

The emergence of abstract ideas: evidence from networks and babies

Eliana Colunga* and Linda B. Smith

Department of Psychology, University of Indiana, 1101 East Tenth Street, Bloomington, IN 47405-7007, USA

What is abstraction? In our view, abstraction is generalization. Specifically, we propose that abstract concepts emerge as the natural product of associative learning and generalization by similarity. We support this proposal by presenting evidence for two ideas: first, that children's knowledge about how categories are organized and how words refer to them can be explained as learned generalizations over specific experiences of words referring to categories; and second, that the path of concepts from concrete to more abstract can be observed throughout development and that even in their more abstract form, concepts retain some of their original sensory basis. We illustrate these two facts by examining, in two kinds of learners—networks and young children—the development of three abstract ideas: (i) the idea of word; (ii) the idea of object; and (iii) the idea of substance.

Keywords: connectionism; word learning; object and substance distinction

1. INTRODUCTION

Generalization is a form of abstraction. Webster's Dictionary defines abstraction as follows: '1. a generalized idea or theory developed from specific concrete examples of events; 2. the forming of general ideas or concepts from specific concrete examples'. We pursue the idea that abstraction is generalization and that the processes that create abstract concepts are no different from the processes that create concrete ones. In both cases, knowledge is acquired from specific concrete instances and then generalized to new instances by similarity. What we demonstrate is that these processes can lead to quite abstract forms of knowledge. This account thus makes abstraction a natural by-product of the most ordinary processes of learning about specific concrete instances. We make our case by considering evidence on early word learning. Children learn their first names for things by ostensive definition—by mapping a specific heard word to a specific seen object. But this highly concrete learning leads eventually to ontology, for example, to a partition of entities into overarching categories of objects and substances, and the partition of sounds into the general classes of words and non-words. All this, we will show, is the product of generalizations over specific learned instances. We begin with an overview of several key phenomena in early word learning.

2. THE PHENOMENA

Children learn their first words with considerable difficulty. They need to hear a particular label many times, in many contexts, and used to refer to a variety of

instances before they know the range of things to which the label applies. It is as if they must individually learn each word-object pairing. However, after children have slowly acquired a number of labels, word learning accelerates to the point where children acquire as many as four or five words a day and seem to know, from hearing a single instance labelled, the whole range of instances that fall into the category. Word learning looks very different at this point, as if children had general rules about the nature of words and lexical categories.

One experimental task that documents this newly found knowledge consists of presenting the child with categorization tasks and either labelling or not labelling the to-be-categorized objects. The results of these studies show that labelling has powerful effects for children who are rapid word learners. More specifically, there is a strong correlation between labelling and more abstract thought, such that children are more likely to form taxonomic rather than thematic groupings and to make inferences based on deeper and functional properties when the objects are named than when they are not (Gelman 1988; Keil 1989; Baldwin 1995; Waxman & Markow 1995). What is particularly interesting is that these effects are obtained even when the word used to label the object is a word the child has never heard before. For example, if children are presented with a carrot, a tomato and a rabbit and are asked to group them together, they may well group the carrot with the rabbit—a grouping based on spatial and temporal proximity perhaps, but not one based on the *kind* of things tomatoes, carrots and rabbits are. However, if the child is shown the carrot and told that it is 'a dax', and then is asked to find 'another dax', the child will select the tomato as being in the same category as the carrot. This is interesting because 'dax' in, and of itself, can have no meaning to the child, but somehow the child knows that it is a label and the label refers to categories of a specific kind. These effects of labelling on categorization show (i) knowledge as to what counts as a category

* Author for correspondence (ecolunga@cs.indiana.edu).

One contribution of 16 to a Theme Issue 'The abstraction paths: from experience to concept'.




Figure 1. The child was told the name of the exemplar and then asked which of the two test objects has the same name.

(tomato and carrot do), and (ii) knowledge as to what counts as a word (dax does).

(a) *What counts as a category*

One difference between the slow and fast stages of word learning is that fast word learners seem to know what kinds of similarities are relevant to forming a lexical category. More specifically, rapid word learners seem to know that there are different kinds of categories that are labelled by nouns, and these different kinds of categories are organized by qualitatively different kinds of similarities (Landau *et al.* 1988; Jones *et al.* 1991; Soja *et al.* 1991; Gathercole *et al.* 1995). We concentrate, in this paper, on the distinction that children make between categories of objects and categories of substances. The results from a large number of experiments suggest that young children are rapid word learners, in part, because they know that solid objects and non-solid substances are fundamentally different *kinds* that are categorized by different properties. Specifically, children can rapidly learn the name for a novel solid thing because they expect solid objects to be classified by shape, and children can rapidly learn the name for a novel non-solid substance because they expect those categories to be classified by material (Soja *et al.* 1991).

