Learning to Learn Words: A Cross-Linguistic Study of the Shape and Material Biases

Eliana Colunga, Linda B. Smith
Indiana University

1. Introduction

Young word learners seem to know a lot about how nouns map to different kinds of categories. They generalize names for objects they know to new objects in a consistent and seemingly sensible manner according to the kind of object being named. One task that has demonstrated this is the novel noun generalization task. In this task, children are shown an exemplar, it is named, and then the children are asked what other things are called by the same name. Young children extend the name of a solid object to other objects that match the exemplar in shape irrespective of their color, size or material. However, when generalizing the name of a non-solid substance, children consider material to be more important. In brief, young children seem to know that objects and substances are fundamentally different and classified in different ways. This paper is concerned with where this knowledge comes from.

There are several sources of information that correlate with the distinction between objects and substances that could be the basis of the distinction in children. One source of information is perceptual. Solid rigidly shaped things have more angular and more complex shapes than nonsolid things, which cannot hold their shape and thus relax into rounded and simple forms. A second source of information, at least in some languages, is syntax. In English, solid rigidly shaped things are conceptualized as individuals and referred to with count nouns whereas nonsolid substances are conceptualized as continuous masses and referred to with mass nouns. Indeed, Quine (1969) suggested that children learn the ontological distinction between object and substances from the count/mass distinction. A third source of information is the lexicon itself. If children learn many names for solid things that are well organized by shape and many names for nonsolid things that are well-organized by material, then they might **learn** to attend to shape in the context of solid things and to material in the context of nonsolid things. In a content analysis of the first 300 nouns typically learned by children learning English, Samuelson and Smith (1999) found just such a regularity in the early lexicon. Smith (1998) suggested that these regularities among perceptual cues, syntax, and lexical categories create the noun learning biases observed in the novel noun generalization task. The idea is that children's noun generalizations in those experimental tasks (and in the world) are themselves generalizations across the nouns already learned by the child.

One problem with the learning proposals offered by Quine and, also possibly, by Smith is that the shape bias for objects and the material bias for substances may be universal, characteristic of children learning very different languages. Here we consider English and Japanese. Japanese presents an interesting comparison to English because Japanese lacks count/mass syntax and with respect to individuation treats objects and substances the same way. In one important study, Imai and Gentner (1997) compared English-speaking and Japanese-speaking children's generalizations of novel names for complexly shaped solid things, simply shaped solid things, and simply shaped substances. Japanese as well as English speaking children generalized novel names for complexly shaped things by shape but generalized more by material when then named entity was nonsolid. Imai and Gentner suggested from these results that the ontological distinction between objects and substances was universal and did not depend on language learning.

However, Imai and Gentner also observed two differences between children the name generalizations of the two groups of children. First, they differed in how they generalized names for solid but simply shaped things. English-speaking children treated these things like objects and extended their names by shape. In contrast, Japanese-speaking children treated these things like substances and generalized their names by material. Second, Japanese children generalized names for substances more systematically than did English speaking children whose generalizations did not differ from chance. Put another way, the object-substance distinction was stronger in Japanese children, in children learning a language without the presumably supporting mass-count syntax.

We specifically seek to understand these cross-language differences for what they may tell us about how perceptual, syntactic, and lexical information may interact to form biases in a learner. In Experiment 1, we examined the kinds of nouns children know early in the two languages. Do children in both languages learn many names for solid things that name objects in categories well organized by shape and do they learn many names for nonsolid things in categories well organized by material? In Experiments 2 and 3, we ask whether the regularities among early learned nouns in the two languages are computationally sufficient to create the noun learning biases observed in experiments and the differences between children learning English and Japanese.

2. Experiment 1 – Structure of the lexicon

The goal of Experiment 1 is to determine what kinds of nouns children know early and what are the statistical regularities among the early categories. To do this we asked Japanese and English native speakers to make judgments about the structure of the categories corresponding to early nouns.

