Learning cues to category membership: Patterns in children's acquisition of hedges

Author(s): Marisa Casillas and Patrícia Amaral

Proceedings of the 37th Annual Meeting of the Berkeley Linguistics Society (2013), pp. 33-45

Editors: Chundra Cathcart, I-Hsuan Chen, Greg Finley, Shinae Kang, Clare S. Sandy, and Elise Stickles

Please contact BLS regarding any further use of this work. BLS retains copyright for both print and screen forms of the publication. BLS may be contacted via <http://linguistics.berkeley.edu/bls/>.

The Annual Proceedings of the Berkeley Linguistics Society is published online via [eLanguage](#), the Linguistic Society of America's digital publishing platform.

Learning cues to category membership: Patterns in children's acquisition of hedges

MARISA CASILLAS AND PATRÍCIA AMARAL

Stanford University and University of North Carolina at Chapel Hill

Introduction

When we think of children acquiring language, we often think of their acquisition of linguistic structure as separate from their acquisition of knowledge about the world. But it is clear that in the process of learning about language, children consult what they know about the world; and that in learning about the world, children use linguistic cues to discover how items are related to one another. This interaction between the acquisition of linguistic structure and the acquisition of category structure is especially clear in word learning.

0.1 Linguistic Hedges and Category Membership

In order to refer to items in the world, speakers must know about the relations between them: their category structure. An item's category membership is perceived as gradient rather than absolute. Members of semantic categories have different degrees of prototypicality, according to their properties (e.g. Rosch & Mervis 1975). For example, while native speakers might regard a robin as a good instance of the category BIRD, they are likely to see a chicken as a more peripheral member of that category. In language learning, it has been proposed by Heider (1971) that central members of a category are learned earlier than peripheral members.

Linguistic hedges provide the hearer with information about an item's position with respect to the structure of a category. For example, hedges like *par excellence* or *a typical X* require the highest degree of category membership, while other hedges (e.g. *kind of*, *sort of*, and *almost*) encode vagueness or fuzziness in membership (Lakoff 1973). The interpretations of *sort of* and *almost* both involve identifying central and peripheral members of a category, though their meanings differ. The use of *sort of* in (1) indicates that although a croissant can be considered a dessert, it is not the best instance of this category. The use of *almost* in (2) indicates

that the necklace being referred to is not a rainbow, but strongly resembles a typical rainbow, e.g. due to the number of colors it has, or the overall configuration of its strands.

- (1) A croissant is sort of a dessert.
- (2) That necklace is almost a rainbow.

Typically, the meaning of *almost* involves an ordering. To interpret a sentence with *almost*, the hearer must identify a set of ordered elements along a certain dimension, i.e. a scale (Hitzeman 1992; Amaral & Del Prete 2010). In (2), under the interpretation above, *almost* operates on a prototypicality scale, on which we rank possible arrangements of colored lines with respect to how well they resemble a “good” rainbow (the most prototypical one). Another interpretation for (2) is available, namely one that relies on a temporal scale: (2) could also be uttered as a description of an incomplete necklace that is not a (full) rainbow yet, but will eventually become one, once e.g. its blue and violet beads are added. The succession of steps ending in the fabrication of a full rainbow necklace is what underlies this use of *almost*.

The use of these expressions shows that speakers make a range of subtle distinctions within the structure of semantic categories. Learning about similarities and differences between items, and identifying properties associated with a certain category are important skills for the language learner in dealing with vagueness (a property of all natural languages). Semantic categories have fuzzy boundaries, and this is part of the semantic knowledge of adult native speakers. If children are sensitive to the meaning of linguistic hedges, their notions about category structure can be informed by the use of these expressions in the linguistic signal.

0.2 Hedges in Language Acquisition

Previous work on the acquisition of *almost* has shown that children are sensitive to the different scales required by its semantics. They use this knowledge to learn about distinctions in the semantics of different types of gradable adjectives (Syrett 2007). In a corpus study of Child Directed Speech to four children, *almost* systematically occurs as a modifier of expressions involving an ordering on some domain, often as a modifier of expressions denoting endpoints of temporal paths, like “becoming a man” as in (3):

- (3) A: Do men have to shave, and boys, too?
M: Boys when they’re almost men, not little boys like you. (Adam, file 38; Syrett 2007:153)

Almost also operates on a prototypicality scale, sometimes co-occurring with *like*:

(4) M: Tell me what color the bookcase is.

