

Multimodal Control of Reaching—Simulating the Role of Tactile Feedback

Matthew Schlesinger and Domenico Parisi

Abstract—By the onset of reaching, young infants are already able to keep track of the position of their hand by using visual feedback from the target and proprioceptive feedback from the arm. How is this multimodal coordination achieved? We propose that infants learn to coordinate vision and proprioception by using tactile feedback from the target. In order to evaluate this hypothesis, we employ an evolutionary-based learning algorithm as a proxy for trial-and-error sensorimotor development in young infants. A series of simulation studies illustrate how touch: 1) helps coordinate vision and proprioception; 2) facilitates an efficient reaching strategy; and 3) promotes intermodal recalibration when the coordination is perturbed. We present two developmental predictions generated by the model and discuss the relative importance of visual and tactile feedback while learning to reach.

Index Terms—Computational model, multimodal coordination, sensorimotor development, trial-and-error learning.

I. INTRODUCTION

HOW DO infants learn to reach? This is a complicated problem, as it not only involves coordinating information about the position of the target and the hand across several sensory modalities, but also using this information to guide head, eye, and arm movements. A traditional theory proposes that infants learn to reach by guiding their hand visually toward the target [5], [6], [20], [23], [27]. A number of recent studies, however, have discredited this theory by illustrating that young infants do not need to see their hands as they learn to reach [1], [7], [8]. Most striking of these experiments is found in [7], which reported that at the onset of reaching (around age 12 weeks), infants are just as likely to reach for a target in a fully lit room as they are to reach for a luminous target in a dark room.

These findings suggest that by the onset of reaching, infants have already learned to coordinate vision of a target with the felt position of their arm (detected via proprioception). How is this coordination achieved? One possibility is that infants possess an innate multimodal map that computes a two-way transformation between the visual and proprioceptive spaces. Support for this hypothesis comes from studies of newborn infants who show a tendency to generate reflex-like arm movements toward nearby objects in the visual field (“prereaches”; see [10], [12],

and [14]). An alternative hypothesis is that infants learn to coordinate vision and proprioception during the development of reaching. According to this second account, learning to reach creates a task-specific context for infants to coordinate multiple information streams.

We propose that visual–proprioceptive coordination in infants is an acquired skill. In particular, we investigate the hypothesis that the position of the hand can be “triangulated” by combining information from the target and the arm with a third information source.

In principle, at least two sensory channels may provide this third source of information. One possibility is that infants use the sight of their hand to help determine its position relative to a target. The advantage of using visual feedback from the hand to coordinate eye and arm movements is that it is available continuously during a reach. However, vision of the hand is potentially ambiguous (e.g., when the hand occludes the target without touching it) and may require additional computation to determine the hand’s position relative to the target. In addition, monitoring of the hand position also requires attentional resources above and beyond those available for keeping track of the target and controlling arm movements. These reasons may help explain why young infants do not appear to use visual feedback from the hand as an initial reaching strategy.

An alternative information source comes from tactile feedback during contact with the target. Unlike vision of the hand, tactile feedback is unambiguous. That is, the hand’s location can be equated with the seen position of the target only when the target is felt. Thus, tactile feedback provides a unique sensory signal for associating seen locations in space (i.e., the target) with felt positions of the arm. Compared to vision of the hand, however, which provides continuous feedback during a reach, tactile feedback is only available after a successful reach.

The principal goals of the present paper are: 1) to demonstrate how vision and proprioception are coordinated during reaching and 2) to contrast vision-based and tactile-based strategies for achieving this coordination. We employ a computational model of sensorimotor development in young infants [24]. The model uses a genetic algorithm as a stochastic optimization procedure to simulate the gradual emergence of adaptive reaching movements. In Section II, we describe the main features of the Schlesinger *et al.* model. In Section III, we present a series of modeling studies that simulate learning to reach under a variety of conditions. In Section IV, we present two key predictions generated by the model and discuss the how the present simulation results shed light on the potential developmental roles of visual and tactile feedback.

Manuscript received July 28, 2000; revised October 27, 2000.

