

Attention to Different Cues in Noun Learning: The Effect of English vs. Spanish Mass/Count Syntax

Eliana Colunga and Michael Gasser
Indiana University

1. Introduction

Children generalize nouns in ways that are consistent with the referent's ontological and/or grammatical kind. In other words, children generalize a new noun based on both the perceptual properties of the referent and the linguistic properties of the noun (Soja, Carey, & Spelke, 1991; Jones & Smith, 1998, Soja, 1992; Smith, 1995). Cross-linguistic studies have shown that systematic differences in the structures of different languages are reflected in children's novel noun generalizations (Imai & Gentner, 1997; Gathercole & Min, 1997, Yoshida & Smith 1999). One of the differences that has been studied is the ontological object/substance distinction as it relates to the syntactic count/mass distinction. In this paper we look at the effect of mass/count syntax and perceptual cues concerning solidity on English- and Spanish-speaking children's generalization of new nouns.

The task used to study this is the Novel Noun Extension Task. In this task, the child is shown an exemplar and the exemplar is labeled. The child is then asked what other things, matching the exemplar on different dimensions, can be called by the same name. Previous research has shown that children extend the name of a solid object to other objects of the same shape and the name of a nonsolid substance to other shapes made out of the same material. (Soja et al 1991).

This object-substance distinction is at its core about discrete versus continuous quantities. Conceptually, objects are discrete and bounded unitary wholes. Substances, in contrast, are continuous unbounded masses. Thus a cup is an object, and what is the cup is the whole thing with its particular form. A smashed cup is not a cup. Water, however is a substance, continuous and unbounded. Any bit of water, in drops or in puddles, is water. Perceptually, solid things have bounded permanent shapes; nonsolid things are continuous masses with transient shapes. Thus, there is a relation between different solidities and the ontological kinds of object and substance.

There is also a relation between count/mass syntax and the ontological object-substance distinction. In English, nouns are classified as either count or mass nouns. English count nouns such as *cup* and *study* obligatorily take the plural given multiple instances. Count nouns thus refer to entities conceptualized as discrete countable units -- objects. Mass nouns such as *water* and *justice* do not take the plural but instead take continuous quantifiers such as

some and *much*. Mass nouns thus refer to entities conceptualized as continuous masses -- substances.

English and Spanish differ in interesting ways in their count/mass syntax. In English nouns are classified as either count or mass. In this way, the language forces the speaker to view the referent as either an individuated object or a continuous substance. This is not the case in Spanish. In Spanish, nouns are not strictly classified as mass or count. In fact, in principle any noun can co-occur with both count and mass syntactic frames. Mass or count syntax is used according to how the speaker views the referent. Although *taza* (cup) typically appears in a count frame, many nouns such as *hielo* (ice), *esponja* (sponge) or *estudio* (study) frequently appear in either frame (Gathercole 1997). In other words, in English, the ontological status of a referent is given by the noun's grammatical category, while in Spanish, it is the ontological category of the referent that determines which syntactic frame — count or mass — is used.

This difference between English and Spanish count/mass syntax is exemplified in Figure 1. In English a wooden block could be referred to as a block if construed as an object or as some wood if construed as a substance. Furthermore, one cannot say *some block* or *a wood* because *block* is a count noun and *wood* is a mass noun. Thus, both the syntactic frame used and the noun used indicate whether the speaker is referring to an object or a substance. In English all *blocks* are objects and the word *block* is always a count noun; all *wood* is a substance and the word *wood* is always a mass noun.




Figure 1. A block of wood can be construed as an object (block) or as a substance (wood).

In Spanish, the same wooden block could be called *un bloque* a block or *algo de madera* some wood, but it could also be called *una madera* a wood. While *bloque* does refer to the object construal and *madera* does refer to the substance construal, *madera* can be used in both count and mass syntactic frames. Thus, the noun *madera* cannot be said to be either a count noun or a mass noun, as it can occur in both frames. Note that although English has some nouns that work like *madera* in Spanish, (e.g. a muffin would imply a whole muffin-shaped object, while some muffin would refer to a piece of muffin-stuff), in Spanish, **in principle**, all nouns work this way. Does this difference affect children's expectations about the meaning of nouns?

These differences suggest two contrary hypotheses about the influence of count versus mass syntax on children's interpretation of novel nouns.

