Flexibility and variability: Essential to human cognition
and the study of human cognition

Eliana Colunga

Department of Psychology

University of Colorado

Boulder 345 UCB Boulder, Colorado 80309-0345

Linda B. Smith

Department of Psychology

Indiana University

1101 East 10th Street

Bloomington, IN 47405-7007

Abstract

Traditional theories of cognition concentrate on the problem of stability, how it is that people perform the same cognitive act over and over despite different task pressures and varying circumstances. We suggest a different approach that focuses on the flexible and the inventive. From a dynamic systems perspective, intelligence is the adaptive flexibility that integrates the stability of past experience with the idiosyncrasies of the moment. In this paper we demonstrate the value of this perspective in the context of children's novel noun generalizations. We describe D-ALA, a dynamic systems model that embodies the principles of the account and use this model to explain the findings of three different experiments that show how young children flexibly and smartly adapt to novel circumstances in the task of learning new words.

Flexibility and variability: Essential to human cognition and the study of human cognition

Traditional theories of cognition concentrate on the problem of stability, how it is that people perform the same cognitive act over and over despite varying tasks and the vagaries of the moment. How is it, these theories ask, that we understand people, frogs, and trees to all be living things? How is it that we understand the frog in the pond, the frog in the jar, and the frog in the comic strip to all be the same kind of thing –a frog? Traditional theory explains these stabilities by positing stable knowledge structures; one constant internal representation is activated each time we understand some thing to be a living thing, another that is activated each time we understand something to be a frog. (Keil, 1994)

This focus on stability as the principle phenomenon to be explained has led to a set of well-accepted practices in the study of cognitive development. These are:

1. Concentrate on findings that are *generalizable* across tasks because it is the pattern of performance that transcends specific tasks that is most likely to be the result of, and thus informative about, the underlying representational constants.
2. Seek tasks that *minimize variability*. By reducing variability, we can “see through” performance to the knowledge representations underneath.
3. Emphasize *optimal performance* and the best conditions. Performance that comes closest to our theoretical definitions of the underlying knowledge (or to adult standards) present the clearest view

of just what children know, a view unclouded by the noise of performance and task.

These “good research practices” direct us to the stable, the general, the optimal. Using these methods, researchers of cognitive development have learned a lot about early cognitive competencies (e.g., Keil, 1989; Spelke, 1990; Mandler, 1992; Gelman, Coley, & Gottfried, 1994; Kemler Nelson, Frankenfield, Morris, & Blair, 2000; Bloom, 2000; Booth, Huang, & Waxman, 2005; Lavin & Hall, 2001). Nonetheless, this approach may be fundamentally flawed, both theoretically and empirically (see, Thelen & Smith, 1994; Smith, 2005). The theoretical failure lies in the concentration on the stable and the generalizable at the expense of the flexible and inventive. Intelligence is decidedly not doing the same thing over and over again. Intelligence –of the kind that humans have over machines – is an adaptive flexibility that enables us to smartly do **novel** things that integrate the stabilities of past experience with the idiosyncrasies of the moment.

If we are to understand human intelligence we must understand the processes that make this inventiveness. If we are to understand the processes that make inventiveness, we must study more than the generalizable, the stable, and the optimal. In this paper, we demonstrate the value of this perspective in the context of children’s smart novel noun generalizations. The organization of the paper is as follows. First we briefly review the literature on word learning and present the attentional learning account. The central idea is that attention integrates many processes operating over many time scales, creating a momentary

attentional act that is inventive. Next we describe a dynamic systems model that embodies the principles of the account. Then we illustrate the importance of this approach in explaining the findings of three different experiments that show how young children flexibly and smartly adapt to novel circumstances in the task of learning new noun categories.

The phenomenon: Smart word learning in children

Although children initially learn words slowly, by the time they are 2 to 3 years old, they seem to be expert learners. Indeed, many studies have shown that 2- and 3-year old children need to hear only a *single* object named to correctly generalize that name to the whole category (e.g., Golinkoff, Mervis & Hirsh-Pasek, 1994; Markman, 1989; Smith, 1995; Waxman & Markow, 1995). Young children's facility in mapping nouns to categories is particularly remarkable in that not all categories are organized in the same way. Instead, there are different kinds of categories --- animals, objects, and substances --- with fundamentally different organizational structures. The ease with which children learn names for these different kinds suggests that they understand the different organizations. One way to explain this knowledge is to put fixed knowledge representations in the head: knowledge of the kinds of similarities relevant to animals, to artifacts, and to substances. Considerable research has sought to do just this (e.g., Soja et al, 1992; Hall, 1996; Booth & Waxman, 2002; Deisendruck & Bloom, 2003). However, the evidence also hints that there may be no set of fixed representations just many different pieces of knowledge that are softly assembled in the task, smartly directing children's attention to the right properties for that kind of thing. This paper outlines such an account. We begin with a brief review of the literature.

The novel noun generalization task.

The Novel Noun Generalization (NNG) task was originally designed to measure the one-trial category learning that young children seem to naturally show (see Katz, Baker, & Macnamara, 1974; Carey & Bartlett, 1978; Markman, 1989). In this task, the child is shown a single novel object, told its name (e.g., *This is a toma*) and then asked what other things have the same name (e.g., *Where is the toma here?*) With no more information than this – a novel name applied to a single novel thing – 2- and 3-year-olds extend the name in systematic ways that seem right to adults. Experimenters have studied three kinds of entities (as shown in Figure 1) and found three different patterns of generalization. Given objects with features typical of animates (e.g., eyes), children extend the name narrowly to things that are similar in multiple properties. Given a solid inanimate thing, children extend the name broadly to all things that match in shape. Given a nonsolid substance, children extend the name by material. These are highly reliable and replicable results – obtained by many researchers – and in their broad outline characteristic of children learning a variety of languages (e.g., Jones, Smith & Landau, 1991; Kobayashi, 1998; Jones & Smith, 2002; Yoshida & Smith, 2001; Markman, 1989; Booth and Waxman, 2002; Gathercole & Min, 1997; Imai & Gentner, 1997; Landau, Smith & Jones, 1988, 1992, 1998; Soja, Carey, & Spelke, 1991; see also, Gelman & Coley, 1991; Keil, 1994). However, children’s specific patterns of performances in this task are also highly dependent on the task, the stimuli, and specific language cues.

