Psychological Review
1997, Vol. 104, No. 4, 686–713

Rethinking Infant Knowledge: Toward an Adaptive Process Account of Successes and Failures in Object Permanence Tasks

Yuko Munakata and James L. McClelland
Carnegie Mellon University and the Center for the Neural Basis of Cognition

Mark H. Johnson
Medical Research Council Cognitive Development Unit

Robert S. Siegler
Carnegie Mellon University

Infants seem sensitive to hidden objects in habituation tasks at 3.5 months but fail to retrieve hidden objects until 8 months. The authors first consider principle-based accounts of these successes and failures, in which early successes imply knowledge of principles and failures are attributed to ancillary deficits. One account is that infants younger than 8 months have the object permanence principle but lack means–ends abilities. To test this, 7-month-olds were trained on means–ends behaviors and were tested on retrieval of visible and occluded toys. Means–ends demands were the same, yet infants made more toy-guided retrievals in the visible case. The authors offer an adaptive process account in which knowledge is graded and embedded in specific behavioral processes. Simulation models that learn gradually to represent occluded objects show how this approach can account for success and failure in object permanence tasks without assuming principles and ancillary deficits.

We infer what infants know from their behavior. Researchers design tasks to tap various kinds of knowledge and test whether infants succeed or fail. This approach has been used over many years, in an effort to increase our understanding of infants' knowledge. Puzzles often arise, however. Two tasks supposedly tap the same knowledge, but the same infants succeed on one and fail the other. For example, young infants in visual habituation experiments have demonstrated behavior consistent with knowledge of several physical principles (Baillargeon, 1993; Leslie, 1988; Spelke, Breinlinger, Macomber, & Jacobson, 1992). And yet, infants simultaneously, and often for many months following such demonstrations of knowledge, fail other seemingly basic measures designed to tap the same knowledge of physical principles. Such patterns of simultaneous successes and failures have been documented across a wide variety of ages, domains, and task conditions, perhaps most notably in the contrast between recent demonstrations of early competence and Piaget's (1952, 1954) robust findings of infants' limitations in cognitive tasks. However, relatively little discussion has been devoted to the resulting implications for what infants really know. Instead, attention has typically been focused on infants' successes, with explanations of their failures relegated to factors considered to be outside the theoretical domain of interest (see discussions in Johnson & Morton, 1991; McClelland, 1994; Siegler, 1996; Smith & Thelen, 1993; Thelen & Smith, 1994).

However, given that the same developing system produces both the successes and the failures during the learning process, accounting for both seems likely to be critical in understanding the origins of knowledge (Braine, 1959; Brown, 1976; Flavell, 1985). Why do infants simultaneously fail and succeed on different tasks meant to measure the same knowledge? What might this tell us about the nature of cognitive development? How can we understand the changes that underlie these developmental patterns?

The Piagetian concept of object permanence provides a rich experimental domain in which to explore these questions. Piaget defined the object permanence concept as the understanding that objects continue to exist independent of our percepts of them and that objects maintain their identity through changes in location (Piaget, 1954). Piaget used the successful retrieval of a com-

---

Yuko Munakata and James L. McClelland, Department of Psychology, Carnegie Mellon University (CMU), and the Center for the Neural Basis of Cognition, Pittsburgh, Pennsylvania; Mark H. Johnson, Medical Research Council Cognitive Development Unit, London, England; Robert S. Siegler, Department of Psychology, Carnegie Mellon University.

A version of this article was presented at the International Conference on Infant Studies, Paris, April 1994, and submitted by Yuko Munakata to CMU in partial fulfillment of the requirements for the degree of doctor of philosophy.

The research was supported by McDonnell-Pew Grant 93-29, an Office of Naval Research graduate fellowship to Yuko Munakata, National Institutes of Health Grant MH-00385, National Science Foundation Grant DBS-9120433, direct funding from the United Kingdom Medical Research Council, and National Institute of Child Health and Human Development Grant HD-19011.

We thank Annette Karmiloff-Smith, Randy O'Reilly, and Liz Spelke for extensive comments on drafts of this article, and we thank them and members of the CMU Developmental and Parallel Distributed Processing Research Groups for many useful discussions. We thank Christine Gundlach, Leslie Tucker, and Robin Wilcox for lab assistance.

Correspondence concerning this article should be addressed to Yuko Munakata, who is now at the Department of Psychology, University of Denver, 2155 South Race Street, Denver, Colorado 80208. Electronic mail may be sent via Internet to munakata@du.edu.

pletely hidden object as one measure of the object permanence concept. Infants succeed with this task only around 8 months of age and even then show incomplete mastery by making errors in the $A\bar{B}$ task. In this task, devised by Piaget (1954), infants watch as an object is hidden in one location (A). Those infants who are able to successfully retrieve the object may nonetheless fail to retrieve the object when it is hidden in a new location (B), showing a perseverative error (the $A\bar{B}$ error) of reaching to the original hiding location (A).

Many researchers have followed Piaget's tradition of studying infant behavior to evaluate an underlying object permanence concept. The common assumption is that sensitivity to hidden objects takes the form of such an underlying concept; the task is then one of devising a test for such sensitivity to assess the presence of the concept. In this vein, clever experiments have been designed to evaluate how early infants "have" the object permanence concept. Baillargeon (1987a, 1993), for example, measured looking-times to possible and impossible events involving occluded objects and found that infants as young as 3.5 months looked longer at the impossible events. These longer "looking times" are taken as an indication of the infants' perception of an unusual event and therefore of the understanding of the continued existence of the occluded object. Earlier signs of competence have also been demonstrated in the $A\bar{B}$ task. Diamond (1985) has noted that infants occasionally reach to the wrong location while looking at the correct location, and often reach to the wrong location without looking in it and then immediately reach to the correct location. Hofstadter and Reznick (1996) confirmed that when infants' looking and reaching behaviors differ in this task, the looking response is more accurate. Infants show differences on looking and reaching measures on the $A\bar{B}$ task in two other ways. First, in versions of this task in which the infant is allowed only to look (e.g., Hofstadter & Reznick, 1996; Lecuyer, Abgueguen, & Lemarie, 1992; Matthews, 1992), results again suggest greater sensitivity to a change in hiding location than indicated by reaching measures. In addition, Ahmed and Ruffman (1997) and Baillargeon and colleagues (Baillargeon, DeVos, & Graber, 1989; Baillargeon & Graber, 1988) have demonstrated through looking-time analogs of the $A\bar{B}$ task that 8–12-month-old infants can detect impossible events after delays at which the infants would fail to search correctly. From these types of findings, many researchers now attribute a concept of object permanence to infants as young as 3.5 months. In fact, some of Spelke's findings using the looking-time method (Spelke et al., 1992) have suggested the existence of such a concept at an even earlier age.

The key question then becomes: Why do infants fail to retrieve hidden objects until 8 months and even then show the $A\bar{B}$ error, if they have a concept of object permanence many months earlier? The way in which one answers this question depends critically on one's conception of what it means to know something. What does it mean to say that infants know that an occluded object is still there? What form does this knowledge take? How is this knowledge accessed and used?

These are fundamental questions, but there has been relatively little theorizing to address them or to explain what infants can and cannot do. Instead, there has been a common tendency to treat infants' knowledge as taking the form of all-or-none, proposition-like entities (see discussions in Fischer & Bidell, 1991 [specifically in the context of object permanence]; Flavell, 1971, 1984; Karmiloff-Smith, 1991, 1992; Siegler, 1989, 1993; Siegler & Munakata, 1993; Smith & Thelen, 1993; Thelen & Smith, 1994). For example, Baillargeon (1994) stated:

> The first developmental pattern is that, when learning about a new physical phenomenon, infants first form a preliminary, all-or-none concept that captures the essence of the phenomenon but few of its details. (p. 133)

Although the assumption that knowledge has these features is not always so explicit, we believe it is implicit in many theoretical discussions of cognitive development. Often, discussions are framed in terms of the idea that knowledge takes the form of principles that function like propositions: that is, the principles are construed as generally accessible inputs to a reasoning process (Diamond, 1991; Spelke et al., 1992). We refer to this type of approach as a principle-based approach. Although accounts within this principle-based approach sometimes allow for the elaboration or enriching of initial concepts, these elaborations are often overlooked in theorizing about infant behavior. In particular, such elaborations are not typically used in explanations of the context-dependent nature of infant competence.

Because infants seem to behave in accord with principles at times, there might be some use to describing their behavior in these terms. The danger, we believe, comes in the tendency to accept these descriptions of behavior as mental entities that are explicitly accessed and used in the production of behavior (for a similar discussion of linguistic rules, see Rumelhart & McClelland, 1986). That is, one could say that infants' behavior in a looking-time task accords with a principle of object permanence, in the same way that one could say that the motions of the planets accord with Kepler's laws. However, it is a further—and we argue, unfounded—step to then conclude that infants actually access and reason with an explicit representation of the principle itself. In the same way, one would not want to explain the motions of the planets by claiming that the planets derive their next location in space on the basis of reasoning with Kepler's laws. We present an alternative approach that focuses on the adaptive mechanisms that may give rise to behavior and on the processes that may underlie change in these mechanisms. We show that one might characterize these mechanisms as behaving in accordance with particular principles (under certain conditions); however, such characterizations would serve more as a shorthand description of the mechanisms' behavior, not as a claim that the mechanisms explicitly consult and reason with these principles. We believe that progress in our understanding of cognitive development depends on the specification of such processing mechanisms, rather than on the attribution of principles as explanations of behavior.

In our approach, the knowledge underlying infants' behaviors is best viewed as graded in nature, evolving with experience, and embedded in specific processes underlying overt behavior. We call this approach the *adaptive process approach.* Our approach is motivated by general views of the nature and development of cognitive competence from the frameworks of parallel distributed processing (PDP) and cognitive neuroscience. In the PDP approach, behaviors are expressed through the activation of processing units actually engaged by a task (Rumelhart & McClelland, 1986; for developmental applications see McClelland, 1989, 1992). These activations are determined by the strengths of the connections linking the processing units. Such connections are graded in nature and evolve gradually in response to experience. Graded, embedded, and evolving processes are also evident across a wide variety of domains in cognitive neuroscience (e.g., Bachevalier & Mishkin, 1984; Greenough, Black, & Wallace, 1987; Morton & Johnson, 1991). Morton and Johnson (1991), for example, have reconciled a large body of data on infant face recognition by proposing that a subcortical system directs infants to attend to faces and that these experiences then drive gradual cortical learning that takes place over many months or years. They proposed that the increasing specificity of cortical representations of faces allows the infant to make increasingly detailed judgments about individual faces, expressions, and eye-gaze direction.

Our adaptive process approach is similar in some ways to the skills approach taken by Fischer and Bidell (1991) and the dynamic systems approach detailed by Thelen and Smith (Smith & Thelen, 1993; Thelen & Smith, 1994). In the skills approach, behavior is the expression of skills (context-sensitive procedures) that evolve with practice. In the dynamic systems framework, behaviors are viewed as emergent patterns of activity dependent on an individual's situation and history and embodied in physical processing systems. Our adaptive process approach is consistent with these notions of evolving, experience-based, embodied knowledge, although we emphasize the possibility that processing is guided by representations that the system learns to form through experience.

The adaptive process and principle-based approaches lead to contrasting ways to think about infants' successes and failures. Within the principle-based approach, infants' early successes lead to the attribution of principles within the first few months of life. Such attributions raise the possibility that some of these principles are present at birth (Spelke et al., 1992); indeed, they have sometimes been taken as supporting nativist views of the origins of knowledge in a range of domains (Keil, 1981).

To account for infants' failures to retrieve hidden objects up until 8 months given the principle-based framework, one needs to look outside the concept of object permanence to some sort of *ancillary deficit* because infants seem to have this principle months earlier. Several researchers (Baillargeon, Graber, DeVos, & Black, 1990; Diamond, 1991; Willatts, 1990) have turned to an account based on deficits in means–ends behaviors, arguing that infants cannot act on one object as a means to retrieving another.[1] For example, Diamond (1991) hypothesized that infants might fail the retrieval task because

> infants cannot organize a means–end action sequence at 4–5 months, but they can at 7½–8 months, and the actions which infants have been required to make to demonstrate that they understand object permanence have always involved a sequence of actions (e.g., removing a cloth as the means to retrieving the toy underneath it). (p. 80)

The principle-based framework also leads to specific ways of thinking about why infants make the $A\bar{B}$ error. Again, explanations are relegated to factors external to infants' knowledge representations. Diamond (1985) has suggested that infants' inability to inhibit the conditioned reaching response causes the $A\bar{B}$ error. Baillargeon and colleagues (Baillargeon et al., 1989; Baillargeon & Graber, 1988) have attributed the $A\bar{B}$ error to deficits in search behaviors and problem-solving abilities.

Although some findings appear to be consistent with the principle-based approach, most of the evidence is equivocal, and there are some contrary findings. For example, the means–ends explanation of failures to reach for occluded objects seems to be supported by the finding that 5-month-old infants reach for objects in the dark (Clifton, Rochat, Litovsky, & Perris, 1991; Hood & Willatts, 1986). Infants cannot see the objects in the dark, so object permanence representations are thought to be required. But they can reach for them directly, so means–ends behaviors are presumed to be unnecessary. If one views the occluded-object and object-in-the-dark tasks as differing only in their reliance on means–ends behaviors, then earlier successes in the dark suggest that failures with occluded objects are in fact based on means–ends deficits. However, the occluded-object and object-in-the-dark tasks differ in another way; in the occluded-object condition, the visual input may suggest that no object is present whereas in the dark there is no visual input at all. Later in this article, we use this difference to explain these data within our adaptive process approach, without invoking an ancillary deficit.

The inhibition theory of the $A\bar{B}$ error seems to be supported by the finding that infants occasionally reach to Location A even when the object is visible at Location B (Bremner & Knowles, 1984; Butterworth, 1977; Harris, 1974). However, such reaches may be random incorrect responses rather than true $A\bar{B}$ errors. In support of this, Sophian and Yengo (1985) ran an $A\bar{B}$ experiment with a third, control location. They found that when the toy was visible and infants erred, they were as likely to search to the control location as to the previous location. Several other findings also call into question the inhibition theory for the $A\bar{B}$ error. Infants show the $A\bar{B}$ error even after merely seeing (but not retrieving) an object hidden and revealed at Location A (Butterworth, 1974; Diamond, 1983; Evans, 1973). Also, the extent to which infants show the $A\bar{B}$ error is influenced by factors apparently unrelated to inhibition, such as the presence of a cover at Location A (Bremner & Knowles, 1984) and the distinctiveness of available location cues (see Wellman, Cross, & Bartsch, 1986, for meta-analysis). Additionally, looking and reaching responses would presumably receive similar conditioning in the $A\bar{B}$ task, and yet looking measures have revealed earlier sensitivity than reaching measures (Baillargeon et al., 1989; Baillargeon & Graber, 1988). For all of these reasons, a simple inability to inhibit a conditioned response explanation seems insufficient.

Taking an adaptive process approach can lead to a very different perspective on infants' successes and failures in object permanence tasks. An initial question that arises within this approach is "What kinds of processes might support infants' longer looking times to impossible events involving occluded

---

[1] This notion of means–ends behavior differs from the information-processing notion of determining the difference between the current and goal states and finding an operator to reduce this difference (Newell & Simon, 1972). In the context of object permanence, "means–ends behavior" has been used to refer to the process of acting on one object in relation to another. For consistency with the object permanence literature, we adopt the latter usage in this article.

objects?" We suggest that longer looking times are driven by a mismatch between an infant's expectations about the world and the events that actually transpire. A sketch of the model that we use to illustrate this process is shown in Figure 1. In this mechanism, visual input drives lower level representations capturing spatial relations among visible objects in the world, and these in turn provide one source of input to higher level representations capturing spatial relations between represented objects that may or may not be visible. A second source of input to these higher level representations comes from their own prior state such that, for example, objects that were visible at one time and then occluded can be represented at this level even though the lower level visual representation no longer indicates their presence. These higher level representations can then serve as the basis for implicit predictions about subsequent states of the lower level visual representations. For example, if an occluding object begins to move from its station in front of a hidden object, a representation of the hidden object at the higher level could trigger the prediction that the occluded object will reappear. A discrepancy between these predictions and events in the world provides a signal that causes an increase in looking when unexpected events occur.

Our approach relies on maintained activation to represent objects that are no longer visible for short periods of time. Evidence relevant to the idea that maintained activation underlies sensitivity to such objects comes from data on the inferotemporal and prefrontal cortices. Miyashita (Miyashita, 1988; Miyashita & Chang, 1988) described neurons in inferotemporal cortex that selectively code for visible stimuli. Their activation is partially maintained through a delay period during which the (no longer visible) stimuli must be remembered. Perrett (Perrett, Rolls, & Caan, 1982; Perrett et al., 1984, 1985) and many others (see Maunsell & Newsome, 1987, for a review) have shown that inferotemporal neurons also respond preferentially to various kinds of more complex stimuli. Several researchers have demonstrated similar maintained activity to to-be-remembered stimuli in neurons in the prefrontal cortex (Fuster, 1989; Goldman-Rakic, 1987).

These data suggest that the maintenance of the activation of neurons that respond preferentially to certain objects in the world, after the removal of the sensory stimuli, can support appropriate behaviors in object permanence tasks. Farah (1988) and numerous others (Davidson & Schwarz, 1977; Goldenberg, Podreka, Steiner, & Willmes, 1987; see Farah, 1988, for a thorough review) present evidence supporting analogous arguments for common neural substrates underlying the perception of stimuli and imagery in the absence of actual stimulation. In the specific adaptive process account explored here, we posit that such neural substrates—shared for the representation of visible and occluded or absent objects—are also shared across tasks, such as implicit prediction formation and reaching.

