

Quantifying Generalization from Trial-by-Trial Behavior of Adaptive Systems that Learn with Basis Functions: Theory and Experiments in Human Motor Control

Opher Donchin, Joseph T. Francis, and Reza Shadmehr

Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, Maryland 21205

During reaching movements, the brain's internal models map desired limb motion into predicted forces. When the forces in the task change, these models adapt. Adaptation is guided by generalization: errors in one movement influence prediction in other types of movement. If the mapping is accomplished with population coding, combining basis elements that encode different regions of movement space, then generalization can reveal the encoding of the basis elements. We present a theory that relates encoding to generalization using trial-by-trial changes in behavior during adaptation. We consider adaptation during reaching movements in various velocity-dependent force fields and quantify how errors generalize across direction. We find that the measurement of error is critical to the theory. A typical assumption in motor control is that error is the difference between a current trajectory and a desired trajectory (DJ) that does not change during adaptation. Under this assumption, in all force fields that we examined, including one in which force randomly changes from trial to trial, we found a bimodal generalization pattern, perhaps reflecting basis elements that encode direction bimodally. If the DJ was allowed to vary, bimodality was reduced or eliminated, but the generalization function accounted for nearly twice as much variance. We suggest, therefore, that basis elements representing the internal model of dynamics are sensitive to limb velocity with bimodal tuning; however, it is also possible that during adaptation the error metric itself adapts, which affects the implied shape of the basis elements.

Key words: adaptation; motor learning; system identification; receptive field; psychophysics; model

Introduction

The brain appears to predict forces that are necessary for an upcoming movement and programs muscle activations accordingly (Lackner and DiZio, 1994; Shadmehr and Mussa-Ivaldi, 1994). When force prediction is incorrect, movement errors drive adaptation of the motor commands (Thoroughman and Shadmehr, 1999). Interestingly, error experienced in a single movement has broad effects. It influences subsequent movements to other directions (Sainburg et al., 1999), at other arm configurations (Shadmehr and Moussavi, 2000), with different trajectories (Conditt et al., 1997; Goodbody and Wolpert, 1998), and even movements of the other arm (Criscimagna-Hemminger, et al., 2002). One hypothesis is that through practice, the brain builds an internal model of the novel forces. The internal model is a mapping from the arm's position and velocity space into force (Conditt and Mussa-Ivaldi, 1999). Therefore, the patterns of generalization are a reflection of how this map is encoded neurally. For example, learning forces in one workspace generalize to an arm configura-

tion 80 cm away (Shadmehr and Moussavi, 2000), suggesting that the neural elements of the internal model encode arm position very broadly. Cells in the motor cortex, somatosensory cortex, and spinocerebellar tract are typically modulated globally and often linearly as a function of position of the limb (Bosco et al., 1996; Tillery et al., 1996).

Generalization is a key psychophysical tool to infer the computations underlying adaptation. Indeed, a number of studies have quantified generalization to infer the shape of the receptive fields underlying a transformation learned by the brain (Shadmehr and Mussa-Ivaldi, 1994; Ahissar and Hochstein, 1997; Conditt et al., 1997; Goodbody and Wolpert, 1998; Imamizu et al., 1998; Lee et al., 1999; Sainburg et al., 1999). Typically, subjects train with one set of inputs and are then tested in another set; however, if we assume that in every trial, including the test trials, errors update the representation, then we can only probe a few test points before the test points themselves become a significant fraction of the training trials. In human psychophysics, we can get around this by recruiting a large number of subjects and after training provide different test points to each volunteer. This becomes impractical in animal studies. The problems associated with separating training from test trials is perhaps the main reason why in motor control, generalization functions are qualitatively inferred from psychophysical data, but their exact shape is often unknown.

Recently, we used a different approach in which there were no test trials. Instead, generalization was estimated from perfor-

Received April 1, 2003; revised July 23, 2003; accepted July 30, 2003.

This work was supported by a grant from the National Institutes of Health (NIH) (R01-NS37422). O.D. was supported by a grant from the NIH (F32-NS11163) and a Distinguished Postdoctoral Fellowship from Johns Hopkins University Department of Biomedical Engineering. We are very grateful to Kurt A. Thoroughman. Many of the ideas presented in this paper are extensions of the analysis that he developed in his thesis. We also thank Maurice A. Smith for discussions.

Correspondence should be addressed to Opher Donchin, Johns Hopkins School of Medicine, 416 Traylor Building, 720 Rutland Avenue, Baltimore, MD 21205-2195. E-mail: opher@bme.jhu.edu.

Copyright © 2003 Society for Neuroscience 0270-6474/03/239032-14\$15.00/0

mance changes during training (Thoroughman and Shadmehr, 2000). The approach focused on trial-by-trial changes in performance. The idea was to quantify how error in one movement affected performance on the subsequent movement. In the current report, we advance the mathematical foundations of the theory and then test them in a number of experiments. We assume that learning of internal models can be thought of as adaptation with basis functions. Although the bases are unknown, we derive a set of equations that relate trial-by-trial changes in performance with a generalization function that is dependent on these bases. We apply the equations to the movements of people and quantify a generalization function. We assert that if this generalization function is truly a reflection of the basis functions, then it should remain invariant as task parameters change. We test the theory with subjects who learned different kinds of force fields, including a field that was unpredictable.

Theory

This paper takes as its starting point the ideas in Thoroughman and Shadmehr (2000). That paper described adaptation of reaching movements to curl fields (fields in which forces are perpendicular to hand velocity). Its conclusion was that “errors in learning dynamics of arm movements suggest that the brain composes motor commands with computational elements that are broadly tuned to arm velocity.” This result rests on two theoretical ideas. First, adaptation of reaching movements to a force field imposed on the hand was viewed as forming a sensorimotor map with a set of bases that encoded limb velocity. The map changed in response to errors experienced in a movement. Thoroughman and Shadmehr (2000) showed that error on one movement was linearly related to performance change on a subsequent movement. Second, that work introduced the idea that a dynamic system could be used to model the sequence of errors that subjects make during the task. In their dynamic system, a combination of actual (f) and expected forces ($\hat{f}$) generated the measured error (y) on each movement (n). The expected force for the next movement was then a linear combination of the old expectation and the force experienced on this movement. Rewriting their Equation 4, in slightly different notation, we have:

$$\begin{cases} \hat{f}^{(n+1)} &= a\hat{f}^{(n)} + bf^{(n)} \\ y^{(n)} &= \hat{f}^{(n)} + df^{(n)} \end{cases} \quad (1)$$

Although additional subtleties were involved, it is broadly correct to say that the three parameters a , b , and d , were used to fit this system of equations to the actual sequence of errors generated by the subject. The dynamic system fit the sequence of errors quite well, and an analysis of the parameters of fit led to the conclusion about broad tuning quoted above. Thoroughman and Shadmehr (2000) chose to use a dynamic system with a canonical form that has been thoroughly analyzed in the literature of system identification in part because the canonical form was numerically tractable.

We began our research by noting that despite the success of this approach, Equation 1 does not have any privileged relationship to the idea of an adapting internal model. Therefore, we attempted to derive a rigorous connection between the theories of adapting internal models and dynamic systems. Starting from simple assumptions about the internal model, this mathematical approach led us to a very different dynamic system than the one used by Thoroughman and Shadmehr (2000). We outline this

result below, but provide the actual derivation in the supplemental material, section S.1 and S.2 (available at www.jneurosci.org).