These conclusions derive from children's performances in the Novel Noun Generalization task (Landau *et al.* 1988; Soja *et al.* 1991). This task consists of showing the child an exemplar, labelling that exemplar 'This is a dax', and then asking the child to indicate which other things can be called by the same name. All the objects and names used are novel; thus rather than children's knowledge about specific categories, this task measures children's generalized expectations about how novel names are attached to categories and about how categories are to be formed. In one experiment, Soja *et al.* (1991) presented children with three solid things as illustrated in figure 1a. The children in this experiment were 24 and 30 months of age, already in the period of fast word learning. In the experiment, the child was told the name of the exemplar and then asked which of the two test objects had the same name. Soja found that children consistently chose the test object matching in shape. In a second condition, Soja presented the children with the three non-solid objects as illustrated in figure 1b. These were made of such non-solid substances as face cream, sand and so forth that children were allowed to touch and thus confirm they were non-solid. When the non-solid exemplar was named and




Figure 2. Children were asked to select an object (exemplar or distractor) by referring to the object with the earlier paired word ('get the toma') or with the earlier paired sound ('get the tone').

the children were asked to indicate the test object that had the same name, they chose the material match.

Importantly, this knowledge is a product of development. Very young children in the slow stages of word learning do not generalize names for solids and non-solids differently. Instead, this distinction becomes increasingly robust during the period from 18 to 36 months. This fact is consistent with the idea that children learn the distinction between objects and substances as they slowly learn their first object and substance terms. However, if this is so, just what are children learning? The knowledge that children exhibit in these experiments is like a generalized rule: they know that the names for solid things in general—whatever their shape, whatever their colour, whatever their material—refer to the shapes of those things. They also know that names for non-solid things in general—whatever their material, whatever their colour, whatever their shape—refer to the materials of those things. This knowledge is an abstraction in at least two ways. First, it transcends the specific properties of specific things. Second, children use this knowledge in an abstract, rule-like way, such that for all categories of solid things shape matters and for all categories of non-solid things material matters, thus dividing the world into two generalized kinds each with a characteristic category organization. This phenomenon thus seems a perfect case for addressing the question of the nature of abstraction and its underlying processes.

(b) *What counts as a word*

Rapid word learners must also know what sounds count as words. The evidence suggests that when children are rapid word learners, they partition sounds into two classes, words and non-words. However, during the period of slow word learning, children do not seem to know how words differ from other sounds. In one critical experiment, Woodward & Hoyne (1999) presented children with two novel objects as seen in figure 2. One, the exemplar, was paired with a sound. In the Word condition, the exemplar was paired with a word ('this is a toma'). In the Sound condition, the exemplar was paired with a non-linguistic sound, such as a tone. The experiment asked whether children would take both the word and the tone as names for the exemplar's category. Woodward & Hoyne (1999) reasoned that if children took the sounds (word or tone) as names, then those names should refer. Accordingly, in the test phase of the experiment children were asked to select an object (exemplar or distractor) by referring to the object with the earlier paired word ('get the toma') or with

the earlier paired sound ('get the *tone*'). In the Word condition, children correctly selected the exemplar, taking the pairing of the exemplar and the word as a naming event. The key results concern the performance of children in the Sound condition. Very young word learners, 13-month-old infants, did take the tone as referring to the exemplar, readily selecting the exemplar when asked to 'get the *tone*'. In contrast, more experienced and faster word learners, 20-month-olds, did not interpret the tone as a labelling event and failed to select the exemplar when asked to retrieve the *tone*. It is as if the older children, but not the younger children, know that only words with their characteristic form can be used to label things. This result is particularly remarkable because the older children with, presumably, greater cognitive skills actually perform worse than the younger children in the Sound condition, choosing randomly whereas the younger children choose correctly.

Knowing what sounds count as words and what sounds do not requires a form of abstract knowledge. It is abstract because it transcends the specific instances such that children recognize even novel word forms—like 'dax'—as words, but do not treat other sounds as words. Again, this knowledge appears to develop during the time that children move from slow to fast word learning. Again, we ask: what processes of abstraction create this knowledge?

The answer we propose, both for learning about objects versus substances and for learning about words versus non-words, is that these abstractions are generalizations over specifically learned instances. More specifically, we propose that abstract concepts are the product of very simple associative processes—processes that consist only of associations between specific words and specific objects but that nonetheless give rise to behaviours usually taken as implying the existence of abstract knowledge. In developmental terms, we suggest that children slowly and painstakingly acquire individual words, learning about the specific sounds that make them up and learning about the specific entities to which those words refer. One natural product of learning these specific associations is generalization, or as Webster's dictionary defines it, abstraction.

To this end, the structure of this paper is as follows. First, we describe experiments with young children and connectionist networks that show that children's knowledge about the features that determine object and substance categories can be created as generalizations over specific instances in children's experiences, and that these generalizations are indeed very abstract, forming a partition that divides categories into two kinds. Second, we present evidence from experiments with younger children suggesting that the notion of what counts as a word is also learned as a generalization over concrete experiences with naming situations, and that because of this, children's abstract knowledge retains some of its original perceptual basis, that is, traces of the specific experiences that created it. Finally, we return to the question that concerns this special issue: what is abstraction?