N: Red.

M: Almost. It's pink. (Nina, file 17; Syrett 2007:153)

(5) E: This is almost like tapioca. (Eve, file 7; Syrett 2007:153)

Other uses of *almost* as a hedge presuppose an awareness of category structure and the identification of properties associated with a certain category:

(6) A: Look this is almost a beehive it got little holes. (Abe, 3;10.14)

In a comprehension study on the acquisition of *almost* by 3- to 5-year-olds, Amaral (2010) shows that children understand the meaning of *almost* when they grasp the underlying scale required by the meaning of the adverb. For example, 3-year-old children understand an expression of the type 'almost X' about 60% of the time when *almost* operates on a numeric scale (e.g. *He has almost 5 blocks*) or on a directional spatial path (e.g. *The frog jumped almost to the lily pad*). A consistent, adult-like interpretation of the adverb as a modifier of different syntactic categories seems to be reached only at age five (Amaral 2010).

To our knowledge no study has focused on the acquisition of *almost* or other expressions as hedges, which encode information about degrees of category membership. Adults use hedges when offering children information about categories (e.g. "A moth is sort of a butterfly, but..."), and in that context, the hedged description introduces a new word-referent pair by building on existing knowledge. The child who understands the meaning of the hedge phrase *sort of* is invited to make the inference that the object in question shares some properties with butterflies but is not, strictly speaking, a butterfly.

Little is known about the acquisition of hedges and their role in word learning, and yet learning about differences between members of a category is important in mapping words to their referents. This study investigates whether children are sensitive to the use of the linguistic hedges *sort of* and *almost* and, if so, whether they associate these hedges to non-prototypical members of a category or incomplete objects/objects-to-be.

We ask the following questions:

1. Are children sensitive to the use of hedges by other speakers?
2. Do children associate hedges to objects that are incomplete or to non-prototypical category members?
3. Are children sensitive to differences in meaning among hedges?

Our results indicate that from age three children can identify the use of a hedge by another speaker, and by age five they can reliably use this information to distinguish between more and less prototypical members of a category. The structure of the paper is as follows: In section 1 we present our experiment. Section 2 describes the results and section 3 presents a discussion of our results and future work.

1 The Experiment

Our participants were 3-, 4-, and 5-year-old children, all native English speakers. We collected data from twelve participants in each age group for a total of 36 children.

1.1 The Task

Participants took part in a four-alternative forced choice task that lasted approximately 10–15 minutes and comprised a total of 12 trials (4 fillers and 8 critical trials). At the beginning of each trial, the participant was given a set of four semantically related pictures to look at. Two of the pictures were competitors for a category, one prototypical member and a non-prototypical member (e.g. “butterfly” and “moth” for the category BUTTERFLY), and the two other pictures were semantically related non-competitors (e.g. “bee” and “mantis”). After the children had a moment to look at each picture, they were told that a copy of one of the pictures would be hidden inside of a Mystery Box, and were asked to close or avert their eyes while the experimenter hid the picture inside. Then a puppet was allowed to peek inside the box and give the children a clue about which picture was hidden. Following the puppet’s clue, children were asked to point to which of the four pictures they thought was most likely to be hidden. Sessions were recorded on a camcorder for later review and coding.

1.2 Manipulations

1.2.1 Use of Hedge

Knowledge of hedges was tested through the clue that the puppet gave. Clues took the form of a linguistic frame, in which the noun phrase always occurred utterance-finally. In the 8 critical trials, which all featured competing pictures (e.g. “butterfly” and “moth”), the puppet used both hedged and unhedged frames (see Table 1): “It’s almost a ____,” “It’s sort of a ____,” “It’s a ____,” and “In here there’s a ____.” This last frame (“Long unhedged”) was added to match the syllable length of the hedged frames, while leaving the *unhedged* category label in utterance-final position. Clues were played through a small speaker inserted inside of the puppet’s head so that each participant heard the same acoustic signal.

