Knowledge embedded in process:
The self-organization of skilled noun learning

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

Young children's skilled generalization of newly learned nouns to new instances has become the battleground for two very different approaches to cognition. This debate is a proxy for a larger dispute in cognitive science and cognitive development: cognition as rule-like amodal propositions, on the one hand, or as embodied, modal, and dynamic processes on the other. After a brief consideration of this theoretical back drop, we turn to the specific task set before us: an overview of the Attentional Learning Account (ALA) of children's novel noun generalizations, the constrained set of experimental results to be explained, and our explanation of them. We conclude with a consideration of what all of this implies for a theory of cognitive development.

Knowledge embedded in process: The self-organization of skilled noun learning

In the course of science, there are phenomena that temporarily (for years and even decades) seem to attract more than their fair share of attention. Young children's skilled generalization of a newly learned noun to new instances, a phenomenon (and experimental paradigm) first introduced by Katz, Baker, and Macnamara in 1974, is one of these cases. Children are so skilled and systematic in generalizing newly learned names of things that this basic task is used to study a wide variety of issues, including category formation, syntactic development, object recognition, social cognition, and attention (e.g., Prasada and Haskell, 2002; Hall, Quantz, and Personage, 2000; Soja, 1992; Baldwin & Baird, 2001). Because of the broad reach of the method, children's novel noun generalizations have also become the battleground for two very different approaches to cognitive development. The editors of this special issue propose to advance the field by asking researchers associated with the two different sides to consider and explain, each from their own perspective, a constrained set of experimental results, and to answer the question, again each from their own perspective, of what counts as an explanation of cognitive development.

We begin, not with what counts as a theory of development, but with what counts as cognition. In contemporary cognitive science, there is a sharp divide between two all encompassing views of cognition as rule-like amodal propositions, on the one hand, or as embodied, modal, and dynamic processes on the other. How children generalize names for things (and the perception-conception debate embedded within it) is the proxy for this larger dispute in the developmental literature. After a brief consideration of this theoretical back drop, we turn to the specific task set before us: an overview of the

Attentional Learning Account (ALA) of children's novel noun generalizations, the four findings to be explained, and our explanation of them. We conclude with a consideration of what all of this implies for a theory of cognitive development.

What counts as cognition

The traditional view divides mental life into discrete steps of “sense-think-act.” Cognition, by definition, is about the “think” part, the knowledge that *mediates* between perceiving and acting. Knowledge, in this view is amodal and propositional, consisting of relatively fixed representations. Knowledge is thus profoundly different in kind, and theoretically separable, from the real time processes of perceiving, remembering, attending, and acting.

The main idea on the opposing side is that knowledge has no existence separate from process, but is instead *embedded* in, *distributed* across, and thus *inseparable* from the real time processes of perceiving, remembering, attending, and acting (see Samuelson & Smith, 2000). In this view, knowledge *just is* these processes bound to each other *and to the world through perception and action* in real time (see, for example, O'Regan & Noe, 2001; Samuelson & Smith, 2000) with no fixed and segregated representation of anything (see also, Barsalou, 1993; 2003; Smith & Jones, 1993; Port & van Gelder, 1995).

In the literature beyond the study of children's novel noun generalizations, the core issues relevant to these two approaches all concern the special properties of propositional representations such as compositionality, rules and variables, evidence (or non-evidence) for these properties in human cognition, and the ability of process models to successfully mimic these processes without propositional representations. These issues have not been so central in the literature on children's novel noun generalizations. Instead, the

discussion has been ill-defined, taking its form from within the sense-think-act tradition, and more specifically from Piaget's theory of developmental progression from sensory-motor (sense-act) to representational (sense-think-act) thought, wherein unitary proposition-like symbols intervene between perception and action. Within this definition of the debate, the empirical question has been defined straightforwardly as whether conceptual representations intervene between perceiving and acting in creating children's generalizations. For example, if a child is shown an object that has properties that make it look like an artifact, but if they are told that it can "be happy," do they go by the perceptual appearance or do they reason from a conceptual understanding about the kinds of things that can be happy?

The problem with this construal of the debate --- sense-act versus sense-think-act --- is that newer ideas of embodied and embedded cognition do not fall straightforwardly on either side of the divide. Theories about embodied cognition do share aspects with Piaget's ideas about sensory-motor thought (see Thelen and Smith, 1994; Clark, 2001; Barsalou, 1999; Brooks, 1991; Pfeifer & Scheier, 1999), but they are also fundamentally different in that they propose that even clearly abstract forms of thought (in both children and adults) emerge from the very same processes that give rise to more obviously perceptually-based forms of thought (see Barsalou, 1999; Dale & Spivey, 2005; Lakoff, 1994; Vittorio & Lakoff, 2005; Colunga & Smith, 2003). This view does not deny conception but instead says it is fundamentally different in form from that presumed by the sense-think-act tradition. In this view, conception is not propositional and not different in kind from perceiving, attending, remembering and acting, but is instead continuous with and made in those very processes.