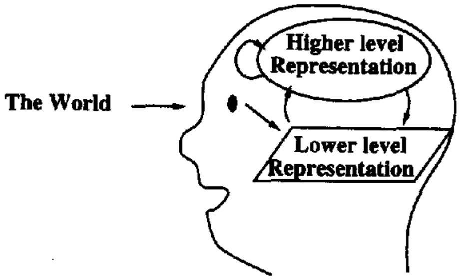

Figure 1. Simplified mechanism for implementing an adaptive process account of longer looking times to impossible events: Input from the world produces lower level representations. These in turn produce higher level representations that can be used to form predictions about subsequent inputs from the world. When there is a mismatch between predictions and observed events, the discrepancy serves as a signal that causes infants to increase looking when unexpected events occur.

What accounts for infants' successes within the adaptive process approach? In this framework, the ability to make perceptual predictions does not imply that infants have a concept of object permanence. Infants' knowledge is not viewed as simply present or absent. Instead, knowledge is viewed as embedded in the underlying processing systems that give rise to behavior. Specifically, we suggest that the ability to represent occluded objects depends on the connections among relevant neurons and that the ability is acquired through a process of strengthening these connections. This in turn leads to a gradual strengthening of the representations of occluded objects such that infants become increasingly able to behave in ways that demonstrate sensitivity to hidden objects.[2] In simulations presented later, we illustrate these points.

What might account for infants' failures on some tasks but not others within the adaptive process approach? One possibility—the one we stress in this article—is that different behaviors may require different degrees of development in the relevant underlying processing systems and the resulting internal representations. In particular, a weak internal representation of an occluded object might be sufficient to guide perceptual predictions and therefore longer looking times to impossible events. However, these weak representations might not be strong enough to drive reaching behaviors, perhaps because of a greater complexity and effort level of reaching behaviors, their lower frequency, or both. Thus, reaching behaviors may require more fully developed internal representations.

Similarly, stronger internal representations might be required for the retrieval of objects occluded in the light versus hidden by darkness. A somewhat weak internal representation of an occluded object might not be able to overcome the interference produced by the visual stimulus of an occluder where the object used to be. This same representation, however, might be strong enough to guide a reach in the dark when there is no direct visual information conflicting with the weak internal representation.

In the neuropsychological literature, there are numerous ana-

---

[2] An understanding of the permanence of objects might exist independent of the actual persistence of visual object representations. For example, as adults, we seem to have an explicit, stateable understanding of objects continuing to exist independent of our percepts of them, without needing to base this understanding on the persistence of specific object representations. Whether or when infants develop knowledge of such an explicit form is an open question. In the arguments presented here, we focus on a more implicit object permanence understanding that we believe develops primarily prior to and as support for later, more explicit understanding (for relevant discussion, see Karmiloff-Smith, 1992; Mandler, 1992).

logs to the argument that various behaviors can be differentially affected by the state of a processing system; these arguments rely on the notion of graded strength of underlying representations (see Farah, O'Reilly, & Vecera, 1993, for a review), as our arguments do. For example, certain patients with a deficit known as extinction have been shown to make accurate same–different judgments about pairs of stimuli, although they are impaired in reporting the identity of the stimulus opposite the side of their brain lesion (Volpe, LeDoux, & Gazzaniga, 1979). Farah, Monheit, and Wallace (1991) have argued that this task-dependency is based on the poor perception of extinguished stimuli, with identification requiring more visual information than same–different judgments. Similar arguments have been made to explain task dependency in prosopagnosia (Farah et al., 1993), the selective deficit in the overt recognition of faces. Although prosopagnosic patients show impairments on a variety of overt measures of face recognition, they are nonetheless able to show signs of covert recognition (e.g., in consistently relearning correct face–name and face–occupation pairings more quickly than learning incorrect pairings; De Haan, Young, & Newcombe, 1987a, 1987b). Farah et al. (1993) have proposed that a single system subserves overt and covert visual recognition, but degraded representations following damage to this system are only strong enough to support covert recognition. We are applying such a graded representation argument to characterize the performance of developing neural systems in object permanence tasks. Specifically, we suggest that the neural substrates subserving processing of occluded objects in 7-month-old infants may sustain representations that are strong enough to drive perceptual expectations and also longer looking times to impossible events, but the representations are not so strong as to be able to drive reaching responses. With gradual changes in these substrates in response to experience, these internal representations might become increasingly able to guide a greater range of behaviors and overcome stronger interference.

We stress that in the adaptive process approach, performance is a function of the state of development of both task-specific mechanisms and representational systems that may be shared across tasks. However, we believe that principle-based approaches have called sufficient attention to the importance of developments in task-specific factors, such as reaching and means–ends behaviors. Thus, a primary goal of this article is to demonstrate the potentially powerful (and typically overlooked) contributions of the development of graded representations to task-dependent behavior.

A schematic outline of our approach to understanding infant behavior in object permanence tasks is presented in Table 1. We present three studies that challenge the means–ends deficit account of infants' failures to retrieve occluded objects. Our data rule out what is arguably the most commonly accepted explanation of the task-dependent nature of infants' behavior in object permanence tasks. However, the principle-based account could be salvaged by invoking deficits in global ancillary factors such as capacity, memory, attention, and motivation, all of which are assumed to improve with development. More difficult tasks require fuller development of the ancillary factor, so infants can demonstrate their underlying knowledge only as the factor improves. Although some such proposals may be consistent with available data, we argue that they suffer from several limitations that make it worthwhile to consider alternatives. Our adaptive process approach provides this, as we show through simulations that instantiate our adaptive process approach in PDP networks. We show how the PDP simulations can account for the ability to behave in accordance with principles, and we illustrate how this ability may be acquired gradually in response to experiences that accord with the principles. These effects reflect the operation of the gradual strengthening of connections in the PDP network. Thus, the simulations show that appeals to global ancillary factors may be unnecessary for understanding development.

Evaluating the Means–Ends Deficit Theories

Is it reasonable to view infants' knowledge as taking the form of principles, requiring means–ends explanations for why infants fail to retrieve hidden objects up until 8 months? To explore this question, one can test infants on two types of trials that require the same means–ends abilities but different object concept knowledge. If the infants fail more often on one type of trial than the other, we cannot attribute this failure to a mere means–ends deficit.

Table 1
Consideration of Accounts of Infants' Looking and Reaching Behaviors in Object Permanence Tasks

	Theory	
Aspect	Principle-based	Adaptive process
Premise	Infants have object permanence principle within first few months of life.	Connections underlying formation and use of representations needed to perform object permanence tasks strengthen with experience.
Explanation of task-dependency	Reaching makes demands on means–ends abilities that infants lack.	Reaching requires stronger representations.
Evaluation	Our experiments show more toy-guided reaching under visible versus occluded conditions even when means–ends demands are equated, contrary to means–ends account.	Our simulations show how the adaptive-process approach can account for infants' task-dependent behaviors and show graded changes in performance.
Extension	Other ancillary deficit accounts are possible, but the adaptive process approach renders appeals to such accounts unnecessary.	Graded changes in simulations capture phenomena motivating appeals to ancillary deficits in principle-based accounts.

Piaget (1954) conducted relevant demonstrations with his own infants when they were between 6.5–9 months of age, showing that they would search for partially occluded objects but not fully occluded objects. Many researchers (Gratch, 1972; Gratch & Landers, 1971; Miller, Cohen, & Hill, 1970; Uzgiris & Hunt, 1975) have replicated Piaget's observation that infants search for partially occluded objects at younger ages than they search for fully occluded objects. These findings might argue against the means–ends hypothesis; infants' successes with partially occluded objects suggest that they in fact have the requisite means–ends skills needed to retrieve fully occluded objects. In a similar vein, Bower and Wishart (1972) and others (e.g., Yonas, cited in Bower & Paterson, 1972) tested infants with transparent and opaque covers over toys. Five-month-old infants successfully retrieved a toy from underneath a transparent cover, but not from underneath an opaque cover. Again, these findings might indicate that infants' failures in the opaque case were not based on means–ends deficits because infants could carry out the required means–ends behaviors (lift the cover to get the toy) in the transparent case.

However, there is an alternative interpretation to these results. The partially hidden and transparent conditions may not require the same means–ends abilities as the completely hidden–opaque conditions. Infants could succeed in the partially hidden and transparent conditions by reaching directly for the toy. Given a partially hidden toy, infants may reach directly for the unobstructed part of the toy. In the transparent cover case, infants might inadvertently retrieve the toy by attempting to reach directly for it. Diamond (1981, 1991) has shown that infants exhibit a remarkably strong tendency to reach directly along the line of sight for an object under a transparent box rather than reaching through an open face of the box. In the context of experiments with transparent and opaque covers, such direct reaching might make it more likely that infants would knock over or grasp the transparent cover and thus more likely that they would retrieve the toy in the transparent condition. In contrast, the opaque cover does not elicit direct reaching behavior and so requires more developed means–ends behaviors for success. Given this alternative interpretation that can hold for any means–ends task requiring handling of the cover, the greater success in the partially hidden and transparent conditions over the fully occluded condition does not necessarily challenge the means–ends explanation.

To avoid the alternative interpretation of nonequivalent means–ends abilities required for success in partially hidden–transparent versus completely hidden–opaque cases, we used tasks in which no advantage was produced by direct reaching for the toy. These tasks required 7-month-old infants to either pull a towel or push a button to retrieve a distant toy. Prior to testing, infants were trained on the means–ends behaviors required for toy retrieval. Although these behaviors might be considered relatively advanced for infants of this age according to certain means–ends scales (Uzgiris & Hunt, 1975), training proved to be effective for many of the infants. During the test, a screen was placed between infants and the toy; this screen was either transparent or opaque. Our prediction was that infants would fail more often in the opaque condition than in the transparent condition, indicating that their difficulties in the opaque condition were not based simply on means–ends deficits.

Experiment 1

Method

Participants

Twelve full-term 7-month-olds (7 months, 5 days to 7 months, 20 days, mean age = 7 months, 10 days) participated in the experiment. There were 9 boys and 3 girls in the group. An additional 12 infants were excluded because they either failed to pass the criteria to move into the testing phase of the experiment (10 participants)[3] or because they became upset during testing (2 participants). The participants had no known or suspected abnormalities and came from predominantly middle-class suburban families.

Stimuli

The stimuli were baby toys of various sizes, shapes, and colors such as toy phones, hammers, and cars. During the training phase of the study, the experimenter used toys from a set of 10 toys, in no set order. The infant's interest in certain types of toys or parents' comments about the infant's preferences were considered in toy selection during this training period. Using a different fixed set of toys in the test phase, the experimenter presented the toys in the same sequence to each infant. This consistency in presentation allowed for proper counterbalancing (described in the following section).

Design

The experiment involved a 2 × 2 within-subjects design, with screen type (opaque or transparent) crossed with toy presence (toy or no-toy). The no-toy condition served as a comparison for the toy condition and ensured that infants would not simply learn to pull the towel on every trial as a conditioned response that always yielded rewards. The experiment included 7 of each of the four types of trials—opaque toy, opaque no-toy, transparent toy, and transparent no-toy—for a total of 28 trials for each infant. The trial types were randomly ordered by blocks. To ensure that the desirability of certain toys would not be confounded with screen type for reaching behavior, toys were counterbalanced by screen type. That is, for any given toy, each participant saw it on the same trial; however, for half of the participants, the toy was occluded by the opaque screen, and for the other half, the transparent screen was placed in front of the toy.

Apparatus

Toys were placed on a 51 in. × 28 in. lavender towel sewn in several places across its width to produce ridges for easier grasping (Figure 2). Thick black poster boards enclosed the table on three sides, rising 30 in. above the table surface, to focus the infant's attention on the toy. A transparent screen was cut from acetate and an opaque screen from

---

[3] One might argue that the number of infants failing to pull the towel to criterion to retrieve the toy supports the means–ends explanation for infants' failure to retrieve occluded objects. However, it should be noted that the infants' failure to pull the towel could be due to many factors (one infant, for example, failed to pull the towel to retrieve toys during the training period, but afterwards quickly and easily pulled the towel to retrieve her bottle!). It thus becomes difficult to draw conclusions from the behaviors of the nonretrieving infants. More important, the drop-out rate (which is not unusual for studies of infancy) does not weaken our argument that even after infants pass the relevant means–ends criteria, other factors prevent them from successfully retrieving occluded toys.

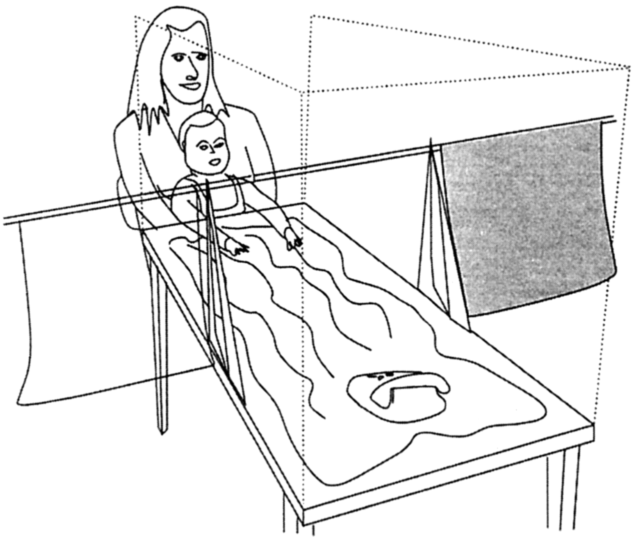

Figure 2. Apparatus for testing means–ends account. Infants were trained to pull the towel to retrieve the distant toy. Infants were then tested on retrieving toys from behind the transparent and opaque screens.

denril; both of these screens were 24 in. × 18 in. The screens were taped to a wooden dowel, with 30 in. (the width of the table) between them. Triangular wooden structures clamped to the table supported the dowel. To pull screens in and out of view of the infant, the dowel was moved horizontally. All of the wooden components of the apparatus were spray-painted black, to blend in with the black poster boards.

Procedure

The experiment was videotaped by a ceiling-mounted camera for later analysis. There were two phases to the experiment: a means-ends training phase and an object retrieval testing phase.

Means–ends training phase. Each infant was tested individually in the laboratory, while seated on the parent's lap. The experimenter explained to the parent that one goal of the study was to teach the infant to pull the towel to retrieve the distant toy. Parents were encouraged to help the infant achieve this goal in any way that they felt was appropriate. The experimenter suggested that the parents might demonstrate to the infant how to retrieve the toy, put their hands over their infant's hands and guide the infant through the proper motions, and verbally cheer on the infant. The experimenter then placed a toy on the far end of the towel, directed the infant's attention to it, and observed the infant's reaching behavior. A new toy was introduced when the infant seemed to lose interest in the previous one. Infants were trained on the toy-retrieval task in this way until they pulled the towel to retrieve the distant toy within 20 s, on two consecutive trials. Training took 9 min on average.

Object retrieval testing phase. The experimenter explained that the parent should not encourage or assist the infant in any way during the subsequent part of the study. In addition, the parent was asked to prevent the infant from reaching for the towel until the screen was in place. Without this momentary constraint on the infant's movements, the infant might have retrieved toys from behind the opaque screen by merely following through with reaching behaviors initiated prior to the occlusion of the toy, independent of any kind of understanding that the toy continued to exist behind the screen. Analysis of the videotapes revealed that the parent constrained the infant for an average of 4 s (from toy–hand introduction to screen in final position), resulting in an average delay of half a second for the infant between toy viewing and towel-pulling opportunity.

At the start of each trial, the infant's attention was directed to the distant end of the towel, with the shaking of a toy or, on the no-toy trials, with the tapping of the experimenter's hand. Then, either the opaque or the transparent screen was placed between the infant and the toy. Each toy trial continued until either the infant had successfully retrieved the toy or until 20 s had passed, whichever came first. Similarly, each no-toy trial continued either until the infant had pulled the entire towel from behind the screen, or until 20 s had passed, whichever came first. If at 20 s the infant was engaged in retrieval behavior, the trial was prolonged until the behavior was completed.

It should be noted that although parents in this study were not made blind to the trial types, it is unlikely that any parental bias would have favored our hypothesis that toy-guided retrieval would be greater in the transparent versus opaque condition. The difference between parents' and infants' knowledge was presumably greater in the opaque condition, so parental bias probably would have had its greatest effects on efforts to retrieve toys in the opaque condition. If parents acted on their knowledge, their actions would have biased the data against our hypothesis.

Coding

Each trial was coded for whether a retrieval was completed, and if so, the time to complete the retrieval. The trials were timed from the point at which the screen was in place. In the no-toy trials, a (pseudo-) retrieval was considered completed when the entire towel had been pulled from behind the screen. In the toy trials, a retrieval was considered completed either when the entire towel had been pulled from behind the screen or when the infant had first touched the toy. Using the metric of pulling the entire towel from behind the screen provided consistency in the measure of completed retrievals across toy and no-toy trials. However, in the toy trials, infants sometimes retrieved the toy without pulling the entire towel from behind the screen; coding of the toy trials thus included the first touch of the toy as another criterion for a completed retrieval.

The difference between number of retrievals on toy and no-toy trials was used as a measure of toy-guided retrieval. This measure of toy-guided retrieval reflects infants' discriminating retrieval responses; it also controls for simple preferences for screen type. Toy-guided retrieval was predicted to be higher in the transparent condition than in the opaque condition.

Results

See Figure 3 for a graph of the retrieval results. The average number of retrievals completed out of seven possible was 5.1 (73%, SE = 8.5) in the transparent toy condition, 2.3 (33%, SE = 9.2) in the transparent no-toy condition, 3.2 (46%, SE = 9.4) in the opaque toy condition, and 2.5 (36%, SE = 8.3) in the opaque no-toy condition. As predicted, infants completed more toy-guided retrievals in the transparent condition than in the opaque condition. That is, infants showed differentially more retrieval responses on toy versus no-toy trials in the transparent condition (M = 2.75) than in the opaque condition (M = .67), $t(11)$ = 2.076, p = .03, one-tailed. A coder who was blind to the purposes and hypotheses of the study coded half of the data. Interrater reliability for these items was .99, with coders agreeing on 166 out of 168 trials.

