

Optimal Control Methods for Simulating the Perception of Causality in Young Infants

Matthew Schlesinger (matthew@cs.umass.edu)

Department of Computer Science; University of Massachusetts
Amherst, MA 01003 USA

Andrew Barto (barto@cs.umass.edu)

Department of Computer Science; University of Massachusetts
Amherst, MA 01003 USA

Abstract

There is a growing debate among developmental theorists concerning the perception of causality in young infants. Some theorists advocate a top-down view, e.g., that infants reason about causal events on the basis of intuitive physical principles. Others argue instead for a bottom-up view of infant causal knowledge, in which causal perception emerges from a simple set of associative learning rules. In order to test the limits of the bottom-up view, we propose an optimal control model (OCM) of infant causal perception. OCM is trained to find an optimal pattern of eye movements for maintaining sight of a target object. We first present a series of simulations which illustrate OCM's ability to anticipate the outcome of novel, occluded causal events, and then compare OCM's performance with that of 9-month-old infants. The implications for developmental theory and research are discussed.

Introduction

How does the perception of causality develop? Do we perceive cause-and-effect relations at birth, or are months of experience necessary? Developmental researchers have approached these questions by studying infants' perceptual reactions to causal events (e.g., Baillargeon, 1986; Keil, 1979; Leslie, 1982; Oakes & Cohen, 1990). Much of this research depends on the tendency for infants to anticipate the outcomes of causal events, often showing surprise to unexpected outcomes (as inferred by measures of attention).

Consider the pair of causal events presented in Figure 1. The first (1a) is a simple, occluded movement display; by age 6 months, infants will quickly learn to anticipate the block's reappearance (Bower, Broughton, & Moore, 1971; Rutkowska, 1993). The second event (1b), however, is more complex. A wall obstructs the path of the block; note that the wall is partially occluded by the screen, revealing only




Figure 1: Occluded causal events. In (a), the block passes behind the occluding screen and reappears on the opposite side. In (b), a partially-visible wall obstructs the path of the block; after passing behind the screen, the block fails to reappear.

the upper and lower portions of the wall. While both events begin in a similar manner, they end differently, depending on the presence of the wall.

Two broad theoretical views have been proposed to explain infants' reactions to events like those in Figure 1. First, several researchers advocate a top-down view of infant causal knowledge (Baillargeon, 1994; Spelke, 1998). According to this view, infants use naive or intuitive physical principles to predict, reason about, or deduce the outcomes of occluded causal events. Two recent computational models help illustrate how the representations underlying this type of prediction system might develop (Mareschal, Plunkett, & Harris, in press; Munakata, McClelland, Johnson, & Siegler, 1997).

Alternatively, several infant causal perception studies have drawn attention to the role of simple perceptual preferences and associative learning rules (Bogartz & Shinskey, 1998; Rivera, Wakeley, & Langer, in press; Schilling, 1997). These researchers argue for a bottom-up view of causal perception. According to this approach, prediction is not an *a priori* goal, nor is representation of hidden objects necessary for the perception of causality in occluded events.

It is theoretically possible, if not likely, that both top-down and bottom-up factors play a role in infants' causal perception. How should the two views be reconciled? The strategy that we propose is to construct a model based on the bottom-up view, and then to test the extent of its perceptual "abilities" when presented with causal events like those shown to young infants. Any gaps or limitations in the performance of the model could then be addressed, we assume, by using the top-down approach.

Rather than simulating causal perception as a *representational* task (cf., Mareschal et al., in press; Munakata et al., 1997), we model the phenomenon as an optimal control problem. The optimal control model (OCM) is a *sensorimotor* model of infant causal perception. Unlike human infants, OCM: (1) has no intuitive knowledge, (2) cannot generate predictions, and (3) learns only by trial-and-error. OCM's objective is to learn a sequence of eye movements that best maintain a target object in view. After training OCM to track a target, we then test OCM's reactions to novel, occluded causal events like the one presented in Figure 1b. We next briefly describe OCM.

The Optimal Control Model

The Tracking Display

Figure 2a presents a snapshot of the 2-dimensional tracking display used to train OCM. During each trial, the block (rep-




Figure 2: The OCM: (a) the tracking display (the target is represented by the solid black square, while the visual field is indicated by the black frame); (b) OCM's visual input for the corresponding display.

represented by the black square) moves from left to right, at the rate of 1 unit per timestep. At the start of the trial, OCM's visual field (the large, square frame) is positioned with the block in the center of the field. Each trial lasts 50 timesteps.

