

Internal Models and Contextual Cues: Encoding Serial Order and Direction of Movement

Stephanie K. Wainscott, Opher Donchin, and Reza Shadmehr

Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, Maryland

Submitted 10 March 2004; accepted in final form 19 September 2004

Wainscott, Stephanie K., Opher Donchin, and Reza Shadmehr. Internal models and contextual cues: encoding serial order and direction of movement. *J Neurophysiol* 93: 786–800, 2005. First published September 22, 2004; doi:10.1152/jn.00240.2004. During reaching, the brain may rely on internal models to transform desired sensory outcomes into motor commands. This transformation depends on both the state of the limb and the cues that can identify the context of the movement. How are contextual cues and information about state of the limb combined in the computations of internal models? We considered a reaching task where forces on the hand depended on both the direction of movement (state of the limb) and order of that movement in a predefined sequence (contextual cue). When the cue was available, the motor system formed an internal model that used both serial order and target direction to program motor commands. Assuming that the internal model was formed by a population code through a combination of unknown basis elements, the sensitivity of the bases with respect to state of the limb and contextual cue should dictate how error in one type of movement affected all other movement types. Using a state–space theory, we estimated this generalization function and identified the adaptive system from trial-by-trial changes in performance. The results implied that the basis elements were tuned to direction of movement but output of each basis at its preferred direction was multiplicatively modulated by a weak tuning with respect to the contextual cue. Activity fields that multiplicatively encode diverse sources of information may serve as a general mechanism for a single network to produce context-dependent motor output.

INTRODUCTION

One way to link activity fields (e.g., receptive fields or motor fields) of neurons in a particular region of the brain to perception or action is to assume that these fields are akin to mathematical bases with which a function is computed. For example, during a reach, some neurons in motor regions of the brain exhibit a tuning with respect to movement direction φ (Georgopoulos et al. 1982). In population coding of movement direction (Georgopoulos et al. 1986), tuning of each cell is a basis function $g_i(\varphi)$, and a weighted combination provides an estimate of direction. This “internal model” simply produces an identity map, $\varphi \rightarrow \hat{\varphi}$, but with different weighting functions any other function can be modeled (Pouget et al. 2000). This perspective has implications for our understanding of motor learning because the shape of the basis functions will determine how training in one part of the space generalizes to another part of the space (Pouget and Sejnowski 1997; Thoroughman and Shadmehr 2000). The tuning of neurons can be

very different in different parts of the brain. This means that the generalization patterns recorded in a task may be related to the tuning of the neurons in the parts of the brain that are involved in learning to compute the function of interest (Poggio et al. 1992).

In control of limb movements, it has been suggested that the brain computes a function that approximates inverse dynamics of the limb $(\theta, \dot{\theta}, \ddot{\theta}) \rightarrow \hat{f}$, estimating the forces \hat{f} that are necessary to achieve a particular desired limb state $\theta, \dot{\theta}, \ddot{\theta}$ (Conditt and Mussa-Ivaldi 1999; Shadmehr and Mussa-Ivaldi 1994). Assuming that the approximation of inverse dynamics is linear, one can represent this computation with the expression $\hat{f} = Wg(\theta, \dot{\theta}, \ddot{\theta})$, where $g = [g_1, g_2, \dots, g_n]^T$ is a vector of bases that encodes state of the limb and $W = [p_1 | p_2 | \dots | p_n]$ associates each basis with a preferred force vector p_i . In theory, movement errors produce a change in the preferred force vector of each basis, which in turn produces a performance change in the subsequent movement. If the subsequent movement explores a different part of the state space than the effect of the error on this new part of the state space is a measure of generalization.

Thus, theory predicts that there is a relationship between trial-by-trial change in performance and the shape of the bases (Donchin et al. 2003; Thoroughman and Shadmehr 2000). For example, patterns of generalization have been used to show that bases elements are tuned to movement direction (Donchin et al. 2003); that directional tuning is multiplicatively modulated by a linear function of limb position (Hwang et al. 2003); and that the encoding of limb position and velocity is in terms of the intrinsic coordinates of muscles and joints (Malfait et al. 2002; Shadmehr and Moussavi 2000). These are also properties of some cells in the primary motor cortex (Georgopoulos et al. 1984) and the cerebellum (Bosco et al. 1996; Coltz et al. 1999). However, there has been very little exploration of the representation of cues other than the current state of the limb. For example, if one needs to reach a target while holding a tennis racquet in one instance and a badminton racquet in another instance, a contextual cue must predict different forces for similar states of the limb (Osu et al. 2004; Wada et al. 2003). Our ability to interact with various objects in our environment predicts that contextual cues must have a significant influence on the activity fields of the bases. How are “high level” contextual cues and “low level” information about the state of the limb combined in the bases that compute internal models of dynamics? Here we use a system identifi-

Address for reprint requests and other correspondence: R. Shadmehr, Johns Hopkins School of Medicine, 419 Traylor Bldg., 720 Rutland Ave., Baltimore, MD 21205 (E-mail: reza@bme.jhu.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

cation technique to quantify generalization patterns across a multidimensional space that includes both contextual cues and state of the limb.

This work is part of an MS thesis submitted by S. K. Wainscatt to Johns Hopkins Biomedical Engineering Department.

THEORY

In our task, subjects reach to targets while interacting with a force field that depends on both the state of the limb (its velocity) and the serial order of that movement within a sequence. Assume that this task requires computation of a function $(s, \theta) \rightarrow \hat{f}$, where s is a contextual cue, θ is limb state (its velocity), and \hat{f} is an estimate of force on the limb. The cue is the ordinal number of a movement within a sequence. We assume that the map is computed with a collection of bases g_i , each a scalar function representing the activity of the basis element in different parts of the input space (s, θ) . Each basis is associated with a preferred force vector p_i . Adaptation results in changes in these preferred force vectors. We have $\hat{f} = Wg(s, \theta)$, where $g = [g_1, g_2, \dots, g_m]^T$ and $W = [p_1, p_2, \dots, p_m]$. \hat{f} and p_i are 2×1 vectors. If in trial n a reaching movement is made, the force error in that trial is

$$\tilde{f}^{(n)} \equiv f^{(n)} - \hat{f}^{(n)}$$

where $f^{(n)}$ is the actual force (e.g., at max velocity) in that trial. The “input state” in trial n is defined as $\Omega^{(n)} \equiv [s^{(n)}, \theta^{(n)}]$, where $\theta^{(n)}$ is evaluated at maximum velocity during that trial. Error in trial n results in a change in the preferred force vector associated with each basis. If we define scalar quantity $e \equiv \frac{1}{2} \tilde{f}^T \tilde{f}$, minimizing this cost function describes how error in trial n should produce a change in p_i

$$\begin{aligned} \frac{de}{dp_i} &= -\tilde{f}g_i \\ p_i^{(n+1)} &= p_i^{(n)} + \eta \tilde{f}^{(n)} g_i(\Omega^{(n)}) \\ W^{(n+1)} &= W^{(n)} + \eta \tilde{f}^{(n)} g^T(\Omega^{(n)}) \end{aligned} \quad (1)$$

In Eq. 1, η is a learning rate. Following Thoroughman and Shadmehr (2000), we multiply both sides of the above expression by $g(\Omega^{(n+1)})$ and note that $g^T(\Omega^{(n)})g(\Omega^{(n+1)})$ is a scalar quantity and arrive at

$$\begin{aligned} \hat{f}^{(n+1)} &= W^{(n+1)}g(\Omega^{(n+1)}) \\ &= W^{(n)}g(\Omega^{(n+1)}) + \eta g^T(\Omega^{(n)})g(\Omega^{(n+1)})\tilde{f}^{(n)} \end{aligned} \quad (2)$$

Therefore, if trial n occurred in state $\Omega^{(n)}$ and trial $n + 1$ occurred in state $\Omega^{(n+1)}$, the error experienced in trial n produces the following *observable* change in the system when the system is evaluated at $\Omega^{(n+1)}$

$$\hat{f}^{(n+1)}(\Omega^{(n+1)}) - \hat{f}^{(n)}(\Omega^{(n+1)}) = \eta g^T(\Omega^{(n)})g(\Omega^{(n+1)})\tilde{f}^{(n)} \quad (3)$$

Because of the error in trial n , the system might have changed everywhere, but because in each trial only one state is observed (e.g., a movement is made to only one direction), the rest of these changes (e.g., forces that are expected for other directions) are not observable. In Eq. 3, the term $g^T(\Omega^{(n)})g(\Omega^{(n+1)})$ is a *generalization function*. This scalar function determines

how much of the error experienced in one state affects any other state.

Our discussion assumed that in a given trial (a movement), the internal model computed a single vector of force. This is a gross simplification, and elsewhere we have derived Eq. 3 with the assumption that a trial is a trajectory of forces and states (Donchin et al. 2003). In that case, the generalization function takes the form of an integral. The result is that generalization from the state in which error was experienced to another state is a correlation between the basis elements as evaluated along the 2 states. When the basis elements are broad, this correlation is high and there is large generalization. When the basis elements are narrow, the correlation quickly drops off as a function of distance between the states.

To give a simplified example, ignore the fact that there may be contextual cues in the task and imagine that our internal model consisted of 6 states, each a direction of movement φ . On any given trial, the particular values of p_i (i.e., preferred force vector) associated with each g_i (i.e., basis function) constitute the motor system’s memory of the task. If we evaluate this memory for all states, the result is a field of forces \hat{F} , as shown in Fig. 1A1. Therefore, at trial n , $\hat{F}^{(n)}$ is a field (represented as a 12×1 vector) that contains a force $\hat{f}^{(n)}$ (a 2×1 vector) for each of the 6 states. The experimenter provides a target and the subject evaluates $\hat{F}^{(n)}$ along that direction and predicts $\hat{f}^{(n)}$. We write this as

$$\hat{f}^{(n)} = K^{(n)}\hat{F}^{(n)}$$

where $K^{(n)}$ is a sparse matrix (2×12) that selects the force along the desired direction from the field \hat{F} . Because the movement is made to that direction, the experimenter has the robot produce a force $f^{(n)}$. The robot force interacts with force produced by the subject, resulting in hand trajectory. Suppose among the points along this trajectory we pick hand position at peak velocity and label it $y^{(n)}$. If there was a difference between $f^{(n)}$ and $\hat{f}^{(n)}$, there will be a movement error

$$\tilde{y}^{(n)} \equiv y^{(n)} - y^*$$

where y^* is the reference hand position that the subject would ideally like to achieve. Because there are 6 possible directions of movement, we use the 12×1 vector Y^* to represent the 6 reference positions and the sparse matrix $L^{(n)}$ (2×12) to select the appropriate reference position for that trial

$$y^* = L^{(n)}Y^*$$

Thus $\tilde{y}^{(n)}$ depends on $\hat{f}^{(n)}$ and the actual forces in that trial $f^{(n)}$. Because errors are generally small (typically < 1.5 cm), we assume that this dependency is linear

$$\tilde{y}^{(n)} \equiv y^{(n)} - y^* = D(f^{(n)} - \hat{f}^{(n)})$$

where D is a 2×2 matrix, and it measures how far the arm will be displaced in different directions for one unit of force error. In this sense, if the arm has some springlike properties during movement, D is the inverse of the strength of the spring. Another way to think of this is that D is the inverse of the gain of an online feedback system that corrects errors during the movement.

