Running head: CONSTRAINING THE MOVEMENT SEARCH SPACE

Developmental Science (2000), *3*, 67-80

Learning to Reach by Constraining the Movement Search Space

Matthew Schlesinger and Domenico Parisi

Italian National Research Council

Jonas Langer

University of California, Berkeley

Abstract

Trial-and-error learning strategies play a central role in sensorimotor development during early infancy. However, learning to reach by trial-and-error normally requires a slow and laborious search through the space of possible movements. We propose a computational model of reaching based on the notion that early sensorimotor control is driven by the generation of exploratory movements, followed by the selection and maintenance of adaptive movement patterns. We find that instead of exhaustively exploring the full search space of movement patterns, the model exploits several emergent constraints that limit the initial size of the movement search space. These constraints exploit both mechanical and kinematic properties of the reaching task. We relate these results to the development of reaching during infancy, and discuss recent findings that have identified similar constraints in young infants.

Learning to Reach by Constraining the Movement Search Space

Several recent reports have focused attention on the development of reaching in young infants (Ashmead, McCarty, Lucas, & Belvedere, 1993; Bushnell, 1985; Clifton, Rochat, Robin, & Berthier, 1994; Ennouri & Bloch, 1996; Konczak & Dichgans, 1997; Lew & Butterworth, 1997; Thelen et al., 1993; von Hofsten & Ronnqvist, 1993). This research is complemented by a wide variety of computational models that simulate learning to reach (Berthier, 1996 & 1997; Berthier, Singh, Barto, & Houk, 1993; Bullock & Grossberg, 1988; Hinton, 1984; Kawato, 1990; Kettner, Marcario, Port, 1993; Kuperstein, 1988; Rosenbaum, Loukopoulos, Meulenbroek, Vaughan, & Engelbrecht, 1995; Sporns & Edelman, 1993; Vos & Scheepstra, 1993). There are several important differences, though, between how these models simulate motor learning and how infants learn to reach. First, many models employ a supervised learning algorithm (e.g., Bullock & Grossberg, 1988; Kettner, Marcario, & Port, 1993; Vos & Scheepstra, 1993). This approach assumes that the learning agent has access to an error feedback signal, presumably through closed-loop visual feedback of hand movements. However, because young infants do not watch their hands as they learn to reach, it is unlikely that visual feedback plays a major role in learning to reach during early infancy (Ashmead, McCarty, Lucas, & Belvedere, 1993; Bushnell, 1985; Clifton, Muir, Ashmead, & Clarkson, 1993; Clifton, Rochat, Robin, & Berthier, 1994; Ennouri & Bloch, 1996; Thelen et al., 1993).

A second difference concerns whether or not the movement trajectory is planned. One strategy, often employed in robotics (see Hollerbach, 1990), is to first plan a reaching trajectory in Cartesian coordinates, and then to use an internal (inverse) model to convert these coordinates into joint-angle changes or muscle activations (Flash, 1987; Kawato, 1990; Saltzman, 1979). However, as several infant researchers have argued, there is little or no evidence to suggest that infants use a pre-planned motor program for reaching (e.g., Thelen, et al., 1993).

Constraining the Movement Search Space

A computational model of reaching which is consistent with infant theory and research should focus on trial-and-error learning strategies. According to this view, reaching develops by generating a wide range of exploratory movements, followed by selection of the most adaptive or functional reaching patterns (Berthier, 1996; von Hofsten, 1984; Thelen et al., 1993). However, because the search process is not directed, it is both slow and inefficient: depending on the level of analysis (see Bernstein, 1967), the resulting movement search space is potentially huge (e.g., consider all possible combinations of arm positions, muscle activation patterns, motor neuron firing rates, etc.).

One method for reducing the size of the search space, as well as learning time, is to selectively restrict the initial range of exploratory movements. For example, as novice adults learn to swing a baseball bat, they tend to hold one or more parts of their body (e.g., shoulders or waist) stiff as they learn to swing (Turvey, Fitch, & Tuller, 1982). This strategy effectively reduces the number of joint degrees of freedom (DOF), decreasing the range of possible swinging motions that are initially produced. With experience, novices gradually learn to "unlock" the stiffened joints, incorporating the rotation of these joints into the previously acquired movement patterns (see Vereijken, Emmerik, Whiting, & Newell, 1992, for a comparable description of learning to ski).

Similarly, computational models of language acquisition (e.g., Elman, 1993) have illustrated that by "starting small", that is, by restricting early learning to a small subset of all possible experiences, overall learning is facilitated. However, this approach has not been studied in the context of simulating the development of reaching. The purpose of the present model is to demonstrate that similar constraints are available to infants as they learn to reach. In order to address this issue, we next introduce a kinematic model of reaching which employs an unsupervised, trial-and-error learning strategy. Despite random exploration in the movement search space, the model does not exhaustively produce all possible reaching movements. Instead, we show how a number of mechanical and kinematic constraints emerge in parallel with the trial-and-error strategy, helping to collectively reduce the size of the initial movement search space during learning.

The Econet Model of Reaching

The goal of our model is to simulate the first major developmental period of reaching, where poorly-organized reaching movements are gradually shaped into relatively direct reaches. During this early stage, the priority is on learning to control simply the movements of the arm, rather than making smooth, precise reaches. Thelen et al. (1993) describe this early period of reaching as learning to make decent, "ball park" reaches: "...it appears that the first task in learning to reach is to get the hand somewhere in the vicinity of the desired object, and that this is primarily a problem of control of the arms through haptic and proprioceptive information" (p. 1059).

In contrast to other computational models of motor learning, we propose that this early period of reaching be simulated by training a population of artificial neural networks. Parisi, Cecconi, and Nolfi (1990) suggest the term econet (i.e., ecological neural network), to capture the idea of an agent that learns to act adaptively in a simple environment (e.g., to find food, avoid obstacles, etc.). The econet model of reaching described in this article is an abstract representation of unsupervised, trial-and-error learning in young infants, and is based on two premises: (1) that infants generate a wide range of alternative reaching movements, and (2) that the most successful movements are retained and improved while the others are eliminated. In a comparable manner, as the population of econets generate reaching movements, they also adapt in response to environmental and selectionist pressures (see Sporns & Edelman, 1993, for a different implementation of the population-based approach to motor learning).

Two limitations of the model should be noted. First, we do not address variants of the reaching task such as learning to reach for moving objects or around obstacles (Diamond & Gilbert, 1989; Hofsten, 1980 & 1984; Lockman, 1984). Indeed, findings from several infant studies suggest that these reaching skills develop in parallel with reaching for stationary, unobstructed objects (van der Meer, van der Weel, & Lee, 1994; von Hofsten, 1979). However, because most infant studies have concentrated on the stationary variant, we have initially restricted our model in a comparable way. Second, although we simulate muscle activity, the econet model of reaching is a kinematic representation of the motor system. In this version of the model, we do not deal with the multiple forces involved in arm movements (i.e., gravity, inertia, centripetal and coriolis forces, etc.).

We divide the description of the model into three sections: (1) the environment, (2) the econet, and (3) the learning algorithm. Each is described briefly here, while additional details are presented in the Appendix for the interested reader.

The Environment

Figure 1 presents a schematic diagram of the econet and its environment. Econets learn to reach in a 2-dimensional workspace that is 100 x 100 units wide. During each trial, a small object is randomly positioned in the environment, within reach of the econet. The position of the econet remains fixed at the center of the workspace.

The Econet

The econet has a monocular visual system, with a 64° visual field. The econet also has a two-segment arm; each segment is 25 units long. While the eye (i.e., head or body-axis) can rotate a full 360° in either direction, the shoulder and elbow joints are limited to 180° of rotation (relative to the segments to which they link). The upper arm is linked to the body-axis, and follows its rotation, while the forearm is linked in a similar manner to the upper arm. Eye movements (i.e., head or body-axis rotation) are controlled by a single muscle, while each arm segment is controlled by two muscles that pull in opposite directions (i.e., as opponent, or agonist-antagonist, pairs). The activation states of the eye and arm muscles are determined by a 4-layer, feed-forward artificial neural network. Figure 2 presents a diagram of the network structure. The input layer is divided into visual, proprioceptive, and tactile modalities or sensory channels. First, each visual input unit covers 8° of the visual field. Second, the proprioceptive input units encode the amount of stretch in the muscles of the arm. Finally, the tactile input unit becomes activated when the hand (i.e., arm endpoint) makes contact with the object.

Before all of the sensory signals converge, the visual and proprioceptive systems each project to a separate intramodal hidden layer. An advantage of this 4-layer architecture is that information in the visual and proprioceptive channels is internally re-represented at the intramodal hidden layer, before being integrated with the tactile sensory input. In addition, this architecture is also based on both neurophysiological and computational findings which suggest that partial modularization of sensory input may facilitate feature detection, while minimizing crosstalk or interference effects across sensory modalities (Jacobs, Jordan, & Barto, 1991; Mishkin, Ungerleider, & Macko, 1983). In order to verify this result, we compare the performance of the 4-layer network to a 3-layer network in which the sensory signals are not initially separated.