3. ABSTRACTING THE CONCEPTS OF OBJECT AND SUBSTANCE

The evidence on children's acquisition of an object/substance distinction—on generalizing names for

solids and non-solids differently in the Novel Noun Generalization task—suggests that this knowledge is a consequence of learning, emerging as children learn more words. We propose that the learned associations that constitute early word learning *create*—are in fact the very processes of—abstraction. Figure 3 presents our analysis of the learning task. Step 1 consists of associating names to objects—the name 'ball' to a particular ball and the name 'cup' to a particular cup, for example. A child will do this multiple times for each name as the child encounters new instances. What the child might learn beyond these specific associations depends on what regularities exist across the instances that are learned. If the objects that acquire the same name, for example 'ball', are similar in shape (by being all round), then the child might make a generalization that 'ball' refers to round things. This is shown in step 2, here the child knows that balls are round and cups are cup-shaped. These *first-order* generalizations allow the learner to correctly categorize novel balls and novel cups.

Critically, a higher-order generalization is also possible. This higher-order generalization is over first-order generalizations that balls are round and cups are cup-shaped. This generalization requires regularities, not across specific associations between a word and the concrete objects that the word refers to, but across categories and the way they are organized. For example, if the words that refer to solid things refer also to categories organized by similarity in shape (round, cup-shape), then children could also learn the *second-order* generalization that names for solid things *in general* name categories that span things of similar shapes. As illustrated in step 3, this second-order abstraction requires generalizations over specific names and specific category structures. Making this higher-order generalization would enable the learner to extend the name for any solid thing, even one encountered for the first time, to new instances by shape. Once children make such a generalization, they would clearly have the means to be a fast one-trial learner. Presented with a novel solid thing and a name, they would immediately know that that name referred to a category of things with shapes similar to that of the novel thing.

Can we explain these second-order generalizations from only simple associations and generalization by similarity? Can an associative learner, trained on specific associations, form second-order generalizations of the kind illustrated in the last step of the three-step process? There are many who doubt that associative processes and generalization by similarity are capable of building abstract ideas of this kind. In a series of connectionist simulations, we have shown they can. Connectionist networks provide a framework within which to study the process of generalization. Connectionist networks are simple associative devices that learn specific correlations and then generalize those correlations by similarity. Thus, they are well suited as a test-bed for the idea that abstract concepts are the natural product of more specific, concrete experiences. In this section, we briefly summarize our simulations, showing that generalization of specific associations creates abstract concepts. We do so in a simulation in which we teach a connectionist network a higher-order rule by presenting that network with a series of specific words paired with specific patterns.




Figure 3. Analysis of the learning task. (a) Step 1 consists of associating names to objects. (b) In Step 2, the child knows that balls are round and cups are cup-shaped. (c) In Step 3, the child could also learn the second-order generalization that names for solid things in general name categories that span things of similar shapes.

The series of instances in which we presented the network were based on the actual noun categories that young children learn. More specifically, the training sets for the networks were structured to conform to the first 300 object names that children learning English typically learn. Samuelson & Smith (1999) examined the similarity structure of these categories and found pervasive regularities of the kind that could teach a rudimentary object–substance distinction. Specifically, in this early corpus, nouns that name solid objects mostly refer to categories of things similar in shape, and nouns that name non-solid substances mostly refer to categories of things similar in material. These regularities were robust and pervasive, but naturally, not all-inclusive. For example, ‘soap’ refers to things that are solid, but similar in material, ‘bubble’ refers to things that are non-solid but similar in shape. Nonetheless, for the most part there are regularities across the early lexical categories that children learn. The question is this: are the regularities present in the early lexicon enough, in and of themselves, for an associative learner to develop abstract concepts of object and substance?

(a) *Architecture*

We addressed this question by presenting these regularities to an associative learner, a neural network. The network is an attractor network of the generalized Hopfield type (Hopfield 1982, 1984). The network was trained using Contrastive Hebbian Learning (Movellan 1990), an algorithm that adjusts weights on the basis of correlations between unit activations. The architecture of the network is shown in figure 4. There is a Word Layer, in which each unit corresponds to one word in the training vocabulary. Individual entities are represented on the Perceptual Layer. Activation patterns on this layer represent the solidity, 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. Representing the perceptual properties of objects in a distributed manner captures the graded similarities of objects between categories and enables generalization by similarity. Solidity is represented locally; there is one unit that is solid and another that stands 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 Perceptual Layer

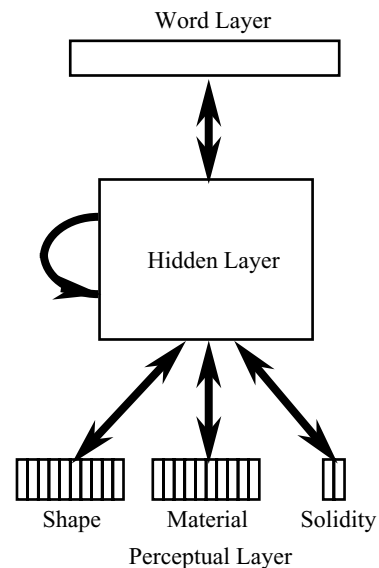


Figure 4. Architecture of the network used in all simulations.

are only connected through the Hidden Layer; there are no direct connections between them.