Hypothesis one: If what matters is *strength of correlation*, Spanish will have a greater effect of solidity and English will have a greater effect of Syntax because in English the syntactic cue is correlated more tightly with ontology and with individual categories. When English-speaking children hear a novel noun used with count syntax (this is a dax), they can be sure that it is the object that is the referent, and thus that its shape is relevant. Similarly, when English-speaking children hear a novel noun used with mass syntax (this is some dax) they can expect the referent to be the substance and thus the material to be relevant. In contrast, when a Spanish-speaking children encounter a novel noun in a count frame (este es un dax) the noun could mean either the shape of the object or to the material it is made out of. In the same way, when Spanish-speaking children encounters a novel noun in a mass frame (esto es (un poco de) dax) the noun could refer to either the object or the substance.

Hypothesis two: If what matters is *informativeness*, Syntax will have a greater effect for Spanish speakers because the syntactic cue is more informative, carrying information independent of the object's category and solidity. When English-speaking children hear this is a cup or this is some wood , they can ignore the syntax, since the noun itself tells which dimension is relevant. Thus, English-speaking children could learn to rely less on syntax, because it does not add much information to what the noun already carries. In contrast, when Spanish-speaking children hears una madera the use of count syntax indicates that the speaker construes the referent as an object. In the same way, when Spanish-speaking children hear a noun in a mass frame they can infer the speaker construes the referent as a substance.

Accordingly, in our first experiment we manipulated both count/mass syntax and solidity perceptual cues of ontological kind. We showed children novel objects and named them with novel nouns in either count or mass syntax. We included objects of different solidities. All the objects presented to the children were solid, but one set consisted of rigid things (like wood and hardened clay) and the other set consisted of flexible things (like sponge and acrylic stuffing). We chose to use a flexible set because it presents ambiguous solidity cues for ontological kind --- solid things tend to be objects and nonsolids tend to be substances, but there is no strong association of flexible things to either objects or substances. Therefore, we expected the effect of syntax to be more noticeable in the flexible set than in the rigid set, because the flexible set presented more ambiguous perceptual cues. In the same vein, we included two kinds of material matches ---one with a particular, definite shape (what we call a whole) and one with an irregular, accidental-looking shape (what we call a piece). Again, we thought perhaps the piece material match --with its ambiguous perceptual information as to whether it should be construed as an object or a substance-- could be more dramatically affected by count/mass syntax than the whole material match, which more strongly cues object.

2. Experiment 1

The goal of this experiment was to see if children's linguistic history influences how children weigh syntactic and perceptual cues to determine the ontological kind of the referent and thus extend the name of an exemplar. We use a yes/no Novel Noun Extension task. We manipulate three factors: Solidity (Rigid vs. Flexible) and Syntax (Count vs. Mass) and Language (Spanish vs. English). Solidity was manipulated within subject, Syntax and Language were manipulated between subjects

2.1. Method

Participants. Thirty-two monolingual English-speaking 2-3 year-olds (range 2.13, 3.84 M=3.05) were recruited in Bloomington, IN; thirty-two monolingual Spanish-speaking 2-3-year-olds (range 2.24, 3.31 M=2.91) were recruited in Monterrey, NL, Mexico.

Stimuli. The training stimuli consisted of familiar objects: For the Count Syntax condition, the exemplar was a spoon and the test objects were an identical spoon, a lemon, a pencil and a pair of glasses; for the Mass Syntax condition the exemplar was a chocolate bar and the test objects were an identical chocolate bar, a lemon, a biscuit and a slice of bread.

The experimental stimuli consisted of the two sets shown in Figure 2. The exemplar of the Rigid set was a barbell shape made out of red clay. The exemplar of the Flexible set was a square-U shape made out of blue sponge. Each set had eight test items. The test items were (1) an identical match, (2) a shape match, (3) a whole material match and (4) a piece material match, (5) a whole color+material match and (6) a piece color+material match, (7) a color match, and (8) an object did not match the exemplar in shape, material nor color. The novel nouns used were Dugo and Zup in English and Dugo and Mepa in Spanish.



















	Exemplar	ID	Shape	Material Whole	Material Piece	Col+Mat Whole	Col+Mat Piece	Color	None
FLEXIBLE	blue sponge 	blue sponge 	white stuff 	green sponge 	pink sponge 	blue sponge 	blue sponge 	blue feather 	orange scrub 
RIGID	red clay 	red clay 	gray granitex 	yellow clay 	blue clay 	red clay 	red clay 	read leather 	blonde wood 

Figure 2. Stimuli for Experiment 1.

Design. Children were randomly assigned to either the Mass Syntax condition or the Count Syntax condition. Each child saw the Deformable set

and the Rigid set. The order in which these sets were presented was counterbalanced across children in both conditions.