Next, the visual, proprioceptive, and tactile signals are integrated at the intermodal hidden layer. The output of the final layer controls eye and arm movements. The eye moves a maximum of 15° left or right per timestep as a function of the first output unit. The remaining four output units are grouped into two pairs. The first pair of units contract the extensor and flexor muscles in the upper arm, causing the shoulder joint to rotate a maximum of 15° degrees left or right. Movement of the forearm is determined in the same manner, using the remaining pair of output units.

The Learning Algorithm

In order to simulate the development of reaching, an initial population of econets is generated, each with the same body structure and neural network architecture. Each econet is roughly analogous to a class of motor action patterns (i.e., "styles" or strategies of visuomotor coordination); the population of econets represent a collection of possible movement patterns from which the most adaptive reaching movements are selected. Econets differ only in the values of the connection weights in their neural networks. During the first generation, all weights are assigned random real-number values between -10 and 10. Each econet is then tested by placing it in the environment, and giving it 20 trials to reach the object. Each trial lasts 100 movement (i.e., input-output) cycles. The econet is randomly positioned at the start of each trial. All subsequent positions -- and thus, sensory input patterns -- are determined as a function of the movement of the econet.

For each cycle that the econet spends with its hand in contact with the object, it receives a 1-point reward. Otherwise, it receives a reward of 0. Each econet's fitness score is the sum of cycles spent in contact with the object. Kinematic studies of reaching development in humans suggest that faster reaches (i.e., reaches with no deceleration component) are associated with batting or pushing the target object, rather than grasping it (e.g., Berthier, 1996). Thus, we chose the present reward function in order to study reaching movement patterns that provide the opportunity for grasping and retrieving objects (and consequently, object exploration and manipulation). In addition, this method of rewarding contact with the object also creates a bias toward minimum-time reaching movements (Berthier, 1996; Harris & Wolpert, 1998; see Engelbrecht, in press, for a review of minimum principles in motor control).

A variation-and-selection learning algorithm, analogous to evolutionary learning (Holland, 1975), is employed as a training regime. After testing, the econets are sorted as a function of their fitness and the twenty best are then selected. Each of the twenty are duplicated five times (for a new population of 100 econets); the five copies are not identical to the original, however. Variability is introduced into the next generation of econets by strengthening or weakening 2% of the connection weights, selected at random, by a small random amount. The process of testing, selecting, and reproducing econets continues for 200 generations.

Results

We divide the results into two major sets of simulations and analyses. In the first section, we briefly review the overall performance of the model during learning. In the second set of results, we identify and describe four different emergent constraints that help to reduce the size of the movement search space.

Overall Performance

We consider four specific performance measures. Fitness is the sum of all movement cycles spent with the hand on the target object. Successful trials are defined as the number of trials in which the target object is contacted at least once. Latency is time of first contact during a trial. Finally, separations are the number of times that the hand leaves the target object, after successfully reaching it.

In order to evaluate performance in the normal condition, we present the results from two additional training conditions. The first additional condition ("3-layer") is a population of econets trained with a 3-layer architecture. Three-layer econets have the same input and output systems as normal econets; the two hidden layers, however, are replaced by a single hidden layer with 8 units. In this case, all sensory input units project in parallel to the single hidden layer. This condition helps us to evaluate the prediction that partially modularizing the sensory signals will facilitate learning (see "The Econet", above). In the second additional condition ("random-movement"), all of the econets are programmed to generate random movements. This condition provides us with an estimate of the performance level expected by chance.

Figure 3 presents the average fitness (3a), number of successful trials (3b), latency (3c), and number of separations (3d), as a function of training condition. The data represent the average of ten independent runs (i.e., randomized initial values) for each condition (recall that during each run, 100 econets x 200 generations are trained). In this and all subsequent figures, unless otherwise noted, we adopt the convention of presenting the average results of the top 20 econets from each generation (i.e., those that were selected for reproduction).

As predicted, econets with the 4-layer network architecture learn more quickly than 3-layer econets, generating higher fitness scores from virtually the start of training (see Figure 3a). Although there is only a small difference in the two conditions with respect to reaching latency (3c), 3-layer econets average fewer successful trials than their 4-layer counterparts (3b). In addition, after reaching the target, 3-layer econets tend to lose contact with it more often than 4-layer econets (3d). Indeed, 3-layer econets average roughly as many separations as would be expected by chance.

Emergent Constraints

While analyzing learning and performance in the normal condition, we identified four specific learning patterns that function as constraints on the movement search space. Each constraint facilitates learning by reducing the range of possible movements. While some of the constraints appear to be only present during early learning (e.g., the first 20 generations), others function across a longer time-span. We next describe each of the constraints. Note that except when mentioned otherwise, all analyses refer to learning in the normal condition.

1. Joint-Locking. The econet has three joint DOF: the eye/head/body-axis, the shoulder joint, and the elbow joint. Because the arm workspace is 2-dimensional, one of the three joints is redundant. In other words, only two of the three joints are needed to reach any given target in the workspace. An adaptive solution to this DOF problem is to eliminate one of the redundant joints

by locking it in place (Bernstein, 1967; Sporns & Edelman, 1993; Vereijken, Emmerik, Whiting, & Newell, 1992). This strategy creates a primitive joint synergy, by joining two adjacent components of the biomechanical system into a single link.

It is important, however, not to eliminate any necessary joints. In our model, locking the elbow or body-axis would make some portions of the arm workspace unreachable. Thus, the optimal strategy in this case is to lock the shoulder joint, while leaving the body-axis and elbow joints mobile. Indeed, this is the strategy that emerged in all of the runs that we performed. In eight of the ten runs, econets learned to hold their shoulder in the fully extended position; in the remaining two runs, the shoulder was held in the fully flexed position.

We decided to analyze this phenomenon more closely, in order to identify the specific mechanism through which the joint-locking strategy is achieved. Our analyses led us to the pair of muscles controlling the shoulder joint: we found that in all ten runs, one of the two muscles quickly became dominant, consistently pulling the shoulder joint until the upper or lower joint limit was reached. Figure 4 presents the percent of cycles in which there was shoulder flexion, for each of the ten runs. An average of 50% means that on half of all cycles, the flexor was activated more than the extensor, causing flexion of the upper arm (i.e., rotation toward the left). We present the results from all econets in each generation, including all of the cycles in a trial up until the target object was reached.

In eight of the ten runs, the shoulder extensor became the dominant muscle (i.e., virtually 100% of all upper arm movements were rotations away from the body); in the remaining two runs (1 and 7), shoulder flexion was dominant. Figure 5 presents the comparable data for the elbow. In contrast to the shoulder joint, the elbow remains compliant (i.e., avoids locking into place); in nine of ten runs, neither of the two elbow muscles consistently controlled the movement of the forearm.

How does this joint-locking strategy help to constrain the search space of reaching movements? By locking the shoulder, the body-axis and upper arm become effectively fused: wherever the body-axis turns, the upper arm follows. Consequently, only rotations of the body-axis and elbow need to be controlled. This strategy eliminates one dimension of the search space (at the joint level), reducing it from a 3- to a 2-dimensional search problem.

2. Muscle Coactivation. Following the work of Bernstein (1967), Kawato (1990) discusses the muscle DOF problem. Comparable to the level of joints, there are an infinite number of different muscle activation patterns that can be used to produce a given arm movement. How should a particular muscle activation pattern be selected by the learning agent? Figure 6 presents a schematic diagram of muscle activation space: flexor activation is plotted against extensor activation. We can distinguish between two major classes of muscle activation patterns. The line labeled "ideal coactivation" represents the set of activation patterns with equal muscle activations in both muscles (i.e., co-contractions). This type of pattern creates muscle stiffness, and causes the joint angle to remain fixed. Alternatively, the line labeled "ideal reciprocal activation" represents optimal patterns of arm movement, where a given muscle relaxes in proportion to the activity of its opponent muscle.

In order to see whether a particular muscle activation pattern tends to emerge during learning, we took each flexor-extensor activation pair and computed its deviation (i.e., minimum distance) from the "reciprocal activation" and "coactivation" curves. Figures 7a and 7b present the average deviations, for the shoulder and elbow joints, respectively, prior to first contact with the object in each trial. There are two major results. First, during early learning, the coactivation pattern is predominant in both the shoulder and elbow muscle pairs (remember that lower average values indicate a closer fit to the ideal curve). Second, as the shoulder muscles shift to using the reciprocal activation pattern, the elbow muscles remain consistently coactivated. Although coactivation of the elbow muscles slowly diminishes over learning, it continues to be the dominant pattern.