(b) *Training*

The theoretical idea is that children, at first slowly, learn specific associations between words and individual objects, but when enough of these have been learned, generalizations of that learning reflect the second-order correlations that partition categories into solid and shape-based on the one hand and non-solid and material-based on the other. The goal of the training phase was to teach the network these initial formative associations. To this end, the statistical regularities found in young children’s lexicons and shown in figure 5 were built into the network’s training set in the following way. First, for each noun that the network was to be taught, a pattern was generated to represent its value along the relevant dimension—the dimension adults said characterized the similarities of entities named by that noun. Second, at each presentation of the noun, the value along the irrelevant dimension for that lexical category was varied randomly. For example, the word ‘ball’ was judged 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.

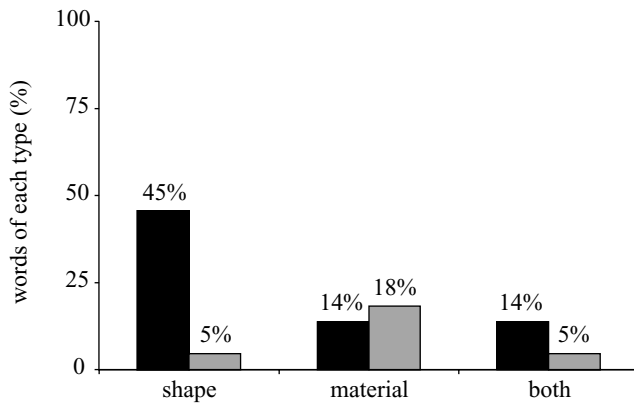


Figure 5. Regularities in the training set for the networks in the simulations. Black bars, words for solids; grey bars, words for non-solids.

All balls presented to the network were defined as having this shape. In addition, each ball presented to the network also consisted of a unique, randomly generated pattern defining the material. So, whenever the unit representing the word ‘ball’ was activated, the pattern representing ball-shape was activated along the Shape Layer, and a different—randomly chosen for this particular instance of ‘ball’—pattern was randomly activated along the Material Layer. The set of specific associations and the categories they represented conformed to regularities reported by Samuelson & Smith (1999). There were solid, shape-based categories, non-solid material-based categories, and so on, in the same proportion as in the children’s corpus, *and* there were exceptions to these overarching regularities in the training set in the same proportion as in the child’s corpus. If learning specific associations that represent overarching regularities is enough to create novel noun generalizations that are different for solids and non-solids, then the networks after training should, like the children who already know many nouns, attend to the shape of novel solid things and to the material of novel non-solid things.

(c) Testing and results

To test this prediction, we presented the networks with novel input patterns—novel shapes and materials—and examined the resulting patterns of activations on the Hidden Layer—the network’s ‘internal representations’. If the network has learned to highlight information about shape in the context of solidity, then the pattern of activation on the Hidden Layer, given an input pattern marked as solid, should mainly represent the shape information from the input pattern and not the material information. If, in addition, the network has learned to highlight information about material in the context of non-solidity, then the pattern of activation on the Hidden Layer, given an input pattern marked as non-solid, should represent mostly the material information from the input pattern and not the shape information. Thus, the patterns of activation on the Hidden Layer for two solid things of the same shape but different materials should be highly similar. By contrast, the patterns of activation on the Hidden Layer for two *non-solid* things of the same material but different shapes should be highly similar (Smith 1995; Smith *et al.* 1997). Accordingly, on each simulated test trial, we measured the




Figure 6. Results from the network simulations using the training set based on the regularities in children’s early object and substance terms.

similarity of the internal patterns of representation for two test objects, one matching the exemplar in shape and one matching the exemplar in material, and calculated the probability of forced choice from these similarity measures.

Figure 6 shows the proportion of shape choices in our simulated version of the Novel Noun Generalization Task after being trained for 32 epochs. The trained connectionist networks, like the 2-year-olds in the experiments of Soja *et al.* (1991), show a preference for shape matches for solids and a preference for material matches for non-solids. In addition, the proportion of shape choices predicted for the solid trials was significantly greater than chance and for the non-solid trials was significantly less than chance. That is, the networks are more likely to choose the shape match for a novel solid exemplar, but more likely to choose the material match for a novel non-solid exemplar. This supports the idea that the statistical regularities in the lexicon are sufficient to create behaviour that conforms to an abstract distinction between objects and substances.