Because we are interested in the children’s knowledge of hedge phrases themselves, and not in (related) cues to uncertainty such as pausing, filler usage, intonation, or lengthening, the recordings of the frames were made with a confident declarative tone and their acoustic properties were kept constant across all the frames. These other cues to uncertainty or hedging are most likely relevant to word learning in the same way that the hedge phrases are, but we do not look into them

Figure 1: Experimental setup: Participants were given a four-picture set to review. A copy of one picture was hidden by the experimenter in the "Mystery Box." Children were asked to guess which picture was hidden after hearing a verbal clue from the puppet, who had peeked into the box.

Table 1: Clue "frames" used by the puppet. Each frame appeared twice on critical trials. Filler trials only used the Default frame.

Unhedged frames	Hedged frames
<u>Default</u> : It's [a butterfly].	<u>Almost</u> : It's almost [a butterfly].
<u>Long</u> : In here there's [a butterfly].	<u>Sort of</u> : It's sort of [a butterfly].

here. Previous work on children's perceptions of speaker reliability has demonstrated that children are highly sensitive to these cues (Koenig & Harris 2005; Sabagh & Baldwin 2001), but little work has tried to control what factors are actually contributing to these perceptions.

Our prediction is that, if children understand the meaning of *sort of* and *almost*, then they will be more likely to choose the non-prototypical competitor when they hear a hedged frame in the clue. In the case that they hear an unhedged frame, they should be more likely to choose the prototypical competitor.

Further, we predict that there will be a developmental effect of the interpretation of hedges. Namely, that younger children will show a strong prototypicality bias, more often choosing the prototypical competitor in response to a hedged frame.

1.2.2 Prototypicality vs. Completeness

As we said above, *almost* may operate both on a prototypicality and on a temporal scale, while *sort of* only denotes fuzziness of membership related to prototypicality. When describing a moth that strongly resembles a butterfly, we may say *It's almost a butterfly*, relying on a scale of prototypicality. We might also indicate the relation of a moth to a butterfly by saying *It's sort of a butterfly*: although a moth is not a good exemplar of the category BUTTERFLY, it shares certain significant properties with a prototypical one (e.g. having wings). On the other hand, if we want to refer to the life cycle of a butterfly, in which there is a temporal path culminating in the transition to a winged form, a fully mature larva breaking out of its cocoon can be described as *almost a butterfly*. But this interpretation is not available to the expression *sort of a butterfly*.

Children may perform differently on these different dimensions of category membership, and so we decided to test them separately. In the following, we refer to these dimensions by the terms Prototypicality and Completeness.¹ In the Prototypicality condition, the two competitors differed with respect to degree of category centrality. For each category, a picture of a good exemplar and a peripheral exemplar were both provided (e.g. butterfly/moth, sock/Christmas stocking). In the Completeness condition, competitors were related on a temporal scale, one picture representing the fully grown or complete exemplar and the other an object-to-be or an incomplete version of the object (e.g. butterfly/cocoon, sock/partially knit sock). The possible competing pictures for the category BUTTERFLY are shown in Figure 2.

Each participant was assigned to one of two between-subject conditions: Prototypicality or Completeness. Participants were split evenly within age groups so that each condition had a total of 18 participants. If children have acquired the difference in meaning between the two hedges, we predict that they will show a benefit for their comprehension of *sort of* on the Prototypicality condition and a benefit for *almost* on the Completeness condition. While *sort of an X* describes a non-prototypical member of category X, *almost an X* may indicate either that the described item resembles a prototypical member of category X or that it is an incomplete X/an X-to-be. The fact that the latter instances are more frequent in corpus data leads us to expect that *almost* be more readily understood in the Completeness condition.

¹ This is a general term that we chose to refer to the temporal scale that can be involved in the meaning of *almost*.

Figure 2: The prototypical, non-prototypical, and incomplete stimuli for the category BUTTERFLY used as competitors in the picture sets. Each participant only saw either Prototypicality relations or Pcompleteness relations. The other seven categories were: TREE, FROG, ROOSTER, SANDWICH, CASTLE, SOCK, and DOLL.