The shift to the reciprocal activation pattern in the shoulder muscles raises several questions concerning potential dependencies between the muscular and joint levels. First, is the tendency to lock the shoulder joint simply a result of the emerging reciprocal activation pattern? This is unlikely, since reciprocal activation does not guarantee that either the flexor or extensor will consistently predominate, but instead constrains the mutual activity of the muscle pair. Second, we might also wonder if "relaxing" the shoulder muscles is inconsistent with the tendency to lock the shoulder joint in position. As our model does not simulate the dynamic properties of the arm, we cannot yet fully answer this question. A preliminary explanation, though, is that as the joint constraint is imposed on the shoulder, the need for the muscular constraint is reduced; consequently, the shoulder muscles soon discover an activation pattern which moves the shoulder to its preferred position as efficiently as possible.

Thus, the muscular level appears to produce constraints on the movement search space that independently complement those found at the joint level. Initial muscle coactivation (or co-contractions) "stiffen" both the shoulder and elbow joints, reducing the range of arm movements. This finding replicates the developmental pattern reported by Thelen et al. (1993). While the elbow muscles appear to maintain the coactivation pattern for the extent of training, however, the shoulder muscles gradually "relax". Extrapolating from the present results, we might predict that the elbow muscles will also switch to the reciprocal activation pattern with further training. This learning sequence would provide support for a proximodistal direction of motor control, from the shoulder to the elbow (von Hofsten, 1984; Trevarthen, Murray, & Hubley, 1981).

3. Stereotyped Approach Movements. As we studied the reaching movement patterns produced by our model during learning, we identified a third emergent constraint on reaching

movements. Unlike the first two constraints, that operate at the joint and muscular levels, the third constraint exerts its primary influence at the kinematic level, by shaping the direction of the hand path during approach movements. We originally expected to find that approach movements would become progressively more direct. While this is generally true (i.e., the probability of approaching the object when it is fixated increases with learning; see Figure 8), we were disappointed to discover that even near the end of training, roughly 25% of the movements are directed away from rather than toward the target. However, as we analyzed these movements in more detail, we began to understand their kinematic function more clearly. In most cases, relatively direct reaches are preceded by a stereotyped preparatory movement, during which the hand is first brought back toward the econet, to the same initial position. We suspected that the proprioceptive sensory system played a role in this behavior. In order to investigate our suspicion, we "lesioned" the proprioceptive sensory system after training was completed, by "severing" all connections from the proprioceptive input units. Using this method, we found that the approach movements became severely disrupted. In particular, initial reaching movements appeared jerky and disorganized, and were no longer initiated from the same position.

What is the adaptive significance of these stereotyped approach movements? There are several possible advantages. First, in parallel with the first two constraints, stereotyped approach movements also reduce the range of possible reaching patterns, by eliminating variability in the starting position of each reach. This type of kinematic strategy may also facilitate the development of a proprioceptive map, by anchoring the map at the preferred starting point.

4. <u>Movement Slowdown</u>. We identified a fourth emergent constraint, while analyzing the frequency of <u>prereaching</u> movements in our model. Prereaches are poorly-controlled, ballistic-like reaching movements (Bower, Broughton, and Moore, 1970; Bruner, 1973; von Hofsten, 1982 & 1984; von Hofsten and Ronnqvist, 1993; Ennouri & Bloch, 1996; Trevarthen, 1975). Von Hofsten

(1984) describes three stages in prereaching development. During the first stage, prereaches are visually elicited and appear relatively reflexive. Next, prereaching movements become temporarily suppressed, and then in the third stage they reappear (near the onset of reaching) as more adaptively organized approach movements (Bushnell, 1985; Fetters & Todd, 1987; Konczak & Dichgans, 1997; Mathew & Cook, 1990).

Following von Hofsten's (1984) analysis of forward extensions of the hand, we defined an approach movement as a reduction of the distance between the hand and target. We then computed the percent of approach movements during object fixation, as well as during non-fixation. Figure 8 presents the probability of an approach movement in the normal condition, as a function of target fixation. For purposes of comparison, we also present the probability of an approach movement during fixation in the random-movement condition.

As Figure 8 indicates, the probability of an approach movement occurring at random (i.e., in the random-movement condition) is 0.5. Initially, the likelihood of producing an approach movement during non-fixation in the normal condition is greater than chance. While perhaps counterintuitive, this result is a consequence of the fact that during the early stages of learning, the econet spends more time visually searching for the object than actually reaching for it. Because the arm movements during this period of learning are not random, but instead often entrained with the rotation of the visual field, the probability of moving the hand toward the object is greater than chance. However, as the movements of the eye and arm are gradually dissociated, the probability of an approach movement during non-fixation declines to chance level. In contrast, fixated approach movements remain below chance during the first 50 generations. Thus, during early learning, sight of the object appears to suppress the likelihood of moving the hand toward the target object. After the initial decline, fixated approach movements steadily increase. This learning pattern closely mirrors the developmental sequence described by von Hofsten (1984). Given the present analyses, though, there are at least two ways to interpret the initial decline in approach movements. Either: (a) it all fixated movements become less likely (i.e., the arm spends more time in place), or (b) the probability of moving away from the target temporarily increases.

Additional analyses lend support to the first explanation. Figure 9 presents the average speed of hand (9a) and head/eye (9b) movements during fixation of the target object. In both the visual and reaching motor systems, sight of the object is associated with an initial movement slowdown. There are several possible causes for this early slowdown. For example, one of the first abilities that appears to emerge is the capacity to "hold still" once the target object has been successfully reached (see Figure 3d; recall that fitness is a function of reaching and then maintaining contact with the target). A consequence of this ability is that sight of the object is temporarily sufficient to decrease motor activity. In addition, the movement slowdown may also be the result of learning to maintain sight of the target, once it is fixated.

At first glance, the movement slowdown appears to be a non-adaptive change. For example, slower movements (and fewer approach movements, in particular) mean that the target object is reached less often. However, there are several potential advantages to a temporary movement slowdown. First, slower hand movements may facilitate learning a vision-proprioception intermodal mapping; because the hand covers less space over time, fewer visuo-proprioceptive states are encountered (the stereotyped approach movements would complement this effect). In addition, by slowing down the speed of hand movements, the number of different positions assumed by the arm is decreased. Like the other constraints, this may also reduce the movement search space.

Discussion

The principal goal of our model has been to demonstrate that several constraints on reaching can emerge in parallel with a trial-and-error learning strategy, and play a critical role during learning. Two major findings are discussed in order below. First, we discuss four constraints that collectively reduce the size of the movement search space. Next, we discuss the emergence of the four constraints. The first two constraints are mechanical. At the shoulder, arm movements are limited by rotating and locking the shoulder (normally in the fully extended position), while recruiting body-axis and elbow rotation for the purpose of reaching. Interestingly, an optimal activation pattern also emerges in the shoulder muscles, in conjunction with this mechanical constraint; presumably, this activation pattern helps to lock the shoulder into position as quickly and efficiently as possible. At the elbow, meanwhile, muscle coactivation predominates throughout training, maintaining the elbow in a relatively "stiffened" state. Thus, the shoulder exploits joint mechanics while the elbow exploits muscle mechanics; nevertheless, both constraints limit not only the possible positions of the arm, but also the size of arm movements.

The second pair of constraints operate at the kinematic level. First, we observed that regardless of initial arm and object position, a consistent, stereotyped movement pattern is generated. This pattern of approach movements may be related to learning to use the proprioceptive sensory channel, as we found that eliminating proprioceptive input causes these approach movements to become disorganized. Second, we also observed a temporary decline in hand (and "head") speed. Together, these constraints decrease the full range of alternative reaching movements.

The second major finding from our model concerns how these constraints appear. In past models, two approaches have been studied. First, some models have constrained the search process by employing a supervised error-feedback learning algorithm. However, as we have already noted, it is unlikely that young infants learn to reach with a closed-loop, error-feedback strategy (though it may play a more prominent role after the onset of reaching; see Ashmead, McCarty, Lucas, & Belvedere, 1993; Bushnell, 1985; McDonnell & Abraham, 1979). Meanwhile, other computational models limit the search space by imposing built-in constraints (e.g., Elman, 1993), that are designed to represent the effects of maturational factors. Our results indicate that neither error-feedback nor preprogrammed maturational control is necessary. Instead, several constraints appear to "fall out" as a consequence of a relatively simple trial-and-error learning algorithm.

While our model attempts to represent many of the important characteristics of reaching development in infants, we have simplified or accommodated some features for the purposes of simulation. Nevertheless, the predictive value of the model can be assessed directly by comparing our results with the empirical database. In particular, what evidence, if any, is there for the existence of these constraints during the development of reaching in infants? First, a recent kinematic analysis of young infants during the onset of reaching provides support for the joint-locking phenomenon (Berthier, Clifton, McCall, & Robin, in press). In this study, it was found that young infants tend to use shoulder and torso rotations as they begin to reach, while holding the elbow in place. This reaching pattern is similar to the strategy observed in our model (i.e., combining body-axis rotation with elbow movements). However, because there are different numbers of total joint DOF (i.e., at least 5 major joints in infants, and three in our model), different specific combinations of "locked" and rotating joints are observed in each case.