Trials on which a retrieval was completed were analyzed for time to complete the retrieval. The average time to complete retrievals was 11.5 s (SE = 1.1) in the transparent toy condition, 12.2 s (SE = 2.0) in the transparent no-toy condition, 10.7 s

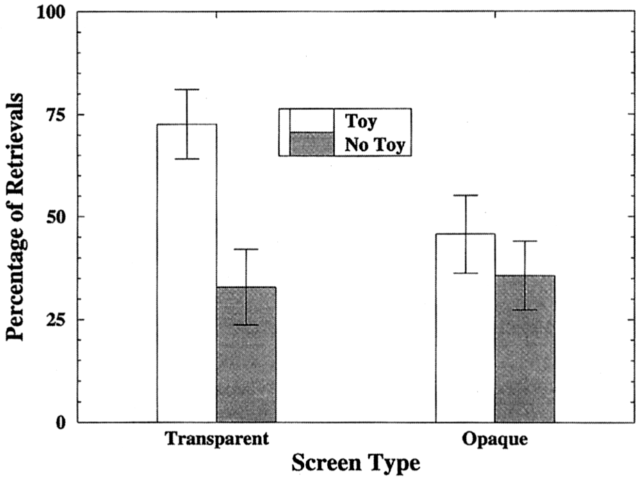

Figure 3. Retrievals by screen type and toy presence in Experiment 1. Infants completed more toy-guided retrievals (toy–no-toy) in the transparent condition than in the opaque condition. The descending lines indicate standard error.

($SE = 1.5$) in the opaque toy condition, and 12.8 s ($SE = 1.0$) in the opaque no-toy condition. There was no difference in toy-guided time to complete retrievals between the transparent condition ($M = -.73$) and opaque condition ($M = -2.04$), $t(7)^4 = 0.612$, $p = .56$. That is, infants did not differ between transparent and opaque conditions in the effect of toy presence on time to retrieve. This pattern of results was consistent across subsequent experiments and will not be reported on further.

Discussion

The means–ends deficit theory predicts that infants should show similar toy-guided retrieval in the transparent and opaque conditions because the means–ends abilities required for success in the two conditions are identical. However, infants showed more toy-guided retrievals in the transparent condition, indicating that their difficulties in the opaque condition were not based simply on means–ends deficits.

One might argue, however, that the differences in the infants' behaviors in the opaque and transparent conditions were due to differences between the conditions other than their dependence on object permanence understanding. Two such possible factors, the *wall* and *different means–ends* alternatives, are described next.

One might argue that the infants failed to retrieve the toy in the opaque task because they mistakenly thought of the opaque screen as a solid wall (Baillargeon, personal communication, March 1993). According to this explanation, the infant knows that the toy is behind the opaque wall but decides that the toy cannot penetrate the wall and so does not reach. It should be noted that such a belief about the wall's penetrability would not necessarily stop the infant from trying to retrieve the toy. As previously mentioned, when infants try to retrieve an object from underneath a transparent box, they exhibit a remarkably strong tendency to reach directly along the line of sight (coming into contact with a wall of the box), rather than reaching through an open face of the box (Diamond, 1981, 1991). With the transparent box, the infant receives direct tactile feedback about the presence of the wall yet still perseveres in trying to retrieve the toy with direct reaching. Thus, even if the infants in the current experiment believed that the opaque screen were a wall, it is not clear that this belief would prevent them from trying to retrieve the toy.

Alternatively, the process of pulling the towel to retrieve the toy may not be a means–ends process but actually more of a hill-climbing process, whereby infants evaluate the towel-pulling process through feedback about the movement of the toy. In this case, infants might have a strategy of "Pull the towel once. If I receive positive feedback, pull again," rather than "Pull the towel to retrieve the toy." According to this argument, infants' internal representations of the toy might actually be the same in the occluded and transparent cases, but the external feedback about the success of retrieval behaviors is not. That is, infants may have complete knowledge about the presence of the toy but need continual feedback about the success of their towel pulling to continue retrieval behaviors. Infants can receive such feedback in the transparent, but not opaque, screen conditions and so reach differentially in the two types of trials.

Experiment 2

To systematically address the wall and feedback alternatives, we ran a second experiment in which infants were required to push a button to retrieve a distant toy (Figure 4). As in Experiment 1, each infant was seated on the parent's lap throughout the experiment. The toy sat on a ledge that was too far from the infant to reach directly. When the nearby button was pushed, the ledge dropped and the toy slid down a ramp to the infant. In this apparatus, the transparent and opaque screens were both rigid. When the ledge dropped, the toy slid under the screen, rather than through it. Thus, in this case, the perceived attainability of the toy should not have been affected by whether or not the infant realized that the screens were wall-like. In addition, the means–ends behavior and response were immediate: If the button was pushed, the toy was obtained. Thus, there was no advantage of continuous feedback in the transparent condition.

Method

Participants

Twelve full-term 7-month-olds (7 months, 1 day to 7 months, 16 days, mean age = 7 months, 8 days) participated in the experiment. There were 6 boys and 6 girls in the group. An additional 10 infants were excluded either because they failed to pass the criteria to move into the testing phase of the experiment (9 participants)[5] or due to experimenter

---

[4] The degrees of freedom are reduced because infants with missing data points (indicating no reaches in at least one condition) cannot be included in the analysis.

[5] Again, it should be noted that the infants' failure to push the button could be due to many factors, so that it is difficult to draw conclusions from the behaviors of the nonpushing infants. For example, several of these infants seemed to be afraid of the noise of the ledge dropping and the speed of the approaching toy and thus appeared to be avoiding the button rather than failing to link the button to the ledge drop.

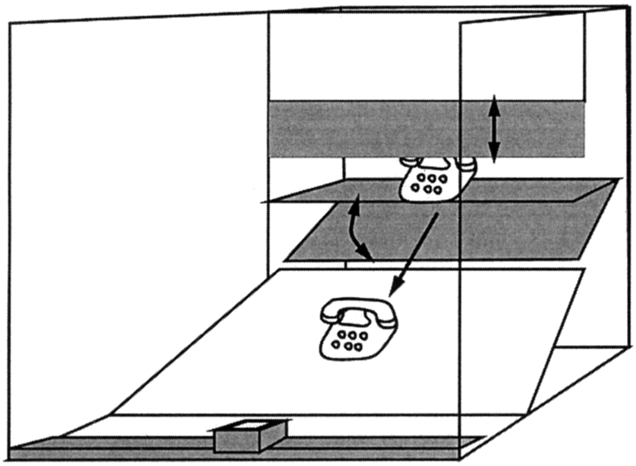

Figure 4. Apparatus for testing wall and feedback alternatives. Infants were trained to push the button to retrieve the distant toy on the ledge. When the button was pushed, the ledge dropped and the toy slid down the ramp to the infant. Infants were then tested on retrieving toys from behind the transparent and opaque screens that were lowered in front of the ledge. Toys slid under the screens, and feedback was immediate.

error (1 participant). The participants had no known or suspected abnormalities and came from predominantly middle-class suburban families.

Stimuli

The stimulus set consisted of the toys from Experiment 1 that were both tall enough to be seen on the ledge and smooth enough to slide down the ramp, along with additional toys fitting these criteria.

Design

The design was identical to that in Experiment 1, with 7 of each of the four types of trials—opaque toy, opaque no-toy, transparent toy, and transparent no-toy—for a total of 28 trials for each infant.

Apparatus

Toys were placed on a 29 in. × 12 in. ledge (Figure 4). Black wooden boards enclosed the table on three sides, rising 36 in above the table surface, to focus the infant's attention on the toy. The transparent and opaque screens were cut from glass, with the opaque screen painted white. Both of these screens were 28 in. × 10 in. The screens were supported by fishing wire attached to wooden handles that rested on the rear wall of the apparatus. To pull screens in and out of view of the infant, these handles were raised and lowered. In order to keep parents blind to the trial types, a black screen measuring 30 in. × 10 in. was attached to the front of the apparatus. The screen blocked the parent's view of the ledge area and was situated behind the infant's head so that the infant's view was unobstructed. Parents were also asked to wear headphones that connected to a tape player located behind the apparatus. The tape had music and toy sounds on it to mask the sounds of toys handled by the experimenter. With the exception of the transparent and opaque screens, all of the visible components of the apparatus were black to blend in with the black wooden boards.

The ledge for the toy was supported by solenoids. Pushing the button completed the electric circuit and caused the solenoids to retract and the ledge to fall, creating a 29 in. × 4.5 in. aperture under the screen, through which the toy could slide. An additional switch at the rear of the apparatus allowed the experimenter to effectively turn the power of the button on or off (see *Procedure*). A small light bulb was wired to the solenoid circuit and attached to the side of the apparatus, to indicate when the button had been pressed. This light bulb was out of view of the infant and parent, but it was detectable by the ceiling-mounted camera.

Procedure

The experiment was videotaped by a ceiling-mounted camera for later analysis. There were two phases to the experiment: a means–ends training phase and an object-retrieval testing phase.

Means–ends training phase. The first part of the training phase was the same as that in Experiment 1 except that the toy was placed on the ledge rather than on the end of a towel. Infants were trained on the toy-retrieval task until they were able to push the button to retrieve the distant toy within 10 s on two consecutive trials. Infants were then presented with alternating no-toy and toy trials until they retrieved the toys in two toy trials within 10 s. On no-toy trials, the experimenter tapped the ledge with her hand. Pilot testing had indicated that the dropping of the ledge itself, without any toy on it, was interesting enough to some of the babies that they pushed the button equally often in toy and no-toy conditions. To encourage greater discrimination between these conditions, the switch at the rear of the apparatus was set so that the button had no power (i.e., pushing the button had no effect on the ledge, though it did illuminate the side bulb) during this and all subsequent no-toy trials. No-toy trials were included in the training phase to demonstrate to the infants that pushing the button did nothing when there was no toy present and thus to possibly reduce their exploratory button-pushing in the no-toy trials during the test phase of the experiment. Training took 12 min on average.

Object-retrieval phase. The parent was asked to listen to the tape of music and toy sounds through headphones so that the sounds of toys handled by the experimenter were masked. The tape also served to cue the parent to the start of a trial; the tape played while the trial was set up and stopped when the screen was fully in place. As in Experiment 1, the experimenter explained that the parent should not encourage or assist the infant in any way and asked the parent to prevent the infant from reaching for the button until the screen was fully in place (i.e., when the tape player stopped). Analysis of the videotapes revealed that the parent constrained the infant for an average of 5 s (from toy–hand introduction to screen in final position), resulting in an average delay of 2 s for the infant between toy viewing and button-pushing opportunity.

At the start of each toy trial, the experimenter placed a toy on the ledge and then directed the infant's attention to the ledge by tapping her hand along it. On no-toy trials, the experimenter simply tapped her hand along the ledge. Then, either the opaque or the transparent screen was placed between the infant and the toy. Each toy trial continued until either the infant had successfully retrieved the toy, or until 10 s had passed, whichever came first. Each no-toy trial continued until 10 s had passed.

Coding

Each trial was coded for whether the button was pushed. A button push was measured by the illumination of the side light that indicated that the button had been pressed to complete the circuit (see *Apparatus*).

As in Experiment 1, the difference in number of retrieval responses between toy and no-toy trials was used as a measure of toy-guided retrieval. Toy-guided retrieval was predicted to be higher in the transparent condition than in the opaque condition.

Results

The pattern of results was similar to that from Experiment 1 (Figure 5). The average number of retrievals completed out of seven possible was 3.8 (54%, SE = 8.8) in the transparent toy condition, 2.8 (40%, SE = 8.0) in the transparent no-toy condition, 2.8 (40%, SE = 7.8) in the opaque toy condition, and 3.0 (43%, SE = 6.3) in the opaque no-toy condition. As predicted, infants completed more toy-guided retrievals in the transparent condition than in the opaque condition. That is, infants showed differentially more retrieval responses on toy versus no-toy trials in the transparent condition (M = .92) than in the opaque condition (M = −.17), $t(11) = 2.00, p < .04$, one-tailed. A new coder who was blind to the purposes and hypotheses of the study coded half of the data. Interrater reliability for these items was exactly the same as that in Experiment 1, .99, with coders agreeing on 166 out of 168 trials.

Discussion

Again, the means–ends deficit theory predicts that infants should show similar toy-guided retrieval in the transparent and opaque conditions because the means–ends abilities required for success in the two conditions are identical. However, infants showed more toy-guided retrievals in the transparent condition, indicating that their difficulties in the opaque condition were not due to means–ends deficits alone, a belief that toys could not be pulled through the opaque wall, or the lack of continuous feedback about the effects of their behaviors on the toy. The results from Experiments 1 and 2 are thus consistent with the idea that infants' difficulties in the opaque condition are not based simply on means–ends deficits.

However, one might argue that training infants on the retrieval of visible toys in Experiments 1 and 2 led them to generalize better to the transparent condition. According to this argument, when the toy is occluded behind the opaque screen at test, infants know that the toy is there but do not know that they can retrieve it with a towel pull or button push because these means–ends behaviors are associated only with visible toys.[6] Experiment 3 tests this preferential training hypothesis.

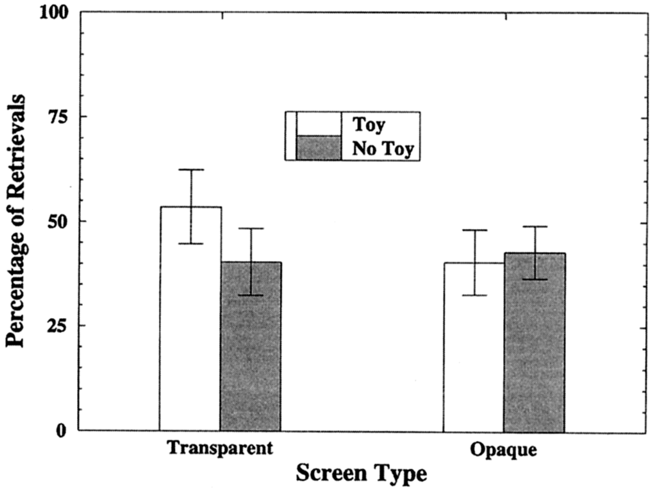

Figure 5. Retrievals by screen type and toy presence in Experiment 2: Infants completed more toy-guided retrievals (toy–no-toy) in the transparent condition than in the opaque condition. The descending lines indicate standard error.

Experiment 3

To test whether infants' difficulties with opaque screens in Experiments 1 and 2 were based on a failure to generalize from training, we ran a third experiment in which infants were trained on the button-push apparatus, without toys, to push the button to release the ledge. Equal exposure was given to an opaque and a transparent screen in front of the ledge. Following a brief demonstration of the effects of a button push on both visible and hidden toys, infants were then tested on toy-guided retrieval. Infants thus received equal exposure to no-toy trials under transparent and opaque conditions as well as to toy trials under transparent and opaque conditions. The training in the transparent and opaque conditions was equivalent, so that differences at test cannot be explained in terms of a failure to generalize from training.

Method

Participants

Twelve full-term 7-month-olds (7 months, 5 days to 7 months, 15 days, mean age = 7 months, 10 days) participated in the experiment. There were 4 boys and 8 girls in the group. An additional 11 infants were excluded because they either failed to pass the criteria to move into the testing phase of the experiment (9 participants)[7] or because they became upset during testing (2 participants). The participants had no known or suspected abnormalities and came from predominantly middle-class suburban families.

Stimuli, Design, and Coding

The stimuli, design, and coding were identical to those in Experiment 2.

Apparatus

The apparatus from Experiment 2 was modified for this study. Because infants were to be trained without toys on the ledge release, an attempt was made to increase the salience of this release by attaching a shiny red cloth to the ledge and its supports. All other modifications were made in an attempt to reduce the noise in the data-collection procedure. In Experiment 2, infants frequently seemed to push the button inadvertently, often as they were in the process of turning away from the apparatus. Infants also showed a tendency to get distracted from the task by the black screen behind them that served to keep parents blind to the trial condition and by the headphones worn by the parents during test. To reduce the number of inadvertent button pushes, the entire apparatus shown in Figure 4 was moved 2 in. back from the edge of the table, away from the infant. The distracting black screen was removed, and parents instead wore a blindfold during the test trials. To reduce the novelty of the headphones at test and to avoid similar problems with the blindfold, the headphones and blindfold were introduced at the start of the experiment. Parents were not made blind during training (no sounds were played through the headphones, and parents wore the blind-

---

[6] We thank Andrew Meltzoff and John Flavell for independently suggesting this possibility.

[7] Again, it should be noted that the infants' failure to push the button could be due to many factors, so that it is difficult to draw conclusions from the behaviors of the nonpushing infants. For example, the ledge drop did not seem intrinsically interesting to several of these infants.

fold over their foreheads); infants were simply given the opportunity to become accustomed to the headphones and blindfold in the environment.

Procedure

There were two phases to the experiment: a means–ends training phase and an object-retrieval testing phase.

Means–ends training phase. The means–ends training phase in this experiment involved learning, habituation, and demonstration components. The learning component was designed to allow the infant to learn the appropriate means–ends behavior (push the button to make the ledge fall). The habituation component was designed to allow the infant to subsequently habituate to the button-pushing behavior so that the button would not be pushed indiscriminatingly during test. The demonstration component was designed to simply show the infant that the ledge's fall would bring a toy on the ledge within reach.