Insert Figure 1 about here

1. Children’s attention to the different properties of different kinds is most robust in naming tasks. Many experiments have included control tasks that are

identical to the NNG task, except the object is not named. Instead, children are shown the exemplar and then are asked what other objects are “like” or “go with” the exemplar. With these cues, children do not systematically attend to the different properties of different kinds. (e.g., Imai & Gentner, 1997; Landau, Smith & Jones, 1988, 1992; 1998; Jones, Smith & Landau, 1991; Soja, Carey, & Spelke, 1991). This fact suggests a link between naming and knowledge about category specific organizations (see, Deisendruck & Bloom and also Colunga & Smith, in press, for a further discussion).

2. Kind-specific name generalizations emerge *with* vocabulary growth

(Landau, Smith & Jones, 1988; Soja, Carey & Spelke, 1991; Jones, Smith & Landau, 1991; Samuelson & Smith, 1999, 2000; Smith, 1995). The evidence indicates that the tendency to attend to shape in the context of naming emerges only after children already know some 50 to 150 nouns. For children learning English, a bias to extend names for animates by multiple similarities and a bias to extend names for nonsolid substances by material emerge even later (see, especially, Jones et al, 1991; Samuelson & Smith, 2000). Thus, biases to attend to different properties when extending names for different kinds co-develops with increasing vocabulary, a fact consistent with the idea that children’s word learning helps create these generalized expectations about different kinds.

3. Kind-specific name generalizations are modulated by linguistic cues.

One area of relevant research concerns the influence of count and mass syntactic frames on English-speaking children’s interpretations of novel object and substance. Count nouns are nouns that take the plural and can be preceded by words such as *a*, *another*, *several*,

and *few*, as well as numerals. Count nouns thus label things we think of as discrete individuals --- chairs, trucks, shirts, studies, and hopes. Mass nouns, in contrast, cannot be pluralized but instead are preceded by words such as *some*, *much*, and *little*. Mass nouns thus label things that are conceptualized as unbounded continuous masses --- water, sand, applesauce, research, and justice. Past research shows that count syntactic frames (e.g., *a mel*, *another mel*) push children's attention to the shape of the named thing whereas mass syntactic frames (e.g., *some mel*, *more mel*) push attention to material (e.g., Gathercole, Cramer, Somerville, & Haar, 1995; McPherson, 1991, Soja, 1994). Other studies show that linguistic cues associated with animacy (descriptors such as *happy* or the personal pronouns *he* and *she*) push attention to multiple similarities even when the objects do not possess perceptual cues such as eyes that are diagnostic of animacy (Yoshida & Smith, 2005; Yoshida et al, 2001; Booth & Waxman, 2002) and even when these words are not referring to the object in question (Colunga & Smith, 2004, Colunga, 2006). In brief, the evidence shows that language exerts an on-line influence on children's category formation in the NNG task.

4. The particular language learned influences children's novel noun generalizations. Although there are overarching universals in the name generalizations of children learning different languages --- solid rigid things tend to be named by shape, nonsolid things by material, and things with features suggesting animacy by multiple similarities --- there are differences as well (Colunga & Smith, 2005; Gathercole, Thomas, & Evans, 2000; Imai & Gentner, 1997; Yoshida & Smith, 2005; Kobayashi, 1998). For example, Japanese speaking children are more sensitive to animacy cues than English speaking children (Yoshida & Smith, 2003); Spanish-speaking children are more

sensitive to count mass syntax cues than English speaking children (Colunga & Smith, 2002); and Japanese- and English-speaking children place the boundaries between animals and objects and between objects and substances in different places (Yoshida & Smith, 2003b; Colunga & Smith, 2005).

Summary

The results show that children consistently use different kinds of similarities to organize different kinds of categories. Because this knowledge is manifest in a task in which novel things are named by a novel name and then generalized to other novel things, this is a task that demands inventiveness, that demands doing something that has not been done before. The task effects documented in the literature provide clues as to how this feat of inventiveness might be accomplished. A complete explanation of the phenomenon clearly requires the integration of many different sources of information—all relevant to different kinds of categories. These include the in-task cues of the properties of the named thing and the words said; these also include the regularities among words, properties and category organization in the learners past.

One way to understand this inventiveness is through the process of attention—what properties the child attends to when the object is named and what similarities between this exemplar and the generalization objects that the child attends to when deciding whether that name applies to the new instance. That is, the child's problem in the NNG task may be construed as one of attending—in the moment—to the right kind of similarities for the specific kind of thing. Under this conceptualization, a unified theory needs to explain how the allocation of attention at any given moment is product of a smart system that integrates different forces on attention into a single coherent and

adaptive behavioral outcome. Moreover, such an integration will need to occur over nested time scales. The extant evidence suggests that the immediate cues in the task matter. These are processes operating on fast time scales; the immediate perceptual input (the object, its properties, the heard name) orient attention in the moment. But they do so in context sensitive *habitual* ways. Thus, the longer term history of the learners also matters. The regularities among words, object properties and category organization are acquired over slower time scales. The regularities include the learner's already acquired vocabulary, the noun categories and linguistic correlates of different kinds of categories, and the correlations of words and syntactic devices in that particular language. These also influence the system by adjusting the weights between cues and attention to certain properties and dimensions of category organization.

The Dynamic- Attentional Learning Account

The Attentional Learning Account (ALA, cite) explains children's novel noun generalizations through learned cues that shift dimension weights among potentially relevant dimensions. The core idea is that attention is dynamically tied to contextual cues in the moment, to learned associations in the past, and through these processes selects some similarities over others, directing the learners attention to just the right information for the kind of thing being named. Attention is a powerful process because it is inherently multi-causal, integrating in the moment influences over multiple time scales. Further, attention strongly constrains the future by determining what will be learned in the moment. In brief, attention is one important mechanism through which prior knowledge is brought into the present, through which the past constrains learning and thus constrains future knowledge, building –moment to moment – the developmental trajectory.

