Submitted to *Epigenetic Robotics*, 2002

A lesson from robotics: Modeling infants as autonomous agents

Matthew Schlesinger
Department of Psychology
Southern Illinois University
Carbondale, IL 62901 USA
matthews@siu.edu

Abstract

While computational models are playing an increasingly important role in developmental psychology, at least one lesson from robotics is still being learned: modeling epigenetic processes often requires simulating an embodied, autonomous organism. This paper first contrasts prevailing models of infant cognition with an agent-based approach. A series of infant studies by Baillargeon (1986; Baillargeon & DeVos, 1991) is described, and an eye-movement model is then used to simulate infants' visual activity in this study. I conclude by describing three behavioral predictions of the eye-movement model, and discussing the implications of this work for infant cognition research.

1. Introduction

During the last decade, researchers within robotics and developmental psychology have identified a number of common goals. Parallel work in the two fields has benefited both disciplines. For example, many robotics researchers have begun to move away from heavily pre-designed or hand-built systems, advocating instead naïve agents that acquire adaptive behaviors by interacting with their environment (e.g., “developmental engineering” in Metta, Sandini, and Konczak, 1999). This approach assumes an epigenetic view of development, in which both the organism and the environment play a critical role.

Developmental psychologists, meanwhile, have begun to recognize the value of computational models for investigating developmental processes, and in particular, infant cognitive development, (e.g., Mareschal & French, 2000; Mareschal, Plunkett, & Harris, 1999; Munakata, McClelland, Johnson, & Siegler, 1997; Simon, 1998; Thelen, Schöner, Scheier, & Smith, 2001). A common theme across much of this work is the description of adaptive behavior in infants by means of a compact set of computational principles (e.g., learning by prediction of future states, knowledge as graded representations, etc.)

Despite the fact that these models illustrate an impressive range of theoretical perspectives, modeling architectures, and learning algorithms, many overlook a central element of robotics research: the notion of an embodied, autonomous agent that interacts with a real or virtual environment (Schlesinger, 2001; Schlesinger & Parisi, 2001).

In this paper, I argue that developmental psychologists still have much to learn from work in robotics. In particular, I propose that by modeling the infant not just as a computational system, but more generally as an agent—that perceives its world via sensors and changes its world via effectors—we are able to investigate development as an epigenetic process. And perhaps more importantly, a variety of new insights on how young infants learn may be revealed.

In the next section, I contrast conventional modeling approaches with an emerging perspective often described as an *agent-based approach*. In Section 3, I highlight a series of infant studies conducted by Baillargeon (1986; Baillargeon & DeVos, 1991) to illustrate a critical debate concerning early infant knowledge. Section 4 introduces an eye-movement model, inspired by the agent-based approach, which I have developed to address the debate. Section 5 presents two simulations of Baillargeon's study with the model. In section 6, I conclude by presenting some of the novel behavioral predictions generated by the eye-movement model, and discuss the implications of the model for infant cognition research.

2. The importance of autonomy

Conventional models of infant cognition tend to focus on the development of internal information processing systems (e.g., recognition or categorization of visual stimuli). As a result, many models do not explicitly simulate either a sensory system that receives sensory data (e.g., a visual array), or a motor system that performs overt behaviors (e.g., a reaching movement, a gaze shift).

For example, Munakata et al. (1997) propose a multi-layer recurrent network for simulating an infant that tracks moving objects. On the input side, a visual display is preprocessed and parsed into discrete objects. Similarly, instead of producing motor behaviors, the output of the model is a prediction of the sensory input expected during the following timestep.

In contrast, physical robots are by definition embodied, and “inhabit” a real environment. In a similar manner, robotic simulations capture quasi-realistic features of the physical world (e.g., dynamic features such as gravity and inertia, perceptual features such as visual perspective, etc.). In general, robots are not buffered from their environment, but instead interface or make contact with it in at least two ways, first through sensory systems, and second through effector systems.

Another important feature of robotics is that because autonomous robots both sense *and* act on their environment, they are "free" to select their own sensory inputs (Nolfi & Parisi, 1993). As I have illustrated elsewhere (Schlesinger, Parisi, & Langer, 2000), an important consequence of "self-selection" of sensory inputs is that autonomous agents explore computational search spaces in a highly efficient manner. These learning trajectories often reproduce important patterns of development found in human infants.

