Published in final edited form as:

Hum Mov Sci. 2004 November ; 23(5): 543–568. doi:10.1016/j.humov.2004.04.003.

Generalization as a Behavioral Window to the Neural Mechanisms of Learning Internal Models

Reza Shadmehr

Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, 419 Traylor Building, 720 Rutland Ave, Baltimore, MD 21205 USA

Abstract

In generating motor commands, the brain seems to rely on internal models that predict physical dynamics of the limb and the external world. How does the brain compute an internal model? Which neural structures are involved? We consider a task where a force field is applied to the hand, altering the physical dynamics of reaching. Behavioral measures suggest that as the brain adapts to the field, it maps desired sensory states of the arm into estimates of force. If this neural computation is performed via a population code, i.e., via a set of bases, then activity fields of the bases dictate a generalization function that uses errors experienced in a given state to influence performance in any other state. The patterns of generalization suggest that the bases have activity fields that are directionally tuned, but directional tuning may be bimodal. Limb positions as well as contextual cues multiplicatively modulate the gain of tuning. These properties are consistent with the activity fields of cells in the motor cortex and the cerebellum. We suggest that activity fields of cells in these motor regions dictate the way we learn to represent internal models of limb dynamics.

Keywords

reaching movements; motor control; adaptation; motor learning; motor cortex; cerebellum; system identification; trial-to-trial analysis

1. Introduction

Our arm has inertial dynamics that describe a complex relationship between motion and forces. In order to reliably produce even the simplest movements, for example, flexion of the elbow, the brain must activate not only elbow flexors, but also shoulder flexors that counter the shoulder extension torque that is produced by the acceleration of the elbow. The importance of these interaction forces was quite apparent when engineers were trying to control motion of robots (Hollerbach and Flash, 1982). Yet the principle is the same for control of biological limbs, as has been recently confirmed in EMG recordings from the human arm (Gribble and Ostry, 1999). This has led to the idea that contrary to earlier hypotheses (Flash, 1987); passive properties of muscles are not enough to compensate for the complex physics of our limbs. Rather, with every movement, the brain must *predict* the specific force requirements of the task.

To illustrate this idea, suppose we ask a volunteer to pick up an opaque carton of milk that appears full but has been drained empty. The visual appearance of the bottle retrieves an internal

Send correspondence to: Reza Shadmehr at the above address. Email: reza@bme.jhu.edu, Phone: 410-614-2458, Fax: 410-614-9890..

PsychINFO classification categories: Motor Processes 2330, Learning & Memory 2343.

model that predicts the forces that are necessary to lift the bottle. The brain overestimates the mass of the carton by only about a pound (the weight of the missing milk), yet the motor commands that are constructed based on this prediction make the arm dramatically veer off-course. The on-line error correcting mechanisms inherent in the length-tension properties of the muscles and the spinal and supra-spinal reflexes reduce the effects of the prediction error. However, the fact that small, unexpected changes in limb dynamics produce large errors suggests that first, the gains associated with the on-line error correcting pathways are typically small, and second, motor commands are largely constructed based on predictions about the dynamics of the task.

The accuracy of internal models is particularly important for control of our arm because our hands evolved in large part to support manipulation --- holding an object changes dynamics of reaching. Furthermore, because our arm is connected to a potentially moving trunk, interaction torques due to motion of the trunk or rotation of the body will also change dynamics of reaching (Lackner and Dizio, 1994). The neural system that predicts force properties of reaching would have to be able to accommodate this variability and adapt.

But how does the brain compute internal models of dynamics? To approach this question, consider how the error in the lifting of the empty milk carton might influence a subsequent movement with the same carton. Suppose that on the subsequent trial, we ask the volunteer to slide the carton to the right. In effect, we are asking the brain to *generalize* its experience in the lifting movement to a movement 90° away. Intuitively, if we assume that the internal model is composed via a set of elements, then the pattern of generalization should have something to do with the sensitivity of the elements to movement direction. If the elements are about as sensitive to an upward movement as to a movement at 90°, then the error that was experienced for the upward lift should generalize. If on the other hand one group of elements are activated for lifting and another group are activated for the 90° movement, then there should be little generalization.

Here I review the theoretical framework that has been employed to express these ideas and consider the experimental results regarding patterns of generalization. I argue that the generalization patterns represent a fundamental signature of the neural system that computes internal models.

2. Internal models of limb dynamics associate sensory states of the arm to forces

To study the properties of the neural system with which the brain learns to predict forces, we use a paradigm (Fig. 1) where arm dynamics is systematically changed through imposition of forces by a robotic manipulandum (Shadmehr and Mussa-Ivaldi, 1994). The subject is provided with a target and asked to reach while holding the handle of the robot. When the robot's motors are disengaged (null field condition), movements are straight (Fig. 1A). The forces in the field typically depend on the velocity of the hand (Fig. 1B). When the field is applied, movements are perturbed (Fig. 1C). With practice, hand trajectories become nearly straight. The brain's ability to modify motor commands and predict the novel forces is revealed in *catch trials* where there is a sudden removal of the forces. Very early in training, the hand's trajectory in the catch trial is a straight path to the target. With further training in the field, as trajectories in *field trials* become straight, trajectories in catch trials (Fig. 1E) become approximately a mirror image of the early, field trials (Fig. 1C). The trajectories in these catch trials are after effects of adaptation of the internal model.

Improvement in performance occurs because training results in a change in the motor commands. One possibility is that movements improve because subjects co-contract antagonist

muscle groups. This motor strategy effectively increases the gain of the error correcting feedback system. However, in catch trials, this kind of adaptation would not produce any after effects.

An alternate hypothesis is that the composition of motor commands by the brain relies on a neural system that for any given movement, predicts the motor commands that are necessary to counter the forces that will be imposed by the robot. One way to do this is to imagine a tape that is played out as a function of time as the limb moves toward a target. This tape may be an average record of forces that were sensed in the previous movements in that direction. Mathematically, the inputs to this system are direction and time and the output is force. To test this idea, Conditt et al. (1997) trained subjects to reach to a small number of targets in a force field and then asked them to draw a circle in the same field. They reasoned that if what was learned was like a tape recording of the forces encountered in reaching to each target, then the neural system that had been trained to predict forces in short, straight reaching movements should contribute little to longer, circular movements. However, they found that performance was quite good in circular movements when the field was on and, importantly, the subjects showed after effects when the field was off.

This suggested that the neural system did not predict forces explicitly as a function of time. Rather, during the reaching trials the neural system had learned to associate the sensory states of the hand—its position and velocity—to forces. The temporal order in which those states were visited (i.e., the temporal component of the trajectory) was apparently immaterial. What was important was the region of the state space—the hand's velocity at a given position—that the reaching movements had visited and the forces that were encountered at that state space. If the temporal order of the states were changed from the “training set” in which the system had experienced the forces, the neural system could still predict forces because the states themselves were part of the initial training set.