1.2.3 Man-made and Natural Kinds

The critical item categories were grouped into man-made (SANDWICH, CASTLE, SOCK, and DOLL) and natural kinds (BUTTERFLY, TREE, FROG, and ROOSTER). The man-made and natural kinds distinction was included to avoid any possible bias associated to differing knowledge of, or attention to attributes across these two category types. For example, a child who is not aware that a tadpole will become a frog will not be well-equipped to respond to the clue *It's almost a frog*. That same child, however, may find the clue *It's almost a sandwich* easy to respond to, since she is familiar with the process of sandwich-making. Thus, in addition to making our categories as familiar and culturally appropriate as possible, the man-made vs. natural kinds distinction allows us to ensure that children's behavior wasn't limited to a single category type.

A second reason for this manipulation comes from previous studies on word learning. Andersen (1975) provides evidence for a developmental trend in the acquisition of vague category boundaries within the semantic field of artifacts. In her study, young children (around age 3) tend to rely on salient perceptual properties of objects to make decisions about category membership. Only later do children combine these physical properties with functional properties as criteria for drawing category boundaries. In another study also focusing on labels for artifacts, it was found that in naming contexts, categorization by function is more likely only in older preschool children, at least 4-year-olds (Kemler Nelson et al. 2000).

Given that functional properties only apply to man-made objects, the inclusion of both man-made and natural kinds would allow us to see if the integration of functional properties creates any bias in learning the category structure of man-

made items. If older preschool children take into account functional properties in their decisions about categorization, this may play a role in identifying prototypical and non-prototypical members of man-made categories.

Returning to the Completeness condition, on a temporal scale, a picture of a chick fulfills the description “It’s almost a rooster” better than a picture of a rooster. But, this is not the case with the description “It’s sort of a rooster,” in which the best option (rooster or chick) is less clear, and adults may even be more likely to choose the rooster, despite the use of a hedge. This distinction between the meanings of *almost* and *sort of* applies more weakly to the man-made items: a picture of a sock in the process of being knit is a better answer than a prototypical sock whether the clue is, “It’s almost a sock” or “It’s sort of a sock.” If children are picking up on the subtle differences in meaning between these hedges, it should be evident in their differential responses in the Completeness condition, and may even be stronger for natural kinds than man-made objects.

Since there were only 12 trials, randomization was constrained to keep the man-made and natural kind trials evenly spread throughout the session. The order of trials was shuffled for each participant so that fillers were dispersed evenly throughout the 12 trials, starting at the first trial. The placement of images in the 2×2 picture sets were randomized for each trial, with the caveat that the competitors always appeared diagonal from each other (not in the same row or column).

Figure 3: Response averages to hedged frames by age group. Age-collapsed response averages to the default (unhedged) frame is added on the right for comparison. Children in all three age-groups consistently chose the prototypical item in response to the default frame.

2 Results

Children and adults were at or near ceiling on the Default frame trials, almost always choosing the prototypical competitor in response to a clue like *It's a butterfly*. The effect of hedge was significant ($p < 0.05$), meaning that overall, children and adults were significantly more likely to choose the non-prototypical or incomplete competitor in response to a hedged frame. From Figure 3, it is clear that though this effect holds for all age groups of children, there is a strong developmental effect. Three-year-olds take the hedge phrase into account in their responses, but only about 25% of the time. Five-year-olds, on the other hand, are much more likely to treat the hedge phrase in an adult-like way (about 70% of the time). T-tests indicate a significant difference between 3- to 4-year-olds and 5-year-olds ($p < 0.05$), but no difference between 3- and 4-year-olds. This suggests that the 4–5 year age span is an important developmental time for the acquisition of the meaning of hedges.

A mixed-effects model was fit to analyze the effects of the separate manipulations: age, use of hedge, “type of frame”, Prototypicality vs. Completeness, and Man-made vs. Natural kinds. In addition to a significant effect of age (by year and age in months), there was a significant interaction between age and frame, and age and the Long unhedged frame. These interactions stem from two results. First, performance with *almost* was slightly better than with *sort of*, though the effect is not a significant one. Second, and primarily, children treated the Default unhedged frame and the Long unhedged frame differently.