A considerable literature on attentional learning in a variety of different kinds of tasks shows that attention at any given moment is influenced by (1) the system's previous history about the effectiveness of allocating attention to one or another cue in similar contexts, (2) by recent actions, memories and attention allocation, and (3) by percepts and internal state. Previous versions of the attentional learning account of children's novel noun generalizations have concentrated on changes on one time-scale (the time scale of development and learning, see Colunga & Smith, 2005; Smith, Jones, Landau, Gershkoff-Stowe, and Samuelson, 2003). Here, we outline an extended version of this account, the Dynamic – Attentional Learning Account (D-ALA) that integrates experience at at least three different timescales. These are illustrated in Figure 2 and discussed below.

Insert Figure 2 about here

Consider a child in a NNG. That child brings to the experiment some already known words and categories as well as associations among linguistic cues (e.g., count syntax), properties (e.g., solidity), and the similarities relevant to specific nominal categories (e.g., shape). This knowledge is acquired slowly. Over the first two years, a child acquires a vocabulary consisting mostly of object names for categories of solid man-made things organized by shape (e.g., ball, spoon, cup). This learning sculpts attentional biases leading to a tendency to attend to shape that is likely to be particularly strong in the presence of an object name and a solid man-made thing. At the next time-scale down, the child's recent allocation of attention, for example, to material and taste if the child was just eating a snack, or to form and movement, if the child was just rolling a ball, will also influence attention. Finally, there are the faster processes of attending in

the moment. At this time scale, the immediate input, the characteristics of the object named, the wording used, and the goals of the child add another layer of influence to attention allocation. Importantly, these in-the-moment influences on attention may work either by activating memories and processes at slower time scales or through the intrinsic attention grabbing properties of specific things (e.g., glittery or noisy things). Thus, these are nested processes and time scales with past learning influencing the faster attentional processes that underlie learning in the moment and with attention in the moment influencing the regularities that are learned over longer time scales. In these ways, attention in the moment, is a blended product of multiple processes and nested within processes operating over slower time scales.

The beauty of a system that integrates processes and inputs over different timescales is that it can generate behaviors that are optimal, stable, generalizable, the behaviors that seems so essential to higher cognition, and it can also generate behaviors that are flexible and inventive. Through such multi-layered and nested processes, the child's noun learning can effectively adapt to different environments (say, English-speaking or Japanese-speaking), a changing world (moving from one activity to the next, one place to the next), and changing inputs and task goals (biting the cookie, but rolling the ball). Thus, this soft-assembly of attention out of unthinking and nested processes may lead to robust, seemingly knowledge-based behavior and also to smart, flexible adaptability.

If this Dynamic Attentional Learning Account is even near right, it means we need a different way of studying the system than one that merely documents the stable, the generalizable, the optimal. Instead, to understand a system such as that outlined

above, we need to examine its adaptability. Rather than probe the system at the points of most stability, we need to query it at points of variability. Only by looking at variability, at how the system performs when perturbed, can we disentangle the contributing (and potentially tightly coupled) component processes. (Thelen, 1989). Thus as researchers, we need to push the limits of the system under *novel circumstances*. In the remainder of the paper we describe a dynamic model of the processes involved in learning new words and test it in three situations that ask the system to interpolate between known regularities, to extrapolate, and to combine information in novel ways.

The Dynamic Attentional Learning Account (D-ALA)

The mathematical model described here –as all models– is a simplification. The goal is to capture and explore the main features of a dynamical attentional system: integration of a variety of cues, integration at different timescales, and emergent stable but flexible behavior. The model consists of equations representing three nested timescales: developmental time, attention allocation, and decision making. We describe each of these in turn.

Developmental time. This time scale captures the regularities in the long term experiences of children, knowledge that they bring to any encounter with a novel word and novel object. According to the model, these regularities lead to the increased weighting of attentional cues (e.g., solidity) that have been associated with category relevance of some property (e.g., shape) in some context (e.g., naming). To capture these regularities, the model receives as input the characteristics of the child's noun vocabulary – the number of words they know, the dimensions that define them, and the similarities that organize known noun categories. From this, the model generates attention maps for

different dimensions. In this version of the model, we concentrate on maps for the dimensions of shape, material, and texture, given the cues of eyes and solidity. (For this version of the model, we are making a distinction between cue properties such as eyes that signal the kind of category and the dimensions such as shape or texture along with lexical categories are organized and to which attention is allocated in making a decision about whether a noun applies to some novel instance. A more complete version of the model does not make this distinction, but lets it emerge from the statistical properties of the lexical corpus.)

The attentional landscape for any dimension is constructed based on the number of lexical categories that are organized by that dimension, and defined over a hyperspace formed by the cue dimensions. Intuitively, for each dimension, we calculate the number of words in the intersection of each cue (e.g., solid-eyes-shape, solid-noneyes-shape, etc) and use these numbers to define points in the cue space. Then we fit a hyperplane (2D in this case) over these points on the cue space to make the attention map for that dimension. More specifically, the evidence gleaned from the vocabulary is turned into an attention map for each dimension as follows:

$$E_{dim} = \begin{bmatrix} \sigma\left(\frac{N_{dim \cap sol \cap neye}}{N_{sol \cap neye}}\right) & \sigma\left(\frac{N_{dim \cap sol \cap neye}}{N_{sol \cap neye}}\right) \\ \sigma\left(\frac{N_{dim \cap sol \cap neye}}{N_{sol \cap neye}}\right) & \sigma\left(\frac{N_{dim \cap sol \cap neye}}{N_{sol \cap neye}}\right) \end{bmatrix}$$

Where the factors E_{dim} are calculated as a normalized logistic sigmoid function of the proportion of words for that cue organized by that dimension ($N_{dim \cap cue}$) divided by the

number of words with that cue value. Thus, the output of this module is then the attention maps for each dimension defined as:

$$\begin{aligned} \dim(sol, eye) &= E_{\dim} \bullet f(sol, eye) \\ f(sol, eye) &= \begin{bmatrix} (1-sol)*(1-eye) & sol*(1-eye) \\ (1-sol)*eye & sol*eye \end{bmatrix} \end{aligned}$$

Where $\dim(sol, eye)$ is the function that determines how much attention should be allocated to dimension $\dim$, when the cue solidity is sol and the cue eyes is eye .

