To appear in *Developmental Review*

The Agent-Based Approach: A New Direction for Computational Models of Development

Matthew Schlesinger[1] and Domenico Parisi[2]

[1]University of Massachusetts, Amherst

[2]Italian National Research Council, Rome

Abstract

The agent-based approach emphasizes the importance of learning through organism-environment interaction. This approach is part of a recent trend in computational models of learning and development, toward studying autonomous organisms that are embedded in virtual or real environments. In this paper we introduce the concepts of online and offline sampling, and highlight the role of online sampling in agent-based models. After comparing the strengths of each approach for modeling particular developmental phenomena and research questions, we describe a recent agent-based model of infant causal perception. We conclude by discussing some of the present limitations of agent-based models, and suggesting how these challenges may be addressed.

The Agent-Based Approach: A New Direction for

Computational Models of Development

Computational models of learning and development are playing an increasingly critical role in child development research (Cassidy, 1990; Clark & Karmiloff-Smith; 1993; Edelman, 1987; Elman et al., 1996; McClelland, 1995; Newell & Molenaar, 1998; Parisi, 1996; Plunkett & Sinha, 1992; Quartz & Sejnowski, 1998; Ramsey & Stich, 1991; Rutkowska, 1987 & 1993; Simon & Halford, 1995). As a class of methodological tools, these models (1) reveal many of our fundamental assumptions regarding organism-environment interaction, (2) provide insight into multi-timescale and multi-level dynamic systems, and more generally, (3) both inform and are informed by concurrent behavioral research. To help illustrate the wide range of developmental phenomena that are open to simulation, Table 1 lists a brief sample of recent computational models.

Rather than re-present the basic features of computational models in psychology, as well as their importance for studying developmental phenomena, we refer the interested reader to excellent introductory reviews by Elman et al. (1996) and Simon and Halford (1995). Instead, our primary goal is to propose a new conceptual approach that expands on the existing array of computational models available to developmental researchers.

The **agent-based** approach spans a multidisciplinary group of researchers and theorists in cognitive science, incorporating ideas from evolutionary biology, robotics, linguistics, philosophy, neuroscience, psychology, and numerous other disciplines. A general theme across much of this work is the **embodied** or **situated agent**, an artificial or real organism that possesses a physical body and interacts with its environment by receiving sensory signals and producing motor outputs. From the perspective of developmental processes, the critical feature of agent-based models is the concept of online sampling, a computational procedure that enables autonomous agents to effectively control the sequence of sensory patterns they receive from the environment.

We begin by introducing the concepts of online and offline sampling. In the second section, we identify some of the developmental issues that each approach is ideally suited to address. In particular, we focus on the role of online sampling in agent-based models, and how employing an agent-based perspective can advance our understanding of development. Section three follows with an in-depth discussion and analysis of a recent agent-based model of infant causal perception, which illustrates many of the key features of the agent-based approach. In the final section, we discuss some of the present limitations of agent-based models, and suggest how these challenges may be addressed.

Offline and Online Sampling

The general goal of most computational models of development is to investigate how experience changes behavior. To this end, a ubiquitous feature of most models is the idea of an organism that receives sensory data from the environment, produces a response, and uses environmental feedback to improve its response. Typically, models equate "experience" with short exchanges or interactions with the environment, generated according to a sampling rule that dictates how sensory data varies over time.

In this section, we introduce a formal system for describing an arbitrary sampling rule, and use the system to illustrate two general classes of sampling rules. First, we define the input space $\boldsymbol{I}$ as the set of m possible input patterns, and i_t as a sequence of patterns sampled from $\boldsymbol{I}$ over a time interval t. More formally, $i_t \in \boldsymbol{I}$. Similarly, $o_t \in \boldsymbol{O}$ for the sequence of outputs produced by the model in response to the corresponding inputs. When $t = m$, the entire set of input patterns is exhaustively sampled (a single sweep through $\boldsymbol{I}$ is normally referred to as an **epoch**). An important question concerns how the input sequences i_t are generated. Typically, i_t is selected according to some well-defined function:

$$i_t = f(\boldsymbol{I}, t)$$

In the conventional case (e.g., supervised learning), f selects input patterns at random from $\boldsymbol{I}$ without replacement:

$$i_t = random(t)$$

Figure 1a presents a hypothetical sequence of input-output pairs, following this selection regime, while Figure 1b illustrates the corresponding trajectory through a 2D input space $\boldsymbol{I}$.

Because many model architectures and learning algorithms are sensitive to initial learning conditions, one side effect of random sampling is that two independent samples of $\boldsymbol{I}$ may result in different learning outcomes. This is due in part to the fact that the first few input patterns encountered may create a response bias that facilitates or interferes with later learning (e.g., the model may become trapped in a local minimum). One strategy for dealing with this problem is to train many instances of a model, thereby suppressing the effects of specific input sequences by averaging the results over a large number of input samples.

We refer to the general approach of sampling from $\boldsymbol{I}$ at random as **offline sampling**. Most agent-based models, however, do not employ this approach. Instead, an **online sampling** (or *active* sampling, see Parisi, 1997) procedure is used, which samples from the input space over time as a function of the model's output:

$$i(\mathrm{t}) = i_{t-1} + g(o_{t-1})$$

where i(t) is the model's input at time t, i_{t-1} and o_{t-1} are the model's input and output, respectively, during the previous timestep, and g is a function which computes the incremental change in the input state given the model's response. In other words, each response generated by the model has an iterative effect on its position in the input space ***I***. Rather than discrete input-output pairs, online sampling generates time-ordered, action-dependent sequences over a continuous episode.[1, 2] Figures 2a and 2b illustrate a hypothetical sequence of input-output states generated according to this scheme, and the corresponding trajectory through ***I***.

The key difference between offline and online sampling concerns how time is experienced by the model. During offline sampling, each input pattern constitutes a discrete, independent moment in time[3] (see Figure 1b). In addition, although the model may receive feedback after each output (e.g., a training signal), it does not experience any short-term or long-term consequences of the response. In many respects, the model's experience during offline sampling is analogous to a human subject in a typical psychology experiment: (1) a random sequence of discrete training exemplars is presented, (2) responses are independent (i.e., one response does not affect another), and (3) interference effects between successive exemplars are treated as noise, and consequently minimized by repeated training runs.

If we think of offline sampling as a procedure analogous to a laboratory experiment, then online sampling is akin to an ethological study: we first design a virtual organism and a quasi-natural learning environment (or perhaps, a robotic organism in a real environment), and then leave the organism free to explore and learn through interaction with its environment (we refer elsewhere to this as an **ecological** approach, Parisi, 1996; Parisi, Cecconi, & Nolfi, 1990). Because each behavior is followed by not only a potential training signal (where appropriate), but also the sensory consequences, our organism experiences both the immediate and long-term effects of its behavior.

Sampling and Development

In designing a computational model of development, the model-builder must decide whether to employ an online or offline sampling procedure. Fortunately, this decision can often be made independently of the specific learning algorithm, modeling framework, and theoretical orientation (e.g., constructivism versus nativism) that are being studied. Perhaps the most important consideration in choosing between online and offline sampling is whether temporal structure is a relevant aspect of the developmental phenomenon under investigation. In addition, the specific kinds of questions that the developmental model-builder hopes to answer will also impact on this choice. In order to highlight the particular strengths of each approach, we next briefly identify some of the general developmental issues that offline and online sampling are each ideally suited to address.

Offline Sampling and Development

Most of the models presented in Table 1 employ an offline sampling procedure. Many of these simulations focus on a related set of questions. First, how can we use computational models of continuous processes to help explain discontinuity in behavior over time? Second, how can computational models identify structure and change in internal representations? And third, what environmental constraints on development are revealed through modeling?

1. Transition mechanisms. Several developmental models have used offline sampling to study transition mechanisms in development. For example, McClelland (1989; McClelland & Jenkins, 1991), Raijmakers, van Koten, and Molenaar (1996), and Shultz, Mareschal, and Schmidt (1994) each present a connectionist model which replicates the four stages of reasoning observed in children on Piaget's balance scale task (Siegler, 1976). Each model captures the process of qualitative, stagewise change with a different connectionist approach (e.g., modular information processing, generative networks). Shultz has also used generative networks (i.e., the

cascade correlation learning algorithm) to study other stagewise developmental patterns, including children's reasoning on seriation and conservation tasks (Mareschal & Shultz, 1993; Shultz, 1998). Similarly, in the domain of language acquisition, Rumelhart and McClelland (1986) and Plunkett and Marchman (1991) show how U-shaped learning trajectories emerge in networks that learn to generate past-tense forms of English verbs (e.g., talk ⇒ talked, go ⇒ went).

2. Internal representations. Second, offline sampling has also been used to study the development of adaptive internal representations. Several descriptive and analytical methods are available (Elman, 1993; Sejnowski & Rosenberg, 1986). For example, Shultz (1998) presents a principal component analysis (PCA) of a generative network as it learns to conserve number. A related, graphical method is implemented in simulations of perceptual categorization (Mareschal & French, in press; Quinn & Johnson, 1997); these models are used to produce 3D maps of the internal representation space during training. Similarly, Mareschal, Plunkett, & Harris (1999) and Munakata, McClelland, Johnson, & Siegler (1997) use recurrent networks to model the infant's ability to predict the reappearance of an occluded object. An intriguing aspect of these models is how they employ a computational framework to study the putative representations underlying infants' sensorimotor actions (e.g., gaze control, reaching). This work not only coincides with a growing interest in studying representational processes in young infants (Goubet & Clifton, 1998; Leslie, Xu, Tremoulet, & Scholl, 1998; Wilcox & Baillargeon, 1998), but also offers a new and promising set of research tools that complement behavioral studies of infants' expectations (Needham & Baillargeon, 1997; Schlesinger & Langer, 1999; Wynn, 1992).