As a starting point, we used a simulation that included dynamics of a human-like arm holding a robotic arm (Shadmehr and Brashers-Krug, 1997). The task was to make 10 cm reaching movements toward targets positioned at various directions. The robotic arm produced forces that depended on hand velocity, and the objective was to adapt to these forces so that reaching movements were once again smooth and straight. The simulation adapted to the field by building an internal model that was implemented as a sensorimotor map. This map was represented as a combination of basis elements. Basis elements are a common way to implement a mapping function (Poggio et al., 1992; Sanner and Kosha, 1999), in particular if we assume that the internal model is represented as a population code (Pouget and Sejnowski, 1997). The internal model in our case used basis elements (g_i) to encode limb velocity ($\dot{x}$). Each basis was associated with a preferred force vector, labeled as W_i . Using a population code, the bases were combined to predict the force acting on the hand ($\hat{F}$):

$$\begin{aligned} \hat{F}(\dot{x}) &= Wg(\dot{x}) \\ g &= [g_1(\dot{x}), \dots, g_m(\dot{x})]^T \\ W &= \begin{bmatrix} w_{x1} & \dots & w_{xm} \\ w_{y1} & \dots & w_{ym} \end{bmatrix} \end{aligned} \quad (2)$$

The internal model adapted by changing the preferred force vector associated with each basis. The process was as follows. A movement resulted in a trajectory of errors, each experienced at some velocity during that movement. We discretized this trajectory to 1 msec intervals and for each interval updated the force vectors W_i associated with each basis using gradient descent:

$$w_{ij}^{(n+1)} = w_{ij}^{(n)} - \eta \frac{\partial (F - \hat{F})}{\partial w_{ij}}.$$

The weights were updated after each movement by the force errors experienced in that movement. As expected, the adaptive system exhibited changes in performance from trial-to-trial. In section S.2 of the supplemental material we show that in theory, despite the nonlinear dynamics of the arm, trial-to-trial performance of reaching movements during adaptation generally follows the constraints of a specific linear dynamic system and that this dynamic system has parameters that depend on the shape of the bases.

Here we present the linear dynamic system in some detail, because we will be using it to analyze behavioral data from humans and draw conclusions about the underlying representation. Let us say that on movement n , the subject moved in direction $k^{(n)}$ and made an error $y^{(n)}$. This error is the displacement of the hand from some nominal “desired trajectory,” evaluated at peak velocity. Through a compliance matrix, D , this displacement is related to the difference between the force applied by the robot, $F^{(n)}$, also measured at peak velocity, and the force expected by the internal model, $\hat{F}_{k^{(n)}}^{(n)}$. That is:

$$y^{(n)} = D(F^{(n)} - \hat{F}_{k^{(n)}}^{(n)}).$$

Note that the expected force depends on the direction of movement. The experienced error changes the force expectation for that direction and for all other possible directions of movement. It will be more convenient for us to represent the internal model in units of position rather than force, so we make a change of variable that allows us to accomplish this:

$$z_l^{(n)} \equiv D\hat{F}_l^{(n)}.$$

Because we have eight directions of movement, we have eight $\mathbf{z}^{(n)}$ variables, each of which is a vector. The amount that each of these variables is updated after each movement will depend on the error, $\mathbf{y}^{(n)}$, and on a generalization function that relates its direction to the direction of movement. We will call the generalization function $B_{l,k}$. The value of $B_{l,k}$ gives the effect of error experienced in direction k on the expectation in direction l . Putting all of this together, we get the following dynamic system:

$$\begin{cases} \mathbf{y}^{(n)} = D\mathbf{F}^{(n)} - \mathbf{z}_{k^{(n)}}^{(n)} \\ \mathbf{z}_l^{(n+1)} = \mathbf{z}_l^{(n)} + B_{l,k^{(n)}}\mathbf{y}^{(n)} \quad l = 1, \dots, 8. \end{cases} \quad (3)$$

These equations have two multivalued parameters. The first, D , is a 2×2 compliance matrix with units of meters per Newton. Such a compliance matrix can be shown to accurately describe the errors generated by our simulation (Donchin and Shadmehr, 2002). The second parameter, $B_{l,k}$, an 8×8 matrix, is the generalization function and is dimensionless.

Because this dynamic model was derived from our theoretical adaptive controller, it can be used to uncover the generalization function of the adaptive controller. Indeed, we can use the derivation to write an equation for the generalization function in terms of the basis elements: writing $\hat{\mathbf{x}}_k$ for maximum velocity during movements in direction k and $\mathbf{u}_k$ for unit vectors in direction k , we can show:

$$B_{l,k} \approx \alpha \int_{\mathbf{x}'=0}^{\mathbf{x}'=\|\hat{\mathbf{x}}_k\|} [\mathbf{g}(\mathbf{x}'\mathbf{u}_k)^T \mathbf{g}(\hat{\mathbf{x}}_l)] \mathbf{x}' d\mathbf{x}', \quad (4)$$

where α is a constant defined by Equation S5 of the supplemental material and $\mathbf{g}$ represents the vector of basis functions (Eq. 2). This means that the generalization function approximates an average of correlations of basis elements evaluated at different velocities. It further suggests that, given a generalization function, we may be able to work back to make a guess at the basis elements that would generate it.

Fitting Equation 3 to a sequence of movements is a nonlinear optimization problem. To prevent over-fitting, we opted to reduce the 64 parameters that represent the generalization function $B_{l,k}$ to only 8. This was accomplished by assuming that the generalization function depended only on the angular difference between two directions of movement. That is:

$$\begin{aligned} B_{0^\circ,0^\circ} &= B_{45^\circ,45^\circ} = \dots = B_{315^\circ,315^\circ} \\ B_{0^\circ,45^\circ} &= B_{45^\circ,90^\circ} = \dots = B_{315^\circ,0^\circ}, \end{aligned}$$

and, more generally, $B_{l,k} = B_{l+\theta, k+\theta}$. This simplifying assumption certainly holds for adaptive systems in which basis functions are distributed symmetrically in velocity space. It may not necessarily hold for subjects, but we will be able to judge the model and the validity of this assumption by the quality of fit.

Comparison of theory with Thoroughman and Shadmehr (2000)

Unlike our previous work, here we began with a simulated adaptive controller that learned dynamics of a robotic arm with basis functions and found that in theory, its trial-to-trial performance was governed by a specific linear dynamic system. The parameters of this linear system explicitly depended on the shape of the basis functions used by the adaptive controller. As a result, there are a number of differences between the new model (Eq. 3) and the old one (Eq. 1). A full discussion is presented in section S.4 of the supplemental material; however, three of the differences are

fundamental. First, in the current model, error and force are measured as vectors. This allows us to apply the model to fields other than the curl field and to account for parallel as well as perpendicular error. We will show that with the new model we can account for human behaviors in a wide range of force fields. Second, the scalar representation that was used previously assumed a rotation of the error vector because it was generalized to other directions. Here, our model makes no such assumption. In the results, we show that generalization without a rotation is a better match for the data. Finally, in the current model it is the error of the previous movement that causes learning, whereas in the old model learning was driven by an arbitrary combination of force and the old state.

A fixed desired trajectory

One of the assumptions in our adaptive controller, and hence in the dynamic model of Equation 3, is that the desired trajectory is revealed by movements in the null field (i.e., a baseline condition before application of forces) and does not change as a result of the introduction of the force field. There have been a number of proposals regarding the desired trajectory. Some, like the minimum jerk trajectory that we used in our adaptive controller (Flash and Hogan, 1985), do indeed have these properties. Others, such as the minimum torque change model (Nakano et al., 1999), do not. A recent comparison of a number of competing theories suggested that behavior in force fields was best explained by an invariant desired trajectory (Thoroughman, 2002), somewhat justifying our choice to use this assumption in our theoretical development. If this assumption is wrong, however, then the consequences are quite serious, and we explore this possibility as well.

Materials and Methods

Experimental task. Subjects sat facing a vertical screen, holding the handle of a two-joint, planar robotic arm that was used to apply forces (Shadmehr and Mussa-Ivaldi, 1994). Targets were presented on the screen, and the subject acquired these targets by using the robot to control the position of a cursor (Fig. 1A). Each subject made at least 380 movements during which no forces were applied by the robot (null field training). This was followed by one to three sets of 192 movements during which the robot applied forces during most movements (field trials), but occasional trials were without applied forces (catch trials). Forces were always linearly dependent on hand velocity: $\mathbf{F} = V\dot{\mathbf{x}}$. For most of our subjects, the robot applied a curl field. In the curl field forces are perpendicular to the instantaneous velocity of the hand. We used both clockwise curl fields ($V = \{0 \ 13; -13 \ 0\}$) (Fig. 1B) and counterclockwise curl fields ($V = \{0 \ -13; 13 \ 0\}$) (Fig. 1C).

We also used two additional fields in which forces were applied both parallel and perpendicular to the instantaneous velocity. One was the curl-assist field: $V = \{11 \ -11; 11 \ 11\}$ (Fig. 1D). This field combined a clockwise curl field with an assistive field that acted to increase the subject's velocity in whatever direction the subject was moving. An additional field was a saddle field: $V = \{11 \ 11; 11 \ -11\}$ (Fig. 1E).