Task time. At a faster timescale, the model receives the attention maps resulting from the regularities in the organization of the learned vocabulary. It also receives task input and the perceptual properties of the objects in view. In turn, this nested process calculates the allocation of attention, represented as a number between 0 and 1 for each of the three dimensions involved in the model (shape, material and texture), using the output of the Developmental time and the values along the cue dimensions of the objects in view. (This is also where linguistic input such as count and mass syntax would enter the model as cues; the present simulation does not include these linguistic cues). Thus, the attention vector during task time is given by:

$$A(sol, eye) = [\dim_1(sol, eye), \dots, \dim_i(sol, eye), \dots, \dim_n(sol, eye)]$$

where $A_{\dim}$ is the attention allocation to the different dimensions given the cues present in the task at hand.

Decision/Action time. The attention allocation numbers are fed into the final element of the model, together with task demand information such as whether the objects are presented as part of a forced choice or a yes/no task. The similarity of the test items to the standard is computed by weighing attention to the different matching and mismatching dimensions according to the attention allocation predicted by the task level

time. Thus, the similarity of the test item to the standard, depending on whether they match or mismatch along the different dimensions ($sim(testItem)$) is defined as:

$$sim(testItem) = A(solidity(testItem), eyes(testItem)) \cdot Match(testItem)$$

$$Match(testItem) = \left[\dots, \left\{ \begin{array}{ll} match & \text{if } testItem \text{ matches } dim_i \\ mismatch & \text{if } testItem \text{ doesn't match } dim_i \end{array} \right. , \dots \right]$$

For these simulations, the *match* factor was set at 0.95, and the mismatch factor at 0.2.

If modeling a yes/no task, the similarity between the exemplar and the test item directly predicts the probability of making a “yes” choice, $prob(testItem)$, as defined below. In a forced choice task, the similarities to the different test items are calculated and then used to predict the probability of choice, $prob(testItem_i)$.

$$prob(testItem) = sim(testItem)$$

$$prob(testItemA) = \frac{sim(testItemA)}{sim(testItemA) + sim(testItemB)}$$

The model was used to make predictions on three novel tasks of considerable importance to understanding the inventive nature of human cognition: interpolation, extrapolation, and the blending of multiple, potentially contradicting knowledge sources. The study of how a system operates in novel task contexts such as these, how it brings information from the past into the present to generate *new* behavioral outcomes, is crucial to understanding its internal dynamics and the source of its inventiveness.

Interpolation

The first task was designed to explore the problem of interpolations or how the cognitive system, through attention, finds solutions to tasks that fall between two known anchors on some continuous dimension. Here, we explore this idea in the context of

solidity. The typical vocabulary for young English-speaking children has many words for clearly solid things like bricks (63% of early nouns) and clearly non-solid things like sand (24% of early nouns, see Samuelson & Smith, 1999). But the world also is made up of things that are somewhere between bricks and sand, things such as pillows, play-dough, and pudding, falling at different points along this continuum (Samuelson & Smith, Spelke, Wynn CITE). In early vocabularies, children only know the names of a few such things (6% of the first 312 nouns typically learned by children learning English, see Colunga, 2002). If presented with such things of “intermediate” solidity, how would the cognitive system resolve the ambiguity? There are a number of possibilities: For example, if the current problem is just too far from known regularities, it could show a kind of brittleness and just not generalize sensibly at all. Or it could use one solution (e.g., shape) as the default solution. Alternatively, it could interpolate, forming novel but coherent generalizations that fall between known solutions.

Insert Figure 3 about here

In a series of simulations, we have examined how the model of nested influences on attentional resolves this interpolation problem. In one simulation, a model was created on the basis of the words for solids and non-solids that children know. The question is, how will the model respond to items that are in-between the two extremes of solidity? Thus the model was tested with extreme solidity values at 0.95 (solid) and 0.05 (non-solid) as well as with values in between (0.25, 0.50, 0.75). The attention allocation as calculated by the developmental scale of the model, based on these regularities is shown in Figure 3. The proportion of shape choices as predicted by the model for a forced choice task between a shape match and a material match are shown in Figure 4 for

each of these values of solidity. The model predicts that children will *gradually* shift attention to shape as perceptual information changes incrementally from that indicating a solid to a nonsolid thing (Figure 4). Thus, this simulation suggests a graded adaptation between two known clusters of regularities.

Insert Figure 4 about here

Behavioral tests of young children show a similar pattern. In this experiment, new words were taught to 30-36-month-old children in a NNG using exemplars at 4 degrees of solidity: 1) rigid, does not change shape when pressed, for example a brick, 2) dough, changes shape when pressed, but doesn't take shape of its container, for example playdough, 3) "goop", viscous material that flows when touched and takes shape of its container and is contiguous, for example pudding, 4) powder - takes shape of its container, but is not contiguous, for example rice (Colunga, under review).

All the shapes and materials used in the experiment were novel to the children. Each child saw one exemplar at each of the four levels of solidity, and was told its unique name, "Look at the dax." The child was then asked to select from a choice of a material match or a shape match to the exemplar, both at the same level of solidity as the exemplar, "Where is the dax here?"