In the learning component of the training phase, the experimenter tapped her hand along the ledge and then the transparent or opaque screen was lowered. The sequence of the screens was randomly ordered by blocks, so that any odd-numbered trial together with the subsequent trial would involve both types of screen. These odd–even pairs are referred to as both-screen pairs. Each trial ran until the infant, parent, or both pushed the button, in an attempt to give the infant equal exposure to the button–ledge relation in the opaque and transparent conditions. Infants were trained to push the button to make the ledge fall until they were able to do so within 10 s, on two trials in a both-screen pair. From this point, an attempt was made to habituate infants to the button push. Cheering following the button push was no longer permitted. Each trial was run for 10 s. The sequence of habituation trials continued until either eight trials had passed or the infant had shown signs of habituation by failing to push the button within 10 s on two consecutive trials in a both-screen pair, whichever came first. This habituation phase was followed by a brief demonstration of the effects of the ledge's fall on the toy. The experimenter placed a toy on the ledge, tapped along the ledge, and lowered a screen. The parent then pushed the button, causing the toy to slide down the ramp. The toy was quickly returned to the experimenter and the demonstration repeated three times, for a total of four demonstration trials. Two of these demonstrations occurred with an opaque screen, two with transparent, so that the infant was given equal exposure to toy retrieval following a button push in the opaque and transparent conditions. The ordering of the screen types for the entire means–ends training phase was counterbalanced across participants. Total training time was 6 min on average.

Object-retrieval phase. The object-retrieval phase was similar to that of Experiment 2, with one exception. In Experiment 2, a button push had no effect on the ledge in the no-toy trials. This manipulation was made in an attempt to reduce infants' exploratory button pushing in the no-toy trials. In the current experiment, because infants had been trained to expect the ledge to drop in no-toy trials, such a manipulation might instead have increased exploratory button pushing in the no-toy trials as infants tested their violated predictions. Thus, in the current experiment, the button push caused the ledge to drop in both toy and no-toy trials. Analysis of the videotapes revealed that the parent constrained the infant for an average of five seconds (from toy–hand introduction to screen in final position), resulting in an average delay of 2 s for the infant between toy viewing and button-pushing opportunity.

Results

The pattern of results was similar to that from Experiments 1 and 2 (Figure 6). The average number of retrievals completed out of seven possible was 4.7 (67%, $SE = 8.5$) in the transparent toy condition, 2.9 (41%, $SE = 8.3$) in the transparent no-toy condition, 3.6 or (51%, $SE = 9.2$) in the opaque toy condition, and 3.5 or (50%, $SE = 8.5$) in the opaque no-toy condition. As predicted, infants completed more toy-guided retrievals in the transparent condition than in the opaque condition. That is, infants showed differentially more retrieval responses on toy versus no-toy trials in the transparent condition ($M = 1.75$) than in the opaque condition ($M = .083$), $t(11) = 3.46$, $p < .003$, one-tailed. Another coder who was blind to the purposes and hypotheses of the study coded half of the data. Once again, interrater reliability for these items was .99, with coders agreeing on 166 out of 168 trials.

Discussion

Again, the means–ends deficit theory predicts that infants should show similar toy-guided retrieval in the transparent and opaque conditions because the means–ends abilities required for success in the two conditions are identical. However, after receiving equal training with the transparent and opaque screens without toys, and equal demonstration with the transparent and opaque screens with toys, infants again completed more toy-guided retrievals in the transparent condition than in the opaque. These results demonstrate that infants' difficulties with opaque screens were not based on a failure to generalize from training. This single experiment thus demonstrates that infants' greater success in toy-guided retrievals in the transparent condition was not based on means–ends deficits alone, a belief about the opaque screen as a wall, a continuous feedback advantage to the transparent condition, or differential generalization from training on visible toys to testing in the transparent condition. These results are completely consistent with the findings from Experiments 1 and 2. The effects are clearest in Experiment 3, perhaps due to the simple changes made to the apparatus to reduce distractions and inadvertent button pushes, as discussed in the *Apparatus* section.

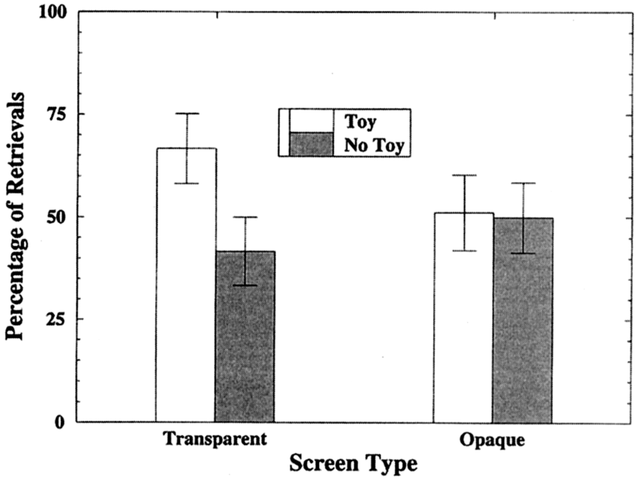

Figure 6. Retrievals by screen type and toy presence in Experiment 3: Infants completed more toy-guided retrievals (toy–no-toy) in the transparent condition than in the opaque condition. The descending lines indicate standard error.

Discussion of Experiments

In this series of three experiments, we have attempted to demonstrate the insufficiency of the principle-based, means–ends explanation to account for infants' retrieval behavior. Infants in all three experiments failed to carry out the same means–ends behavior required to retrieve occluded toys that they employed in the retrieval of visible toys. Both experiments with the button apparatus demonstrate that the greater success with the toy visible was not based on either a belief about the rigidity of the opaque screen or continuous feedback in the visible condition. Furthermore, Experiment 3 indicates that the greater success in the visible case was not based on preferential training to visible toys. These findings call into question the standard conclusion from looking-time studies that infants have a principle of object permanence from at least the first few months of life and that only means–ends deficits prevent infants from acting on this knowledge to retrieve hidden toys. Our studies demonstrate that deficits in means–ends behaviors do not pose problems when the processing of hidden objects is not required.

We do not claim that means–ends abilities are static during the period when infants progress from showing longer looking times to impossible events with hidden objects to eventually reaching for hidden objects. Improvements in means–ends behaviors may well contribute to improvements in infant search—after all, infants in our experiments required training to learn the relevant means-ends behaviors. In addition, development in means–ends behavior is perfectly compatible with our adaptive processing approach, in which performance is a function of both task-specific factors (such as means–ends behaviors) and underlying representational systems. What our experiments show is that means–ends development alone cannot account for the fact that infants eventually retrieve hidden objects. Even after infants learned the relevant means–ends behaviors, they failed to demonstrate toy-guided retrieval under occluded conditions. Thus, the means–ends account is insufficient for understanding infants' retrieval behavior. The adaptive process approach allows us to consider how the state of underlying representations may contribute to infant behavior. We turn to this approach in the next section.

Although our experiments demonstrate that means–ends deficits alone cannot explain the looking–reaching task dependency in object permanence, one could still argue that other factors ancillary to a concept of object permanence play a role in infants' failures to reach for occluded objects. We have ruled out several such possibilities with Experiments 2 and 3. In this discussion, we consider two additional types of ancillary deficit accounts.

One possibility is that this ancillary factor is motivation to reach. According to this account, infants in our experiments had the same knowledge about the toys' presence in the opaque and transparent conditions but were more motivated to retrieve when they could see the toys. Such an interpretation could account for increased likelihood of responding in the toy-present–transparent condition versus the toy-present–opaque condition. However, under the motivation interpretation, it is difficult to explain why infants responded in the two opaque conditions of Experiment 3 (tending to press the button more often in the two opaque conditions relative to the toy absent–transparent condition), while they showed no difference in the probability of button pressing between the toy-present and toy-absent conditions when the screen was opaque. These data seem to suggest that infants were at least somewhat motivated to retrieve toys in the opaque conditions, but they were insensitive to whether a toy was present and so made retrieval responses indiscriminately, independent of toy presence. In a similar vein, it is difficult to see how a motivational account could be applied to infants' perseverative responses in the $A\bar{B}$ task. Even after infants successfully retrieve hidden objects, they still often show the perseverative error of reaching to an original hiding location (A) when the object is hidden in a new location (B). Once infants can successfully retrieve hidden objects in a single hiding location, motivational limitations regarding actions toward hidden objects have presumably been overcome. Infants at this point should still have a concept of object permanence, as well as their newfound capabilities to act on unseen objects, and so it becomes difficult to explain why they reach to incorrect locations.

One could always introduce new abilities or deficits to explain away each inconsistency. For example, one might argue that infants have the object permanence concept early in life, fail to retrieve hidden objects because of motivational deficits, fail to discriminate toy presence and absence behind an opaque screen because of some other limitation, and then later fail the $A\bar{B}$ task because of yet another limitation that is overcome when the infants succeed on this task. This type of "one finding, one explanation" approach can be ad hoc, making the phenomena seem like a list of unrelated facts. Rather than relying on a host of ancillary deficits to explain behavior, we think it is more worthwhile to seek a unified framework for understanding infants' developing sensitivity to the permanence of objects.

One might argue for a unified framework in which infants have a principle of object permanence from the first few months of life, together with a general ancillary capability that develops gradually. For example, improvements in memory, attention, resources or processing capacity might underlie infants' increasing ability to show object permanence knowledge. The notion is that performance in object permanence tasks requires both the knowledge of the principle of object permanence and some ancillary factor. To make this idea work when the proposed ancillary factor is memory, for example, one could assume that the basic principle of object permanence is present from birth but that performance in object permanence tasks requires both this principle of object permanence and the ability to retain in memory a representation of whether a toy was present before a screen dropped on a particular trial. Knowledge of the principle of object permanence (objects continue to exist even after occluded), together with memory that a toy was present before the screen dropped, would allow the infant to infer that a toy is present and thus to respond. However, if the infant were to forget whether in fact a toy had been present before the screen dropped, the principle of object permanence by itself would be of little help. On this view, memory improves with age such that infants gradually get better at remembering hidden objects over longer delays and in greater detail. One might explain infants' successes in looking-time tasks before reaching tasks because the latter somehow require more memory. Similarly, one might suppose that attention improves with age so that infants gradually get better at paying attention to what they know about hidden toys. Perhaps infants in our experiments had the same knowledge about the toys' presence in the visible and occluded conditions but were unable to attend to this knowledge in the occluded condition. One could explain success on looking-time tasks before reaching tasks because the looking-time tasks cue infants to attend to their knowledge of hidden objects in a way that the reaching tasks do not. Finally, one might also propose that some general resource or processing capacity increases with age. In work on language functions in adults, Just and Carpenter (1992) have suggested that there is a general resource pool that affects both the ability to carry out cognitive processes and the ability to maintain information in memory. If we apply these ideas to object permanence, with the added assumption that this general resource pool grows with age, we could explain the fact that infants succeed on looking-time measures of object permanence at an earlier age than they succeed in reaching measures by suggesting that the latter require more resources. Reaching requires maintenance of memory of the object, inference that the object is still there, and execution of reaching behavior. Although looking measures may require memory and inference, they require only a relatively passive response that might arguably require fewer resources than reaching.

It seems likely that memory, attention, or capacity-based accounts could be constructed that would account for the looking–reaching task dependency as well as for other aspects of behavior in object permanence tasks, such as the $A\bar{B}$ error. Thus, accounts based on gradual development of one or more of these factors could be viewed as having the potential to provide a unified framework for understanding infants' development in object permanence tasks. However, we believe that principle-based accounts that attribute developmental changes to ancillary factors tend to finesse critical questions about cognitive development (see Thelen & Smith, 1994, for related arguments). Although the focus of theory and experiment is directed toward the underlying principles, such as object permanence, these principles do not ultimately carry the explanatory burden of accounting for developmental change. As a result, these accounts remain underspecified in several ways. First, the ancillary systems that hold all of the explanatory power can go unspecified because these systems are outside the theoretical domain of interest. Second, even when attempts are made to specify the nature of the ancillary capabilities (e.g., Case, 1985; Halford, 1993; Just & Carpenter, 1992; Pascual-Leone, 1970), there is little discussion of mechanisms leading to changes in these capabilities. Although some accounts point to the maturation of various brain regions (e.g., Case, 1992; Diamond, 1991), they do not specify the mechanisms by which such maturation might result in the global improvements that are posited. These accounts also tend to obscure the role that experience may play in causing developmental change. Change in the ancillary factor may explain change in behavior, but what explains the change in the ancillary factor itself? Often, the factors are so global that it is tempting to view them as changing through an experience-independent process akin to growth. Although such explanations cannot, of course, be ruled out a priori, we believe there is sufficient evidence for the role of experience in developmental change to consider alternatives.

Exploring Adaptive Process Accounts

In what follows, we consider what new kinds of accounts of developmental change become accessible if we abandon principle-based characterizations of the knowledge that underlies performance in cognitive tasks. We put forth one possible account of the way in which developmental changes in performance in object permanence tasks may arise from changes in a system that learns to represent objects. We demonstrate through simulations how changes in the connections underlying processing in our system result in stronger representations of occluded objects. We show how such changes might be driven by experiences conforming to the principle of object permanence. In so doing, we hope to demonstrate that developmental changes in ancillary factors such as capacity, memory, and attention need not be assumed to account for the looking–reaching task dependency. Improvements attributed to these factors may be a natural consequence of the strengthening of underlying connections.

Simulation Modeling

We explore a specific adaptive process model of performance in object permanence tasks in which patterns of activity representing objects must be maintained across delay periods during which there is no perceptual support for the representations. We use simulations to demonstrate the following points:

1. An adaptive processing system can gradually improve its ability to retain information about occluded objects through experiences with objects that conform to the principle of object permanence (i.e., experiences in which objects that disappear when occluded reappear when the occluder is removed).
2. Such a system's predictions of reappearance of occluded objects are graded in nature and become weaker with longer delay.
3. Improvements in the performance in this system depend on the strengthening of the system's ability to maintain internal representations of occluded objects.
4. Learning in such a system can support generalization, in that the ability to make predictions for objects used in training can be applied to novel objects.

We then use this model to explore a possible account of the contrast between looking time and reaching measures of object permanence. We show how the representations in the model can be strong enough to support predictions based on hidden objects and reaching based on visible objects but not reaching based on hidden objects. The simulations thus illustrate one further point:

5. An adaptive processing system can behave in task-dependent ways because tasks depend differentially on the system's ability to maintain strong internal representations of occluded objects.

The simulations are based on parallel distributed processing (PDP) models, in which processing occurs through the propagation of activation among simple processing units. The processing capabilities of such models depend on the connections between the units, which are inherently graded in strength. Experience leads to gradual adaptive changes in these connections, thereby leading to gradual increases in the capabilities of the processing system.

All of the simulations reported here involve simple recurrent networks (Elman, 1990; Jordan, 1986) that are trained to anticipate the future positions of objects in a very simple, simulated visual display. The network architecture is shown in Figure 7. Details on the simulations are provided in the Appendix. The network "sees" sequences of inputs corresponding to simple events. On each time step, the input specifies the identity and location of one or two objects—a "barrier," a "ball," or both.[8] When the ball is present, it sits at a discrete point in the network's visual field. When the barrier is present, it moves back and forth across the network's visual field. When both the barrier and the ball are present (Figure 8), the barrier passes in front of the ball and occludes it (time4), then moves one step further (time5). The barrier may remain in that position for up to four time steps or, as shown in Figure 8, may begin to move back in the other direction immediately. As the barrier begins to move (time6), the network can anticipate the reappearance of the ball at the next time step. It is on the basis of learning to correctly predict this reappearance when the ball was actually present in the earlier time steps that the network comes to exhibit knowledge of the principle of object permanence.

Clearly, these simulated events were not meant to capture the full range of infants' visual experiences. Instead, our aim was to explore new ways of thinking about the mechanisms underlying infants' behaviors in object permanence tasks and the processes through which these mechanisms develop. The simple causal structure of the simulated events used here is sufficient to induce a rudimentary model based on the reappearance of occluded objects when the occluder is removed. In real infants, we would expect that the richer experience available to them leads to the formation of richer models of the causal structure of events involving occluded objects.

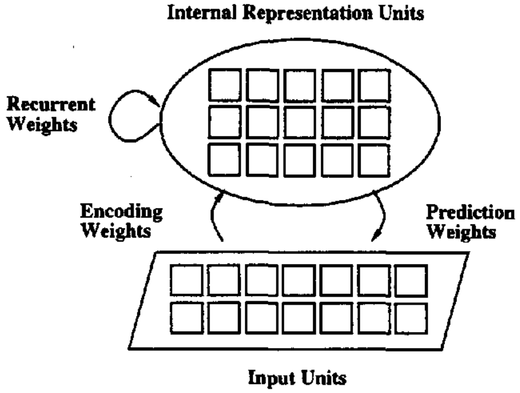

Figure 7. Recurrent network for learning to anticipate the future position of objects. The pattern of activation on the internal representation units is determined by the current input and by the previous state of the representation units by means of the encoding weights and the recurrent weights respectively. The network sends a prediction back to the input units to predict the next state of the input. The stimulus input determines the pattern of activation on the input units, but the difference between the pattern predicted and the stimulus input is the signal that drives learning.

A single ball unit was used in the training and testing of these networks; we later present simulations in which the network was trained with more complex objects to test its ability to generalize. The stimuli were designed so that the duration of an occlusion period, determined by the number of time steps in which the barrier sat in front of the ball, was not confounded with the display's appearance. That is, during occlusion periods of all durations, the input information was identical, with only a stationary barrier visible. With all else equal in this way, each additional time step in the occlusion period places additional stress on the network's ability to maintain internal representations about the preocclusion display and thus to form predictions about the ball's reappearance.