Procedure. The experiment began with a series of training trials. A stuffed bear was introduced and the exemplar was introduced with the appropriate syntax. In the Count condition the experimenter showed the child the spoon and said the bear wants more spoons. The child was then shown one of the familiar training items and asked is this a spoon? . Analogously, in the Mass condition the child was told the bear wants more chocolate and then asked about the training items is this some chocolate? . The complete scripts in the two languages are shown in Table 1. There were a total of 8 training trials; during these trials the children were given feedback on their mistakes.

Table 1. Count and Mass Condition instructions in English and Spanish

English	
Count	Mass
This is Sam. Sam has a spoon. Sam likes spoons. Sam wants more spoons. Let's find a few more spoons for Sam! I may have another spoon in here... Is this a spoon?	This is Sam. Sam has some chocolate. Sam likes chocolate. Sam would like to have a bit more chocolate. Let's find some more chocolate for Sam! I may have a little more chocolate in here... Is this some chocolate?
Spanish	
Count	Mass
Esta es Ana. Ana tiene una cuchara. A Ana le gustan las cucharas. Ana quiere m s cucharas. Vamos a encontrar m s cucharas para Ana! Tal vez tenga otra cuchara aqu Esto es una cuchara?	Esta es Ana. Ana tiene un poco de chocolate. A Ana le gusta su chocolate. A Ana le gustar a tener mucho chocolate. Vamos a encontrar m s chocolate para Ana! Tal vez tenga algo de chocolate por aqu Esto es chocolate?

The experimental trials followed the same procedure. Children were introduced to a new stuffed animal, the exemplar was presented and the children were asked for help in finding more. Then the test items were presented twice in a random order and the children were asked for each of them is this a dugo? or is this some dugo? depending on the syntax condition.

2.2. Results and Discussion

The results, coded as proportion of yes responses, for English- and Spanish-speaking children are shown in Figures 3 and 4 respectively. Since it

has already been established that children extend novel names for solid things to other things of the same shape, we look at the results for the identical and shape matches separately. Children of both language groups extended the name of the exemplar to an identical test object (93% for the Flexible set, 97% for the Rigid set). Children of both language groups extended the name of the exemplar to an object matching in shape for the Rigid set (91% for English speakers, 95% for Spanish speakers), but did so at chance for the Flexible set (52% for English speakers, 56% for Spanish speakers).

A Solidity (rigid vs. flexible) X Syntax (count vs. mass) X TestObject (whole vs. piece vs. none) ANOVA was conducted for each language group. The ANOVA for the English group revealed a significant main effect of TestObject ($F = 19.814, p < .001$) and a significant Solidity by TestObject interaction ($F = 24.538, p < .001$). The ANOVA revealed no significant main effects or interactions involving Syntax, therefore the results depicted in Figure 4 are collapsed across this factor. As can be seen in the figure, English-speaking children are more likely to generalize by material when the exemplar is flexible than when the exemplar is rigid.

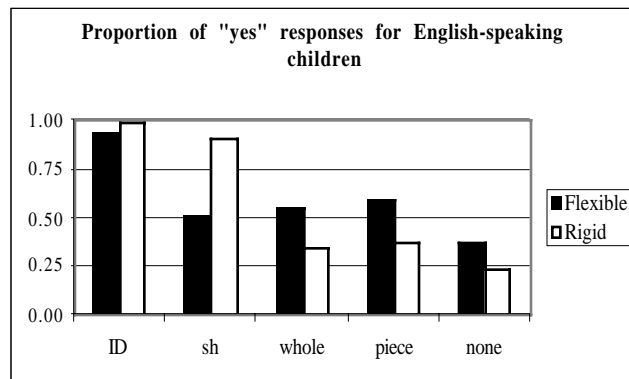


Figure 3. Results for English-speaking children

The ANOVA for the Spanish group revealed significant main effects of TestObject ($F = 59.582, p < .001$) and Syntax ($F = 6.49, p < .02$). These effects were subsumed by the significant interactions between Solidity and TestObject ($F = 21.61, p < .001$) and Syntax and TestObject ($F = 3.5, p < .02$). As can be seen in Figures 5a and 5b, Spanish-speaking children, like English-speaking children, are more likely to generalize by material when the exemplar is flexible than when the exemplar is rigid. In addition, Spanish-speaking children are more likely to generalize a novel noun to a material match when the noun is presented in a mass syntactic frame than when it is presented in a count syntactic frame.

These results show that speakers of both languages attend to solidity properties of the referent, generalizing by shape more in the rigid set than in the

flexible set, and by material and color more in the flexible than in the rigid set. However, the use of the different syntactic frames only has a significant effect in the novel noun generalizations of Spanish-speaking children, and not in the generalizations of English-speaking children. This suggests that Spanish speakers attend more to mass/count syntactic cues than English speakers.