Therefore, at least one reason to simulate cognitive development in infants with an agent-based approach is that the notion of an autonomous agent represents the infant as an active organism that learns by interacting with its world.

There are, of course, a number of additional advantages for adopting an agent-based perspective. In the next section, I briefly describe a major debate in the field of infant cognition that has reached an impasse. I suggest that this debate can be addressed by implementing an agent-based model of infants' visual tracking, which simulates infants' moment-to-moment visual activity. The model not only provides several new ways to measure infants' visual expectations, but also offers a novel perspective on cognitive development in young infants.

3. The "car study"

Baillargeon (1986; Baillargeon & DeVos, 1991) presented young infants with a simple mechanical display, in which a car rolls down a ramp, behind a screen, and out the other side. Figure 1A presents a schematic display of this *Habituation* event, so named because infants watch this event repeat several times until they gradually lose interest in it. Note that at the start of the Habituation event, the screen is raised to show the infant that nothing is behind it.

Once habituated, infants then see two test events in alternation (see Figures 1B and 1C). During both the *Possible* and *Impossible* test events, a box is revealed behind the screen. During the Impossible event, however, the box is placed on the track, in the path of the car. Nevertheless, during both test events the car reappears after passing behind the screen.

Baillargeon found that by at least age 6 months, and perhaps even earlier, infants look significantly longer at the Impossible event than the Possible event. How did she interpret these findings? First, she suggested that infants mentally represent both the occluded box and the car as it passes behind the screen. Second, she proposed that infants use these representations to "compute" when the car should reappear, and are consequently surprised to see the car reappear during the Impossible event even though its path is obstructed by the box. Thus, because the Impossible event is surprising or anomalous to infants, they spend more time looking at it.

3.1. The "competent infant" debate

Experiments such as Baillargeon's car study have sparked a broad debate among infant cognition researchers. Some researchers agree with Baillargeon's conclusions, arguing that developmental psychologists have tended to underestimate the infant's ability to represent the physical world, as well as their capacity to reason or think systematically about events in the world (Baillargeon, 1999; Spelke, 1998).

This *representational account* has been challenged by a group of theorists who advocate a *perceptual-processing* account, arguing instead that conclusions about infants' knowledge of the physical world should not be based solely on the amount of time an infant spends looking at possible or impossible displays (Haith, 1998; Smith, 1999). These researchers propose that other measures of infants' visual activity, and particular, of their expectations during possible and impossible events,

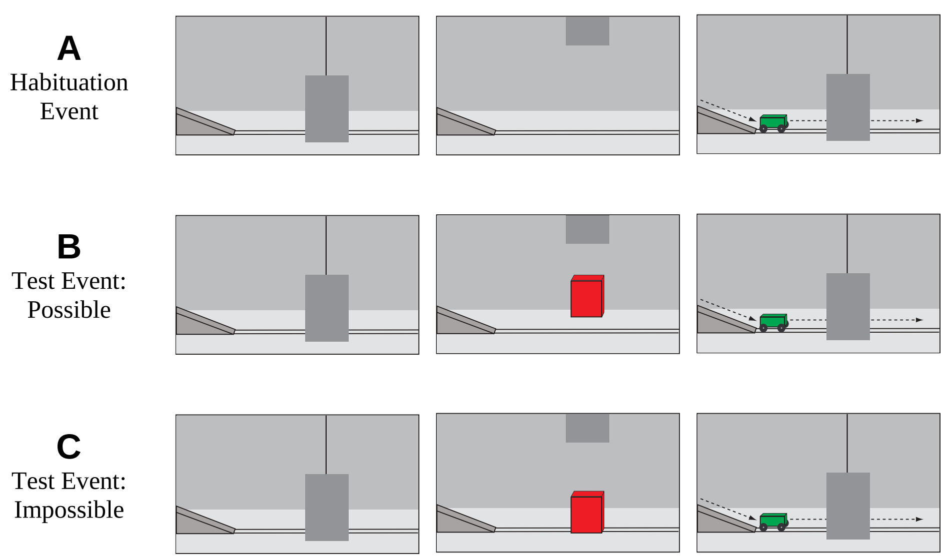

Figure 1: Schematic display of the Habituation (A), Possible (B), and Impossible (C) events studied by Baillargeon (1986; Baillargeon & DeVos, 1991).

should be studied in order to corroborate standard looking-time measures.

