

Cognition Without Concepts

Linda B. Smith

Susan S. Jones

Indiana University

There is an intellectual unease in our article and in the commentaries on it. Although the article and the commentaries are all about the role of perception in children's concepts, the subtext concerns the validity of our most fundamental assumptions about cognition. All of the articles show the tension. For us, and for Mervis, Johnson, and Scott (1993), the strain is evident in data showing dramatic changes and context sensitivity in categories that are neither explained by, nor of interest to, current theory. Mandler's (1993) and Gelman and Medin's (1993) commentaries show the strain in acknowledging interactions between perception, task, and various knowledge sources, yet at the same time also arguing that some kinds of knowledge—intensions, or essences, or functions, or ontological beliefs—are more fundamental than others and, indeed, are *the explanation of why behavior is the way it is*. Barsalou (1993) explicitly recognizes that the tension derives from discontent with the founding assumptions of a structural approach to cognition: "We have no accounts of how propositional representations arise in the cognitive system, either innately or through experience. We haven't the faintest idea of how biological mechanisms could produce abstract propositions, nor of how information processing mechanisms might transduce propositions from perceptual experience" (p. 173). As Barsalou argues, when we cannot imagine how our basic ideas about cognition can possibly be realized, we ought to consider the possibility that they are wrong.

Our original article also asked whether the traditional view of concepts is wrong. We approached this possibility via a long argument that perception matters. Our claim, however, was not about the relative importance of perception and conception. The issue was not whether perceptual features belong at the cores or at the peripheries of concepts. The point was not that there is one early word-learning or object recognition strategy that favors perceptual information. We presented evidence showing that children's novel word extensions vary across different contexts in a way that is systematic, intelligent, and complex. This kind of intelligent variability in cognition means that perception always

Correspondence and requests for reprints should be sent to Linda B. Smith, Psychology 360, Indiana University, Bloomington, IN 47405.

matters. Intelligent variability in cognition also forces us to confront the real issue: namely, do cognitive structures like concepts actually exist?

We believe that the tension between our original article and the commentaries on it derives from the possibility that stable concepts might not exist, and from the difficulty of imagining what cognition could be without represented concepts. Accordingly, in this reply we seek to make the subtext explicit. We consider first the traditional approaches to stability and variability, and why, given the shape-bias data, we find them wanting. Next, we consider just what cognition without represented concepts could be like.

THE TRADITIONAL VIEW

One goal of cognitive psychology is to explain the stability of cognition—the fact, for example, that each time we hear the word *cat*, we think about the same kind of object. When we hear the sentence *That is an odd cat* and when we hear *That is a fat cat* we understand both as referring to a common kind because, according to the traditional view, we possess a single concept of *cat*. This single concept, this representation of what it is to be a cat, sits in our head and is activated by occurrences of the word *cat* (or perceptions of the object). In this traditional view, the reason there is stability in cognition is because we have static, unchanging representations that are repeatedly activated in different contexts.

Stability, however, is only one fact about cognition. There is also variability. In the two cases of the *fat cat* and the *odd cat*, we do not really understand *exactly* the same thing. We have different ideas in the two contexts about the cats referred to. Because by the traditional view, there is a single concept of what it means to be a cat, this context-specific variability cannot be caused by the concept “cat” but rather must be caused by context-specific modifications added to that “core” meaning of cat. In the traditional account, cognitive stability, the fact that we understand (almost) the same thing each time we hear the word *cat*, and cognitive variability, the fact that our understanding adaptively fits the specific task at hand, have separate causes. Stability is caused by constant structure and context variability is caused by real-time processes that access the constant structure and somehow adjust it to the task at hand. This traditional view of cognition is summarized in the cartoon in Figure 1. There is the concept of “cat” and then real-time processes that access, and select from and adjust that constant structure for use in individual real-time tasks.

This traditional partition of cognition into structure (stability) and process (variability) has played a major role in how research is done. The empirical goal is to describe the underlying structures, and the strategy is to discern those structures by finding out what is the same across different instances of a single cognitive act. We do this by collecting data from different individuals in different contexts and then trying to strip away the variability—the context specificities.




Figure 1. A traditional view of cognition: Individual cognitive acts concerning cats vary with context (*fat cat*, *white cat*, *striped cat*, *odd cat*) due to the differential actions of real time processes on a single, stable, represented concept *cat*.

In the traditional view, what is constant across individual cognitive acts, is what is important; and what is variable is noise and error. As Gelman & Medin (1993) remind us, the empirical task of finding the constant structures amidst the variability of real performance is not easy. The very same data can be explained by a theory positing simple representations and complex processes or by a theory positing complex representations and simple processes. The intellectual unease that pervades the current discussion, however, does not derive from disagreement about simple versus complex representations, nor from disagreements about how much of behavior is due to constant structures and how much reflects process. Rather the tension arises from the possibility that the whole idea of structure is wrong—that cognitive structures cannot explain the intelligence of cognition.

WHAT IS SMART ABOUT COGNITION?