Experiments with the adaptive controller. We initially tested our theory on the data generated by various simulated adaptive controllers (see Theory and supplemental materials section S.1). We used an adaptive controller in which the basic elements (Eq. 2) were Gaussians ($\exp[-\|\hat{\mathbf{x}} - \hat{\mathbf{x}}_j\|^2/2\sigma^2]$). The "centers" ($\hat{\mathbf{x}}_j$) of the Gaussians were distributed in a square lattice with a separation of 0.05 m/sec. We found that with reasonably dense tiling, the separation affected only the learning rate and not the generalization function. Width of the Gaussians, σ , was varied in the range 0.1–0.4 m/sec. This affected the patterns of generalization and the trial-to-trial performance of the controller. The objective was to fit this trial-to-trial performance to our linear dynamic system, quantify the generalization function, and then ask whether this function matched the bases used in the controller.

The adaptation in the controller was as follows. After each simulated




Figure 1. A, Experimental setup and the coordinate system for simulations of human arm and robot arm dynamics. B–F, Force fields that we examined in this report. In each case, $F = V\dot{x}$. Units of V are in kilograms per second. B, Standard curl field, $V = \begin{bmatrix} 0 & 13 \\ -13 & 0 \end{bmatrix}$; C, opposite curl field, $V = \begin{bmatrix} 0 & -13 \\ 13 & 0 \end{bmatrix}$; D, curl-assist field, $V = \begin{bmatrix} 11 & 11 \\ -11 & 11 \end{bmatrix}$; E, saddle field, $V = \begin{bmatrix} 11 & -11 \\ 11 & -11 \end{bmatrix}$.

Table 1. This table gives the learning rate (η) used for different widths of the basis elements (σ) in the simulation to get a peak in the generalization function (B) of Equation 3 equal to 0.18

σ (m/sec)	η
0.1	9.918×10^{-5}
0.15	3.787×10^{-5}
0.2	2.010×10^{-5}
0.3	9.279×10^{-6}

movement, the associated force vectors ($w_{xy,i}$) for each basis were updated using gradient descent as explained above. Experience with the simulation proved that the learning rate, η , varied with σ . We used trial and error to find η so that learning was at the same rate (and consistent with subject learning rates) for all σ . This gave us the values shown in Table 1.

From each movement performed by the simulation, we extracted a measure of error ($y^{(n)}$) by comparing the position of that movement at peak speed with the position of the desired trajectory at peak velocity. The force produced by the robot at that velocity was $F = V\dot{x}$. Therefore, from each simulated movement, we extracted two vectors: a movement error $y^{(n)}$ and the force acting on the hand $F^{(n)}$. We fit Equation 3 to the

sequence of errors, forces, and movement directions ($\{y, F, k\}^{(n)}$) and found the parameters B , D , and initial conditions $z_l^{(1)}$ that produced the best fit.

Human experiments. Over the past few years, our laboratory has collected a large body of data from naive subjects who were all exposed to the same pattern of forces in their first two or three target sets and were then tested in various other paradigms. The data for these subjects ($n = 75$) has been published previously (Shadmehr and Brashers-Krug, 1997; Thoroughman and Shadmehr, 1999; Thoroughman and Shadmehr, 2000) and is reanalyzed here with our new approach. To this data set, we added new subjects ($n = 43$) who were recruited specifically for the current report. Subjects were healthy and their age range was between 18 and 55. Position, velocity, and force data were collected from the robot's sensors at 100 Hz. The subject's arm was supported by a sling in an effort to restrict the problem to two dimensions. All subjects gave written consent using a form approved by the Johns Hopkins University Institutional Review Board.

We estimated the desired trajectory separately for each subject in each direction by averaging movements made during the last 200 movements of null field training. We then calculated $y^{(s,n)}$ for each movement, n , and for each subject, s , by subtracting the position at peak velocity during that movement from the position at peak velocity of the desired trajectory in the appropriate direction. These errors were then averaged across subjects to generate the average error on each movement, $y^{(n)}$. We then fit Equation 3 to the sequence of errors, $y^{(n)}$, and extracted the parameters D , B , and $z_l^{(1)}$.

Our development and our data analysis assume that the desired trajectory is fixed and does not change with training in a field. This is why the trajectories during the null field can be used as an estimate of the desired trajectory; however, we also explored the consequences to our model of a violation of this assumption. In this case, we used the movements at the end of training in the field, when performance had stabilized, to estimate the desired trajectory. The details of this calculation are presented in section S.7 of the supplemental material.

Measuring statistical significance of the model. Once we determined parameters B , D , and $z_l^{(1)}$, we assessed goodness of fit. We used standard bootstrap techniques (Efron and Tibshirani, 1993) to determine confidence limits for our parameters (details in section S.8 of the supplemental material). We used 200 bootstrap resamplings of our subjects to generate these confidence limits and used the same resamplings for the confidence limits of both B and D .

Additionally, using B , D , $z^{(n)}$, and $F^{(n)}$, we could regenerate trial-by-trial performance of the system, $\hat{y}^{(n)}$. We compared this with the actual measured sequence $y^{(n)}$ to determine the quality of fit of the model. Our goodness of fit measure is on the basis of the statistic r^2 , the percentage of the variance explained. Usually, measuring quality of a linear fit to data in a vector space, r^2 is defined as:

$$r^2 = \frac{\|\hat{y} - \bar{y}\|^2}{\|y - \bar{y}\|^2}.$$

To generalize this definition to our data, we require an appropriate replacement for $\bar{y}$. Because $\bar{y}$ plays the role of a "baseline model" and determines the amount of variance that needs explaining, we replaced it with a baseline model, y_0 , that included only the $z_k^{(1)}$ term, whereas B and D are fixed at 0:

$$\begin{cases} y_0^{(n)} = -z_k^{(n)} \\ z_l^{(n+1)} = z_l^{(n)} \quad l = 1, \dots, 8. \end{cases}$$

We fit this model to the data to find $y_0^{(n)}$. We could then calculate:

$$r^2 = \frac{\sum_{n=1}^N \|\hat{y}^{(n)} - y_0^{(n)}\|^2}{\sum_{n=1}^N \|y^{(n)} - y_0^{(n)}\|^2}.$$

This works well as an estimate if we are guaranteed that the model, $\hat{y}^{(n)}$, is optimal in the least-squares sense. If it is not, r^2 can easily be > 1 , and it is more convenient to define r^2 as:

$$r^2 = 1 - \frac{\sum_{n=1}^N \|\hat{\mathbf{y}}^{(n)} - \mathbf{y}^{(n)}\|^2}{\sum_{n=1}^N \|\mathbf{y}^{(n)} - \mathbf{y}_0^{(n)}\|^2}.$$

We used this equation to define our goodness of fit. Note that this equation is equivalent to the previous definition for an optimal fit.

We were also interested in the relative importance of B and D in the fit. This question can be addressed by calculating the partial r^2 . One does this by including more parameters in the baseline model but still leaving out the parameter of interest. For example, a baseline model that does not include the B matrix would be (we label output predicted by a model that does not include a B matrix as $\mathbf{y}_B^{(n)}$):

$$\begin{cases} \mathbf{y}_B^{(n)} = D\mathbf{F}^{(n)} - \mathbf{z}_{k(n)}^{(n)} \\ \mathbf{z}_l^{(n+1)} = \mathbf{z}_l^{(n)} \quad l = 1, \dots, 8. \end{cases}$$

The difference between $\mathbf{y}_B^{(n)}$ and $\mathbf{y}^{(n)}$ represents the total additional variance that could be explained by the introduction of the parameter B . We can calculate the percentage of this residual variance that B actually does explain by using the equation:

$$r_B^2 = 1 - \frac{\sum_{n=1}^N \|\hat{\mathbf{y}}^{(n)} - \mathbf{y}_B^{(n)}\|^2}{\sum_{n=1}^N \|\mathbf{y}^{(n)} - \mathbf{y}_B^{(n)}\|^2}.$$

A similar method was used to generate $\mathbf{y}_D^{(n)}$ and r_D^2 . We used a test similar to the bootstrap called the randomization test (Manly, 1997) to test the significance of our r^2 parameters (details in section S.9 of the supplemental material).