3. Environmental regularities. A third, related area of modeling work that employs offline sampling has helped developmental researchers to not only identify and study many of the

environmental regularities that are available to infants and young children, but also analyze how these regularities act as constraints on development. Models of language acquisition, in particular, provide a number of examples (Karmiloff-Smith, 1992; Plunkett & Marchman, 1991; Rumelhart & McClelland, 1986). Plunkett and Sinha (1992), for instance, show how a variety of language acquisition phenomena depend on the statistical properties of the training set (e.g., type-token frequency and rule overgeneralization). Similar methodological approaches (e.g., environmental biases, selective pretraining, differential learning rates) have also been implemented in models of object and numerical cognition (Shultz, 1998; Shultz, Schmidt, Buckingham, & Mareschal, 1995; Munakata et al., 1997) as well as perceptual categorization (Mareschal & French, in press; Quinn & Johnson, 1997).

Online Sampling and Development

As the key feature of the agent-based approach, online sampling allows researchers to simulate and investigate developmental phenomena that span continuous stretches of time. We describe here three broad developmental questions that depend in one way or another on the concept of time, and which are appropriate areas of study from the agent-based approach. First, how can we simulate processes of organism-environment interaction that unfold over more than a single timestep? Second, how might self-selection of sensory states act as an emergent constraint on developmental processes? And third, how can agents learn to exploit perception-action linkages?

1. Organism-environment interaction. Because agents exist in time, they can participate in exchanges or “dialogues” with their environment, in which multiple actions are attempted, evaluated, and selected by the agent. This makes the agent-based approach particularly relevant for studying the emergence and organization of sequential problem-solving strategies, such as

tool-use development in young infants (Bates, Carlson-Luden, & Bretherton, 1980; Koslowski & Bruner, 1972; Schlesinger & Langer, 1999; Willats, 1984). What is critical about online sampling as a modeling approach for these kinds of phenomena is that it supports directed trial-and-error learning, in which a strategy is progressively refined (Koslowski & Bruner, 1972; Willats, 1990). In addition, online sampling allows the agent to both explore the consequences of various primitive actions (e.g., pull tool, lift tool), while also experimenting with alternative ways of assembling these components into goal-directed sequences. In contrast, offline sampling typically captures learning processes at the timescale of single, primitive actions (however, see Bertsekas & Tsitsiklis, 1996, for a discussion of dynamic-programming approaches to this problem).

Interestingly, a difficulty in computational learning theory arises when an agent is allowed to generate sequences of actions over time that are not evaluated until the sequence is complete (e.g., a target is reached, a game is won or lost, a perceptual input is decomposed and categorized). The **credit-assignment problem** (Minsky, 1995) results from the fact that while many primitive actions are chosen during the sequence, not all contribute equally to the success of the action sequence as a whole. How much credit does each action receive? Fortunately, there are a number of learning algorithms that address this problem (e.g., temporal-difference learning, Sutton & Barto, 1998).

More generally, agent-based models are an appropriate tool for simulating developmental phenomena that depend critically on experiencing how one's actions transform perceptual input over time. For instance, Hiraki, Sashima, and Phillips (1998) employed an agent-based model in order to investigate the effect of locomotion on the developmental shift in young infants from egocentric to allocentric spatial perception. Similarly, agent-based models might be used to investigate the development of perceptual constancies, as well as the relation between optical flow and postural control.

In addition to interaction with the physical world, agents can also interact with their social environment. For example, Figure 3 illustrates a multi-agent model, in which two agents experience a shared input space. One application of this type of model is the study of language evolution and development in groups of agents that acquire a shared communication system (e.g., agents that both produce and perceive speech-like signals, Cangelosi & Parisi, 1998; Steels, 1998). These agent-based models extend the range of offline-sampling models of language acquisition (e.g., Elman, 1993; Plunkett & Marchman, 1991; Rumelhart & McClelland, 1986), by not only simulating the acquisition of semantic and syntactic properties of language, but also by incorporating some of the social and pragmatic factors that motivate early language development. Other aspects of social interaction open to investigation with agent-based models include the evolution and development of attachment, social-referencing, perspective taking, theory of mind, and cooperative social strategies.

2. <u>Sensory self-selection</u>. Because agents self-select their own input states, they effectively choose their own trajectories through the learning space (see Figure 2b). One important consequence is that rather than sampling the entire space of inputs uniformly, agents often limit their early learning to small portions of the sensory space (Nolfi & Parisi, 1993; Schlesinger, Parisi, & Langer, 2000). For example, Schlesinger, Parisi, and Langer (2000) observed an early, persistent pattern of stereotyped reaching movements in an agent-based model of hand-eye control, which was recently identified in human infants as well (Berthier, Clifton, McCall, & Robin, 1999).

From a developmental perspective, sensory self-selection reveals a number of important phenomena. First, it illustrates how the agent can effectively decompose a large problem into a sequence of smaller, tractable steps by concentrating early learning in a particular region of the input space. In contrast, by exhaustively sweeping through the input space, offline sampling often prevents such "divide-and-conquer" solutions from emerging. Second, it creates a source of constraint on developmental trajectories that, unlike the kinds of preprogrammed or preexisting constraints studied with the offline sampling approach (e.g., maturational factors, environmental regularities), is both emergent and unique to each organism's history of experience.

What kinds of developmental phenomena might be related to sensory self-selection in agent-based models, and amenable to simulation? One analogous phenomenon is the development of overt shifts of attention in infants (for a recent review, see Posner, Rothbart, Thomas-Thrapp, & Gerardi, 1998). More generally, sensory self-selection is appropriate for studying how infants and young children learn to process to complex, multi-dimensional stimuli, such as the development of speech segmentation, cross-modal perception, and perceptual categorizing.

3. <u>Perception-action linkages</u>. As Figure 2a illustrates, online sampling replaces the idea of feedforward information processing (i.e., from sensation to action) with the image of a chain or loop; sensation and action both influence and are influenced by interaction with the environment. As the broad theoretical relevance of such perception-action linkages has been argued elsewhere (e.g., embodied knowledge, circular reactions, reentrant mappings; Bechtel, 1997; Edelman, 1987; Parisi, 1995; Piaget, 1952; Smith & Scheier, 1998; Varela, Thompson, & Rosch, 1991), we focus our attention here on implications for development.

First, a robust finding among a variety of agent-based models is that adaptive agents rapidly learn to discover and exploit contingencies between their actions and subsequent perceptions. In many cases, these models rely on relatively simple associative learning strategies. For example, Treves, Miglino, and Parisi (1992) describe a series of simulations with an artificial neural network that learns under selective pressure to solve a navigation task by composing a sequence of adaptive sensorimotor behaviors (e.g., orienting, wall-following, etc.). Similarly, Scheier & Pfeifer (1995) study "body-based", sensorimotor classifying strategies in a mobile robot. Finally, Nolfi and Parisi (1999) present several examples of the "perceptual aliasing" problem (where two or more identical sensory input patterns require different responses, Whitehead & Ballard, 1991), and show how associative learning mechanisms can be used to link together sequences of perception-action mappings that either solve or avoid the problem.

Second, the fact that these sensorimotor strategies emerge early and often in learning suggests that they may also play a fundamental role in infant sensorimotor development. Consequently, one of the places that agent-based models can make a major impact is in the current debate on early infant knowledge (Baillargeon, 1999; Haith, 1998; Smith, 1999; Spelke, 1998). Of particular interest in this debate are the questions of how and when infants learn to maintain and operate on mental representations of the physical world. The prominent availability and use of sensorimotor or body-based representations in agent-based models support the idea that sensorimotor representations may developmentally precede and provide a structural foundation for later symbolic representations[4] (Ballard, Hayhoe, Pook, & Rao, 1997; Brooks, 1991; Gibson, 1988; Johnson, 1987; Keijzer, 1998; Lakoff, 1987; MacDorman, 1999; Rutkowska, 1996; Piaget, 1952). In the next section, we discuss in greater detail a recent model of infant causal perception that illustrates this conceptual approach (Schlesinger & Barto, 1999).

Modeling Infant Causal Perception

Thus far, much of our discussion has focused on the kinds of developmental questions that agent-based models are appropriate for investigating. In particular, we have argued that some types of developmental phenomena, especially those that unfold over multiple input-output exchanges with the environment, are difficult to capture within an offline-sampling framework. We now turn our attention to a specific area of development, with three goals in mind. First, we demonstrate how an agent-based model of infant causal perception is constructed. Second, we contrast both the structure and performance of the model with two similar models that employ an offline sampling approach. Finally, we highlight some of the specific ways in which such a dialogue between these alternative models can help advance our understanding of development.