Method. As our list of early nouns we used common nouns on the Macarthur Communicative Development Inventory and the Japanese Communicative Development Inventory. The CDI for each language is a parental checklist used to measure children's productive vocabulary. These checklists were developed from independent and extensive study of the common words and phrases known by children learning the two languages (Fenson et al, 1993; Ogura and Watamaki, 1997). The words in these lists are known by 50% of large samples of children at 30 months of age. Thus, the lists are good proxies for the kinds of words commonly known by young English-speaking children and by young Japanese-speaking children. We specifically selected for study all object and substance terms on the two lists, excluding names for people, animals, and abstract objects (e.g., wind). The remaining items consisted of common objects and food terms. There were 150 such nouns on the English list and 167 such nouns on the Japanese list.

What we want to know are the kinds of things to which these nouns refer --- solid things organized by shape or nonsolid things organized by material. We know of no precedent for characterizing a large number of categories other than Rosch's (1973) seminal work on family resemblance structure. Accordingly, we followed her methodological approach and asked adults to make judgments about the perceptual properties characteristic of the instances of each early noun (see also Samuelson and Smith, 1999). Native speakers of English were asked to judge the English and native speakers of Japanese were asked to judge the Japanese nouns.

Solidity judgments were made by two native speakers of each language. They answered the following three questions

1. Do items in the category change shape when pressed?
2. Do they return to their original shape after being pressed?
3. Do they take the shape of their container?

Each noun was then classified as referring to **Solid** things**,** if all three questions were answered "no" or as referring to **Non-solid** things, if all three questions were answered "yes". Nouns with any other pattern of answers, which referred to things that were neither clearly solid nor clearly non solid (like playdough and shirt) were categorized as ambiguous.

The judgments of within-category similarity were made by a larger sample of native speakers of the two languages, because of the greater variability of the judgments. 13 English speakers (as part of Samuelson and Smith, 1999) and 10 Japanese speakers judged each of the nouns in their corresponding list. For each noun, they answered:

1. Are the items in the category similar in shape?
2. Are they similar in color?
3. Are they similar in material?

We then classified a noun as naming by **Shape** if at least 70% of the subjects agreed that the things named by that noun were judged to be similar in shape. We classified a noun as naming by **Material**, if subjects agreed the things named by that word were similar in material (and/or color)**.** Note that each word could either be classified as naming by Shape, Material, Both or Neither.

Results. The results, as Venn diagrams, are shown in Figure 1a for English and Figure 1b for Japanese. These Venn diagrams represent the number of nouns of each kind by the area enclosed in a space. Thus, the larger outline area represents all the nouns that were judged in the language ---including those that were ambiguous with respect to solidity or within-category similarity. The area of the black horizontally striped rectangle represents the number of noun that were judged to refer to solid thing by shape, and so forth. In both languages about half the nouns refer to solid objects (42% in English, 48% in Japanese) and there are very few (24 in English, 21 in Japanese) nouns for non-solids in either list of early vocabulary items. Both languages also had a similar proportion of ambiguously solid items (58% in English, 60% in Japanese). In both languages more nouns were judged to refer to things similar in shape (38% in English, 49% in Japanese) than to things similar in material and/or color (31% in English, 20% in Japanese).

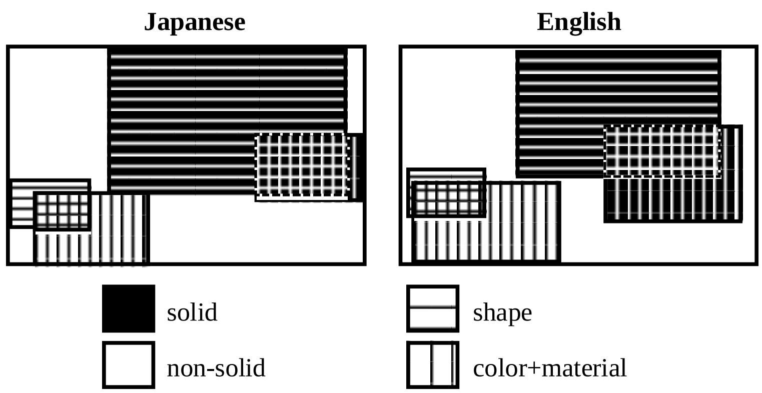

Figure 1. Results from adult judgments of solidity and within-category similarity of early vocabulary.