Thus, the two sides of the current debate about children's novel noun generalizations are often at cross-purposes with contemporary ideas about embedded/embodied cognition, confused with Piaget's definition of sensory-motor versus representational thought. The empirical question is thus confused with the kind of experimental tasks that Piaget used to contrast his view of sensory-motor thought and his view of symbolic thought:

Perceptual (non-conceptual) processes are operationally defined as dependent on the immediate sensory input, whereas representational thought (conception) is operationally defined as dependent on words (e.g., Waxman & Markow, 1998; Soja, 1992, Gelman & Bloom, 2000), on remembered events such as actions or "hidden" properties that were perceived several seconds earlier (e.g., Kemler Nelson, Russell, Duke & Jones, 2000; Kobayashi, 1997), on perceptible but subtle properties of the things rather than overall similarity (e.g., Keil, 1994; Gelman & Koenig, 2003), or on the longer term history of the learner with the specific instances (e.g., Mandler, 1992, Gelman, 1988).

These operational definitions are contestable on several grounds (Ahn & Luhman, 2005). Moreover, they do not line-up at all with the embedded cognition approach, which makes no such distinction between perceptual and conceptual processes at all (Smith & Gasser, 2005; Samuelson & Smith, 2000; Smith & Jones, 1998). By the embedded cognition view, all of the results will be explainable without recourse to unitary or proposition-like representations but instead will be explainable in processes of attention, memory, learning, perception, and action. Thus in the embedded-cognition view, children's understanding of hidden properties, their use of transient events in making decisions, their long-term knowledge of the regularities in the world are all

grounded in the very processes that also underlie perceiving, remembering, attending, and acting.

The incommensurate nature of the two views on what counts as cognition leads to the bizarre outcome that proponents of the two sides can conduct nearly identical experiments and each see the same patterns of results as strongly supporting their own position (compare Cimpian & Markman, 2005, to Yoshida & Smith, 2003a; Booth & Waxman, 2002 to Yoshida & Smith, 2003b; and Diesendruck & Bloom, 2003 to Samuelson & Smith, 2000).

The Attentional Learning Account

The Attentional Learning Account (ALA) of children's novel noun generalizations is firmly in the embedded/embodied cognition camp. It specifically seeks to explain an expansive set of data concerning developmental changes in early noun learning, including the accelerating pace of new noun acquisitions during the period between 12 and 30 months (Fenson, Dale, Reznick, Bates, & Thal, 1994), the developmental emergence of systematic biases in the generalization of names for animals versus objects versus substances (Jones & Smith, 2002; Soja, Carey, & Spelke, 1991, Landau, Smith & Jones, 1988), cross-linguistic differences in these biases (Imai & Gentner, 1997; Yoshida & Smith, 2003b), and the lack of these biases in children with delayed language acquisition (Jones, 2003; Jones & Smith, 2005).

The main idea is that attentional learning is an ongoing continuous process such that attention is dynamically shifted *in the moment* to properties, features and dimensions that have *historically* been relevant for the task context. The mechanism of change is a simple correlational learning system which, by internalizing the systematic patterns

(statistical relations) present in the environment, instantiates much intelligence. This kind of ongoing, unconscious learning has been widely demonstrated in experimental psychology (Chun & Jiang, 1998, 1999, 2001, 2003; Krushke, 2001; Regier, 2005) and is well understood mechanistically and theoretically. There are three core claims relevant to applying these general cognitive processes of attentional learning to the developmental problem of early noun acquisitions:

(1) The learning environment presents correlations among linguistic devices, object properties, and perceptual category organization. Studies of statistical structure of the first 300 nouns (in English and in Japanese, Samuelson & Smith, 1999; Jones & Smith, 2002; Yoshida & Smith, 2001; Smith, Colunga & Yoshida, 2002; Colunga & Smith, 2005; and to a lesser degree, Mandarin, see, Sandhofer, Luo & Smith, 2001) show that artifacts tend to be rigid, angular, solid things in categories organized by shape, that animals tend to have features such as eyes, legs and heads, and to be in categories organized by multiple similarities, and that substances tend to be nonsolid and in categories organized by material. Further, these statistical regularities among perceptual properties and perceptual category organizations also correlate with a variety of words (beyond the specific names of specific things) such as determiners, classifiers, and verbs (Samuelson & Smith, 1999, Yoshida & Smith, 2001).

(2) Children learn the statistical regularities that characterize individual categories and *the whole system of acquired categories*. Young children learn names for specific categories; as a consequence, they will learn, as first-order generalizations, the *many* specific properties relevant to those specific categories (Yoshida & Smith, 1999; Rosch & Mervis, 1975; Samuelson & Smith, 1999; McRae, de Sa, & Seidenberg, 1997).

