternatives from different subordinate-level categories were present and linguistically negated during the learning phase, compared to when only the target referent was visible. This effect held even though participants in both conditions saw identical target exemplars and learned from a single labeling event, suggesting that learners were able to reason about the speaker's intent by interpreting the linguistically-negated labels as semantic alternatives to

the target word's meaning. Complementing this interpretation of the results, in Experiment 2 we found that this contrastive inference was further modulated by the speaker's knowledge state: learners were more likely to generalize novel labels to the subordinate-level meaning when the speaker had visual access to the alternatives, compared to when those were hidden from the speaker's view. Since learners themselves always observed the contrasting alternatives in both conditions, this finding suggests that word meaning inferences were guided by reasoning about what the speaker knew and what level of specificity was appropriate for the speaker to adopt.

Combined, these results provide strong evidence for a pragmatic account of subordinate noun learning. Specifically, our data support the role of contrast in the acquisition of subordinate meanings in accordance with much prior literature (see especially Waxman, 1990; Waxman et al., 1991, 1997; Waxman & Hatch, 1992, among others); furthermore, they break new ground by clarifying the mechanism whereby contrast exerts its effects. Rather than contrast being a simple surface-level phenomenon involving two competing form-to-meaning mappings or immediate perceptual comparisons, we demonstrate that contrast is fundamentally a cue to speaker intent. As such, contrastive inferences are also available when visual alternatives are linguistically negated (Experiment 1) and when they are merely present for the speaker to consider (Experiment 2).

One limitation of the present study is the use of relatively familiar categories. We chose such categories deliberately because our experiments (especially Experiment 2) were already long and demanding for young children. Given these constraints, we selected items that closely followed the materials used in prior studies on subordinate noun learning with children (e.g., Jenkins et al., 2015; Wang & Trueswell, 2019; Waxman et al., 1997; Xu & Tenenbaum, 2007a). As in these past studies, only some of the subordinate-level categories we used had a single existing label (e.g., dalmatian, schoolbus), whereas others did not (e.g., wooden chair, green pepper). Thus, it is unlikely that the children simply translated the new words into words they already knew. Furthermore, and more importantly, in a past study involving individual differences (Jenkins et al., 2015), children's prior category knowledge was shown to affect generalization rates by shifting the baseline toward more subordinate-level (vs. basic-level) interpretations for the specific items that were more familiar (the shift did not interact with the experimental manipulation). We controlled for this effect by counterbalancing items with conditions (e.g., Contrast, Speaker Knowledge) across experiments. Given that any item-level familiarity effects were thus distributed equally across conditions, we take our findings to reflect the contribution of pragmatic reasoning over and above whatever baseline familiarity with the categories children brought to the task. Future work using a greater variety of categories, including novel categories, would further strengthen this conclusion.

4.2. Speaker knowledge and pragmatics

Our findings contribute to a well-studied but complex picture regarding young children's emerging perspective-taking abilities. Sensitivity to others' knowledge and communicative intent has roots in early development: by 12 months, infants communicate informatively by pointing to objects whose location the interlocutor does not know about (Liszkowski et al., 2008), and by 19 months, infants can use the speaker's gaze direction to guide word-object mappings between a label and a salient but unintended referent (Baldwin, 1993). However, integrating speaker knowledge into inferences about the precise *meaning* of a word, rather than simply a possible referent, poses additional challenges: the learner must not only track what the speaker can see, but also use that information to constrain an internal hypothesis space over possible word meanings. On this and other more challenging problems in referential communication, the literature has long documented that, while children around 5 years of age can successfully integrate information about a speaker's perspective in disambiguating the intended referent, this capacity is highly sensitive to task demands throughout the preschool years (e.g., Carpenter et al., 2002; Epley et al., 2004; Houston-Price et al., 2011; Nilsen & Graham, 2009; Papafragou et al., 2017; Southgate et al., 2010; Srinivasan et al., 2019; Wilson et al., 2022), with reports that even adults sometimes act egocentrically (e.g., Horton & Gerrig, 2005; Wardlow Lane & Ferreira, 2008). On one interpretation, children's difficulties stem not from an inability to represent others' mental states *per se*, but from the considerable cognitive effort required for tracking and integrating this knowledge while simultaneously managing competing information sources. This is further corroborated by findings that performance correlates strongly with executive function skills, particularly inhibition (e.g., Nilsen & Graham, 2009). Our own experimental design incorporated cues that may have helped children overcome these cognitive factors by establishing a clear visual link to speaker knowledge (via the presence vs. absence of the speaker in the scene) and giving children explicit goals (to help the speaker "find more gavais"; cf. Nadig & Sedivy, 2002; Kampa & Papafragou, 2020). Children's success in our study offers independent evidence for children's perspective-taking abilities beyond classic referential paradigms.