The display is 100 units wide and 20 units high. OCM's visual field covers a 20-unit square region of the display, permitting only lateral eye movements. The block is an 8-unit square, while the screen (when present) is 20 units wide and 12 units high.

Objects in OCM's visual field activate corresponding input units on its retina (see Figure 2b). Because the block is OCM's target object, its activation on the retina is 1 (i.e., maximum salience). The background has an activation of 0, while the wall and screen's activation levels are 0.6 and 0.2, respectively.

Model Architecture

OCM's sensorimotor "knowledge" is represented by a multi-layer, artificial neural network. There are two input systems. First, OCM receives visual input from its 20-by-20 unit retina. Figure 2b illustrates a typical visual input pattern. OCM also receives an additional input indicating the position of the visual field with respect to the display, normalized from 0 to 1.

There are 20 hidden units, and 5 output (motor) units. The network is fully connected, with only feedforward connections. Each of the motor units controls one of 5 possible eye movements: $\langle -4, -1, 0, 1, 4 \rangle$. On each timestep, the movement corresponding to the most active motor unit is performed.

Learning Algorithm

OCM is rewarded for generating eye movements which keep the block within the visual field; OCM learns by trial-and-error to find a pattern of eye movements which optimize sight of the block (i.e., maximize the total reward). Any movement which is followed by sight of the block is rewarded; the reward ranges from 0 to 1, as a function of the proportion of the block in the visual field after the eye movement (e.g., 1 when fully visible, 0.5 when half visible, etc.).

The output of each motor unit is an estimate of the value (i.e., probability of reward) for performing the corresponding eye movement. We employed the Sarsa learning algorithm, an unsupervised, online version of reinforcement learning methods (see Sutton & Barto, 1997) to train OCM. Using standard gradient descent methods, the Sarsa algorithm attempts to minimize the difference between the estimated and observed rewards after each eye movement.

Consequently, the direction and magnitude of the weight changes for the output layer depend on the eye movement chosen, and the corresponding reward, during a given timestep. These weight changes are then propagated backwards to the hidden layer (see Lin, 1991, for a discussion of reinforcement learning and back-prop hybrid models).

Simulation Overview

We conducted a series of simulation studies which assess OCM's ability to learn to track visible and occluded targets. In each study, OCM was first trained to track a target during two types of events. In the occluded event, the block passed behind a screen and reappeared on the other side. In the other (fully visible) event, the block encountered a wall and then remained in place. After OCM learned to optimally track the block during these events, we then tested OCM's tracking during a novel, occluded causal event which included both the screen and the wall. In Studies 1 and 2, the wall was partially occluded by the screen, while it was completely occluded in Study 3.

Study 1: Tall Wall

Study 1 addresses the question of how OCM will respond to a partially occluded causal event. Figure 3 displays the events used to train and then test OCM.

Method

Training. During training, OCM was presented with two causal events. On Screen trials, a screen occluded the central portion of the display. On Wall trials, an obstacle was positioned in the center of the display; the block remained in place after making contact with the wall.

Screen and Wall trials alternated randomly. Training continued until OCM's total rewards during both Screen and Wall trials were at least 95% optimal over 10 consecutive trials (i.e., maximum total rewards were 30 and 50 points for Screen and Wall trials, respectively). If criterion was not reached by 300 trials, the run was terminated, the data were discarded, and a new set of random initial weights were generated.

Testing. After training, all weights in the network were frozen (i.e., learning was turned off¹). OCM was then presented with 10 Wall-Screen trials. During Wall-Screen trials, the wall was positioned behind the screen; when the block passed behind the screen, its path was obstructed by the wall

¹This was done to prevent OCM's responses during early test trials from contaminating later trials.

Study 1




Figure 3: Training and test events presented to OCM in Study 1. Note that the Screen trial type was identical during training and testing.

(as during Wall trials), and consequently did not reappear. In addition, OCM was also presented with 10 Screen trials, in order to assess OCM's ability to track an occluded, unobstructed object.

Results

Training. Figure 4a presents the average number of trials to criterion in Study 1, averaged across 50 runs (36 additional runs were discarded). OCM reached criterion on Wall and Screen trials after 65.5 and 93.2 training trials, respectively. The difference is statistically significant ($t(98) = 1.78, p < .05$). Like human infants, OCM learns to track a fully visible target before it learns to track an occluded target. However, an average of 148.8 trials were necessary before reaching criterion on *both* trial types concurrently.