Second, electromyographic (EMG) data from young infants' muscle activity near the onset of reaching are consistent with the second constraint that we observed: co-contraction of flexor and extensor muscles is the most common pattern in infants as they begin to learn to reach (Thelen et al., 1993). One explanation for this pattern is that elevated levels of muscle stiffness are initially used to control the reactive forces of the arm (e.g., inertia, torque). However, we obtain a comparable result with a kinematic model. This finding suggests that increased arm stiffness during early infancy may be due to both kinematic and dynamic factors. While Thelen et al. (1993) do not report any age-wise changes in muscle activity as a function of agonist-antagonist pair, the results of our simulations suggest the hypothesis that more central or proximal muscle pairs should transition to the reciprocal inhibition pattern prior to more distal pairs (e.g., shoulder before elbow). Third, we also observed a pattern of stereotyped approach movements as reaches are initiated. Berthier et al. (in press) report a comparable strategy: regardless of pre-trial hand-position, many infants consistently moved their hand back and/or up to their side, before reaching for the target. Our analyses have implicated the proprioceptive sensory system as being at least partially responsible for this phenomenon. Besides reducing the movement search space, we also suggested that these stereotyped movements may help to create a spatial "anchor", linking visual and proprioceptive systems. This explanation generates a number of developmental hypotheses. For example, we should expect young infants to be better at discriminating between postures near this preferred starting position, compared to other arm postures. In addition, we also predict more accurate reaches when infants begin from the preferred position, compared to reaches that are initiated from other positions.

Finally, we found that an early movement slowdown led to a temporary drop in approach movements. Although no developmental studies have yet identified a movement slowdown prior to the onset of reaching, (however, Berthier and McCarty, 1996, report a temporary slowdown in 7-month-olds), our model replicates the initial decline in infants' prereaching approach movements reported by von Hofsten (1984). It is possible that the temporary decline in infants' prereaches is also due to an overall slowdown in reaching movements. Because von Hofsten (1984) did not include a measure of hand speed, though, this remains an open question. It is important to remember that in our model, the slowdown occurs relatively early and is short-lived. Similarly, von Hofsten (1984) reports a drop in prereaches at age 7 weeks, followed by a rapid recovery. Unfortunately, because most infant studies have concentrated on the onset of reaching -- around the age of 14 to 20 weeks -- there is comparatively little developmental data on the early weeks of infancy. Thus, additional research is necessary to confirm whether, like our model, young infants exploit a movement slowdown as a constraint on learning to reach.

As we noted earlier, the econet model of reaching is directed toward simulating the first major developmental period of reaching. During this stage, "just getting there" is the putative goal for infants. Our simulation results suggest a number of constraints available to infants for initially simplifying this task. There are two important advantages to employing this constraint-based approach, in contrast to learning without any initial constraints. First, early learning is accelerated by reducing the size of the movement search space. It is clear, though, that in order for more sophisticated reaching movements to develop (e.g., around obstacles), these constraints must ultimately be relaxed or perhaps even eliminated (Turvey, Fitch, & Tuller, 1982; Vereijken, Emmerik, Whiting, & Newell, 1992). This observation draws attention to the second (and perhaps more interesting) advantage of the constraint-based approach. As the initial movement patterns are learned, they create a potential repertoire of movement primitives which can then be combined (either sequentially or in parallel) in more difficult task contexts (see Greeno & Simon, 1974, for an introduction to the sequence production problem).

References

Ashmead, D.H., McCarty, M.E., Lucas, L.S., & Belvedere, M.C. (1993). Visual guidance in infants' reaching toward suddenly displaced targets. Child Development, 64, 1111-1127.

Bernstein, N.A. (1967). The coordination and regulation of movements. Oxford: Pergamon.

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

Berthier, N.E. (1997). Analysis of reaching for stationary and moving objects in the human infant. In J.W. Donahoe & V.P. Dorsel (Eds.), Neural-network models of cognition: Biobehavioral foundations, pp. 283-301. New York: Elsevier.

Berthier, N.E., & McCarty, M.E. (1996, April). Speed of infant reaching during the first year: Confirmation of a prediction. Presented at the International Conference on Infant Studies, Providence, RI.

Berthier, N.E., Clifton, R.K., McCall, D.D., & Robin, D.J. (in press). Structure of initial reaching in human infants. Experimental Brain Research.

Berthier, N.E., Singh, S.P., Barto, A.G., & Houk, J.C. (1993). Distributed representation of limb motor programs in arrays of adjustable pattern generators. Journal of Cognitive Neuroscience, 5, 56-78.

Bower, T.G.R., Broughton, J.M., & Moore, M.K. (1970). Demonstration of the intention in the reaching behavior of neonate humans. Nature, 228, 679-681.

Bruner, J.S. (1973). Organization of early skilled action. Child Development, 44, 1-11.

Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: Emergent invariants and speed-accuracy properties during trajectory formation. Psychological Review, 95, 49-90.

Bushnell, E.W. (1985). The decline of visually guided reaching during infancy. Infant Behavior and Development, 8, 139-155.

Clifton, R.K., Muir, D.W., Ashmead, D.H., & Clarkson, M.G. (1993). Is visually guided reaching in early infancy a myth? Child Development, 64, 1099-1110.

Clifton, R.K., Rochat, P, Robin, D.J., & Berthier, N.E. (1994). Multimodal perception in the control of infant reaching. Journal of Experimental Psychology: Human Perception and Performance, 20, 876-886.

Diamond, A., & Gilbert, J. (1989). Development as progressive inhibitory control of action: Retrieval of a contiguous object. Cognitive Development, 4, 223-249.

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

Engelbrecht, S.E. (in press). Minimum principles in motor control. Journal of Mathematical Psychology.

Ennouri, K., & Bloch, H. (1996). Visual control of hand approach movements in new-borns. British Journal of Developmental Psychology, 14, 327-338.

Fetters, L., & Todd, J. (1987). Quantitative assessment of infant reaching movements. Journal of Motor Behavior, 19, 147-166.

Flash, T. (1987). The control of hand equilibrium trajectories in multi-joint arm movements. Biological Cybernetics, 57, 257-274.

Greeno, J.G., & Simon, H. (1974). Processes for sequence production. Psychological Review, 81, 187-197.

Harris, C.M., & Wolpert, D.M. (1998). Signal-dependent noise determines motor planning. Nature, 394, 780-784.

Hinton, G. (1984). Parallel computations for controlling an arm. Journal of Motor Behavior, 16, 171-194.

Hofsten, C. von. (1979). Development of visually directed reaching: The approach phase. Journal of Movement Studies, 5, 160-178.

Hofsten, C. von. (1980). Predictive reaching for moving objects by human infants. Journal of Experimental Child Psychology, 30, 369-382.

Hofsten, C. von. (1982). Eye-hand coordination in the newborn. Developmental Psychology, 18, 450-461.

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

Hofsten, C. von, & Ronnqvist, L. (1993). The structuring of neonatal arm movements. Child Development, 64, 1-46-1057.

Holland, J.H. (1975). Adaptation in natural and artificial systems. Ann Arbor, MI: University of Michigan Press.

Hollerbach, J.M. (1990). Planning of arm movements. In D.N. Osherson, S.M. Kosslyn, & J.M. Hollerbach (Eds.), An invitation to cognitive science: Visual cognition and action, Vol. 2, pp. 183-211. Boston: MIT Press.

Jacobs, R.A., Jordan, M.I., & Barto, A.G. (1991). Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. Cognitive Science, 15, 219-250.

Kawato, M. (1990). Computational schemes and neural network models for formation and control of multijoint arm trajectory. In W.T. Miller, R.S. Sutton, & P.J. Werbos (Eds.), Neural networks for control, pp. 197-228. Boston: MIT Press.

Kettner, R., Marcario, J., & Port, N. (1993). A neural network model of cortical activity during reaching. Journal of Cognitive Neuroscience, 5, 14-33.

Konczak, J., & Dichgans, J. (1997). The development toward stereotypic arm kinematics during reaching in the first 3 years of life. Experimental Brain Research, 117, 346-354.

Kuperstein, M. (1988). Neural model of adaptive hand-eye coordination for single postures. Science, 239, 1308-1311.

Lew, A.R., & Butterworth, G. (1997). The development of hand-mouth coordination in 2- to 5-month-old infants: Similarities with reaching and grasping. Infant Behavior and Development, 20, 59-69.

Lockman, J.J. (1984). The development of detour ability during infancy. Child Development, 55, 482-491.

Mathew, A., & Cook, M. (1990). The control of reaching movements by young infants. Child Development, 64, 953-959.

McDonnell, P.M., & Abraham, W.C. (1979). Adaptation to displacing prisms in human infants. Perception, 8, 175-185.