Additionally, the present data offer one of the earliest demonstrations of children's success with informativeness computations in the domain of word learning. Recall that the subordinate vs. basic-level distinction is logically defined in terms of informativeness; in addition, the basic level typically satisfies the informational needs of a generic addressee, while the subordinate level is more appropriate in contexts where being more informative than usual is called for. On this perspective, taxonomic nouns can be understood as a scalar phenomenon (Choe & Papafragou, 2023), and while taxonomic specificity is unlike classical scales (e.g., quantification) in many ways, we observe similar themes. For example, children's success depends critically on the relevance of the stronger alternatives: just as children compute the pragmatic under-informativeness of *some* more successfully when the relevant alternative *all* or even *none* is made salient in the task (Skordos & Papafragou, 2016), children in our experiments increased subordinate (vs. basic-level) generalizations when the presence of semantic alternatives highlighted a contrast at the subordinate level of categorization. Moreover, the prior finding that children expect informative descriptions from speakers with full knowledge (Kampa & Papafragou, 2020) aligns with our finding that children recognized that knowledgeable speakers must have intended to be more informative in their choice of a (subordinate) word, given the relevant alternatives. These converging sources of evidence point to a more general capacity for sophisticated pragmatic computations in young children, one that integrates considerations of informativeness and relevance with the epistemic state of the interlocutors.

4.3. Pragmatics and word learning

Finally, our findings contribute to a rich line of research on the psychological mechanisms that guide early word learning. Classic work in language acquisition has long recognized the fundamental problem of mapping words to their intended meanings (Fisher, 1996; Gleitman, 1990; Landau & Gleitman, 1985; Lidz et al., 2003; Naigles, 1990; Papafragou et al., 2007). This is a seemingly insurmountable challenge when most word learning contexts in natural parent-child interactions are referentially ambiguous (e.g., Cartmill et al., 2013; Gillette et al., 1999; Kako, 2005). Even in the rare cases of referential certainty involving familiar entities, ambiguity in word meaning persists, as in the "subset problem" of identifying the intended level of a taxonomic description (Quine, 1960), as we have discussed at length.

In navigating this challenge, children from the earliest stages of life leverage various filtering mechanisms on the noisy and indeterminate input to identify the rare moments of clarity that hold the richest clues to a word's meaning (see review in Gleitman & Trueswell, 2020). We posit that children's pragmatic reasoning about the speaker's perspective constitutes an important filtering mechanism: by tracking information in the common ground and available to the speaker, children can constrain the space of possible word meanings that could have been intended by the speaker's choice of a specific word. Our findings support the broader perspective that word learning is fundamentally driven by principles of communication (see also Clark, 2017; Clark & Wong, 2002; Diesendruck, 2005; Diesendruck et al., 2004; Grigoroglou & Papafragou, 2021). The domain of hierarchical noun meanings is thus one instance of the pervasive involvement of pragmatic mechanisms in language acquisition, comprehension, and use.