Solidity and category organization interact in the way we predicted. The correlation was very strong for solid and shape, with most of the words that were classified as solid also judged to refer to things that were similar in shape (79% in English, 93% in Japanese) and most words that were classified as referring to things similar in shape also classified as solid (88% in English, 90% in Japanese). The correlation was weaker for non-solid and material. While words that were classified as non-solid were judged to refer to things that were similar in material (96% in English, 81% in Japanese), the correlation didn't hold in the opposite direction (49% in English, 52% in Japanese).

For the majority of early nouns, solidity and with-in category similarity agree in a way consistent with the shape and material biases. However, there are also many exceptions to this "rule". This leads us to the next questions: Are the regularities in these early vocabularies enough – at least computationally – to create the learning biases observed in novel noun learning tasks?

3. Experiment 2 – Network simulations

The goal of Experiment 2 is to determine how potent the regularities found in Experiment 1 are. In other words, are they sufficient to create the similarities and differences between English- and Japanese-speaking children in the novel noun generalization task? If a simple statistical learner, given the regularities in the lexicon is capable of developing the same learning biases children show, this would support the idea that the biases come from the regularities in the lexicon and not the other way around. Thus in Experiment 2, we train a simple connectionist network on the regularities found in Experiment 1 and compare the network's performance to data on 2 year-olds from Imai and Gentner (1997).

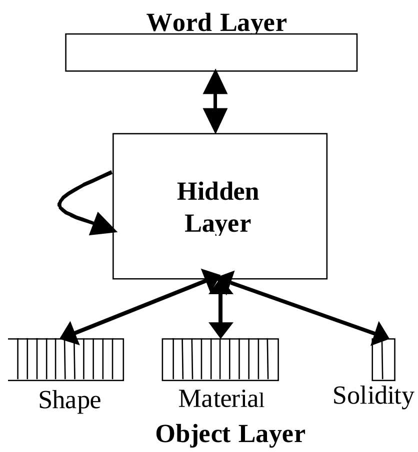

Figure 2. Architecture of the network used for the simulations in Experiment 2.

Method. We used a Hopfield network, which is a simple recurrent network. The networks were trained using Contrastive Hebbian Learning, an algorithm which adjusts weights on the basis of the correlations between unit activations. Figure 2 shows the architecture of the network. The networks have a Word Layer, in which each unit corresponds to one word in the vocabulary. Individual objects are represented on what we call the Dimension layer. Activation patterns on this layer represent the shape and material of each individual object or substance presented to the network. More specifically, the shape and material of an object (say the roundness of a particular ball and its yellow rubbery material) are represented by an activation pattern along the whole layer, in a distributed fashion. In the Solidity layer one unit stands for Solid and another for Non-Solid. Finally, there is a hidden layer that is connected to all the other layers and recurrently with itself. Note that the Word Layer and the Dimension and Solidity layers are only connected through the Hidden layer, there are no direct connections among them.

First, we trained networks on the English or Japanese nouns studied in Experiment 1. This mimics the vocabulary learning that a child brings into a novel noun generalization experiment. The statistical regularities in the early vocabularies that we observed in Experiment 1 were built into in the network's training set in the following way. First, for each word that the network was to be taught, a pattern was generated to represent its value along the *relevant* dimension ---the dimension that the adults in Experiment 1 said characterized the similarities of objects named by the noun. Second, at each presentation of the word, the value along the *irrelevant* dimension was varied randomly. For example, the word "ball" was judged by the adults in Experiment 1 to refer to things that were similar in shape; thus, a particular pattern of activation was randomly chosen and then assigned to represent ball-shape. All balls presented to the network were defined as having this shape, although each ball presented to the network also consisted of a unique and randomly generated pattern defining the material and color. So, whenever the network got the unit representing the word "ball" on, it also got the pattern representing ball-shape along the Shape dimension and a different pattern along the Material dimension.