Critically, the network was trained only on specific instances. The only knowledge the network really has is the specific associations—all blended together through common connections within the network—and generalization by similarity. Nonetheless, the network acts as if it has generalizable knowledge about solids versus non-solids. This is because the abstract distinction *exists in* the training set, albeit as a second-order generalization. This means that any overarching regularity in the data should lead to ‘abstraction’. In constructing the training set, we noticed one such regularity in the structure of early lexical categories that seems to distinguish solids and non-solids as distinct types. Specifically, entities that share a name also share their solidity value. In the first 300 nouns learned by children, names refer either to only solid things or to only non-solid things; names do not refer to categories that span across the solid–non-solid boundary. This is true for all words in the early corpus we studied except for one—egg, which adults judged to have both solid and non-solid forms. This fact about the regularities in the early lexicon makes a clear prediction: if noun generalizations are abstractions over specific learned categories, then the generalizations of new names for novel things should adhere to this constraint. That is, given a solid exemplar, the learner should not choose the shape match if this is non-solid, and given a non-solid exemplar the learner




Figure 7. Stimuli and design for the solid exemplar version of the ‘ontology bias’ experiment.

should not choose the material match if it is solid. We tested this idea in simulations with the network and confirmed the prediction. To test this prediction we prepared a new testing set for the network that we had trained on the regularities in children’s early noun vocabulary. In this new training set, we crossed the solidity boundaries. That is, instead of comparing each novel exemplar pattern with shape and material matches of the same solidity as we had done before, we compared the solid exemplar with a non-solid shape match (and a solid material match) and the non-solid exemplar with a solid material match (and a non-solid shape match). Given these new cross-solidity tests, the trained networks did not take the shape (but non-solid) match for the solid exemplar, and did not take the material (but solid) match for the non-solid exemplar.

We also asked if the prediction held for children. In one experiment, we examined how children generalize names for solid things. The participants were children between the ages of 30 and 36 months. The procedure used was a Novel Noun Generalization Task with forced choice. The children were shown an exemplar (i.e. the Teema) and told its name (‘this is the Teema’). The child was then presented with pairs of objects, a shape match and a material match, and asked ‘Can you show me the Teema?’

The critical experimental manipulations are illustrated in figure 7. There were two exemplar objects. The exemplar for one set, the Teema, was a ‘U’ shape covered with red sand-paint. The exemplar for the other set, the Wazzle, was an irregular ‘M’ shape covered with blue cheesecloth. For each exemplar there were three objects matching in material and two sets of items matching in shape. The same-solidity set consisted of three solid objects that matched the exemplar in shape and differed

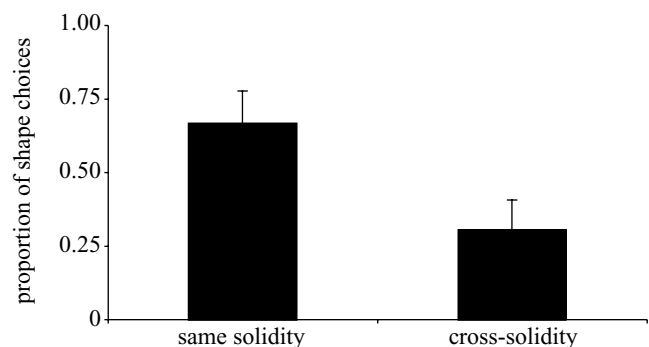


Figure 8. Proportion of shape choices for the solid exemplar version of the ‘ontology bias’ experiment.

in material (e.g. metallic clay, styrofoam covered with fur). Here, we expected children to choose the shape match, generalizing the name for the solid exemplar by shape. The cross-solidity set consisted of shape matches made out of non-solid materials (e.g. shaving cream, hair gel). This is the critical test set. Here, the exemplar is solid and thus generalizing the name by shape is expected. However, extending the name by shape would require the child to form a category containing both solids and non-solids, a kind of category that does not exist in the early lexicon. If children’s novel noun generalizations are generalizations from the regularities over specific learned associations, then children should not extend the name to the non-solid shape-matching test object.

Figure 8 shows the proportion of shape choices for the solid exemplar cross-solidity and same-solidity sets, respectively. In the same-solidity sets, children’s performance replicates previous findings: they consistently chose

the shape match in the trials with solid exemplars. In the cross-solidity sets, children did not choose the shape match over the material match. Children apparently do not want to put solids and non-solids in the same lexical category. This is a finding consistent with the idea of ontology, that is, with knowledge about two different kinds. But remember, the networks also did this and the networks' performance was based solely on generalizations over specific learned pairings of words and individual objects. Regularities across these associations, however, yield second-order generalizations—knowledge that names span categories of things that are all solid, or categories of things that are non-solid.