3.2. Modeling infants' eye movements

In order to address this debate, I have developed an oculomotor control model that simulates the tracking behavior of an infant (Schlesinger & Barto, 1999; Schlesinger & Parisi, 2001). Like human infants, the model watches simple mechanical displays and learns to track salient moving objects.

It should be noted that the eye-movement model employs a bottom-up approach, consistent with the perceptual-processing account of infant cognition. Accordingly, the model has: (1) no prior knowledge of the physical world (i.e., no internal model), (2) no explicit (e.g., declarative) memory systems, and (3) no built-in capacity for prediction. Nevertheless, the model quickly learns to track moving objects, and like human infants, also learns to correctly anticipate the future location of objects that are temporarily occluded.

Because the eye-movement model simulates visual behavior on a variety of levels (e.g., eye-movements, gaze-shifts, scanpaths, etc.), it is an ideal tool for developing novel measures of infants' visual activity that complement the conventional looking-time methods. Consequently, the primary goal of the model is to present it with a series of events like those in Baillargeon's car study, and to use the behavior of the model to suggest new ways to study infants' expectations in comparable situations.

4. The eye-movement model

I present here a brief description of the stimuli used to train and test the eye-movement model, as well as the structure of the model itself. For additional details on a previous version of the model, the interested reader may refer to Schlesinger and Barto (1999) and Schlesinger and Parisi (2001).

4.1. Training & test displays

The training and testing of the model is designed to mimic the experiences of an infant in Baillargeon's car study. Consequently, three computer-animation events were constructed as analogs to the Habituation, Possible, and Impossible events. However, note that because the model is explicitly trained rather than habituated (see Section 4.4, below), the Habituation event is renamed as the Training event in the model.

Each event is rendered in grayscale, with a duration of 82 frames. Figure 2 presents selected frames from each of the events, corresponding to the respective events in Figure 1 (frame number is noted in the upper right corner). Although the animations simplify many aspects of the real events (e.g., they are 2D rather than 3D), they were designed to capture the most relevant perceptual features of the car study (e.g., occlusion of the "car" behind the screen; relative salience of the car, screen, and box, etc.).

During all three events, the screen moves up then down. Next, the car (i.e., the black square) appears on the left of the display, and passes behind the screen and out the other side. During the Training event, there is nothing behind the screen; during the Possible and Impossible events, the box (i.e., the small, gray rectangle) is revealed as the screen moves up. The box is above the path of the car during the Possible event, while it is within the path of the car during the Impossible event.

4.2. Model architecture

The oculomotor control system is composed of a 3-layer feedforward neural network. The input layer is divided into three sensory channels: a low-resolution, peripheral visual system (33 units), a high-resolution fovea (144 units), and an eye-position system (2 units). The input layer is fully connected to the hidden layer (20 units), which is in turn fully connected to the output layer (10 units).

Each of the animation events is "projected" onto the retina. While the position of the peripheral system is fixed, it spans the entire event display. The fovea, meanwhile, fixates no more than 12% of the display at a time, and can be moved from one part of the display to

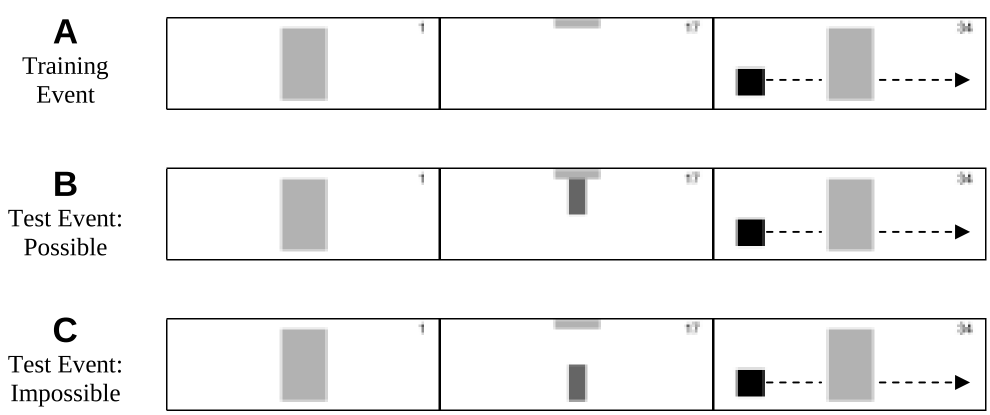

Figure 2: Schematic display of selected frames from the animation events used in Study 1 to train (A) and test (B-C) the eye-movement model (frame number displayed in upper right corner).

another.