Meer, A.L.H. van der, Weel, F.R. van der, Lee, D.N. (1994). Prospective control in catching by infants. Perception, 23, 287-302.

Mishkin, M., Ungerleider, L.G., & Macko, K.A. (1983). Object vision and spatial vision: Two cortical pathways. Trends in Neuroscience, 6, 414-417.

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

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

Saltzman, E. (1979). Levels of sensorimotor representation. Journal of Mathematical Biology, 20, 91-163.

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

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

Trevarthen, C. (1975). Growth of visuomotor coordination in infants. Journal of Human Movement Studies, 1, 57.

Trevarthen, C., Murray, L, & Hubley, P. (1981). Psychology of infants. In J.A. Davies & J. Dobbing (Eds.), Scientific Foundations of Pediatrics, pp. 211-274. London: William Heinemann.

Turvey, M.T., Fitch, H.L., & Tuller, B. (1982). The Bernstein perspective: I. The problems of degrees of freedom and context-conditioned variability. In J.A.S. Kelso (Ed.), Human motor behavior: An introduction, pp. 239-252. Hillsdale, NJ: Lawrence Erlbaum Associates.

Vereijken, B., Emmerik, R.E.A. van, Whiting, H.T.A., & Newell, K.M. (1992). Free(z)ing degrees of freedom in skill acquisition. Journal of Motor Behavior, 24, 133-142.

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

Author Note

Matthew Schlesinger and Domenico Parisi, Institute of Psychology, Rome; Jonas Langer, Department of Psychology.

This model was developed during the first author's visit to the Italian National Research Council in Rome, and was supported in part by a Fulbright Junior Research Grant. Subsequent support was provided by NSF IRI-9720345 and NSF CDA-9703217, at the University of Massachusetts. We wish to thank Stefano Nolfi, Daniele Denaro, Raffaele Calabretta, and the other researchers at the CNR for their help and encouragement, and Sascha Engelbrecht and Michael McCarty at UMASS for their comments and constructive feedback.

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

Appendix

We present here additional details of the econet model: (a) the network architecture and function, and (b) the learning algorithm.

Network Architecture and Function

General Features. The visual and proprioceptive input layers were fully connected to their respective intramodal hidden layers. Meanwhile, the tactile input unit, as well as the intramodal hidden units, were fully connected to the intermodal layer (see Figure 2). Similarly, the intermodal layer was fully connected to the output layer. The activation of the intramodal, intermodal, and output layer units was computed using the standard logistic function of the total weighted input. All bias terms were fixed at 0.

Sensory Input. There were 8 visual input units, each covering 8° of the visual field. Each visual input unit became active when the hand or object entered its respective portion of the visual field; otherwise, its activation level was 0. Note that because the arm workspace is 2-dimensional, only the front "edges" of the hand and object were visible. Consequently, the back sides of the object, as well as any occluded portions of the hand and object were not seen.

The visual field was divided into 64 "windows", each 1° wide. During each time step, each window was checked to see whether the target object or the econet's hand was inside it (the rest of the arm was made transparent to prevent it from filling the visual field). If a window was occupied by the hand or object, the distance from the econet to the nearest portion of the object in the window was calculated. This distance was then normalized by dividing it by 50 (i.e., the length of the arm), and then subtracting the result from 1; high values (i.e., close to 1) corresponded to near objects, while low values (i.e., close to 0) corresponded to objects at arm's length. Each visual input unit then received as activation the average of the eight values across its respective 1° windows.

In Figure 1, for instance, the hand is visible while the target object is outside the visual field. Given this sensory state, the vector of activation values for the 8 visual input units is <0 0 0 0 0 0.19 0.17 0>.

The activation of the proprioceptive input units was a function of the relative inner and outer angles of the shoulder and elbow joints, respectively (see Figure 1). For the upper arm muscles, the (inner) angle of the shoulder joint was first scaled from 0 to 1. This value was then passed to the upper arm flexor proprioceptive input unit, while its complement (1 - value) was passed to the upper arm extensor input unit. The forearm proprioceptive input activation values were computed in the same manner, as a function of the elbow joint angle.

The tactile input unit took binary activation values: 1 when the "hand" (i.e., the last 5 units of the forearm segment) made contact with the object, and 0 otherwise.

Motor Output. The visual field/body-axis rotates a maximum of 15° left or right during a time step. The activations values of the visual field/body-axis motor unit was scaled from -15° to +15°, and this value was then used to update the subsequent position of the visual field. Similarly, the difference between each respective pair of arm output units was also scaled to the range of -15°, +15°. For example, if the output of the upper arm flexor was 1 and the output of the upper arm extensor 0, a 15° rotation of the shoulder joint to the left was performed. Relative joint movements were not permitted beyond the upper and lower joint limits (i.e., 0° and 180°, see Figure 1). However, since the arm segments were hierarchically linked, absolute movements could occur for a full 360° (e.g., the upper arm could swing in a complete circle if "dragged" by the body-axis).

Learning Algorithm

For each generation, all connection weights were initially encoded using a binary representation (i.e., 8 bits per weight). When an individual econet was selected for reproduction, its connection weights were copied five times in binary form. However, during the copying process, 2% of the bits were randomly selected and switched (e.g., from xx1xxxxx to xx0xxxxx). Before testing, the connection weights were converted back to a real-valued representation.

Figure Captions

Figure 1. The econet in its 2-dimensional workspace. (Left and right edges of the visual field are shown, as well as the rotational limits of the shoulder and elbow joints.)

Figure 2. Diagram of the 4-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. Sensory input signals are determined at each time step as a function of the previous eye and arm movements.

Figure 3. Average fitness (3a), number of successful trials (3b), latency (3c), and number of separations (3d) in the normal, 3-layer, and random-movement training conditions (N = 10 runs for each condition). Note that average fitness in the random condition remains near zero.

Figure 4. Percent of shoulder flexion during learning, for each of the ten runs. Shoulder extensions predominated in eight of the ten runs, while flexions predominated in the other two (i.e., 1 and 7). In each run, N = 100 econets for each generation.

Figure 5. Percent of elbow flexion during learning, for each of the ten runs. There was a clear predominance for the elbow flexor in only one of the runs (i.e., 8). In each run, N = 100 econets for each generation.

Figure 6. Schematic diagram of the 2-dimensional activation space for an agonist-antagonist muscle pair. The diagonal lines correspond to ideal coactivation and reciprocal activation patterns. A sample muscle activation is presented, as well as the minimal deviation from each ideal activation pattern.

Figure 7. Average deviations from the ideal coactivation and reciprocal activation patterns, in the shoulder (7a) and elbow (7b) muscles, respectively (N = 10 runs). (Lower values represent a closer fit to the ideal activation pattern.)

Figure 8. Probability of an approach movement, during target fixation versus non-fixation in the normal condition, and during fixation in the random-movement condition (N = 10 runs for each condition).

Figure 9. Average speed of hand (9a) and head/eye (9b) movement during fixation of the target, in the normal and random-movement conditions (N = 10 runs for each condition).