In the example of Fig. 1A2, error $\tilde{y}^{(n)}$ was experienced along state $\varphi^{(n)} = 90^\circ$, but because of generalization, it affected all states $i = 1 \dots 6$. Generalization is a scalar function b that

depends on the distance between the state in which the error was experienced $\varphi^{(n)}$ and all other states $\varphi^{(i)}$

$$b(\varphi^{(i)} - \varphi^{(n)}) \equiv \eta g^T(\varphi^{(i)})g(\varphi^{(n)}) \quad \forall i$$

Accordingly, if there are q directions of movement, there are q distances between these directions and the generalization function can be parameterized with q parameters. The field of forces in trial $n + 1$ is described as

$$\hat{F}^{(n+1)}(\varphi^{(i)}) = \hat{F}^{(n)}(\varphi^{(i)}) + b(\varphi^{(i)} - \varphi^{(n)})\tilde{f}^{(n)} \quad i = 1 \dots 6$$

For an arbitrary generalization function, the resulting field is shown in Fig. 1A3. Generalization affected the field for every direction. It is convenient to write the above equations in a more compact form

$$\hat{F}^{(n+1)} = \hat{F}^{(n)} + BK^{(n)T}\tilde{f}^{(n)}$$

where B is a 12×12 matrix that contains the 6 parameters of the generalization function (because there are 6 directions of movement, the function b is evaluated at 6 discrete angular distances)

$$M = \begin{bmatrix} b(0) & b(5\pi/3) & \dots & b(\pi/3) \\ b(\pi/3) & b(0) & \dots & b(2\pi/3) \\ b(2\pi/3) & b(\pi/3) & \dots & b(\pi) \\ \vdots & \vdots & \ddots & \vdots \\ b(5\pi/3) & b(4\pi/3) & \dots & b(0) \end{bmatrix}$$

$$B \equiv \begin{bmatrix} M & 0 \\ 0 & M \end{bmatrix}$$

In the above description for matrix B , 0 represents a 6×6 matrix where all elements are zero.

In an ideal adaptive system, transition from trial to trial is dictated by the error in the previous trial and the generalization function. However, other factors may also affect the system. For example, memory may decay from one trial to the next. We can represent this effect with another parameter a

$$\hat{F}^{(n+1)} = a\hat{F}^{(n)} + BK^{(n)T}\tilde{f}^{(n)}$$

Figure 1B provides an example of this system. In this example, we have only one direction of movement but the sequence consists of 2 movements (odd and even, i.e., $s = 1$ or $s = 2$) and only 2 states. To perform movement n , we have $\hat{f}^{(n)}$, resulting in trajectory error $\tilde{y}^{(n)}$. $\tilde{f}^{(n)}$ updates contents of both states by a generalization function b . Because we have 2 states, we assume that b can be represented with 2 parameters: distance between movements of same ordinal number (i.e., 0) and distance between movements of different ordinal number (i.e., 1).

To compute a map that depends on both the direction of movement and the order of that movement in a sequence, the bases must be sensitive to both variables. Therefore, state of a trial is $\Omega^{(n)} \equiv [s^{(n)}, \varphi^{(n)}]$. Distances in “sequence” space are not




FIG. 1. System identification of internal models. A: in this example, the internal model consists of 6 states, each a direction of movement. (1) If the internal model could be evaluated at all its states, it would produce a field of forces $\hat{F}^{(n)}$, i.e., a snapshot of the contents of the internal model on trial n . (2) However, on trial n , the content of only one state can be (indirectly) measured by making a movement toward a target. This produces a trajectory $y^{(n)}$ that, in comparison to a reference trajectory y^* , produces an error $\tilde{y}^{(n)}$. That error is related to a force error $\tilde{f}^{(n)}$ by an online error correcting gain D . (3) Error $\tilde{f}^{(n)}$ is broadcast to all directions by a scalar generalization function b , which depends on the distance of the state in which the error was experienced and the state that is updated. After all states are updated, the result is $\hat{F}^{(n+1)}$ (drawn in black). For comparison, $\hat{F}^{(n)}$ is drawn in gray. B: in this example, movements are performed in only one direction, but there are 2 movements in a sequence and therefore movement order is a contextual cue. There are 2 hidden states, $\hat{f}_o^{(n)}$ and $\hat{f}_e^{(n)}$, expressing the force predicted for an odd-numbered and an even-numbered movement. \hat{f} interacts with the actual force in that movement $f^{(n)}$, producing an error $\tilde{f}^{(n)}$. Force error interacts with D and produces a movement error $\tilde{y}^{(n)}$. $\tilde{f}^{(n)}$ updates both hidden states with a generalization function b , which is a function of the distance between the state in which the error was experienced and the state that is updated (in this case, the distance is either 0 or 1). C: when the internal model is represented as a population code by a set of basis elements, the generalization function b depends on the shape of the bases. Here, the task depends on both movement direction and order within a sequence and therefore the bases encode both variables. Figure shows the generalization functions when the bases encode these 2 variables additively vs. multiplicatively. Multiplicative nature of generalization is also a feature of a mixture of experts.

equal to distances in “direction” space. As a result the generalization function will be parameterized as

$$b(|s^{(i)} - s^{(j)}|, \varphi^{(i)} - \varphi^{(j)}) \equiv \eta g^T(s^{(i)}, \varphi^{(i)}) g(s^{(j)}, \varphi^{(j)}) \quad (4)$$

For example, when there are 2 movements in a sequence and 6 directions of movement, the generalization function b will be represented with 2×6 parameters. Matrix B becomes a 24×24 matrix where each entry is one of these 12 parameters. Alternatively, one can dispense with the assumption of symmetry; that is, one can assume that distance from state 1 to state 2 is not the same as the distance between state 2 to state 1. In that case, for q states the generalization function will be parameterized with q^2 number of parameters. Here we assumed that distances could be symmetrically measured between states.

In summary, an internal model that learns to compute a function $(s, \theta) \rightarrow \hat{f}$ may be thought of as a dynamical system with many hidden states. At any given time, the experimenter can indirectly examine the contents of only one of those states by having the subject make a movement. The result is a movement that may have some error. The error is *broadcast* to all states. The effect of the broadcast is observed as a generalization function that depends on the tuning of the bases with respect to those states. We represent this process with the following dynamical system

$$\begin{aligned} \hat{y}^{(n)} &= D(f^{(n)} - K^{(n)}\hat{F}^{(n)}) + L^{(n)}Y^* \\ \hat{F}^{(n+1)} &= a\hat{F}^{(n)} + BK^{(n)T}(f^{(n)} - K^{(n)}\hat{F}^{(n)}) \end{aligned} \quad (5)$$

where $\hat{y}^{(n)}$ is the model’s estimate of hand position in trial n . In that trial, the robot produced a force $f^{(n)}$, and we measure the resulting hand position $y^{(n)}$. For that trial we also know the “state” of the task [i.e., $L^{(n)}$] and $K^{(n)}$. The rest of the variables are unknown and will need to be estimated by fitting the $\hat{y}^{(n)}$ to the measured data $y^{(n)}$.

HYPOTHESIS

There are a number of ways to learn a task that depends on both the ordinal number of a movement (s) within a sequence and movement velocity θ . We considered 3 ways of encoding this information: a model with separate basis functions for direction and serial order that interacted additively, a model that combined coding of direction and serial order multiplicatively, and a mixture of 2 models resembling a mixture of experts. Here we show the patterns of generalization that each type of encoding predicts.

First consider an additive model that uses separate bases for encoding ordinal and velocity information. A weighted sum of the bases approximates the function of interest

$$\hat{f}^{(n)} = \sum_i w_{i,1}^{(n)} g_{1,i}(s^{(n)}) + w_{2,i}^{(n)} g_{2,i}(\theta^{(n)})$$

To compute the generalization properties of this system, we find the derivative of $\hat{f}^{(n)}$ with respect to the weights and write an expression like Eq. 3 (i.e., an estimate of forces in the next trial as a function of the errors in the previous trial). If movement n was in state $[s^{(n)}, \theta^{(n)}]$, then generalization from that state to the next movement in the same direction will be $b(1, 0) = \eta \sum_i g_{1,i}(s^{(n)}) g_{1,i}(s^{(n+1)}) + g_{2,i}^2(\theta^{(n)})$. We note that the term $b(0, 0) - b(1, 0)$ is independent of velocity. Therefore, in

this model the differences in the serial order of movements in the sequence will have an additive effect on the generalization function. A second point to note is that with an additive model, it will not be possible to compute a function where movement velocity and serial order are combined nonlinearly in the production of force. We will be presenting our subjects with such a nonlinear function.

To combine both cues in each basis functions, we can have one variable act as a gain to modulate sensitivity of the other variable

$$\hat{f}^{(n)} = \sum_i w_i^{(n)} g_{1,i}(s^{(n)}) g_{2,i}(\theta^{(n)})$$

Generalization from the movement state in trial n to the next movement in the sequence in the same direction will be: $b(1, 0) = \eta \sum_i g_{1,i}(s^{(n)}) g_{1,i}(s^{(n+1)}) g_{2,i}^2(\theta^{(n)})$. Generalization to the same movement number in the sequence and the same movement direction will be: $b(0, 0) = \eta \sum_i g_{1,i}^2(s^{(n)}) g_{2,i}^2(\theta^{(n)})$. Unlike the additive model, in this model the term $b(0, 0) - b(1, 0)$ is not independent of velocity. Rather, the change in the contextual cue multiplicatively modulates the generalization function with respect to direction, and so we can call this the multiplicative model.

In a mixture of experts (Ghahramani and Wolpert 1997), for each possible ordinal position in the sequence there is a separate basis (or expert) that encodes velocity during those movements. A separate model selector chooses which of the experts to use based on the serial order of the current movement. For example, if there are 2 movements in a sequence, we have

$$\hat{f}^{(n)} = r^{(s^{(n)})} \sum_i w_{1,i}^{(n)} g_i(\theta^{(n)}) + \{1 - r^{(s^{(n)})}\} \sum_i w_{2,i}^{(n)} g_i(\theta^{(n)})$$

It turns out that the generalization pattern of the mixture of experts model is identical to the multiplicative model if the encoding of movement order is similar. Therefore, both the multiplicative model and the mixture of experts predict that the generalization function should exhibit a multiplicative interaction between movement order and movement direction.

For example, imagine that each basis has a preferred movement number s_i and a preferred velocity θ_i . The coding of movement number may be with a Gaussian

$$g_{i,1}(s) \equiv \exp\left(\frac{-|s - s_i|^2}{2\sigma_s^2}\right) \quad (6)$$

To represent encoding of velocity, one can also assume a Gaussian with a center randomly positioned in limb velocity space. However, our earlier work suggests that a better guess for velocity encoding may be a bimodal function (Donchin et al. 2003)

$$g_{i,2}(\theta) \equiv c + \exp\left(\frac{-|\theta - \theta_i|^2}{2\sigma_\theta^2}\right) + \frac{1}{k} \exp\left(\frac{-|\theta + \theta_i|^2}{2\sigma_\theta^2}\right) \quad (7)$$

To compute a generalization function, we assumed that there were 2 movements in a sequence. We distributed the preferred velocity θ_i uniformly in the 2D space (with maximum values at ± 0.7 m/s). We assumed that movement n was in direction $\varphi^{(n)} = 90^\circ$ with ordinal value $s^{(n)} = 1$. We evaluated the bases at peak hand velocity typical for a reaching movement (0.3 m/s), and then computed the generalization function $b[|s^{(i)} - s^{(n)}|, \varphi^{(i)} - \varphi^{(n)}]$ for all possible movement directions $\varphi^{(i)} = 30, 90, \dots, 330^\circ$, and ordinal values $s^{(i)} = 1$ and 2.