All these properties, jointly and alone, depending on the systematicity of their correlations, have the potential to dynamically shift attention. The key – and more powerful – claim of ALA is that children do not just learn these first-order correlations but also learn higher (second, third) order correlations that arise *over* the learned correlational patterns of *many different categories*, that solid things with angular shapes tend to be categorized by shape, that things with eyes tend to be categorized by multiple similarities, that the determiners “a” or the word “another” tend to be correlated with things in categories organized by shape, that the subjects of verbs such “eat” or “loves” tend to have eyes and be in categories organized by multiple similarities. These higher order correlations (correlations across systems of categories) enable dynamic *intelligent* shifts in attention to the appropriate kinds of similarities even given novel things and novel names, creating highly abstract knowledge that approximates a variablized rule (see Colunga & Smith, 2005). These higher order regularities reflect the statistical regularities not of any one noun category but across a system of categories and as a consequence are highly useful in the *first stage* of learning a new object name, by constraining attention to similarities statistically likely to be relevant.

(3) Children’s learning of the statistical regularities and their *application of that learning* in the task of generalizing a name to a new instance are mechanistically realized through learned associations that yield *contextually cued dynamic shifts in attention*. ALA proposes that children’s attention is automatically directed (without deliberative thought) to similarities that have been systematically relevant in those linguistic and perceptual contexts in the child’s past. The core mechanism, then, is the *top-down* control of attention in the moment by past experience (see especially, Smith,

2001, Yoshida & Smith, 2005). This is a potentially powerful learning mechanism in several ways: (1) it is exquisitely tied to and integrates multiple (perceptual *and* linguistic) contextual cues in the moment, and is therefore always graded and task dependent; (2) it enables the learner to attend to (and construe) the same perceptual object in different ways depending on context (e.g., with count syntax, a muffin is construed as an object of a particular shape, but with mass syntax the same perceptual muffin can be seen as a substance of a particular material); and (3) through it, attention and learning in the moment are strongly guided by the history of regularities in the learner's past.

The data to be explained

The four assigned papers (Booth, Waxman & Hwang, 2005; Diesendruck & Bloom, 2003; Smith, Jones, Landau, Gershkoff-Stowe, Samuelson, 2002; Samuelson, 2002) all concern variants of the novel noun generalization task. In the prototypic version, children are presented with a novel exemplar object, and in some conditions told its name and/or facts about it, then are shown novel test objects and asked (in various different ways) which of these is in the same category. The main results (as we see them) that need to be explained are these:

(1) *Booth, Waxman & Huang (BWH)*. The experiments in this paper show that 20- and 30-month-old children's novel noun generalizations are influenced by the words that experimenters say when they talk about the exemplar. For example, saying that the object is "happy" changes children's name extensions – so that objects with perceptual properties commonly associated with artifacts (angularity, no eyes, no legs) and categories organized by shape are treated by children as if they were animals; or more, specifically, names for these things are extended to new instances by texture as well as

shape. The pattern of generalizations for the 20-month-olds is much weaker overall than those for the 30-month-olds.

(2) *Diesendruck & Bloom, (D&B)*. The experiments in this paper show that 2- and 3-year-old children systematically categorize artifacts by shape even in non-naming tasks. Three-year-olds show a shape bias when asked to “Get the dax” but also when asked “to get one of the same kind” as well as when asked to generalize a hypothesized category-relevant property (it comes in a special box) or, to a lesser extent, a hypothesized category-irrelevant property (my uncle gave me this). Two year olds show a shape bias when asked to generalize a name (get the dax) or to indicate the “same kind.” These results thus indicate a shape bias for solid artifact-like things that becomes more pervasive across linguistic contexts with age.

(3) *Smith et al.(SJLGS)*. In these experiments, very young children (17-month-olds) are intensively taught (over an 8 week period) names for pairs of specifically shaped solid things that are alike in their shape. This training yields a generalized bias to extend the names of solid things (even very differently shaped solid things) by shape. This shape bias for solid angular things which is well-documented in the literature in older children is not evident in untrained 17-month-olds nor in 17-month-olds in control training regimens who were taught names for things that were not systematically alike in shape or who were taught un-named shape categories. Teaching children a generalized bias to name solid things by shape also accelerates real world noun learning causing a dramatic increase in vocabulary size over the experimental sessions for children in the Experimental training group but not for those in the Control groups.

(4) *Samuelson, 2002(S)*. Among early learned nouns, there is a stronger correlation among words associated with count nouns, solidity, angularity and shape based categories than there is among words associated with mass nouns, nonsolidity, less angular (or constructed) shapes, and material based categories. The count noun-solidity-shape correlation is stronger in two senses: there are many more of these kinds of nouns in early noun vocabularies and the reliability of linguistic cues (e.g., determiners) and solidity as predictors of the relevance of shape is greater than is the reliability of linguistic cues and nonsolidity as predictors of the relevance of material. Intensively teaching (over an 8 week period) very young (15 to 20 month old) children new categories that reflects the statistical properties of English creates a generalized shape bias (and accelerated vocabulary growth outside of the laboratory) but not a generalized material bias. Further, a formally instantiated model of ALA (as a neural network) when given the statistical regularities of early English vocabularies (mimicking the words the children know at the start of the experiment) and then given the experimental training regimen given to the children, closely simulated the children's performances in the generalization test.