M. Schlesinger was with the Department of Computer Science, University of Massachusetts, Amherst, MA 01003 USA. He is now with the Department of Psychology, Southern Illinois University, Carbondale, IL 62901 USA (e-mail: matthews@siu.edu).

D. Parisi is with the Institute of Psychology, National Research Council, 00137 Rome, Italy (e-mail: parisi@www.ip.rm.cnr.it).

Publisher Item Identifier S 1089-778X(01)03284-2.

II. A MODEL OF HAND-EYE CONTROL IN INFANTS

The Schlesinger *et al.* model is designed to represent two core features of infant sensorimotor development. First, infants learn to reach by trial-and-error, that is, by generating a wide range of exploratory movements, followed by selection of the most adaptive or functional reaching patterns [3], [13], [26]. Second, because infants do not use visual feedback to guide their initial reaches, only coarse information (i.e., the success or failure of a reach) is likely to be available for selecting and improving reaching movements.

In order to incorporate these features, the Schlesinger *et al.* model implements an unsupervised variation-and-selection learning algorithm analogous to evolutionary learning [2], [15]. Specifically, a population of artificial neural networks is used to represent the diverse set of reaching movements available to an infant. Because the networks inhabit a simple ecology (including their bodies), we refer to them as “econets” (see [22]). Like the movements generated by infants, econets adapt under selective pressure to produce the most direct efficient-reaching patterns.

It should be stressed that because the model focuses on the key features of learning to reach (i.e., a multimodal sensory system, a motor system, and a reward, evaluation, or performance feedback system), it simplifies many aspects of sensorimotor development in infants (e.g., a monocular visual system, a single arm). In particular, our use of an evolutionary-based learning algorithm is not intended to suggest that homologous learning mechanisms are necessarily involved in infant sensorimotor development (however, see [25] for an alternative neodarwinist account of learning to reach). Rather, the genetic algorithm is employed as a proxy for the comparable kinds of trial-and-error sensorimotor learning in which infants engage.

Despite the necessary simplifications of our model, as well as the major differences between it and human infants, its explanatory and predictive value can be evaluated according to three criteria. First, to what extent does the model reproduce significant features of sensorimotor development in infants? Second, is the model able to explain these phenomena in a plausible and empirically verifiable way? Finally, does the model suggest any novel developmental hypotheses that may help us choose between alternative theories or explanations?

With respect to the first question, the Schlesinger *et al.* model reproduces several key aspects of reaching development in infants (see [24] for a complete description) including: 1) limited use of redundant joints [4]; 2) a predominance of muscle cocontractions during early reaching [4], [26], [28]; 3) the early emergence of stereotyped reaching movements [3], [17]; and 4) a temporary decline in prereaches [13]. In addition, the model suggests potential explanations for several characteristics of reaching development in infants (e.g., the U-shaped developmental pattern of prereaching movements; see [13] for a longitudinal study of prereaching development). Finally and most importantly, the model also generates a number of developmental predictions (two of which are described in the present paper).

In order to focus on the question of multimodal coordination, we now present a brief overview of the Schlesinger *et al.* model. Interested readers may refer to [24] for a full description, including the details of the learning algorithm and a complete list of all major parameter values.




Fig. 1. Econet in its 2-D workspace. (Left and right edges of the visual field are shown, as well as the rotational limits of the shoulder and elbow joints.)




Fig. 2. Diagram of the four-layer feedforward network (not all connections shown). Visual and proprioceptive signals first pass through an intramodal hidden layer and then converge with the tactile input signal at the second (intermodal) hidden layer.

A. Reaching Workspace

Fig. 1 presents a schematic diagram of the econet in its two-dimensional (2-D) reaching workspace. The econet remains at the center of the 100×100 unit workspace. The target is a small object (triangle) placed randomly in the workspace at the start of each trial.

B. Econet

The econet has a monocular visual system with a 64° visual field (see Fig. 1). The econet also has a two-segment arm that spans the workspace. The trunk, head, and eye are positioned along the same axis and so move as a single unit with the rotation of the eye. As Fig. 1 illustrates, the eye can rotate freely in either direction while the arm’s movement is limited to 180° rotation at the shoulder and elbow joints. A multilayer feedforward neural network (see Fig. 2) uses visual input from the eye, proprioceptive input from the arm, and tactile input from the hand to control the movements of the eye and arm. Note that the visual and proprioceptive sensory inputs first pass through a unimodal hidden layer before combining at the second intermodal hidden layer. This partially modularized architecture helps minimize crosstalk or interference effects between sensory modalities and results in improved performance [16], [24].