Tracking Occluded Objects

Figure 4a presents a schematic display of an object that passes behind a screen and reappears on the other side. By age 6 months, infants will not only track these kinds of movements, but also anticipate the object's reappearance once it is occluded (Aslin, 1987; Bremner, 1985; Bower, 1979; Bower, Broughton, & Moore, 1971; Goldberg, 1976; Meicler & Gratch, 1980; Rutkowska, 1993). What cognitive mechanisms are necessary to explain the development of this ability?

According to one approach, we might assume that infants construct a mental representation of the occluded object, and mentally compute the future location of the object based on its observed trajectory. Two recent computational models employing this approach (Mareschal et al., 1999; Munakata et al., 1997) provide insight into how such representations might emerge, and how they might be recruited to guide prospective behavior (i.e., anticipating the object's reappearance).

Alternatively, it is possible that neither mental representations nor computations of future states are necessary to account for this kind of prospective behavior in infants. Instead, anticipation may be an adaptive by-product of learning to track the object. Schlesinger and Barto (1999) investigated this second account by constructing an agent-based model of tracking and gaze control in infants.

The simulated infant in the Schlesinger and Barto model is presented with movement events like those presented in Figure 4a. While the display is composed of 100x20 pixels, the infant's visual field covers a 20x20 pixel region, and can be moved laterally. Figure 5 re-presents the previous tracking display, with the infant's visual field (the square grid) superimposed over the display (note how the reappearance of the target is anticipated by the model). On each timestep, the 400 pixels in the visual field serve as the infant's sensory input. These pixel values propagate through a feedforward artificial neural network, resulting in a motor signal that drives the visual field to the left or right (thus contributing to changes in visual input).

The Schlesinger and Barto model simulates an infant that learns to track a moving object during both visible and occluded movement events. In the occluded or **screen** event (Figure 4a), the target object moves from the left to the right side of the display, passing behind a rectangular occluder in the center of the display. The simulated infant is also presented with a fully-visible **wall** event (Figure 4b), in which the target remains in place after encountering a "wall" near the center of the display.

The principal assumption of the Schlesinger and Barto model is that sight of the target is reinforcing to infants. Because the simulated infant's visual field spans 20% of the display, the infant must learn to generate a sequence of (lateral) eye movements that keep the target in the visual field. Consequently, a reinforcement learning algorithm was used to strengthen the reoccurrence of eye movements that were followed by sight of the target (see Schlesinger and Barto, 1999, for a detailed description of the neural network architecture and learning algorithm). In addition to the screen event, the model was also presented with the wall event during training, in order to insure that it learned to track the target rather than generate a stereotyped sequence of eye movements.

Note that different patterns of eye movements over the same tracking event will produce different input sequences. Thus, the Schlesinger and Barto model employs an online sampling procedure. This is in contrast to the Munakata and Mareschal models, which present analogous occluded event sequences, but train with an offline sampling procedure (i.e., event sequences that are both fixed and independent of the model's output).

In one set of simulations, we trained the model on 200 trials each of the screen and wall events. The model replicated several important developmental phenomena. First, the model learns to track a fully visible object before it tracks an occluded one (Bremner, 1985; Bower, 1979; Bower, Broughton, & Moore, 1971). Second, as the model learns to track the occluded target, it also anticipates its reappearance. Third, as the width of the obstacle increases (i.e., occlusion time), the model requires more training time to learn to track the target.

The model also replicates an unusual phenomenon in the development of infant occluded tracking: the *visual* A-not-B behavior. That is, successful tracking of occluded objects is preceded by a stage in which infants respond to the disappearance of the target by fixating the left edge of the screen, where it was last seen (Bower, 1979; Bower, Broughton, & Moore, 1971). Figure 6 helps illustrate an analogous tracking strategy in the model. Each horizontal dotted line represents the average position of the (left and right edges of the) visual field during the 13 timesteps of occlusion in the screen event (Figure 6a), or during the corresponding timespan in the wall event (Figure 6b). The horizontal axis of the figure represents the width of the tracking display, while the vertical axis represents the extent of training time in trials. The light gray regions represent the positions of the screen and wall within their respective events.

Figure 6a tells us where the simulated infant is “looking” during the occluded portion of the screen event. As in human infants, the model initially responds to the disappearance of the target by fixating the left edge of the screen (trials 1-60, screen event). By 80 trials, however, the model begins to successfully track the occluded target, consistently anticipating its reappearance on the right side of the screen. In contrast, the model learns to fixate the fully visible target in the wall event by virtually the start of training.

Thus, the Schlesinger and Barto model matches the performance of the Mareschal and Munakata models, employing a behavior-based rather than a representation-based strategy for tracking occluded objects. However, an important limitation of all of these models is that while they can anticipate the outcome of an occluded *movement* event, it is not clear how they would respond to an occluded *casual* event. For example, what if the target encountered an obstacle while behind the screen? Recent experimental findings that young infants anticipate the outcomes of occluded causal events (e.g., occluded collisions, Baillargeon, 1986; Baillargeon & DeVos, 1991; Lucksinger, Cohen, & Madole, 1992) strongly suggest a capacity for explicit representational (and perhaps also physical reasoning) mechanisms in early infancy. However, we next show how the Schlesinger and Barto model can begin to provide a more parsimonious account of these findings.

Occluded Causal Perception

It is possible to take our model one step further by constructing an occluded causal event that includes the perceptual features of both the screen and wall events. Figure 7 illustrates the novel **wall-screen** event. Notice that while the target reappears in the screen event, it does not in the wall-screen event, because of the obstructing wall.

How do human infants respond to the wall-screen event? In a recent study of occluded tracking in infants (Berthier, Bertenthal, Seaks, Sylvia, Johnson, & Clifton, 1999), 9-month-olds exhibited disrupted tracking of the target during the wall-screen event. We can explain this reaction in one of two ways. According to the representational account discussed above, because infants mentally reconstruct the target's path along with the new obstacle, they "know" the target will be blocked and therefore stop tracking it. Alternatively, we might argue that the disruption in tracking is a simple generalization from the earlier experience of watching the screen and wall events: infants learn to associate sight of the exposed wall with the obstruction of the target.

We conducted a series of simulations with the Schlesinger and Barto model to investigate these two accounts. In Study 1, the model was first trained on the screen and wall events (see Figure 4). After training, the model was then presented with both the novel wall-screen event, as well as the familiar screen event[5]. In order to compare the model's tracking behavior during the screen and wall-screen events, we recorded when the model first fixated the right edge of the screen. For the screen event, the model's **tracking latency** was then defined as the difference between this time-point and the time of the target's reappearance. Thus, a negative tracking latency indicates anticipation (i.e., fixation of the target's point of reappearance before it arrived there), while a positive latency indicates delayed tracking. The same temporal indices were also used to compute tracking latency during the wall-screen event, assuming reappearance of the block had it not been obstructed.

Figure 8a presents the mean tracking latencies for both the screen and wall-screen events in

Study 1, averaged over 50 runs (replications of the model). During the screen event, the model anticipates the reappearance of the target by nearly 10 timesteps. In contrast, when a partially hidden wall obstructs the target's path, tracking of the target is delayed by nearly 20 timesteps. Thus, like 9-month-old infants, the model generalizes appropriately to the novel wall-screen event, only exhibiting anticipatory tracking when the target is unobstructed[6]. However, in the case of the model it is clear that this response pattern is based on associative perceptual cues (e.g., the exposed portion of the wall), and not a "mental" reconstruction of the occluded target and its trajectory.

An important question is whether these findings are limited to partially hidden causal events. In order to evaluate this question, we modified the Schlesinger and Barto model in Study 2 by making the wall *shorter* than the screen (see Figure 9a), so that it was fully hidden when occluded[7]. Figure 8b shows that in this condition, the model appears to "forget" about the occluded wall, by anticipating the target's reappearance during both screen and wall-screen events. While this appears to be a failure of the model, notice that there is no penalty for always waiting for the target to reappear on the right of the screen, even if it only reappears during the screen event.

However, what if the target *did* reappear during both screen and wall-screen events, but reappeared on *different* sides of the screen depending on whether or not the wall was present? In Study 3, we modified the physics of our simple world by allowing the target to "rebound" off the short wall (see Figure 9b). This change meant that the target reappears on the right of the screen during screen events, and on the left of the screen during wall-screen events. Thus, it allowed us to investigate if the presence or absence of the completely occluded wall would affect the model's anticipatory tracking of the target. (Recall that as in Study 2, both the screen and wall-screen test events were preceded by a short exposure to what was behind the screen.) Figure 8c presents the mean tracking latencies for this third rebound condition. As in Study 1, the model only anticipates the reappearance of the target on the right side of the screen when its path is unobstructed.

Summarizing the Results

While learning to follow a moving object, the Schlesinger and Barto tracking model also learns to anticipate the outcomes of both partially and fully occluded causal events. Rather than explicitly representing the event features (e.g., the trajectory of the target, the location of the wall), the model implicitly encodes these features in its tracking behavior (e.g., "stop tracking if wall is seen"). Interestingly, these sensorimotor representations are able to capture and exploit both perceptually current and past features of the occluded events. Taken together, these simulation findings are consistent with a growing body of evidence that implicate perceptual processes as a critical factor in infants' reactions to occluded causal events (Bogartz & Shinskey, 1998; Rivera, Wakeley, & Langer, 1999; Schilling, 1997).