The structural approach concentrates on the stability of cognition and does a plausible job explaining it. The structural approach pays less attention to variability (indeed, under a traditional approach, we design experiments to minimize variability) and not surprisingly, it does a poor job explaining the variability and context sensitivity of individual cognitive acts. This is a crucial flaw. According to the traditional view, intelligence—what is smart about human cognition—is the represented knowledge. To be useful, static representations must apply across contexts—across individual acts of cognition and individual categories. Thus, what is constant in concepts is highly abstract—a “notion of kind” for Mandler (1993), abstract perceptual features for Barsalou (1993), an “essence placeholder” for Gelman & Medin (1993). The highest forms of intelligence in this view are the most abstract, the most removed from the messy (perceptual) details of here-and-now reality. This is why the fact of context-sensitive cognitions challenges the traditional structuralist view. If what is represented is far removed from here-and-now details and universal, then the abstract representations cannot explain the adaptiveness of individual cognitive acts.

The data on the shape bias present a different picture of what is smart about human cognition. The intelligence of children’s novel word interpretations lies in the ability to do something unique—something that has never been done before but that fits the specifics of the moment. If there is an abstract represented structure of some kind that sits behind the shape bias, it is the least interesting, least intelligent part, of the child’s word-learning behavior. All the work that makes novel word interpretations smart is done by those processes that involve the specific objects and specific words at hand. It is these real-time, real-task processes that flexibly adjust attention to find the most likely referent of a specific utterance of some unknown word. Novel word interpretation is not smart because it is stable; novel word interpretation is smart because it is creatively adaptive. And for this kind of intelligence, perception—information about the here-and-now—always matters.

Our claim that perception matters, however, is a deeper theoretical claim than that novel word interpretation is creatively adaptive. In our view, what is smart about all of human cognition is its continuous exquisite adaptiveness to the particulars of real-time external (and also internal) events. Perception—contact with here-and-now reality—matters not just for naming but in individual acts of believing, intending, feeling, and realizing. If these cognitive acts in their individual occurrences are adaptively intelligent, then they too must depend on contact with the here-and-now.

Herein lies our discontent: If structures control what is constant about cognition, but if individual cognitive acts are smartly unique and adaptive to the context, structures cannot be the cause of the adaptiveness of individual cognitions. Why, then, are structures so theoretically important? If the intelligence—and the cause of real-time individual cognitive acts—is outside the constant structures, what is the value of postulating such structures?

But here lies the tension. Although it is difficult to imagine how represented structures like “essence placeholders” or a “notion of kind” can ever be realized by our biology, it is also difficult to imagine cognition without such representations. The metaphor behind the traditional approach to cognition is the computer metaphor: Cognition is computation. There is no computation without representations. What could cognition be if it were not computation, if there were not representations?

COGNITION AS PROCESS

van Gelder (1992) has offered a metaphoric alternative for thinking about cognition and specifically for thinking about what cognition might be if it were not computation. van Gelder’s insight derives from thinking about two possible solutions to a 19th century engineering problem. The problem is this: You have a steam engine that drives a flywheel to which machinery is connected. It is important that the speed of the flywheel remain constant and this is a practical problem because both the workload on the engine and the steam pressure vary irregularly and continuously. How could one design a device, called a governor, to maintain a constant flywheel speed? One solution is a computational one. A computational device might, as van Gelder proposes, contain the following: a tachometer for measuring the speed of the wheel, a device for calculating speed discrepancy, a steam pressure meter, a device for calculating the throttle valve adjustment, a throttle valve adjuster, and an executive to handle the sequencing of operations. Notice that within these devices are representations; indeed this computational governor works—just like the traditional metaphor of cognition—by the manipulation and passing of representations from one component to the next.

A computational governor might well work but probably would not do as well or adapt as intelligently and fluidly to changes in workload and pressure as does the simple and elegant device invented by James Watt in the early 1800s. A

version of Watt's device, the centrifugal governor, is illustrated in Figure 2. The centrifugal governor consists of a vertical spindle geared into the main flywheel so that it rotates at a speed directly dependent upon that of the flywheel itself. As van Gelder writes:

Attached to the spindle by hinges were two arms, and on the end of each arm was a metal ball. As the spindle turned, centrifugal force drove the balls outwards and hence upwards. By clever arrangement, the arm motion was linked directly to the throttle valve. The result was that as the speed of the main wheel increased, the arms raised, closing the valve and restricting the flow of steam; as the speed decreased, the arms fell, opening the valve and allowing more steam to flow. The result was that the engine adopted a constant speed, maintained with extraordinary swiftness and smoothness in the presence of large fluctuations in pressure and load. (van Gelder, 1992, p. 3)

Watt's governor accomplishes its tasks in a radically different way than the computational device—in a way that allows for intelligent and continuous context sensitivity and a global order (constant speed) *without representations or computations*. Watt's governor does not represent anything; it just does the job. It does its near perfect job because its activity at every point in time is sensitively




Figure 2. A centrifugal governor like that invented by Watt to control the speed of a steam engine. Redrawn from van Gelder (1992).

and simultaneously dependent on everything. The “sensory input,” the pressure and load fluctuations, always matter; there is no purpose to the centrifugal governor without them. Finally, the device is smart; it *embodies* (although it does not represent) as much knowledge as the computational governor.