The network's ability to form expectations is subserved by its connection weights. These weights are adjusted in the course of learning to make predictions from observed events. As the encoding weights from the input layer to the internal representation layer and the recurrent connections within the internal representation layer are adjusted, the network becomes increasingly able to represent occluded objects, not part of the input itself, as patterns of activity on the internal representation units. These patterns of activity thus provide a signal for an occluded object's continued existence. Connections from the internal representation layer can transform this signal into specific behaviors or predictions. In the case of the prediction weights from the internal representation layer back to the input layer, the signal that an occluded object continues to exist is transformed into a prediction about the world.[9] The network can thus gradually learn to predict an occluded object's reappearance. This is illustrated in Figure 9, where we graph the magnitude of a network's sensitivity to an occluded object reappearing from behind a barrier as a function of training experience and length of occlusion period. The sensitivity depends on the extent to which the network distinguishes between events with and without balls and is defined as the network's predicted activation for the ball unit at the time step when the ball should reappear (when there is in fact a ball behind the occluder), minus the network's predicted activation for the ball unit at the same time step, when a ball should not reappear (when in fact no ball was present before the occluder moved in).

The length of the occlusion period varies from three to seven time steps. At first, the network's connections do not support meaningful predictions and so the network is not sensitive to an occluded object's reappearance, but with experience seeing the occluded balls reappear, the network comes to have these expectations. If the object does not reappear, there is a discrepancy between the prediction and the observed event. Such a discrepancy, we suggest, is the signal that causes infants to increase looking when surprising events occur. At any point in

---

[8] The input world of the network was parsed into separate objects. We made this simplification for purposes of simulation and are not assuming that this information is innately available to the infant.

[9] We do not mean to claim that representational and predictive systems can necessarily be so cleanly divided in the real system. The goal in these simplified systems is simply to concretize aspects of development that are typically not considered in explanations of infant behavior.

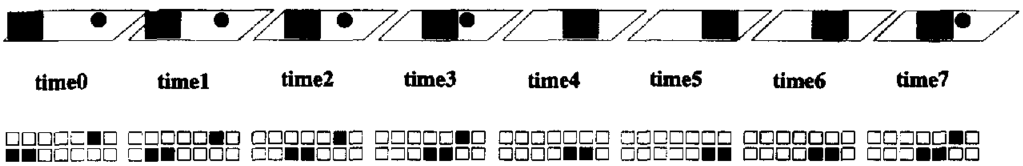

Figure 8. A series of inputs to the network as a barrier moves in front of a ball and then back to its original location. The top row shows a schematic drawing of an event in the network's visual field; the bottom row indicates the corresponding pattern of activation presented to the network's input units, with each square representing one unit. Learning in the network is driven by discrepancies between the predictions that the network makes at each time step and the input it receives at the next time step. The correct prediction at one time step corresponds to the input that arrives at the next time step.

development, the network exhibits a greater sensitivity to the continued existence of occluded objects when tested with shorter occlusion periods. These curves are reminiscent of the developmental data presented by Diamond (1985) for the *A$\bar{B}$* task. Diamond showed that infants become increasingly able to withstand longer delay periods before producing the *A$\bar{B}$* error. The delay needed to produce the error increased at a rate of approximately 2 s per month, with infants younger than 7.5 months producing the error with delays of less than 2 s while 12-month-olds could withstand delays of over 10 s. In the simulation data, one can see a similar pattern of increasing ability to withstand delays.

It is evident from the simulation that the overall course of development is quite extended, but that even at a relatively early point there is some degree of sensitivity. Thus, the simulation is consistent with the idea that differential looking times might show sensitivity to hidden objects at a relatively early age, even though the full development of the representation may span a much longer period. These results were replicated in 20 different runs of this simulation, each beginning with networks with different random starting weights. Although there was some variability in time to show sensitivity to occluded objects,[10] the general pattern of learning was the same across simulations.

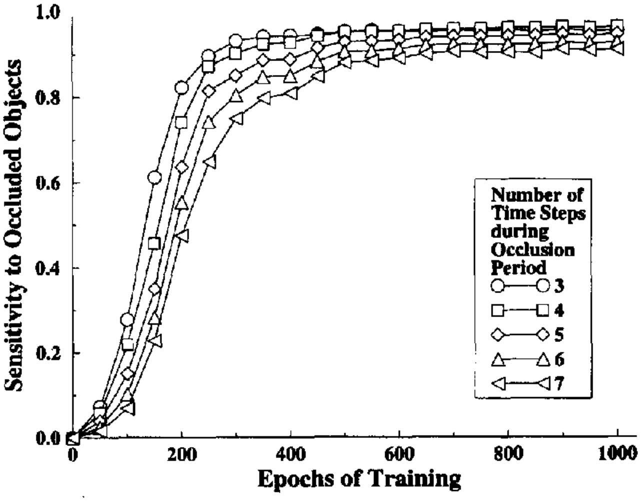

Figure 9. Gradual learning curves indicating the network's increasing sensitivity to the reappearance of occluded objects: The network's sensitivity is computed as the difference between the network's expectations for events with and without occluded objects.

Analysis of Internal Representations

How does the network solve the task of making correct predictions? As we shall see, it does so by learning to represent objects that are no longer visible as patterns of activation in its hidden units. Such patterns of activation may correspond to the active representations that infants form for hidden objects, and the processes that the network uses to learn such representations may correspond to the processes that give rise to infants' tendency to maintain representations of occluded objects.

Method

To demonstrate that the network learns to represent the continued existence of occluded objects, we can record patterns of activity across the network's internal representation units during various occlusion and nonocclusion events. Because these patterns of activity must represent other things in addition to the occluded or nonoccluded ball (e.g., the location of the barrier if one is present, which direction it is moving in, etc.), we look at the differences between particular pairs of patterns to isolate the representation of the ball. For example, to isolate the network representation of the ball during events involving a barrier, we record the pattern of activity across the network's internal representation units at a particular time step in a particular "ball–barrier" event and subtract from it the pattern of activity from the corresponding time step in the corresponding "barrier-only" event (Figure 10). Similarly, to isolate the network representation of the ball during events without a barrier, we record the pattern of activity for the "ball-only" stimulus and subtract from it the pattern of activity recorded for the analogous "nothing" stimulus.

Results and Discussion

When we compare these various representations for the object, we see that the network gradually learns to represent the object in similar ways when it is visible (whether alone, or prior to or following occlusion) and when it is occluded. Consider the network after 100 epochs of training, when it demonstrates

---

[10] The 20 simulation runs also varied somewhat in the smoothness and slope of their learning curves and in the size of sensitivity differences between different delay periods.

limited sensitivity to occluded objects (as shown previously in Figure 9). The representations for the ball are shown in Figure 11. Each column in the figure corresponds to a particular internal representation unit, and each row corresponds to a particular stimulus subtraction. For example, the top row corresponds to the ball–barrier-minus-barrier-only subtraction for time0, when the barrier is in the left-most position in the display. At this time step, the ball is not yet occluded, so the result of the stimulus subtraction to isolate the network's representation of the ball is labeled "ball visible preocclusion." The ball-only-minus-nothing subtractions in the lower half of the figure show the network's representation for the ball when it is visible alone. The shading of the boxes represents the sign of the stimulus differences, with white for positive and black for negative values, and the size of the boxes represents absolute magnitude of stimulus differences. Large boxes thus indicate units that code for the ball's presence, because these units are activated differentially for stimuli that are similar in all regards except for the presence or absence of the ball.

Units 1, 8, 10, 11, and 15 code for the ball most strongly. That is, averaged across all visible conditions, these units have the greatest absolute magnitude in the difference between their activations for stimuli with and without balls. Of course, there is no structural resemblance between the existence of the ball in the network's world and the representation of the ball across these internal representation units. Instead, the relation between the ball in various states and the internal representation of the ball is a second-order isomorphism (Shepard, 1970). That is, rather than the ball mapping to an internal representation that resembles it, the ball in various states maps to internal representations that resemble one another; the network shows similar patterns of activation across the internal representation units for the ball when it is visible alone and when it is visible prior to and following occlusion. At this point in development, the second-order isomorphism between the ball and the internal representation of it does not clearly hold when the ball is occluded. The signal for the ball seems to be only weakly maintained during this period.

The network's developing ability to represent the ball during the occlusion period can be traced through the activity of the internal representation units. In Figure 12, these representations are shown for the network after 200 epochs of training and after 1,000 epochs of training. The network's representation of the occluded ball becomes more similar to its representation of the visible ball. This pattern is particularly clear in four of the five units that code for the ball most strongly (1, 10, 11, and 15). At a given level of training, the longer the period of occlusion, the weaker the representation becomes. These observations are confirmed by quantitative comparisons of the ball representations during visible and occluded conditions. To make these comparisons, the network's isolated representation of the ball at each time step (as shown in Figures 11 and 12) was treated as a 15-element vector. Each vector for each of the seven time steps during which the ball was occluded was compared to the vector for the corresponding time step during the ball visible alone condition. The similarity between the visible and occluded ball representations was computed as the dot product between the two vectors (Figure 13). The dot products decrease with each occlusion time step, indicating that the representations of the occluded ball weaken with longer delays. The dot products increase with experience, indicating that the network becomes increasingly able to maintain the representation of the ball during the occlusion period.

Generalization to Novel Objects

In the previous set of simulations, a network's sensitivity to hidden objects was measured through its responses to the presence or absence of a single ball. In infants however, the notion of a concept of object permanence is not limited to a single object; infants show sensitivity to many different kinds of occluded objects. If this sensitivity is to be learned, it surely must depend upon generalization. To test the ability of networks to generalize to novel objects, we expanded the network stimulus environment to 35 objects. Networks were trained on a subset of these objects, and were then tested on their ability to maintain representations of specific novel and familiar objects across delays.

Method

The stimulus environment was expanded through the use of distributed representations for the objects. Rather than a single unit indicating the presence of a ball, each object was uniquely specified by the activation of three out of seven possible units, for a total of 35 objects. Example occlusion sequences with 2 of these objects are shown in Figure 14.

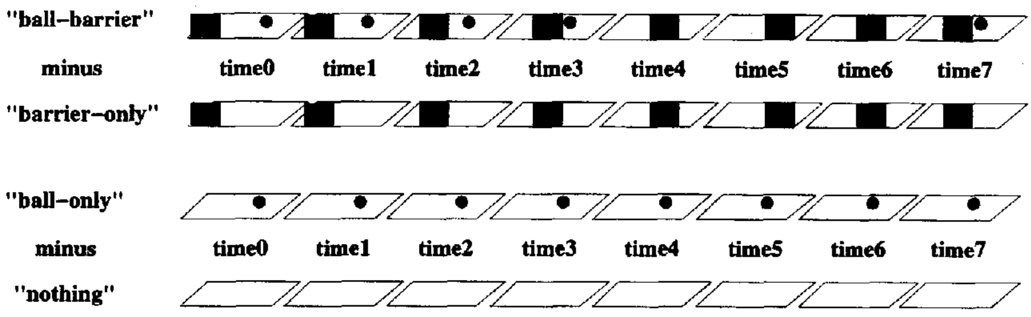

Figure 10. "Ball–barrier," "barrier-only," "ball-only," and "nothing" events: To isolate the network's representation of the ball, the pattern of activity across the network's internal representation units is recorded for these events. The "barrier-only" pattern is subtracted from the "ball–barrier" pattern, and the "nothing" pattern is subtracted from the "ball-only" pattern.

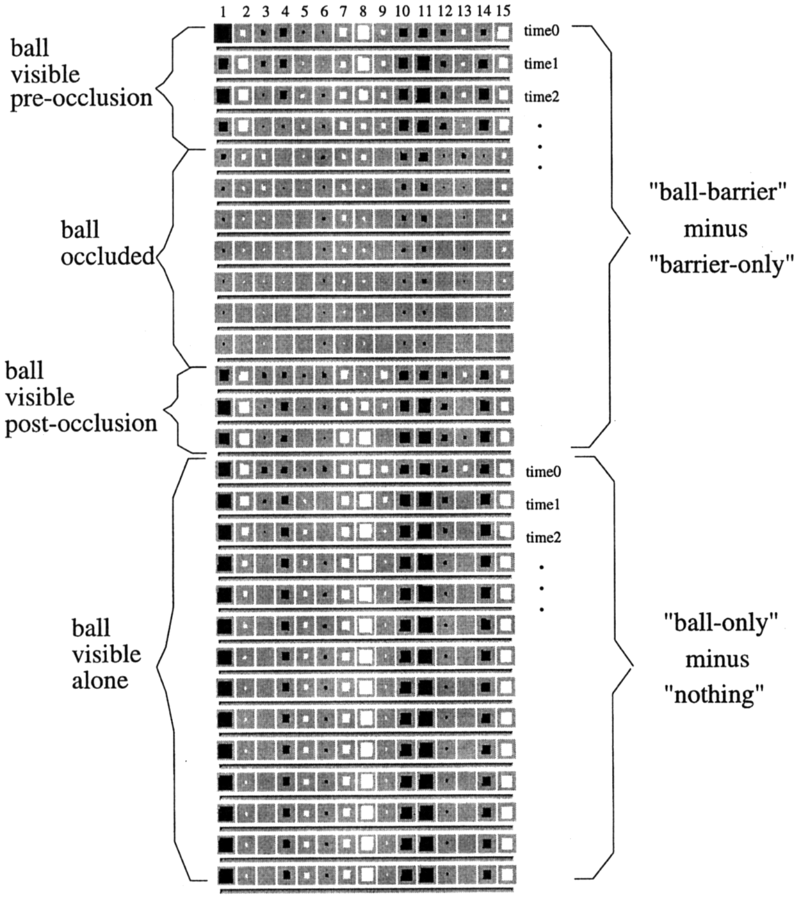

Figure 11. Internal representation analysis isolating the network's representation of the ball after 100 epochs of training. The 15 units in the internal representation layer are represented across each row in the table, with each row corresponding to a particular stimulus subtraction. The resulting values of the subtractions for each unit are represented by the shading and size of the boxes: white for positive values and black for negative, with size indicating absolute magnitude. Units 1, 8, 10, 11, and 15 code for the ball most strongly. The signal from these units seems to be maintained only weakly during the occlusion period.

The length of occlusion periods was again determined by the number of time steps in which the barrier blocked the view of an object, and ranged from three to seven time steps.

As in the previous simulations, the network's sensitivity to occluded objects depended on the extent to which the network distinguished between events with and without the objects. In the previous simulations, sensitivity was computed through the network's prediction just prior to the ball's reappearance, as the difference between the activations of the single ball unit during the ball-present and ball-absent events. In the current simulations, the network must do more than maintain that an object is present—the network must maintain which particular object is present as well. The extent to which the network predicted a particular object over other possible objects was measured by recording from all seven of the object units in the network's prediction. Object-specific predictions were computed as the average activity of the three units that did code for the object minus the average activity of the four units that did not code for the object. The network's sensitivity to the continued existence of the occluded object was then computed as the difference between the network's object-specific predictions for the object during the object-present and object-absent events.

In order to test the network's ability to generalize to novel objects, a random 15% of the 35 objects in the stimulus environment (5 objects) were not presented during training. The network was tested at various points in its learning with both familiar and novel objects. During the testing, learning was turned off, so that later tests were not contaminated by learning during earlier testing.

Results and Discussion

Ten different runs of this simulation, each beginning with networks with different random starting weights, exhibited a strong ability to generalize to novel objects (Figure 15). That is, the network's ability to maintain a representation of an occluded object in order to predict its reappearance was only minimally affected by whether the object was familiar (having been presented numerous times on each epoch of learning) or completely novel.

It is important to note that the network's capacity to generalize to novel objects is critically dependent on the overlap in the distributed representations of the novel and familiar objects.

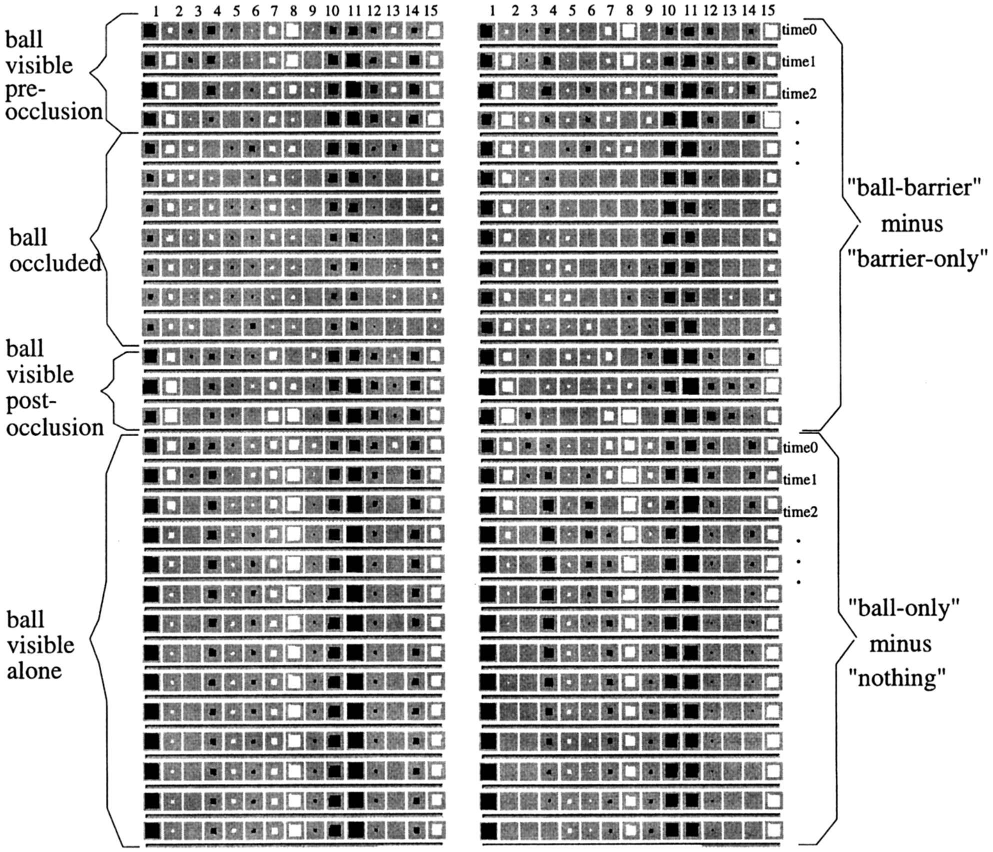

Figure 12. Internal representation analysis isolating the network's representation of the ball after 200 epochs and after 1,000 epochs of training. The network becomes increasingly able to maintain the signal for the ball during the occlusion period.