Figure 1C shows the generalization functions for the additive and multiplicative models. The free parameters in the simulations were c , k , σ_q , and σ_s . The values for this simulation were $c = 0.06$, $k = 3.5$, $\sigma_q = 0.15$, and $\sigma_s = 0.8$. The solid line is $b(0, \Delta\varphi)$ (i.e., generalization to the same ordinal number but potentially different direction). The dashed line is $b(1, \Delta\varphi)$. In both models, generalization as a function of direction is bimodal, which is a property of the bases (Eq. 7). However, note that in the additive model, when ordinal distance is 1, the generalization function is shifted down by a constant value, whereas in the multiplicative model, the generalization function is scaled. Therefore, in the multiplicative encoding of sequence and direction, a gain modulates the generalization function. This property is not observed in the additive model.

METHODS

Healthy volunteers ($n = 30$) performed a series of reaching movements to targets at 10 cm while holding the handle of a robotic arm (Fig. 2A). Targets were presented on an LCD monitor positioned approximately 50 cm in front of the subject at eye level. A *trial* was defined as a reach toward a target. On trial n , a force field was produced that depended linearly on hand velocity but nonlinearly on ordinal position of the movement in a 2-movement sequence (i.e., whether the movement was odd or even)

$$f = (-1)^n \begin{bmatrix} 0 & -13 \\ 13 & 0 \end{bmatrix} \dot{x}$$

In this expression force has units of Newton and velocity has units of m/s. The direction-dependent forces flipped from a clockwise curl

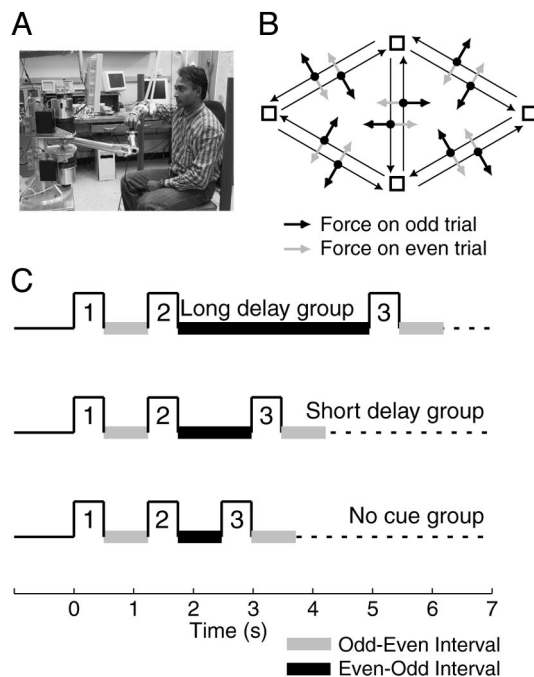


FIG. 2. Experiment setup. A: participants held the handle of a robot and reached to a series of targets displayed on a vertical monitor. B: targets, presented one at a time, were arranged in a diamond and were 10 cm apart. A velocity-dependent field imposed forces on the hand that were perpendicular to the direction of motion. Direction of forces depended on both the numerical order of the movement and the direction of movement. C: participants were divided into 3 groups (no-cue NC, short-delay SD, and long-delay LD). In the LD and SD groups, the relatively longer delay between specific movements served as the cue that the sequence was restarting.

pattern to a counterclockwise pattern in alternating trials. Movements were made in 6 directions and arranged in a pattern shown in Fig. 2B. Targets were arranged as a random walk about the nodes of this pattern. Therefore, for a given start position, there were at least 2 potential directions of movement.

Task description

The task began with presentation of a green target. This instructed the subject to reach. After completion of the reach, the next target was immediately presented in red. This instructed the subject to wait until the target turned green, and then the movement was performed. The trial was discarded if movement began before the target turned green. Upon completion of each reach, feedback was given to indicate movement success. The target would “explode” if the period of the movement was 0.5 ± 0.05 s. If the movement was either too long or too short, the target turned blue or red, respectively.

In our main groups of subjects, a pattern was imposed on the intertrial interval so that movements were temporally but not spatially organized into pairs: 1–2, 3–4, 5–6, 7–8, and so on. That is, after completion of an odd-numbered movement, there was a short time delay Δ_1 before the next target turned green. However, after completion of an even-numbered movement, there was a longer time delay Δ_2 . The long delay Δ_2 served as a cue that signaled start of the 2-movement sequence. In these groups, the subjects were coached to associate the movement sequence as pairs of movements. The coaching took the form of observing the experimenter perform movements in a null field while she sounded out “movement 1,” “movement 2,” “movement 1,” and so forth. Therefore, the subject had ordinal information about the movements. However, the direction of each movement was randomly chosen and could not be predicted.

Training took place in 2 consecutive days and was divided into target sets. Each set contained 240 trials, divided into 2 subsets of 120 trials. Subjects rested for a minute between the subsets. On Day 1, training began in a null field (robot in a passive mode) with 240 trials, and was then followed by 960 trials in the field. The field training continued on Day 2, where subjects trained in another 960 trials. For any given target, there was a 1/6 probability of a catch trial. Catch trials are movements for which the field is temporarily turned off.

Previous work on a similar task (Karniel and Mussa-Ivaldi 2002) had found that when $\Delta_1 = \Delta_2$ (i.e., no temporal cues were present to provide information about the sequential order of the force fields), subjects could not learn this task (as measured over a few hundred training trials). We hypothesized that if the temporal cues were to be a useful aide in adaptation, their effectiveness might be a function of their disparity: that is, learning may improve as a function of $\Delta_2 - \Delta_1$. To test this, we considered 3 conditions (Fig. 2C). In our control task, termed *no-cue group* (NC) ($n = 10$ subjects), the delays were equal (i.e., $\Delta_1 \approx \Delta_2$). This group was neither coached nor had temporal cues to identify ordinal structure of the task. Our main group was divided into 2 subgroups. In the *short-delay group* (SD) ($n = 10$), we had $\Delta_1 < \Delta_2$. In the *long-delay group* (LD) ($n = 10$), we had $\Delta_1 \ll \Delta_2$. The values for the intertrial intervals were as follows: Δ_1 for all 3 groups was 0.5 s; Δ_2 values for the no-cue, short-delay, and long-delay groups were 0.5, 1.0, and 3.0 s, respectively.

Pilot studies indicated that this was a difficult task that required many hundreds of trials. To encourage the participants, in addition to the target explosions we also provided a score that was continuously displayed at the bottom right of the video monitor. The score R was computed as

$$R^{(n)} = \frac{[5 \times R^{(n-1)} + |\Delta t^{(n)} - 0.5|]}{6}$$

where $\Delta t^{(n)}$ is movement time recorded for trial n . The score was for the participant’s encouragement and entertainment. Our performance measures did not consider it.

Performance measures

We focused on the errors that were caused by the force field and the errors that were produced in catch trials. For each movement we computed hand position at maximum hand velocity and found its distance from a straight line to the target. The result was labeled as PD, hand displacement perpendicular to the direction of target (PD). With learning, PD in field trials should decrease, whereas PD in catch trials should increase. Furthermore, the direction of PD in catch trials should be a function of the field appropriate for that trial. For example, for a given target direction, catch trial PDs measured on odd-numbered trials should be in the opposite direction of PDs for even-numbered catch trials. We labeled average PD in odd field trials as \overline{PD}_{fo} , in even field trials as \overline{PD}_{co} , in odd catch trials as \overline{PD}_{co} , and in even catch trials as \overline{PD}_{ce} . Because errors in both field and catch trials are important indicators of adaptation, we used a learning index (LI) that considered relative changes in these values (Criscimagna-Hemminger et al. 2003). We used the following equations to calculate an LI for odd- and even-numbered movements, as well as an overall learning index

$$\begin{aligned} LI_o &= \frac{\overline{PD}_{co}}{|\overline{PD}_{co}| + |\overline{PD}_{fo}|} \\ LI_e &= \frac{\overline{PD}_{ce}}{|\overline{PD}_{ce}| + |\overline{PD}_{fe}|} \\ LI &= \frac{\overline{PD}_{co} - \overline{PD}_{ce}}{|\overline{PD}_{co} - \overline{PD}_{ce}| + |\overline{PD}_{fo} - \overline{PD}_{fe}|} \end{aligned} \quad (8)$$

As the training begins, \overline{PD}_c is close to zero, so LI_o and LI_e are near zero. With training, LI_o and LI_e should converge toward 1 and -1 . The combined value, LI , should increase toward 1.

System identification procedure

We fit the system in Eq. 5 to runs of 240 trials of $y^{(n)}$. In our initial analysis, $y^{(n)}$ was generated by averaging the data across subjects for each movement. However, we also fit the model to $y^{(n)}$ generated by each subject. We report the resulting generalization function for both approaches.

For our 2 main groups, where movements consisted of 6 directions and 2 ordinal numbers within the sequence, Eq. 4 produced a generalization function b with 12 parameters. For the control group (NC), where no information about sequencing was available, a 6-parameter generalization function was sufficient but we continued to use the 12-parameter model as a sanity check, expecting that sequence information should be irrelevant for the b in the NC group.

The only observable quantities are $y^{(n)}$ (i.e., hand position at maximum velocity recorded for trial n) and $f^{(n)}$, the force produced by the robot at maximum velocity. For each run of 240 trials, the unknown parameters are: a , b , D , $\hat{F}^{(1)}$, and Y^* , where a is a scalar quantity and represents trial-to-trial memory loss in the system; b is the generalization function and contains 12 unknown parameters; D is the gain of online feedback during each trial and contains 4 unknown parameters; and $\hat{F}^{(1)}$ is the internal model's field of forces at the very first trial (i.e., initial condition of the dynamical system). Because there are 12 states (6 directions and 2 ordinal numbers), $\hat{F}^{(1)}$ contains 12 vectors, each a 2×1 , for a total of 24 parameters. Y^* is the reference trajectory for each movement. Because there are 10 different kinds of movements (Fig. 2B), Y^* has 10 vectors each representing the reference position (a 2×1 vector) for each movement. Therefore, the total number of unknown parameters is 61. However, it is not necessary to estimate $\hat{F}^{(1)}$. Rather, $\hat{F}^{(1)}$ for a given set is simply $\hat{F}^{(240)}$ for the previous set. In case of the first field set, $\hat{F}^{(240)}$ comes from the previous null set and for the null set $\hat{F}^{(1)}$ is simply zero. This reduced the number of unknown parameters in the system to 37.

To estimate these parameters, we minimized the sum of squared errors between the model and the data

$$\varepsilon = \frac{1}{2} \sum_i [y^{(n)} - \hat{y}^{(n)}]^T [y^{(n)} - \hat{y}^{(n)}]$$

To minimize this cost, we used a version of Newton's method for gradient descent. For each model parameter, we evaluated the derivatives of the cost function with respect to that parameter. For example, we updated the current estimate of a by amount Δa

$$\Delta a = - \left(\frac{d^2 \varepsilon}{da^2} \right)^{-1} \frac{d\varepsilon}{da}$$

To help compute the derivatives, we "unfolded" the recursive Eq. 5 and wrote it as follows

$$\begin{aligned} \hat{y}^{(n)} &= Df^{(n)} + L^{(n)}Y^* - DK^{(n)}\Gamma^{(n-1)}f_h^{(n-1)} - DK^{(n)} \prod_{i=n-1}^1 A^{(i)}\hat{F}^{(1)} \\ A^{(n)} &= aI - BK^{(n)T}K^{(n)} \\ \Gamma^{(n-1)} &= \begin{bmatrix} A^{(n-1)}A^{(n-2)} \dots A^{(2)}BK^{(1)T} \\ A^{(n-1)}A^{(n-2)} \dots A^{(3)}BK^{(2)T} \\ \vdots \\ A^{(n-1)}BK^{(n-2)T} \\ BK^{(n-1)T} \end{bmatrix}^T \\ f_h^{(n-1)} &= [f^{(1)} \quad f^{(2)} \quad \dots \quad f^{(n-1)}]^T \end{aligned}$$

where I is an identity matrix and T is the transpose operator.