The econet's 64° visual field is divided evenly into eight "windows," each 8° wide; each window corresponds to one of the eight visual input units. Whenever the target or hand enters a window, the corresponding visual input unit is activated. The activation value of the visual input unit ranges from zero to one, relative to the distance of the object from the econet (e.g., very close objects have activations near one, while distant objects have activations near zero). Using this input scheme, both the hand and target are represented by the visual input system in a common computational form. The proprioceptive input units encode the amount of stretch in the arm muscles, represented by a linear mapping from the shoulder and elbow joint angles. Finally, the tactile input unit encodes contact of the hand with the target as a binary signal. Movements of the econet's eye and arm are produced by activation on its output units. A single unit moves the eye either left or right, while the remaining four units activate the flexor-extensor muscle pairs of the upper arm and forearm.^{1,2}

C. Learning Algorithm

We employ an evolutionary-based learning approach as a proxy for trial-and-error sensorimotor development in infants. A population of econets is generated and evaluated on the reaching task. Each econet corresponds to a class of motor action patterns (i.e., a reaching "style" or a visuomotor coordination strategy), while the population represents the collection of possible movement patterns from which the most adaptive reaching movements are selected. Subsequent generations are produced by selecting and copying the most successful econets; behavioral variation is produced by mutation (i.e., copy errors) during the reproduction process.

The initial population consists of 100 individuals with random connection weights. Each connection weight is initially encoded as an 8-bit binary string (e.g., 10100111), and then transformed to a real-valued number between +10 and -10. Each generation of econets is evaluated by presenting each of the econets with 20 trials to reach the target. At the start of each trial, the econet's body (i.e., eye and arm) is oriented randomly and the target is positioned randomly in the workspace. Trials last for 100 timesteps. Fitness is incremented by one point for each timestep spent with the hand touching the target.

The choice of the fitness (or reward) function was guided by two criteria. Specifically, we wanted to impose a bias toward a minimum-time solution while also maximizing time spent in contact with the target after reaching it [3], [11] (see [9] for a review of minimum principles in motor control). These two constraints were combined in order to encourage reaching patterns with both acceleration and deceleration components, which are consistent with reaching and grasping an object rather than purely accelerating reaches, which are consistent with hitting or batting an object [3]. The second constraint (reward for time spent in contact with the target) was also included to simulate

¹The Schlesinger *et al.* implementation is a kinematic model of reaching (i.e., it tracks the position and velocity of the eye and arm) because: 1) the muscles are represented by a linear force function and 2) no other forces (e.g., gravity, inertia, coriolis forces, etc.) are modeled explicitly.

²Note that the movements of an individual econet are completely deterministic. Random exploration is achieved transgenerationally through the mutation of successful econets.

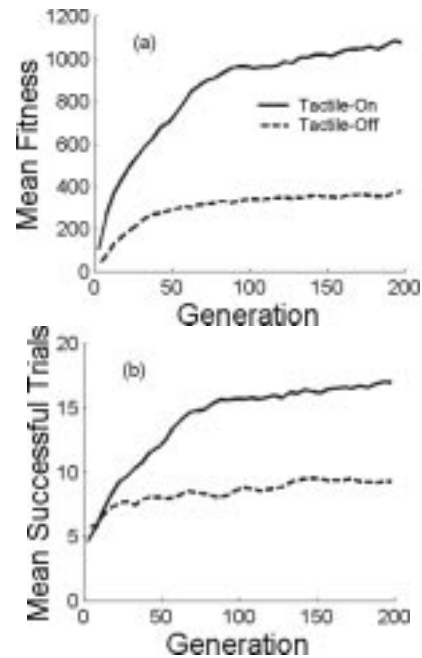


Fig. 3. Average fitness. (a) Time spent in contact with the target and (b) number of successful trials, i.e., number of trials out of 20 in which the target was successfully reached in two populations of econets ($N = 10$ runs for each population). "Tactile-on" econets receive tactile feedback when they touch the target, while "tactile-off" econets receive no tactile feedback.

the infants' goal of not only reaching target objects, but also exploring and manipulating these objects after reaching them.