However, one could argue that the reason why the subjects learned to associate states to forces, rather than some other input that explicitly included time, was because the force field that was imposed on the hand was itself not explicitly time dependent. Rather, it was dependent on hand velocity. Conditt and Mussa-Ivaldi (1999) tested this by asking whether subjects could adapt to force fields that explicitly depended on time. Remarkably, the results indicated that they could not. When a predictable, time-dependent pattern of force was imposed during reaching, behavior in circular movements suggested that subjects still learned to associate states of the arm to forces. Therefore, the brain's ability to predict force did not explicitly depend on movement time. Rather, the experiments of Conditt and Mussa-Ivaldi suggested that with practice, participants learned a sensory to motor transformation where a position and velocity-like input signal indicating state of the arm was transformed into a force-like output signal. The term that is used to label this computation is an *internal model* of the force field.

3. Interaction between internal models and on-line error correcting mechanisms

The internal model, i.e., the neural system that predicts force as a function of a given desired state of the limb, is our tentative answer to the question: “What is being learned with training?” We have not specified how information is represented in this internal model, or how this information is acquired through experience. All we can say at this point is that at the start of training the internal model is “empty” (i.e., it predicts zero force for all input states) and after a period of training, it adapts in the sense that it correctly predicts forces for typical states. However, there is sufficient information in this statement to allow us to test whether our formulation thus far is consistent with measurements.

If a simulation of an adaptive controller attached to the dynamics of the arm acquires an internal model of a force field, what will its trajectories of motion look like? The dynamics of the arm (in this case, a two-joint planar system) are derived from Newton's laws and are written as equations that describe how the limb's acceleration depends on forces. They describe how the mass of the limb responds to force input from the muscles. To represent the error feedback system of the muscles and the reflexes, we add to the equations a simple low-gain spring-damper element that stabilizes the limb about the desired trajectory. To produce a movement, we assume that the joint torques are commanded based on knowledge of the inverse dynamics of the limb, i.e., a map that transforms the desired sensory state of the limb into torques so that it compensates for the arm's inertial dynamics. This is an internal model of the arm's physical dynamics. The equations have been detailed in Shadmehr and Mussa-Ivaldi (1994).

Initially in training, the simulated internal model has no knowledge of the robot-imposed forces. Because of this, the simulated arm does not move straight to the target (Fig. 1D). Rather, it moves along a hooking trajectory that is similar to what we have recorded in our participants, (Fig. 1C) (Shadmehr and Mussa-Ivaldi, 1994; Shadmehr and Brashers-Krug, 1997d). Note how the trajectory errors (deviations from a straight line) are larger for reaches toward 90° than for 0°. The simulated internal model underestimates the forces by the same magnitude for all directions, yet the trajectory errors are direction dependent. The simulation demonstrates that this difference is due to direction dependent anisotropies in limb inertia and stiffness. Therefore, errors in reaching are not directly related to errors in the internal model. Rather, limb biomechanics transforms the errors in force prediction to produce errors in limb trajectory. This is an important idea that we will come back to when we consider the question of generalization.

Now we change the internal model so that it completely takes into account the added dynamics of the force field. If we introduce a catch trial, the resulting movement (Fig. 1F) is approximately mirror image of the field trials early in training. Therefore, the trajectories that we had recorded in the reaching movements of our subjects suggest a motor control system that utilizes both a feed-forward internal model and a feedback error correcting mechanism.

4. Neural correlates of learning an internal model of a force field

The equations of motion that produced trajectories in Fig. 1 included torque generators but not specific muscles. Thoroughman and Shadmehr (1999) added to these equations constraints regarding function of arm muscles and translated expected forces on the hand to changes in muscle activation patterns. To visualize the changes, for each muscle they plotted the magnitude of activation, averaged for a given reach, as a function of movement direction (Fig. 2B). The result was a tuning function that had a preferred direction (PD) for each muscle (Flanders and Soechting, 1990; Sergio and Kalaska, 1998). The PD for a muscle was the direction of movement in Cartesian coordinates centered on the hand for which the modeled muscle was most active. They observed that when the model adapted to certain force fields, the PD of certain muscles rotated.

For example, they considered adapting to a curl force field (Fig. 1B). In a curl field, forces push the hand perpendicular to its current velocity. In the clockwise version of this field, force vectors point in a clockwise direction perpendicular to the direction of hand motion. The simulations predicted and experiments confirmed that in adapting to a clockwise curl field, the PD of elbow muscles rotated by ~27° and PD of shoulder muscles rotated by ~18° (Thoroughman and Shadmehr, 1999). Therefore, one way to represent the change in motor commands due to adaptation of an internal model is via a rotation in the PD of the directional tuning of arm muscles.

In a sense, this is a trivial result because of course the commands to the muscles must change if forces are to counter the robot imposed field. However, the results are useful for two reasons. First, because the simulation predicts that for any given field, formation of an internal model should accompany a specific rotation in the PD of certain simulated muscles, it provides a compact way to quantitatively predict the experience-dependent change that should occur in the motor commands. Second, because in the monkey motor cortex, in certain conditions where PD of muscle activation functions had changed, so had the PD of some cells in the primary motor cortex M1 (Sergio and Kalaska, 1997), one can suggest that the rotation in EMG is echoing a similar change in PD of some motor cortical cells.

In an experiment where monkeys learned reaching movements in a clockwise curl field, Bizzi and colleagues (Li et al., 2001) observed that about 20% of the task related cells in M1 underwent a median clockwise shift in PD of 16° . This was close to the average change in PDs in the arm muscles of the monkeys. These “dynamic” or “muscle-like” cells had PDs that rotated back to the null condition when the field was washed out. The results suggested that perhaps the motor cortex was involved in computing the internal model. However, what might be a testable behavioral consequence of this hypothesis?

One of the consistent properties of reach related cells in M1 is that if a cell is directionally tuned at one arm configuration, it is also likely to be directionally tuned when the configuration of the arm is changed (Caminiti et al., 1990). However, the change in the workspace often results in a change in the PD of cells. When reaching is examined in the horizontal plane, PD of a typical M1 cell will rotate approximately with the shoulder angle (Caminiti et al., 1990; Caminiti et al., 1991). The reason for this rotation is possibly the observation that many M1 cells are sensitive to force requirements of the task. Sergio and Kalaska (1997) trained monkeys to generate isometric force ramps in 8 spatially constant directions in a horizontal plane while holding the arm in 9 different locations in a 16 cm diameter workspace. Typically, M1 activity was directionally tuned for the direction of isometric force in any given arm location in the workspace. However, many cells showed small but systematic shifts of directional tuning at different workspace locations even though the output force was in a constant spatial direction. On average, there was a significant clockwise rotation of cell PDs from the central hand location to locations to the right, and a significant counter clockwise rotation of cell PDs for hand position to the left. These rotations were consistent with the rotation of PDs in the shoulder and elbow muscles of the arm in the same task.