Measuring the goodness of fit

Once system parameters were estimated, Eq. 5 was evaluated by providing it a sequence of inputs $f^{(n)}$. The output of the system was a sequence $\hat{y}^{(n)}$. Therefore, the model was evaluated in an autonomous mode without access to the measured data $y^{(n)}$. To evaluate goodness of fit we followed a procedure for testing significance of nonlinear models (Glantz and Slinker 2001). We computed the sum of squares explained by the model over the sequence of trials

$$ss_{\text{mod}} = \sum_{i=1}^{240} [\hat{y}^{(i)} - L^{(i)}\bar{Y}]^T [\hat{y}^{(i)} - L^{(i)}\bar{Y}]$$

where \bar{Y} is a 20×1 vector that contains the mean of hand positions (over that set of trials) for each of the 10 types of movement (Fig. 2B). The total sum of square was

$$ss_{\text{tot}} = \sum_{i=1}^{240} [y^{(i)} - L^{(i)}\bar{Y}]^T [y^{(i)} - L^{(i)}\bar{Y}]$$

The difference is the residual variation of the measurement about the model (i.e., $ss_{\text{res}} = ss_{\text{tot}} - ss_{\text{mod}}$). Standard error of the model is

$$se_{y|f} = \sqrt{\frac{ss_{\text{res}}}{n-p}} = \sqrt{ms_{\text{res}}}$$

where n is the number of trials, p is the number of model parameters, and ms_{res} is the mean of squared variations about the model. An F -statistic was calculated by the ratio

$$F = \frac{ms_{\text{mod}}}{ms_{\text{res}}}$$

where $ms_{\text{mod}} = (ss_{\text{tot}} - ss_{\text{res}})/p$. For this F -statistic, the degrees of freedom in the numerator and denominator are p and $n - p$. We also report coefficient of determination $r^2 = 1 - ss_{\text{res}}/ss_{\text{tot}}$.

The number of parameters in any model has a strong influence on the goodness of fit. As the size increases, the model will fit the data better. To balance this improved fit versus risk of overfitting, we used the Akaike information criterion (AIC) to estimate how the size of the model affected the fit (Akaike 1974). If there are m unknown parameters and N trials, then

$$AIC = \ln \left\{ \frac{\sum_n [y^{(n)} - \hat{y}^{(n)}]^2}{N} \right\} + 2 \frac{m}{N}$$

As the number of parameters in the model increases, the second term increases; as the fit improves, the first term decreases. If AIC increased after a model was enlarged, then the model may be overfitting the data.

RESULTS

We found that the no-cue (NC) group did not adapt to the sequence of force fields, whereas the short (SD) and long delay (LD) groups both adapted. The state-space analysis of the data showed that the contextual cue (i.e., movement order) modulated directional tuning of the generalization function. This modulation was significantly larger in the LD group as compared to the SD group.

Psychophysical results

Figure 3 shows representative trajectories in catch and field trials early and late in training from a single subject in each group. As the training in the field began, the robot produced a counterclockwise field on odd trials, and a clockwise field on even trials. Early in training (trials 1–120), there were large errors in field trials but little or no errors in catch trials. Near the end of training (trials 1,800–1,920), the subject in the NC group showed only a slight improvement on field trials and no consistent errors in catch trials. The lack of errors in catch trials is our most reliable indicator that there were little or no changes in the ability of this subject to predict the forces on a given trial.

Figure 3 also shows data for typical subjects in the SD and LD groups. Training in these subjects resulted in improved performance, and the improvement was somewhat larger in the LD group. Late in training the catch trials showed errors that were different for odd and even trials. That is, if an upward movement was odd numbered in the sequence, the error in a catch trial was opposite to the error observed for the same odd-numbered movement in a field trial. The direction of error in the catch trial reversed when the upward movement happened to be even numbered. Therefore, in the cued group, subjects learned to partially predict forces that depended on both direction of movement and the order of that movement in the sequence.

Figure 4A displays average performance measures for each group. All groups reached to the same sequence of targets and except for the null set in day 1, target sequence was identical in days 1 and 2. Perpendicular displacement (PD) at max velocity was calculated for each trial for each subject and then averaged within the target set and then across subjects. In the NC group, training produced significant changes in the field trials. The average change from set 1 to set 8 for odd trials was +1.3 mm ($P = 0.01$) and for even trials was +2.9 mm ($P < 0.001$). Therefore, performance improved in the even field trials but worsened in the odd field trials. By the end of training

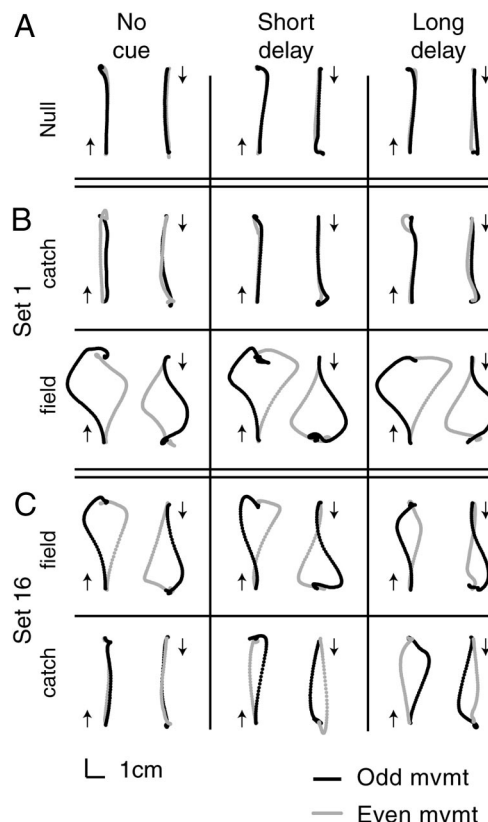


FIG. 3. Performance of 3 typical subjects, one from each group. Only the movements made at 90° and 270° are shown. Arrows indicate movement direction. Odd-numbered movements are displayed in black; even-numbered movements are displayed in gray. Plots show the average hand trajectories in field and catch trials in a target set. A: trajectories in the null set before start of field training. B: catch trials and field trials in the first set of training. C: catch trials and field trials in the last set of training. Although in the no-cue group catch trials did not differentiate between odd and even movements, the individuals in the SD and LD groups had catch trials that were opposite to the field trials and specific to the ordinal number of the movement.

(set 8), errors in catch trials were not significantly different from zero (odd trials, $P = 0.47$; even trial, $P = 0.06$) but the specific sequence of targets in each set appeared to be responsible for some of the variability in the catch trials. To account for this, we compared within-subject change in catch trials from day 1 to day 2 for each target set. We found that although there was a small change in catch trial errors (+0.6 mm for odd trials, $P = 0.07$; +0.96 mm for even trials, $P = 0.06$), the direction of change was not opposite of the field trials. Consequently, in the NC group the learning indices for the odd and even movements (LI_o and LI_e), as well as the total learning index LI , remained near zero.

In the SD group there was a general reduction in errors in field trials and a clear tendency for errors in catch trials to be opposite the field trials. In the field trials, average changes from set 1 to set 8 for odd trials was -2.1 mm ($P < 0.005$) and for even trials was $+2.4$ mm ($P < 0.001$). Therefore with training, performance improved in both the even and odd field trials. By set 8, errors in odd catch trials were on average -4.9 mm (significantly different from zero, $P < 0.001$), and for even catch trials $+2.1$ mm (significantly different from zero, $P = 0.01$). With respect to the catch trial errors in set 1 by set 8 odd catch trials had changed by -3.1 mm ($P < 0.001$) and even catch trials had changed by $+3.8$ mm ($P < 0.001$). Conse-

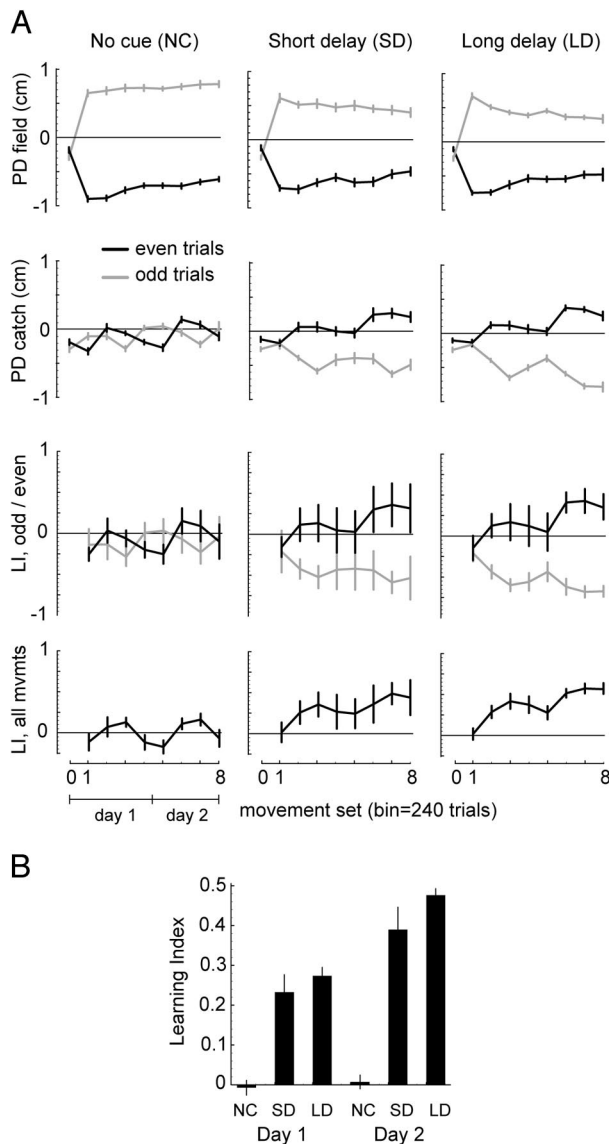


FIG. 4. Average performance in each group. A: PD is the perpendicular distance of the hand position at maximum velocity with respect to a straight line to the target. Sequence of targets was random, but identical between groups and identical in days 1 and 2. In set 0, all trials were in the null field. To compute the performance in field and catch trials, we computed average PD for each subject for each target set (240 trials) and then averaged across subjects. PD plot for field and catch trials represents mean \pm SE. LI is a learning index (Eq. 8), which was calculated separately for odd and even movements and then recalculated for all movements for each subject in each set. LI data points are mean \pm SD. B: average LI for each day of training for the subjects in each group (mean \pm SD).

quently, LI_o became positive, LI_e became negative, and the total learning index LI became consistently positive.

In the LD group, field trials changed from set 1 to set 8 for odd trials by -3.95 mm ($P < 0.001$) and for even trials by $+3.4$ mm ($P < 0.001$). This improvement in performance (errors in set 8 vs. set 1) was somewhat larger than that seen in the SD group, although the change was only marginally significant (odd trials, $P = 0.006$; even trials, $P = 0.05$). By set 8, errors in odd catch trials were on average -6.5 mm (significantly different from zero, $P < 0.001$), and for even catch trials $+1.9$ mm (significantly different from zero, $P = 0.001$). With respect to the catch trial in set 1, by set 8 odd catch trials

had changed by -4.7 mm ($P < 0.001$) and even catch trials had changed by $+3.3$ mm ($P < 0.001$).