CRediT authorship contribution statement

June Choe: Writing – original draft, Methodology, Formal analysis, Conceptualization. **Anna Papafragou:** Writing – review & editing, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors are grateful to the members of the Language and Cognition lab at the University of Pennsylvania for providing valuable feedback on this project. We are especially grateful to Olivia Maltz, Madison Paron and Jamie Yuen for their help with data collection. We also thank all the preschools and families that participated in the experiments.

Data availability

The data that support the findings of this study are openly available in OSF at https://osf.io/ure43/overview?view_only=bad14fdbf1bf48b990d95023a900ebf3

References

Alday, P. M., & Bates, D. (2025). *Juliastats/mixedmodels.jl: v4.35.0. [Computer software]*. https://doi.org/10.5281/zenodo.8337373
Au, T. K., & Glusman, M. (1990). The principle of mutual exclusivity in word learning: To honor or not to honor? *Child Development, 61*, 1474–1490.
Baldwin, D. A. (1993). Early referential understanding: Infants' ability to recognize referential acts for what they are. *Developmental Psychology, 29*(5), 832–843.
Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition, 118*, 84–93.
Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2018). Parsimonious mixed models. *ArXiv e-prints*. https://doi.org/10.48550/ARXIV.1506.04967
Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review, 59*(1), 65–98.
Blewitt, P. (1994). Understanding categorical hierarchies: The earliest levels of skill. *Child Development, 65*, 1279–1298.
Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*(6), 1482–1493.
Brown, P., & Dell, G. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology, 19*(4), 441–472.
Callanan, M. (1985). How parents label objects for young children: The role of input in the acquisition of category hierarchies. *Child Development, 56*(2), 508–523.
Callanan, M. (1989). Development of object categories and inclusion relations: Preschoolers' hypotheses about word meanings. *Developmental Psychology, 25*(2), 207–216.
Callanan, M., Repp, A., McCarthy, M., & Latzke, M. (1994). Children's hypotheses about word meanings: Is there a basic level constraint? *Journal of Experimental Child Psychology, 57*, 108–138.
Callanan, M. A., & Markman, E. M. (1982). Principles of organization in young children's natural language hierarchies. *Child Development, 53*(4), 1093–1101.
Caplan, S. (2022). *The immediacy of linguistic computation* [Doctoral dissertation, University of Pennsylvania].
Carpenter, M., Call, J., & Tomasello, M. (2002). A new false belief test for 36-month-olds. *British Journal of Developmental Psychology, 20*, 393–420.
Cartmill, E. A., Armstrong, B. F., Gleitman, L. R., Goldin-Meadow, S., Medina, T. N., & Trueswell, J. C. (2013). Quality of early parent input predicts child vocabulary 3 years later. *Proceedings of the National Academy of Sciences of the United States of America, 110*(28), 11278–11283.
Choe, J., & Papafragou, A. (2023). The acquisition of subordinate nouns as pragmatic inference. *Journal of Memory and Language, 132*, Article 104432.
Choe, J., & Papafragou, A. (2025). Distributional signatures of superordinate nouns. *Language Acquisition*, 1–21.
Clark, E. (1987). The principle of contrast: A constraint on language acquisition. In B. MacWhinney (Ed.), *Mechanisms of language acquisition*. Hillsdale, NJ: Lawrence Earlbaum Assoc.
Clark, E. (1990). On the pragmatics of contrast. *Journal of Child Language, 17*, 417–431.
Clark, E. (1997). Conceptual perspective and lexical choice in acquisition. *Cognition, 64*, 1–37.