The output system is composed of 2 banks of 5 units; each bank controls movement of the fovea in either the vertical or horizontal direction, respectively. Motor signals from the 2 banks are superimposed, producing a net movement in any of 8 directions. Within a bank of output units, four of the units encode either a small (i.e., smooth pursuit) or large (i.e., saccade) movement, in either a positive or negative direction. The fifth unit in each bank produces no movement in the respective direction.

During training and testing, the network is presented with an appropriate animation event, one frame at a time. On each timestep, a single animation frame is projected onto the retina (i.e., periphery and fovea), and activation values are propagated forward. The movement of the fovea is computed by selecting the output unit within each bank with the highest activation (i.e., "winner takes all" selection rule), and updating the fovea's position according to the movement encoded by the two winning units. After the fovea's position (i.e., the fixation point) is updated, the next animation frame is presented.

4.3. Learning algorithm

Two key assumptions of the eye-movement model are: (1) the car in Baillargeon's study is the most salient object, and (2) that infants learn to track the movement of the car. Accordingly, the model employs a reinforcement-learning algorithm, in which the network is rewarded for each timestep that it succeeds in fixating the car.

Specifically, the network receives a scalar reward between 0 and 1 on each timestep, for the proportion of the car that is visible within the fovea. (Note that no reward is possible before the car appears, and while it is occluded behind the screen) Standard temporal-difference learning was employed, including Q-learning at the output layer, followed by back-propagation of prediction errors to the hidden layer (see Sutton & Barto, 1998).

In less formal terms, each output unit encodes a specific eye movement. The activation of each unit is an estimate of the reward expected to follow by producing that unit's particular movement. Thus, a *greedy action-selection* rule is employed, in which the unit within each bank that estimates the highest reward is chosen to produce an eye movement. Exploration of non-optimal movements is achieved by selecting a random eye movement 1% of the time (i.e., ε-greedy action selection, with $\varepsilon = 0.99$).

4.4. Simulation overview

In contrast to infants in Baillargeon's car study, the model is *trained rather than habituated* during the Habituation event. Thus, the first event experienced by the model is called the Training event.

Note that optimal tracking of the car generates a reward of 40 points. In order to avoid overtraining the model, which may lead to highly stereotyped tracking strategies, training only continues until average performance is at least 75% optimal (i.e., average reward is 30 or more points). This training criterion is also in line with the assumption that infants have several goals during the car study, including tracking the car, and therefore they may not track the car optimally.

After the training criteria is reached, learning is turned off (i.e., connection weights are frozen; no exploratory actions are selected), and the Possible and Impossible events are presented to the model. In the following studies, the results of each simulation represent the average performance over a population of 50 networks that are initialized randomly, trained, and then tested.

5. Simulation studies

Two simulation studies are described here. In both studies, the model first learns to track the car during the Training event. After training, the Possible and Impossible test events are presented.

5.1. Study 1: On vs. behind the track

Study 1 simulates the events presented in Figure 1. In this condition, infants see the box placed either behind the track (Possible event) or on the track (Impossible event) during the test phase. In the animation events, these relative positions are translated into *above* (Possible) or *within* the path of the car (Impossible event, see Figure 2).

5.1.1. Results, Study 1

Recall that 50 networks were trained and tested, and that the training criteria was at least 75% optimal tracking (i.e., a total reward of 30 points out of 40 per trial). On average, 145 training trials were required per network to reach criteria.

After training, connection weights were frozen and the exploration parameter was set to 0 (i.e., only optimal eye-movements were chosen). In order to establish a performance benchmark, the model was first re-presented with the Training event, now referred to as the Control event since no learning occurred during this phase. The Possible and Impossible test events were presented next.

Tracking performance was defined as the sum of rewards obtained over the entire event duration (i.e., 82 frames). Figure 3 presents the average total reward as a function of event type (error bars plot the standard error of the mean). Average total reward during the Control event was 32.20 points, while it was 24.69 and 18.01 for the Possible and Impossible test events, respectively.

Tracking was significantly lower during the Possible and Impossible events than during the Control event. In particular, tracking in the Possible event was significantly lower than the Control event ($t(49) = 11.65$, $p < .001$), and tracking in the Impossible event was significantly lower than the Possible event ($t(49) = 4.51$, $p < .001$).