Generalization and rotation. Our adaptive controller assumes that force errors generalize from one direction to another in Cartesian coordinates without rotation. This assumption is implicitly carried into our dynamic model, because $\mathbf{y}^{(n)}$ is only multiplied by a scalar $B_{l,k}$ in updating the internal model state, $\mathbf{z}_k^{(n)}$. Thus, our dynamic model imposes on the data an assumption that generalization never involves rotation. In contrast, Thoroughman and Shadmehr (2000) assumed that the error vector would be rotated during generalization in proportion to the extent of generalization. We tested these two assumptions against the data using a model that allowed rotation. This model is identical to the one in Equation 3 except that $B_{l,k}$ is a 2×2 matrix for each l and k . To limit the number of parameters we used a 2×2 matrix of the form:

$$B_{l,k} = \begin{bmatrix} \alpha_{l,k} & -\beta_{l,k} \\ \beta_{l,k} & \alpha_{l,k} \end{bmatrix}.$$

Using the transformations $\text{Mag}_{l,k} = \sqrt{\alpha_{l,k}^2 + \beta_{l,k}^2}$ and $\theta_{l,k} = \tan^{-1}(\beta_{l,k}/\alpha_{l,k})$, we can extract a magnitude (equivalent to the scalar B from Eq. 3) and a rotation angle from this 2×2 matrix. This allowed us to ask whether errors generalize with or without rotation.

Task variations. We designed a number of variants of the experiment so we could answer two questions. (1) How well does the theory agree with trial-to-trial variations in performance under different conditions? (2) If the theory is robust, do the parameters of the system that it identifies remain invariant across repeated measures and across different conditions? In the first condition, 75 subjects trained in the clockwise curl field. Data for 40 of these subjects was presented previously in Thoroughman and Shadmehr (2000) and is reanalyzed here using our new methods. In all of these 75 subjects, the target set was out and back; that is, a target at 45° was always followed by a target at 225° . The odd-numbered targets were at 0, 45, 90, or 135° , whereas the even-numbered targets were at 180, 225, 270, and 315° . In addition, the same sequence of 192 targets was used in three successive sets performed by the subjects.

To test whether serial correlations in the original out-and-back arrangement influenced our results, subjects ($n = 9$) trained in a target set in which movement directions were selected pseudorandomly from the eight possible directions. Before each movement, the robot's motors were used to move the subject's arm passively to the appropriate starting point for that movement. We also tested the applicability of our model

when movement errors were not only perpendicular to the direction of target but also parallel to it. We had subjects train in a curl-assist field ($n = 8$) or in a saddle field ($n = 9$). Target directions did not include serial correlations, and the motors were used to move the hand to the appropriate starting point before each movement.

Finally, note that in our theory there is no requirement that errors converge. In fact, our approach to system identification is at its best when the system of interest is producing large errors. This suggested that if the bases with which the brain adapts are consistent across tasks, the bases should be invariant even in a task where forces randomly change from trial to trial. Subjects ($n = 12$) reached in a field in which forces were selected pseudorandomly to be either the clockwise curl field (42% of trials), the counterclockwise curl field (42% of trials), or a null field (16% of trials). The target directions were arranged in an out-and-back structure as in the original paradigm.

In addition to estimating generalization in each task variation, we also asked how well our estimates of B from the first experiment (variation 1, curl field correlated set) could explain the data from all other experiments (variations 2–5). For this validation, we fit the data for variations 2–5 by restricting the eight parameters of B in Equation 3 to be a scaled version of the parameters that we found in variation 1. Thus, we required the shape of B to remain invariant across conditions but allowed it to be scaled. Therefore, rather than eight parameters, a single parameter, the scaling factor, was used. The 12 unknown parameters (4 for D , 8 for B) of the fit are reduced to 5 in this case.

Results

Estimating generalization of an artificial adaptive controller

We began with an artificial adaptive controller that learned to move a model of a human arm attached to a robotic arm that produced a force field. The adaptive controller improved its performance by building an internal model. The internal model was a collection of Gaussian basis functions, where each basis encoded arm velocity and was associated with a preferred force vector. Adaptation resulted in a change in the preferred force vector. The coding of the velocity space was with Gaussians of various width σ (Eq. 2). Thus, we expected the generalization function to be progressively wider as the width of the basis functions increased.

To test this idea, we fit the dynamic system of Equation 3 to the trial-by-trial performance of the adaptive controller. We tried controllers that had basis elements with widths varying between 0.1 and 0.4 m/sec and examined the quality of the fit and the parameters of the fit. Figures showing the trial-by-trial performance of the controller and the fit to the model are shown in Figures S1 and S2 of the supplemental material. The r^2 values for these fits were always >0.9 . This suggests that despite the nonlinear dynamics of the simulated arm and its adaptive controller, an appropriate linear dynamic system could effectively describe the trial-by-trial data during adaptation.

The estimated compliance matrix D generated from these fits is shown graphically in Figure 2A for simulations with different values of σ . As would be expected, when basis elements change, no change is observed in D . We estimate that D captures 0.974 of the otherwise unexplained variance in fitting the model to a simulation with $\sigma = 0.1$, and 0.954 of the otherwise unexplained variance in fitting a simulation with $\sigma = 0.3$ m/sec.

Figure 2B shows the generalization function, B , estimated from simulations with different σ . This function describes what percentage of error in one movement direction is generalized to another direction at some angular distance. Its width is a reflection of the width of the basis functions in the simulation. When bases are narrow, there is little generalization to neighboring directions. A wider σ in the simulation produces a wider generalization function. The generalization function, B , captures 0.94 of




Figure 2. *A, B*, Parameters of the fit using Equation 3 on data from simulations with different σ . *A*, A graphical representation of the matrix D . D is a compliance matrix that transforms force on the hand into displacement. To represent this, we multiply D by a force vector of length 1 N as the vector rotates about a circle. The result is hand displacement. The coordinates of the force vector are Cartesian, centered at the position of the hand (data not shown). The intersecting lines show the effect of D on the coordinate axes of force. *B*, The generalization function, B , characterizes the effects of error in one direction on all other directions of movement. The x-axis shows the difference (in degrees) between the direction in which the error was experienced and the other directions of movement. The y-axis is unitless gain. It describes what portion of the error vector (in terms of a ratio) was distributed to the state in which the movement occurred (0°) and neighboring states. Note how wider Gaussian basis elements (larger σ) produce a wider generalization function. *C*, A comparison of the generalization function found by two different methods for basis elements of different widths. Solid lines come from fitting the dynamic model (Eq. 3); dashed lines were calculated using the numerical derivation (Eq. 4). The results from the numerical derivation have been scaled down by a factor of 1.3. The factor was apparently introduced by the approximations made in the derivation.

otherwise unexplained variance for $\sigma = 0.1$ and 0.912 for $\sigma = 0.3$ m/sec.

Our theory predicts the values of B shown in Figure 2*B*: Equation 4 gives the value of the generalization function in terms of the parameters of the adaptive controller. We test this prediction by calculating these predicted values for the generalization function for different σ . Our numerically computed values were larger by a factor of ~ 1.3 relative to the values found by extracting the parameters from the sequence of errors. Figure 2*C*, however, shows that after adjusting for this scaling factor, there is good agreement between the two methods. The scaling factor probably results from the approximations made in the derivation. The reader is referred to the supplemental material for a full derivation as well as a summary of these approximations. We suspect that the scaling factor results from a simplification of the movement trajectory that we used to aid in estimating the value of an integral. What is clear, however, is that despite the approximation, we are able to use Equation 3 to find the shape of the generalization function.

The generalization function in human data

We next applied our theoretical framework to reaching movements of humans. Our first group of subjects performed either two or three target sets, each consisting of 192 movements. Every movement began from the center of the workspace outward followed by a movement back to the center. All subjects experienced the clockwise curl field (Fig. 1*B*).

In Figure 3 we show movement errors for the first target set, averaged across subjects. The figure shows the X and Y components of the movement error for each direction separately; it shows the actual sequence of errors, combining all directions, for the first 75 movements in the middle plot (Fig. 3*I*). The percentage variance explained by the model (r^2) is 0.77, which is significant at $p < 0.01$. Movement errors for all three target sets can be seen in Figure 4*A, C, E*. Here, we display the parallel and perpendicular components of the error, rather than the X and Y components. Movements to 180° are shown separately (Fig. 4*B, D, F*). The r^2 values for the model fit are 0.77, 0.80, and 0.77, respectively, for sets 1–3 (all of which are significant at $p < 0.01$).