Clark, E., & Wong, A. (2002). Pragmatic directions about language use: Offers of words and relations. *Language in Society, 31*(2), 181–212.
Clark, E. V. (2017). Semantic categories in acquisition. In H. Cohen, & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 397–421). Elsevier.
Clark, E. V., Gelman, S. A., & Lane, N. M. (1985). Compound nouns and category structure in young children. *Child Development, 56*(1), 84–94.
Cruse, D. (1977). The pragmatics of lexical specificity. *Journal of Linguistics, 13*(2), 153–164.
Degen, J., Hawkins, R. D., Graf, C., Kreiss, E., & Goodman, N. D. (2020). When redundancy is useful: A Bayesian approach to "overinformative" referring expressions. *Psychological Review, 127*(4), 591–621.
Diesendruck, G. (2005). The principle of conventionality and contrast in word learning: an empirical examination. *Developmental Psychology, 41*(3), 451–463.
Diesendruck, G., Carmel, N., & Markson, L. (2010). Children's sensitivity to the conventionality of sources. *Child Development, 81*(2), 652–668.
Diesendruck, G., Markson, L., Akhtar, N., & Reudor, A. (2004). Two-year-olds' sensitivity to speakers' intent: An alternative account of Samuelson and Smith. *Developmental Science, 7*(1), 33–41.
Emberson, L. L., Loncar, N., Mazzei, C., Treves, I., & Goldberg, A. E. (2019). The blowfish effect: Children and adults use atypical exemplars to infer more narrow categories during word learning. *Journal of Child Language, 46*(5), 938–954.
Epley, N., Morewedge, C. K., & Keysar, B. (2004). Perspective taking in children and adults: Equivalent egocentrism but differential correction. *Journal of Experimental Social Psychology, 40*(6), 760–768.
Fisher, C. (1996). Structural limits on verb mapping: The role of analogy in children's interpretations of sentences. *Cognitive Psychology, 31*(1), 41–81.
Gelman, S. A., Wilcox, S. A., & Clark, E. V. (1989). Conceptual and lexical hierarchies in young children. *Cognitive Development, 4*, 309–326.
Gillette, J., Gleitman, L., Gleitman, H., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition, 73*, 135–176.
Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition, 1*, 3–55.
Gleitman, L., & Trueswell, J. (2020). Easy words: Reference resolution in a malevolent referent world. *Topics in Cognitive Science, 12*(1), 22–47.
Golinkoff, R. M., Hirsh-Pasek, K., Bailey, L. M., & Wenger, N. R. (1992). Young children and adults use lexical principles to learn new nouns. *Developmental Psychology, 28*(1), 99–108.
Golinkoff, R. M., Mervis, C. B., & Hirsh-Pasek, K. (1994). Early object labels: The case for a developmental lexical principles framework. *Journal of Child Language, 21*(1), 125–155.
Golinkoff, R. M., Shuff-Bailey, M., Olguin, R., & Ruan, W. (1995). Young children extend novel words at the basic level: Evidence for the principle of categorical scope. *Developmental Psychology, 31*(3), 494–507.
Grice, P. (1975). Logic and conversation. In P. Cole, & J. Morgan (Eds.), *Syntax and semantics (Volume 3)*. New York: Academic Press.
Grigoroglou, M., & Papafragou, A. (2019). Children's (and adults') production adjustments to generic and particular listener needs. *Cognitive Science, 43*, 12790.
Grigoroglou, M., & Papafragou, A. (2021). Contributions of pragmatics to word learning and interpretation. In A. Papafragou, J. Trueswell, & L. Gleitman (Eds.), *The Oxford handbook of the mental lexicon*. Oxford: Oxford University Press.
Horn, L. R. (1972). *On the semantic properties of logical operators in English* [Doctoral dissertation, UCLA].
Horton, W. S., & Gerrig, R. J. (2005). The impact of memory demands on audience design during language production. *Cognition, 96*(2), 127–142.
Houston-Price, C., Goddard, K., Séclier, C., Grant, S. C., Reid, C. J., Boyden, L. E., & Williams, R. (2011). Tracking speakers' false beliefs: Is theory of mind available earlier for word learning? *Developmental Science, 14*, 623–634.
Jenkins, G. W., Samuelson, L. K., Penny, W., & Spencer, J. P. (2021). Learning words in space and time: Contrasting models of the suspicious coincidence effect. *Cognition, 210*.
Jenkins, G. W., Samuelson, L. K., Smith, J. R., & Spencer, J. P. (2015). Non-Bayesian noun generalization in 3-to 5-year-old children: Probing the role of prior knowledge in the suspicious coincidence effect. *Cognitive Science, 39*(2), 268–306.