Unfortunately, we cannot present a "head-to-head" comparison between models employing online and offline sampling, in part because the Mareschal and Munakata models have only been used thus far to study occluded object tracking (but not yet occluded causal perception). Nevertheless, one of the key insights offered by the Schlesinger and Barto model is how prospective or future-oriented behavior can emerge from simple associative learning mechanisms. While a comparable level of explanation may not necessarily apply to human infants, the results can be used as an existence proof to demonstrate that in addition to the representation-based approach advocated by the Mareschal and Munakata models, a behavior-based solution is also available to infants. However, one intriguing theoretical account that may integrate these two levels of explanation developmentally is based on the idea that both strategies (e.g., fully representational and behavior-based tracking) are available to infants, but that infants are biased toward employing sensorimotor representations earlier and more often than abstract mental representations due to task and processing constraints (see Ballard et al., 1997, for a related account of representational processes in adults).

However, it is clear that some aspects of infant occluded-object tracking cannot be captured by sensorimotor representation alone. For example, infants respond with surprise, search behavior, and other reactions when an occluded object reappears with new features (e.g., color, shape, etc.; Wilcox & Baillargeon, 1998; Meicler & Gratch, 1980; Meltzoff & Moore, 1998). Although the Schlesinger and Barto model is capable of maintaining a sensory trace of past inputs (e.g., Studies 2 and 3), it is not designed to compare that trace with current sensory inputs (see Simon, 1998, for a model that is explicitly constructed to make such a comparison). In addition, because the Mareschal and Munakata models learn by comparing predicted future inputs with observed inputs, they also have a built-in mechanism for simulating causal expectations and surprise in infants (Baillargeon, 1993; Spelke, 1985). Thus, sensorimotor representations may be limited to capturing specific kinds of environmental features, especially those that serve as perceptual cues to adaptive behavior (e.g., geometric properties such as object positions and trajectories, but not colors or shapes; see Meltzoff & Moore, 1998, for a developmental account of occluded-object tracking with respect to these features).

Conclusions

Online sampling introduces time as a critical dimension of organism-environment interaction. Unlike models that employ offline sampling, where interaction is limited to single input-output exchanges, agent-based models simulate organisms that exist in their environment over time, experiencing both the immediate and long-term consequences of their behavior. Our goal in advocating the agent-based approach has been to show that it offers a number of opportunities and potential insights to developmental researchers.

We began by describing offline and online sampling procedures, and illustrating the key implementation differences between these two modeling approaches. As a rough analogy, we noted a correspondence between offline sampling and psychology learning experiments, and between online sampling and ethological studies. Thus, online and offline sampling provide complementary modeling approaches. While offline sampling simulates learning and development under highly controlled experimental conditions (e.g., stimulus order), online sampling captures many of the characteristics of naturalistic conditions.

We then described some of the relative strengths of online and offline sampling for modeling developmental phenomena. Offline-sampling models have been used by developmental researchers to study transition mechanisms in development, as well as the detection and representation of structural regularities in the environment. Because temporal structure plays a critical role in online sampling, agent-based models of development are ideal for simulating temporally extended sequences of organism-environment interaction. We mentioned some of the possible (social, cognitive, linguistic, etc.) developmental phenomena that an agent-based approach might be used to study.

Two key insights for developmental researchers are revealed by recent work in agent-based models. First, an interesting consequence of being able to control their sensory inputs over time is that agents often concentrate their early developmental experiences in small portions of the input space. We suggested one possible use of this sensory self-selection effect in studying how infants segment or decompose complex perceptual inputs. Second, agents rapidly learn to detect contingencies or regularities between their actions and perceptions, and exploit these perception-action linkages to simplify complex problems. We argued that this phenomenon might help to inform the current debate among cognitive developmentalists on the origins of representation in infancy.

The latter point was then explored in detail by presenting findings from a recent agent-based model of infant causal perception. We described the Schlesinger and Barto tracking model, and compared its performance to two similar models that employ an offline sampling method for studying the development of occluded-object tracking. The point of contrasting these two modeling approaches was not to demonstrate that certain developmental characteristics can only be captured with online sampling, but rather to illustrate how each approach provides a unique perspective on a given phenomenon. In the case of occluded-object tracking and causal perception, one of the principle lessons of the agent-based model was that representational explanations for prospective behavior in infants need not be invoked when more parsimonious accounts are available.

Thus, the agent-based approach complements the wide variety of computational models that are already available to developmental researchers. However, because it has only recently begun to emerge as a coherent framework for developmental model-builders, the agent-based approach will continue to evolve as new ideas and features are introduced. We conclude with three of the major challenges that agent-based models will face.

First, are agent-based models limited to simulating sensorimotor development (e.g., motor control)? Because agents are typically situated or embodied (i.e., sensory inputs and motor outputs), it is perhaps natural that most agent-based modeling work thus far has focused on learning and development in sensorimotor systems. While this characteristic may complement some theoretical approaches (e.g., that early knowledge is grounded in physical experience, Ballard et al., 1997; Johnson, 1987; Keijzer, 1998; Lakoff, 1987; Rutkowska, 1994; Varela, Thompson, & Rosch, 1991), it may also create a difficulty for developmental researchers if the psychological phenomenon being studied is not body-based or conveniently described in a physical form (e.g., development of temperament, gender roles, moral reasoning, etc.). However, the initial success of agent-based models of language acquisition (e.g., Cangelosi & Parisi, 1998; Steels, 1998) suggests that some types of developmental phenomena are not completely reducible to sensorimotor activity, and yet are still plausibly modeled with an agent-based approach.

A second and related question is whether agent-based models can capture "high-level" or abstract cognition (e.g., counterfactual or hypothetical reasoning, planning, metaphors). To a large extent, the answer to this question depends less on the fact that agents are almost by definition embodied, and more on what the model-builder places inside the agent's body (or more specifically, in the simulated brain of the organism). While the agent's interface with the physical and social world is necessarily sensorimotor (as it is for all biological organisms), there are no *a priori* structural limitations on the kinds of cognitive machinery that can be designed inside an autonomous agent. However, agents need not be physical. Indeed, game-theory models of social behavior (e.g., Maynard Smith, 1976) often simulate organisms who themselves are not only abstract entities (i.e., they have no body, do not occupy any space, etc.), but who also engage in abstract activities (e.g., buying or selling commodities) via abstract encounters. Nevertheless, one of the ways we anticipate that agent-based models may evolve to address this question is for them to "hybridize" with fully internalized models of abstract cognitive processes (especially with explicitly symbolic models, e.g., ACT-R; see Clark & Karmiloff-Smith, 1993, for a discussion of hybrid connectionist-symbolic models).

Finally, agent-based models are no better than their offline-sampling counterparts at escaping the “brittleness” problem. That is, while capable of discovering elegant solutions to a given problem, most models fail miserably when the task structure or context is changed. Thus, regardless of the sampling regime, task- and context-specific solutions predominate. Nevertheless this is a critical limitation of agent-based models, which online sampling is unable to address. It is clear that computational models will need to evolve away from simulating performance on small tasks (e.g., the balance scale) toward investigating development across tasks, contexts, and knowledge domains (Karmiloff-Smith, 1992). We believe that by placing an emphasis on situated organisms, the agent-based approach will help developmental model-builders to realize this change.

References

Aslin, R.N. (1987). Visual and auditory development in infancy. In J.D. Osofsky (Ed.), Handbook of infant development, 2nd ed., pp. 5-97. New York: Wiley and Sons.

Baillargeon, R. (1986). Representing the existence and the location of hidden objects: Object permanence in 6- and 8-month-old infants. Cognition, 23, 21-41.

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

Baillargeon, R. (1999). Young infants' expectations about hidden objects: A reply to three challenges. Developmental Science, 2, 115-132.

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

Ballard, D.H., Hayhoe, M.M., Pook, P.K., & Rao, R.P.N. (1997). Deictic codes for the embodiment of cognition. Behavioral & Brain Sciences, 20, 723-767.

Bates, E., Carlson-Luden, V., & Bretherton, I. (1980). Perceptual aspects of tool using in infancy. Infant Behavior and Development, 3, 127-140.

Bechtel, W. (1997). Embodied connectionism. In D.M. Johnson & C.E. Erneling (Eds.), The future of the cognitive revolution, pp. 187-208. New York: Oxford University Press.

Berthier, N.E. (1996). Learning to reach: A mathematical model. Developmental Psychology, 32, 811-823.

Berthier, N.E., Bertenthal, B.I., Seaks, J., Sylvia, M., Johnson, R., & Clifton, R.K. (1999). Catching a moving object: Infant reaching when trajectory information is removed. Manuscript in preparation.

Berthier, N.E., Clifton, R.K., McCall, D.D., & Robin, D.J. (1999). Proximodistal structure of initial reaching in human infants. Experimental Brain Research, 127, 259-269.

Bertsekas, D.P., & Tsitsiklis, J.N. (1996). Neuro-dynamic programming. Belmont, MA: Athena Scientific.

Bogartz, R.S., & Shinskey, J.L. (1998). On perception of a partially occluded object in 6-month-olds. Cognitive Development, 13, 141-163.

Bower, T.G.R. (1979). Human development. New York: W.H. Freeman.

Bower, T.G.R., Broughton, J.M., \& Moore, M.K. (1971). Development of the object concept as manifested in changes in the tracking behavior of infants between 7 and 20 weeks. Journal of Experimental Child Psychology, 11, 182-193.