The results show that children's attention to shape and material varies depending on the degree of solidity –on average, the more solid the exemplar, the more shape match choices will be made; the more non-solid the exemplar, the more material match choices will be made (Fig 5) These results are not easily explainable by traditional views that explain the stabilities in cognition in terms of internal knowledge structures about two categorically distinct kinds (e.g., objects and substances, see Soja et al, 1992). A simple

rule –shape if solid, material if nonsolid -- cannot by itself explain the data. Instead, the results fit the idea that children's generalizations are the product of graded *processes* and emergent *in the task*. The results also provide specific support for the model, given the statistics that characterize the nouns that young children know, generalizations are graded from more solid to less solid materials.

Insert Figure 5 about here

What if the statistics were different? One value of this model –and of conceptualizing the problem of inventive generalization as interpolation – is that it provides the means to ask systematic questions about how different kinds of regularities in the learner's environment might interact with the specific tasks. We can change the statistical regularities fed into the system (making nonsolid categories more frequent than solid ones, making them equal, or enhancing the lopsided nature of the regularities even more by increasing the number of solid categories over nonsolid ones). One could mimic these perturbations in the statistics in training experiments with children and thus test the model in greater detail than we have done here (see Samuelson, 2004, as an example). Thus, the present work provides an example of dynamic systems ideas may go beyond metaphor (Thelen, 1989) to make specific testable (and mechanistic) models of developmental process.

Extrapolation

Extrapolation is a second way that cognitive systems extend what they know and make novel generalizations. Consider, the case of animacy. Past research on children's novel noun generalizations has shown that children tend to generalize names for things with one or two properties –eyes or limbs– suggestive of animacy by attending to

multiple similarities and particularly to texture as well as shape, often requiring members of the same category to match on both dimensions, and thus generalizing more conservatively, extending the name to fewer new instances. What happens if one exaggerates animacy properties? Is animacy like pregnancy (you either are or you are not)? Or, in the human conceptual system, is “animacy” a matter of degree? Past research, in both linguistics and in children’s novel noun generalizations (cite Lucy, 1992; Colunga & Smith, 2005; Imai and Gentner, 2002; Yoshida & Smith, 2005) suggests that animacy might well be a matter of degree. Arguably, intelligent machines and some plants are close to the boundary and, again, arguably, the degree of “animacy” may be seen as increasing from slime-molds to slugs to mammals to people.

In the simulations, degree of animacy –in the sense of number or intensity of cues characteristic of animals – was manipulated to explore extrapolation. The model was given a novel problem that did not sit between known problems but instead presented a case that is outside the bounds of past experience. To simulate this, the model (and the 3 year old children in the behavioral experiments) were given exemplars and test objects that had the feature “eyes” outside of the regular values. Specifically, the model was given a value of 1.5 for eyes, putting it clearly outside of its (0.1) range; for children, this was done by increasing the number of eyes. In young children’s experiences, instances of animal categories have two eyes. To present an exaggerated version of this regularity, in the Many Eyes condition, children were shown things with five eyes, arranged in a symmetric cluster, as shown in Figure 6a. Note, increasing the number of eyes is not equivalent to increasing “animacy”. Instead, it is increasing (perhaps) the intensity of one cue that is associated with animal categories and attention to multiple similarities.

Thus, the items given to the children and to the model, are not “super-animate”, or particularly good or idealized exemplars of animates. Quite the opposite, they are odd and outside of that which has been experienced. But critically, these exemplars with exaggerated cues stand in a very specific relation to known exemplars; they present an exaggerated value on one highly predictive property.

Insert Figure 6 about here

There are several behavioral outcomes that could be informative about the underlying system. First, if the pattern of generalization shown by children when given animacy cues is about a fixed and represented concept of animacy, a thing with five eyes could be a poor example of an animate thing because of its dissimilarity to known instances. Alternatively, if there is a higher order rule that is used to define animates (e.g., “things that can see”), it may be a perfectly fine example of an animate. In the first case, we would expect the generalization pattern to be unlike that shown for things with two eyes, with children perhaps resorting to some default strategy such as shape. In the second case, we would expect the generalization pattern to be just like that for things with two eyes. However, if attention to the different dimensions is being soft-assembled based on the past history of the predictiveness of each cue, and underlying processes through which the strength of activation of the cue is directly related to the strength of attention to the predicted dimension, then an exaggerated cue might lead to patterns of generalizations that also exaggerate the pattern typically exhibited by young children for depictions of animates.

What might such an exaggeration of attentional cues look like in children’s novel noun generalizations? Two hallmarks of children’s generalizations of names for things

with eyes is (1) increased conservatism such that the category name is extended only to things highly similar on all varying dimensions and (2) a particular increase in attention to texture. Accordingly, extrapolation in this case might mean smaller categories as well as more generalizations on the basis of matches in texture.

Extrapolation was simulated in the model by giving it a value in the eyes dimension that was outside of its established range, 1.5. The model's predictions and children's data are shown in Figure 7, for solid things (solidity =.95) with No Eyes (eyes=.05), Two Eyes (eyes=.95), and Many Eyes (eyes=1.5.) Overall, model and children are more likely to say "no" to test items in the Many Eyes than in the Two Eyes or No Eyes conditions. More specifically, the model predicts greater attention to texture as the value of Eyes increases, this translates into both a decrease in "yes" responses to the shape only matches, and an increase in the number of "yes" responses to texture matches. That is, when given a value of eyes outside of its regular range, both 3-year-olds and the model generalize the name of the exemplar more conservatively, weighing texture information more heavily.

Insert Figure 7 about here

Importantly, when generalizing names of things with five eyes, children do not do the same thing as when they generalize names of things with 2 eyes. They do not seem to apply some sort of abstract rule that applies to all things with animacy features. They also are not just confused or resorting to a default strategy. They are systematically extending what they have gleaned from the regularities in the input to this novel situation with attentional tendencies exaggerated in the context of exaggerated perceptual cues. This

kind of result is explainable by real-time attentional processes in which the strength of a cue is directly related to the strength of activation of attentional biases.