Some children appear to use the length of the frame (e.g. in syllables) as a cue to modification, corresponding in this case to non-prototypicality of the category member. For them, longer-than-expected descriptions are less likely to apply to typical category members, falling in line with Gricean maxims of Manner and Quantity. This behavioral pattern was most common among 3-year-olds, who show a marginal difference ($p = .08$) between the Long and (short) Default unhedged frames.

While responding to each trial, children were encouraged to share any comments they had about the set of pictures or their reasoning for selecting one over the others. These data allow us to peek beyond their forced-choice responses, giving us a richer interpretation of their behavior. The video recording of each child's participation was reviewed by two paid, naive coders who transcribed the children's spontaneous commentary, any hesitation in their response (in comparison to the easier filler trials), and any coping strategies the children used in approaching the task.² In this paper, we will not address these coping strategies, but leave a more in-depth analysis of these behaviors for a later date.

² E.g. choosing two pictures that might **both** be valid answers. In these cases, children were reminded that they could only choose *one* picture.

Figure 4: Comparison of the two unhedged frames.

2.1 Spontaneous Commentary and Hesitations

From their forced-choice response data alone, we cannot tell whether children are more likely to choose non-prototypical category members in response to hedged frames because (1) the hedged frames are longer, and so the Gricean inference applies to all but the Default frame, or (2) the hedged frames contain a hedge phrase, or (3) both. Here the children's spontaneous commentary provides a great deal of insight about how they interpreted the use of hedges. Children's explanations and commentary often focused on the lexicalized hedge phrase itself or the category features it brought to light:

About the hedge:

1. He didn't say it was a sock, he said it was sort of a sock.
2. He told me it's sort of like a rooster.
3. Sort of a castle? Because it has the same top as a castle. It's almost a castle but it's just one big tower.

About category features:

1. It was something like a butterfly. This is like a butterfly, it has wings.
2. It's not big, it doesn't have a tail. The grownup [rooster] and the baby [chick].
3. Almost a frog... It's green like a frog, it has feet like a frog.

Although there were no significant differences between the two hedges in this task, in the Completeness condition children’s comments on the frame “It’s almost a ____.” were revealing. Children refer to temporal paths to explain their choice of the incomplete or object-to-be competitor:

About temporal paths:

1. Because this tadpole will grow into a frog it’s almost a frog.
2. Oh it’s almost a sock but not a sock yet; because it’s not a sock it’s knitted.
3. When butterflies are in cocoons they turn into this. Butterflies aren’t in cocoons.

Coder reliability for the presence of hesitation compared to filler trials was 75.5%. Taking only those trials which both coders marked as a “hesitated” response, we found that participants hesitated on 54% of hedged trials and 12% of unhedged trials. Over two-thirds of the hesitations on unhedged trials were in response to the Long unhedged frame, and the rest were dispersed evenly.

Table 2: Comparing results from the forced-choice responses and children’s hesitations, it becomes clear that 4-year-olds fall into an intermediate category in which they superficially look like 3-year-olds, but appear to incur processing costs when presenting with a hedge, like 5-year-olds.

Response measure	3-year-olds	4-year-olds	5-year-olds
“Prototypical” response to hedged frames	77.1%	68.8%	* 37.5%
Hesitations during response	36.2%	* 59.6%	64.0%

The forced-choice response data showed 5-year-olds diverging from 3- and 4-year-olds in their use of hedged phrases to select non-prototypical options. This suggests that four-year-olds have not yet advanced in their knowledge of hedges from their performance at three years of age. However, when we compare hesitation data, as in Table 2, we see that 4- and 5-year-olds *both* diverge from 3-year-olds, hesitating more often in response to the prompts.

The increase in hesitation with age indicates that the older children are more likely to take their time to process all of the information given in the frame, while younger children are more likely to jump ahead with their (prototypicality-biased) guess. Taking these data into account, then, it appears that 4-year-olds fit nicely into an intermediate category between 3- and 5-year-olds, in which they are more likely to be sensitive to hedge use, but may not have mastered its association with non-prototypicality yet.