5.1.2. Discussion, Study 1

The eye-movement model is more successful at tracking the car during the Training event, than during either of the test events. These results suggest the conclusion that it is the appearance of the box, during the Possible and Impossible events, that specifically disrupts tracking.

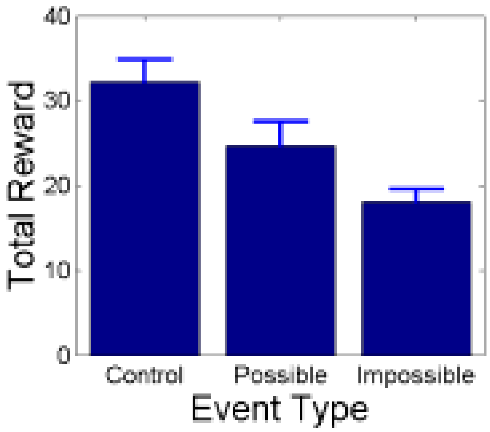

Figure 3: Tracking performance (average total reward) in Study 1 during the Control, Possible, and Impossible events (error bars plot the standard error of the mean).

This conclusion is supported by an inspection of the model's tracking behavior during the Control event.

Figure 4 presents a typical set of scanpaths produced by the model, during the Control, Possible, and Impossible events (the green "x" indicates the center of the fovea, while the trailing dots indicate recent fixations). When no box is present, the model generates at least 2 anticipatory behaviors, including: (1) movement of the fovea toward the left side of the display at the start of the event, *before the car appears* ("Control event", Frame 17), and (2) an anticipatory saccade from the left to the right of the screen *while the car is occluded* ("Control event", Frame 55).

As Figure 4 illustrates, the first behavior, anticipation of the car before it appears, is disrupted during both the Possible and Impossible events. In addition, the second behavior, anticipatory tracking of the car while it is occluded, is also disrupted during the Impossible event. This helps explain why tracking performance is lower in the Impossible than the Possible event.

Why does the box's appearance behind the screen interfere with tracking, and more importantly, why is the disruption greater during the Impossible event? There are two likely explanations.

First, it may be that the model "confuses" the box with the car. Since the box appears earlier during the Impossible event (and for a longer duration, see Figures 2B-C), it may have a greater disruptive effect on the model's tracking behavior. Alternatively, it may not be the timing of the box's appearance, but its position relative to the car's path that is important. According to this second explanation, it is because the box appears in the car's path, where it is has historically been rewarded for looking, that tracking is disrupted during the Impossible event.

Note that the data from Study 1 do not allow us to distinguish between these two accounts. In particular, both accounts predict a greater disruption of tracking during the Impossible event. However, by shifting the car's trajectory to the upper half of the display, the two effects can be teased apart. In this case, the box appears sooner and for more time during the Possible event, but it appears within the car's path during the Impossible event.

Indeed, this condition parallels a similar condition studied by Baillargeon, in which the box appears either on (Impossible) or in front of the track (Possible). As before, Baillargeon (1986; Baillargeon & DeVos, 1991) found that infants looked significantly longer at the Impossible event. Study 2 investigates a comparable simulation condition.