Kako, E. (2005). Information sources for noun learning. *Cognitive Science, 29*, 223–260.
Kampa, A., & Papafragou, A. (2020). Four-year-olds incorporate speaker knowledge into pragmatic inferences. *Developmental Science, 23*(3), 12920.
Kampa, A., & Papafragou, A. (2023). Children and adults use pragmatic principles to interpret non-linguistic symbols. *Journal of Memory and Language, 131*, Article 104429.
Kampa, A., Richards, C., & Papafragou, A. (2024). Preschool children generate quantity inferences from both words and pictures. *Journal of Experimental Child Psychology, 239*, Article 105805.
Katsos, N., & Bishop, D. (2011). Pragmatic tolerance: Implications for the acquisition of informativeness and implicature. *Cognition, 120*(1), 67–81.
Klibanoff, R. S., & Waxman, S. R. (2000). Basic level object categories support the acquisition of novel adjectives: Evidence from preschool-aged children. *Child Development, 71*(3), 649–659.
Landau, B., & Gleitman, L. (1985). *Language and experience: Evidence from the blind child*. Cambridge, MA: Harvard University Press.
Lewis, M., & Frank, M. (2018). Still suspicious: The suspicious-coincidence effect revisited. *Psychological Science, 29*(12), 2039–2047.
Lewis, M. L., & Frank, M. C. (2016). Understanding the effect of social context on learning: A replication of Xu and Tenenbaum (2007b). *Journal of Experimental Psychology: General, 145*(9), 72–80.
Lidz, J., Gleitman, H., & Gleitman, L. (2003). Understanding how input matters: The footprint of universal grammar on verb learning. *Cognition, 87*(3), 151–178.
Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition, 108*(3), 732–739.
Lockridge, C., & Brennan, S. (2002). Addressees' needs influence speakers' early syntactic choices. *Psychonomic Bulletin and Review, 9*(3), 550–557.
Markman, E. (1989). *Categorization and naming in children: Problems in induction*. Cambridge, MA: MIT Press.
Markman, E. M. (1987). How children constrain the possible meanings of words. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 255–287). Cambridge, MA: Cambridge University Press.
Markman, E. M., Horton, M. S., & McLanahan, A. G. (1980). Classes and collections: Principles of organization in the learning of hierarchical relations. *Cognition, 8*, 227–241.
Mervis, C., & Crisafi, M. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development, 53*(1), 258–266.
Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
Murphy, G. L. (2016). Explaining the basic-level concept advantage in infants. or is it the superordinate-level advantage? *Psychology of Learning and Motivation, 64*, 57–92.
Nadig, A. S., & Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science, 13*, 329–336.
Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language, 17*, 357–374.
Nilsen, E. S., & Graham, S. A. (2009). The relations between children's communicative perspective-taking and executive functioning. *Cognitive Psychology, 58*(2), 220–249.
Papafragou, A., Cassidy, K., & Gleitman, L. (2007). When we think about thinking: The acquisition of belief verbs. *Cognition, 105*(1), 125–165.
Papafragou, A., Fairchild, S., Cohen, M., & Friedberg, C. (2017). Learning words from speakers with false beliefs. *Journal of Child Language, 44*, 905–923.
Papafragou, A., Friedberg, C., & Cohen, M. (2018). The role of speaker knowledge in children's pragmatic inferences. *Child Development, 89*, 1642–1656.
Quine, W. (1960). *Word and object*. Cambridge, MA: MIT Press.
Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.
Samuelson, L. K., Schutte, A. R., & Horst, J. S. (2009). The dynamic nature of knowledge: insights from a dynamic field model of children's novel noun generalization. *Cognition, 110*(3), 322–345.
Skordos, D., & Papafragou, A. (2016). Children's derivation of scalar implicature: alternatives and relevance. *Cognition, 153*, 6–18.
Sloutsky, V. (2010). From perceptual categories to concepts: What develops? *Cognitive Science, 34*, 1244–1286.
Sloutsky, V., & Fisher, A. (2004). When development and learning decrease memory: Evidence against category-based induction in children. *Psychological Science, 15*(8), 553–558.