3 Discussion

From age three children are just beginning to take linguistic hedges into account, and they tend to choose the prototypical object in response to both hedged and unhedged cues. Younger children are more likely to rely on length of description as a cue to category membership: shorter descriptions are interpreted as unmodified and

hence as referring to a central member of the category. Heider (1971) assumes that central members of a category are learned earlier and hypothesizes that children would make more errors on peripheral category members than adults. We believe that this may explain the higher percent of hesitations found overall with hedged frames. In this case, identifying the appropriate referent involves mapping the linguistic description to a non-central category member (that the child is likely to be less familiar with), incurring processing costs that result in greater response delay.

By age 5, children distinguish hedged from unhedged frames in an adult-like manner. This finding shows that 5-year-olds understand the contribution of this type of modification. Furthermore, this result supports previous findings on the acquisition of the meaning of *almost*: only 5-year-olds displayed an adult-like pattern in the interpretation of the adverb with all types of scales considered in Amaral (2010). This suggests that at age 5 children have generalized the semantic value of *almost* across different syntactic categories. In the present study, age 4 presents a middle-point for the developmental trend since children this age pair with 3-year-olds in the forced-choice task, but with 5-year-olds with respect to their pattern of hesitations.

Frames and frame-like units are relevant in acquisition (e.g. in learning morphological patterns and in forming syntactic categories, cf. among others Mintz 2003). However, little is known about the role that frames play in learning semantic categories, and in particular about the role that hedged frames play in learning about category structure.

In future studies, we intend to expand this research in two directions. First, we will analyze the adult baseline, in particular the adults' sensitivity to differences in meaning between the two hedges considered in this task. Second, in line with our interest in the role of hedges in word learning, we intend to investigate how children use their knowledge of hedges to make inferences about category membership of unfamiliar objects.

References

- Amaral, Patrícia. 2010. Almost means 'less than': preschoolers' comprehension of scalar adverbs. In *On-line Proceedings Supplement of BUCLD 34*, Boston.
- Amaral, Patrícia and Fabio Del Prete. 2010. Approximating the Limit: The Interaction Between *Quasi* 'Almost' and Some Temporal Connectives in Italian. *Linguistics and Philosophy* 33(2):51–115.
- Andersen, E. S. 1975. Cups and Glasses: Learning that Boundaries are Vague. *Journal of Child Language* 2(1):79–103.
- Heider, E R. 1971. On the internal structure of perceptual and semantic categories. Unpublished paper, Psychology Department, University of California, Berkeley.

Patterns in children's acquisition of hedges

- Hitzeman, J. 1992. The selectional properties and entailments of “Almost”. In & J. Marshall Denton C. P. Canakis, G. P. Chan, ed., *Papers from the 28th regional meeting of the CLS*, Chicago: CLS.
- Kemler Nelson, D.G., A. Frankenfield, C. Morris, and E. Blair. 2000. Young Children's Use of Functional Information to Recognize Artifacts: Three Factors that Matter. *Cognition* 77(2):133–168.
- Koenig, M A and P L Harris. 2005. Preschoolers mistrust ignorant and inaccurate speakers. *Child Development* 76:1261–1277.
- Lakoff, George. 1973. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic* 2:458–508.
- Mintz, Toben. 2003. Frequent frames as a cue for grammatical categories in child directed speech. *Cognition* 90(2):91–117.
- Rosch, Eleanor and Carolyn Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7:573–605.
- Sabbagh, M A and D A Baldwin. 2001. Learning words from knowledgeable versus ignorant speakers: Links between preschoolers theory of mind and semantic development. *Child Development* 72:1054–1070.
- Syrett, Kristen. 2007. Learning about the Structure of Scales: Adverbial Modification and the Acquisition of the Semantics of Gradable Adjectives. Ph.D. thesis, Northwestern University.

Marisa Casillas
Department of Linguistics
Margaret Jacks Hall, Stanford University
Stanford, CA 94305-2150

Patrícia Amaral
Department of Romance Languages & Literatures
320 Dey Hall, CB 3170
Chapel Hill, NC 27599

middyp@stanford.edu, pamaral@unc.edu