Figure 4: Trials to criterion during training in Studies 1, 2, and 3. See text for details.

Testing. Our analyses of the test trials focus on OCM's tracking behavior once the block disappears behind the screen, and how the presence or absence of the wall affects this behavior. In particular, we are interested in whether or not OCM moves its visual field to the right edge of the screen *before* or *after* the block reappears, during Screen trials. Consequently, we define *tracking latency* as the difference in time between OCM's first fixation of the right edge of the screen, and the block's reappearance at the right edge during Screen trials. Although the block does not reappear during Wall-Screen trials, we can use the same temporal index to compute OCM's tracking latency (i.e., assuming reappearance of the block, had it not been obstructed). A positive latency (or delay) means that OCM fixates the right edge of the screen after the block has (or would have) reappeared, while a negative latency means that OCM anticipates the reappearance of the block.

Figure 5 presents OCM's tracking latencies during the test phase of Study 1. During Screen trials, OCM anticipated the block's reappearance, fixating the right edge of the screen 8.9 timesteps sooner than the reappearance of the block ($t(49) = 7.03, p < .01$). In contrast, OCM's average tracking latency was significantly delayed by the presence of the tall wall during Wall-Screen trials; on average, OCM fixated the right edge of the screen 18.9 timesteps after an unobstructed block would have reappeared ($t(49) = 4.21, p < .01$).




Figure 5: Mean tracking latencies in the test phase of Study 1, for Wall-Screen and Screen trials. OCM anticipated the reappearance of the block during Screen trials, while tracking was delayed during Wall-Screen trials.

Discussion

OCM's learning trajectory parallels that of human infants. OCM learns to track a fully visible object before it learns to track the movements of an occluded object. After training, OCM appears to react as if it "knows" when the occluded path of the block will, or will not be obstructed. OCM anticipates the reappearance of the occluded object during Screen trials, but not Wall-Screen trials.

It is tempting to conclude that OCM learns to use the presence of the wall as a cue for tracking the occluded block. However, there is more than one way to explain OCM's behavior. One explanation is that OCM learns nothing about the wall when training on Wall trials; rather, it only learns to hold the visual field in place when the block stops moving. According to this explanation, the presence of the partially visible wall, during Wall-Screen trials, simply disrupts the tracking pattern learned during Screen trials. Alternatively, we might argue that OCM learns to associate the sight of the wall with its effect on the block.

These two explanations can be tested by placing the wall

Study 2




Figure 6: Training and test events presented to OCM in Study 2. Unlike Study 1, a thinner wall was included in the display, representing a wall which has been moved back relative to its position in Study 1. During Wall and Wall-Screen trials, the block passed in front of the wall.

“back” (from the perspective of OCM), beyond the path of the block. Thus, the movement of the block is identical during Screen and Wall-Screen trials. If sight of the wall is used as a cue, OCM should anticipate the block on *both* Screen and Wall-Screen trials; otherwise, if the wall disrupts tracking during Wall-Screen trials, then OCM should only anticipate the block during Screen trials. Study 2 tests these alternative hypotheses.

Study 2: Back Wall

Study 1 was repeated, replacing the tall wall which obstructs the block’s movement with a wall placed “back” (i.e., farther from the observer’s point of view), beyond the path of the block. Because the display is a 2-dimensional projection of a 3-dimensional world, we represented the perceptual effect of moving the wall back by decreasing its width (from 4 to 2 units). Consequently, the block passed in front of the wall during both Wall and Wall-Screen trials (see Figure 6).

Results

Training. Figure 4b presents the mean number of trials to criterion, during training in Study 2, across 50 runs (28 additional runs were discarded). Compared to Study 1, fewer trials were needed to independently reach criterion on Wall and Screen trials (41.0 and 61.3, respectively).

Testing. Figure 7 presents OCM’s mean tracking latencies, during testing, for the Wall-Screen and Wall events. Placing the wall back significantly reduced OCM’s tracking latency during Wall-Screen trials, compared to the tall-wall condition in Study 1 (-3.12 versus 18.9 timesteps; $t(98) = 4.11, p < .01$). However, OCM’s anticipatory tracking was slightly slower on Wall-Screen trials, than during Screen trials (see Figure 7).