An explanation in terms of learned correlations and attention

The general form of the explanation is outlined with reference to Figure 1. This figure shows three classes of input (linguistic context, labels, and perceptual properties) that may be correlated with the relevance of different kinds of similarities to a category decision. These different inputs, by the history of their past correlations with each other and with attention in category decision tasks serve as context cues, increasing and decreasing attentional weights to different kinds of similarities. The three kinds of inputs

listed are not the only kinds that children learn. Gestures, facial expressions, tones of voice, locations are all likely correlates and learned contextual cues. We concentrate on perceptual properties and linguistics cues because these are the contextual cues to attention studied in the four assigned papers.

By perceptual cues, we refer to the properties characteristic of specific things –having eyes, being angular, moving in a particular way, being nonsolid. Labels refer to the names, common nouns, given to individual things. Linguistic cues refer to the linguistic contexts in which those names and things occur and include determiners, pronouns, verbs and so forth. There are learnable statistical relations among all of these (see especially Yoshida & Smith, 2003a, 2003b). Perceptual cues correlate with each other such that things with eyes typically have mouths and move in a certain way and such that nonsolid things take a limited range of shapes. Linguistic cues correlate with each other and with specific names and with perceptual properties such that “eat” and “sleep” are associated with things with eyes and with the labels “dog” and “cat.”.

Finally, all these cues –individually and as clusters – predict the relevant relations for categorization (and for naming): things with eyes that “eat” and “sleep” are correlated with categories organized by multiple similarities, things with angular complex shapes that are “broken” and “made” are correlated with categories organized by shape, things that are nonsolid and “spill” are correlated with categories organized by material.

Past research shows that in large systems of correlations such as these, there is considerable latent structure pertinent to syntactic categories (Farkas & Li, 2001; Landauer & Dumais, 1997; Mintz, 2003; Monaghan, Chater & Christiansen, 2005), to taxonomic category organization (Rogers & McClelland, 2005; McRae, de Sa, &

Seidenberg, 1997), and, we believe, to children's understanding of different ontological kinds (Colunga & Smith, 2005; Yoshida and Smith, 2003b) and their intelligence in systematically generalizing nouns to new instances.

Explaining BWH

One of the powerful aspects of ALA is that attention is dynamically shifted in tasks, in the moment, as the consequence of the specific consortium of cues in the task. In support of this idea, we have conducted experiments nearly identical to BWH, and found, like them, that we can shift children's name extensions for the very same objects by the verbs used in conjunction with the named object (Yoshida & Smith, 2003a, 2003b). In their specific experiments, BWH provided children with many linguistic cues in competition with a few predictive perceptual cues, and the linguistic cues won out. By ALA, the relative strength of cues may be predicted *a priori* from their prevalence and reliability as predictors of categories. We (along with Hanako Yoshida) have begun analyses of large corpora of parent speech to children in order to make and experimentally test such fine-grained predictions. The results so far suggest that the class of words that correlate with nouns for animates and those that correlate with nouns for artifacts are distinct. In other experiments we have found that using these animacy- or artifact-correlated words prior to doing the noun generalization task, *even in the absence of the exemplar and without being used to refer to the exemplar*, shifts children's responses in the same way as BWH's vignettes (Colunga and Smith, 2004; Colunga, 2006). How does this account differ from BWH? BWH suggest that words (as well as properties such as eyes) activate unitary represented concepts of what it means to be an animate or an artifact. ALA suggests, in contrast, that this knowledge can be explained

without recourse to unitary concepts, that it can be explained as knowledge embedded in processes of attention and in the system of learned cues that organize attention in the real-time task of deciding whether or not a name applies to some thing.

Explaining D&B

Our explanation of D&B is similar to our explanation of BWH. D&B's main result is that 2- and 3-year-old children show a shape bias, even in non-naming tasks, for artifact-like things. In our work, we have also reported the shape bias 1) in non-naming tasks, in adults (Landau, Smith & Jones, 1988), 2) when the stimuli have highly complex shapes in 4 year olds (Sandhofer & Smith, 2004), and 3) in contexts in which very young children spontaneous name objects on their own (Samuelson & Smith, 2005). Attention to shape in these non-explicit naming tasks may be explained by the perceptual properties of the stimuli –angular artifact like shapes – that may cue attention to shape. In addition, the task context, and the specific words used in the task will also play a role. Indeed, we suspect that D&B's specific results across their various conditions might be readily modeled by the statistical properties of the learning environment. Specifically, “goes with” may be associated with thematic relations in children's experiences (socks go with shoes, milk goes with cookies) and thus not strongly push attention to shape, “same,” “one”, “this”, “gave”, “made in factory” and “special” may be more strongly associated with count nouns, basic level categories, and shape than with other properties. In this way, developmental differences would be explained by the children's learning the most pervasive and statistically reliable correlations before the less pervasive and less robust correlations. Although, at this point these ideas are speculative, they are directly testable and we have begun the relevant analyses of corpora of parent speech to children.