Exploration is accomplished by first selecting the 20 econets in each generation with the highest fitness (i.e., time spent in contact with the target). Each parent then produces five offspring. During the reproduction process, variation is introduced by mutating the offspring's connection weights. Specifically, each bit (in the 8-bit string representing each weight) has a 2% chance of being switched during reproduction. No other genetic operators (e.g., crossover) were employed. The sequence of testing, selecting, and reproducing econets continues for 200 generations. This stopping criteria was employed based on the observation that the average rate of change in fitness approached zero after approximately 100–150 generations.

III. SIMULATING VISUAL-PROPRIOCEPTIVE COORDINATION

We now present a series of simulations that focus on the development of visual-proprioceptive coordination while learning to reach. The first set of simulations addresses the question of how tactile feedback facilitates coordination between vision of the target and proprioception of the arm. The second set highlights different behavioral strategies that emerge depending on whether or not tactile feedback is available. In the final set, we perturb the visual-proprioceptive relation and investigate how the model learns to recalibrate the two sensory systems.

A. Touch Helps Coordinate Vision and Proprioception

Fig. 3(a) presents the average fitness (i.e., time spent in contact with the target) in two populations of econets over ten simulation runs. (In all analyses and figures, we adopt the convention of presenting the results of the 20 most fit individuals from




Fig. 4. Average fitness in the tactile-on, tactile-off, and tactile-on-luminous population. Tactile-on-luminous econets can see the target, but their hand is invisible ($N = 10$ runs for each population).

each generation, as these are the econets that are selected for reproduction.) Both populations receive visual and proprioceptive sensory inputs. In the “tactile-on” population, sensory input is also received from the tactile input unit (when the target is contacted). In the “tactile-off” population, the tactile input unit in the “tactile-off” population is fixed at zero. Fig. 3(b) presents the average number of trials (out of 20) in which the target was reached. As Fig. 3 illustrates, tactile feedback not only results in reaching the target more often, but also in maintaining contact with it for more time.

Before we can conclude that the advantage of the tactile-on population is due to better coordination of vision and proprioception, several additional analyses are necessary. First, we can ask what information sources are available to each population for determining the hand’s position? Note that while the tactile-on population can “feel” the target, both populations can also see their hand. Thus, it is not clear if the tactile-on population relies on touch to coordinate vision and proprioception or if sight of the hand also plays a role. To address this question, we simulated a variant of the tactile-on population that could not see their hands (tactile-on-luminous). This condition corresponds to an infant that learns to reach for luminous objects in a dark room and allows us to compare reaching when visual information about the location of the hand is either available (tactile-on) or unavailable (tactile-on-luminous) to infants.

Fig. 4 presents the results of the previous analysis, along with the average fitness in the tactile-on-luminous population. There are two important findings. First, in both of the tactile-on populations, performance and rate of learning are near equal. Thus, visibility of the hand does not appear to play a major role in coordinating vision and proprioception when tactile feedback is available. Second, this result also parallels the findings of [7], who observed similar onset ages for reaching in normal and luminous-object conditions. Taken together, these results provide additional support for the conclusion that tactile feedback helps coordinate vision of the target with proprioception of the arm.

However, a second question concerns the informational “content” of the tactile sensory signal. Specifically, the tactile system not only indicates when the hand is touching the target, but also when the econet’s fitness changes (recall that touching the target increases fitness). Because touch and fitness are confounded, it is not clear if touch specifically improves visual–proprioceptive coordination, or if it simply causes a more general change