A closer analysis revealed that during 6 of the 50 runs, tracking of the block was in fact completely interrupted by the partially visible back wall, during Wall-Screen trials. However, when the remaining 44 runs are analyzed, OCM’s average tracking latencies during Wall-Screen and Screen trials are -10.37 and -10.5 timesteps, respectively. During the majority of the runs in Study 2, therefore, sight of the wall did not disrupt OCM’s anticipatory tracking.

Discussion

Study 2 replicates and extends the findings of Study 1. In both studies, OCM spontaneously learns to anticipate the reappearance of the occluded block. Further, when the wall is




Figure 7: Mean tracking latencies in the test phase of Study 2, for Wall-Screen and Screen trials. OCM anticipated the reappearance of the block during both Screen and Wall-Screen trials.

positioned so as to have no effect on the movement of the block, it does not disrupt OCM’s anticipatory tracking. Taken together, the results of Studies 1 and 2 support the conclusion that OCM learns to use both the screen and the wall as cues for perceptual action.

In contrast to Studies 1 and 2, a number of infant causal perception studies present perceptual cues to infants *prior to*, rather than *during* the occlusion event (e.g., occluded collision events, studied by Baillargeon, 1986; Lucksinger, Cohen, & Madole, 1992). Because the pairs of test events are identical in these studies, infants must *remember and recruit* information made available to them *before* each occluded event is presented.

We can simulate this type of causal event by reducing the height of the wall; when occluded, a short wall is no longer visible. While Studies 1 and 2 presented OCM with *partially* occluded causal events, Study 3 simulates OCM’s reaction to a *completely* occluded causal event.

Study 3: Short Wall

Figure 8 presents a display of the training and test events used in Study 3. Three modifications were made to the method employed in Study 1. First, the height of the wall was reduced from 16 to 10 units. Second, 20 new input units were added to OCM’s neural network. These “context” units were activated via recurrent connections from OCM’s hidden layer, providing a functional memory of past internal states (Elman, 1990).

Third, each trial was preceded by a preview. During the

Study 3




Figure 8: Training and test events presented to OCM in Study 3. Unlike Study 1, a short wall was included in the display, which was fully occluded by the screen during Wall-Screen trials. Note that after the preview, Screen and Wall-Screen trials are perceptually identical; past state information is necessary to differentiate these two trial types.

preview, OCM's visual field was held at the center of the display for 10 timesteps. During Screen and Wall-Screen trials, the screen was not included in the preview (i.e., OCM saw what was "behind" the screen). Learning was turned off during the preview. After the preview, each trial proceeded as in Studies 1 and 2.

Results

Training. Figure 4c presents the mean trials to criterion, during training in Study 3, across 50 runs (21 additional runs were discarded). When compared with Study 1, there were no significant differences in training time after changing the tall wall to the short wall.

Testing. OCM appears to "forget" about the short wall once it is occluded by the screen. As Figure 9 indicates, OCM's tracking latencies during Screen and Wall-Screen trials were identical; regardless of whether or not the short wall was present, OCM anticipated the reappearance of the block by 7.34 timesteps ($t(49) = 4.12, p < .01$).




Figure 9: Mean tracking latencies in the test phase of Study 3, for Wall-Screen and Screen trials. Unlike Study 1, OCM anticipated the reappearance of the block during both Screen and Wall-Screen trials.

Discussion

In contrast to the results of Study 1, OCM's tracking behavior was not affected by the presence of a short wall. The findings from Study 3 suggest that OCM relied on its immediate perceptual input, while ignoring or failing to use its memory of the short wall.

However, it is important to remember that there was no pressure on OCM during training to learn to use memory.

First, during the fully visible wall trials, memory is unnecessary. Second, during the Screen trials, OCM learns to use the sight of the screen (rather than an internal representation of the occluded block) as a perceptual cue for anticipating the block's reappearance. Thus, the task constraints operating during training make the use of memory redundant.

General Discussion

The results from the three sets of simulations highlight both the strengths and limitations of the optimal control model of infant causal perception. There are two major findings. First, OCM quickly learns a set of optimal tracking strategies for following a moving object. Second, when presented with a novel causal event, OCM appropriately anticipates the outcome of partially occluded, but not fully occluded, versions of the event.