D&B designed their study to show that a shape bias emerges in many different contexts, not just naming. The underlying logic of their experiment and conclusion appears to be this: If many different task contexts and cues yield the same behavioral outcome (attention to shape), it must be because all these contexts and cues activate the same underlying concept of kind (see Keil, 1994). But this needs not be the case; constancy in a behavioral outcome does not mean a single constant cause on the inside (see Thelen and Smith 1994). The general processes of ALA will learn (and blend) a whole system of predictive cues and do so without a unitary intervening concept (see Yoshida and Smith, 2003b; Colunga & Smith, 2005). In brief, naming is just one cue and not a necessary one by ALA.

Still, we have suggested in a number of prior papers that learning object names –and the cues present in the act of naming a thing – may be a particularly powerful influence on attentional learning and attention in a task. This idea is based on our original finding that young children showed especially robust attention to shape in naming but not in non-naming tasks (Landau, Smith & Jones, 1988) and on additional findings that naming shifted attention to shape and away from other salient properties (Jones, Smith & Landau, 1992; Samuelson & Smith, 1999), that developmental increases in attention to overall shape were tightly linked to nominal vocabulary growth (Samuelson, 2002; also Gershkoff-Stow & Smith, 2004), that in training studies, teaching names taught a shape bias but teaching un-named shape categories did not (SJLGSS), and that in our formal models of the acquisition of the shape bias, learning names for things appeared to be computationally important to forming higher order generalizations (Colunga, 2001). None of this means that learning object names is necessary to the development of

contextually cued attention to shape but these results do fit the idea that learning object names may be a strong, and perhaps even special force on real world attentional learning.

Explaining SJLGSS

One way in which we have pursued the relation between real-world noun learning and the development of a shape bias in novel noun generalization tasks is through a series of training studies. These studies provide the strongest evidence for links between attention to shape and early noun learning: teaching very young children to attend to shape when naming things causes dramatic increases in the rate of new noun acquisitions beyond the experiment. There were a variety of different control conditions. Two critical ones were these: (1) Language control: Children were taught names for things in non-shape based categories and (2) Category control: Children were taught the same 4 shape-based categories as in the Experimental condition but they were not taught names for these categories. At the end of training, we tested children in two name extension tasks: we asked 1) whether the children could recognize new instances of the trained categories; and 2) whether after training, these children would extend a *novel* name for a newly-encountered *novel* thing by shape. We also measured children's productive vocabulary at the start and end of the experiment.

Here are the main results: (1) Intensively teaching 17 month olds lexical categories well-organized by shape creates a generalized bias to extend the names of even novel things by shape; (2) This training also accelerates children's learning of English names for objects outside the laboratory (as measured by parental report of productive vocabulary.); (3) Teaching children names for non-shape based categories and teaching them shape based categories without names does not lead to increased rates of new noun

acquisitions. These results suggest a developmental feedback loop between learning object names and attention to shape: Learning names provides a context in which children can learn the relevance of shape for object categories. The contextual cues associated with naming things (and perhaps most importantly linguistic cues such as determiners) progressively create a generalized bias to extend names to new instances by shape, and as a consequence the more rapid learning of common noun categories.

Although there are many correlations in the learning environment that children may learn about and that may guide attention, we suspect learning object names may be especially potent in the development of the kind-specific attentional biases. Language is – without a doubt – a very special form of regularity in the world in that it is pervasive and shared. We also have proposed that language is special because it is a symbol system and as such conveys special computational properties within an associative learning system that enhance the learning of higher order regularities, but that is an issue for other discussions (see Smith & Gasser, 2005; Colunga & Smith, 2003, Colunga, 2001; see also Yoshida & Smith, 2005).

Explaining S.

The training studies by SJLGSS sought to intensively teach one regularity hypothesized to be relevant to forming new noun categories – that solid artifactual things tend to be in categories containing things similar in shape. The training technique was highly focused – just four lexical categories of unambiguously same-shaped things that matched on no other properties. This method was remarkably effective and potentially relevant to intervening in cases of language delay (see Jones, 2003; also Johnston & Wong, 2002). However, the procedure did not mimic the natural statistics in the world

—which are much messier. The regularities in the natural statistics are over-lapping and probabilistic. Further there are regularities not just relevant to artifacts and shape but also relevant to other kinds of categories such as animates and substances that will simultaneously influence the attentional learning system. ALA suggests that children learn the statistical regularities embedded in the complexities of real world experiences with categories and words. As a first step in testing this idea, Samuelson trained 15 to 20 month old children with two different types of real world artifact categories (mostly organized by shape) and real world substance categories (mostly nonsolid and mostly organized by material) conforming to the natural proportions found in young children's vocabularies. Very young children trained with these more realistic categories and more realistic stimuli acquired a generalized shape bias for solids (and showed accelerated growth in count noun acquisitions outside of the laboratory) but did not develop a generalized material bias. This pattern of development in this micro-genetic experiment thus mimicked the observed developmental trend in novel noun generalizations. Samuelson also implemented a neural net model of the ALA and in her simulations showed that the model's learning closely simulated that of children in the experiments, lending support to the idea that the processes in the model may capture the important aspects of children's learning processes.