Solid objects were assumed to have a bigger range of values along the shape dimension. This assumption is in line with the fact that solid things can hold more varied and complex shapes than non-solid things.

Second, we tested the networks in an analog of the novel noun generalization task used with children. Our approach is based on our conceptualization of the novel noun generalization task. In that task, the child sees and exemplar and hears its name. If, for example, the child attends exclusively to the shape of the named exemplar, then a test object that matches the exemplar in shape (although different from the exemplar in material) should be perceived as highly similar to the exemplar. Thus, we asked if the network's internal representations of a named exemplar and a test object were similar. More specifically, we asked if the patterns of activations on the hidden layer for the named exemplar and the test object were similar.

The novel noun generalization task used with children typically is a forced choice task in which the child must choose between an object matching the named exemplar in shape and one matching in material. Accordingly, on each simulated test trial, we measured the similarity of the internal patterns of representation for two test objects –one matching the exemplar in shape and one matching the exemplar in material.

More specifically, on each test trial, a novel exemplar object was generated by randomly creating an activation pattern along the shape and material dimensions. Then a novel shape-matching test object was generated by combining the exemplar's shape pattern with a novel randomly generated material pattern. A similarity measure of the exemplar and the shape match was computed using the Euclidean distance between the activation patterns in the Hidden Layer after the exemplar and its shape match were presented.

Similarly, a novel material-matching test object was generated by combining the exemplar's material pattern with a new randomly generated shape pattern and the similarity between exemplar and material match was computed. Finally, we used these similarity measures between the emergent patterns of activation on the hidden layer to calculate the probability of choosing the shape and the material match using Luce's Forced Choice Rule.

In this way, we trained 10 networks (with 10 different randomly generated initial connection weights) with the object and substance terms young English children know. During training, we presented multiple instances of each trained noun until the network stably produced the right noun when presented an instance of each kind. We then tested each of these English networks in the novel noun generalization task ---with 30 novel exemplars and 30 novel names. These 30 test trials were divided evenly into three kinds: exemplars defined by patterns activation representing solid and complexly shaped things, solid and simply shaped things, and nonsolid and simply shaped things. Similarly, we trained 10 networks with all the words in the Japanese corpus, and at the end of this training, tested those ten Japanese networks with the same 30 novel noun generalization trials. If the statistical regularities in the two vocabularies are sufficient to create learning biases and the cross-language differences, then the performances of these networks should look like the performances of the children in Imai and Gentner's study.

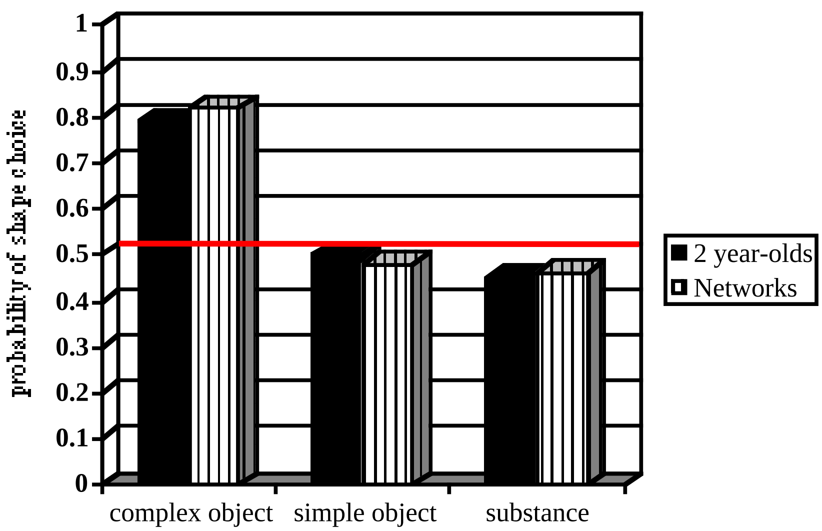

Figure 3. Results for networks trained on Japanese-based vocabularies compared to Japanese 2 year-olds in Imai and Gentner (1997).