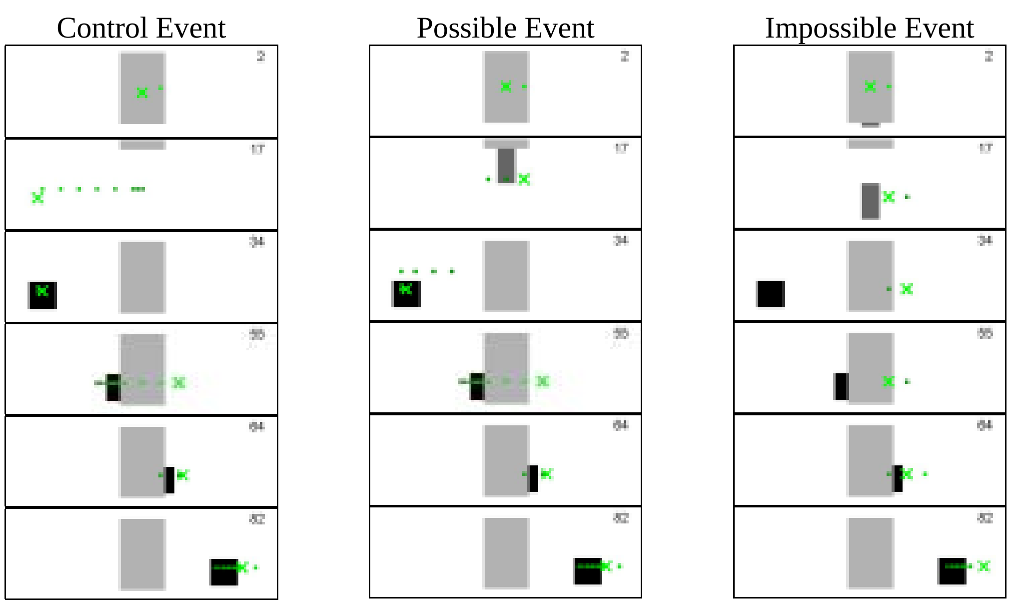

Figure 4: Scanpaths produced by a typical network during the test phase of Study 1. The green "x" indicates the center of the fovea, while the trailing dots indicate recent fixations.

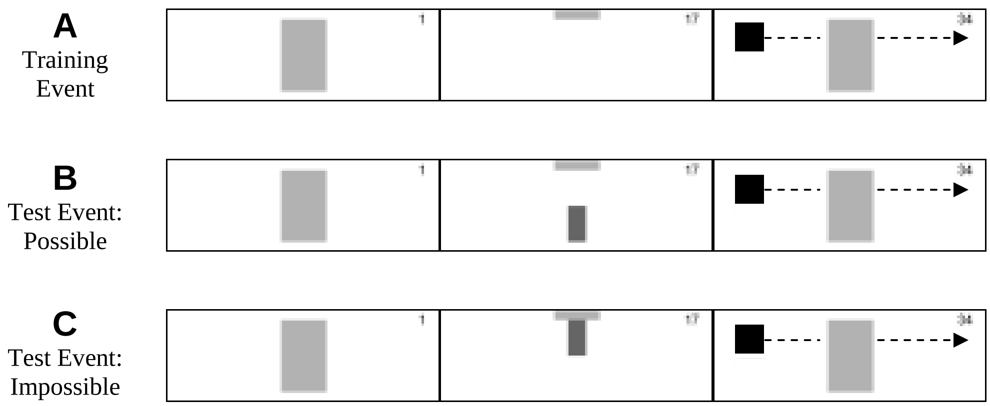

Figure 5: Schematic display of selected frames from the animation events used in Study 2 to train (A) and test (B-C) the eye-movement model (frame number displayed in upper right corner). Note that in contrast to Study 1, the "car" moves along the upper half of the display.

5.2. *Study 2: On vs. in front of the track*

Figure 5 presents selected frames from the animation used to test and train the eye-movement model in Study 2. In contrast to Study 1, the "car" moves along the upper half of the display in Study 2. Thus, in the Possible event the box is revealed sooner (and for more time), while during the Impossible event the box is located in the car's trajectory. Therefore, if tracking performance is lowest during the Possible event, it is the timing of the box's appearance, and not its location, that affects tracking. Alternatively, if tracking is lowest during the Impossible event, than it is the location of the box relative to the car's path that is critical.

Except for a minor change in the trajectory of the car, note that the method of Studies 1 and 2 is virtually identical. As before, 50 replications of the model were trained and tested.

5.2.1. *Results, Study 2*

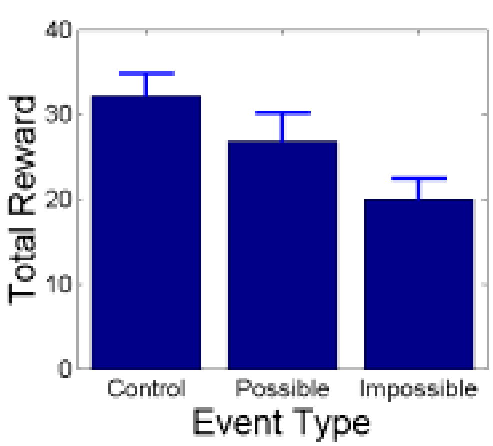

Figure 6: Tracking performance (average total reward) in Study 2 during the Control, Possible, and Impossible events (error bars plot the standard error of the mean).