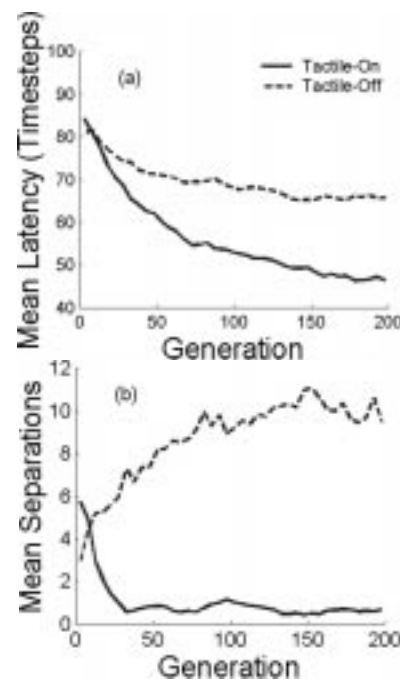


Fig. 5. Average latency. (a) Time until first contact with target per trial and (b) number of separations (number of times contact with target is lost) after reaching the target in the tactile-on and tactile-off populations.

in reaching behavior. One possibility is that both the tactile-on and tactile-off populations reach the target equally well, but that the tactile-on population is better able to maintain contact with the target because they “know” when they are touching it. This is unlikely, as the tactile-on population averages more successful trials that the tactile-off population [see Fig. 3(b)]. Nevertheless, we can evaluate this hypothesis by observing how long it takes each population to reach the target on a successful trial [average latency, see Fig. 5(a)] and how often the hand leaves the target after reaching it [average number of separations, see Fig. 5(b)]. First, note that the average time to reach the target decreases more rapidly in the tactile-on population. Second, tactile-on econets experience fewer separations from the target after reaching it. Tactile-off econets actually experience a roughly constant rate of separations. However, the increase in separations shown in Fig. 5(b) is an artifact due to the fact that tactile-off econets are also increasing the number of times they reach the target [see Fig. 3(b)].




Fig. 6. Average fitness in the “hover” condition, where econets are rewarded for keeping their hand near (but not on) the target ($N = 10$ runs for each population).




Fig. 7. Relative time spent fixating the hand or the object in the tactile-on and tactile-off populations ($N = 10$ runs for each population). Tactile-on econets find the target first and bring the hand toward it while tactile-off econets find the hand first and then search for the target.

Thus, tactile feedback *does* appear to improve visual–proprioceptive coordination. We can test this conclusion more rigorously by dissociating touch and fitness completely. Imagine, for example, a world where econets are not rewarded for contact, but instead for maintaining their hand within a small distance of the target. If the confounding of touch and fitness explains the better performance of the tactile-on population, then we should expect equal performance in the tactile-on and tactile-off populations when touching the target is no longer rewarded. Fig. 6 presents the average fitness for the two populations in the “hover” condition (i.e., 0 fitness when touching the target, +1 for remaining near). While the advantage of the tactile-on population is diminished [compare to Fig. 3(a)], reaching is still significantly better with tactile feedback.

B. Touch Facilitates an Efficient Reaching Strategy

The results thus far suggest that sight of the hand makes a negligible contribution toward visuo-proprioceptive coordination. Instead, it appears that feeling the target while looking at the same location in space facilitates learning to coordinate eye and arm movements. We can explore this coordination more closely by analyzing the specific reaching strategies employed by the tactile-on and tactile-off populations.

Reaching the target requires first determining the positions of the target and the hand and then moving the hand toward the target. We might expect different sequencing strategies to

emerge depending on how well vision and proprioception are coordinated. Specifically, if they are well coordinated, econets should employ a “target-first” strategy in which the target is located and then the hand is brought to it as the target is fixated. However, the target-first strategy depends on knowing the position of the hand, especially when it is outside the visual field. If vision and proprioception are coordinated poorly, we should instead expect a “hand-first” strategy, where econets first bring their hand into the visual field and then search for the target.

The target-first reaching strategy is typical of the tactile-on population. At the start of each trial, econets in this population search for the target as their hand is held outside the visual field. Once the target is located, the eye remains fixated on it as the hand is brought into the visual field and toward the target. We can quantify this reaching strategy by calculating the relative proportion of time spent viewing the hand and the target prior to contacting the target (Fig. 7). In the tactile-on population, the majority of time is spent fixating the target (59%), while the hand is fixated less often (41%). Alternatively, without the sense of touch, the hand-first reaching strategy is employed predominantly. Thus, in the tactile-off population (Fig. 7), the hand is fixated more often than the target (62% and 38%, respectively). However, it is important to note that differences in reaching strategies may not be solely responsible for the two fixation patterns presented in Fig. 7. That is, other behavioral differences in the tactile-on and tactile-off populations (e.g., reliance on visual feedback) may also be contributing to this pattern.