Because the “dynamic” cells that Bizzi and colleagues found (Li et al., 2001) were sensitive to changes in force properties of the task, i.e. their PDs rotated as the task was changed from a null field to a curl field, we can hypothesize that these cells might be “muscle-like”. By this, we mean that their PDs might change as the configuration of the arm changes. We would expect that their PDs will rotate with the shoulder in a way similar to rotation of PDs in arm muscles.

Imagine that the change in PD due to learning of the field and changes in PD due to shoulder rotation are cumulative. Then training in one workspace should result in the rotation of PDs by a certain amount, and translation of the arm to a new workspace should result in an additional rotation by an amount approximately equal to the rotation in the shoulder joint. At the new workspace, despite the fact that no prior training had taken place there, an effect of the prior training should be observed, i.e., we should observe generalization. This is a behavioral prediction of the model.

5. Generalization as a function of arm's position

Consider adapting to a force field shown in Fig. 2A. If the right arm is near the horizontal plane and the shoulder is flexed so that the hand is at a “left” workspace, preferred direction of biceps (an elbow flexor) is about -90° (Fig. 2B). When a subject trains in the field, one observes a

~35° clockwise rotation in the PD of biceps. Now imagine that there are cells in the motor cortex that rotate their PD by an amount similar to this. If we now take the subject's arm and extend the shoulder so that the hand is at the "right" workspace, we would expect the muscle-like M1 cells that were directionally tuned with the arm in the left workspace to also be directionally tuned when the hand is at the right workspace. Furthermore, we would expect that on average, the 90° clockwise rotation in the shoulder joint should cause the PD of these cells to rotate by an average of 90°. So for a motor cortical cell that was "muscle-like" and had a PD of say 180° at the left workspace, adaptation to the field at that workspace should cause the PD to change to 150° (i.e., 180°-30°), and movement of the hand to the right workspace should bring the PD to 60° (i.e., 180°-30°-90°). If the subject had not practiced movements in the field, at right this cell would have a PD of 90°. Therefore, the effect of training at the left workspace should be observable in terms of the behavior of the hand at the right workspace if the muscle-like cells that rotated their PD at the left workspace maintained their relative rotation at the right workspace. In terms of forces, this corresponds to a field where the relative rotation of the muscle PDs is maintained as a function of arm configuration.

One can approximate such a force field by transforming forces on the hand at the left workspace to joint torques, and then transforming the torques back to hand forces at the right workspace (Shadmehr and Moussavi, 2000). For the field shown in Fig. 2A, this procedure produces 180° rotation of the field (Fig. 2A). This theoretical result means that if one trains at left in the field shown, at the right workspace one should generalize to negative of that field. We were intrigued by this prediction because we had earlier observed that if one adapts to field *B* and then is given field *-B* in the same workspace, performance in *-B* is absolutely terrible. In fact, performance in *-B* for these subjects is far worse than performance of naïve subjects in the same field (Brashers-Krug et al., 1996). The model now predicted that if after training in *B* we simply moved the subject's arm to a new location, we would see that performance in *-B* is quite good. Experiments confirmed this prediction (Fig. 2C) (Shadmehr and Moussavi, 2000; Malfait et al., 2002). The results are consistent with the idea that the tuning properties of the muscle-like cells in M1 dictate some of the patterns of behavioral generalization.

However, it is certainly not the case that all M1 cells are muscle-like in their tuning properties. In many instances, experiments have demonstrated that a significant portion of cells in M1 code for parameters of reaching movements in extrinsic coordinates (Moran and Schwartz, 1999; Kakei et al., 1999). Indeed, in their force field learning experiment, Bizzi and colleagues (Li et al. 2001) found that 34% of M1 cells that they recorded had tuning properties that remained invariant despite the changes in force properties of the task (labeled as "kinematic" cells). One would predict that these kinematic cells would not change their PD with the configuration of the arm. Therefore, our hypothesis predicts that M1 cells that have more muscle-like properties, i.e., change their discharge patterns in a way that correlates with changes in muscle activations, are the ones that contribute most to the representation of internal models for dynamics of reaching movements.

6. Computing an internal model with a population code

How does one quantitatively test the idea that tuning properties of certain cells influence patterns of generalization during learning? To answer this question, we need to advance beyond a description of the input-output variables that are encoded by the internal model (sensory state of the arm and force, respectively) and consider how the transformation from input to output might take place. That is, we must first consider how the central nervous system might compute an internal model.

One of the most widely used models of neural computation is the population code. While the idea of using populations of neurons to code variables of interest is old (Humphrey et al.,

1970), it has become a compelling tool since it was combined with a simple decoding strategy called a population vector to reconstruct the direction of reaching movements from cells in M1 (Georgopoulos et al., 1982). To motivate our approach, let us put aside for now the problem of estimating force fields and consider the simpler problem of representing direction of movement of the hand.

Georgopoulos et al. (1982) recorded from a collection of cells in M1 and asked whether one could estimate direction of a reaching movement from the discharge of cells. Each cell had a preferred direction of movement that we represent as a vector of unit length w_i . The movements were in a plane. Therefore, w_i is a 2D vector that might point along any direction about a unit circle. In a given trial, imagine that the movement direction is α , and each cell i discharges by amount r_i . This discharge can be decomposed into two terms. The first term is an average response $g_i(\alpha)$ which represents the cell's tuning curve as computed over many movements to various directions. The second term is noise n_i that we might encounter at any given trial i :

$$r_i = g_i(\alpha) + n_i$$

Experiments show that the tuning curve is typically a cosine-like function of movement direction and has a half-width at half-height value of approximately 56° (Amirikian and Georgopoulos, 2000). The second term is noise that cannot be accounted for by the "input" (target direction). Experiments suggest that this noise term (for neurons in the visual cortex) is often normally distributed with a variance that is proportional to the mean value of the tuning function (Tolhurst and Thompson, 1982). If cells did not have this noise and we could record from a large number of cells at the same time, we could simply note that cell j happened to fire most during some movement and estimate the movement direction $\hat{\alpha}$ to be the angle of the preferred direction vector of that cell, $\hat{\alpha} = \angle w_j$. This is a winner-take-all coding. However, because cells are noisy, our estimate would have a large variance from trial to trial, even though the actual direction of movement did not change. A better approach is a population code (Georgopoulos et al., 1982). In this approach, each cell's discharge is weighted by its preferred direction vector. The sum of these vectors produces the estimate of movement direction:

$$\hat{\alpha} = \angle \sum_i w_i r_i = \angle \left(\sum_i w_i g_i(\alpha) + \sum_i w_i n_i \right)$$

This approach is better in the sense that when the movement direction is fixed, it produces a smaller variance in its estimate from trial to trial than the winner-take-all approach. In fact, if the tuning curves were exactly cosine functions, the estimate would be optimal in the sense that its variance would be as small as possible (Seung and Sompolinsky, 1993). Therefore, the success of population coding depends on computing with neurons that broadly encode the input variable. Where this condition has been approximately met, experiments have generally demonstrated that a population code could successfully be used to estimate the input variable from noisy neuronal discharge (Georgopoulos et al., 1986; Schwartz, 1994).