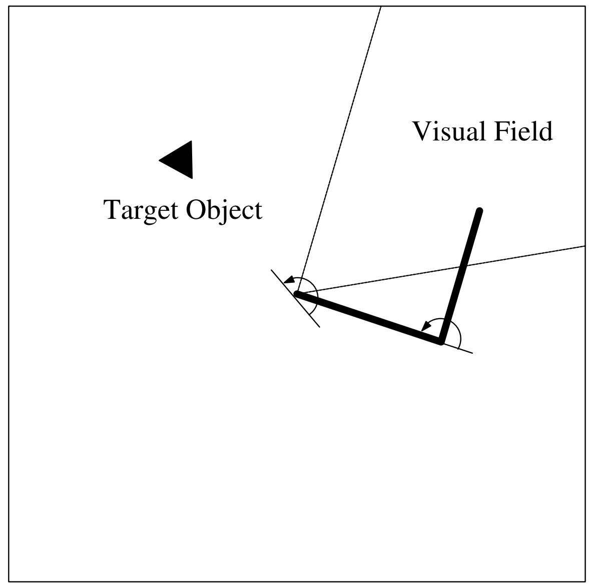

Visual Field
Target Object

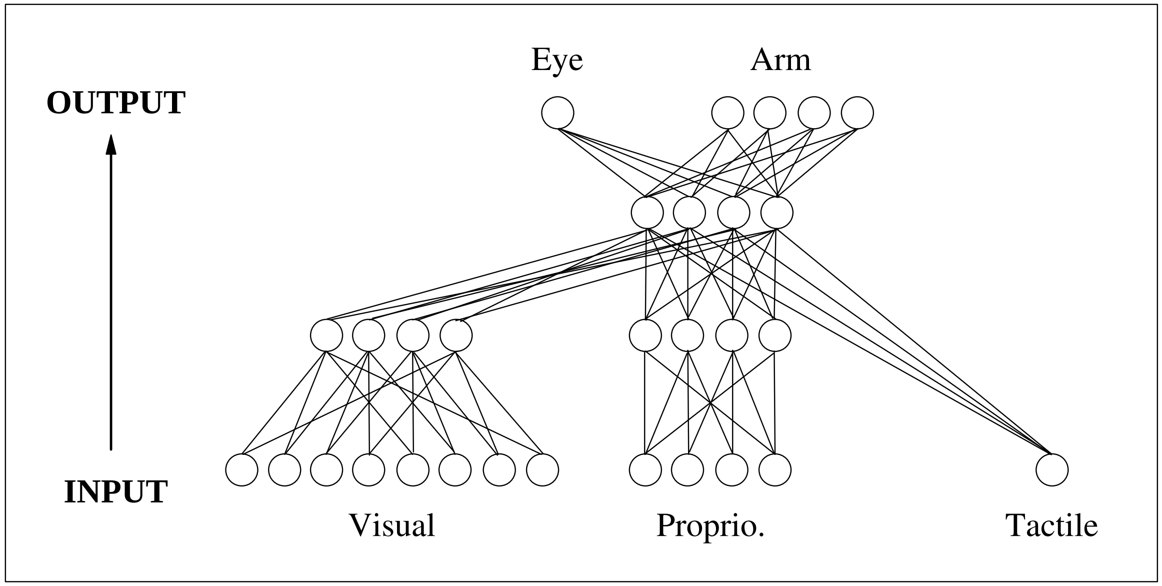

OUTPUT
INPUT
Eye
Arm
Visual
Proprio.
Tactile

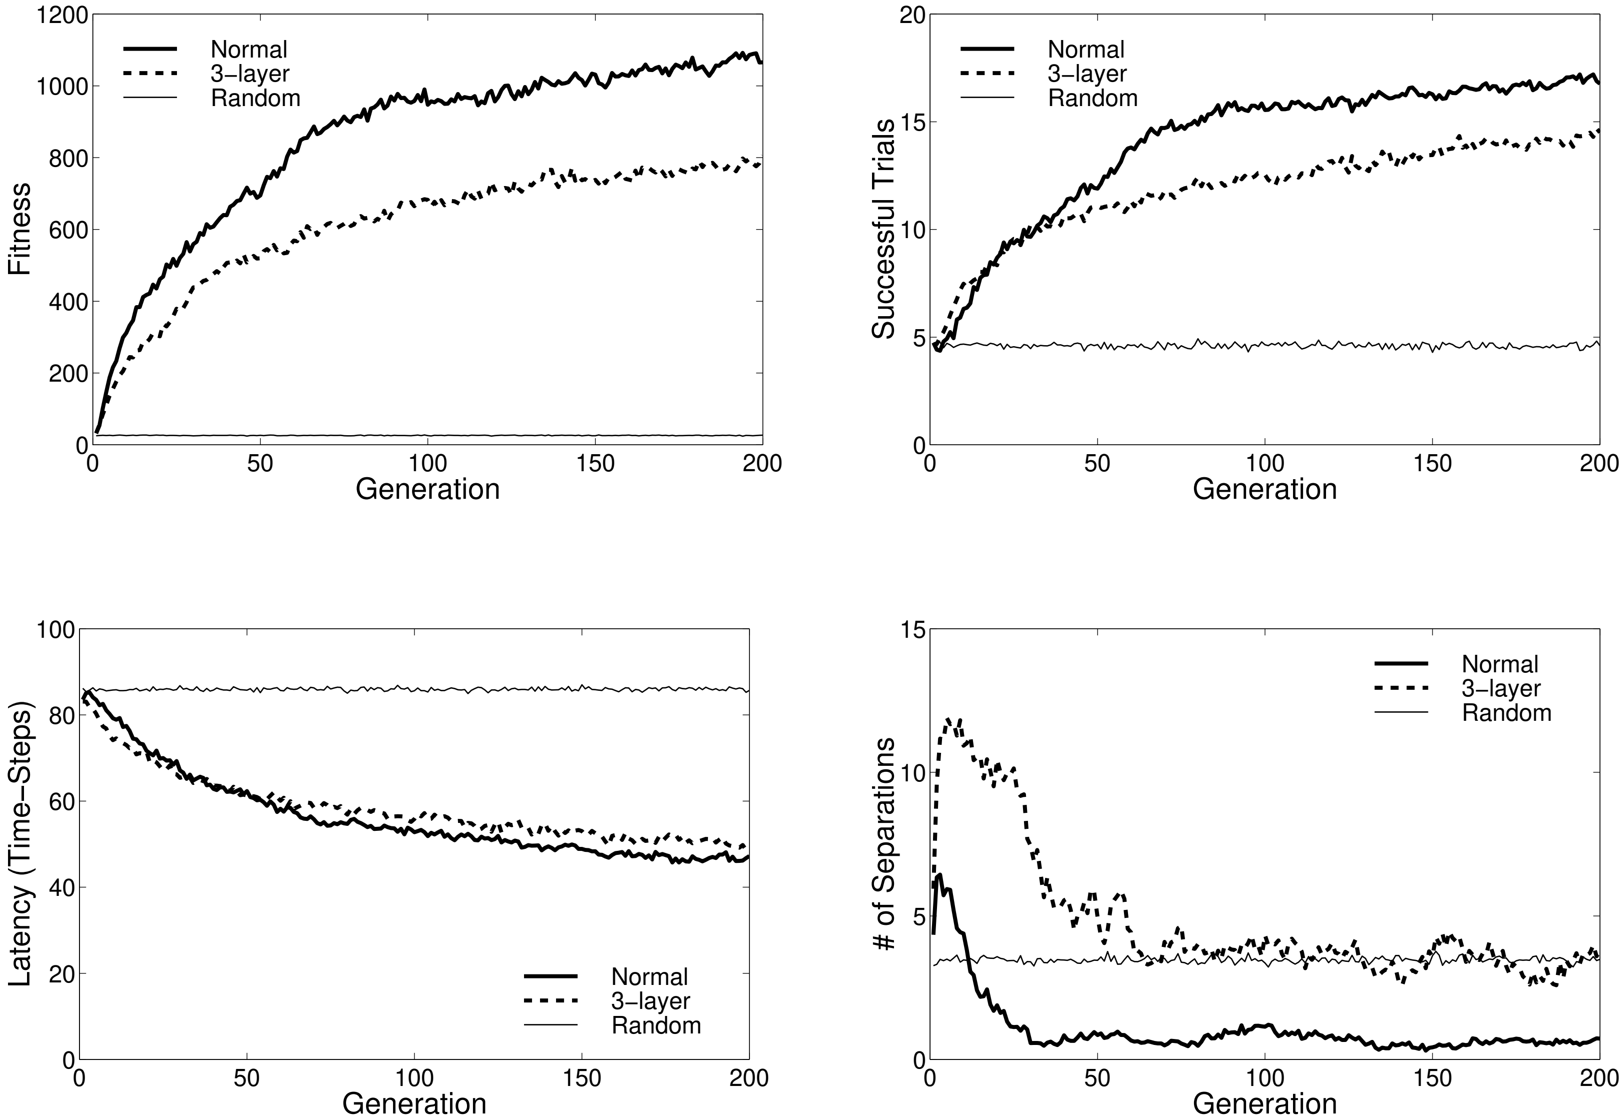

Fitness
Generation
Normal
3-layer
Random
Successful Trials
Generation
Normal
3-layer
Random
Latency (Time-Steps)
Generation
Normal
3-layer
Random
# of Separations
Generation
Normal
3-layer
Random