Comparable to Study 1, an average of 176 training trials were necessary to reach criterion. Tracking performance during the test phase was also comparable to Study 1. Specifically, average total reward was 32.19, 26.83, and 20.04 during the Control, Possible, and Impossible events (see Figure 6). Paired comparisons of the three events resulted in the same qualitative pattern of results as obtained in Study 1. Thus, while tracking was significantly lower during both of the test events than the Control event, it was also significantly lower during the Impossible than the Possible event.

5.2.2. *Discussion, Study 2*

Study 2 replicates the findings of Study 1 in two key ways. First, as before, the appearance of the box during the test phase disrupts the model's ability to track the car. Second, this disruption is greater during the Impossible event. In addition, the results are also consistent with the conclusion that the timing of the box's appearance does not have a critical effect on tracking the car, while the position of the car—relative to the car's trajectory—does significantly affect tracking.

6. Conclusions

Taken together, the findings from the two simulation studies inform the debate on early infant cognition in three important ways.

First, why do infants look longer at impossible events? Baillargeon proposes that when infants are surprised or puzzled by an impossible event, they pay more attention to it. Notice that this representational account presupposes not only the ability to mentally represent the physical world, but also prior knowledge of the physical world that allows infants to reason about occluded events.

In contrast, simulation results from the car study suggest an alternative, more parsimonious account: when the box appears in the car's trajectory (i.e., the Impossible event), infants' tracking is disrupted, and thus they pay more attention to the Impossible event as they search for the car to continue tracking it. I discuss below the implications of this kind of account for infant cognition research.

Before we accept this alternative, perceptual-processing account, it must be empirically verified. How can it be tested? Answering this question suggests a second major consequence of the eye-movement model: because the model produces overt behaviors (i.e., eye movements) in a quasi-realistic world, we can draw an analogy between qualitative behavior patterns in the model, and those produced by human infants in the car study. Therefore, the model suggests at least 3 specific qualitative predictions:

(1) Infants should scan the Possible and Impossible events in different ways (see Figure 4).

(2) Infants should be more successful at tracking the car during the Possible event (see Figures 3 and 6).

(3) Infants' anticipatory eye-movements should be disrupted during the Impossible event.

Note that these predictions are valuable for a number of reasons. First, they provide a direct test of the perceptual-processing account. Second, they can be measured in parallel with infants' global looking time during possible and impossible events, and so offer the means to integrate multiple measures of infants' visual activity across different spatiotemporal scales (e.g., fixations, gaze shifts, scanpaths, etc.).

Most importantly, the predictions generated by the eye-movement model are novel behavioral measures that have not been investigated by infant cognition researchers in looking-time studies such as Baillargeon's. By forcing the representational and perceptual-processing accounts to specify the details of infants' visual behavior at increasingly finer levels, we diminish the likelihood that both accounts will generate a similar pattern of predictions.

Finally, what if the eye-movement model's predictions are confirmed? What are the implications of the model for infant cognition research?

As I noted at the outset, the eye-movement model is motivated by the perceptual-processing account of infant cognition. Recall that the model has no prior knowledge of the physical world, and lacks an explicit memory or prediction system. Therefore, the model suggests the minimal perceptual and cognitive mechanisms necessary for explaining how infants learn to track the car in the car study, and consequently, respond differentially to the Possible and Impossible events.

To conclude, at least one implication of the eye-movement model, then, is that before researchers assume that top-down, knowledge-based, or representational accounts explain infants' visual activity, they should systematically investigate and eliminate bottom-up or perceptual-processing explanations.

Acknowledgements

This work was supported by a grant from the National Institute of Child Health and Human Development (1R03 HD40780-02).

References

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

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

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

Mareschal, D. & French, R.M. (2000). Mechanisms of categorization in infancy. *Infancy*, *1*, 59-76.

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

Metta, G., Sandini, G., & Konczak, J. (1999). A developmental approach to visually-guided reaching in artificial systems. *Neural Networks*, *12*, 1413-1417.

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.

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.

Schlesinger, M. (2001) Building a better baby: Embodied models of infant cognition. *Trends in Cognitive Sciences*, *5*, 139.

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.

Schlesinger, M, & Parisi, D. (2001). The agent-based approach: A new direction for computational models of development. *Developmental Review*, *21*, 121-146.

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

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

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

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

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

Thelen, E., Schöner, G., Scheier, C. & Smith, L., B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 1-34.