Our assumption is that such overlap exists in the representations of objects in the world, such that even never-before-seen objects share higher level visual object representations with familiar objects to some degree. One might test the effects of degree of novelty of various objects, to the extent that it can be measured, on the ability to maintain representations of these objects when occluded; perhaps a developing system has more difficulty maintaining representations of highly unusual stimuli. The point of the current simulations is simply to demonstrate that the network has the capacity to generalize to novel objects that overlap with familiar objects but are not identical to anything that the network has seen before.

We should also stress that although the networks generalized to novel objects, their limited experience may limit their ability to generalize to other kinds of novelty. For example, whenever the barrier was presented to the networks, it was always moving. One might expect that the networks would thus come to consistently predict that the barrier would move whenever it was presented, even given a novel event in which the barrier in fact remained still (suggested in simulations described in Marcus, 1997). We expect that for both networks and infants, a richer experience base would allow them to avoid or correct such erroneous predictions. Moreover, we would predict that infants would face the same difficulties as the networks in responding to certain kinds of novelty, if the infants were raised—like the networks—with an extremely limited range of visual experiences. This prediction is supported by various selective rearing experiments, demonstrating that animals with limited visual experience have difficulty processing novel stimuli. For example, animals exposed to horizontal lines alone seem blind—both behaviorally and physiologically—to vertical lines, and vice versa (Blakemore & Cooper, 1970; Hirsch & Spinelli, 1970), and animals exposed to lines moving in one direction show greatest sensitivity to lines moving in that direction (Tretter, Cynader, & Singer, 1975). Similarly, networks exposed to a highly restricted set of visual stimuli inevitably face some difficulty in responding to certain kinds of novelty. The simulations nonetheless allow us to explore certain critical issues regarding the potential nature of the representations and mechanisms underlying the development of sensitivity to hidden objects.

Accounting for Task-Dependency in Reaching and Looking-Time Measures of Object Permanence

We have illustrated how a network that learns from experience can gradually develop the ability to maintain representations of occluded objects on the basis of experiences that conform to the principle of object permanence. These simulations

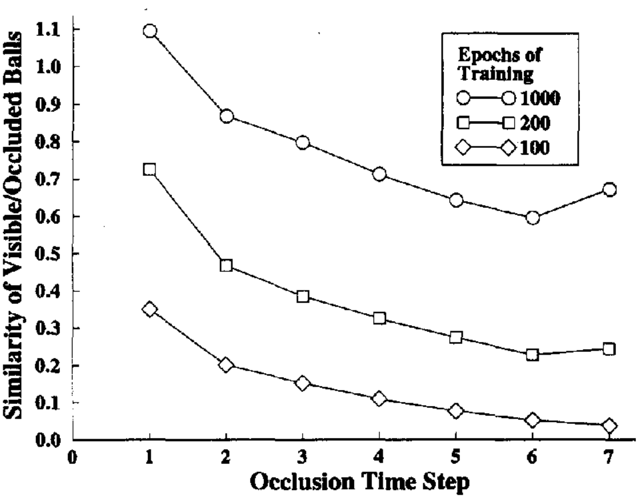

Figure 13. Dot product similarity comparisons of representations for visible and occluded balls: The network's internal representation of a ball on a given time step was treated as a 15-element vector. Similarity was computed as the dot product between the vector for an occluded ball and the vector for the visible ball on the corresponding time step. Similarity between representations for occluded and visible balls decreases across delays and increases with experience.

suggest how the gradual strengthening of the connections underlying the ability to maintain these representations could provide one way of accounting for the fact that infants' abilities to withstand delays between presentation and test increases with age. But how might we understand why infants succeed in looking-time tests of object permanence and fail on reaching and other action-based measures of the sort used in the present experiments?

One possibility is that successful reaching requires a stronger internal representation than is required to exhibit surprise through longer looking times. We have suggested that this difference in required strength of representation might be due to a greater complexity or effort required for reaching or lower frequency of reaching behaviors. Another possible reason for the difference in required strength might be that the reaching system develops later than the systems that underlie looking at interesting events. Infants fixate interesting stimuli from birth (Banks & Salapatek, 1983), whereas they first reach consistently only around 3–4 months of age (Thelen et al., 1993; Hofsten, 1984). For these reasons, making predictions might be possible with relatively weak representations, whereas overt reaching responses might depend on stronger internal representations.

The idea that reaching measures require stronger internal representations than looking measures is not the only way to account for the looking-time–reaching task dependency within an adaptive processes approach. Later we consider alternatives. We have chosen to simulate the differential strength possibility because it may at first seem counterintuitive, especially in view of particular aspects of the data. It appears that, at one and the same time (i.e., 7 months), infants are well past the age where they first show sensitivity to occluded objects in looking-time tasks, and they are able to retrieve visible objects, but they are not yet able to retrieve occluded objects. Evidence for the first point is provided by all of the evidence reviewed previously of sensitivity to hidden objects in looking-time tasks. Evidence for the second and third point comes from all three of our studies, in which infants retrieve visible objects (and make fewer retrieval responses when they can see that there is no object to retrieve) but show no sensitivity to occluded objects in their retrieval responses. These findings seem to create something of a contradiction. If infants' representations are sufficient to exhibit sensitivity to occluded objects in looking time tasks and their response capabilities are sufficient for retrieving visible objects, then why should they not be able to reach for occluded objects? The contradiction arises if we think of the capabilities of the

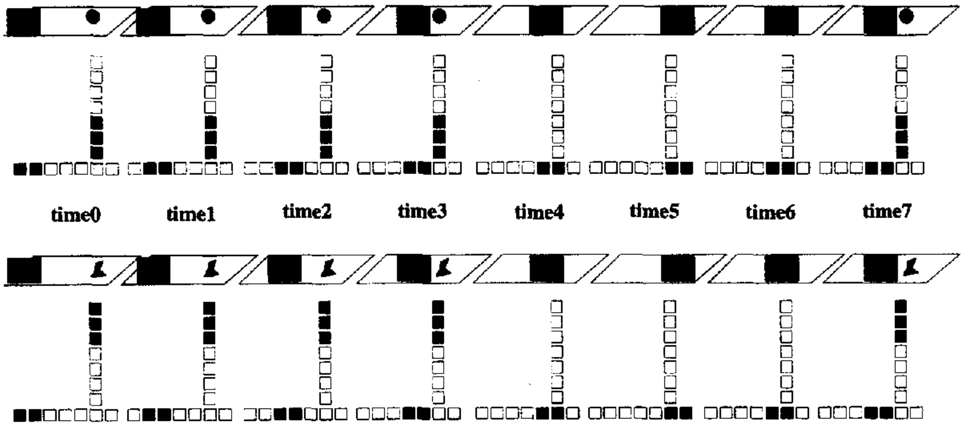

Figure 14. A series of inputs to the network, as a barrier moves in front of two objects and then back to its original location. The two rows above the time markers correspond to Object 1; the two rows below the time markers correspond to Object 35. The top row in each pair shows a schematic drawing of an event in the network's visual field; the bottom row indicates the corresponding pattern of activation presented to the network's input units.

representation and reaching systems in all or nothing terms. But, as our simulations illustrate, if we think of these capabilities in graded terms, as functions that are gradually acquired as a result of experience in an adaptive processing system, the contradiction disappears.

We now present a simulation that uses a single internal representation to drive two different outputs. One of these outputs is assumed to correspond to the prediction output, and the other to an overt reaching response. The simulations demonstrate how (a) a system can show sensitivity to occluded objects through prediction and to visible objects but not occluded objects through reaching; (b) the system could become able to show sensitivity to occluded objects through reaching, on the basis of developments in the representational system alone; and (c) developments in the representational and output systems can provide unique contributions to improvements in the system's performance.

In some ways, it might be preferable if our simulation incorporated a system that carried out a plausible analog of overt reaching, together with the system already described that makes predictions. We have not followed this course, however, for two reasons. First, visually guided reaching is a highly complex task requiring sensory-motor integration, and an adequate computational understanding of the reaching process is only just beginning to emerge (Rosenbaum, Loukopoulos, Meulenbroek, Vaughan, & Engelbrecht, 1995). Second, any attempt to implement a reaching system would require several specific assumptions, but the argument we wish to make is far more general and would apply whenever output systems differ in their strength for any reason. We have chosen, therefore, to capture our point in a system that has the crucial general property, without actually simulating reaching.

The model we use in the present simulations is illustrated in Figure 16. It has a single representation system, as before, but two different outputs. One of these is identical to the prediction output as before. The other was set up to be a system in which the task structure is basically the same as the task of the prediction system, but two manipulations were employed to ensure that its development was delayed relative to the prediction output. The first manipulation was to reduce the rate of learning within this second system to one tenth the rate of learning in the looking system. This manipulation was meant to serve as a proxy for the mechanisms responsible for the later mastery of reaching behaviors. The second difference between the reaching and looking systems was in their onset of training; the reaching system began developing only after the network had partially learned to form predictions. We call this second system the delayed output system, noting that the delay is due both to differences in the onset of training and the rate of learning. These manipulations allow the delayed system to capture the assumption that reaching behaviors are delayed relative to predictions, without specifying the detailed mechanistic basis for this. The model thus incorporates the two facets of the adaptive processing approach—developing representations and developing task-specific output behaviors—and allows us to explore how different output systems might be differentially sensitive to the level of strength of the internal representation. The results reported here were typical of 20 different runs of this simulation, each beginning with networks with different random starting weights.[11]

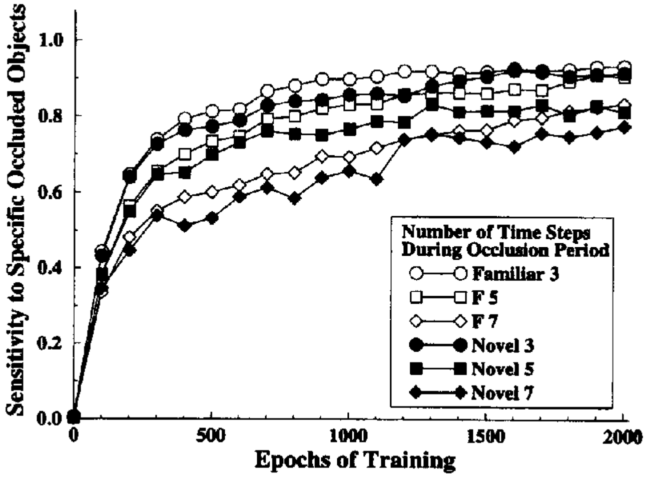

Figure 15. Learning curves averaged across 10 different simulation runs, indicating the networks' increasing sensitivity to specific objects and ability to generalize from familiar to novel objects. The learning curves for the familiar and novel objects are averaged across the 30 familiar and 5 novel objects, respectively.

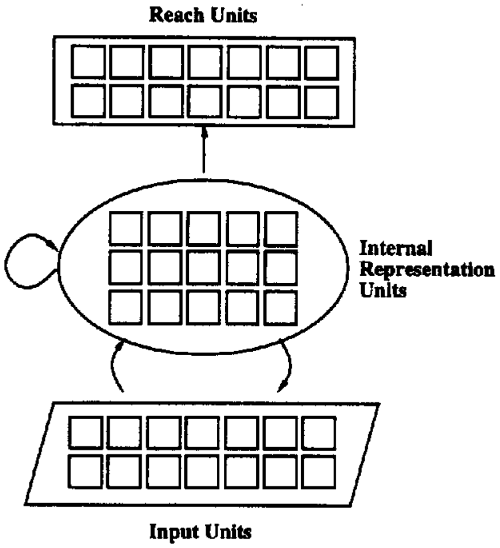

Figure 16. Recurrent network for learning to anticipate the future position of objects. The delayed output system is identical to the standard prediction output system, but training of the weights in the delayed output system begins after the network has begun to make predictions, and the learning rate parameter for these connections is smaller than for the prediction connections.

The two output systems were each trained with the same stimuli used in the original set of simulations, in which objects

---

[11] Again, the 20 simulation runs showed variability in time to show sensitivity to occluded objects, smoothness and slope of learning curves, and size of sensitivity differences between different delay periods.

moved across a simple, simulated visual display. The goal was again to predict the future positions of the objects. The network's sensitivity to occluded objects was measured separately for the normal and delayed systems and was computed in the same way as in the original model, based on the network's ability to form different predictions for events involving barriers with and without balls.

Like the original model, the network gradually came to acquire internal representations that allowed it to track the continued existence of the object behind the barrier. Analyses of the units in the internal representation layer revealed the same developmental pattern shown by the original model: The network gradually learned to represent the occluded object in similar ways to the visible object. The magnitude of the network's sensitivity when the occluded object failed to reappear from behind the barrier is shown for the normal and delayed output systems in in Figure 17. The delayed output system shows the expected delayed course of development with occluded objects because of its reduced learning rate and delayed onset of training.

Also shown in Figure 17 is the developmental course of sensitivity of the delayed output system to the presence of the ball when it is not occluded by a barrier. It is seen that at the same time that the delayed output system is showing virtually no sensitivity to occluded objects (i.e., no discrimination between ball presence and absence behind the occluder), the normal output system is showing sensitivity to occluded objects and the delayed output system is showing sensitivity to visible objects. This corresponds to the situation in which infants show sensitivity to occluded objects in looking-time measures and show sensitivity to visible objects in their reaching behavior, but show little or no sensitivity to occluded objects in reaching (i.e., no discrimination in their reaching given an occluder with and without an object behind it).

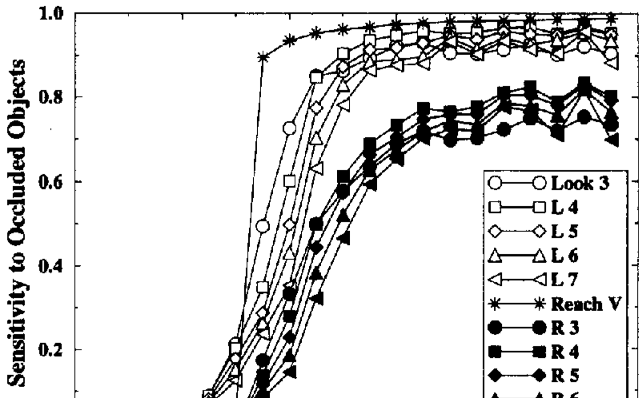

Figure 17. Learning curves indicating the network's increasing sensitivity to the reappearance of occluded objects. The reaching system has been slowed with a reduced learning rate and delayed onset of training, so that the looking system shows greater sensitivity to occluded objects. However, the reaching system can respond appropriately to visible objects early in learning, suggesting that the strengthening of internal representations may play a role in improvements in reaching behavior.

What allows the network to progress from initially failing reaching and looking measures of object permanence, to then passing only the looking measure, to eventually passing both looking and reaching measures? Within a framework in which having some capability—such as the ability to maintain representations of occluded objects—is either all or none, the retrieval behavior of 9-month-old infants should not depend on changes in the ability to represent occluded objects because infants presumably had this ability at 3.5 months. In contrast, within the context of our adaptive process framework in which the systems underlying performance adapt gradually over time, changes in the ability to maintain representations can play a key role in changes in performance over development.

To explore the relative contributions of the representations and output systems to this course of development, one would like to conduct an experiment along the following lines: Take a group of 3.5-month-old infants and freeze the shared representational system underlying looking and reaching. Hold it constant while allowing all other processing systems to develop normally. In a second group of 3.5-month-old infants, freeze the output systems relevant to acting on representations of hidden objects, while allowing the representational system to develop. Periodically bring the two groups of infants back into the lab at later ages and test their ability to retrieve hidden objects. Compare the performance of the two groups to each other and to that of a control group with normally developing systems. Although this cannot be done with real infants, it can be done with models of infants' performance. We can ask to what extent the network's performance can improve on the basis of changes in the connection weights in its representational system. Likewise, we can ask to what extent the network's performance can improve on the basis of changes in the connection weights in its output systems. The relative contributions of the representational and output systems are compared by freezing these systems in networks early in development and comparing subsequent performance to that of intact networks.

Method

In each run of the simulation described above, the network's performance was tested after every 10 epochs of training. For each run, a time point was chosen at which the network showed sensitivity to occluded objects through the normal output system and sensitivity to visible objects through the delayed output system and had a nascent but still underdeveloped sensitivity to occluded objects in the latter system. Specifically, the first time point where the output showed a sensitivity of 0.2 in the delayed output system for the shortest occlusion duration was chosen. This point is referred to as the *early competence point.* At this test point, the weights in the representation system or the weights in the delayed output system were prevented from developing further. The weights in the representation system consist of the weights from the input layer to the representation layer and the recurrent connections at the representation layer. These connection weights determined the pattern of activity formed on the representation layer and thus the extent to which the network was able to represent the existence of occluded objects. The weights in the delayed output system were simply the connections from the internal representation layer to the delayed output layer. These connections are responsible for transforming internal representations into outputs of the delayed output system. Training continued in the networks with these frozen weights, and the networks were tested after every 50 epochs. These tests allow us to see the effect of restricting developmental changes to one part of the system while holding another part constant.

Results and Discussion

The results of applying these procedures to the network from Figure 17 are shown in Figure 18. The pattern of results reported here were typical of the 20 different runs of the simulation, though there was some variability in the level of performance ultimately attained by the frozen networks. The left-most graph shows the results for the normal network at all time points. The early competence point for this network falls at Epoch 360. The middle graph shows the results of testing the network with output weights frozen at the early competence point. The right-most graph shows results from the network with representation weights frozen at the early competence point.