The account of children’s novel word interpretations that we offer in the original article envisions a process like Watt’s centrifugal governor. Children’s attention to object properties in context shift dramatically and intelligently in ways that combine a myriad of different forces because of the simultaneous and continuous dependence of attention on everything. Of course, children’s novel word interpretations are importantly not like the centrifugal governor in one way: Novel word interpretations develop. We must envision a centrifugal governor that changes itself as it operates. Such visions are being theoretically realized under a dynamic systems perspective (e.g., Edelman, 1987, 1992; Saltzman & Munhall, 1992; Thelen & Smith, in press; van Gelder, 1992).

The example of the centrifugal governor suggests that cognition could be all process. Although the centrifugal governor is just a metaphor, the statement that cognition is all process may be more than a metaphor. The brain is all process and there, categories are likely to be dynamic and nonstationary.

DYNAMIC CATEGORIES IN THE BRAIN

What might categories be in a system that is all process? Freeman and his colleagues (Freeman, 1981, 1991; Skarda & Freeman, 1987) suggest a remarkable answer from their study of olfactory categories in rabbits. In their study of the activity across groups of neurons in the olfactory bulb, they found stable behavioral categories that emerged from dynamic and changing patterns of activity. They found that different inhalants did not map to any single neuron or even group of neurons but rather to the spatial pattern of the amplitude of waves across the entire olfactory bulb. Importantly, however, there was no one-to-one relation between odor category and patterns of activity. Thus, after a given point in a rabbit’s experience with smells, say after learning about three different smells, three categories of smells could be discerned in terms of different waves of activity across the olfactory bulb. Each time the rabbit sniffed a particular odorant *under the same conditions*, that odorant produced the same global pattern of activity. But the pattern of activity in the olfactory bulb that corresponded to a particular odor changed with context, the rabbit’s state, and rabbit’s developmental history. The changes in patterns of activity corresponding to particular odorants were particularly dramatic after new experiences. For example, in one study, rabbits were conditioned to associate the scent of sawdust with a particular reinforcement. After this learning there was a new characteristic pattern of olfactory bulb activity associated with the scent of sawdust. When, however, they were taught to recognize the odor of banana, a new characteristic pattern of activity emerged *for sawdust* as well as for banana. The rabbit’s behavioral

response to sawdust still looked the same as in the pre-banana period. But that apparent stability in outward behavior—the categorization of sawdust before and after training about bananas—did not rest on a constant underlying structure or a constant pattern of activity.

CONCLUSION

In our original article, we argued simultaneously on two levels. On one level, we questioned the conclusions that some have drawn about the role (or nonrole) of perception in conceptual development. We questioned these conclusions both because they seem to be unwarranted by the data put forth to support them and because the characterization of concepts and cognition emerging from that literature seemed so at odds with the characterization of cognition suggested by our results on children's novel word interpretations. On a second level, we questioned the theoretical framework—the founding assumptions—that underlie the attempt to define what "concepts really are." We believe that the data on developing novel word interpretations—data showing the creative intelligence of dynamic cognition—seriously challenge the view of cognition as represented knowledge structures. These results suggest that perception always matters *in a deep way*. Perception always matters because cognition is always adaptive to the here-and-now, and perception is our only means of contact with the here-and-now reality.

REFERENCES

- Barsalou, L.W. (1993). Challenging assumptions about concepts. *Cognitive Development, 8*, 169–180.
- Edelman, G.M. (1987). *Neural Darwinism*. New York: Basic Books.
- Edelman, G.M. (1992). *Bright air, brilliant fire. On the matter of mind*. New York: Basic Books.
- Freeman, W.J. (1981). A physiological hypothesis of perception. *Perspectives in Biology and Medicine, 24*, 561–592.
- Freeman, W.J. (1991). The physiology of perception. *Scientific American, 264*, 78–85.
- Gelman, S.A., & Medin, D.L. (1993). What's so essential about essentialism? A different perspective on the interaction of perception, language, and conceptual knowledge. *Cognitive Development, 8*, 157–167.
- Mandler, J.M. (1993). On concepts. *Cognitive Development, 8*, 141–148.
- Mervis, C.B., Johnson, K.E., & Scott, P. (1993). Perceptual knowledge, conceptual knowledge, and expertise: Comment on Jones and Smith. *Cognitive Development, 8*, 149–155.
- Saltzman, E.L., & Munhall, K.G. (1992). Skill acquisition and development: The roles of state-, parameter-, and graph-dynamics. *Journal of Motor Behavior, 24*, 49–57.
- Skarda, C.A., & Freeman, W.J. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences, 10*, 161–195.
- Thelen, E., & Smith, L.B. (in press). *A dynamic systems account of the development of cognition and action*. Cambridge, MA: MIT Press.
- van Gelder, T. (1992). *What might cognition be if not computation?* Research Report No. 75, Program in Cognitive Science, Indiana University, Bloomington.