Summary

BW&H and B&D posit unitary proposition-like concepts that guide children's performance in these generalization tasks; concepts about ontological kinds (animate vs. artifact) in the case of BW&H and about the very general notion of kind itself in the case of B&D. In contrast, ALA is a process account of learning and of real

time performance in noun learning tasks. It is based on fundamental learning processes that have been widely documented and theoretically studied in experimental psychology. ALA offers a comprehensive and unified account of the early growth of noun vocabularies, of the origins of different patterns of categorization for animate, artifact, and substance categories, of the role of linguistic cues in children's early noun learning, of cross-linguistic differences, and of one aspect of early language delay. It makes novel and testable (falsifiable) predictions and is sufficiently well specified that it can be instantiated in formal models (Samuelson, 2002; Colunga and Smith, 2005). It is not a theory about the content of cognition in the sense of propositional concepts nor it is a theory about bottom-up categorization processes. It is a theory about how knowledge is embedded in real time processes such as attentional learning.

Who is right and what does it mean for a theory of development?

The two opposing grand views of cognition – amodal propositions, rules, and variables versus embedded and distributed across (modal) processes all bound to each other and to the world in real time – are generally viewed as direct opposites. Either one or the other is correct (e.g., McClelland & Patterson, 2002; Pinker & Ullman, 2002). However, there is a construal of the developmental debate over novel noun generalizations, to which we are sympathetic, and under which the two approaches may be viewed as not in direct competition but rather as each capturing some truth at different grains or levels of analysis. Viewed in this way, the sense-think-act approach captures in its propositional representations higher-order properties of the cognitive system that may not directly translate into the underlying processes (see Fodor, 1975) but that none-the-less reveal the *structure* of the knowledge embedded in the many processes relevant to

perceiving, attending, remembering and acting. In this sense, the embedded cognition approach seeks to understand at a finer grain the specific processes and mechanisms in which knowledge is instantiated and made manifest in real time performance.

An analogy might be helpful at this point. Consider the phenomenon of someone going to the cupboard, getting food, and eating it. One might explain the behavior by saying that the person was hungry. Or, one might explain the behavior in terms of glucose levels dropping. Hunger does not reduce simply to blood glucose levels (because multiple factors contribute to perceived hunger) and so “hunger” is a useful theoretical construct above and beyond glucose levels to explain eating behavior. However, hunger and glucose levels are also not theoretical competitors. One could not sensibly do experiments to rule out glucose levels as opposed to hunger because glucose levels are one of the underlying causes of hunger. Asking “What is really and truly driving behavior, hunger or glucose levels?” makes little sense. By analogy, “beliefs about object kind” or a “conceptual distinction between animate kinds and artifacts” is not a direct competitor to processes of perceiving, attending and remembering because that knowledge is made from and is embedded in those very processes.

This line of reasoning does not imply that both levels of analysis are equally good for all tasks nor that it is merely a matter of personal preference. Development is fundamentally about change and thus about processes as a function of time. If one wants to understand cognitive development sufficiently well that one can build artificial systems that change over time given real world experiences (see Smith & Gasser, 2005; Breazeal & Scassellati, 2000; Pfeifer & Scheier, 1999); if one wants to understand cognitive development sufficiently well that one can control and influence real world

development (as in the case of children with language delay); if one wants to understand cognitive development sufficiently well that one can specify how moment-by-moment experiences create lasting and long term change, one needs to understand *process* -- perceiving, remembering, attending and acting.

References

- Ahn, W. & Luhmann, C. C. (2005). Demystifying Theory-Based Categorization. In Gershkoff-Stowe, Lisa (Ed); Rakison, David H. (Ed). (2005). *Building object categories in developmental time*. Carnegie Mellon Symposia on cognition. (pp. 277-300). Mahwah, NJ, US: Lawrence Erlbaum Associates, Publishers. xvii, 463 pp.
- Baldwin, D. A. & Baird, J. A (2001) Discerning intentions in dynamic human action *Trends in Cognitive Sciences*. 5(4), pp. 171-178
- Barsalou, L. W. (1993) Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In Collins, A. F. (Ed); Gathercole, S. E. (Ed); Conway, M. A. (Ed); Morris, P. E. (Ed). *Theories of memory*. (pp. 29-101). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc. xii, 428 pp.
- Barsalou, L. W. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences*. 22(4), pp. 577-660
- Barsalou, L. W., Simmons, W.K., Barbey, A. K. & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*. 7(2), pp. 84-91
- Breazeal, C. & Scassellati, B. (2000) Infant-like social interactions between a robot and a human caregiver. *Adaptive Behavior*. 8(1), Win 2000, pp. 49-74.