Results. We compare the performances of the networks to the patterns reported by Imai and Gentner (1997) in Figures 3 and 4. The solid bars show children's performance; the striped bars show the networks' performance. As is apparent, the connectionist networks were successful in qualitatively modeling children's behavior in the extremes: complex objects vs. substances. However, they do not predict the behavior of English-speaking children in the simple object trials.

While both the Japanese and English networks and the Japanese-speaking children treat simply-shaped objects like substances, English-speaking children treat simple objects and complex objects in the same way. To explain this difference we need something in the English training that puts both complex and simple objects together and differentiates them from substances. The obvious place to look is count/mass syntax. Therefore, in Experiment 3 we add the count/mass syntax correlations to the English-trained networks.

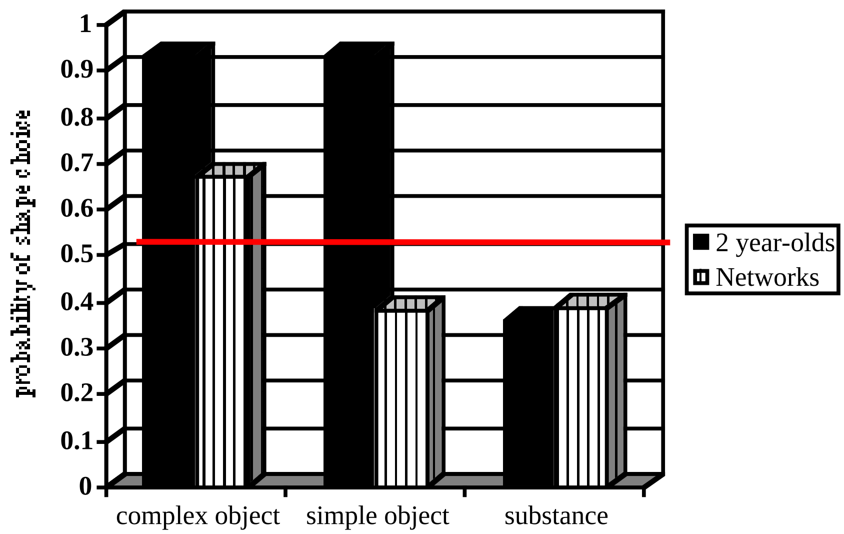

Figure 4. Results for networks trained on English-based vocabularies compared to English 2 year-olds in Imai and Gentner (1997).

4. Experiment 3 – Adding syntax

The goal of Experiment 3 is to see how consideration of the regularities of English count/mass syntax might alter the network's development of learning biases and thus account for children's behavior in the case of simple solid objects. To pursue this issue, we first had English speakers make syntactic judgments and then we trained a connectionist network with the English-like vocabulary used in Experiment 2 plus these syntactic regularities.

Method. Two native English speakers answered the following questions for each of the nouns in the English list used in Experiment 1.

1. Cannot be used with "much" or "less"?
2. Can be used with "many", "several" or numerals?

3. Has a plural form?
4. Can be preceded with “a” or “an”?

Then each noun was classified as **Count,** if all four questions were answered “yes” or as **Mass**, if all four questions were answered “no”. The analysis of these judgments revealed that, indeed, almost all solid, shape-based categories are named by count nouns (95%), and that while most mass nouns do refer to material-based categories (91%), only about half of the non-solid, material-shaped categories are named by mass nouns (47%).

The networks we used in this experiment had the same architecture as in Experiment 2 except for an added Syntax Layer. The syntax layer had two units, one to represent count syntax, one for mass syntax. The networks were trained on the English vocabulary of Experiment 2 plus the new count/mass syntax information and then tested as in Experiment 2.