To compare performance between groups, we estimated an average learning index for each subject for each day of training (Fig. 4B), and then performed a 2-way ANOVA (main effects of time and group) and found a significant difference between the groups ($P < 0.01$). Both SD and LD groups had significantly larger learning indices (LI) than the NC group on days 1 and 2 (all cases $P < 0.001$). However, although there was a trend for the LD group to have a larger learning index than the SD group, this trend was not significant (day 1, $P = 0.20$; day 2, $P = 0.09$).

Goodness of model fits

We averaged hand positions (at maximum velocity) across subjects within a group and arrived at a 240-element series of hand positions $y^{(n)}$ for each set. Figure 5A shows the 2 components of the vector $y^{(n)}$ for the LD group in the last target set (blank lines). We fit Eq. 5 to these trials. Figure 5A displays the model's fit $\hat{y}^{(n)}$ (gray lines). Among all the sets, this was one of the worst fits (Fig. 5B): the model accounted for 72.9% of the variance in the data [$F(37,203) = 14.8$, $P < 0.0001$]. On average the model accounted for 82.4% of variance in the NC group, 75.5% of the variance in the SD group, and 76% of the variance in the LD group.

We cross-validated each fit by testing it on the subsequent target set. We fixed all parameters of the model and set $\hat{F}^{(1)}$ for the test set to be equal to the estimate for the last movement in the previous set. We found that, on average, the r^2 values declined by 6.8%. The decline was largest in the LD group (8.7%) and smallest in the NC group (4.3%). In all cases the fit remained highly significant ($P < 0.001$).

We analyzed the sequence of $\hat{y}^{(n)}$ produced by the model using the same procedure that we used to compute PD and learning indices in Fig. 4. The plots in Fig. 6 show output of the model in terms of PD in field and catch trials and learning indices. The solid lines are the measured data from subjects (same data as in Fig. 4) and the dotted lines are computed from the model. In field trials, the model tracked the measured data so closely that the dotted and solid lines are exactly on top of each other. In catch trials, the model tracked the data well but appeared to slightly underestimate the size of the catch trials in the SD and LD groups. This produced slightly smaller learning indices for the model versus the measured data.

The generalization function

Figure 7A plots the generalization function estimated by the model for each set and each group. In the NC group, the model allowed for the generalization function b to have 12 degrees of freedom (DoF). The result of the fit was a b value that was consistently positive and often tuned to angular distance, but was insensitive to ordinal position of movements, i.e., $b(0, \Delta\varphi) \approx b(1, \Delta\varphi)$. Average b across the sets is plotted in Fig. 7B. The peak of b is near 0.14, which implies that the internal model's estimate of force \hat{f} for any given direction of movement changed by 14% of error experienced in that direction. The function is bimodal, which means that error experienced in a given direction affected the opposite direction of movement nearly as much as the same direction of movement. Therefore,

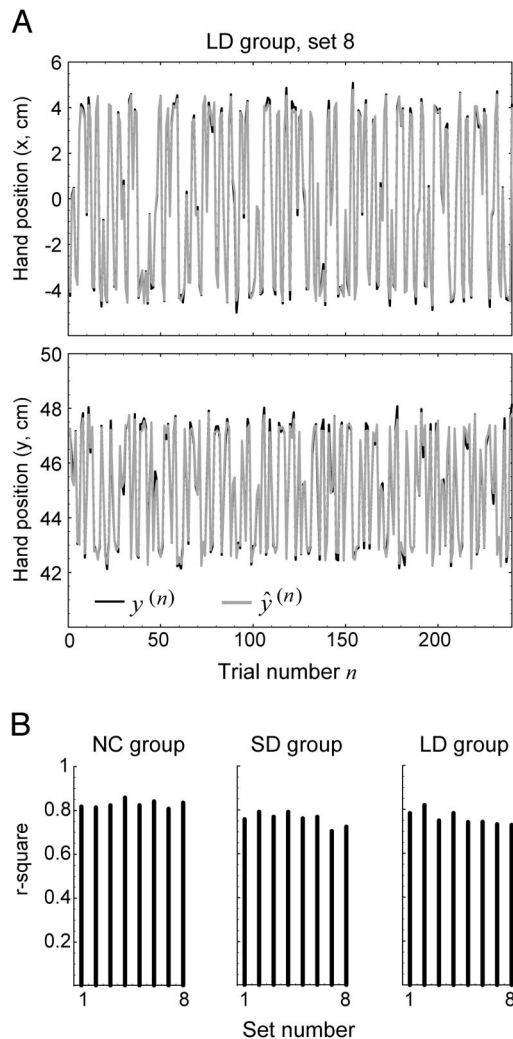


FIG. 5. An example of the fit of the model to across-subject averaged sequence of hand positions during one set (240 trials). *A*: plot shows the sequence of hand positions $y^{(n)}$ in the LD group (averaged across subjects) during the last set of training (set 8). Because $y^{(n)}$ is a 2×1 vector, its x - and y -components are plotted separately. Equation 5 was fit to the data set containing a sequence of 240 input–output vectors $\{f^{(n)}, y^{(n)}\}$. Once the model parameters were estimated, the resulting dynamical system was provided a sequence of inputs $f^{(n)}$, resulting in a sequence of output $\hat{y}^{(n)}$. *B*: goodness of fit of the model. Sets 1–4 were performed on day 1 and the remaining sets were performed on day 2. For each set, the fit as quantified with an F -statistic (see METHODS) is significant at $P < 0.0001$.

in the NC group, error in a given movement produced a consistent adaptation that was reflected as a change in the subsequent movement. However, because the generalization function could not differentiate between the movements in the sequence, the result was trial-by-trial adaptation but no net learning.

It is easy to imagine that the brain might stop responding to the errors in situations where performance does not improve with training. If so, then in the NC group the generalization function should go to zero. The data in Fig. 7*A* suggest that this was not the case. Even in the last set errors in a given trial influenced performance on the subsequent trial.

In the NC group, the insensitivity of the generalization function to movement order suggested that a 12 DoF b was unnecessary. Figure 7*B* shows the average generalization func-

tion when b had only 6 DoF. To quantify the effect of the number of parameters in b with respect to goodness of fit, we used the Akaike information criterion, which indicated that a 12-parameter model was an overfitting of the NC group but not the SD and LD groups. In the NC group, AIC increased from 7.49 to 7.57 when the number of hidden states increased from 6 to 12. This increase indicated that the cost of increasing the number of parameters was not balanced by an improvement in fitting the data. However, in the SD and LD groups, AIC decreased from 7.48 to 7.40 and 7.59 to 7.13, respectively, when the number of hidden states was increased from 6 to 12. This agrees with the intuition that sensitivity to movement order is necessary to account for the performance in the SD and LD groups.

However, the plots in Fig. 7*A* for the SD and LD groups suggest that that, whereas b is directionally tuned, $b(0, \Delta\varphi)$ is very similar to $b(1, \Delta\varphi)$. The level of similarity between these 2 functions is a measure of the interference between 2 movements of the sequence. Across-set average of the generalization functions is plotted in Fig. 7*B*. Because $b(0, \Delta\varphi)$ and $b(1, \Delta\varphi)$ are very similar, the brain is generalizing the errors across the

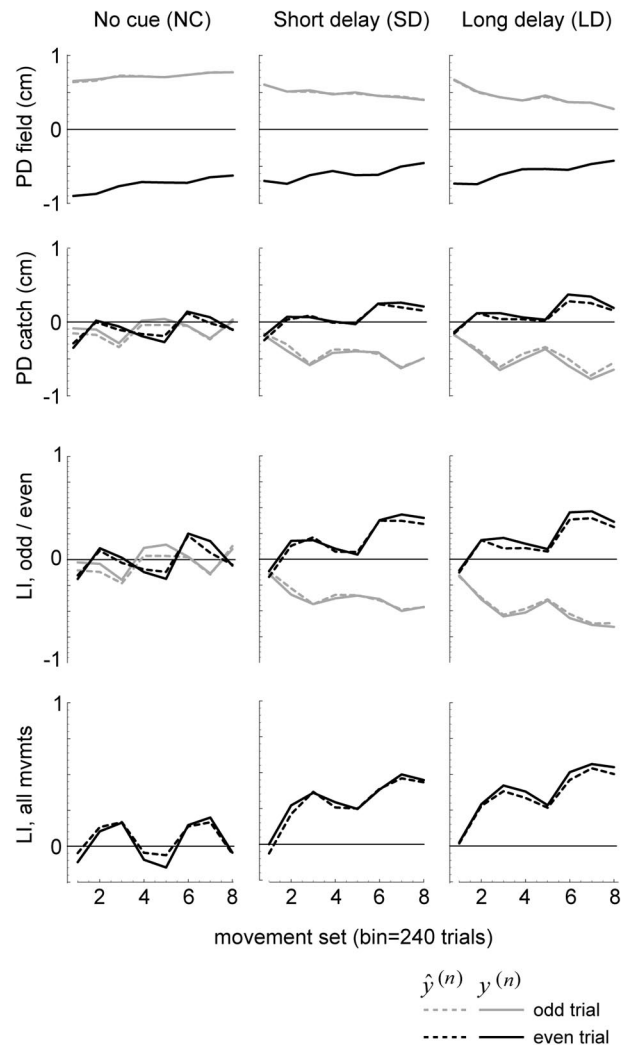


FIG. 6. Model's output $\hat{y}^{(n)}$ is displayed in a format identical to that of Fig. 4. Model's output is plotted with the dashed lines, whereas the measured data (same as in Fig. 4) is plotted with solid lines. For the PD measure in field trials, the dotted and dashed lines are exactly on top of each other.