The example of population coding above is an instance of neural computation of an "identity mapping", i.e., a map where the output is an estimate of the input variable (movement direction). In general, a population code could also be used to map an input variable x into any other variable y (Poggio, 1990; Pouget et al., 2000). In this case, the tuning curves of the neurons that participate in this computation become the *basis functions* with which the output is approximated. When basis functions are linearly combined, they can approximate almost any linear or non-linear function. For example, Pouget and Sejnowski (1997) suggested that

neurons in the parietal cortex might serve as basis functions with which the brain computes position of a visual target with respect to the head. Cells in this region of the brain typically have a discharge r that is modulated by both position of the eye x_e in the orbit and position of the target on the retina x_r . These cells have a preferred position on the retina where discharge is maximum, and this discharge is modulated approximately linearly with the position of the eye (Andersen et al., 1985). The tuning function of a cell i can be labeled as $g_i(x_e, x_r)$. Using a weighted sum of these functions, one could estimate position of the target with respect to the head:

$$\hat{y} = \sum_i w_i g_i(x_e, x_r) + w_i n_i$$

The appropriate weighting w_i would have to be learned to form this map. However, Pouget and Sejnowski (1997) point out that because the tuning functions are the bases with which the map is constructed, the same bases can be used to form any other representation, for example, a shoulder centered representation of the target. This idea is important because it demonstrates that population coding, a method that can be used to form neural computation of identity maps, is equally suited for the more general problem of computing nonlinear maps. However, whereas in the population code described for decoding of movement direction the weights were vectors that were static and pointed in the preferred direction of a cell, here if the bases are to be used for learning arbitrary maps, then the weight vectors will change and will have no specific relationship with the tuning function.

Let us now return to our problem of how the brain might compute an internal model. One can think of an internal model as a map that transforms sensory input regarding the desired state of the arm θ (i.e., an ordered set of vectors representing the joint position and velocity of the arm) into joint torques τ . We choose a joint centered representation of limb states and forces because of the pattern of generalization in Fig. 2. Assume this computational is performed via a population code. Each basis that participates in this computation has a tuning curve g_i that describes activation as a function of limb position and velocity. Each basis has a preferred torque vector w_i . The population vector response of the network is:

$$\hat{\tau} = \sum_i w_i g_i(\theta) + w_i n_i \quad (1)$$

We now have a framework to relate tuning g_i with behavioral generalization. Consider the following experiment: force trajectory $\mathbf{f}$ acts on the hand as the hand moves along trajectory $\mathbf{x}_1$. The hand forces results in joint torques τ at arm state θ_1 . The difference between τ and $\hat{\tau}$ is an error signal that can change weights w . Assuming Hebbian learning rules, the weight change will be maximum for those bases that happened to be most active about state θ_1 . The subject is then asked to make a movement with the arm along state θ_2 , an arm position (or velocity) where the subject has not been trained. If performance is different from naïve, then the function g_i for which the weights adapted for movements along θ_1 must have been broad enough to not only be active for θ_1 , but also θ_2 . Therefore, if the internal model is represented via a population code, then generalization is affected by the shape of the tuning curves of the bases.

From the experiment in which we observed generalization from one arm configuration to another (Fig. 2) we can now conclude that the basis functions could not have had sharply tuned activity fields around a preferred hand position. If they did, then experiencing force at one hand position could not have generalized to another hand position 80 cm away. Our idea is that the tuning properties of muscle-like cells in M1 may be the function g in this population coding

(Eq. 1). To mathematically describe how activity of the bases might vary with arm position and velocity, we note that cell activity in M1 is modulated globally and often linearly as a function of static limb position (Georgopoulos et al., 1984; Sergio and Kalaska, 1997), and cells have preferred directions of movement that often changes as a function of the shoulder angle (Caminiti et al., 1990). To capture these observations, let us hypothesize that cells that are involved in representing the internal model have tuning functions that are described as follows.

$$g_i(\theta, \dot{\theta}) = (k_i^T \theta + c) \exp\left(-\frac{\|\dot{\theta} - \dot{\theta}_i\|^2}{2\sigma^2}\right) \quad (2)$$

The above function represents the output of a basis in the internal model. The position dependent term is a linear function that encodes joint angles. The gradient vector k reflects sensitivity of the basis to changes in static position of the shoulder and elbow. Preferred limb velocity is specified by $\dot{\theta}_i$ and c is a constant. The output for various limb positions and movement directions is plotted in Fig. 3. The basis is directionally tuned, but the preferred direction changes as a function of hand location. Furthermore, change in hand position alters both the activation at static hold time and the depth of tuning of the basis.

Consider a situation in which the internal model is constructed as a linear combination of these bases. With training at a given arm configuration, the preferred torque vector for some of these elements will change. The change will occur in those elements that happen to be most active at that arm configuration. The way that these active elements change their static discharge with arm position dictates how far in position space the learning will generalize. The gain k dictates how close two workspaces have to be (in position space) before learning of conflicting fields becomes impossible. When the gain is high, output of the bases changes greatly as a function of hand position. This results in poor generalization between neighboring positions of the hand, making it possible to learn two different patterns of force at two different hand positions. When the gain is low, output of the bases changes slowly as a function of hand position. At its limit, output changes not at all as a function of hand position and effectively there is no coding of hand position. In this situation, forces generalize globally in hand position space and it is not possible to learn two different forces for the same direction of movement in two different spatial locations.

Hwang et al. (2003) estimated the gain k by measuring generalization patterns. The results of their estimate produced the basis function that was plotted in Fig. 3. Using these bases, they predicted learning patterns in a number of tasks (Fig. 4). For example, the model predicted poor learning when the force field was translation invariant in Cartesian coordinates (Fig. 4A). The model predicted poor learning when three movements that were in the same direction had a null field on the outside as opposed to the middle (Fig. 4B). The model predicted poor learning when a position dependent field presented opposite forces for movements that were in the same direction but started from slightly different positions (Fig. 4C). Behavioral measures have confirmed these predictions. Therefore, behavioral measures suggest that the internal model of limb dynamics is constructed with bases that broadly and perhaps linearly encode static limb position and this encoding act as a gain that multiplicatively modulates tuning with respect to direction.

The reason why we assumed that the bases linearly coded hand position space was because discharge of cells in the spinocerebellar tract (Bosco et al., 1996), somatosensory cortex (Helms Tillery et al., 1996), and M1 (Georgopoulos et al., 1984) is modulated linearly with hand position. The reason for assuming that the bases encoded joint velocity (rather than hand velocity) was because such encoding can account for the observation that preferred direction

of many task related cells rotates with a rotation in the shoulder angle (Scott and Kalaska, 1997; Ajemian et al., 2001). Indeed, it appears that a linear-like modulation of discharge with respect to limb position and a Gaussian-like tuning with respect to arm velocity may be a fundamental property of many cells in the motor system. Simulations suggest that if the brain computes the internal model with elements that have these tuning properties then one can account for spatial generalization patterns that we and others (Ghez et al., 2000) have recorded in human subjects.