Although this manipulation of exaggerated cues is clearly informative about the nature of the underlying system, one might also ask whether a system that responds in this way is also useful in children's lives. Although 5-eyed things are not something that will ever be encountered, different kinds of animates do differ in the strength of predictive features that they present. For example, categories of furry and four-legged animals are very frequent; scaly or 8-legged are comparatively infrequent. Yet the cognitive system described here would nonetheless demonstrate a coherent attentional bias. Clearly this is to the benefit of children's learning and to adapt behavior in novel situations. Further, children's generalizations appear based on an attentional system in which the strength of attentional biases is determined by the (combined) strength of cues. Knowledge that is embedded in graded processes (rather than in fixed representations) will lead naturally to graded behavior, and to extrapolation. Thus, the results strongly suggest that the system extends information through attentional processes rather than through higher-order conceptual rules.

Integration

Interpolating and extrapolating are straightforward mathematical operations and apparently also straightforward cognitive ones. A more challenging test of the model is integration. Can the cognitive system generate a novel and coherent response to a new instance that is not systematically related to any single known category? For this simulation, we combined features characteristic (nonsolidity and eyes) of two ontological kinds that would seem far apart in conceptual space, creating a substance with animacy

features. That is, in this simulation, we tested the model in a task in which it was the *combination* of features that was novel –a viscous solid (like shaving cream) with eyes. If children’s generalization of words proceeds from some structural representations or fixed concepts, such a novel instance would seem unlikely to activate that knowledge, because such an exemplar should be quite distant from prototypical examples of either kind. This line of reasoning predicts haphazard or perhaps no generalizations at all. Alternatively, the fixed representations might predict a coherent response on the following reasoning: an eyed-substance could lead to a coherent response if the system decided that it was either an animate or substance. The children would seem to have to construe the stimulus as a member of an ontological kind, and then decide (based on their knowledge about this kind in particular) which dimensions are relevant and which are not. From this view, two patterns thus seem possible: construe the eyed substance as an animate and attend to multiple similarities including texture or they construe it as a substance and attend to material.

The more radical possibility is that there are no fixed representations at all. If children’s attentional solution is the integrated blend of multiple processes soft-assembled in the task, then novel –intermediate but coherent – solutions may emerge. The specific solution will, of course, depend on the relative frequencies of lexical categories for solid and nonsolid things, and for things with and without eyes in the learner’s history of experiences. Cues should be weighted depending on their past predictability with attention deployed to the different dimensions accordingly.

The model was tested on items marked as non-solid (solidity = .05) that either had eyes (eyes = .95) or not (eyes = .05) in a yes/no task on items matching on shape, texture,

material, and combinations of these features. The predictions of the model are shown in Figure 8 for each test item. First, overall, the model predicts more “yes” responses to substances with eyes, than to substances with no eyes. That is, in the case of non-solid substances, adding eyes makes the system *less* conservative. Second, the model predicts greater attention to shape in the Eyes than in the NoEyes condition; a test item matching in shape alone is considered more similar to the exemplar when exemplar and test items both have eyes than when they do not. This is not the case for matches in material and texture, which are considered equally acceptable whether the substances have eyes or not. These predictions, not obvious from intuitions about eyes and nonsolidity, arise from the *interaction* of the cues in the allocation of attention.

Insert Figure 8 about here

Three-year-olds were tested in an equivalent task. The stimuli used was made out of materials such as colored sand, frosting, and hair gel; the texture of the substances was manipulated by making them smooth, or with holes, spikes or glitter (Figure 6b). The results of this experiment are a qualitative match to the model’s predictions, as shown in Figure 9. Three-year-olds in this task will extend the name of the exemplar to those matching in shape and surface properties in the NoEyes condition, but will liberally generalize it to anything that matches in shape when eyes are added to the exemplars and test items. This pattern of generalization is unlike that characterized as indicative of substance or animate kinds, and yet it is robust across children (8/12 children in the Eyes condition generalize by shape alone, compared to 3/12 in the NoEyes condition).

Insert Figure 9 about here

The results of this study show that children are able to show consistent, robust behavior in a novel situation that that is neither an in-between case from what has been learned and can be interpolated, nor an exaggerated version of what has been experienced and can be extrapolated. Moreover, the results do not quite fit with the classic view of cognition or with the concept-based accounts of novel noun generalizations. Three-year-olds do not appear to construe non-solids with eyes as either members of the substance kind or of the animate kind. Instead, the system finds another solution –broader generalizations and increased attention to shape. Further, a model that encompasses attentional learning over nested time scales, that finds attentional solutions emergent in the regularities of past experience, its history of attentional deployment, and the current attentional demands elegantly predicts the otherwise unexpected result.

Inventive Cognition

Children's performance in novel noun generalization tasks has attracted considerable interest from researchers studying language learning from a variety of theoretical perspectives. In part this is because, children's smart performances in this task offer insight into how it is that they learn words so rapidly. Given a novel thing, they attend coherently and selectively to certain similarities. In part the interest in this task also arises because children's performances suggest knowledge about different ontological kinds –about animals, about objects, about substances. Children's kind specific generalizations suggest stabilities in the cognitive system, a capturing of the regularity that animals within a single basic level category are similar in many ways, that artifacts within a similar basic level category are similar in shape and that substances within a single basic level category are similar in material. Children's novel noun

generalizations are also of interest because they are inherently *inventive*. In this task, children form brand new categories given a single *novel* thing labeled with a *novel* name. D-ALA explains the stability and the inventiveness through processes that include neither fixed concepts of ontological kinds nor fixed rules about the kinds of similarities that matter for different kinds. Instead, the model explains the stabilities and the inventiveness through attentional processes that are influenced by many different learned cues and by the immediate sensory input, but that must converge on a single behavioral outcome. Because of this, the system behaves –and attends –coherently even when confronted with highly novel cases.

This is a system that through its own self-organizing processes builds the knowledge that makes for stable categorization patterns and also dynamically finds new solutions, a system that can interpolate, extrapolate and blend information and thus generate a coherent attentional response even in what might seem the strangest situations. It is this dynamic creativity that makes children's word learning so robust and so smart.