Fig. 8. Normalized fitness in the tactile-on and tactile-off populations during the final 20 generations of training (“preadaptation”) and for 20 generations of simulated prismatic displacement of visual input (“adaptation,” $N = 10$ runs for each population).

C. Touch Promotes Intermodal Recalibration

For the final set of analyses, we ask how econets adapt to a perturbation of the visuo-proprioceptive coordination. In particular, we create an intermodal conflict by simulating a 10° prismatic displacement of the visual field. Thus, while touch and proprioception remain veridical, the visual sensory system now encodes objects in the visual field as 10° to the left of their actual position. This condition is analogous to prismatic adaptation studies with infants, who quickly learn to accommodate their reaching movements while wearing displacement prisms [20], [21].

In order to establish an intermodal coordination, we first train a population of networks. As before, training continues for 200 generations. After the initial training phase (“preadaptation”), prismatic displacement is simulated by shifting the apparent position of all objects in the visual field 10° to the left. Training during the adaptation phase then continues for an additional 20 generations (“adaptation”). Fig. 8 presents the mean normalized fitness in the tactile-on and tactile-off populations during the preadaptation and adaptation phases. (The break in the curves between generations 200 and 201 is exaggerated to illustrate the transition from preadaptation to adaptation.) We present normalized fitness (i.e., scaled according to the maximum in each population) in order to highlight the relative effect of prismatic displacement on each population.

There are two major findings. First, prismatic displacement has a larger initial effect on the tactile-off population. Second, the tactile-on population recovers a higher level of performance during the adaptation phase. These results are consistent with the earlier analyses and provide further support for the hypothesis that touching the target helps coordinate vision of the target with the felt position of the arm.

IV. DISCUSSION

Our first goal was to demonstrate that intermodal coordination of vision and proprioception need not be preprogrammed, but can instead emerge while learning to reach. In all of the simulation conditions that we studied, at least some rudimentary form of visuo-proprioceptive coordination appeared gradually as reaching improved. We also found that vision and proprioception are best coordinated when tactile feedback from the target is available.

Our second goal was to evaluate and compare two mechanisms for achieving intermodal coordination. One candidate mechanism is the use of visual feedback from the hand. We found that when econets cannot feel the target, they employ a reaching strategy that keeps the hand in the visual field, thereby providing continuous feedback of the hand’s position. However, when tactile feedback is available, econets tend to keep the hand out of the visual field until the object is first located. Thus, use of tactile feedback is preferred over visual feedback when both information sources are available. In addition, echoing the results of [7], we also found that as long as the target can be felt, making the hand “invisible” does not interfere with learning to reach.

It is important to note that the low resolution of the visual input system in the Schlesinger *et al.* model may bias econets against exploiting visual feedback from the hand. In order to evaluate the hypothesis that increasing the visual acuity of the model will encourage the use of visual feedback, we conducted a follow-up simulation in which the number of visual input units was doubled from 8 to 16 (while the size of the visual field remained at 64°). Nevertheless, after 200 generations of training, tactile-off econets averaged only 20% of the fitness generated by tactile-on econets in the same condition. Thus, tactile feedback appears to maintain its advantage when visual acuity is increased.

A second possible bias against employing visual feedback stems from the fact that the model represents visual inputs from the target and hand in the same way (i.e., when both the hand and target are in the visual field, there is a strong degree of perceptual aliasing, i.e., perceptual similarity or confusion between the hand and target). Consequently, econets avoid the visual discrimination problem altogether by keeping the hand out of the visual field as much as possible. This observation suggests the developmental hypothesis that infants may learn to employ comparable strategies in order to avoid being distracted by the hand as they maintain sight of the target.