Brooks, R.A. (1991). Intelligence without representation. Artificial Intelligence, 47, 139-159.

Cangelosi, A., & Parisi, D. (1998). The emergence of a "language" in an evolving population of neural networks. Connection Science, 10, 83-97.

Cassidy, S. (1990). When is a developmental model not a developmental model? Cognitive Systems, 2, 329-344.

Christiansen, M.H., Allen, J., & Seidenberg, M.S. (1998). Learning to segment speech using multiple cues: A connectionist model. Language & Cognitive Processes, 13, 221-268.

Clark, A., & Karmiloff-Smith, A. (1993). The cognizer's innards: A psychological and philosophical perspective on the development of thought. Mind and Language, 8, 487-519.

Dehaene, S., & Changeux, J.P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. Journal of Cognitive Neuroscience, 1, 244-261.

Dehaene, S., & Changeux, J.P. (1993). Development of elementary numerical abilities: A neuronal model. *Journal of Cognitive Neuroscience*, *5*, 390-407.

Edelman, G.M. (1987). *Neural darwinism*. New York: Basic.

Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*, 71-99.

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

Gibson, E.J. (1988). Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Review of Psychology*, *39*, 1-41.

Goldberg, S. (1976). Visual tracking and object constancy in 5-month-old infants. *Journal of Experimental Child Psychology*, *22*, 478-491.

Goubet, N., & Clifton, R.K. (1998). Object and event representation in 6 month-old infants. *Developmental Psychology*, *34*, 63-76.

Haith, M.M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behavior & Development*, *21*, 167-179.

Hartelman, P.A., van der Maas, H.L.J., & Molenaar, P.C.M. (1998). Detecting and modeling developmental transitions. *British Journal of Developmental Psychology*, *16*, 97-122.

Hiraki, K., Sashima, A., & Phillips, S. (1998). From egocentric to allocentric spatial behavior: A computational model of spatial development. *Adaptive Behavior*, *6*, 371-391.

Johnson, M. (1987). *The body in the mind: The bodily basis of imagination, reason, and meaning*. Chicago: University of Chicago Press.

Karmiloff-Smith, A. (1992). Nature, nurture, and PDP: Preposterous developmental postulates? *Connection Science*, *4*, 253-269.

Koslowski, B., & Bruner, J.S. (1972). Learning to use a lever. Child Development, 43, 790-799.

Kuperstein, M. (1991). INFANT neural controller for adaptive sensory-motor coordination. Neural Networks, 4, 131-145.

Lakoff, G. (1987). Women, fire and dangerous things: What categories reveal about the mind. Chicago: University of Chicago Press.

Leslie, A.M., Xu, F., Tremuoulet, P.D., & Scholl, B.J. (1998). Indexing and the object concept: Developing "what" and "where" systems. Trends in Cognitive Science, 2, 10-18.

Lucksinger, K.L., Cohen, L.B., and Madole, K.L. (1992, May). What infants infer about hidden objects and events. Presented at the International Conference for Infant Studies, Miami, FL.

Luger, G.F., Wishart, J.G., & Bower, T.G. (1984). Modeling the stages of the identity theory of object-concept development in infancy. Perception, 13, 97-115.

MacDorman, K. F. (1999). Grounding symbols through sensorimotor integration. Journal of the Robotics Society of Japan, 17, 20-24.

MacWhinney, B., & Chang, F. (1995). Connectionism and language learning. In C.A. Nelson (Ed.)., The Minnesota symposium on child psychology: Basic and applied perspectives on learning, cognition, and development, Vol. 28, pp 33-57. NJ: Lawrence Erlbaum Associates.

Mareschal, D., & French, R. (in press). Mechanisms of categorization in infancy. Infancy.

Mareschal, D., Plunkett, K., & Harris, P. (1999). A computational and neuropsychological account of object-oriented behaviours in infancy. Developmental Science, 2, 306-317

Mareschal, D., & Shultz, T.R. (1993). A connectionist model of the development of seriation. Proceedings of the fifteenth annual Conference of the Cognitive Science Society, pp.

170-175. Hillsdale, NJ: Lawrence Erlbaum.

Maynard Smith, J. (1976). Evolution and the theory of games. American Scientist, 64, 41-45.

McClelland, J.L. (1989). Parallel distributed processing: Implications for cognition and development. In R. Morris (Ed.), Parallel distributed processing: Implications for psychology and neurobiology, pp. 9-45. New York: Oxford University Press.

McClelland, J.L. (1995). A connectionist perspective on knowledge and development. In T.J. Simon & G.S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling, pp. 157-204. Hillsdale, NJ: Lawrence Erlbaum.

McClelland, J.L, & Jenkins, E. (1991). Nature, nurture, and connections: Implications of connectionist models for cognitive development. In K. van Lehn (Ed.), Architectures for intelligence, pp. 41-73. Hillsdale, NJ: Lawrence Erlbaum Associates.

Meicler, M., & Gratch, G. (1980). Do 5-month-olds show object conception in Piaget's sense? Infant Behavior and Development, 3, 265-282.

Meltzoff, A.N., & Moore, M.K. (1998). Object representation, identity, and the paradox of early permanence: Steps toward a new framework. Infant Behavior and Development, 21, 201-235.

Minsky, M. (1995). Steps toward artificial intelligence. In E.A. Feigenbaum & J. Feldman (Eds.), Computers and thought, pp. 406-450. Cambridge, MA: MIT Press. (Originally printed 1963).

Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: A PDP model of the $A\overline{B}$ task. Developmental Science, 1, 161-211.

Munakata, Y., McClelland, J.L., Johnson, M.H., & Siegler, R.S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. Psychological Review, 104, 686-713.

Needham, A., & Baillargeon, R. (1997). Object segregation in 8-month-old infants. Cognition, 62, 121-149.

Newell, K.M., & Molenaar, P.C.M. (1998). Applications of nonlinear dynamics to developmental process modeling. Mahway, NJ: Lawrence Erlbaum.

Nolfi, S., & Parisi, D. (1993). Self-selection of input stimuli for improving performance. In G. A. Bekey & K.Y. Goldberg (Eds.), Neural Networks in Robotics, pp. 403-418. Boston: Kluwer.

Nolfi, S., & Parisi, D. (1999). Exploiting the power of sensory-motor coordination. In: D. Floreano, J-D. Nicoud, & F. Mondada (Eds.), Advances in Artificial Life: Proceedings of the Fifth European Conference on Artificial Life, pp. 173-182. Berlin: Springer-Verlag.

Parisi, D. (1995, June). Learning in the active mode. Presented at the European conference on artificial life, Granada, Spain.

Parisi, D. (1996). Computational models of developmental mechanisms. In R. Gelman & T. Au (Eds.), Perceptual and Cognitive Development, pp. 373-412. San Diego: Academic Press.

Parisi, D. (1997). Active sampling in evolving neural networks. Human Development, 40, 320-324.

Parisi, D, Cecconi, F, & Nolfi, S. (1990). Econets: Neural networks that learn in an environment. Network, 1, 149-168.

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

Plunkett, K., & Marchman, V. (1991). U-shaped learning and frequency effects in a multilayered perceptron: Implications for child language acquisition. Cognition, 38, 1-60.

Plunkett, K., & Sinha, C.G. (1992). Connectionism and developmental theory. British Journal of Developmental Psychology, 10, 209-254.

Plunkett, K., Sinha, C.G., Moller, M.F., & Strandsby, O. (1992). Symbol grounding or the emergence of symbols? Vocabulary growth in children and a connectionist net. Connection Science, 4, 293-312.

Posner, M.I., Rothbart, M.K., Thomas-Thrapp, L., & Gerardi, G. (1998). The development of orienting to locations and objects. In R.D. Wright (Ed.), Visual attention: Vancouver studies in cognitive science, Vol. 8, pp. 269-288. NY: Oxford University Press.

Quartz, S., & Sejnowski, T.J. (1998). The neural basis of cognitive development: A constructivist manifesto. Behavioral & Brain Sciences, 20, 537-596.

Quinn, P.C., & Johnson, M.H. (1997). The emergence of perceptual category representations in young infants: A connectionist analysis. Journal of Experimental Child Psychology, 66, 236-263.

Raijmakers, M.E.J., van Koten, S., & Molenaar, P.C.M. (1996). On the validity of simulating stagewise development by means of PDP networks: Applications of catastrophe analysis and an experimental test of rule-like network performance. Cognitive Science, 20, 101-136.

Ramsey, W., & Stich, S.P. (1991). Connectionism and three levels of nativism. In W. Ramsey, S.P. Stich, & D.E. Rumelhart (Eds.), Developments in connectionist theory, pp. 287-310. Hillsdale, NJ: Lawrence Erlbaum.

Regier, T. (1996). The human semantic potential: Spatial language and constrained connectionism. Cambridge, MA: MIT Press.

Rivera, S.M., Wakeley, A., & Langer, J. (1999). The drawbridge phenomenon: Representational reasoning or perceptual preference? Developmental Psychology, 35, 427-435.

Rumelhart, D.E., & McClelland, J.L. (1986). On learning the past tense of English verbs. In J.L. McClelland, D.E. Rumelhart & PDP Research Group (Eds.), Parallel distributed processing: Vol. 2. Psychological and biological models, pp. 216-271. Cambridge, MA: MIT Press.