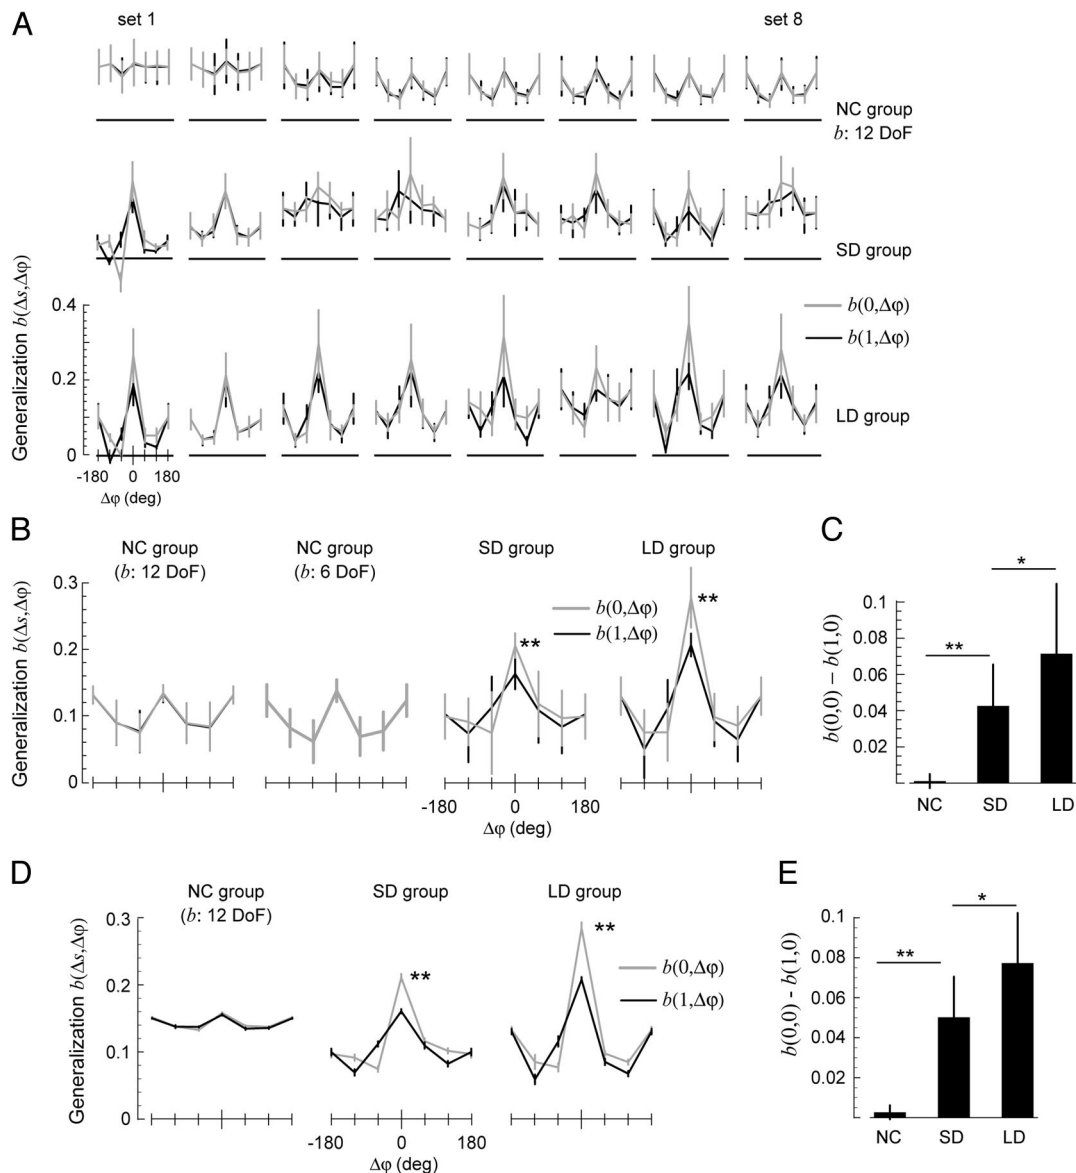


FIG 7. Generalization function $b(\Delta s, \Delta\varphi)$: an estimate of the proportion of error experienced in a given movement that affects any other movement, as a function of the distance between the states of the 2 movements. Distance in terms of ordinal number is noted by Δs . Distance in terms of movement direction is noted by $\Delta\varphi$. There are 6 distances in direction space and 2 distances in ordinal space, producing 12 degrees of freedom (DoF) in the generalization function. Results in A–C are from model fits to the across-subject averaged sequence $y^{(m)}$ (as in Fig. 6). Results in D–E are from model fits to $y^{(m)}$ of individual subjects. A: $b(\Delta s, \Delta\varphi)$ (mean \pm SD, bootstrapped) is plotted for each target set and each group. B: average b over the course of the experiment for each group (mean \pm SD). NC group was analyzed with both a 12 DoF model and a 6 DoF model. In the 12 DoF model, the lines for the $b(0, \Delta\varphi)$ and $b(1, \Delta\varphi)$ are on top of each other. ** indicates significant ($P < 0.01$) modulation of the generalization function at 0° . C: within-group effect of contextual cue on the generalization function at 0° , measured within each set. Error bars are SD. * $P < 0.05$, ** $P < 0.01$. D: within-subject estimate of $b(\Delta s, \Delta\varphi)$. Generalization function was estimated for each subject in each set and then averaged across sets and then averaged across subjects (plotted as mean \pm SE). E: within-subject effect of the contextual cue on the generalization function at 0° (mean \pm SD). * $P < 0.05$, ** $P < 0.01$.

movements almost irrespective of their numerical order in the sequence. This explains why the task is so hard to learn.

In the SD and LD groups, the only significant within-group difference between $b(0, \Delta\varphi)$ and $b(1, \Delta\varphi)$ is at $\Delta\varphi = 0^\circ$ ($P < 0.01$ for both groups). At this angular distance, movement of a given ordinal number had a significantly larger effect on the contents of the internal model for that same movement number than for the subsequent movement. This implies that contextual information about the movement affected the directional tuning of each basis only along its preferred direction.

In our hypothesis, we formulated 3 theoretical scenarios in which a system might learn to compute a function that depends

on direction of movement and context. In the multiplicative model and in the mixture of experts, the ordinal number multiplicatively influences the directional tuning of the bases. Both scenarios predict that the difference between $b(0, \Delta\varphi)$ and $b(1, \Delta\varphi)$ should be largest at $\Delta\varphi = 0^\circ$. This prediction is in agreement with our data. In an additive model, this difference should be equal for all angular differences. Therefore, the data in Fig. 7B clearly reject the additive model. However, the results are not entirely consistent with a multiplicative model. Because b is bimodal, the multiplicative hypothesis predicts a nonzero difference between the generalization functions both at $\Delta\varphi = 0^\circ$ and at $\Delta\varphi = 180^\circ$ (the 2 main peaks of the

generalization function). We observed a significant nonzero difference only at $\Delta\varphi = 0^\circ$.

Figure 7C plots the within-group quantity $b(0, \Delta\varphi) - b(1, \Delta\varphi)$ at $\Delta\varphi = 0^\circ$. The measure is marginally larger in LD versus SD ($P = 0.045$), but much larger in SD versus NC ($P < 0.001$). The small difference between the LD and SD groups is consistent with the small performance differences in the 2 groups (Fig. 4). This suggests that the contextual cue in both LD and SD groups increased output of the bases at the preferred direction of movement. However, in the LD group the increase was marginally larger.

When we compare the generalization functions between groups, we do not find that $b(1, \Delta\varphi)$ in the LD or SD groups is smaller than $b(0, \Delta\varphi)$ in the NC group. Therefore, availability of the contextual cue did not result in a reduction of generalization between incompatible movements. This would have been a reasonable way to learn the task. Rather, the contextual cue preferably increased generalization between compatible movements. Therefore, the presence of the contextual cue (in comparison to its absence or irrelevance) appeared to produce increased activity among some of the bases without greatly affecting the remaining bases.

These results were arrived at by fitting the model to the across-subject averaged time series of hand positions. We also fit the model to the time series produced by each subject in each set. The fits produced 8 estimates of b for each subject (one for each set). We averaged these 8 estimates to produce a mean b for each subject and plotted the across-subject distribution (mean \pm SE) in Fig. 7D. The NC group again showed no difference between $b(0, \Delta\varphi)$ and $b(1, \Delta\varphi)$, whereas the SD and LD groups showed modulation at $\Delta\varphi = 0^\circ$. The within-subject modulation $b(0, \Delta\varphi) - b(1, \Delta\varphi)$ at $\Delta\varphi = 0^\circ$ is plotted in Fig. 7E. This gain was significantly larger in the LD group versus the SD group ($P = 0.01$), and significantly larger in the SD group versus the NC group ($P < 0.001$).

In the NC group, model fits to individual subjects produced an average b that was somewhat flatter (less directionally tuned) than that observed in the fits to the mean of across-subject time series (Fig. 7D vs. B). The reason for this is unclear to us but we note that, in general, our ability to fit data from single subjects was poor. Average r^2 values for single-subject fits were 0.44. Although this is still a highly significant fit [$F(37,203) = 4.54, P < 0.001$], it explains only half of the variance that the model could account for when the times series was an across-subject average.

Gain of online feedback and other model parameters

Other than generalization, there were 4 parameters in the model: a , D , Y^* , and $\hat{F}^{(1)}$. Here we report on the values of these parameters when the model was fit to individual subjects. The parameter a represented trial-to-trial changes in \hat{F} that could not be accounted for by error in the previous trial. For example, if there was memory decay in the system from trial to trial, then a should be < 1 . We found that in all groups and all sets, a was extremely close to 1: in the NC group, $a = 1.00 \pm 0.002$ (mean \pm SD); in the SD group, $a = 0.99 \pm 0.011$; in the LD group $a = 0.99 \pm 0.009$. Therefore, we could not observe any decay in the state transitions.

The D matrix represented the inverse of the gain of an online feedback control system that corrected errors during a move-

ment. In Fig. 8A, the between-subject averaged D matrix in each set is plotted by multiplying it by a unit vector as it rotated about a circle. The shape of the matrix is very similar to the compliance of the arm (Mussa-Ivaldi et al. 1985). We examined the main effect of group and set number on the orientation and size of D . Matrix size (determinant) was not different among the groups [$F(2,23) = 0.15, P > 0.8$]. With training, the determinant tended to become smaller [$F(7,23) = 0.44, P = 0.03$], which implies that the gain of the online feedback system appeared to increase. The D matrix was oriented at a mean angle of 18.3, 29.2, and 14.5° for the NC, SD, and LD groups, respectively [$F(2,23) = 15.4, P < 0.01$]. Training did not produce a significant change in this orientation [$F(7,23) = 0.09, P > 0.8$]. Post hoc analysis of orientations did not find a significant difference between the NC and the LD groups [$F(1,15) = 2.16, P = 0.2$], although the SD group had a D with an orientation significantly different from that of the other groups.

The vectors in Y^* represent the reference points for each movement type, i.e., the location at maximum velocity planned for that movement. This interpretation makes sense because when $y^{(n)} - L^{(n)} Y^*$ is zero, there is zero error in that trial and no changes take place in the internal model. The gray dots in Fig. 8B are the between-subject averaged Y^* for each of the 10 movements in the task in each target set. The black dot in Fig. 8B is the average hand position at maximum velocity in the null set. In general, movements in the null field were not straight, but tended to reside inside the diamond shape. In the field sets, Y^* remained close to the reference location in the null set. Therefore, the planned trajectories were generally not along a straight line to the target and the nonstraight pattern was consistent between groups.

The vector $\hat{F}^{(1)}$ represents the initial condition of the internal model (i.e., expected force) at the start of the target set (trial 1). Because the internal model has 12 hidden states (6 directions of movement and 2 ordinal numbers), there is a force predicted for each one of these states. We did not fit $\hat{F}^{(1)}$ explicitly but set it equal to the estimate provided from the final movement in the preceding target set. Figure 8C plots $\hat{F}^{(1)}$ for set 1 in each movement direction for each group. For each direction there are 2 vectors. The gray and black vectors represent the components of $\hat{F}^{(1)}$ expected for an odd- and an even-numbered trial, respectively. At the onset of training, the gray and black vectors are generally small and point to the same direction in all groups. By the start of the final set of training (Fig. 8D), $\hat{F}^{(1)}$ values in the SD and LD groups point nearly perpendicular to each direction of movement. In these 2 groups $\hat{F}^{(1)}$ is clockwise for even trials and counterclockwise for odd trials, i.e., opposite the force field produced by the robot in these trials (Fig. 2B). By contrast, $\hat{F}^{(1)}$ shows no obvious pattern for the NC group.