- Brooks, R.A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, pp. 139-159.
- Chun MM, Jiang Y (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36, 28-71.
- Chun MM, Jiang Y (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, 10, 360-365.
- Chun MM, Jiang Y (2003). Implicit, long-term spatial contextual memory. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 29(2), 224-234.
- Cimpian, A. & Markman, E. M. (2005). The Absence of a Shape Bias in Children's Word Learning. *Developmental Psychology*. 41(6), Nov 2005, pp. 1003-1019
- Clark, A. (2001) Reasons, robots and the extended mind. *Mind & Language*. 16(2), pp. 121-145
- Colunga, E., (2001). Local vs. distributed representations: Implications for language learning. *Proceedings of the Annual Conference of the Cognitive Science Society*, 23. Edinburgh, UK.
- Colunga, E. (in press). The effect of priming on preschooler's extensions of novel words : how far can "dumb" processes go? To appear in *Proceedings of the 30th Annual Boston University Conference on Language Development*.
- 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.

Dale, R. & Spivey, M. J. (2005). From apples and oranges to symbolic dynamics: A
framework for conciliating notions of cognitive representation. *Journal of
Experimental & Theoretical Artificial Intelligence*. Special Issue: Theoretical
cognitive science. Vol 17(4), pp. 317-342

Dennis, S. & Kruschke, J. K. (1998). Shifting attention in cued recall. *Australian Journal
of Psychology*, 50, 131-138.

Diesendruck, G. & Bloom, P. (2003). How specific is the shape bias? *Child
Development*. 74(1), pp. 168-178

Farkas, I., & Li, P. (2001). A self-organizing neural network model of the acquisition of
word meaning. In M.E.Altmann & A.Cleeremans (Eds.), *Proceedings of the 2001
Fourth International Conference on Cognitive Modeling* (pp.67–72).Mahwah, NJ:
Erlbaum.

- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E. & Thal, D. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development*, 59, (Serial no. 242).
- Gallese, V. & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge *Cognitive Neuropsychology*, 22(3-4), pp. 455-479
- Gelman, S.A. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology*, 20, 65–95.
- Gelman, S.A., & Bloom, P. (2000). Young children are sensitive to how an object was created when deciding what to name it. *Cognition*, 76, 91–103. 950.
- Gelman, S. A. & Koenig, M. A. (2003). Theory-based categorization in early childhood
In Rakison, David H. (Ed); Oakes, Lisa M. (Ed). (2003). *Early category and concept development: Making sense of the blooming, buzzing confusion*. (pp. 330-359). New York, NY, US: Oxford University Press. xxi, 442 pp.
- Hall, D.G., Quantz, D.H., & Personage, K.A. (2000). Preschoolers' use of form class cues in word learning. *Developmental Psychology*, 36, 449–462.
- Imai, M., & Gentner, D. (1997). A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. *Cognition*, 62, 169–200.

- Jiang Y, Chun MM (2001). Selective attention modulates implicit learning. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*. 54A(4), Nov 2001, pp. 1105-1124
- Johnston, J. R. & Wong, M-Y A. (2002). Cultural differences in beliefs and practices concerning talk to children. *Journal of Speech, Language, and Hearing Research*. 45(5), pp. 916-926
- Jones, S. S. (2003). Late talkers show no shape bias in object naming. *Developmental Science* 6, 477–83.
- Jones, Susan S; Smith, Linda B. How children know the relevant properties for generalizing object names *Developmental Science*. 5(2), pp. 219-232
- Jones, S. S. & Smith, L. B. (2005). Object name learning and object perception: A deficit in late talkers. *Journal of Child Language*. 32(1), pp. 223-240
- 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. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In L.Hirschfeld & S.Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 234–254). New York: Cambridge University Press.
- Kemler Nelson, D. G., Russell, R., Duke, N. & Jones, K. (2000). Two-year-olds will name artifacts by their functions. *Child Development*. 71(5), pp. 1271-1288

- 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.
- Kruschke, J. K. (2001). The inverse base rate effect is not explained by eliminative inference. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 27, 1385-1400.
- Lakoff, George (1994). What is a conceptual system? In Overton, Willis F. (Ed); Palermo, David Stuart (Ed). *The nature and ontogenesis of meaning. The Jean Piaget symposium series*. (pp. 41-90). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc. xiii, 307 pp.
- Landau, B., Smith, L.B., & Jones, S.S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 299–321.
- Landauer, T. K. & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*. 104(2), pp. 211-240
- Mandler, J. M. (1992). The foundations of conceptual thought in infancy. *Cognitive Development*. 7(3), Jul-Sep 1992, pp. 273-285
- McClelland, J. L. & Patterson, K. (2002). Rules or connections in past-tense inflections: What does the evidence rule out? *Trends in Cognitive Sciences*. 6(11), pp. 465-472