References

- Bloom, P. (1996). Intention, history and artifact concepts. Cognition, 60(1), 1–29.
- Bloom, P. (2000). How children learn the meanings of words. Cambridge, MA: MIT Press.
- Booth, A. & Waxman, S. (2002). Word learning is 'smart': Evidence that conceptual information affects preschoolers' extension of novel words. Cognition, 84, B11-B22.
- Booth, A. & Waxman, S. (2003) Mapping words to the world in infancy: Infants' expectations for count nouns and adjectives. Journal of Cognition & Development, 4, 357-381.
- Burger, J., Prasada, S. Greenhill, A. (1997). Quantification of solid and nonsolid entities: the role of perceived arbitrariness of structure. Proceedings of the Boston University Conference on Child Language, Cascadilla Press.
- Carey, S (
- Chiang, W. & Wynn, K. (2000). Infants' tracking of objects and collections. Cognition, 77, 169-195
- Colunga, E., Smith, L. B. (2003) The emergence of abstract ideas: evidence from networks and babies. *Philosophical Transactions by the Royal Society B*. Theme Issue: “The abstraction paths: from experience to concept”. L. Saitta (Ed) 358(1435) pp. 1205-1214.
- Colunga, E. & Smith, L. B. (2004). Dumb Mechanisms make Smart Concepts. *Proceedings of the Annual Conference of the Cognitive Science Society*, 26. pp. 239—244.
- Colunga, E., Smith, L. B. (2005) From the Lexicon to Expectations About Kinds: A Role

for Associative Learning. *Psychological Review*, 112(2) pp.347-382.

Colunga, E. (2006). The effect of priming on preschooler's extensions of novel words :

how far can “dumb” processes go? In *Proceedings of the 30th Annual Boston University Conference on Language Development*.

Colunga, E., & Smith, L. B. (in press) Knowledge embedded in process: The self-organization of skilled noun learning. *Developmental Science*.

Dickinson, D. K. (1988) Learning names for materials: Factors constraining and limiting hypotheses about word meaning. *Cognitive Development*, 3, 15-35.

Diesendruck, G. & Bloom, P. (2003) How specific is the shape bias? *Child Development*, 74, 168-178.

Feigenson, L., Carey, S. Spelke, E. (2002). Infants' discrimination of number vs. continuous extent. *Cognitive Psychology*, 44(1), 33–66.

Gathercole, V. & Min, H. (1997). Word meaning biases or language-specific effects? Evidence from English, Spanish and Korean. *First Language*, 17(49 Pt 1), 31-56.

Gathercole, V. C. M., Cramer, L., Somerville, S. C. Haar, M. J. O. D. (1995). Ontological categories and function: acquisition of new names. *Cognitive Development*, 10, 225–251.

Gathercole, V., Thomas, E. M., & Ecans, D. (2000) What's in a noun? Welsh-, English-, and Spanish-speaking children see it differently. *First Language*, 20(58 Pt 1), 55-90.

Gelman, S. A. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology*, 20, 65–95.

Gelman, S. A. & Coley, J. D. (1991). Language and categorization: The acquisition of natural kind terms. In Gelman, S. A. & Byrnes, J. P. (Eds) *Perspectives on language*

and thought: Interrelations in development. New York, NY: Cambridge University Press.

- Gelman, S. A. Bloom, P. (2000). Young children are sensitive to how an object was created when deciding what to name it. Cognition, 76(2), 91–103.
- Golinkoff, R., Mervis, C. B., & Hirsh-Pasek, K (1994). Early object labels: The case for a developmental lexical principles framework. Journal of Child Language, 21(1), 125-155.
- Golinkoff, R., Hirsh-Pasek, K., Bloom, L., Smith, L. B., Woodward, A., Akhtar, N., Tomasello, M., & Hollich, G. (2000) Becoming a word learner: A debate on lexical acquisition. London: Oxford University Press.
- Hall, D. G. (1996). Naming solids and non-solids: children's default construals. Cognitive Development, 11, 229–264.
- Imai, M. Gentner, D. (1997). A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. Cognition, 62, 169–200.
- Huntley-Fenner, G., Carey, S. & Solimando, A. (2002) Objects are individuals but stuff doesn't count: Perceived rigidity and cohesiveness influence infants' representations of small groups of discrete entities. Cognition, 85, 203-221.
- Jones, S. S., Smith, L. B. Landau, B. (1991). Object properties and knowledge in early lexical learning. Child Development, 62, 499–516.
- Katz, N., Baker, E., & Macnamara, J. (1974). What's in a name? A study of how children learn common and proper names. *Child Development*. 45(2), pp. 469-473
- Keil, F. C. (1989). Concepts, kinds, and cognitive development. Cambridge, MA: MIT Press.

- Keil, F. (1994) The birth and nurturance of concepts by domains: The origins of concepts of living things. In Hirschfeld, L. Gelman, S. (Eds.) Mapping the mind: Domain specificity in cognition and culture (pp. 234-254). New York, NY, US: Cambridge University Press.
- Kemler Nelson, D. G., Frankenfield, A., Morris, C. & Blair, E. (2000). Young children's use of functional information to categorize artifacts: Three factors that matter. Cognition, 77, 133-168.
- Kobayashi, H. (1997). The role of actions in making inferences about the shape and material of solid objects among 2-year-old children. Cognition, 63, 251–269.
- Landau, B., Smith, L. B. Jones, S. S. (1988). The importance of shape in early lexical learning. Cognitive Development, 3, 299–321.
- Landau, B., Smith, L. B. Jones, S. S. (1992). Syntactic context and the shape bias in children's and adults' lexical learning. Journal of Memory and Language, 31, 807–825.
- Lavin, T. & Hall, D. G. (2001) Domain effects in lexical development: Learning words for foods and toys. Cognitive Development, 16, 929-950.
- Lucy, J. A. (1992). Grammatical categories and cognition: a case study of the linguistic relativity hypothesis. Cambridge: Cambridge University Press.
- Lucy, J. A. Gaskins, S. (2001). Grammatical categories and the development of classification preferences: a comparative approach. In M. Bowerman S. C. Levinson (Eds.), Language acquisition and conceptual development, 257–283, Cambridge: Cambridge University Press.