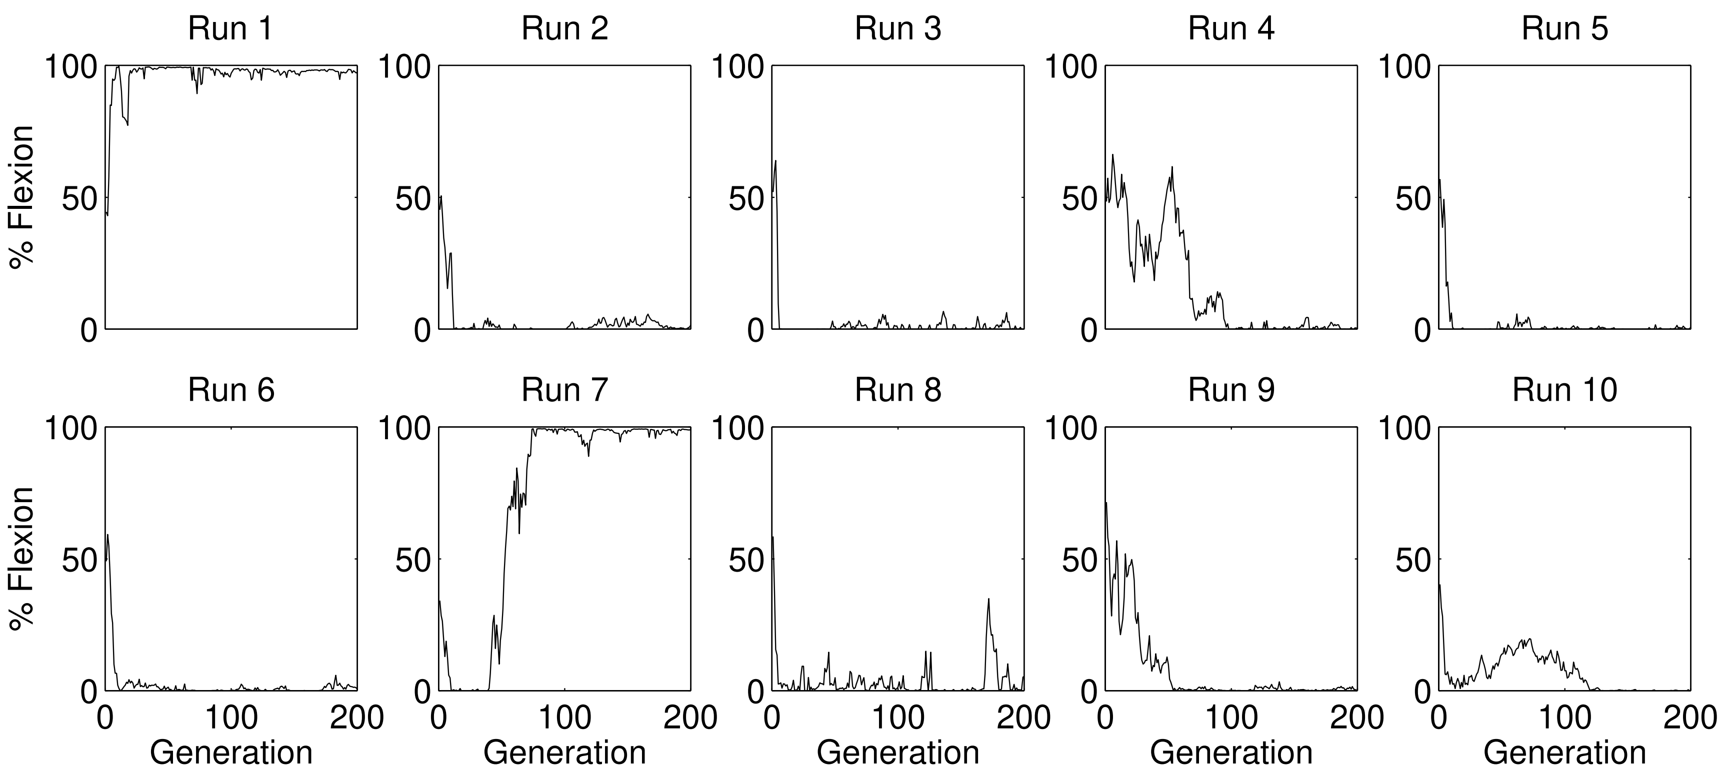
Run 1
Run 2
Run 3
Run 4
Run 5
Run 6
Run 7
Run 8
Run 9
Run 10
% Flexion
Generation
0
50
100
200

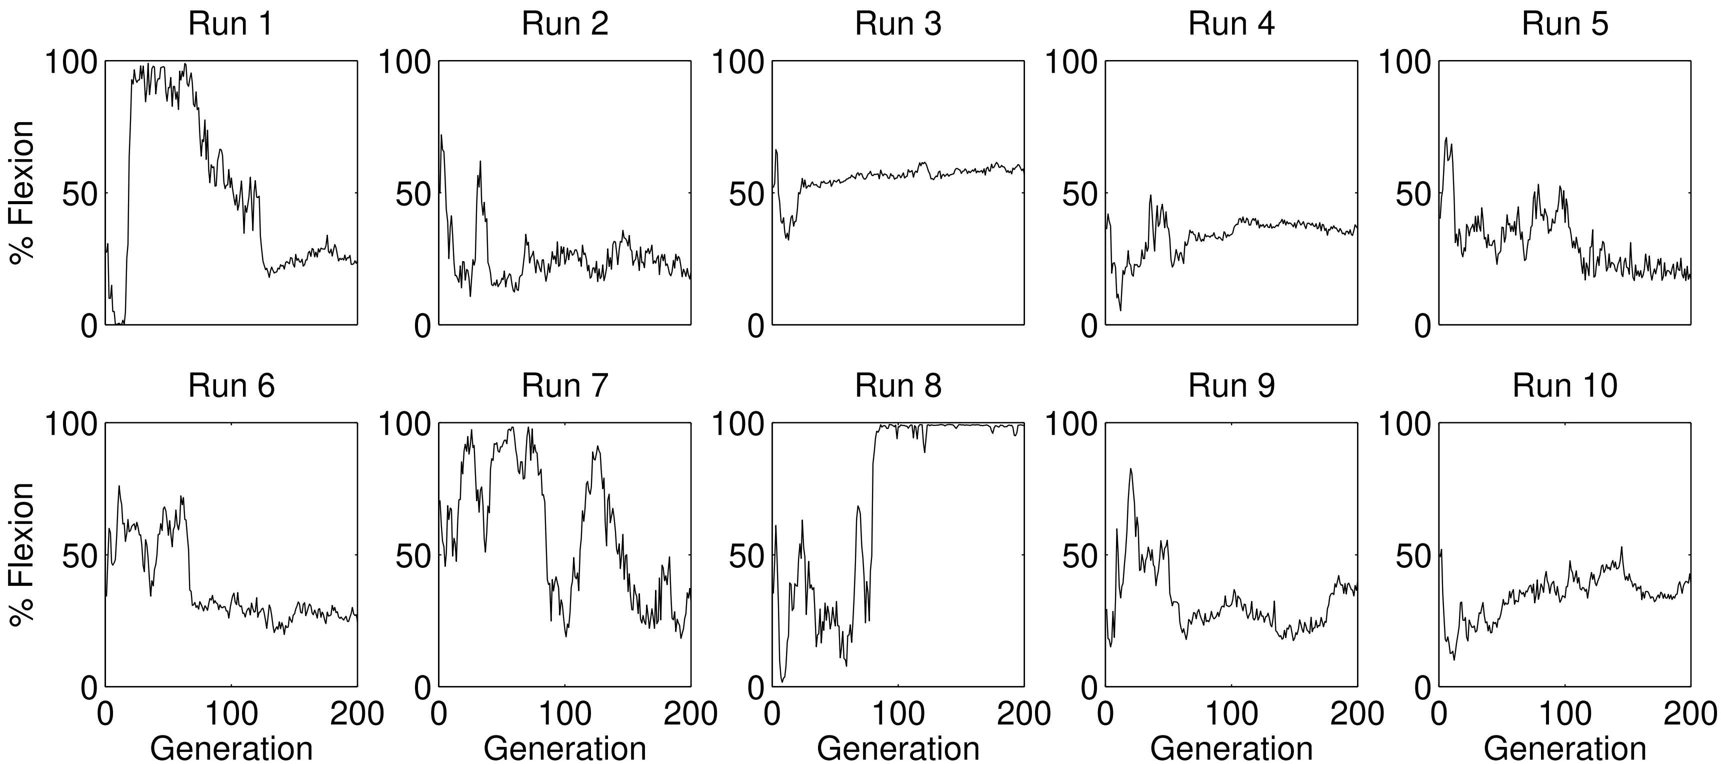
Run 1
Run 2
Run 3
Run 4
Run 5
Run 6
Run 7
Run 8
Run 9
Run 10
% Flexion
Generation
0
50
100
200

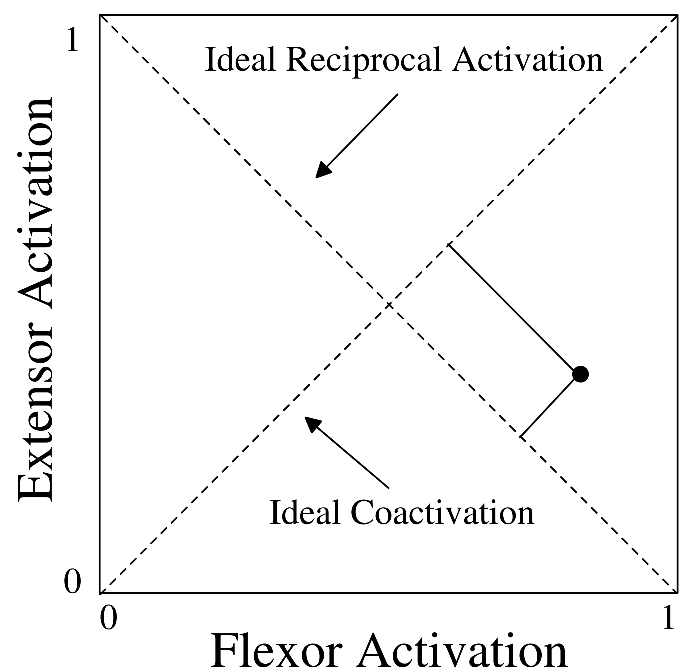
1
Ideal Reciprocal Activation
Extensor Activation
Ideal Coactivation
0
0
1
Flexor Activation