We can think of the three sets in this experiment as a kind of

repeated measure: the same subject is being retested with the same target set. This tests whether the underlying adaptive system continues to use the same basis set, although performance varies from set to set. For example, note that in set 1, performance improves over the course of the set, whereas in sets 2 and 3 the performance has reached a plateau. Therefore, although the theory can fit the data quite well in each case, it is more interesting to ask whether the generalization function remains constant. Figure 5 allows comparison of B and D for the models generating the fits shown in Figure 4. The limb's compliance, as estimated by D , did not change appreciably. Interestingly, the shape of the generalization function B also remained similar in the three data sets. Error bars in Figure 5*B* represent the bootstrapped SEM of each parameter of B . We found that the peak value of B (at 0°) in the first set was just outside of the 95% confidence interval for the B in the third set. This suggests that from the first to third set there may have been a reduction in the degree to which errors generalized to the same direction (i.e., 0°).

What is the relative importance of B and D in producing the fit shown in Figure 4? The partial r^2 shows that D explains 0.74, 0.80, and 0.76 of the otherwise unexplained variance in the three sets, respectively. B explains 0.28, 0.23, and 0.22 of the otherwise unexplained variance. For both B and D , bootstrap statistics show that the partial r^2 values are significant at $p < 0.01$. Thus, the performance of subjects is well explained by a fixed compliance matrix and a fixed generalization function. Another way of demonstrating the contribution of B to the quality of the fit is to plot the model against the data (as in Fig. 4) but to subtract out those components of the fit that come from the other parameters. This is what we show in Figure 6, where we plot:

$$\begin{aligned} \mathbf{y}_B &= \mathbf{y}^{(n)} - D\mathbf{F}^{(n)} - \mathbf{z}_{k(n)}^{(0)} \\ \hat{\mathbf{y}}_B &= \hat{\mathbf{y}}^{(n)} - D\mathbf{F}^{(n)} - \mathbf{z}_{k(n)}^{(0)} \end{aligned}$$

for sets 1–3. These plots eliminate the effects caused by differences between field trials and catch trials and any effects caused by differences between different directions. They demonstrate that many features of the data are quite well captured by the model.




Figure 3. Averaged performance of subjects ($n = 75$) in a clockwise curl field with occasional catch trials. X and Y components of the data are shown in gray; fit for the model (Eq. 3) is in black. Circles indicate catch trials. A–H, Movements in each of the eight directions during a 192-movement set. I, All directions of the first 75 movements of the set. Standard errors (data not shown) for the measured data are on the order of 1 mm. The r^2 for the fit of the model over the entire 192-movement set is 0.77.

Remarkably, the generalization function is significantly bimodal: it drops off toward 45° but rises back by 180° ; however, although the bimodality is significant, the shape of the generalization function also seems compatible (judging from the error bars on the values of B) with a constant, i.e., a straight line at a value of ~ 0.12 . We tested whether a model like this could explain the data. We changed our model so that the B matrix had only a single parameter. This meant that the effect of errors in any movement on movements in all other directions was identical. We fit this model to our data in the first set. The fit produced r^2 of 0.71 and a partial r^2 for the B parameter of 0.10. Therefore, although the partial r^2 for D changed by $<5\%$, the partial r^2 for B dropped by 64%. Bootstrap shows that the partial r^2 is signifi-

cantly less with this approach than in our original model; however, it appears to us that perhaps the best evidence for bimodality is the consist shape of the B function across repeated measures in these three sets. As we will see, this bimodality is maintained in other variations of the task.

If we compare the compliance of our simulated adaptive controller (Fig. 2A) with the compliance that we measured from our subjects (Fig. 5A), we see that there is similarity in shape, although the human arm is not as stiff as the model. This suggests that the stiffness and inertial parameters that were used in the simulation are realistic but too large in magnitude to accurately represent the real human arm.

It is striking, however, that the shape of the generalization




Figure 4. Averaged performance of subjects in a clockwise curl field. Errors parallel (Par) to the movement direction and perpendicular (Perp) direction are shown. Gray is subject data, and black is model fit. Circles indicate catch trials. *A, B*, Errors in the first fielded set ($n = 75$ subjects). *C, D*, Errors in the second fielded set ($n = 75$ subjects). *E, F*, Errors in the third fielded set ($n = 44$ subjects). The sequence of targets and the catch trials are the same in each set. r^2 values for all movements within a set are 0.77, 0.80, and 0.77, respectively, for sets 1–3. Note that a “learning curve” can be seen clearly only in the first half of the first set.




Figure 5. Parameters of the models that produced the fit to subject data shown in Figure 4. *A*, The D matrixes are illustrated following the format of Figure 2*A* as the effect of the matrixes on the unit circle. *B*, The generalization functions for the three successive sets performed by the subject. Error bars are bootstrapped standard errors.

function B in human data is distinctly different from the one generated by an adaptive controller using our Gaussian bases. Symmetrical Gaussian bases generalize unimodally instead of bimodally (Fig. 2*B*). Therefore, the Gaussians that we used for coding of velocity space are incompatible with the pattern of generalization demonstrated by subjects (Fig. 5*B*). Because our paradigms are designed to test only generalization across directions and not across speeds, we cannot know exactly what the shape of the basis elements should be; however, we can try to find a basis set that would be compatible with our results. In an effort to find such a basis set, we modified the simulated controller to use bases that encoded limb velocity with a bimodal function:

$$g_j(\dot{\mathbf{x}}) = \exp\left(-\frac{\|\dot{\mathbf{x}} - \dot{\mathbf{x}}_j\|^2}{2\sigma^2}\right) + \frac{1}{K} \exp\left(-\frac{\|\dot{\mathbf{x}} - (-\dot{\mathbf{x}}_j)\|^2}{2\sigma^2}\right).$$

This function has a preferred velocity at $\dot{\mathbf{x}}_j$ where activation is maximum, and a second but smaller peak at $-\dot{\mathbf{x}}_j$. The unitless

constant K controls the relative size of the secondary peak. Figure 7*A–D* compares the generalization functions generated using the modified simulation with those generated using subject data. The best fit to the subject data seems to be where $K = 2$ and $\sigma = 0.15$ (Fig. 7*B*, dark line). Using nonlinear fitting in Equation 4, we worked backward to estimate the σ and K that most closely matched the subject data and found $\sigma = 0.20$ and $K = 1.7$. This is quite close to the parameters that we found by searching the parameter space. The predicted shape of the basis is shown in Figure 7*E*. The basis has a peak of activation for one direction (its preferred direction) and a second peak with approximately half the magnitude in the opposite direction.

Therefore, from the pattern of generalization that we estimated from human data we infer that the bases of the internal model have this bimodal shape; however, we note again that this is only one possible shape that could produce the same generalization function, and the current results are not sufficient to determine the exact shape of the basis elements.

Are error vectors rotated as they are generalized?

One of the fundamental differences between our approach and that of Thoroughman and Shadmehr (2000) is that we used a vector rather than a scalar measure to describe error in each movement. In the earlier work, the error measure was perpendicular displacement of hand trajectory with respect to the target direction. The error was positive if it was clockwise with respect to movement direction and negative otherwise. It was assumed that this error is generalized to all other directions also as a perpendicular displacement, which seemed reasonable because the field




Figure 6. The quality of the fit caused by the generalization function (B) is demonstrated by plotting $y - Df - z^{(0)}$ (gray) and $\hat{y} - Df - z^{(0)}$ (black), thereby eliminating from the fit the contribution of both the compliance matrix D and any average difference between the directions. As can be seen, many of the details of the remaining variation are still being captured. Format is the same as Figure 4. Par, Parallel; Perp, perpendicular.

was always perpendicular to the direction of movement. Therefore, that measure of generalization implicitly assumed a rotation of the error vector to other directions. The result was a generalization function that was broad but became negative at 180° . This is consistent with behavior of wide Gaussian-like bases when movement errors are represented as scalars.