7. Generalization from one arm to the other

We can take this argument a step further and predict generalization patterns from one arm to another based on tuning properties of cells in M1. One of the remarkable properties of many of these cells is that if their discharge is modulated as a function of movement direction for the contralateral arm, that discharge is also modulated when the reaches are performed with the ipsilateral arm. In fact, it was recently reported that many cells in M1 maintained their preferred direction of motion (calculated with the contralateral arm) even when the ipsilateral arm was performing the reaching movements and the contralateral arm was at rest (Steinberg et al., 2002). Kalaska and colleagues also observed similar properties of tuning functions for cells in the premotor cortex (Kalaska et al., 2000). An important technical point in both of these experiments is that the left and the right arms performed reaching movements directly in front of the animal in the same workspace.

Imagine that this invariance of PD with respect to the right and left arms is also a property of the muscle-like cells that we hypothesized were involved in representing this task. Consider a subject who trained with her right hand in a curl force field, resulting in a rotation of PDs in the dynamic cells by approximately 30° . We now ask the subject to use her left hand and make reaching movements. Because the neurons in the left hemisphere changed their PD due to training with the right hand, and because these same cells are also tuned for movements with the left hand, they could potentially influence movements with the left hand. In fact, the model predicts that there should be generalization from right to left. Furthermore, it predicts that the coordinate system of generalization should be in an extrinsic, Cartesian-like coordinates.

This is a very surprising prediction because we noted before that PDs rotate with the shoulder of the trained arm, causing the training to generalize in an intrinsic coordinates within the same arm. The theory now predicts that because the PDs are invariant to the arms, if we looked for generalization between arms, we would see transfer in extrinsic coordinates.

We tested this by asking whether acquisition of an internal model of arm dynamics allowed the brain to generalize the acquired information to the other arm, and if so, what was the coordinate system of representation (Criscimagna-Hemminger et al., 2003). We used the standard reaching movement paradigm with curl fields. Hand position was directly in front of the subject centered on the midline. We considered two coordinate systems for generalization: intrinsic (joint) coordinates and extrinsic (Cartesian) coordinates. In the intrinsic coordinate system, if a movement to a given direction required increased activity in the biceps (for example), then the same movement direction with the other arm should also require increased activity in biceps. Mathematically, this results in a mirror transformation of the force field to the other hand. In the extrinsic coordinate representation, the system would expect the same forces to act on the other hand in terms of direction of movement in Cartesian space. We explored generalization from right hand to left hand in right-handed subjects. We found that in right-handed individuals, generalization occurred only from the dominant right to the left arm, and its coordinate system was extrinsic in the workspace that we tested.

8. Tuning properties with respect to color of the target

The main claim of the hypothesis is that tuning properties of cells in the “muscle-like” cells in the motor cortex can strongly influence behavior. In particular, the tuning properties predict how we learn dynamics of reaching movements. In most of the examples that we have considered thus far, forces that were imposed on reaching movements explicitly depended on the state of the arm. Let us now consider a task where the forces do not depend on state of the arm. For example, imagine a reaching movement where limb position or velocity does not uniquely describe the forces in the task. A very simple case is one where a target is presented at a given direction, but the forces that will be presented during that movement depend both on hand velocity and the color of the target.

If the bases that take part in learning this task are strongly tuned with respect to position or velocity of the arm and not to the color of the target, then this apparently simple task should be in fact extremely difficult to learn. Gandolfo et al. (1996) asked subjects to make movements to various directions. For the initial 48 movements, a velocity dependent field, labeled B_1 , was present. For the next block of 48 movements, field $-B_1$ was present, and so on. During presentation of each field, the room was flooded with a specific color of light. Despite hundreds of movements, subjects never learned to use the color as a cue to predict the pattern of forces.

We recently simplified this experiment by limiting movements to only one direction (Rao and Shadmehr, 2001). On any given trial, the color of the target was randomly chosen as either red or green. Red implied a clockwise curl field and green implied a counter clockwise curl field. Because the movement was always in the same direction, the pattern of forces on that movement depended exclusively on this cue. We trained subjects extensively on this task, providing them with over 3000 trials, spread over three days. Remarkably, in catch trials we found no evidence of color-specific after effects. Therefore, over the course of this training our volunteers were not able to use color as a cue to predict the pattern of forces. However, under some conditions, color can be a salient cue for learning of fields. For example, Krouchev and Kalaska (2003) demonstrated that a highly trained monkey could differentiate between certain force fields using a color cue. Osu et al. (2004) found that with cognitive instructions and certain training schedules color cues could be used for learning of different fields. In all these reports, however, it is apparent that color is simply not as salient a cue as movement direction. This predicts that while the bases that take part in representing dynamics of reaching movements are significantly affected by movement direction, their activation is not as fundamentally altered by the color of the reaching target.

9. Problems in measuring a generalization function

The basic idea then is that by measuring generalization, one can infer some of the properties of the bases with which the internal model is computed. However, it is useful to outline the problems that are inherent in this approach.

To measure generalization, subjects are trained with an input $\mathbf{x}_1$ and are then tested with a new input $\mathbf{x}_2$. The first problem with this approach is that it requires an experimentally naïve set of participants to be trained in each pairing of $\mathbf{x}_1$ and $\mathbf{x}_2$. As a result, behavioral experiments are often limited to training and testing with one or two pairs of inputs, and conclusions are in terms of qualitative statements regarding the shape of the bases, i.e., wide or narrow. The approach is simply not feasible for non-human primate studies.

The second problem is that in motor control, we have to consider coordinate systems. Generalization not only depends on the distance between training and test locations, but also on the coordinate system in which that space is encoded. For example, a force that is experienced at a given location may be generalized in terms of torques on the joints or forces

on the hand. These two coordinates predict different patterns of generalization in terms of position of the hand.

The third problem is that the bases that are inferred from one generalization experiment might not be consistent with those that are inferred in another. In other words, adaptation to one force field might result in a pattern of generalization that is inconsistent with the pattern observed in adaptation to another field. It would indeed be remarkable if behavioral data from a wide variety of force adaptation experiments suggested a consistent shape to the bases. If this were the case, then one could argue that one has estimated the basic motor primitives with which internal models are computed.

Finally, even if we are lucky enough to solve all of these problems, we would still have the problem of interpretation: we would hope that the bases that are inferred by this abstract model not only explain behavior, but also are interpretable in terms of the neurophysiology of the motor system. In the next section, I suggest that all four questions can be approached.