There are other regularities in that early corpus as well. In addition, if our account of the origins of children's seemingly abstract knowledge about kinds is right, then these regularities should also be apparent in children's novel noun generalizations. In one set of experiments, we specifically examined the regularities that characterize categories of non-solids. One of these, as we have seen, is that categories of non-solids contain only non-solid things, but a second regularity is that non-solids have a particular kind of shape. Non-solid things tend to be flat, rounded and irregular (a splat of oatmeal, a mound of sand), in contrast to solid things, which can be quite complex in shape, taller than they are wide, and angular. If children's abstract knowledge is merely a generalization over specific learned instances, and if these regularities exist across those specific learned instances, then these regularities should also exist in children's novel noun generalizations.

We tested these ideas in one experiment on children's name generalizations for non-solids. Again, the main prediction is that children should extend a name for non-solids to a same material match, but should do so only if that material match is also non-solid. To test this, we needed to present children with material matches that could be both solid and non-solid. For example, plastic in a solid and a gel-like form; paint in a liquid or solidified form. There were two kinds of trials. On same-solidity trials, as illustrated in figure 9, children were presented with a non-solid exemplar, non-solid shape match and a non-solid match. On the cross-solidity trials, children were presented with a non-solid exemplar, a non-solid shape match and a *solid* material match. We also manipulated the shapes of the solid material matches, making the shapes more solid-like or more non-solid-like. For half of our sets, the solid material match was of a non-constructed form, more natural for non-solid things; for the other half of our trials the solid material match was that of a constructed form. Since kind-of-shape is correlated with solidity, it should be part of what defines the partition between kinds, between the correlation cluster of solid, constructed shape, and categorized by shape versus the correlational cluster of non-solid, non-constructed shape, and categorized by material.

The results are shown in figure 10. Children generalized the name for the non-solid exemplar to the material match when it was also non-solid (the same-solidity trials), but were much less likely to do so when the material match was solid (cross-solidity trials). Moreover, children were more likely to generalize the name for a non-solid to a solid material match if it was in the non-constructed shape typical of non-solids, than if it was in the constructed

shape typical of solids. This is as it should be by our account of abstraction as generalization over specific learned instances. If the regularities are in the learned specific associations, then they must also be in the abstraction. These results, however, reveal the correlational origins of children's abstract knowledge about kinds. The complexity of the shapes matters for children's partitions of things into solids and non-solids because shape complexity is correlated with solidity and category structure.

In sum, the experiments summarized in this section suggest the following: first, generalizations over simple associations can give rise to what appears to be abstract knowledge. This will happen as a natural product of learning specific associations if those specific associations present overarching regularities. Second, because the origin is generalization over specific instances, children's knowledge will reflect whatever that pattern of regularities is across those instances. If the regularities are in the second-order generalizations, that is, across the first-order generalizations, as is the case with solids versus non-solids, then children will learn what looks like an abstract partition. If the regularities occur in a more graded way, as is the case of the correlation between shape complexity and solidity, then the knowledge will be more graded and contextual.

This account of abstraction as generalization over specific instances is generally applicable to a variety of kinds of knowledge, not just children's knowledge about objects and substances. We illustrate this in the next section by considering children's discovery of the forms that count as words and non-words.

4. ABSTRACTING WHAT IS A WORD

As children learn specific labels associated with specific instances, they form second-order generalizations about what kinds of features matter for different kinds of categories for, for example, solids and non-solids. We propose that, in the same way, they also form second-order generalizations about the kinds of features that matter for words. The correlational origins of this emerging knowledge are evident in children's developmental homing in on the entities accepted as words. Recall that, in the early slow period of word learning, children seem not to distinguish different kinds of sounds as words versus non-words, but later, in the period of accelerated word learning, children treat only certain sounds as referential labels, but not other sounds. In this section, we consider in finer detail this transition from an open acceptance of all sounds to the narrowing definition of what counts as a word. If children learn the category 'word' as a second-order generalization, then their notion of a 'word' should reflect—just as do their notions of object and substance—whatever regularities are out there between the features of words and categories. That is, if words are formed as a generalization over whatever sound features correlate with categories, then children should take as a possible word any aspect of a labelling event that covaries systematically with categories. For example, for children relying solely on spoken language, these could be properties such as being a speech sound with particular spectral and prosodic forms, being produced by people, coming out of mouths, or co-occurring with pointing and eye gaze to the object. Over



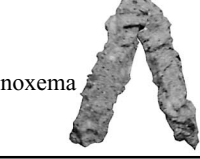








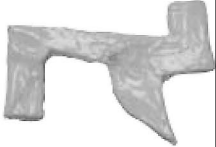
exemplar	shape match	material match		
		same solidity	cross-solidity	
			simple shape	complex shape
 purple hair gel	 shaving cream  noxema	 purple hair gel	 purple plastic	 purple plastic
 hand lotion	 wheateena  toothpaste	 hand lotion	 clay + fabric paint	 clay + fabric paint

Figure 9. Stimuli and design for the non-solid exemplar version of the ‘ontology bias’ experiment.