In contrast, when we represented movement errors as vectors, theory predicted that generalization should involve a scaling of this vector with no rotation. This is a testable prediction of the theory and allows us to directly compare our approach with the earlier approach.

We changed our dynamic model so that its measure of generalization would allow for both scaling and rotation of the error vector. Because of the additional parameters, the model fit the data slightly but not significantly better than the original model, with $r^2 = 0.79, 0.80$, and 0.78 for the three curl field sets, respectively. For all three sets, a consistently small rotation was found to fit the data best, with an average rotation of $10^\circ (\pm 2^\circ \text{SEM})$, $6^\circ (\pm 3^\circ)$, and $9^\circ (\pm 4^\circ)$. The maximum rotation across all three sets was 22° . The results demonstrate that when movement error is represented as a vector, it is generalized to other directions with little or no rotation. Lack of rotation contradicts an assumption of the Thoroughman and Shadmehr (2000) model.

We found an additional way to test this question and also to validate our findings in a manner that is model independent. We searched in our data for examples in which two movements in the




Figure 7. *A–D*, Generalization functions were estimated from data that were generated using simulations of an adaptive controller that uses bimodal primitives in its internal model. Generalization functions are shown for different values of the two parameters: K , the ratio of the peak heights, and σ , the width in meters per second. Superimposed on them (dashed black line) is the generalization function found in the subject data, averaged over the three generalization functions in Figure 5*B*. *E*, The shape of the primitives implied by the generalization functions found in the subject data (Fig. 5). This is the activation function of a single primitive the center of which is at $[0.21, 0.21]$ (approximately peak velocity for a 10 cm movement to 45°). For this primitive, $\sigma = 0.15$ and $K = 2$. The actual prediction of force is computed by an average of the force associated with each primitive, weighted by the activation. The axes of this graph are X velocity and Y velocity, and the limits are from -0.5 to 0.5 m/sec for both X and Y.

same direction were separated (among other movements) by a movement in some specific direction. The intervening movement was a catch trial, and the two movements in the same direction were field trials. We then calculated the difference in position at maximum velocity of these two movements as a function of the direction of movement of the intervening catch trial. The result of this analysis is shown in Figure 8. This figure shows that when the intervening catch trial was in the same direction of movement as the two field trials (vectors at the top of this figure), the effect of the error in the catch trial was an increase in error in the field trial (depicted by the black vector). Because the field is a clockwise curl, the effect of the intervening catch trial on the subsequent field trial is an error to the right. Now, if the intervening catch




Figure 8. A model-free assessment of the generalization function in a clockwise curl field. We found all occurrences where in the intervening trials between two field trials in the same direction there was a catch trial (field-catch-field) in some other direction. This plot shows the vectorial difference between the position at maximum velocity of the first and last field trials. The top vector shows cases in which the middle trial is in the same direction as these trials. The next vector clockwise shows cases in which the middle trial was 45° clockwise from the direction of these trials.

trial is not in the same direction of the two field trials, for example at 45° with respect to the field trials, its effect remains roughly perpendicular to the direction of movement in the intervening catch trial. This means that the error experienced in the catch trial is a vector that undergoes a scaling but little or no rotation as it is generalized to other directions of movement. Notice, however, that this scaling is neither constant nor unimodal: if the catch trial is at a direction 180° to the field trials, its effect appears larger than when the directional difference is 90°. This result demonstrates that bimodality can be found directly in the underlying data.

To summarize, (1) during learning of a curl force field, the compliance matrix D remains invariant across different movement sets, (2) the generalization function B is consistently bimodal as a function of movement direction, (3) error experienced in a given direction generalizes to other directions without a significant rotation, and (4) the shape of B is inconsistent with bases that encode velocity space unimodally as a Gaussian. Rather, the generalization function suggests basis elements that are bimodal with a secondary peak at a direction opposite to the preferred direction and a magnitude of approximately half the magnitude of the main peak.

Generalization function in different force fields and different target sets

We tested the model's robustness by collecting data in a number of variations of the original task. In the first variation, we examined whether the bimodal shape of B was an artifact of the structure of our target sets. It was conceivable that the out-and-back structure of the original paradigm contributed to the secondary peak at 180°. To address this, we tested nine subjects on the same curl field using a target set without the serial structure. In the new paradigm, targets appeared in a pseudorandom sequence of directions (0°, 45°, ..., 315°). After each movement, the robot passively moved the subject's arm to a new start position. The first row of Figure 9 (plots A and B) shows data from these subjects and the fit of the model (Eq. 3) to that data. The fit has an r^2 of 0.72.

The parameters that go with this fit are shown in the first row of Figure 10. The smaller number of subjects leads to a noisier




Figure 9. Errors and the fit of the model (Eq. 3) to a set of 192 movements in variations of the clockwise curl field task. Each row of plots represents a different variation. A, B, Movements in a clockwise curl field using target sets in which the out-and-back structure has been removed; $r^2 = 0.72$ ($n = 9$). C, D, Movements in a curl-assist field; $r^2 = 0.74$ ($n = 8$). E, F, Movements in a saddle field; $r^2 = 0.68$ ($n = 9$). G, H, Movements in a random field; $r^2 = 0.82$ ($n = 12$). Format for plots is as in Figure 4. Par, Parallel; Perp, perpendicular.

generalization function. Nevertheless, the bimodality persists despite the lack of serial structure in the target set. Similarly, the shape of the compliance matrix, as measured by the D parameter, is similar to the values in the standard paradigm.

An additional reservation that could be entertained is that it is possible that the curl field is a very special case. Other velocity-dependent force fields that generate more parallel error may not produce such a clean fit to the model. To address this possibility, we trained subjects on a curl-assist field and a saddle field. These fields produce parallel as well as perpendicular errors in the hand's trajectory. Performance in the curl-assist and saddle fields are shown in Figure 9C–E. The r^2 was 0.74 for the curl-assist field and 0.68 for the saddle field.

The compliance and generalization parameters estimated in each of these fits are shown in Figure 10, C and D, for the curl-assist field and Figure 10, E and F, for the saddle field. For the saddle field, D is significantly more elongated than in the curl field. This may be attributed in part to anisotropic characteristics of the field, which is assistive along the long axis of the compliance matrix and resistive along the short axis. If the viscous response to assistive and resistive fields is not symmetric, this could




Figure 10. The compliance and generalization parameters for the fits of Equation 3 to data collected in the variations on the original task. *A, B*, Clockwise curl field with uncorrelated target set structure. *C, D*, Curl-assist field. *E, F*, Saddle field. *G, H*, Random field. Left column represents the compliance matrix *D*; right column is the generalization function *B*. Format is as in Figure 2.

produce the elongated compliance matrix. A difference that is evident in the curl-assist field is a marked decrease in the magnitude of the generalization function; however, the bimodal shape of the function is still apparent in both fields.

Validating the generalization function

The generalization that we found for all of the task variations seemed bimodal. We validated this impression by fitting the data again but forced the generalization function to be identical in shape to the one that we found in the clockwise curl field task (that is, as described in Materials and Methods, we fit a scaled version of the function). The r^2 of the fit to the target set with serial correlations removed was 0.68, and to achieve this fit the generalization function was scaled by a factor of 0.65. The fit for the curl-assist field gave an r^2 of 0.68, whereas the r^2 for the saddle field fit was 0.66. The scaling factor applied to the generalization function to achieve these fits was 0.15 and 0.91, respectively. This repeats the finding above that the magnitude of the *B* in the curl-assist field was smaller than in the other fields, as suggested by Figure 10.

Adapting in an unpredictable force field

When the task is unpredictable, do subjects continue to adapt their internal model from movement to movement as before? Our theory provides a method to answer this question. We tested subjects in a target set in which the field varied randomly from

movement to movement. Because the field was unpredictable, the task was effectively not learnable in the sense that no long-term improvements in performance were possible, except through increased stiffness of the arm. The important questions are whether the behavior of subjects will continue to agree with the behavior of the model, and if so, whether the system parameters (*B* and *D*) will be similar to those generated in the more traditional learning paradigm.