- Mandler, J. M. (1992). The foundations of conceptual thought in infancy. *Cognitive Development*, 7(3), Jul-Sep 1992, pp. 273-285
- Markman, E. M. (1989). Categorization and naming in children: Problems of induction. MIT Series in learning, development, and conceptual change, Cambridge:MA MIT Press.
- McPherson, L. (1991) A little goes a long way: Evidence for a perceptual basis of learning for the noun categories COUNT and MASS. *Journal of Child Language*, 18, 315-338.
- Prasada, S., Ferenz, K. & Haskell, T. (2002). Conceiving of entities as objects and as stuff. *Cognition*, 83(2), 141–165.
- Samuelson, L. & Smith, L. B. (1999). Early noun vocabularies: Do ontology, category structure and syntax correspond? *Cognition*, 73, 1–33.
- Samuelson, L. K. & Smith, L. B. (2000). Children's attention to rigid and deformable shape in naming and non-naming tasks. *Child Development*, 71(6), 1555–1570.
- Samuelson, L. (2002). Statistical regularities in vocabulary guide language acquisition in connectionist models and 15-20-month-olds. *Developmental Psychology*, 38, 1016-1037.
- Sandhofer, C. M., Smith, L. B. & Luo, J. (2001). Counting nouns and verbs in the input: Differential frequencies, different kind of learning? *Journal of Child Language*, 27, 561–585.
- Smith, L. B. (1995). Self-organizing processes in learning to learn words: Development is not induction. In C. Nelson (Ed.) Basic and applied perspectives on learning, cognition, and development. 28, (1–32). Marwah, NJ: Lawrence Erlbaum

Associates.

Smith, L. B., Jones, S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L., (2002).

Object name learning provides on-the-job training for attention. Psychological Science, 13(1), 13–19.

Smith, L. B., Colunga, E. & Yoshida, H. (2003). Making an ontology: Cross-linguistic evidence. In Rakison, D. & Oakes, L. (Eds.) Early category and concept development: Making sense of the blooming, buzzing confusion (pp. 275-302).

London, Oxford University Press.

Smith, L. B. (2005). Cognition as a dynamic system: Principles from embodiment.

Developmental Review, 25, 278–298.

Soja, N. N. (1992). Inferences about the meaning of nouns; the relationship between perception and syntax. Cognitive Development, 7, 29–45.

Soja, N. N., Carey, S. & Spelke, E. (1992). Perception, ontology, and word meaning.

Cognition, 45, 101–107.

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.

Soja, N. N. (1994). Evidence for a distinct kind of noun. Cognition, 51(3), 267–284.

Spelke, E. & Hespous, S. (2001). Continuity, competence, and the object concept. In

E. Dupoux (Ed.) Language, brain, and cognitive development: Essays in honor of Jacques Mehler. (325–340). Cambridge, MA: The MIT Press.

Spelke, E. S. 1990. Origins of visual knowledge. In D. N. Osherson S. M. Kosslyn (Eds),

Visual cognition and action: An invitation to cognitive science (99–127). Cambridge,

MA: MIT Press.

Subrahmanyam, K., Landau, B., & Gelman, R. (1999) Shape, material, and syntax:

Interacting forces in children's learning in novel words for objects and substances.

Language & Cognitive Processes, 14, 249-281.

Thelen, E., 1989. Self-organization in developmental processes: Can systems approaches

work? In E. Thelen & M. R. Gunnar (Eds), Systems and development: The

Minnesota symposia on child psychology (77-117). Hillsdale, NJ: Lawrence Erlbaum Associates.

Thelen, E. & Smith, L. B. (1994) A dynamic systems approach to the development of cognition and action. MIT Press.

Waxman, S.R., & Markow, D.B. (1995). Words as invitations to form categories.

Cognitive Psychology, 29(3), 257–302.

Xu, F. (1997). From lot's wife to a pillar of salt: evidence that physical object is a sortal concept. Mind and Language, 12, 365–392.

Yoshida, H. & Smith, L.B. (2001) Early noun lexicons in English and Japanese.

Cognition, 82(2), B63-B74.

Yoshida, H. & Smith, L.B. (2003a) Shifting ontological boundaries: How Japanese- and English-speaking children generalize names for animals and artifacts. Developmental

Science, 6, 1-17.

Yoshida, H. & Smith, L.B. (2003b) Correlations, concepts and cross-linguistic differences. Developmental Science, 6, 30-34.

Yoshida, H. & Smith, L. B. (2003) Known and novel noun extensions: Attention at two levels of abstraction. Child Development, 564-577.

Yoshida, H. & Smith, L.B. (2005) Linguistic cues enhance the learning of perceptual cues. Psychological Science, 16(2), 90-95.

Figure Captions

Figure 1. Three kinds of entities studies with the novel noun generalization task, solid objects, objects with eyes, and non-solid substances.

Figure 2. The three nested timescales of the Dynamic-Attentional Learning Account: Developmental time, task time, and decision/action time.

Figure 3. Attention allocation as predicted by the developmental timescale of the model.

Figure 4. Predicted probability of shape choice for different solidity levels.

Figure 5. Mean proportion of shape choices for 3-year-olds in the interpolation study, for each level of solidity.

Figure 6. Novel stimuli for extrapolation and integration experiments. a) solid objects with five eyes and b) non-solid substances with eyes.

Figure 7. Model predictions and 3-year-old's proportion of shape choices for solid things with no eyes (Eyes = 0.05), two eyes (Eyes = 0.95), or five eyes (Eyes = 1.5).

Figure 8. Predicted probability of “yes” responses for each test item (Solidity = 0.05), with and without eyes (Eyes = 0.05 and .95 respectively.)

Figure 9. Three-year-olds' proportion of “yes” responses for each of the test items, for non-solids with and without eyes.

a)

b)