Sloutsky, V., Kloos, H., & Fisher, A. (2007). When looks are everything: Appearance similarity versus kind information in early induction. *Psychological Science, 18*, 179–185.
Southgate, V., Chevallier, C., & Csibra, G. (2010). Seventeen-month-olds appeal to false beliefs to interpret others' referential communication. *Developmental Science, 13*(6), 907–912.
Spencer, J., Perone, S., Smith, L., & Samuelson, L. (2011). Learning words in space and time probing the mechanisms behind the suspicious-coincidence effect. *Psychological Science, 22*(8), 1049–1057.
Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
Srinivasan, M., Foushee, R., Bartnof, A., & Barner, D. (2019). Linguistic conventionality and the role of epistemic reasoning in children's mutual exclusivity inferences. *Cognition, 189*, 193–208.
Wang, F., & Trueswell, J. (2019). Spotting dalmatians: Children's ability to discover subordinate-level word meanings cross-situationally. *Cognitive Psychology, 114*, Article 101226.
Wang, F., & Trueswell, J. (2022). Being suspicious of suspicious coincidences: The case of learning subordinate word meanings. *Cognition, 224*, Article 105028.
Wardlow Lane, L., & Ferreira, V. S. (2008). Speaker-external versus speaker-internal forces on utterance form: Do cognitive demands override threats to referential success? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 1466–1481.
Waxman, S., Lynch, E., Casey, K., & Baer, L. (1997). Setters and Samoyeds: the emergence of subordinate level categories as a basis for inductive inference in preschool-age children. *Developmental Psychology, 33*(6), 1074–1090.
Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. *Cognitive Development, 5*, 123–150.
Waxman, S. R., & Gelman, R. (1986). Preschoolers' use of superordinate relations in classification and language. *Cognitive Development, 1*, 139–156.
Waxman, S. R., & Hatch, T. (1992). Beyond the basics: Preschool children label objects flexibly at multiple hierarchical levels. *Journal of Child Language, 19*(1), 153–166.
Waxman, S. R., & Klibanoff, R. S. (2000). The role of comparison in the extension of novel adjectives. *Developmental Psychology, 36*(5), 571–581.
Waxman, S. R., Lidz, J., Braun, I. E., & Lavin, T. (2009). Twenty-four-month-old infants' interpretations of novel verbs and nouns in dynamic scenes. *Cognitive Psychology, 59*(1), 67–95.
Waxman, S. R., Shipley, E. F., & Shepperson, B. (1991). Establishing new subcategories: The role of labels and existing knowledge. *Child Development, 62*, 127–138.
Wilson, E., Lawrence, R., & Katsos, N. (2022). The role of perspective-taking in children's quantity implicatures. *Language Learning and Development, 19*(2), 167–187.
Xu, F., & Tenenbaum, J. (2007a). Sensitivity to sampling in Bayesian word learning. *Developmental Science, 10*(3), 288–297.
Xu, F., & Tenenbaum, J. (2007b). Word learning as Bayesian inference. *Psychological Review, 114*(2), 245–272.
Zehr, J., & Schwarz, F. (2018). *PennController for Internet Based Experiments (IBEX)*.