These results are in line with hypothesis 2 —that English-speaking children ignore the redundant information offered by count/mass syntax while Spanish-speaking children learn to attend to count/mass syntax to help in determining nature of the referent. The next question is, what kind of learning mechanism would show this effect? In Experiment 2 we use a simple neural network to simulate this effect.

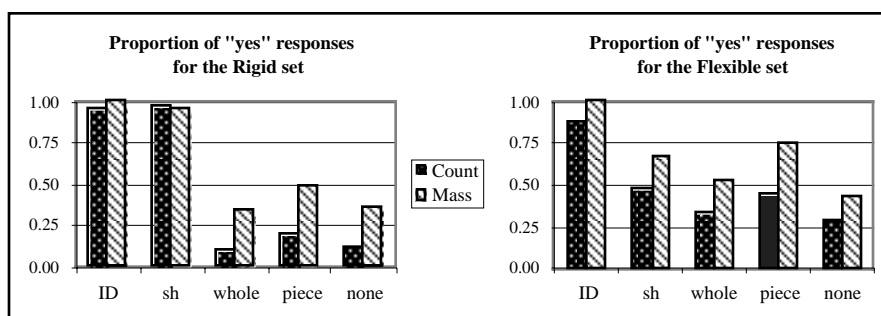


Figure 4. Results for Spanish-speaking children

3. Experiment 2

Next we describe how a simple neural network trained on patterns embodying the gross distinction between Spanish and English nouns predicts this difference in subjects' attention to syntax.

3.1. Method

Each network was designed to model a comprehension task. The input to the network was a word and a syntactic pattern, the output a representation of the referent in terms of two dimensions, which we call Shape and Material. Identical networks were trained on different sets of patterns, which we call English and Spanish, differing in terms of whether the syntactic input agreed with the word input on which output dimension was relevant.

The network architecture is shown in Figure 5. The connectivity is feed-forward. The input layers consist of a Word layer, with a single unit representing each of 12 words, and a Syntax layer, with a single unit representing each of two syntactic patterns (Count and Mass). There is a Hidden layer of 32 units. The output layers consist of a Material and a Shape layer, each containing eight units.




Figure 5. Architecture of the network in Experiment 2.

Training was supervised, using the back-propagation learning algorithm. For each training event, a word and a syntactic pattern were presented to the input layers, and activation was passed through the network, yielding a pattern of activation on the output layers. This was compared to the target output pattern associated with the input word and syntactic pattern, the error was propagated back through the network, and the connection weights were adjusted accordingly.

Each network was trained on either a set of English or a set of Spanish patterns. For the English patterns, each of the twelve words always appeared with the same syntactic pattern, half with the count and half with the mass pattern. Each count word had an associated Shape pattern, a randomly generated pattern consisting of 3 bits on and 5 off. Each mass word had an associated Material pattern, generated in the same fashion. For each English pattern with a count word, the associated target consisted of the Shape pattern assigned to that word and a Material pattern which was generated randomly on each training trial. Similarly, for each English pattern with a mass word, the associated target was made up of the Material pattern assigned to that word and a randomly generated Shape pattern. Thus for English the input word completely specified the output for one of the output layers, but not the other, and the syntactic input redundantly specified which of the output layers was associated with the input word.

For the Spanish patterns, each input word could occur with either of the two syntactic patterns, and each word was assigned both a Shape and a Material pattern. As for the English patterns, each target consisted of a pattern in one layer that was associated with the input word and a randomly generated pattern in the other. But in the case of the Spanish patterns, which output pattern was the one associated with the input word was specified solely by the input syntactic pattern. The count pattern in the input indicated that it was the Shape pattern associated with the word that was to be turned on; the mass pattern in the input indicated that it was the Material pattern that was to be turned on. Example training patterns for two English words and one Spanish word are shown in Figure 6.




Figure 6. Examples of English and Spanish patterns.

Note that since half of the target outputs are always random, the networks were effectively being presented with an impossible task. Yet they could reduce their initial error first, by learning the output patterns (Shape and/or Material) that were associated with particular input words, and second, by learning to produce intermediate, rather than more distinct, values on the output layer with the random target. In order to achieve the latter, the networks would have to learn which output layer was assigned the random target and which the pattern associated with the input word. For the English patterns this information was specified redundantly by the word and syntax inputs: for count words and count syntax, it was the Shape layer that had the non-random target; for mass words and mass syntax, the Material layer. For the Spanish patterns the information was available only in the syntactic input: for count syntax, it was the Shape layer that had the non-random target, regardless of the input word; for mass syntax, the Material layer, again regardless of the word. Apparently because of the redundancy in the English task, it was considerably easier for the networks than the Spanish task.