These findings raise the question, is tactile feedback the primary mechanism for coordinating vision and proprioception in human infants? This is unlikely, as infants engage in a wide variety of sensorimotor activities (e.g., visual regard of the hands) prior to learning to reach. It is possible to incorporate these nonreaching behaviors as part of the learning process. For example, Kuperstein [18], [19] studies a model in which prior to

reaching, the simulated infant first constructs a multimodal map that relates arm movements to changes in multiple sensory inputs (e.g., eye position, vision, proprioception of the arm). This modeling approach is based on the idea that the initially random arm movements generated by infants provide a basis for learning to correlate multiple sensory inputs.

Note that the Kuperstein and Schlesinger *et al.* models generate different developmental predictions. For example, because tactile feedback plays a central role (for coordinating vision and proprioception) in the Schlesinger *et al.* model, it predicts that infants should be better at discriminating arm positions (via proprioception) on or near the reaching trajectory; sensitivity to other postures far from the stereotypical reaching trajectory should develop at a later age. In contrast, because the Kuperstein model constructs a multimodal map prior to reaching, it predicts that infants should be equally good at discriminating arm positions across the entire reaching workspace. Second, the Schlesinger *et al.* model also predicts that infants will adapt more quickly to displacement prisms when they have tactile feedback from the displaced target. If the object is virtual (e.g., a hologram) and generates only visual feedback, adaptation should be slower. Alternatively, the Kuperstein model does not predict a difference in the adaptation rates to real and virtual objects in a displacement condition.

Finally, we note that our findings should not be used to suggest that vision of the hand plays a minimal or unimportant role in the development of reaching. Indeed, as we have already mentioned, vision of the hand occurs in a variety of contexts prior to the onset of reaching. One possible developmental scenario suggested by our model is that while visual feedback from the hand may play a secondary role during the onset of reaching, it may gain central importance as multimodal coordination improves. For example, visual feedback will become critical as infants begin to attempt high-precision reaching tasks (e.g., pincer grasp of a small object). In addition, visual feedback will also play a central role as infants learn to manipulate and visually inspect the objects they have grasped.

REFERENCES

- [1] D. H. Ashmead, M. E. McCarty, L. S. Lucas, and M. C. Belvedere, "Visual guidance in infants' reaching toward suddenly displaced targets," *Child Develop.*, vol. 64, no. 4, pp. 1111–1127, Aug. 1993.
- [2] T. Bäck, U. Hammel, and H.-P. Schwefel, "Evolutionary computation: Some comments on the history and current state," *IEEE Trans. Evol. Comput.*, vol. 1, pp. 3–17, Apr. 1997.
- [3] N. E. Berthier, "Learning to reach: A mathematical model," *Development. Psych.*, vol. 32, no. 5, pp. 811–823, Sept. 1996.
- [4] N. E. Berthier, R. K. Clifton, D. D. McCall, and D. J. Robin, "Proximodistal structure of initial reaching in human infants," *Exp. Brain Res.*, vol. 127, no. 3, pp. 259–269, 1999.
- [5] T. G. R. Bower, *Development in Infancy*. San Francisco, CA: Freeman, 1974.
- [6] E. W. Bushnell, "The decline of visually guided reaching during infancy," *Infant Behavior Develop.*, vol. 8, no. 2, pp. 139–155, Apr./June 1985.
- [7] R. K. Clifton, D. W. Muir, D. H. Ashmead, and M. G. Clarkson, "Is visually guided reaching in early infancy a myth?," *Child Develop.*, vol. 64, no. 4, pp. 1099–1110, Aug. 1993.
- [8] R. K. Clifton, P. Rochat, D. J. Robin, and N. E. Berthier, "Multimodal perception in the control of infant reaching," *J. Exp. Psych.: Human Perception Perform.*, vol. 20, pp. 876–886, 1994.
- [9] S. E. Engelbrecht, "Minimum principles in motor control," *J. Math. Psych.*, to be published.