10. Measuring generalization from trial-to-trial changes in behavior

To approach the first two questions - being limited to naïve subjects and needing to consider coordinate systems - we have been developing a new mathematical method to estimate the shape of the bases from the trial to trial variations in performance (Thoroughman and Shadmehr, 2000; Donchin and Shadmehr, 2002; Donchin et al., 2003; Donchin and Shadmehr, 2002). We leverage the fact that the shape of the bases determines how error experienced during a movement will generalize to the subsequent movement. That is, the weight associated with the basis elements is likely to change most for the elements that are most active where the error was experienced. The weight change is smallest for the elements that were not active where the error was experienced. This means that if error on one movement affects behavior on the next movement in a different direction, then some of the bases must be reasonably active during both of these movements. That is because they must be active in the first movement to be influenced by the error and active in the second movement in order to have an influence on behavior. Thus, generalization of error from one movement to the next can tell us whether the basis elements are wide enough to encompass both movements.

The idea is useful because it suggests that it is not necessary to train in one set and then test in another in order to estimate generalization. Rather, all possible inputs should be presented in a random sequence. From the trial-to-trial changes in performance, one can estimate how error in one movement affected the subsequent movement as a function of the distance of the two movements in the state space (for example, angular distance in directions of the two movements). The result is a *generalization function*. It is a measure of similarity between two states. If the bases are activated by similar amounts in two different states, then the similarity is high and the generalization function will be large. If the bases are uncorrelated in these two states, then similarity is low and there is little generalization. Therefore, the generalization function estimates the similarity between states as seen through the window of the bases.

As an example, consider a task where subjects make movements to 8 directions in a random order. We are interested in estimating how force experienced in a given direction is generalized to all other directions. However, we do not know the bases g . To simplify matters, let us ignore the noise in Eq. 1, assume that the bases are only a function of movement direction α , and rewrite that equation in terms of vector quantities:

$$\hat{\tau} = Wg(\alpha)$$

$$W = \begin{bmatrix} w_{11} & \cdots & w_{1m} \\ w_{21} & \cdots & w_{2m} \end{bmatrix}$$

$$\mathbf{g}(\alpha) = \begin{bmatrix} g_1(\alpha) & \cdots & g_m(\alpha) \end{bmatrix}^T$$

Here $\hat{\tau}$ is a 2×1 vector. It is an estimate of actual robot imposed torques τ . The error in our estimate is:

$$\tilde{\tau} = \tau - \hat{\tau} = \begin{bmatrix} \tilde{\tau}_1 \\ \tilde{\tau}_2 \end{bmatrix}$$

Our objective is to change W so that we minimize the “squared” error e :

$$e = \frac{1}{2} \tilde{\tau}^T \tilde{\tau}$$

To do so, we need the gradient of e with respect to W . After some algebra, we find:

$$\frac{de}{dw_{ij}} = -\tilde{\tau}_i g_j$$

After performing trial n , the error in that trial $\tilde{\tau}^{(n)}$ will be used to change the weights $W^{(n)}$ of the internal model. That change will be in the opposite direction as the gradient, and will be weighted by a small constant η :

$$w_{ij}^{(n+1)} = w_{ij}^{(n)} + \eta \tilde{\tau}_i^{(n)} g_j(\alpha^{(n)})$$

Writing this in vector form we have:

$$W^{(n+1)} = W^{(n)} + \eta \tilde{\tau}^{(n)} \mathbf{g}(\alpha^{(n)})^T$$

If we multiply both sides of this equation by $\mathbf{g}(\alpha^{(n+1)})$, we have:

$$W^{(n+1)} \mathbf{g}(\alpha^{(n+1)}) = W^{(n)} \mathbf{g}(\alpha^{(n+1)}) + \eta \tilde{\tau}^{(n)} \mathbf{g}(\alpha^{(n)})^T \mathbf{g}(\alpha^{(n+1)})$$

This is equivalent to:

$$\tilde{\tau}^{(n+1)}(\alpha^{(n+1)}) - \tilde{\tau}^{(n)}(\alpha^{(n+1)}) = \underbrace{\eta \mathbf{g}(\alpha^{(n)})^T \mathbf{g}(\alpha^{(n+1)})}_{\text{generalization function}} \tilde{\tau}^{(n)} \quad (3)$$

Eq. 3 says that the change in the internal model from trial n to $n+1$ is completely described by the error in trial n times a generalization function. That generalization function is the correlation between the activation of the bases in trial n and trial $n+1$. Intuitively, we see that the correlation will be largest when the two consecutive trials are along the same direction. The shape of the tuning functions will determine the generalization as the distance between the two trials changes in the direction space.

However, Eq. (3) only describes the effect of the generalization function from movement n on the direction of the next movement $n+1$. In principle, error could have been generalized to all directions, but its effect on only one direction is observable in trial $n+1$. Therefore, we need to account for potential generalizations on directions (i.e., states) that cannot be observed in trial $n+1$. We express this idea, we begin by assuming that generalization can be described as a function of the distance between the direction $\alpha^{(n)}$ in which the error was experienced in trial n and all other potential directions $\alpha^{(i)}$:

$$b(|\alpha^{(i)} - \alpha^{(n)}|) = \eta g^T(\alpha^{(i)}) g(\alpha^{(n)}) \quad \forall i$$

The symbol $\|$ is absolute value. Therefore, the internal model in trial $n+1$ is:

$$\tilde{\tau}^{(n+1)}(\alpha^{(i)}) - \tilde{\tau}^{(n)}(\alpha^{(i)}) = b(|\alpha^{(i)} - \alpha^{(n)}|) \tilde{\tau}^{(n)} \quad \forall i$$

We see that if we could estimate the generalization function b from trial-to-trial changes in performance, we could have a reasonable idea of the kind of bases that are being used for computation of the internal model. However, our problem is that we can only record people's reaching movements, not $\hat{\tau}$. Movements will be straight when the internal model has a correct estimate of force. There will be an error in the hand's trajectory when this estimate is incorrect. Let us assume that this error in the hand's trajectory is computed simply as a vector that describes where the hand is at peak velocity $y^{(n)}$ with respect to where it "should be" y^* (perhaps along a straight line to the target). Let us call that position error vector $\tilde{y}^{(n)}$. Let us further assume that it will be related to the force error τ via a compliance matrix D . This matrix relates how force error produces a displacement from the intended trajectory. We now have the following:

$$\begin{aligned} \tilde{y}^{(n)} &= D(\tau^{(n)} - \tilde{\tau}^{(n)}) = D\tilde{\tau}^{(n)} \\ \tilde{\tau}^{(n+1)}(\alpha^{(i)}) &= \tilde{\tau}^{(n)}(\alpha^{(i)}) + b(|\alpha^{(i)} - \alpha^{(n)}|) \tilde{\tau}^{(n)} \quad \forall i \end{aligned} \quad (4)$$

One begins by measuring a sequence of movement errors $\tilde{y}^{(n)}$ and fit them to the system in Eq. (4) in order to find the best fit for matrices D and vector b . If there are 8 directions of movement in the task, b has 8 unknown parameters. D is typically a 2×2 matrix. Therefore, there are 12 unknown parameters in this model. The procedure for fitting these equations to a sequence of movements is provided in Donchin and Shadmehr (2002). If the model is correct, it should describe all the trial-to-trial changes in performance that takes place during adaptation and provide us with an estimate of the generalization function.