First consider the middle graph with output weights frozen. In this simulation, only the representational system is developing. The results of this analysis clearly show that the network can improve its reaching behavior on the basis of developments in the representational system alone; the network's sensitivity to occluded objects is comparable to that of the fully intact network. Thus, this simulation demonstrates that a system could progress from showing sensitivity to hidden objects through looking to showing sensitivity through looking and reaching on the basis of representational developments alone.

Of course, we do not actually believe that infants' output systems are frozen from 3.5 months. The purpose of these freezing analyses is simply to illustrate the potential importance of particular systems. We have held the output systems constant to show that reaching performance can improve through representational changes alone. But this is only part of the picture; the right-most graph shows the contribution from developments in the output system. With frozen representation weights, the curves for the five delay periods tend to diverge from the early competence point. This contrasts with the behavior of the network in the other two cases, in which the curves converge as the representations strengthen. The freezing of the representational system prevents the network from getting any better at retaining information about occluded objects. As the delayed output system develops, the network becomes better able to use weak object permanence representations. However, because the representation weights come from a point in development at which the weights are still very sensitive to delay, the output is likewise sensitive to delay even as the output weights continue to develop.

These results from the network's developing output system make contact with the means–ends training in our experiments. Infants were trained on the means–ends aspect of the tasks only, not on representing hidden objects. This training allowed infants to discriminate at test between toy presence and absence under visible but not occluded conditions. Our freezing analyses of the networks also show that training on the output system alone can at best lead to limited improvements in the system's ability to demonstrate sensitivity to occluded objects. Although it is possible that more training on means–ends behaviors might lead infants to show more sensitivity through retrieval measures, our simulations and experiments are consistent in suggesting that any such improvements might be limited and might require extensive training.

In summary, development of the representational system improves network performance by strengthening the representations of occluded objects so that longer delay periods can be withstood and the relative weakness of the reaching system can be overcome. This representational development may be sufficient to allow a system that could initially reach only for visible objects to then reach for occluded objects as well. In this way, representational developments may be critical to infants' increasing abilities to demonstrate sensitivity to hidden objects across a range of tasks. Holding the output system frozen allows us to see this potential representational contribution. Similarly, holding the representational system constant allows us to see that the development of the output system can improve network performance by expanding what the network is able to do with its internal representations. These freezing analyses thus demonstrate how learning in the representational and output systems can provide unique contributions to the development of reaching behavior. Although explanations of infants' task-dependent behavior have typically focused on deficits in output or ancillary systems alone, our simulations demonstrate the important poten-

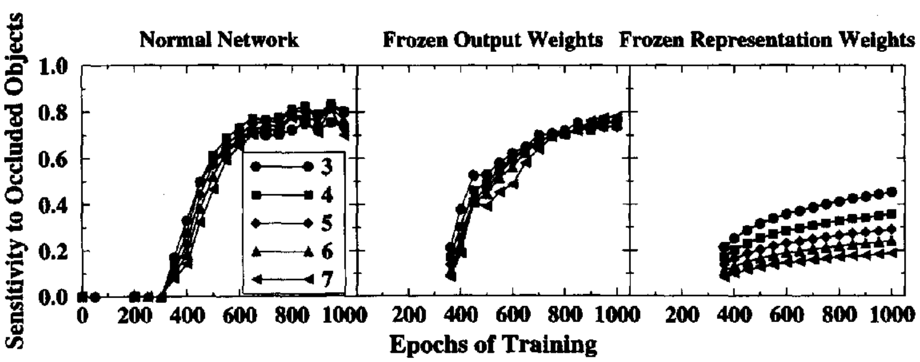

Figure 18. Learning curves for the delayed output system of normal and frozen networks, indicating the distinct contributions of the output and representational systems. The legend indicates the number of time steps during the occlusion period.

tial contribution of representational developments to an increasing ability to demonstrate sensitivity across a range of tasks.

Discussion of Simulations

In these simulations, we have attempted to make concrete one way in which the results of our experiments and the results of other experiments showing early competence can be brought together within an adaptive process account of object permanence. With experience, networks gradually acquired the ability to maintain representations of objects during periods of occlusion. As these representations were strengthened in this way, the networks became increasingly able to demonstrate sensitivity to occluded objects on both looking and reaching measures. We have also considered how these graded representations can lead to success on one measure but not another at intermediate points in their development. The key assumption is that different output systems may differ in their state of development, so that they require different strengths of representations to govern behavior. Even if both output systems function adequately with visible objects, differences between the strengths of representations they require may still account for differences in behavior with occluded objects.

It should be noted that other interpretations of the task dependency in looking and reaching measures of object permanence are possible within our adaptive processing systems framework. One possibility is that infants' looking and reaching depend on different representational systems. According to this hypothesis, infants may succeed in looking-time versions of object permanence tasks but not reaching versions because the systems use different representations of objects, and the ability to maintain representations of occluded objects has developed in the prediction system, but not in the reaching system. This possibility is certainly consistent with the adaptive processing systems framework. In addition, neuropsychological, electrophysiological, and behavioral data provide evidence for two separate anatomical pathways—the dorsal and ventral visual systems—for the processing of visual information (Goodale & Milner, 1992; Maunsell & Newsome, 1987; Ungerleider & Mishkin, 1982). Although the exact characterization of each of these two systems remains a matter of debate, one possibility is that the dorsal visual system is specialized for visual information processing relevant to acting on objects (including reaching for them), whereas the ventral visual system is specialized for perceiving and identifying the objects involved (Goodale & Milner, 1992). Given this, it is conceivable that reaching tasks might tap representations in the dorsal visual system, whereas looking-time measures might reflect predictions made within the ventral visual system. Bertenthal (1996) and Hofsten, Spelke, and colleagues (Hofsten, Spelke, Feng, & Vishton, 1994; Hofsten, Spelke, Vishton, & Feng, 1993; Spelke, 1994) have explored similar ideas of distinct representations governing infant behavior in looking time and reaching tasks.[12]

At present, we have no reason to favor the account we developed in our simulations on the basis of a single representation over this dual-representations account. We have chosen to focus our simulations around the idea that a single system of representation was involved in both tasks for two reasons. First, it did not seem necessary to illustrate the the dual-representation hypothesis with a simulation because it seems relatively straightforward to understand how distinct representation systems could lead to such task dependencies. Second, we think the lessons learned from our simulations based on a single representation system remain relevant to understanding performance in object permanence tasks, even if there are separate representational systems. It seems likely that gradual strengthening of representations within each pathway is still necessary to account for improvements in performance with age within the same behavioral task and to account for differences between tasks that tap common representations. For example, consider the task dependency between looking and reaching to a particular location in the $A\bar{B}$ task. Probability of looking at Location A or B in this task is quite a different measure than the duration of looking at a possible or impossible event as studied by Ahmed and Ruffman (1997) and Baillargeon and colleagues (Baillargeon & Graber, 1988; Baillargeon et al., 1989). It is possible that both looking and reaching to a location in which an object of interest might be hidden is a function of the dorsal visual system. If so, the fact that infants look to the correct location before they reach to that location might be attributed to just the sorts of differences between strength of representations required by output systems considered in the simulations above, even if looking-time effects in habituation experiments are based on different representational systems.

Another alternative explanation of task dependencies in object permanence tasks can be developed within the adaptive process approach on the basis of a distinction between active representations that themselves can govern behavior and latent adaptive changes that affect subsequent processing. Active representations are assumed to take the form of maintained neural firing that can serve as a signal to guide behavioral responses and are very much like the patterns of activity maintained in the networks used in our simulations. Latent adaptive changes, on the other hand, take the form of changes in the connections between neurons or in a neuron's responsiveness; these changes can govern behavior only by altering the course of the subsequent processing of stimuli. The hypothesis is that infants require active representations to reach for hidden objects because they receive no cues to the objects' continued presence in search tasks. In contrast, infants may be able to use latent adaptive changes to show longer looking times to impossible events with occluded objects (see Munakata, 1996, in press, for details). Once again,

---

[12] Mareschal, Plunkett, and Harris (1995; see also Elman et al., 1996) used a connectionist model to explore the possibility that looking measures show earlier sensitivity to hidden objects than reaching measures because the prediction system responsible for longer looking times receives information from only a single "where" pathway, whereas the reaching system depends on coordination of information from this pathway plus a second, object-identity pathway. Although this is an interesting possibility, there are data indicating that infants' looking times are sensitive to object features other than merely their location (Baillargeon, 1987b; Baillargeon, 1991; Baillargeon & DeVos, 1991; Baillargeon & Graber, 1987; Wilcox & Baillargeon, 1997). Thus, Mareschal, Plunkett, and Harris's (1995) coordination-of-representations account may be more relevant to task dependencies between visual tracking and reaching, rather than to the particular puzzle of why infants show longer looking times to impossible events with occluded objects before they reach for occluded objects.

we do not believe the data presently available is sufficient to help us choose between this type of account and the one that we have simulated in this article. As before, we have chosen to focus on graded representations accounts because they illustrate one important ingredient of the active-versus-latent representation account, namely the idea that the representation of hidden objects is experience dependent.

In summary, we have focused on an account of task dependencies that depends on the gradual strengthening of connections underlying the ability to maintain representations because we believe that this idea will play an important role in any full account of infant behavior in object permanence tasks and in our understanding of many aspects of development. We now turn to these more general considerations.

General Discussion

In this section, we evaluate the idea that infants have knowledge that can be characterized as a set of principles and ancillary systems that allow this knowledge to be expressed in particular task situations. This idea appears to be consistent with much developmental theorizing and therefore worthy of close examination. Although elaborations of early knowledge have been acknowledged within principle-based accounts, it is quite common for ancillary factors to carry the explanatory burden for infants' increasing ability to perform tasks thought to tap early principles.

Our simulations suggest that it may not be necessary or even useful to separate knowledge of principles and ancillary factors in this way. In fact, in our simulations, the basis for the early ability to perform in accordance with the principle of object permanence is the same as the basis for the increased robustness of this ability: It is adaptive change to connections among processing units involved in carrying out object permanence tasks. Some early adaptation of these connections provided the basis for early competence; later increases in the strengths of these connections provided the basis of increased ability to exhibit this competence in a broader range of tasks.

Although the attribution of principles may be a useful first step in calling attention to important aspects of cognitive functioning, we believe that we must move beyond this first step for two important reasons. First, infant behavior is highly task dependent; the age at which infants behave in accord with a given principle depends critically on how the principle is measured. Principles that are attributed to infants do little work in making sense of their actual behavior. Second, even principles that may describe behavior under limited conditions may not constitute mechanistic explanations of behavior. Infants may behave in accord with a principle of object permanence in specific tasks without reasoning with a generally accessible principle of object permanence, just as the planets move in accord with Kepler's laws without consulting them. Thus, principles may only serve as descriptions of limited aspects of infant behavior.

A common concern that has been voiced about our simulations and others is that they do not generalize adequately when faced with a broad range of test events. A proponent of principles could then argue, for example, that the only way to achieve sufficient generalization would be to have the principle of object permanence (Marcus, personal communication, April 1996). We have dealt with generalization from one set of training items to a different testing set. However, we have not discussed generalization to different types of events or even different configurations of conceptually quite similar events (e.g., occluder comes from right instead of left; position of occluded item varies across locations—both suggested in simulations in Marcus, 1997). Obviously, our networks do not generalize to most such cases because all of their experiences are tied to one very simple type of event, with one direction of movement and one hiding location. But infants' experiences are not so limited. We would suggest that a broader range of experiences could form the basis for infants' and networks' abilities to generalize, but a proponent of principles could argue that infants' experiences with objects are not sufficiently broad for connection-based changes to form the basis for their ability to generalize across object permanence tests. It is difficult to know which of these perspectives is correct without detailed study of infants' experiences, the extent to which infants actually do exhibit highly general sensitivity and the generalization abilities of networks. The question is an empirical one: whether networks trained with experiences that correspond to those of infants would exhibit as general sensitivity to object permanence as infants do. This is a valid issue for further exploration, but we do not believe that the available evidence is sufficient to cast any real doubt on our connectionist generalization-based approach. Our own conjecture is that network models will prove adequate once the considerations of available experience and actual generalization abilities are factored in but that the networks that prove adequate will be considerably more structured than the simple ones that we have considered here. For example, the encoding of visual experience in the brain is much more richly structured than in our networks, and visual information appears to be coded in terms of multiple frames of reference, including object-centered frames (Olson & Gettner, 1995). An object-centered representation serves to increase the similarity of experiences with objects in different locations in the world and different locations with respect to the retina or the body of the observer. Object-centered representations could thus increase generalization across variations in world- and body-centered locations of experiences.

We do not wish to suggest that there might ultimately be no role for principles in our accounts of the cognitive capabilities of young infants. On the contrary, we believe that there may be constraints—such as the initial patterns of connectivity in a network (Elman et. al., 1996; Rumelhart & McClelland, 1986)—that bias adaptive processing systems to behave in accord with principles. In this way, some form of principles like object permanence, and others that govern the physical and social environment, may be internalized by adaptive processing systems. In fact, it is not inconsistent with an adaptive processing systems approach to suggest that the tendency to behave in accordance with such principles might be in place at birth. However, we would not expect constraints on processing to guarantee behavior in accord with principles because multiple constraints presumably influence the internal representations and the ultimate behavior of the real system. Our essential point is that principles may be embodied in the form of certain processing constraints, but explaining behavior as caused by possession of principles is at best a potentially misleading shorthand. It is crucial to pay more explicit attention to the nature of the mechanisms underlying infant behavior if we are to make real progress in understanding early competence and subsequent cognitive development.

The second role of our simulations is to question the necessity of relegating changes in cognitive functions to ancillary factors. What appears to be an increase in an ancillary factor such as processing capacity can simply fall out of an adaptive processing system as a result of a gradual strengthening process. We have found it useful to capture this process in connectionist models where what is strengthened is the connections. In our model, the system is increasingly able to handle more demanding tasks, but this occurs simply because the connections subserving task-relevant representations become stronger. Similarly, what appears to be a global increase in memory or attention could actually be based on the strengthening of connections that support more robust representations and that in turn drive behavior more readily and are more easily maintained across delays.

Two aspects of this suggestion may deserve some comment. First, the idea of processing capacity is often characterized in extremely general terms, whereas we have focused very specifically on a single task domain. One of the appeals of the capacity idea is that it applies across the board, providing a basis for understanding changing performance across a wide range of different task domains. Our simulations have focused on the development of connections specific to representations of visible and occluded objects; development in other domains could depend upon similar mechanisms. Infants experience a much richer world and respond in more varied and complex ways than our simulations. As a result, the connections subserving performance on a wide range of tasks might be strengthened. Just as our simulations demonstrate generalization of responses from trained stimuli to novel ones, so real infants might generalize from their experiences to a range of novel situations. The resulting advances in behavior might suggest a global increase in capacity, memory, or attention, though there would be no single global factor that plays a causal role in development.

Second, it is necessary to make it clear that our proposals do not obviate a role for control processes in determining task performance. The networks simulated here appear to be autonomous processing systems, but we believe that they should be viewed as components of a larger system in which there may be other parts that control or modulate their function within limits determined by the strengths of the connections inside them. A concrete instantiation of this idea was presented in a model of the control of adaptive processing systems by Cohen, Dunbar, and McClelland (1990). The model was applied to performance in the Stroop task, in which a color word (e.g., the word *green*) is printed with a certain color of ink (e.g., red ink), and the participant must name the ink color or the word. In the model, there were separate pathways for reading words and for naming colors, as well as other modules that modulated processing in these two pathways. The stronger word-reading pathway was less dependent on modulation than the weaker color-naming pathway, accounting for the fact that word reading is overall faster and less error prone and the fact that conflicting words interfere with color naming but conflicting ink color does not interfere with word reading. The model is consistent as well with the fact that the ability of novel shapes to function just as color words do in the Stroop task can be acquired through practice. Attention controls processing in this model to the extent that it determines which aspect of the stimulus (word or ink color) dominates responding, but the strength of the pathways that process these different sources of information are also crucial. This model illustrates the sort of role we see ancillary processes such as attention playing in task performance but underscores the importance of the strengths of the connections within the pathways for accounting for details of task performance, robustness in the face of interference, and so forth.

In summary, our work leads us to suggest that although one might characterize certain behaviors in terms of principles, it may be misleading to treat such principles as entities in the mind that explain behavior. The adaptive processing systems framework instead allows us to consider the mechanisms underlying behaviors that may be described in terms of principles and to understand how developmental changes arise through these mechanisms. Further progress, of course, depends on further elaboration of this framework and the fuller development of explicit mechanistic accounts of behavior in object permanence tasks. For example, the particular mechanisms underlying the development of stronger representations could be further specified and related to known brain systems, and the reasons for certain tasks requiring stronger representations could be explored. The potential for strengthening representations through specific experiences, such as cross-modality presentation of stimuli, could be tested. One might also further evaluate the possibility that separate neural substrates are differentially relevant for maintaining representations for reaching and looking behaviors, in contrast to the common substrate theory explored here. All of these potential directions are clearly indicated by the adaptive process framework; these questions would not arise within a principle-based concept of object permanence. The shift from a principle-based treatment to an adaptive process account is but one step in achieving a deeper understanding of the central mystery of cognitive development—how change occurs.

References

Ahmed, A., & Ruffman, T. (1997). *Why do infants make A not B errors in a search task, yet show memory for the location of hidden objects in a non-search task?* Manuscript submitted for publication.

Bachevalier, J., & Mishkin, M. (1984). An early and a late developing system for learning and retention in infant monkeys. *Behavioral Neuroscience, 98,* 770–778.

Baillargeon, R. (1987a). Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology, 23,* 655–664.

Baillargeon, R. (1987b). Young infants' reasoning about the physical and spatial properties of a hidden object. *Cognitive Development, 2,* 179–200.

Baillargeon, R. (1991). Reasoning about the height and location of a hidden object in 4.5- and 6.5-month-old infants. *Cognition, 38,* 13–42.