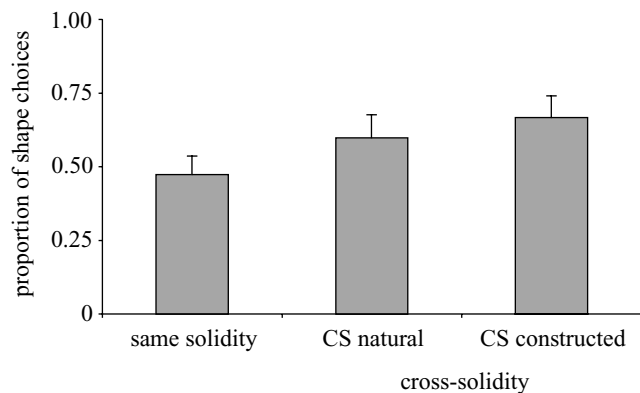


Figure 10. Results for the non-solid exemplar version of the ‘ontology bias’ experiment. CS, cross-solidity.

time, children presumably will abstract away the less essential aspects, but early in development, they should take all of these as defining of a word *if* they co-occur with categories.

In one experiment, we showed this to be true by showing that children treat a spoken word as a word only if it is produced in the normal way by a human speaker. In the experiment, we presented children with a novel object and labelled it. For half the objects, the label was produced in the normal way, spoken by the experimenter, and thus the label emanated from a mouth. For the other half of the objects, the label was also a spoken word, but the source was a hand-held tape recorder. All other aspects of the procedure were the same. The label was associated with the exemplar in a training phase and then,

in a test phase, the child was presented with a choice between the exemplar and a distracter and instructed to ‘get the *word*’, where the word came from the same source—mouth or recorder—as in the training phase. The children (20–26 months old) interpreted the word as a label referring to the exemplar object *only* when it was spoken from the mouth, not when it was produced by the recorder. This is as it should be, if words are second-order generalizations over specifically learned instances that include, among the regularities, characteristic of words—being produced by a mouth and social settings as well as the spectral properties of human speech.

Results from subsequent experiments suggest that mouths may be the most salient early predictor of what counts as a word—more important, in fact, than the spectral properties of the sounds. We tested this by repeating the experiment but replacing the words with non-word-like sounds that were presented, again, either by a human speaker or by a tape recorder. We found that children took the non-word sound as a word when it emanated from the mouth. Apparently, source of sound is a better predictor of a labelling event than is the nature of the sound.

In a third experiment, we tested a perhaps more counterintuitive prediction that follows from this idea of the origin of the category ‘word’. According to our account, any event that systematically predicts category membership, even if not a word, should be taken as a word that refers to, or names, an object. To test this, we studied the domain of animals. This is a domain familiar to children and one in which sounds other than words co-occur systematically with categories. Specifically, animal sounds

correlate with animal category. Dogs bark, cats meow, elephants trumpet and so on. If a label is only a second-order generalization over the sounds that correlate with categories, then animal sounds as a set present the requisite regularities to be category labels. Thus, animal sounds should be taken as names for animals. The experimental task was as follows. Children were presented with a novel animal and that animal was paired with a sound. The sound could be either an animal sound—the kind of sound that points to a category of animals and thus should be taken as a name of an animal category—or the sound could be a mechanical sound—a kind of sound that is not correlated with animal category and thus should not be taken as a label for an animal. We tested whether children took the sound as a label by asking whether they interpret that sound as referring. In the test, children are presented with the original exemplar and a distracter. The experimenter plays the originally paired sound and the question is whether that sound leads the child to select the exemplar. The answer is ‘yes’ for animal sounds used as labels for animals but ‘no’ for mechanical sounds used as labels for animals.

This makes sense by our generalization account. Among the specific instances of sound–object pairings that children have learned are animal sounds that refer to categories of animals. These pairings of frog sounds for various frogs, and barking for dogs, and tweets for birds, set up a second-order generalization that animal sounds are labels for animals. In these experiments, we also showed that it did not matter if the animal sounds or mechanical sounds emanated from the mouth of a speaker or from the animal. This is in contrast to spoken word forms, which must come from the mouth to be interpreted as words by children. We predicted that the source of sound (mouth or recorder) would not matter for animal sounds because the cluster of features defining animal sounds does not include a specific source. In the lives of young children, animal sounds come from the animal (real or toy), as well as from human speakers who imitate these sounds.