Estimating the effects of model assumptions and simplifications

For the purpose of computing a generalization function, we collapsed the 10 movements that were actually produced into 6 movement directions as if they were performed from a single start location. This assumption produces 2 sources of error in fitting the model to the data. First, limb compliance is a function of hand position (Mussa-Ivaldi et al. 1985). The assumption of a constant compliance is a violation of this

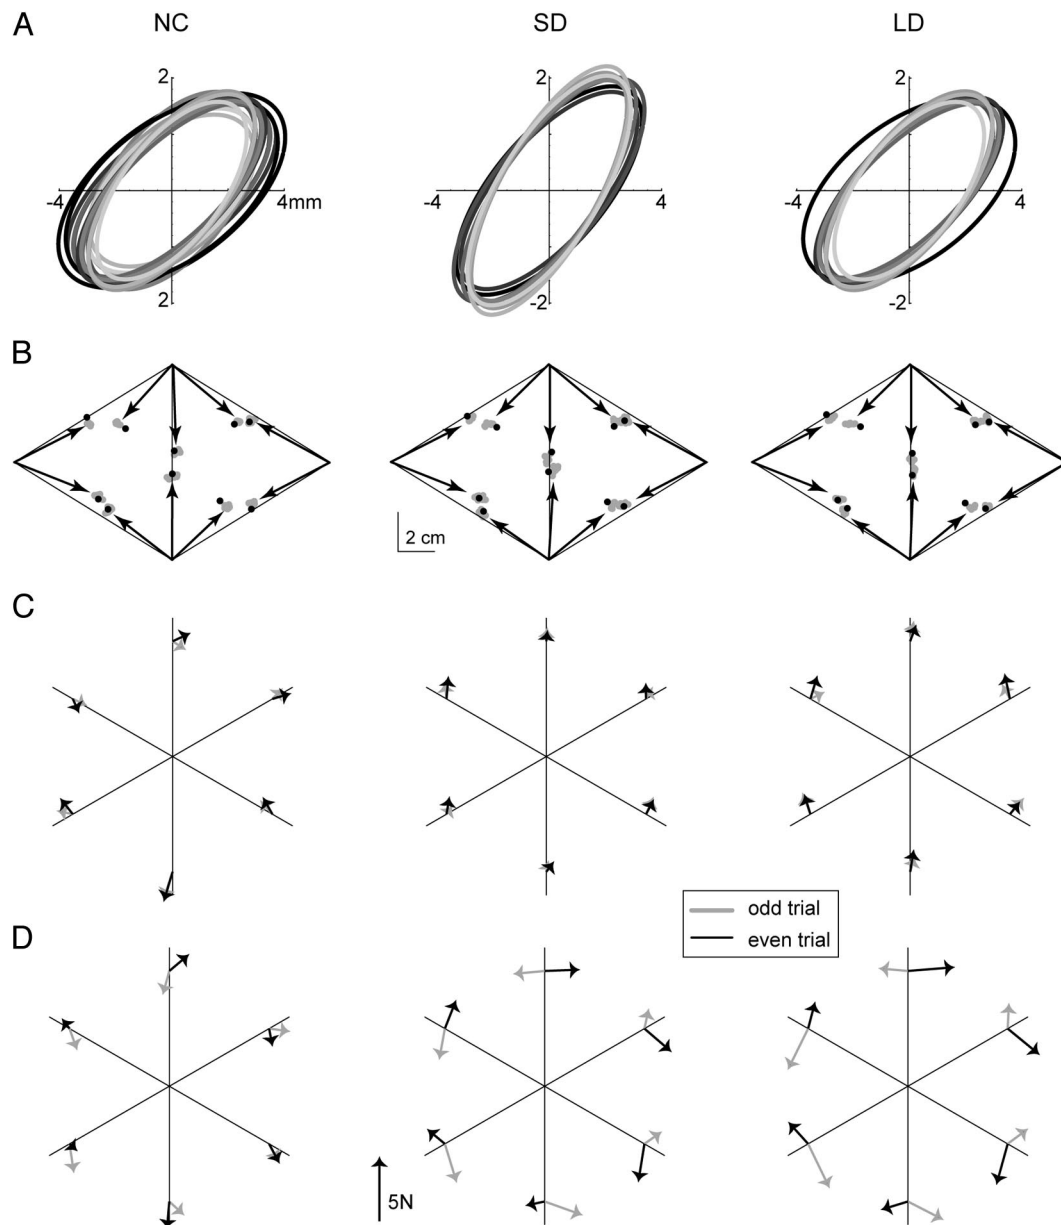


FIG. 8. Estimates of model parameters D , Y^* , and $\hat{F}^{(1)}$ computed from fits to individual subjects and then averaged across groups. *A*: D matrix, representing inverse of the gain of the online error feedback system (i.e., arm compliance). To plot the matrix, we multiplied it by a unit vector that moved about a circle. Darkest line represents the first set of training and the lightest gray line represents the final set. Area of the ellipse tends to become smaller, indicating an increased gain. *B*: Y^* , representing the reference points or “planned trajectory” for each movement. Diamond represents the spatial location of the targets and the arrows indicate the direction of travel for each movement (as in Fig. 2*B*). Black dot for each movement represents the mean hand location at max velocity in the null trials. Gray dots represent Y^* estimated during each of the 8 sets of field trials. *C*: $\hat{F}^{(1)}$, representing the initial condition of the internal model (i.e., expected force) at the first trial of the first force field training set (set 1). *D*: $\hat{F}^{(1)}$ at the first trial of the last training set (set 8).

observation. Second, the generalization function may be smaller when the 2 movements have different origins as compared to when they have the same origin. Hwang et al. (2003) demonstrated that generalization as a function of direction was modulated with a gain that depended on hand position.

To quantify the effects of these sources of error on our model, we performed a set of simulations. We began with a model of the human arm biomechanics with position-dependent stiffness and inertial parameters (Shadmehr and Mussa-Ivaldi 1994), coupled to the dynamics of our robot arm. We added to this model a mechanism of adaptation by basis functions. The bases encoded both limb position and velocity

(in joint coordinates). The encoding of limb velocity was with bimodal Gaussians (Donchin et al. 2003) that were modulated linearly (as in a gain field) with changes in limb position (Hwang et al. 2003). We then performed 2 simulations. In one simulation, the 10-cm movements were along the diamond pattern, i.e., movements to the same direction could start from different start positions (as in the current experiment). In another simulation, all movements were from the same start location at the center. In both simulations, there were 6 possible directions of movement. The idea was to determine to what extent our assumptions regarding position-invariant generalization and compliance degraded our ability to fit the data.

In the simulation, 6 sets of 192 movements were performed in a field that randomly changed from null, to a clockwise curl, to a counterclockwise curl field, producing a performance comparable to that of the no-cue group. In the first 3 sets, all movements began from the center. In the remaining 3 sets, movements had origins that fell on a diamond pattern, identical to our experiment. Values of r^2 for the fits declined from 0.843 to 0.825. That is, the model could account for about 2% less of the variance in the data in the diamond pattern versus in the common-origin pattern. Therefore, the 2 assumptions clearly influenced the fit in a negative way. However, the effect did not appear to be very large. Indeed, this level of fit is quite consistent with the model's performance on actual data (Fig. 5B).

DISCUSSION

In this experiment the serial order of a movement, as well as its trajectory, determined the forces on the hand. This is an example of a situation where contextual cues that are neither proprioceptive nor visual inform the motor systems as to how to program the motor commands. We found that, if given information about serial order, subjects slowly adapted to the force patterns. They learned to produce motor commands that were specific to both the movements' direction and order. We used a system identification approach to track trial-by-trial performance and to estimate a generalization function. The generalization function showed the effect of error in one movement on the subsequent movement as a function of their distance in state space. State space was defined as movement direction and serial order within the sequence. The generalization pattern suggested that the bases that computed the internal model had activity fields that were sensitive to both ordinal number and direction, but that the contextual cue acted as a weak modulating factor on directional sensitivity.

Effect of withholding the contextual cue

We observed little or no learning of the force fields in the no-cue (NC) group where no temporal cues were available to identify an odd or even movement. The results show that the repeating patterns of error and force are not sufficient for the brain to learn the sequence-dependent field. An explicit cue may be necessary. This lack of adaptation without an explicit cue confirms earlier observations (Karniel and Mussa-Ivaldi 2002, 2003). However, one of our key findings is that, although the errors did not converge in this group of subjects, the generalization function was consistently nonzero, implying that subjects continued to respond to the errors in a given trial by altering their motor commands in the subsequent trial. Thus trial-to-trial adaptation continued to take place despite the fact that errors did not converge.

This result, and similar results from Scheidt and colleagues (2001), may have implications for brain imaging paradigms where a random trial is offered as a condition for which no learning is thought to occur. The state-space model suggests that despite behavioral measures that show no consistent improvements in performance, the brain continues to adapt to errors much the same way as in a nonrandom task.

Coding of a contextual cue

With the availability of serial order information, subjects learned to produce motor commands that were sensitive to both

movement direction and serial order. A number of recent studies have found that certain cues can be effective in training subjects to produce different motor commands for the same direction of movement (Krouchev and Kalaska 2003; Osu et al., 2004; Wada et al. 2003). How might the contextual cue and movement direction be encoded by the neural system that learns to represent the internal model? If neural activity fields acted as bases with which the map $(s, \varphi) \rightarrow \hat{f}$ was approximated (where s is a contextual cue and φ is target direction), then encoding will be related to generalization; thus our objective was to quantify how the information about serial order affected patterns of generalization.

We found that the generalization function was largest at directional distance of 0° but had a smaller second peak at 180° . This implies that the bases are directionally tuned but the tuning is possibly bimodal. This kind of tuning is consistent with activity of some Purkinje cells in the cerebellum (Coltz et al. 1999). However, to learn the task, the bases must also be sensitive to the ordinal number of the movement. We found that movement order modulated generalization at 0° but not significantly at other angular distances. To account for this effect, we began with the assumption that the information about movement order was coded as a Gaussian with a center located at a preferred movement number within the sequence. The choice of a Gaussian arose from the observations of Georgopoulos and colleagues regarding the discharge of M1 cells in a task where the animal kept track of potential reach targets as they appeared in a sequence (Carpenter et al. 1999). They found that during the delay period, cells preferred a specific target number within a sequence independent of target location; that is, the neuronal responses were phasic and specific to the preferred target number. The choice of a Gaussian is also attributed to observations of Clower and Alexander (1998), where neurons in SMA discharged maximally for a particular target when it appeared at a particular ordinal number within a sequence.

How might this coding of movement number be combined with directional tuning? Using a model, we contrasted generalization of bases that incorporated the effect of movement order and directional tuning with either an additive or a multiplicative interaction. The generalization patterns in our subjects were clearly inconsistent with the additive model. The multiplicative model predicted that the largest modulation should occur at 0° , which agreed with our data. The multiplicative model also predicted that there should be a smaller but significant modulation at the second peak of the generalization function (i.e., at 180°). We did not observe this in the data. The discrepancy between prediction and observation may be explained by a number of factors; one is the possibility that the bimodality of the generalization function is not an inherent property of a single basis. Rather, the bimodal generalization may reflect the contribution of 2 distinct bases, only one of which is modulated by contextual cues. For example, in the computation of an internal model, the state of the arm may be encoded in both visual and proprioceptive coordinates with separate bases (Hwang et al. *Soc Neurosci Abstr* 822.17, 2003). Multiplicative interaction with one but not the other would produce a generalization pattern similar to the data presented here.

A key characteristic of multiplicative interaction is a change in the depth of tuning. For example, in the posterior parietal cortex there is preliminary evidence that cells that encode a reach plan (i.e., target location with respect to fixation) have a

receptive field that is modulated multiplicatively with reward expectancy (Corneil et al. 2003). It is possible that multiplicative modulation of activity fields is a common mechanism by which information from disparate systems is combined by neurons that take part in computing a transformation (Pouget and Sejnowski 1997). A recent model has found that simulated neurons that multiplicatively code contextual and noncontextual information can allow a single network to switch behavior to respond to changing environmental demands (Salinas 2004). This kind of coding of diverse sources of information is a plausible method with which a single network of neurons can effectively store multiple, context-specific internal models.