Baillargeon, R. (1993). The object concept revisited: New directions in the investigation of infants' physical knowledge. In C. Granrud (Ed.), *Visual perception and cognition* (pp. 265–315). Hillsdale, NJ: Erlbaum.

Baillargeon, R. (1994). How do infants learn about the physical world? *Current Directions in Psychological Science, 3*(5), 133–140.

Baillargeon, R., & DeVos, J. (1991). Object permanence in young infants: Further evidence. *Child Development, 62,* 1227–1246.

Baillargeon, R., DeVos, J., & Graber, M. (1989). Location memory in 8-month-old infants in a non-search AB task: Further evidence. *Cognitive Development, 4,* 345–367.

Baillargeon, R., & Graber, M. (1987). Where is the rabbit? 5.5-month-old infants' representation of the height of a hidden object. *Cognitive Development, 2,* 375–392.

Baillargeon, R., & Graber, M. (1988). Evidence of location memory in 8-month-old infants in a non-search AB task. *Developmental Psychology, 24,* 502–511.

Baillargeon, R., Graber, M., DeVos, J., & Black, J. (1990). Why do young infants fail to search for hidden objects? *Cognition, 36,* 255–284.

Banks, M., & Salapatek, P. (1983). Infant visual perception. In P. H. Mussen (Series Ed.) & M. M. Haith, & J. J. Campos (Vol. Eds.), *Handbook of child psychology: Vol. 2. Infancy and developmental psychobiology* (4th ed., pp. 435–572). New York: Wiley.

Bertenthal, B. I. (1996). Origins and early development of perception, action, and representation. *Annual Review of Psychology, 47,* 431–459.

Blakemore, C., & Cooper, G. (1970). Development of the brain depends on the visual environment. *Nature, 228,* 477–478.

Bower, T. G. R., & Paterson, J. G. (1972). Stages in the development of the object concept. *Cognition, 1,* 47–55.

Bower, T. G. R., & Wishart, J. G. (1972). The effects of motor skill on object permanence. *Cognition, 1,* 165–172.

Braine, M. D. S. (1959). The ontogeny of certain logical operations: Piaget's formulation examined by nonverbal methods. *Psychological Monographs, 73*(475, Whole No. 475).

Bremner, J. G., & Knowles, L. S. (1984). Piagetian stage IV errors with an object that is directly accessible both visually and manually. *Perception, 13,* 307–314.

Brown, A. L. (1976). Semantic integration in children's reconstruction of narrative sequences. *Cognitive Psychology, 8,* 247–262.

Butterworth, G. (1974). *The development of the object concept in human infants.* Unpublished doctoral dissertation, University of Oxford, Oxford, England.

Butterworth, G. (1977). Object disappearance and error in Piaget's stage IV task. *Journal of Experimental Child Psychology, 23,* 391–401.

Case, R. (1985). *Intellectual development: A systematic reinterpretation.* New York: Academic Press.

Case, R. (1992). The role of the frontal lobes in the regulation of cognitive development. *Brain and Cognition, 20,* 51–73.

Clifton, R. K., Rochat, P., Litovsky, R. Y., & Perris, E. E. (1991). Object representation guides infants' reaching in the dark. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 319–323.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the stroop effect. *Psychological Review, 97,* 332–361.

Davidson, R. J., & Schwarz, G. E. (1977). Brain mechanisms subserving self-generated imagery: Electrophysiological specificity and patterning. *Psychophysiology, 14,* 598–601.

De Haan, E. H., Young, A. W., & Newcombe, F. (1987a). Face recognition without awareness. *Cognitive Neuropsychology, 4,* 385–415.

De Haan, E. H., Young, A. W., & Newcombe, F. (1987b). Faces interfere with name classification in a prosopagnosic patient. *Cortex, 23,* 309–316.

Diamond, A. (1981). Retrieval of an object from an open box: The development of visual–tactile control of reaching in the first year of life. *Society for Research in Child Development Abstracts, 3,* 78.

Diamond, A. (1983). Behavior changes between 6–12 months of age: What can they tell us about how the mind of the infant is changing? *Dissertation Abstracts International, 44*(01B), 337. (University Microfilms No. AAD8311882)

Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on AB̄. *Child Development, 56,* 868–883.

Diamond, A. (1991). Neuropsychological insights into the meaning of object concept development. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind* (pp. 67–110). Hillsdale, NJ: Erlbaum.

Elman, J., Bates, E., Karmiloff-Smith, A., Johnson, M., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development.* Cambridge, MA: MIT Press.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14,* 179–211.

Evans, W. F. (1973). The stage-IV error in Piaget's theory of object concept development: An investigation of the role of activity. *Dissertation Abstracts International, 34*(11B), 5651. (University Microfilms No. AAD74-11845)

Farah, M. J. (1988). Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review, 95,* 307–317.

Farah, M. J., Monheit, M. A., & Wallace, M. A. (1991). Unconscious perception of "extinguished" visual stimuli: Reassessing the evidence. *Neuropsychologia, 29,* 949–958.

Farah, M. J., O'Reilly, R. C., & Vecera, S. P. (1993). Dissociated overt and covert recognition as an emergent property of a lesioned neural network. *Psychological Review, 100,* 571–588.

Fischer, K. W., & Bidell, T. (1991). Constraining nativist inferences about cognitive capacities. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind* (pp. 199–236). Hillsdale, NJ: Erlbaum.

Flavell, J. H. (1971). Stage-related properties of cognitive development. *Cognitive Psychology, 2,* 421–453.

Flavell, J. H. (1984). Discussion. In R. J. Sternberg (Ed.), *Mechanisms of cognitive development.* New York: Freeman.

Flavell, J. H. (1985). *Cognitive development.* Englewood Cliffs, NJ: Prentice Hall.

Fuster, J. (1989). *The prefrontal cortex* (2nd ed.). New York: Raven Press.

Goldenberg, G., Podreka, I., Steiner, M., & Willmes, K. (1987). Patterns of regional cerebral blood flow related to memorizing of high and low imagery words: An emission computer tomography study. *Neuropsychologia, 25,* 473–486.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum & V. Mountcastle (Eds.), *Handbook of physiology: The nervous system V* (pp. 373–417). Bethesda, MD: American Physiological Society.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience, 15,* 20–25.

Gratch, G. (1972). A study of the relative dominance of vision and touch in six-month old infants. *Child Development, 43,* 615–623.

Gratch, G., & Landers, W. (1971). Stage IV of Piaget's theory of infants' object concepts. *Child Development, 42,* 359–372.

Greenough, W. T., Black, J. E., & Wallace, C. S. (1987). Experience and brain development. *Child Development, 58,* 539–559.

Halford, G. S. (1993). *Children's understanding: The development of mental models.* Hillsdale, NJ: Erlbaum.

Harris, P. L. (1974). Perseverative search at a visibly empty place by young infants. *Journal of Experimental Child Psychology, 18,* 535–42.

Hirsch, H., & Spinelli, D. (1970). Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science, 168,* 869–871.

Hofstadter, M. C., & Reznick, J. S. (1996). Response modality affects human infant delayed-response performance. *Child Development, 67,* 646–658.

Hofsten, C. von. (1984). Developmental changes in the organization of prereaching movements. *Developmental Psychology, 20,* 378–388.

Hofsten, C. von, Spelke, E., Feng, Q., & Vishton, P. (1994, June). *Predictive reaching for poorly occluded objects.* Paper presented at the Nineteenth Biennial International Conference for Infant Studies, Paris, France.

Hofsten, C. von, Spelke, E., Vishton, P., & Feng, Q. (1993, November). *Principles of predictive catching in 6-month-old infants.* Paper presented at the 34th meeting of the Psychonomic Society, Washington, DC.

Hood, B., & Willatts, P. (1986). Reaching in the dark to an object's remembered position: Evidence for object permanence in 5-month-old infants. *British Journal of Developmental Psychology, 4,* 57–65.

Johnson, M. H., & Morton, J. (1991). *Biology and cognitive development: The case of face recognition.* Oxford: Blackwell.

Jordan, M. I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. In C. Clifton (Ed.), *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* (pp. 531–546). Hillsdale, NJ: Erlbaum.

Just, M., & Carpenter, P. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99,* 122–149.

Karmiloff-Smith, A. (1991). Beyond modularity: Innate constraints and developmental change. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind* (pp. 171–197). Hillsdale, NJ: Erlbaum.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science.* Cambridge, MA: MIT Press.

Keil, F. C. (1981). Constraints on knowledge and cognitive development. *Psychological Review, 88,* 197–227.

Lecuyer, R., Abgueguen, I., & Lemarie, C. (1992). 9- and 5-month-olds do not make the AB error if not required to manipulate objects. In C. Rovee-Collier (Ed.), *Abstracts of papers presented at the eighth international conference on infant studies* (pp. 514). Norwood, NJ: Ablex.

Leslie, A. (1988). The necessity of illusion: Perception and thought in infancy. In L. Weiskrantz (Ed.), *Thought without language* (pp. 185–210). Oxford: Clarendon Press.

Mandler, J. M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review, 99,* 587–604.

Marcus, G. F. (1997). *The development and representation of object permanence: Lessons from connectionism.* Manuscript submitted for publication.

Mareschal, D., Plunkett, K., & Harris, P. (1995). Developing object permanence: A connectionist model. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 170–175). Mahwah, NJ: Erlbaum.

Matthews, A. (1992). Infants' performance on two versions of AB: Is recall memory a critical factor? In C. Rovee-Collier (Ed.), *Abstracts of papers presented at the eighth international conference on infant studies* (pp. 560). Norwood, NJ: Ablex.

Maunsell, J. H., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience, 10,* 363–401.

McClelland, J. L. (1989). Parallel distributed processing: Implications for cognition and development. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 8–45). Oxford University Press.

McClelland, J. L. (1992). *The interaction of nature and nurture in development: A parallel distributed processing perspective* (Tech. Rep. No. 92-6). Pittsburgh, PA: Carnegie Mellon University, Department of Psychology, Parallel Distributed Processing and Cognitive Neuroscience Group.

McClelland, J. L. (1994). The interaction of nature and nurture in development: A parallel distributed processing perspective. In P. Bertelson, P. Eelen, & G. D'Ydewalle (Eds.), *Current advances in psychological science: Ongoing research.* Hillsdale, NJ: Earlbaum.

Miller, D., Cohen, L., & Hill, K. (1970). A methodological investigation of Piaget's theory of object concept development in the sensory-motor period. *Journal of Experimental Child Psychology, 9,* 59–85.

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature, 335,* 817–820.

Miyashita, Y., & Chang, H.-S. (1988). Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature, 331,* 68–70.

Morton, J., & Johnson, M. H. (1991). CONSPEC and CONLERN: A two-process theory of infant face recognition. *Psychological Review, 98,* 164–181.

Munakata, Y. (1996). *Adaptive processes in cognitive change: A unified framework for understanding infants' successes and failures in object permanence tasks.* Unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh, Pennsylvania.

Munakata, Y. (in press). Task-dependency in infant behavior: Toward an understanding of the processes underlying cognitive development. In F. Lacerda, C. von Hofsten, & M. Heimann (Eds.), *Emerging cognitive abilities in early infancy.* Hillsdale, NJ: Erlbaum.

Newell, A., & Simon, H. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice Hall.

Olson, C. R., & Gettner, S. N. (1995). Object-centered direction selectivity in the macaque supplementary eye field. *Science, 269,* 985–988.

Pascual-Leone, J. (1970). A mathematical model for transition in Piaget's developmental stages. *Acta Psychologica, 32,* 301–345.

Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research, 47,* 329–342.

Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., & Jeeves, M. A. (1984). Neurons responsive to faces in the temporal cortex: Studies of functional organization, sensitivity to identity and relation to perception. *Human Neurobiology, 3,* 197–208.

Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., & Jeeves, M. A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society of London: Series B, 223,* 293–317.

Piaget, J. (1952). *The origins of intelligence in childhood.* New York: International Universities Press.

Piaget, J. (1954). *The construction of reality in the child.* New York: Basic Books.

Rosenbaum, D. A., Loukopoulos, L. D., Meulenbroek, R. G. J., Vaughan, J., & Engelbrecht, S. E. (1995). Planning reaches by evaluating stored postures. *Psychological Review, 102,* 28–67.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & PDP Research Group (Eds.), *Parallel distributed processing: Vol. 1. Foundations* (pp. 318–362). Cambridge, MA: MIT Press.

Rumelhart, D. E., & McClelland, J. L. (1986). PDP models and general issues in cognitive science. In D. E. Rumelhart, J. L. McClelland, & PDP Research Group (Eds.), *Parallel distributed processing: Vol. 1. Foundations* (pp. 110–146). Cambridge, MA: MIT Press.

Shepard, R. N. (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology, 1,* 1–17.

Siegler, R. (1989). Mechanisms of cognitive development. *Annual Review of Psychology, 40,* 353–379.

Siegler, R. (1993). Cheers and lamentations. In C. Granrud (Ed.), *Visual perception and cognition in infancy: Carnegie Mellon Symposia on cognition* (pp. 333–344). Hillsdale, NJ: Erlbaum.

Siegler, R. (1996). *Children's thinking: Beyond the immaculate transition.* New York: Oxford University Press.

Siegler, R., & Munakata, Y. (1993, Winter). Beyond the immaculate transition: Advances in the understanding of change. *SRCD Newsletter,* 3–13.

Smith, L. B., & Thelen, E. (Eds.). (1993). *A dynamic systems approach to development: Applications.* Cambridge, MA: MIT Press.

Sophian, C., & Yengo, L. (1985). Infants' search for visible objects: Implications for the interpretation of early search errors. *Journal of Experimental Child Psychology, 40,* 260–278.

Spelke, E. (1994). Initial knowledge: Six suggestions. *Cognition, 50,* 431–446.

Spelke, E., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review, 99,* 605–632.

Thelen, E., Corbetta, D., Kamm, K., Spencer, J. P., Schneider, K., & Zernicke, R. F. (1993). The transition to reaching: Mapping intention and intrinsic dynamics. *Child Development, 64,* 1058–1098.

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

Tretter, F., Cynader, M., & Singer, V. (1975). Modification of direction selectivity of neurons in the visual cortex of kittens. *Brain Research, 84,* 143–149.

Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549–587). Cambridge, MA: MIT Press.

Uzgiris, I. C., & Hunt, J. M. (1975). *Assessment in infancy: Ordinal scales of psychological development.* Urbana: University of Illinois Press.

Volpe, B. T., LeDoux, J. E., & Gazzaniga, M. S. (1979). Information processing of visual stimuli in an "extinguished" field. *Nature, 282,* 722–724.

Wellman, H. M., Cross, D., & Bartsch, K. (1986). Infant search and object permanence: A meta-analysis of the A-Not-B error. *Monographs of the Society for Research in Child Development, 51* (3, Serial No. 214).

Wilcox, T., & Baillargeon, R. (1997). *Object individuation in infancy: The use of featural information in reasoning about occlusion events.* Manuscript submitted for publication.

Willatts, P. (1990). Development of problem-solving strategies in infancy. In D. F. Bjorklund (Ed.), *Children's strategies: Contemporary views of cognitive development* (pp. 23–66). Hillsdale, NJ: Erlbaum.

Appendix

Simulation Details

In all of the reported single-ball simulations, the input and reaching layers contained 14 units each, corresponding to seven locations across the visual field for each of two depth planes. The internal representation layer contained 15 units. The connections shown in Figures 7 and 16 from the internal representation layer back to the input layer, for the networks' predictions about the next state of the input, were implemented using connections to a separate prediction layer with the same layout as the input layer. In the multiple-objects simulations, the input layer contained 14 units, corresponding to seven locations across the visual field for the first depth plane, and 7 potential units for a distributed representation of an object in a single location in the second depth plane. The internal representation layer contained 14 units.

In all of the simulations, units could take on continuous activation values between 0 and 1 and weights values could be any real values. Activation levels were updated according to the logistic activation function, and weights were adjusted after every step in each stimulus sequence according to the back-propagation learning algorithm (Rumelhart, Hinton, & Williams, 1986), using the sum squared error function. The momentum factor was set to .9. The recurrence at the internal representation layer was implemented using a context layer (Elman, 1990) that received a copy of the activity from the internal representation layer. The context layer was in turn fully connected back to the internal representation layer by means of modifiable weights. The activation of the context units was determined by the hysteresis (or self-weight) parameter, μ (Jordan, 1986), according to the following function:

$$act\_c_i(t) = \mu * act\_c_i(t-1) + (1-\mu) * act\_h_i(t), \quad \text{(A1)}$$

where act_c and act_h are the activities of the context and internal representation (hidden) units, respectively. μ was set to .5. All initial weights were given random values between −.25 and +.25. The learning rate for the single output model was .025. The same learning rate was used for the dual output model, except for the connections from internal representation units to reaching output where the learning rate was reduced to .0025. The error from the reaching system in the dual output model was not propagated back to the weights to the internal representation layer, so that the addition of the reaching system did not influence the performance of the prediction system.

For the single-ball simulations, equal numbers of ball-only, barrier-only, ball–barrier, and nothing stimuli were presented on each epoch. Five of each stimulus were presented per epoch, representing each of five possible delay conditions for the stimuli with the moving barrier, with corresponding ball-only and nothing stimuli of matching length. To reduce computational demands for the multiple-object simulations, given the 30-fold increase in number of objects to present, the number of stimulus presentations was reduced. On each epoch, three barrier-only and three ball–barrier events were presented for each object (corresponding to delay conditions of three, five, and seven time steps), along with a single ball-only and a single nothing event matched in length with the longest barrier-only and ball–barrier event. Context units were reset to 0 between stimuli. Networks were trained with a random ordering of stimuli and were tested at the end of epochs.

Received July 17, 1995
Revision received November 18, 1996
Accepted November 18, 1996 ■