During the unpredictable set, errors remained large through the end of the set except on catch trials, in which errors were small (Fig. 9G). There was no evidence for improvement in performance. Remarkably, the model continued to fit the trial-by-trial performance quite reliably ($r^2 = 0.82$). The partial r^2 of the *B* function indicated that this parameter was still significant in explaining the pattern of errors. For these data, *B* explains 0.11 of the otherwise unexplained variance. Also, the shape of the generalization function was similar to the shape found earlier (Fig. 10H). We validated this by fitting a scaled version of the generalization function from the standard curl field learning paradigm. In this case, the r^2 was 0.82 and the scaling factor was 0.47.

Despite the fact that the data in this random field task would suggest no adaptation, in light of the results provided by the theory it appears that the brain is adapting in similar terms as before. The internal model is responding to errors by changing its state, but because the errors are random, the changes do not accumulate.

Allowing change in the desired trajectory

In our theoretical development, we make the crucial assumption that the desired trajectory in the field is the same as the desired trajectory during the baseline, null field training movements. Within the context of current motor control research, this is a reasonable assumption, but we have no way of determining whether it is true. If we relax this assumption, we need some way of determining the effective desired trajectory after the field has been turned on. One option is to estimate the desired trajectories from the movements at the end of each training set. This is like assuming that, at the end of each set, subjects are moving as closely as possible to their desired trajectory. We detail the calculations used to generate this approximation in the supplemental material.

The new desired trajectory changes the reference point for measurement of error, thus changing the values of γ in Equation 3. This, in turn, changes the values of the parameters *B* and *D* that best fit the data. The maximum change for the reference point for the three sets of curl field data were 3.8, 3.2, and 3.6 mm, respectively. For the other fields, the changes were slightly larger but of similar magnitude (4.7–15.3 mm; the largest difference being for the saddle field). The desired trajectories originally estimated from the null movements and the desired trajectory estimated from the end of training are shown in Figure 11. It is clear from the figure that these differences are quite small. Nevertheless, the generalization function was affected by the change (Fig. 12). Although the overall shape of the generalization function is similar, the bimodality has been reduced or eliminated. Importantly, allowing the desired trajectory to change improved the fit of the model. For the overall r^2 , this change was quite small. The overall r^2 values achieved when combining the original curl field data were 0.83, 0.84, and 0.80 for the three sets, compared with 0.77, 0.80, and 0.77 with a fixed desired trajectory. For the other variants on the paradigm, the overall r^2 was between 0.71 and 0.84; however, a significant improvement occurred in the ability of the generalization function to explain the error-dependent changes




Figure 11. A comparison of the desired trajectory as estimated by the movements at the end of null field training (Null; gray) and by the movements at the end of training in a field (Training; black). The first column shows the movement trajectories (with a dot at the position of maximum velocity). The second column shows the speed profiles.

in trial-to-trial performance. For the first set of the standard curl field task, r^2_B went from 0.25 to 0.48, and for the random set the r^2_B went from 0.10 to 0.18. Similar results were found when analyzing the other data sets (r^2_B of ~ 0.35). On average, when we allowed the “desired trajectory” to vary with dynamics of the task, the generalization function accounted for approximately twice as much of the variance as when the desired trajectory was estimated from the null trials.

Discussion

This paper outlines a new approach to understanding human motor learning. We studied the underlying representation used




Figure 12. Allowing change in the desired trajectory improved the ability of the dynamic model to fit the data. Each plot shows the generalization function using a changed desired trajectory for each of the variations of the paradigm. The r^2_B for each of the fits is 0.48, 0.38, and 0.35 (A) for the three sets 0.35 (B), 0.34 (C), 0.35 (D), and 0.18 (E).

by subjects to learn a novel motor task, and we used an approach derived from the theoretical study of an adaptive controller that mimicked human behavior. Our theory allowed us to circumvent the traditional need to first train the system in one state and then test generalization to another state. Instead, we demonstrate that it is possible to measure generalization by visiting all states in a random order. Our central finding was that during reaching movements in velocity-dependent fields and under the assumption of a fixed desired trajectory, error in a given movement direction was generalized in a bimodal pattern to neighboring directions. This suggests the possibility that the internal model is represented with computational elements that encode limb velocity with a bimodal receptive field; however, an important caveat is that this bimodality depends critically on the assumption that the desired trajectory does not change because of the force field.

Our theory advances the idea that one can profitably study mechanisms of adaptation without improvements in performance. In our analysis, what matters is not whether performance is converging, but how performance on one trial depends on the errors experienced on the previous trial. There is a long tradition of modeling the relationship between past and current behavior using linear models (Falmagne et al., 1975; Kowler et al., 1984; Smeets and Brenner, 1995; Thoroughman and Shadmehr, 2000; Scheidt et al., 2001; Witney et al., 2001). In most studies of this sort, it has been found that the current movement is affected by only a small number of previous movements, often only one. Relying on these studies, our theory describes adaptation as a multidimensional hidden state changing because of error experienced in the previous movement. In effect, the system has a memory of one trial. This is consistent with the results of Scheidt et al.

(2001) who tested explicitly for the length of the system memory in a task very similar to ours.

Our theory extends these studies. We demonstrate that it is not only errors that affect trial-to-trial performance changes but also the basis functions used for representation of the task. Our equations characterize adaptation using two parameters: a compliance matrix, D , and a generalization vector, B . We fit these equations to our data to find the shape of the generalization function. We found that under the assumption of a fixed desired trajectory, B was consistently bimodal in subjects experiencing different sequences of fields and different sequences of targets. B remained bimodal even in a randomly changing field in which standard measures of performance would detect no learning.

Although we found that generalization was consistently bimodal, it was not identical across different data sets. In repeated measures across the same subjects, the peak value of B appeared to decline from the first to the third target set. This suggests that on the third set, errors cause a smaller change in the internal model than on the first set. This would indicate an increased resistance of the internal model to change. This change in the peak value of B disappeared when we allowed adaptation of the desired trajectory, however, and further experiments are required to determine how it should best be interpreted. We also found a significantly smaller B in the curl-assist field than in the curl field. Although this is difficult to interpret, we note that our ability to measure error in a subject's movement is far more robust for perpendicular displacements than for parallel displacements. This relative uncertainty in our measures of parallel error might be responsible for the reduced magnitude of B in the curl-assist field.

We believe that our model is a reasonable first step in explaining the human data, but it is by no means optimal. Our approach captured nearly all of the variance in the data generated by the simulated adaptive controller, but, on human data, our model did no better than $\sim 80\%$. Moreover, although the partial r^2 value for B in the human data was always significant, it was never >0.4 , whereas in the simulated data it was ~ 0.9 . Improvements in the model, especially the introduction of nonlinearity, could increase its ability to explain the data; however, to continue on a solid theoretical foundation, these improvements should be linked to an explicit description of the adaptive controller that they would imply. Otherwise, we create a real danger of increasing fit by increasing the number of parameters without actually increasing our understanding of the task. One improvement that is supported by experimental data would be to allow the compliance matrix, D , to change because error seems to be related to increased stiffness (Thorouhman and Shadmehr, 1999). Another improvement that we have explored here is that the desired trajectory may change. Allowing the desired trajectory to change increased the partial r^2 for B by approximately a factor 2. It also increased the consistency of the generalization function between different variants of the paradigm; however, changing the desired trajectory reduced the bimodality of the generalization function. Determining whether the desired trajectory does indeed change, and how it changes, seems to be a key factor in further development of this model.

Any relationship between our results and neurophysiological results is necessarily tentative. Nevertheless, it is intriguing to note that a primate study found that Purkinje cells in the cerebellum were bimodally tuned with respect to hand velocity during reaching movements (Coltz et al., 1999). In the same task, the bimodality was not present in motor cortical cells. The Purkinje cells, unlike neurons in M1, seemed to respond to a preferred

speed as well as a preferred direction (Johnson et al., 1999). It is noteworthy that in our task, cerebellar damage profoundly impairs the ability of subjects to adapt (Smith, 2001). This is consistent with the inference that we made regarding encoding of limb velocity in the internal model. Indeed, these findings motivated our choice of basis elements. Because our paradigm does not measure generalization across different speeds, we cannot know the exact shape of the basis elements. Our data may well have been consistent with other choices for the basis elements that would also have produced similar generalization functions. For instance, we could have used a basis composed of asymmetrical Gaussian elements that are wide along the line connecting their center to the coordinate axis (causing generalization across 0) but narrow in the direction perpendicular to that line. We chose the specific shape that we used to match the data presented in Coltz et al. (1999).