Rutkowska, J.C. (1987). Computational models and developmental psychology. In J.C. Rutkowska & C. Crook (Eds.), Computers, cognition and development: Issues for psychology and education, pp. 187-215. Chichester, England: John Wiley & Sons.

Rutkowska, J.C. (1993). The computational infant: Looking for developmental cognitive science. London: Harvester Wheatsheaf.

Rutkowska, J.C. (1994). Situating representational redescription in infants' pragmatic knowledge. Behavioral and Brain Sciences, 17, 726-727.

Rutkowska, J.C. (1996, September). Computation, dynamics, and sensory-motor development. Presented at the Piaget Centennial Conference, Geneva, Switzerland.

Scheier, C. & Pfeifer, R. (1995). Classification as sensorimotor co-ordination: A case study on autonomous agents. In F. Moran, A Moreno, P. Chacon, & J.J. Merelo (Eds.), Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life, pp. 657-667. New York: Springer Verlag.

Schilling, T.H. (1997, April). Infants' understanding of physical phenomena: A perceptual hypothesis. Presented at the meeting of the Society for Research in Child Development, Washington, D.C.

Schlesinger, M., & Barto, A. (1999). Optimal control methods for simulating the perception of causality in young infants. In M. Hahn & S.C. Stoness (Eds.), Proceedings of the Twenty First Annual Conference of the Cognitive Science Society, pp. 625-630. New Jersey: Erlbaum. (An electronic preprint is available for download from http://www-anw.cs.umass.edu/~matthew/academic.htm.)

Schlesinger, M., & Langer, J. (1999). Infants' developing expectations of possible and impossible tool-use events between ages 8 and 12 months. Developmental Science, 2, 195-205.

Schlesinger, M., Parisi, D., & Langer, J. (2000). Learning to reach by constraining the movement search space. Developmental Science, 3, 67-80.

Sejnowski, T., & Rosenberg, C. (1986). NETalk: A parallel network that learns to read aloud. Johns Hopkins University Electrical Engineering and Computer Science Technical Report, JHU/EEC-86/01.

Shultz, T.R. (1998). A computational analysis of conservation. Developmental Science, 1, 103-126.

Shultz, T.R., Mareschal, D., & Schmidt, W.C. (1994). Modeling cognitive development on balance scale phenomena. Machine Learning, 16, 57-86.

Shultz, T.R., Schmidt, W.C., Buckingham, D., & Mareschal, D. (1995). Modeling cognitive development with a connectionist algorithm. In T.J. Simon & G.S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling, pp. 205-261. Hillsdale, NJ: Lawrence Erlbaum.

Siegler, R.S. (1976). Three aspects of cognitive development. Cognitive Psychology, 8, 481-520.

Simon, T.J. (1998). Computational evidence for the foundations of numerical competence. Developmental Science, 1, 71-78.

Simon, T.J., & Halford, G.S. (1995). Computational models and cognitive change. In T.J. Simon & G.S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling, pp. 1-30. Hillsdale, NJ: Lawrence Erlbaum.

Smith, L.B. (1999). Do infants possess innate knowledge structures? The con side. Developmental Science, 2, 133-144.

Smith, L.B., & Scheier, C. (1998). Babies have bodies: Why Munakata's net fails to meet its own goals. Developmental Science, 1, 161-211.

Spelke, E.S. (1985). Preferential looking methods as a tool for the study of cognition in infancy. In G. Gottlieb & N. Krasnegor (Eds.), Measurement of audition and vision in the first year of postnatal life, pp. 323-363. Norwood, NJ: Ablex.

Spelke, E.S. (1998). Nativism, empiricism, and the origins of knowledge. Infant Behavior and Development, 21, 181-200.

Sporns, O., & Edelman, G.M. (1993). Solving Bernstein's problem: A proposal for the development of coordinated movement by selection. Child Development, 64, 960-981.

Steels, L. (1998). The origins of syntax in visually grounded robotic agents. Artificial Intelligence, 103, 133-156.

Sutton, R.S., & Barto, A.G. (1998). Reinforcement learning: An introduction. Cambridge, MA: MIT Press.

Treves., A., Miglino., O., & Parisi, D. (1992). Rats, nets, maps, and the emergence of place cells. Psychobiology, 20, 1-8.

van Geert, P. (1991). A dynamic model of cognitive and language growth. Psychological Review, 55, 189-208.

Varela, F.J, Thompson, E., & Rosch, E. (1991). The embodied mind: Cognitive science and human experience. Cambridge, MA: MIT Press.

Vos, J.E., & Scheepstra, K.A. (1993). Computer-simulated neural networks: An appropriate model for motor development? Early Human Development, 34, 101-112.

Whitehead, S.D., & Ballard, D. (1991). Learning to perceive and act by trial and error. Machine Learning, 7, 45-83.

Wilcox, T., & Baillargeon, R. (1998). Object individuation in young infants: Further evidence with an event-monitoring paradigm. Developmental Science, 1, 127-142.

Willats, P. (1984). The Stage-IV infant's solution of problems requiring the use of supports. Infant Behavior and Development, 7, 125-134.

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. NJ: Lawrence Erlbaum Associates.

Wynn, K. (1992). Addition and subtraction by human infants. Nature, 358, 749-750.

Author Note

Matthew Schlesinger, Department of Computer Science; Domenico Parisi, Institute of Psychology.

This work was supported in part by NSF IRI-9720345 and NSF CDA-9703217. The authors wish to thank Sascha Engelbrecht, Rod Grupen, Anders Jonsson, and Michael Rosenstein for their helpful comments and suggestions as we prepared this paper.

Correspondence concerning this article should be addressed to Matthew Schlesinger, Department of Computer Science, University of Massachusetts, Amherst, MA 01003. Electronic mail may be sent to matthew@cs.umass.edu.

Footnotes

[1] In some cases (e.g., Elman, 1993, Munakata, McClelland, Johnson, & Siegler, 1997) complex input patterns have temporal structure (e.g., a stream of speech sounds), and are consequently segmented and presented to the model one component at a time. In this context, offline sampling selects each of the complex patterns at random.

[2] Note that the use of online sampling does not alleviate the problem of sensitivity to initial conditions. As with offline sampling, one strategy is to train a model by starting each episode at a random position in the input space. For the remainder of the episode, however, all subsequent input patterns are determined by the actions of the agent.

[3] By *continuous* we mean that online sampling generates a sequence of input patterns from the beginning to the end of an episode, ordered with respect to time. Nevertheless, like offline sampling, online sampling normally samples the input space at *discrete* moments in time, at regular intervals.

[4] Our use of the terms s*ensorimotor* and *symbolic* representations is intended to be analogous with *procedural* and *declarative* knowledge. Sensorimotor representations encode context-specific properties of the environment in an implicit manner, and may not be available to other cognitive processes (e.g., language), while symbolic representations are normally separable from their context and accessible to conscious inspection.

[5] After training on screen and wall events, learning was "turned off" (i.e., connection weights were fixed). This procedural step, typical of most neural network models, eliminated the possibility that exposure to either test event (i.e., the novel wall-screen event and the familiar screen event) would contaminate the model's reaction to the other event.

[6] Control conditions in both the Berthier et al. experiments and the Schlesinger and Barto simulations ruled out the possibility that a simple distraction effect (caused by the exposed portion of the wall) could account for the delay in tracking.

[7] Other necessary changes to the model included: (1) adapting the original feedforward network to include recurrent connections from the hidden to the input layers, and (2) preceding each trial with a "preview" of what was behind the screen (i.e., either nothing, or the short wall).

Table 1

Survey of Recent Computational Models of Development, by Developmental Domain

Developmental Domain	Sub-Domain	Study
Physical Knowledge	Objects	Dehaene & Changeux, 1989
		Luger, Wishart, & Bower, 1984
		Mareschal, Plunkett, & Harris, in press
		Munakata, 1998
		Munakata, McClelland, Johnson, & Siegler, 1997
	Space	Hiraki, Sashima, & Phillips, 1998
Numerical Knowlege	Number Perception	Dehaene & Changeux, 1993
		Peterson & Simon, 2000
	Arithmetic	Simon, 1998
	Balance Scale	McClelland, 1989
		Raijmakers, van Koten, & Molenaar, 1996
		Shultz, Mareschal & Schmidt 1994
	Seriation	Mareschal & Shultz, 1993
	Conservation	Hartelman, van der Maas, & Molenaar, 1998
		Shultz, 1998
Classes & Categories	Perceptual Categorization	Mareschal & French, in press
		Quinn & Johnson, 1997
Language & Cognition	Speech Perception	Christiansen, Allen, & Seidenberg, 1998
	Syntax	Elman, 1993
		Plunkett & Marchman, 1991
		Rumelhart & McClelland, 1986
	Semantics & Lexical Growth	MacWhinney & Chang, 1995
		Plunkett, Sinha, Moller, & Strandsby, 1992
		Regier, 1996
		van Geert, 1991
Motor Control	Reaching	Berthier, 1996
		Kuperstein, 1991
		Schlesinger, Parisi, & Langer, 2000
		Sporns & Edelman, 1993
		Vos & Scheepstra, 1993

Figure Captions

Figure 1. (a) Hypothetical sequence of inputs ("i") and outputs ("o") generated by offline (random) sampling, and (b) the model's corresponding trajectory through a 2-dimensional input state-space. S_i's index discrete states in the input state-space.