- McRae, K., de Sa, V.R., & Seidenberg, M.S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, 126, pp. 99–130.
- Mintz, T. H. (2003) Frequent frames as a cue for grammatical categories in child directed speech. *Cognition*. 90(1), pp. 91-117.
- Monaghan, P., Chater, N., & Christiansen, M. H. (2005). The differential role of phonological and distributional cues in grammatical categorization. *Cognition*. 96(2), Jun 2005, pp. 143-182
- O'Regan, J. K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*. 24(5), pp. 939-1031
- Pfeifer, R & Scheier, C., (1999). *Understanding Intelligence*. Cambridge MA: MIT Press.
- Pinker, S. & Ullman, M. T. (2002). The past and future of the past tense. *Trends in Cognitive Sciences*. 6(11), pp. 456-463
- Port, R. F. & van Gelder, T. (1995) *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA, US: The MIT Press. x, 590 pp.
- Prasada, S., Ferenz, K., & Haskell, T. (2002). Conceiving of entities as objects and as stuff. *Cognition*, 83, 141–165.
- Regier, Terry. (2005). The emergence of words: Attentional learning in form and meaning. *Cognitive Science* 29, 819-865.

- Rogers, T. T. & McClelland, J. L. (2005). A Parallel Distributed Processing Approach to Semantic Cognition: Applications to Conceptual Development In Gershkoff-Stowe, Lisa (Ed); Rakison, David H. (Ed). (2005). *Building object categories in developmental time*. Carnegie Mellon Symposia on cognition. (pp. 335-387). Mahwah, NJ, US: Lawrence Erlbaum Associates, Publishers. xvii, 463 pp.
- Rosch, E. & Mervis, C. B., (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*. 7(4), pp. 573-605
- Samuelson, L. K., & Smith, L. B. (1999). Early noun vocabularies: Do ontology, category organization and syntax correspond? *Cognition*, 73(1), 1-33.
- Samuelson, L. K. & Smith, L. B. (2000) Grounding development in cognitive processes *Child Development*. 71(1), pp. 98-106
- Samuelson, L. K., & Smith, L. B. (2005) They Call It Like They See It: Spontaneous Naming and Attention to Shape. *Developmental Science*, 8:2, 182-198.
- Sandhofer, C. M., Smith, L. B. (2004) Perceptual complexity and form class cues in novel word extension tasks: how 4 year-old children interpret adjectives and count nouns. *Developmental Science* 7:3, pp. 378-388.
- Sandhofer, Catherine M; Smith, Linda B. Perceptual complexity and form class cues in novel word extension tasks: How 4-year-old children interpret adjectives and count nouns *Developmental Science*. 7(3), pp. 378-388

- 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., Jones, S. (1993). Cognition Without Concepts. *Cognitive Development*, 8, 181-188.
- Smith, L. B., Samuelson, L. (2000) Grounding Development in Cognitive Process. *Child Development*, 71(1), 98-106.
- Smith, L.B., Colunga, E., & Yoshida, H. (2003). Making an ontology: Cross-linguistic evidence. In D.Rakison & L.Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (pp.275–302). London: Oxford University Press.
- Smith, L., Gasser, M. (2005) The Development of Embodied Cognition: Six Lessons from Babies. *Artificial Life*, 1113-30.
- 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.
- 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. (1992). Inferences about the meaning of nouns: The relationship between perception and syntax. *Cognitive Development*, 7, 29–45.

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

Waxman, S.R., & Markow, D.B. (1998). Object properties and object kind: twenty-one-month-old infants' extension of novel adjectives. *Child Development*, **69**, 1313–1329.

Yoshida, H. & Smith, L.B. (2001). Early noun lexicons in English and Japanese. *Cognition*, 82, B63-B74

Yoshida, H., Smith, L. B. (2003) Correlations, Concepts and Cross-Linguistic Differences. Blackwell Publishing Ltd.

Yoshida, H. & Smith, L.B. (2003a) Known and novel noun extensions: Attention at two levels of abstraction. *Child Development*, 76 (2) 564-577.

Yoshida, H. & Smith, L. B. (2003b) Shifting Ontological Boundaries: How Japanese- and English-Speaking Children Generalize Names for Animals and Artifacts. *Developmental Science*, 6(1), pp. 1-36.

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

Figure Caption

Figure 1. According to ALA, the child learns how different kinds of input correlate with different kinds of category organization. This figure shows how three of the possible classes of input (linguistic context, labels, and perceptual properties) may be correlated with the relevance of different kinds of similarities to a category decision by the history of their past ability to predict how attention is successfully allocated in a task. For example “cat” co-occurs with determiner “a” and adjective “happy” and with the presence of eyes, fur, and attention to shape and texture.