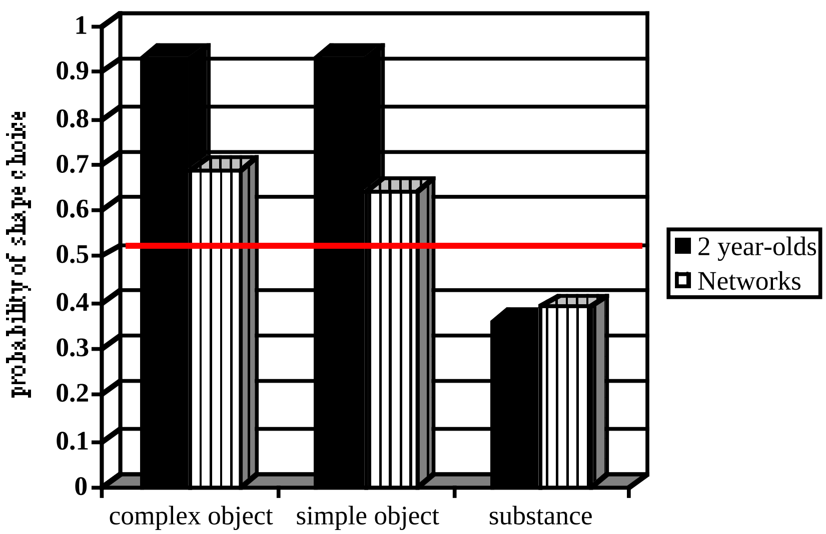

Figure 5. Results for networks trained on English-based vocabularies and count/mass syntax compared to English 2 year-olds in Imai and Gentner (1997).

Results. The results of the network simulations are shown compared to children’s performance in Figure 5. The connectionist networks trained on English with syntax were successful in modeling children’s behavior both in the extremes (complex objects vs. substances) and in the simple object trials. The results of this experiment suggest that English-speaking children treat simply-shaped solid forms as objects because their language (at least the part of their language known to young children) often refers to simple solid forms using the same syntax used to name complexly shaped solid things --- things overwhelmingly named by their shape.

5. **Conclusions**

The kinds of nouns known early by children learning English and by children learning Japanese present an organized structure. Most name solid things and solid things with the same name are similar in shape. A coherent subset of nouns name nonsolid substances and substances with the same name are similar in material (and/or color). The simulations show that these regularities are sufficient for learning a shape bias when naming objects and a material bias when naming substances. Thus, if children are statistical learners like these networks, they may learn noun learning biases.

The similarity of the kinds of nouns known by English-speaking and Japanese-speaking children suggests, contrary to Imai and Gentner's conclusion, the distinction between objects and substances is a product of language learning, and more specifically, lexical learning. Consistent with Imai and Gentner's conclusion, count/mass syntax in English may be responsible for the differences between Japanese- and English-speaking children's interpretation of names for simple solid forms.

Although the present evidence is consistent with a learning account of the shape and material biases --- and indeed demonstrates the computational plausibility of such an account; the evidence does not provide conclusive proof. Indeed, given the coherent and similar structure of early learned nouns in both languages, one has to wonder if there is some filter placed by children on the kinds of nouns they learn. The other possibility is that these regularities --- solid thing named by shape, nonsolid things named by material --- are in the input.

References

Fenson, L., Dale, P., Reznick, J.S., Thal, D., Bates, E., Hartung, J., Pethick, S., Reilly, J. (1993). *The MacArthur Communicative Development Inventories: user's guide and technical manual*. San Diego:Singular Publishing Group

Imai, M & Gentner, D. (1997). A cross-linguistic study of early word meaning: universal ontology and linguistic influence. *Cognition, 62*, 169-200.

Ogura, T. & Watamaki, T. (1997). *Japanese Early Communicative Development Inventory*. Japan.

Quine, W. V. (1969). *Ontological relativity and other essays*. New York: Columbia University Press.

Samuelson, L. K. & Smith, L. B. (1999). Early noun vocabularies: Do ontology, category organization and syntax correspond? *Cognition*, *73* (1), 1-33.

Smith, L. B. (1999). Children's noun learning: How general learning processes make specialized learning mechanisms. In B. MacWhinney (Ed.), *The emergence of language.* (pp 277-303). Mahwah, NJ, USA: Lawrence Erlbaum Associates.

Rosch, E. (1973). Natural categories. *Cognitive Psychology, 4* (3), 328-350.