3.2. Results and Discussion

We trained ten networks, each for 150 repetitions of the 24 training patterns (two for each input word), five on the English and five on the Spanish patterns. With the same length of training, the error in the Spanish networks was still much higher than that in the English networks. We were interested in the extent to which the trained networks were relying on the syntactic patterns rather than in their performance on particular words from the training set, so we tested the networks by presenting them with 12 novel words. We presented each novel word to the network twice, once with the count syntactic pattern, once with the mass, and recorded the activations of the hidden units in each case. We reasoned that if a network is relying relatively heavily on syntax, the two hidden layer patterns for a given input word should differ by a relatively large amount. Thus our dependent measure was the Euclidean distance between the hidden layer activation patterns for a particular novel word input when the count vs. when the mass syntax unit was turned on.

For the English networks, the mean difference between the hidden layer patterns given count and mass input was 0.865; for Spanish the mean difference was 1.007 ($p < .002$). While both sets of training patterns led the networks to make use of the syntactic inputs, the Spanish training patterns resulted in greater reliance on syntax. This agrees with the results of Experiment 1.

4. General Discussion

The results of Experiment 1 suggest that what matters to whether a cue gets attended or not is its informativeness. In Experiment 2 we find a statistical learner for which this is also true. What can we learn from these networks? What do these simulations tell us about the learning mechanism in children? The models in Experiment 2 embody particular assumptions about the learning task and how it is realized in the networks. We believe these assumptions can shed light on how the child approaches word learning.

First, the task for the networks is a purely associative one: lexical and syntactic patterns are associated with semantic patterns through a hidden layers which mediates the association. But beyond this, there are two ways to view the network's (and presumably also the child's) task. It could be purely correlational: patterns, consisting of features of linguistic form as well as features of referents in the environment, are presented and the correlations between these features are learned. Such a task is unsupervised; there is no target provided by the environment, only a pattern. On the other hand, the task could have a particular direction to it: given one set of features, another set is output, and the environment provides a target to be compared to the output. Language comprehension and production are such tasks. In Experiment 2, we took the network's task to be comprehension: the input is linguistic form, and the output is the features of the referent. An explicit target, consisting of a set of Shape and Material features, is provided for every input. Viewed this way, it is not just the correlations among the various features that matters; it is the informativeness of particular features in inferring particular other features, in this case, the informativeness of the lexical and syntactic features in inferring the semantic features. It is ultimately this directionality that leads the networks trained on the Spanish patterns to rely more on the syntactic patterns; as laid out in Hypothesis 2, they are more informative.

In terms of children acquiring language this would mean that the learning of words is not just the learning of correlations between features. Rather it is driven by a particular task, the comprehension of linguistic expressions, and the directionality in the task has implications for what is learned about the language and for how familiar language is used in the learning of unfamiliar language.

5. Conclusions

Upon encountering a new word, children have a variety of cues they can attend to in order to figure out the meaning of the word. In our study we have shown that the relative weight children give to the different cues is, at least in part, a result of the structure of the language they are learning. In particular, cues

that are more informative, that have more predictive power, are going to be preferentially attended. We have also shown that a simple statistical learner, trained on a comprehension task, will learn in this way. This kind of contingent learning may be critical for solving the difficult problem of learning a language: with each word learn the learner becomes better word learner.

References

- Gathercole, V. C. M. (1997). The linguistic mass/count distinction as an indicator of the referent s categorization in monolingual and bilingual children. *Child Development, 22*, 395-415.
- Gathercole, V. C. M., & Min, H. (1997). Word meaning biases or language-specific effects? Evidence from English, Spanish and Korean. *First Language, 17*, 31-56.
- Imai, M., & Gentner, D. (1997). A cross-linguistic study of early word meaning: universal ontology and linguistic influence. *Cognition, 62*, 169-200
- Jones, S. S., & Smith, L.B. (1998). How children name objects with shoes. *Cognitive Development, 3*, 323-334.
- Smith, L.B. (1995). Self-organizing processes in learning to learn words: Development is not induction. In C. A. Nelson (Ed.), *Basic and applied perspectives on learning, cognition, and development. The Minnesota symposia on child psychology* (Vol. 28, pp.1-32). Mahwah, NJ: Lawrence Erlbaum Associates.
- Soja, N. N. (1992). Inferences about the meanings of nouns: The relationship between perception and syntax. *Cognitive Development, 7*, 29-45.
- Soja, N. N., Carey, S., & Spelke, E. S. (1991). Ontological categories guide young children s inductions of word meaning: Object terms and substance terms. *Cognition, 38*, 179-211.