Figure 2. (a) Hypothetical sequence of inputs and outputs generated by online sampling, and (b) the model's corresponding trajectory through a 2-dimensional input state-space (arrows indicate the output chosen for each input).

Figure 3. A two-agent model. Each agent co-determines the sequence of inputs that are perceived in a shared sensory input space.

Figure 4. Two tracking events. In the screen event (a), the target passes behind a screen and reappears on the other side; in the wall event (b), the target encounters a wall and stops moving.

Figure 5. Schematic display illustrating how the Schlesinger and Barto tracking model responds to the screen event. The grid (not shown to scale) represents the model's visual field, which learns to follow the target from left to right.

Figure 6. Average position of the Schlesinger and Barto model's visual field during occlusion of the target in (a) the screen event and the corresponding timespan in the (b) wall event, as a function of training time. The left and right portions of the horizontal lines represent the corresponding edges of the model's visual field.

Figure 7. The wall-screen event. The target passes behind the screen, encounters the partially hidden wall, and stops moving.

Figure 8. Tracking latency in a series of simulation studies. In all studies, the model anticipates the reappearance of the target during the screen event. In Studies 1 and 3, the model's tracking of the target is delayed by (a) a partially hidden wall or (c) a fully hidden wall. See text and

Figure 9 for details.

Figure 9. The wall-screen events simulated in Studies 2-3. In (a) the wall is fully hidden by the screen, and in (b) the wall is fully hidden, and the target "rebounds" off the wall.

Figure 1

(a) (b)

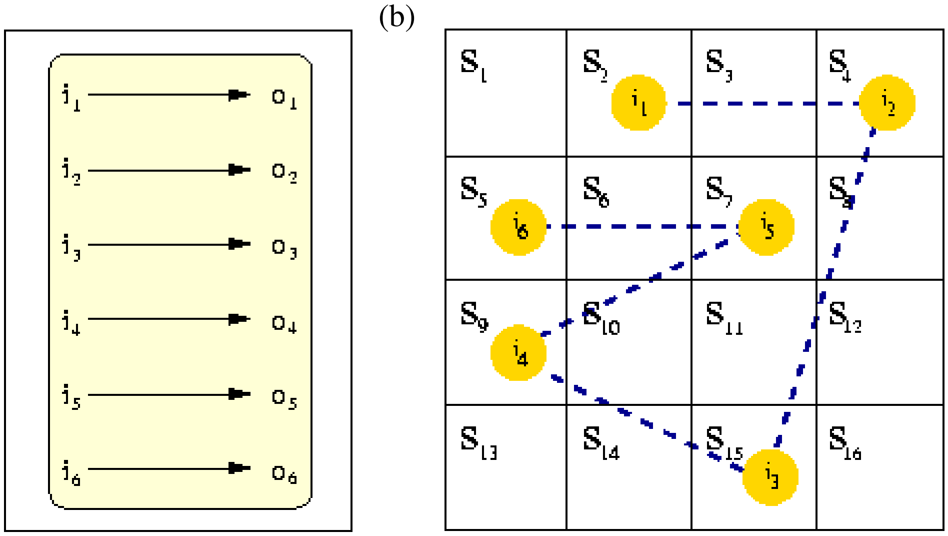
i1 o1
i2 o2
i3 o3
i4 o4
i5 o5
i6 o6
S1 S2 S3 S4
S5 S6 S7 S8
S9 S10 S11 S12
S13 S14 S15 S16
i1 i2 i3 i4 i5 i6

Figure 2

(a) (b)

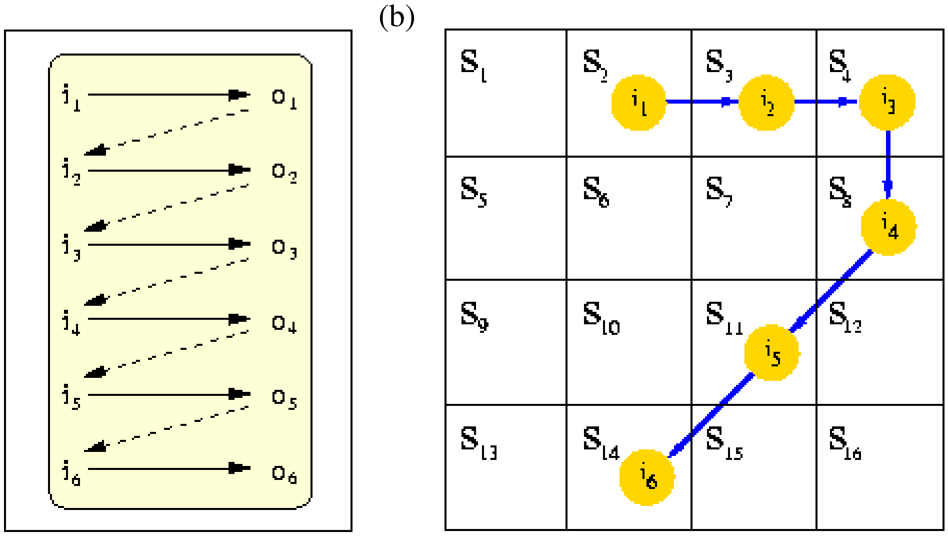
i1
o1
i2
o2
i3
o3
i4
o4
i5
o5
i6
o6
S1
S2
S3
S4
S5
S6
S7
S8
S9
S10
S11
S12
S13
S14
S15
S16
i1
i2
i3
i4
i5
i6

Figure 3

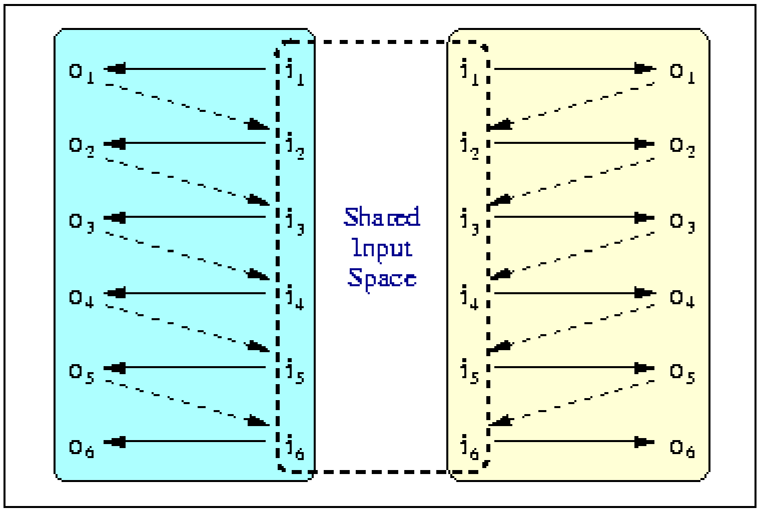
o1
o2
o3
o4
o5
o6
i1
i2
i3
i4
i5
i6
Shared
Input
Space
i1
i2
i3
i4
i5
i6
o1
o2
o3
o4
o5
o6