Donchin et al. (2003) fit these equations to behavioral data (Fig. 5A). Data from a large group of subjects ($n=75$) had been collected as they learned to make movements in a curl force field. The target pattern was out-and-back in a half-pinwheel pattern. That is, movements began at center; a target was presented at 0° , 40° , 90° , or 135° . Upon completion of that movement the center target was lit, and the pattern was repeated. In this way, the movements were to 8

directions but all outward movements were followed by a movement back to the center. We found that: 1) the equations typically explained about 75% of the variance in the data (Fig. 5A); and 2) the generalization function b (Fig. 5B) and compliance matrix D (Fig. 5C) remained consistent across repeated measures. Interestingly, the generalization function was wide and bimodal. That is, generalization dropped off as angular distance of movements increased and reached a minimum at a distance of 90° , but then rose to approximately 50% of its peak value at 180° .

It was possible that this bimodality was an artifact of the out-and-back target sequence. Donchin et al. (2003) tested a new group of subjects in a random target sequence where the robot brought the hand to the start position of each movement (2nd row of Fig. 5). b and D maintained their shape. They also tested another group of subjects in a target sequence where not only the directions of movement were random, but the force field at each direction was also random (third row of Fig. 5). In this condition, at any given trial the field was either null, clockwise curl field, or counter clockwise curl field. As the field was random, adaptation in the traditional sense was not possible. Remarkably, analysis of the trial-to-trial changes in performance produced a generalization function similar to that which was estimated in a “constant” field. In all cases, the generalization function was bimodal, consistent with bases that encode direction of movement with a bimodal activation pattern. The shape of the basis function that is consistent with this generalization pattern is shown in Fig. 5D.

Our finding that the generalization function remains invariant even in a randomly changing force field suggests that the fundamental computational properties of the internal model are approximately the same across repeated measures, across subjects and across a small number of force learning tasks that thus far we have tested. Because the shape of g in Eq. (1) is responsible for generalization, this is our strongest evidence that there exists a single set of bases that encode internal models of limb dynamics.

11. Neurophysiological correlates of the bases

In the 4th question that we posed above, we asked whether the bases that are to be inferred with this abstract model are interpretable in terms of the neurophysiology of the motor system. From the patterns of generalization, we conclude that:

1. The bases encode position of the arm linearly in the horizontal plane. The linear coding of static limb position acts as a gain on directional tuning.
2. The bases encode hand velocity with a function that has a preferred direction and is modulated broadly but is bimodal.
3. The preferred directions of the bases rotate with the shoulder angle.
4. The bases are tuned to movements of the ipsilateral arm such that the preferred direction remains arm invariant if the workspace is near the midline.
5. The bases are either not modulated by color of the target or the modulation is weak.

All of these properties except the bimodality can be found among task related cells in the primary motor cortex, the basal ganglia, and the cerebellum (Georgopoulos et al., 1984; Caminiti et al., 1991; Turner and Anderson, 1997; Johnson and Ebner, 2000; Caminiti et al., 1991; Johnson and Ebner, 2000). The invariance of the preferred direction with respect to movements of the contralateral and ipsilateral arms was recently observed in the cells of the motor cortex (Steinberg et al., 2002), premotor cortex, and the cerebellum (Bradley Greger and Tom Thach, personal communication). However, to our knowledge bimodality has only been observed in the cerebellum during reaching movements: Purkinje cell discharge during reaching movements shows a weak but consistent bimodal activation pattern as a function of

hand velocity (Coltz et al., 1999), whereas no such bimodality is reported in the same task in the primary motor cortex (Johnson and Ebner, 2000).

In reaching movements, a muscle that provides the agonist burst to reach in a particular direction (say 0°) also provides the antagonist burst for a movement in the opposite direction (180°), but is generally not modulated very much when a movement is made to 90° . The antagonist burst is generally significantly smaller than the agonist burst. Therefore, bimodality is a fundamental characteristic of muscle activation functions during reaching and generalization patterns in terms of direction of movement suggest that the bases are likely to have muscle-like tuning functions. We saw earlier that generalization patterns in terms of spatial configuration of the arm also made this suggestion. Taken together, this suggests that the neural computation of the internal model is with neurons that have muscle-like tuning properties with respect to contralateral arm during reaching movements.

The one aspect of the model that is not muscle-like is the encoding of velocity. In Eq. (2), the bases encode velocity with Gaussians. This means that each basis has a preferred velocity of movement. Purkinje cells in the cerebellar cortex appear to encode movement velocity in this way, whereas cells in M1 generally increase their discharge with increased movement speed (Moran and Schwartz, 1999; Johnson and Ebner, 2000). Only one study has considered how the internal model generalizes in terms of speed of movement (Goodbody and Wolpert, 1998). In that study, force adaptation at a given average velocity generalized less than linearly to neighboring velocities. However, the precise shape of the generalization function is not known. If it generalizes globally, then that representation would be muscle-like and consistent with tuning of cells in M1. If it generalizes locally, then that representation implies a coding of velocity that peaks at a particular value and then declines, that is, a preferred velocity. This later generalization would be consistent with tuning of task related cells in the cerebellum.

12. Consolidation

I described learning of internal models using a mathematical framework where acquisition of information was one and the same as memory. In this framework, the preferred torque vector associated with each basis changed to minimize error in the task. Once the task was over, presumably these changes are maintained and that forms the basis of long-term memory.

The scope of my naiveté was plainly demonstrated when we found that acquisition of memory of an internal model was merely the first step in a sequence of events that eventually resulted in a long-term representation of motor memory. Our behavioral measurements suggested that the internal model changed not only during the training session, but also in the hours that followed completion of training (Shadmehr and Brashers-Krug, 1997d). The motor memory appeared to gradually change from an initially fragile state to a state more resistant to change during a period of ~ 5 hours (Brashers-Krug et al., 1996). Some of these results have recently been extended: Ghez and colleagues reported that in a task where subjects learned internal models of an inertial object, motor memory of inertial object 1 could be disrupted if practice was immediately followed by movements with inertial object 2 (Krakauer et al., 1999). Using transcranial magnetic stimulations (TMS), Hallett and colleagues reported that stimulation of M1 immediately after practicing a thumb flexion task resulted in marked retention deficits whereas stimulation of M1 at 5 hours post-practice did not affect retention (Muellbacher et al., 2002). Using functional imaging, we have observed that at comparable levels of motor performance, the map of activation patterns in the brain differed significantly near the end of training on day 1 vs. at 6 hours (Shadmehr and Holcomb, 1997; Shadmehr and Holcomb, 1999), vs. at 2 or 4 weeks after initial practice (Nezafat et al., 2001).

Therefore, passage of time changes the neural representation of the internal model. We currently have no theory to account for this. One hope is that we eventually might be able to track changes in neural representation by measuring their influence on patterns of generalization. For example, one possibility is that with consolidation, the adaptive components of the neural system that represents the internal model will become less plastic. This will be reflected behaviorally in the way that error in one movement affects the subsequent movement.