- [10] K. Ennouri and H. Bloch, "Visual control of hand approach movements in new-borns," *Br. J. Development. Psych.*, vol. 14, no. 3, pp. 327–338, Sept. 1996.
- [11] C. M. Harris and D. M. Wolpert, "Signal-dependent noise determines motor planning," *Nature*, vol. 394, no. 6695, pp. 780–784, Aug. 1998.
- [12] C. von Hofsten, "Eye-hand coordination in the newborn," *Development. Psych.*, vol. 18, no. 3, pp. 450–461, May 1982.
- [13] —, "Developmental changes in the organization of prereaching movements," *Development. Psych.*, vol. 20, no. 3, pp. 378–388, May 1984.
- [14] C. von Hofsten and L. Ronnqvist, "The structuring of neonatal arm movements," *Child Develop.*, vol. 64, no. 4, pp. 1046–1057, Aug. 1993.
- [15] J. H. Holland, *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: Univ. Michigan Press, 1975.
- [16] R. A. Jacobs, M. I. Jordan, and A. G. Barto, "Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks," *Cogn. Sci.*, vol. 15, no. 2, pp. 219–250, Apr./June 1991.
- [17] J. Konczak and J. Dichgans, "The development toward stereotypic arm kinematics during reaching in the first 3 years of life," *Exp. Brain Res.*, vol. 117, no. 2, pp. 346–354, 1997.
- [18] M. Kuperstein, "Neural model of adaptive hand-eye coordination for single postures," *Science*, vol. 239, no. 4845, pp. 1308–1311, 1988.
- [19] —, "INFANT neural controller for adaptive sensorimotor coordination," *Neural Netw.*, vol. 4, no. 2, pp. 131–145, 1991.
- [20] P. M. McDonnell, "The development of visually guided reaching," *Perception Psychophys.*, vol. 19, no. 3, pp. 181–185, Sept. 1975.
- [21] P. M. McDonnell and W. C. Abraham, "Adaptation to displacing prisms in human infants," *Perception*, vol. 8, no. 2, pp. 175–185, 1979.
- [22] D. Parisi, F. Cecconi, and S. Nolfi, "Econets: Neural networks that learn in an environment," *Network*, vol. 1, no. 2, pp. 149–168, Apr. 1990.
- [23] J. Piaget, *The Origins of Intelligence in Children*. New York: Basic Books, 1952.
- [24] M. Schlesinger, D. Parisi, and J. Langer, "Learning to reach by constraining the movement search space," *Development. Sci.*, vol. 3, no. 1, pp. 67–80, Mar. 2000.
- [25] O. Sporns and G. M. Edelman, "Solving Bernstein's problem: A proposal for the development of coordinated movement by selection," *Child Develop.*, vol. 64, no. 4, pp. 960–981, Aug. 1993.
- [26] E. Thelen, D. Corbetta, K. Kamm, J. P. Spencer, K. Schneider, and R. F. Zernicke, "The transition to reaching: Mapping intention and intrinsic dynamics," *Child Develop.*, vol. 64, no. 4, pp. 1058–1098, Aug. 1993.
- [27] B. L. White, P. Castle, and R. Held, "Observations on the development of visually directed reaching," *Child Develop.*, vol. 35, no. 2, pp. 349–364, 1964.
- [28] R. F. Zernicke and K. Schneider, "Biomechanics and developmental motor control," *Child Develop.*, vol. 64, no. 4, pp. 982–1004, Aug. 1993.



Matthew Schlesinger received the Ph.D. degree in cognitive psychology from the University of California, Berkeley, in 1995.

He was a Postdoctoral Researcher at the Italian National Research Council, Rome, Italy, and the Department of Computer Science, University of Massachusetts, Amherst. He is currently an Assistant Professor at the Department of Psychology, Southern Illinois University, Carbondale, IL. His current research interests include problem-solving and sensorimotor development in human infants and nonhuman primates,

as well as agent-based models of sensorimotor development.



Domenico Parisi received the B.A. degree in philosophy from the University of Rome, Italy, the M.A. degree in psychology from the University of Illinois, Urbana, and the Ph.D. degree in psychology from the University of Rome.

He is currently with the Institute of Psychology, National Research Council, Rome, Italy, and also teaches at Libera Università Maria ss. Assunta (LUMSA) University, Rome. He has coauthored *Rethinking Innateness. A Connectionist Perspective on Development* (Cambridge, MA: MIT Press,

1996). His current research interests include simulations of both individual and social behavior using neural networks, genetic algorithms, and artificial life models.