Figure 4

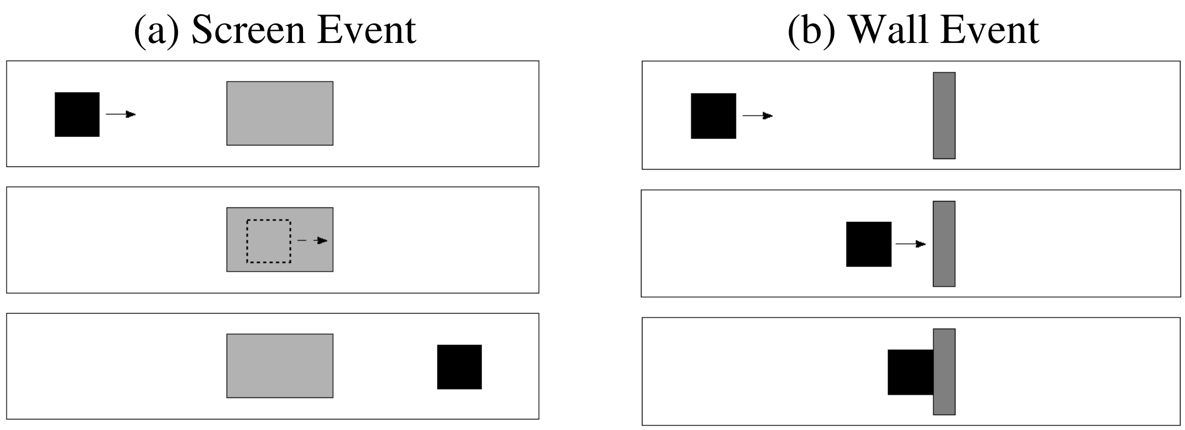

Figure 5

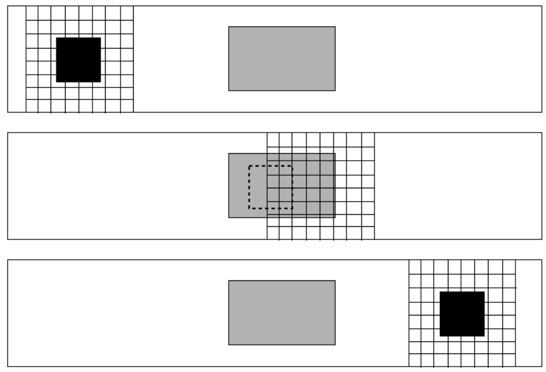

Figure 6

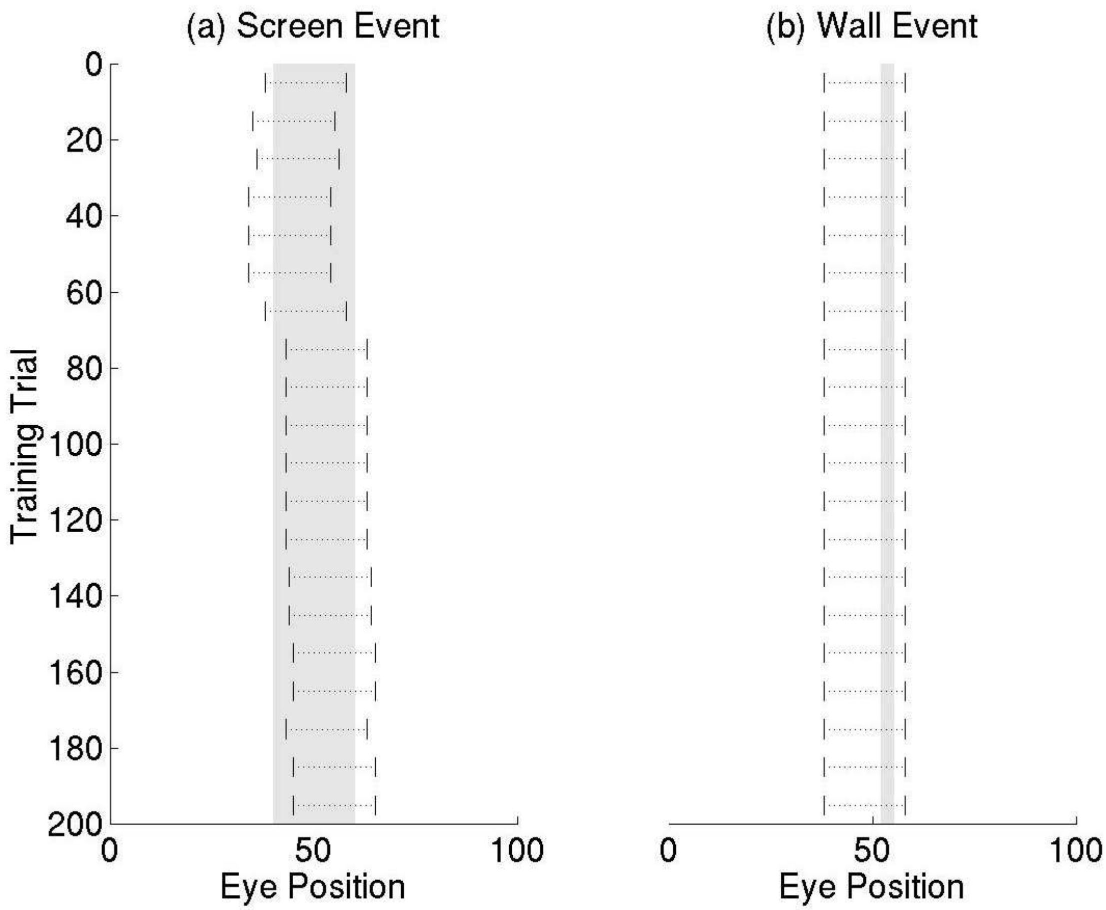
(a) Screen Event
(b) Wall Event
Training Trial
0
20
40
60
80
100
120
140
160
180
200
0
50
100
Eye Position
0
50
100
Eye Position

Figure 7

Wall-Screen Event
(Study 1: Tall Wall)

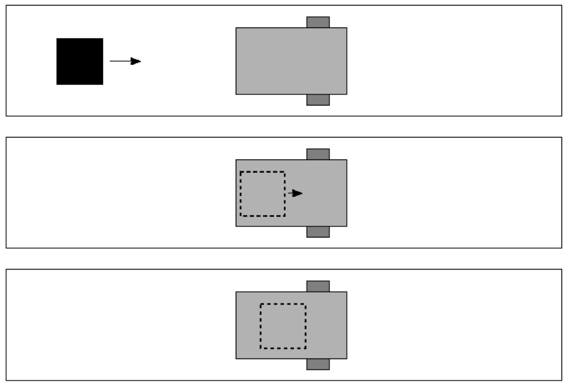

Figure 8

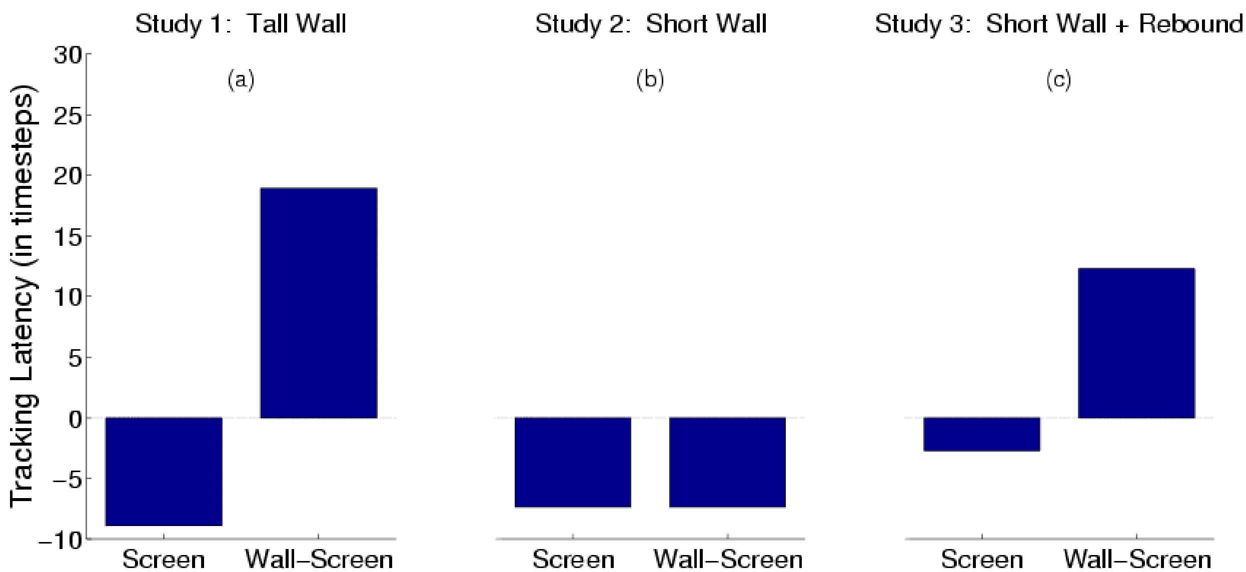
Study 1: Tall Wall
Study 2: Short Wall
Study 3: Short Wall + Rebound
(a)
(b)
(c)
Tracking Latency (in timesteps)
30
25
20
15
10
5
0
−5
−10
Screen
Wall−Screen
Screen
Wall−Screen
Screen
Wall−Screen

Figure 9

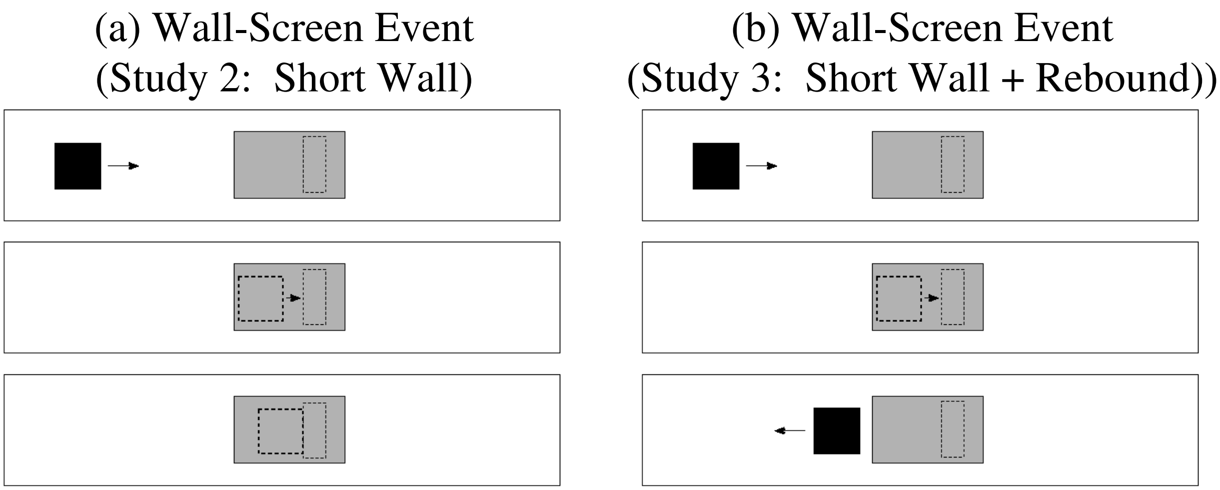
(a) Wall-Screen Event
(Study 2: Short Wall)
(b) Wall-Screen Event
(Study 3: Short Wall + Rebound))