The results of these experiments indicate that the knowledge children have reflects the details in their experiences. First, we found that sounds emanating from mouths are always taken as names because emanating from a mouth is one of the most systematically correlating features of naming situations. Second, source matters for words but not for the animal sounds because animal sounds emanate from the mouths of live animals, from the inside of stuffed animal toys, and from the mouths of people imitating animals. All this fits the idea that children’s knowledge about the form names take seems to be generalized over the specific instances of their experience, and at age 20–26 months, still tied to the statistical regularities in their experience. With more learning, presumably more incidental perceptual features drop out—being produced by a mouth may not be as important for adults as it is for children—creating a more and more abstract concept of word. The point is, even a concept as abstract as what counts as a word seems to have its beginning in correlations, and during the course of development to be tied to the specific instances experienced.

Let us evaluate: these two experiments on children’s emerging knowledge of what counts as a word show the general utility of the idea that abstract knowledge is simply

more specific knowledge and the generalizations that that specific knowledge gives rise to. The two studies on the emerging knowledge of what counts as a word support this by showing, first, that children’s early notions of what is a word include features that are typically present when object categories are labelled with words for young children. Second, the results support the idea of second-order generalizations by showing that whenever regularities over categories—that is, over the first-order generalizations—are sufficiently systematic, they will give rise to abstract rule-like performance. In this case, the rule gleaned from the specifics of naming situations is that animal sounds (but not tones, and not mechanical sounds!) refer to categories. Animal sounds refer, by our account, not because they are intrinsically special, but because they systematically point to a class of categories in a way that tones and mechanical sounds do not.

5. CONCLUSION

What do we mean by abstraction? The idea presented here is that abstraction may often be nothing more than the result of specific learning of specific instances—not a separate process, not a separate kind of knowing—by the natural and very ordinary process of generalization by similarity. At the very least, we have shown how such a notion of abstraction may explain the origins of ‘abstract ideas’ in developmental time. In their word learning, children start as slow instance-by-instance learners. However, after a brief period of such slow learning, they become something else. They become learners who have partitioned the learning space into fundamentally different kinds of problems—learning about words, learning about animal sounds, learning about objects, learning about substances. The knowledge that children bring to bear within these partitioned learning spaces is quite abstract, almost rule-like—knowledge that solid things are named by their shapes, that non-solid things are named by their material, that words refer. This knowledge is properly considered abstract because it applies to never-seen-before instances, transcending particular shapes, materials and sounds. However, our studies also suggest that the transcendence is not complete—at least not for the young children we studied—that the correlational origins are apparent in the graded and contextual nature of children’s behaviour.

This research was supported by a grant from NIMH (R01 MH 60200).

REFERENCES

- Baldwin, D. A. 1995 Understanding the link between joint attention and language. In *Joint attention: its origins and role in development* (ed. C. Moore & P. Dunham), pp. 131–158. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gathercole, V. C. M., Cramer, L., Somerville, S. C. & Haar, M. J. O. D. 1995 Ontological categories and function: acquisition of new names. *Cogn. Dev.* **10**, 225–251.
- Gelman, S. A. 1988 The development of induction within natural kind and artifact categories. *Cogn. Psychol.* **20**, 65–95.
- Hopfield, J. J. 1982 Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl Acad. Sci. USA* **79**, 2554–2558.

- Hopfield, J. J. 1984 Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl Acad. Sci. USA* **81**, 3088–3092.
- Jones, S. S., Smith, L. B. & Landau, B. 1991 Object properties and knowledge in early lexical learning. *Child Dev.* **62**, 499–516.
- Keil, F. C. 1989 *Concepts, kinds and cognitive development*. Cambridge, MA: MIT Press.
- Landau, B., Smith, L. B. & Jones, S. S. 1988 The importance of shape in early lexical learning. *Cogn. Dev.* **3**, 299–321.
- Movellan, J. 1990 Contrastive Hebbian learning in the continuous Hopfield model. In *Proc. 1990 Connectionist Models Summer School* (ed. D. Touretzky, J. Elman, T. Sejnowski & G. Hinton), pp. 10–17. San Mateo, CA: Morgan Kaufmann.
- Samuelson, L. & Smith, L. B. 1999 Early noun vocabularies: do ontology, category structure and syntax correspond? *Cognition* **73**, 1–33.
- Smith, L. B. 1995 Self-organizing process in learning to learn words: development is not induction. In *Basic and applied perspectives on learning, cognition, and development*, vol. 28 (ed. C. A. Nelson), pp. 1–32. Mahwah, NJ: Lawrence Erlbaum.
- Smith, L. B., Gasser, M. & Sandhofer, C. 1997 Learning to talk about the properties of objects: a network model of the development of dimensions. In *Perceptual learning*, vol. 36 (ed. R. L. Goldstone, D. L. Medin & P. G. Schyns), pp. 220–255. San Diego, CA: Academic Press.
- 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.
- Waxman, S. R. & Markow, D. B. 1995 Words as invitations to form categories: evidence from 12- and 13-month-old infants. *Cogn. Psychol.* **29**, 257–302.
- Woodward, A. L. & Hoyne, K. L. 1999 Infants' learning about words and sounds in relation to objects. *Child Dev.* **70**, 65–72.