In contrast to using multiplicative coding within a single network, an alternate approach is to assume existence of multiple modules, called a MOSAIC, which is a collection of competing internal models (Wolpert and Kawato 1998). In MOSAIC, internal models are controllers that receive the goal (or desired trajectory) of the movement. Each controller predicts a different motor command, and a switching mechanism assigns likelihood of success to each potential command. Contextual cues are priors that influence the probability of selection of these weights. MOSAIC predicts that neurons that form each internal model are themselves not modulated by the contextual cue. We found that a simple formulation of this mixture of experts produces the same patterns of generalization found with multiplicative encoding of a single network; thus generalization patterns appear to lack the power to differentiate between these 2 different architectures of adaptation.

Saliency of the contextual cue

Our conjecture that ordinal number multiplicatively modulates directional tuning runs into trouble when we consider that in the long-delay (LD) group, performance was slightly better than that in the short-delay (SD) group. If one is aware of a particular movement's ordinal number within a sequence, why should the temporal delay between the 2 sequences affect learning? One explanation is that in the short-delay group, subjects were simply more fatigued because of the closer proximity of the trials. However, if fatigue has a major role in this task, then in all groups the generalization function should gradually decline during the day of training. The data do not agree with the fatigue hypothesis.

We do not know whether the SD and LD groups could equally identify the 2 movements in the sequence; we quizzed them verbally only during the null field training trials. Our verbal measure suggested equal cognitive awareness, although this was not documented. However, the generalization function in the LD group was significantly more modulated by the contextual cue than in the SD group.

Implicit in the structure of a sequence is its beginning and its end. In our task, the longer delay interval between the sequences was the cue that one sequence had ended and another was about to begin. Fujii and Graybiel (2003) found that in the prefrontal cortex of monkeys, neurons had a phasic peak of activity that signaled the beginning and end of a sequence of movements. For example, the burst peaked about 250 ms after the last movement in the sequence (in this case, a saccade), lasted about 500 ms, and was independent of sequence length, pace, or reward condition, even when the sequence was only 2

movements long. Separate bursts signaled beginning and completion of a sequence.

Imagine that such signals are responsible for starting and resetting a counter that indicates ordinal number of a movement within a sequence. It is possible that in the SD group, where the delay between termination of one sequence and start of the next was 1.0 s; neuronal activities that signal the 2 events interfered temporally, causing an occasional misestimation of movement number within the sequence. There was less chance for this temporal overlap when the delay was 3.0 s (LD group). If we measure the degree of separation of 2 movements in the sequence by the amount of modulation of the generalization function, then in the LD group the 2 movements were functionally better separated. In the LD group, errors in one movement had a smaller interference on the subsequent movement. Given the time course of the prefrontal bursts that identified start and end of a sequence (Fujii and Graybiel 2003), the hypothesis predicts that there should be no further gains in performance if the sequence is separated by longer intervals.

Limitations of the theory

The theory that we developed here is a step toward the goal of representing learning and memory as the accumulation of small trial-to-trial changes. The data that we analyzed were limited to approximately the first 250 ms of movement, a period when motor commands are thought to be dominated by a feed-forward internal model that relies little on online feedback (i.e., from the current movement). However, our earlier work (Bhushan and Shadmehr 1999; Wang et al. 2001) and those of our colleagues (Burdet et al. 2001; Franklin et al. 2003) had found that adaptation in this task involved changes in both the feed-forward inverse model and a feedback-dependent forward model of dynamics. Indeed, we did observe a small but significant increase in the gain of the online feedback system with the progression of training. These changes are probably better measured in later periods of a trajectory, but require random perturbation to the limb. Identification of online mechanisms of control remains an unexplored area for our theory.

The formulation of the theory was in terms of a deterministic system and did not consider noise. In particular, the model assumed that estimates of force were one and the same as motor commands to the limb. If signal-dependent noise (Harris and Wolpert 1998) is an important contributor to generation of motor commands in the range of forces that are typical for reaching, then estimate of force and its production are not equal. To account for this, we would need to introduce a noise term in the output equation. The fact that this term is missing from our formulation may explain why we could fit the across-subject averaged time series almost twice as well as the within-subject time series. Averaging across subjects reduces the variance of the output noise.

Our measure of distance between 2 ordinal numbers was symmetric; that is, the effect of movement 1 on movement 1 was assumed to be the same as the effect of movement 2 on movement 2 (the 2 states were the same). The effect of movement 1 on movement 2 was assumed the same as the effect of movement 2 on movement 1 (the 2 states were different). This resulted in a generalization function that mea-

sured distance in sequence space as being either 0 or 1. The rationale for this was to keep the number of unknown parameters in the generalization function to a minimum. In a longer sequence, "same" and "different" will be insufficient measures of distance between ordinal numbers.

Both adaptation and learning in this task involve more than what we have modeled. However, the behavior that we observed is largely accounted for by a theory where movements are attributed to a population coding where cues regarding state of the limb and context are transformed into a force-like motor command. We predict that this transformation is computed with neurons that are directionally tuned, and the cue associated with movement order weakly modulates this tuning at the preferred direction. This hints at a general solution to the question of how the brain learns multiple internal models. During adaptation, information from disparate sources may be combined multiplicatively in every neuron. Attention, reward, or other cognitive factors serve to highlight specific channels of information, increasing the gain with which the particular information modulates directional tuning.

ACKNOWLEDGMENTS

We are grateful for the insightful comments of three anonymous reviewers.

Present address of O. Donchin: Department of Biomedical Engineering, Ben-Gurion University of the Negev, Be'er Sheva, Israel.

GRANTS

This work was supported by National Institute of Neurological Disorders and Stroke Grants NS-37422 and NS-46033.

REFERENCES

- Akaike H.** A new look at the statistical model identification. *IEEE Trans Automatic Control* AC-19: 716–723, 1974.
- Bhushan N and Shadmehr R.** Computational architecture of human adaptive control during learning of reaching movements in force fields. *Biol Cybern* 81: 39–60, 1999.
- Bosco G, Rankin A, and Poppele RE.** Representation of passive hindlimb postures in cat spinocerebellar activity. *J Neurophysiol* 76: 715–726, 1996.
- Burdet E, Osu R, Franklin DW, Milner TE, and Kawato M.** The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414: 446–449, 2001.
- Carpenter AF, Georgopoulos AP, and Pellizzer G.** Motor cortical encoding of serial order in a context-recall task. *Science* 283: 1752–1757, 1999.
- Clover WT and Alexander GE.** Movement sequence-related activity reflecting numerical order of components in supplementary and presupplementary motor areas. *J Neurophysiol* 80: 1562–1566, 1998.
- Coltz JD, Johnson MTV, and Ebner TJ.** Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor armtracking. *J Neurosci* 19: 1782–1803, 1999.
- Conditt MA and Mussa-Ivaldi FA.** Central representation of time during motor learning. *Proc Natl Acad Sci USA* 96: 11625–11630, 1999.
- Corneil BD, Musallam S, and Andersen RA.** Representation of reward expectancy in the medial bank of the intraparietal sulcus: implications for neural prosthetics. *Soc Neurosci Abstr* 607.8, 2003.
- Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, and Shadmehr R.** Learned dynamics of reaching movements generalize from dominant to nondominant arm. *J Neurophysiol* 89: 168–176, 2003.
- Donchin O, Francis JT, and Shadmehr R.** Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J Neurosci* 23: 9032–9045, 2003.
- Franklin DW, Osu R, Burdet E, Kawato M, and Milner TE.** Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J Neurophysiol* 90: 3270–3282, 2003.
- Fujii N and Graybiel AM.** Representation of action sequence boundaries by macaque prefrontal cortical neurons. *Science* 301: 1246–1249, 2003.
- Georgopoulos AP, Caminiti R, and Kalaska JF.** Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. *Exp Brain Res* 54: 446–454, 1984.
- Georgopoulos AP, Kalaska JF, Caminiti R, and Massey JT.** On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2: 1527–1537, 1982.
- Georgopoulos AP, Schwartz AB, and Kettner RE.** Neural population coding of movement direction. *Science* 233: 1416–1419, 1986.
- Ghahramani Z and Wolpert DM.** Modular decomposition in visuomotor learning. *Nature* 386: 392–395, 1997.
- Glantz S and Slinker BK.** *Primer of Applied Regression and Analysis of Variance*. New York: McGraw-Hill, 2001.
- Harris CM and Wolpert DM.** Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Hwang EJ, Donchin O, Smith MA, and Shadmehr R.** A gain-field encoding of limb position and velocity in the internal model of arm dynamics. *PLoS Biol* 1: 209–220, 2003.
- Karniel A and Mussa-Ivaldi FA.** Does the motor control system use multiple models and context switching to cope with a variable environment? *Exp Brain Res* 143: 520–524, 2002.
- Karniel A and Mussa-Ivaldi FA.** Sequence, time, or state representation: how does the motor control system adapt to variable environments? *Biol Cybern* 89: 10–21, 2003.
- Krouchev NI and Kalaska JF.** Context-dependent anticipation of different task dynamics: rapid recall of appropriate motor skills using visual cues. *J Neurophysiol* 89: 1165–1175, 2003.
- Malfait N, Shiller DM, and Ostry DJ.** Transfer of motor learning across arm configurations. *J Neurosci* 22: 9656–9660, 2002.
- Mussa-Ivaldi FA, Hogan N, and Bizzi E.** Neural, mechanical and geometric factors subserving arm posture in humans. *J Neurosci* 5: 2732–2743, 1985.
- Osu R, Hirai S, Yoshioka T, and Kawato M.** Random presentation enables subjects to adapt to two opposing forces on the hand. *Nat Neurosci* 7: 111–112, 2004.
- Poggio T, Fahle M, and Edelman S.** Fast perceptual learning in visual hyperacuity. *Science* 256: 1018–1021, 1992.
- Pouget A, Dayan P, and Zemel R.** Information processing with population codes. *Nat Rev Neurosci* 1: 125–132, 2000.
- Pouget A and Sejnowski TJ.** Spatial transformations in the parietal cortex using basis functions. *J Cogn Neurosci* 9: 222–237, 1997.
- Salinas E.** Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *J Neurosci* 24: 1113–1118, 2004.
- Scheidt RA, Dingwell JB, and Mussa-Ivaldi FA.** Learning to move amid uncertainty. *J Neurophysiol* 86: 971–985, 2001.
- Shadmehr R and Moussavi ZMK.** Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20: 7807–7815, 2000.
- Shadmehr R and Mussa-Ivaldi FA.** Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208–3224, 1994.
- Thoroughman KA and Shadmehr R.** Learning of action through adaptive combination of motor primitives. *Nature* 407: 742–747, 2000.
- Wada Y, Kawabata Y, Kotosaka S, Yamamoto K, Kitazawa S, and Kawato M.** Acquisition and contextual switching of multiple models for different viscous force fields. *Neurosci Res* 46: 319–331, 2003.
- Wang T, Dordevic GS, and Shadmehr R.** Learning the dynamics of reaching movements results in the modification of arm impedance and long-latency perturbation responses. *Biol Cybern* 85: 437–448, 2001.
- Wolpert DM and Kawato M.** Multiple paired forward and inverse models for motor control. *Neural Networks* 11: 1317–1329, 1998.