The question of whether there exists some simple set of primitives that the CNS might rely on to perform complex actions has been approached in a number of different ways. Some approaches have used principal components or other related techniques to decompose the kinematics of movements. They have asked whether there exist a small number of movement patterns that can be combined linearly to produce the kinematics of the original movements. It has been found that much of the variance in hand writing (Sanger, 2000), grasping of the hand (Santello et al., 1998), typing movements (Soechting and Flanders, 1997), and gait (Olney et al., 1998) can be explained with the first few principal components of the original data set. These results suggest that a low-dimensional description of movement exists, but the approach has yet to address directly questions of representation, planning, or control of movements.

A second approach has been inspired by reports from Bizzi and colleagues (Giszter et al., 1993; Saltiel et al., 2001) on force-field representation in the spinalized frog. Stimulation of distinct sites in the spinal cord has been shown to produce one of four or five distinct types of motor output, quantified as forces as a function of limb position. Because co-stimulation of any two of these neural locations would give rise to a nearly linear combination of the corresponding force fields (Mussa-Ivaldi et al., 1994), it has been suggested that the CNS may control action via a weighted, linear combination of these fields (Mussa-Ivaldi and Giszter, 1992). This theory describes not only how actions might be generated, but also how they may be learned.

Our contribution has been to find a way to test the idea that learning of limb dynamics may be through a linear combination of task-invariant basis elements. The shape of these elements appears to remain fairly invariant across task variations. It is noteworthy that previous work has found evidence for time-dependent changes in functional properties of motor memory (Brashers-Krug et al., 1996; Shadmehr and Brashers-Krug, 1997) and brain activation patterns (Shadmehr and Holcomb, 1997; Nezafat et al., 2001) during consolidation. If the internal model changes in its neural organization as it becomes a part of long-term memory, do the generalization patterns, reflecting the underlying basis, also change?

References

- Ahissar M, Hochstein S (1997) Task difficulty and the specificity of perceptual learning. *Nature* 387:401–406.
- Bosco G, Rankin A, Poppele R (1996) Representation of passive hindlimb postures in cat spinocerebellar activity. *J Neurophysiol* 76:715–726.
- Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. *Nature* 382:252–255.
- Coltz JD, Johnson MT, Ebner TJ (1999) Cerebellar Purkinje cell simple

- spike discharge encodes movement velocity in primates during visuomotor arm tracking. *J Neurosci* 19:1782–1803.
- Conditt MA, Mussa-Ivaldi FA (1999) Central representation of time during motor learning. *Proc Natl Acad Sci USA* 96:11625–11630.
- Conditt MA, Gandolfo F, Mussa-Ivaldi FA (1997) The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78:554–560.
- Crisicimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R (2002) Learning dynamics of reaching movements generalizes from dominant to non-dominant arm. *J Neurophysiol* 89:168–176.
- Donchin O, Shadmehr R (2002) Linking motor learning to function approximation: learning in an unlearnable force field. In: *Advances in neural information processing systems* (Diettrich TG, Becker S, Ghahramani Z, eds), pp 197–203. Cambridge, MA: MIT.
- Efron B, Tibshirani RJ (1993) *An introduction to the bootstrap*. New York: Chapman and Hall.
- Falmagne JC, Cohen SP, Dwivedi A (1975) Two-choice reactions as an ordered memory scanning process. In: *Attention and performance V* (Rabbitt PMA, Dornic S, eds), pp 296–344. New York: Academic.
- Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703.
- Giszter SF, Mussa-Ivaldi FA, Bizzi E (1993) Convergent force fields organized in the frog's spinal cord. *J Neurosci* 13:467–491.
- Goodbody SJ, Wolpert DM (1998) Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79:1825–1838.
- Imamizu H, Uno Y, Kawato M (1998) Adaptive internal model of intrinsic kinematics involved in learning an aiming task. *J Exp Psychol Hum Percept Perform* 24:812–829.
- Johnson MT, Coltz JD, Ebner TJ (1999) Encoding of target direction and speed during visual instruction and arm tracking in dorsal premotor and primary motor cortical neurons. *Eur J Neurosci* 11:4433–4445.
- Kowler E, Martins AJ, Pavel M (1984) The effect of expectations on slow oculomotor control. IV. Anticipatory smooth eye movements depend on prior target motions. *Vision Res* 24:197–210.
- Lackner JR, DiZio P (1994) Rapid adaptation to coriolis force perturbations of arm trajectory. *J Neurophysiol* 72:299–313.
- Lee DK, Koch C, Braun J (1999) Attentional capacity is undifferentiated: concurrent discrimination of form, color, and motion. *Percept Psychophys* 61:1241–1255.
- Manly BFJ (1997) *Randomization, bootstrap and Monte Carlo methods in biology texts in statistical science series*, Ed 2. London: Chapman and Hall.
- Mussa-Ivaldi FA, Giszter SF (1992) Vector field approximation: a computational paradigm for motor control and learning. *Biol Cybern* 67:491–500.
- Mussa-Ivaldi FA, Giszter SF, Bizzi E (1994) Linear combinations of primitives in vertebrate motor control. *Proc Natl Acad Sci USA* 91:7534–7538.
- Nakano E, Imamizu H, Osu R, Uno Y, Gomi H, Yoshioka T, Kawato M (1999) Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. *J Neurophysiol* 81:2140–2155.
- Nezafat R, Shadmehr R, Holcomb HH (2001) Long-term adaptation to dynamics of reaching movements: a PET study. *Exp Brain Res* 140:66–76.
- Olney SJ, Griffin MP, McBride ID (1998) Multivariate examination of data from gait analysis of persons with stroke. *Phys Ther* 78:814–828.
- Poggio T, Fahle M, Edelman S (1992) Fast perceptual learning in visual hyperacuity. *Science* 256:1018–1021.
- Pouget A, Sejnowski TJ (1997) Spatial transformations in the parietal cortex using basis functions. *J Cognit Neurosci* 9:222–237.
- Sainburg RL, Ghez C, Kalakanis D (1999) Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. *J Neurophysiol* 81:1045–1056.
- Saltiel P, Wyler-Duda K, d'Avella A, Tresch MC, Bizzi E (2001) Muscle synergies encoded within the spinal cord: evidence from focal intraspinal NMDA iontophoresis in the frog. *J Neurophysiol* 85:605–619.
- Sanger TD (2000) Human arm movements described by a low-dimensional superposition of principal components. *J Neurosci* 20:1066–1072.
- Sanner RM, Kosha M (1999) A mathematical model of the adaptive control of human arm motions. *Biol Cybern* 80:369–382.
- Santello M, Flanders M, Soechting JF (1998) Postural hand synergies for tool use. *J Neurosci* 18:10105–10115.
- Scheidt RA, Dingwell JB, Mussa-Ivaldi FA (2001) Learning to move amid uncertainty. *J Neurophysiol* 86:971–985.
- Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. *J Neurosci* 17:409–419.
- Shadmehr R, Holcomb HH (1997) Neural correlates of motor memory consolidation. *Science* 277:821–825.
- Shadmehr R, Moussavi ZM (2000) Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20:7807–7815.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224.
- Smeets JB, Brenner E (1995) Perception and action are based on the same visual information: distinction between position and velocity. *J Exp Psychol Hum Percept Perform* 21:19–31.
- Smith MA (2001) *Error correction, the basal ganglia, and the cerebellum*. PhD thesis, Johns Hopkins University.
- Soechting JF, Flanders M (1997) Flexibility and repeatability of finger movements during typing: analysis of multiple degrees of freedom. *J Comput Neurosci* 4:29–46.
- Thoroughman KA (2002) Kinematic, not dynamic, cost minimization by humans reaching in dynamically perturbing environments. *Soc Neurosci Abstr* 24:855.4.
- Thoroughman K, Shadmehr R (1999) Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19:8573–8588.
- Thoroughman K, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. *Nature* 407:742–747.
- Tillery SI, Soechting JF, Ebner TJ (1996) Somatosensory cortical activity in relation to arm posture: nonuniform spatial tuning. *J Neurophysiol* 76:2423–2438.
- Witney AG, Vetter P, Wolpert DM (2001) The influence of previous experience on predictive motor control. *NeuroReport* 12:649–653.