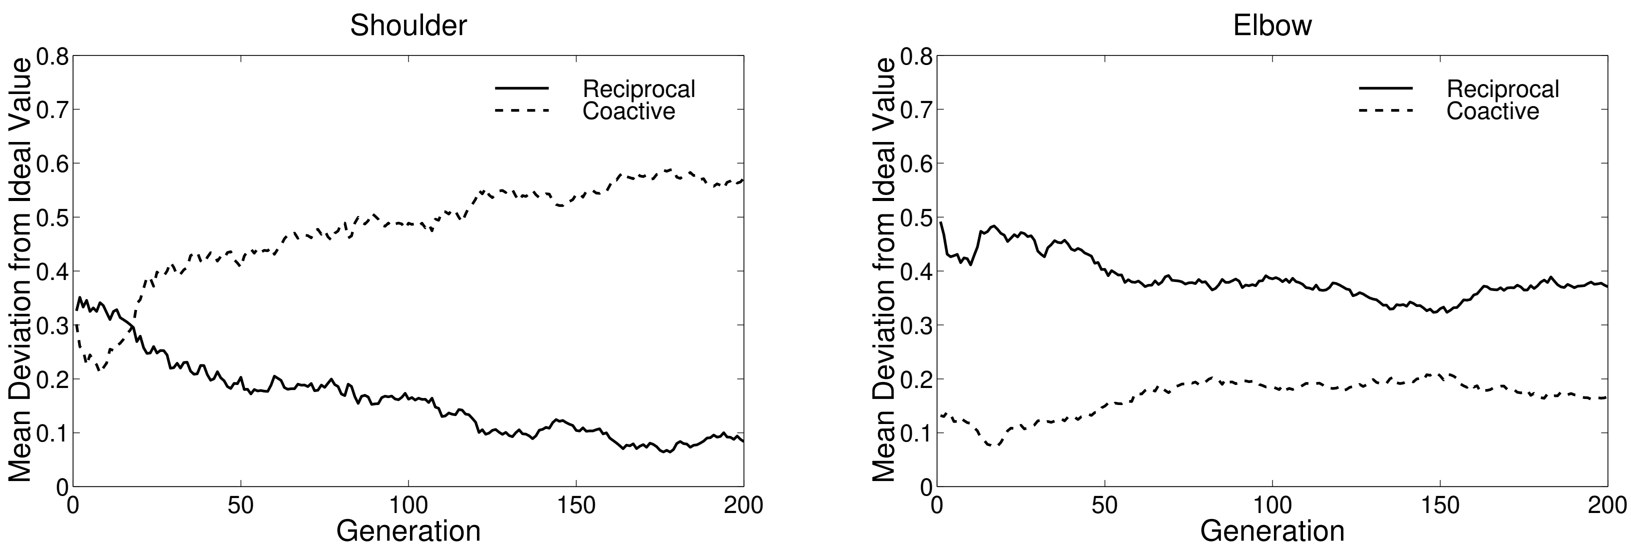

Shoulder
Mean Deviation from Ideal Value
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0
50
100
150
200
Generation
Reciprocal
Coactive
Elbow
Mean Deviation from Ideal Value
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0
50
100
150
200
Generation
Reciprocal
Coactive

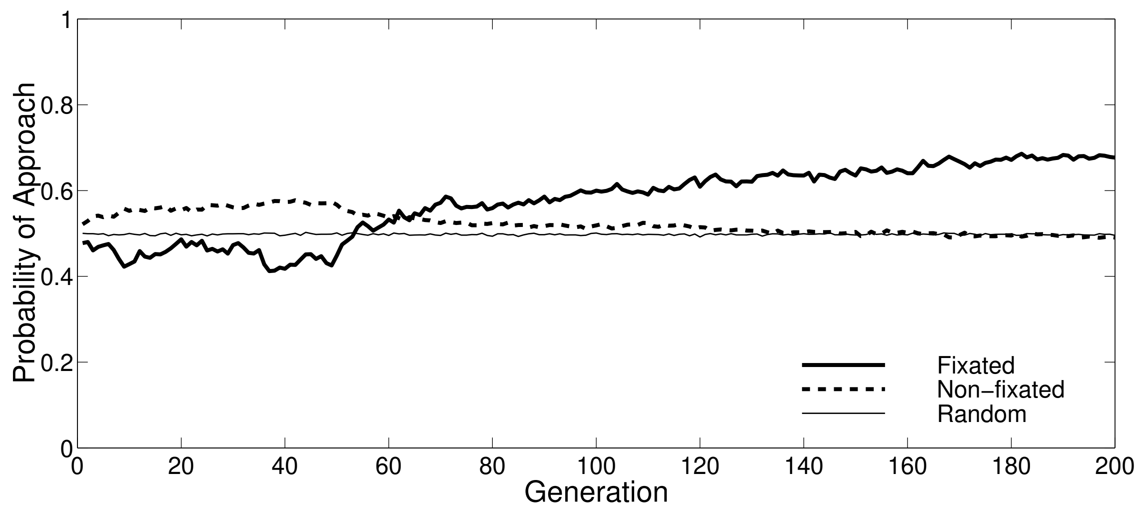
Probability of Approach
1
0.8
0.6
0.4
0.2
0
0
20
40
60
80
100
120
140
160
180
200
Generation
Fixated
Non−fixated
Random

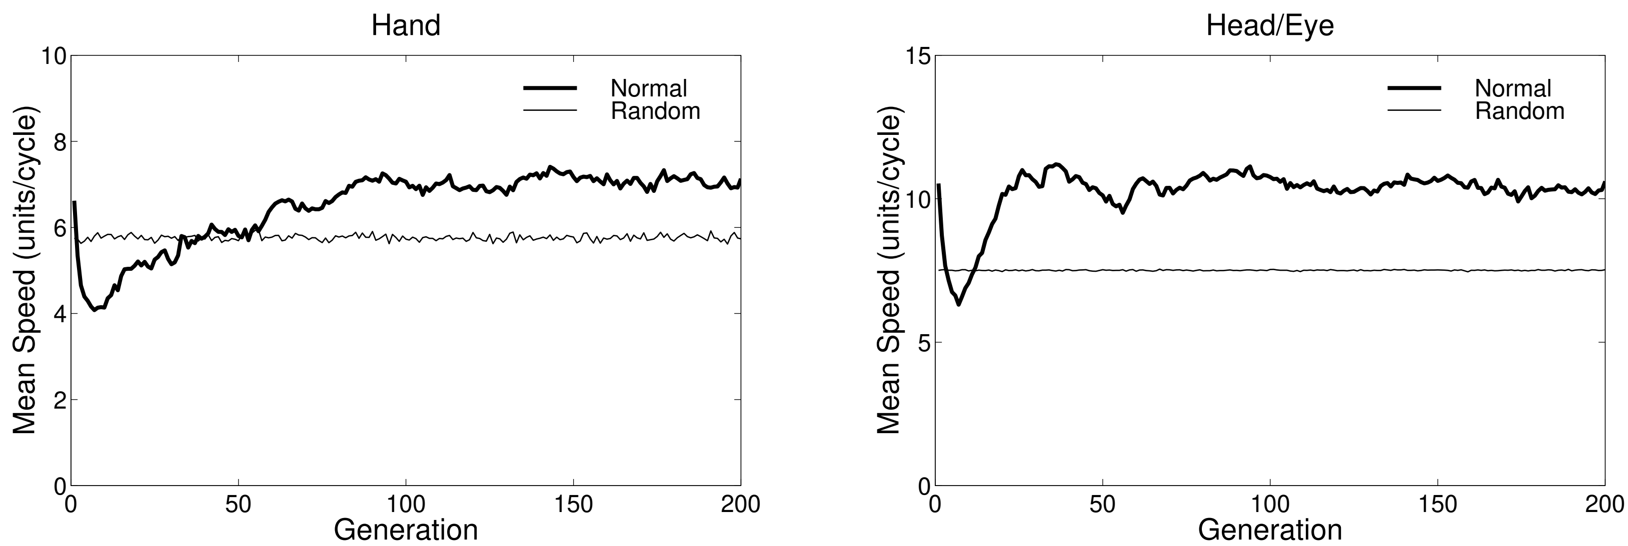
Hand
Normal
Random
Mean Speed (units/cycle)
Generation
0
2
4
6
8
10
0
50
100
150
200
Head/Eye
Normal
Random
Mean Speed (units/cycle)
Generation
0
5
10
15
0
50
100
150
200