We can evaluate the performance of the model by directly comparing the results with data obtained from young infants. Berthier et al. (in preparation) conducted a series of experiments with 9-month-old infants, corresponding to Studies 1 through 3. Figure 10 presents a summary of the test results for OCM, and the comparable average tracking latencies (in seconds) for three groups of 9-month-olds. Across all three




Figure 10: Mean tracking latencies in the test phase of Studies 1-3 for the OCM (top panel) and 9-month-old infants (bottom panel; from Berthier et al., in preparation).

studies, the performance of OCM provides a close qualitative fit to the performance of human infants. Like OCM, 9-month-old infants also use the partially visible wall as a cue, but not the fully occluded wall, to guide their tracking of the occluded target.

When taken together, the human and simulation findings carry a number of implications for developmental theory and research on the perception of causality in young infants. First, many causal perception researchers: (1) assume that infants *explicitly predict* the outcomes of the events they watch, and (2) infer, on the basis of looking-time measures, when infants' predictions are confirmed or violated. While the results from OCM are not intended to provide a reductionist account, they suggest that learning during the habituation or familiarization phase may drive the process of anticipation, helping to shape infants' causal expectations (see Rivera et al., in press; Schilling, 1997).

A second implication concerns the role of internal representations of occluded objects and events. Again, it is often assumed that infants must operate on mental representations, rather than direct perceptions, when the critical objects are occluded or out of sight. However, when tracking an occluded target, OCM relies on *sensorimotor* rather than *representational* strategies for anticipating the target.

For example, the results from Study 1 demonstrate that anticipatory behavior can emerge as a consequence of learning an optimal tracking strategy, without the need for memory or prediction. Indeed, having memory does not seem to facilitate OCM's learning to track the block during Screen training trials (compare Figures 4a and 4c), although there were fewer discarded runs when OCM was trained with a recurrent network (i.e., in Study 3). We suspect that in many causal perception studies, infants employ some combination of sensorimotor and representational strategies. Simulations with models like OCM help to determine if and when the sensorimotor strategies are sufficient to account for the perceptual phenomenon.

This point echoes a question raised in the introduction: what are the performance limits of OCM? On the one hand, there is surprisingly close agreement between the performance of OCM and the recent findings of Berthier et al. Nevertheless, this fit may in part be due to the specific constraints of learning to track, and the way in which this task favors an optimal control solution (e.g., like learning to reach or generate saccades). Thus, two current weaknesses of a bottom-up view in general, and an optimal control approach in particular are: (1) that some tasks may necessarily require predictive, representational strategies, and (2) that OCM may not be able to account for infants' perceptual behavior in other contexts (e.g., preference for "surprising" or impossible events). We are currently exploring an elaborated version of the model which addresses these issues.

Acknowledgments

This work was supported by NSF IRI-9720345 and NSF CDA-9703217.

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. (1994). Physical reasoning in young infants: Seeking explanations for impossible events. *British Journal of Developmental Psychology*, 12, 9-33.
- Berthier, N., Bertenthal, B., Seaks, J., Sylvia, M., Johnson, R., & Clifton, R. (in preparation). Catching a moving object: Infant reaching when trajectory information is removed.
- 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., 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.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Keil, F. (1979). The development of the young child's ability to anticipate the outcomes of simple causal events. *Child Development*, 50, 455-462.
- Leslie, A.M. (1982). The perception of causality in infants. *Perception*, 11, 173-186.
- Lin, L.J. (1991). Self-improvement based on reinforcement learning, planning, and teaching. In L.A. Birnbaum, and G.C. Collins (Eds.), *Machine learning: Proceedings of the eighth international workshop on machine learning*, (pp. 323-327). San Mateo, CA: Morgan Kaufmann.
- Lucksinger, K.L., Cohen, L.B., & Madole, K.L. (1992, May). What infants infer about hidden objects and events. Presented at the *International Conference for Infant Studies*, Miami, FL.
- Mareschal, D., Plunkett, K., & Harris, P. (in press). A computational and neuropsychological account of object-oriented behaviours in infancy. *Developmental Science*.
- 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.
- Oakes, L.M., & Cohen, L.B. (1990). Infant perception of a causal event. *Cognitive Development*, 5, 193-207.
- Rivera, S.M., Wakeley, A., Langer, J. (1999). The drawbridge phenomenon: Representational reasoning or perceptual preference? *Developmental Psychology*, 35, 427-435.
- Rutkowska, J.C. (1993). *The computational infant: Looking for developmental cognitive science*. London: Harvester Wheatsheaf.
- 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.
- 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.