13. Summary

The specific coding of movement parameters in the neurons that compose the motor system may have a significant, measurable influence on behavior. That influence, I suggest, can be observed in how our brain learns to compensate for forces in control of reaching movements. Training to reach in a force field results in a specific, highly reproducible pattern of generalization to other movements. If we assume that the neural computation of an internal model is via a population code, then the tuning curves of the bases that participate in this computation dictate the patterns of generalization. From the patterns of generalization one can infer the shape of these bases, and therefore infer representation.

Acknowledgements

This work was supported by grants from the National Institutes of Health (NS NS37422, NS16375) and the Human Frontiers Science Program.

Reference List

- Ajemian R, Bullock D, Grossberg S. A model of movement coordinates in the motor cortex: posture-dependent changes in the gain and direction of single cell tuning curves. *Cerebral Cortex* 2001;11:1124–1135. [PubMed: 11709483]
- Amirikian B, Georgopoulos AP. Directional tuning profiles of motor cortical cells. *Neuroscience Research* 2000;36:73–79. [PubMed: 10678534]
- Andersen RA, Essick GK, Siegel RM. Encoding of spatial location by posterior parietal neurons. *Science* 1985;230:456–458. [PubMed: 4048942]
- Bosco G, Rankin A, Poppele RE. Representation of passive hindlimb postures in cat spinocerebellar activity. *Journal of Neurophysiology* 1996;76:715–726. [PubMed: 8871193]
- Brashers-Krug T, Shadmehr R, Bizzi E. Consolidation in human motor memory. *Nature* 1996;382:252–255. [PubMed: 8717039]
- Caminiti R, Johnson PB, Galli C, Ferraina S, Burnod Y. Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. *Journal of Neuroscience* 1991;11:1182–1197. [PubMed: 2027042]
- Caminiti R, Johnson PB, Urbano A. Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *Journal of Neuroscience* 1990;10:2039–2058. [PubMed: 2376768]
- Coltz JD, Johnson MTV, Ebner TJ. Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking. *Journal of Neuroscience* 1999;19:1782–1803. [PubMed: 10024363]
- Conditt MA, Gandolfo F, Mussa-Ivaldi FA. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *Journal of Neurophysiology* 1997;78:554–560. [PubMed: 9242306]
- Conditt MA, Mussa-Ivaldi FA. Central representation of time during motor learning. *Proceedings of the National Academy of Science USA* 1999;96:11625–11630.
- Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R. Learned dynamics of reaching movements generalize from dominant to nondominant arm. *Journal of Neurophysiology* 2003;89:168–176. [PubMed: 12522169]

- Donchin O, Francis JT, Shadmehr R. Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *Journal of Neuroscience* 2003;23:9032–9045. [PubMed: 14534237]
- Donchin O, Shadmehr R. Linking motor learning to function approximation: Learning in an unlearnable force field. *Advances in Neural Information Processing Systems* 2002;14:197–203.
- Flanders M, Soechting JF. Arm muscle activation for static forces in three-dimensional space. *Journal of Neurophysiology* 1990;64:1818–1837. [PubMed: 2074466]
- Flash T. The control of hand equilibrium trajectories in multi-joint arm movements. *Biol Cybern* 1987;57:257–274. [PubMed: 3689835]
- Gandolfo F, Mussa-Ivaldi FA, Bizzi E. Motor learning by field approximation. *Proceedings of the National Academy of Science USA* 1996;93:3843–3846.
- Georgopoulos AP, Caminiti R, Kalaska JF. Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. *Experimental Brain Research* 1984;54:446–454.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience* 1982;2:1527–1537. [PubMed: 7143039]
- Georgopoulos AP, Schwartz AB, Kettner RE. Neural population coding of movement direction. *Science* 1986;233:1416–1419. [PubMed: 3749885]
- Ghez, C.; Krakauer, JW.; Sainburg, RL.; Ghilardi, MF. Spatial representation and internal models of limb dynamics in motor learning. In: Gazzaniga, MS.; Gazzaniga, MS., editors. *The New Cognitive Neurosciences*. MIT Press; Cambridge, MA: 2000. p. 501-514.
- Goodbody SJ, Wolpert DM. Temporal and amplitude generalization in motor learning. *Journal of Neurophysiology* 1998;79:1825–1838. [PubMed: 9535951]
- Gribble PL, Ostry DJ. Compensation for interaction torques during single- and multijoint limb movement. *Journal of Neurophysiology* 1999;82:2310–2326. [PubMed: 10561408]
- Helms Tillery SI, Soechting JF, Ebner TJ. Somatosensory cortical activity in relation to arm posture: nonuniform spatial tuning. *Journal of Neurophysiology* 1996;76:2423–2438. [PubMed: 8899615]
- Hollerbach JM, Flash T. Dynamic interactions between limb segments during planar arm movement. *Biological Cybernetics* 1982;44:67–77. [PubMed: 7093370]
- Humphrey DR, Schmidt EM, Thompson WD. Predicting measures of motor performance from multiple cortical spike trains. *Science* 1970;170:758–762. [PubMed: 4991377]
- Hwang EJ, Donchin O, Smith MA, Shadmehr R. A Gain-Field Encoding of Limb Position and Velocity in the Internal Model of Arm Dynamics. *Public Library of Science Biology* 2003;1:209–220.
- Johnson MTV, Ebner TJ. Processing of multiple kinematic signals in the cerebellum and motor cortices. *Brain Research Reviews* 2000;33:155–168. [PubMed: 11011063]
- Kakei S, Hoffman DS, Strick PL. Muscle and movement representations in the primary motor cortex. *Science* 1999;285:2136–2139. [PubMed: 10497133]
- Kalaska JF, Cisek P, Crammond DJ. Effector-independent activity in primate dorsal premotor cortex (PMd) during instructed delay tasks. *Society for Neuroscience Abstracts*. 359.2. 2000
- Krakauer JW, Pine ZM, Ghez C. Visuomotor transformations for reaching are learned in extrinsic coordinates. *Society for Neuroscience Abstracts* 1999;25:2177.
- Krouchev NI, Kalaska JF. Context-dependent anticipation of different task dynamics: rapid recall of appropriate motor skills using visual cues. *Journal of Neurophysiology* 2003;89:1165–1175. [PubMed: 12574490]
- Lackner JR, Dizio P. Rapid adaptation to coriolis force perturbations of arm trajectory. *Journal of Neurophysiology* 1994;72:299–313. [PubMed: 7965013]
- Li CSR, Padoa-Schioppa C, Bizzi E. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 2001;30:593–607. [PubMed: 11395017]
- Malfait N, Shiller DM, Ostry DJ. Transfer of motor learning across arm configurations. *Journal of Neuroscience* 2002;22:9656–9660. [PubMed: 12427820]
- Moran DW, Schwartz AB. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology* 1999;82:2676–2692. [PubMed: 10561437]

- Muellbacher W, Ziemann U, Wissel J, Dang N, Kofler M, Facchini S, Boroojerdi B, Poewe W, Hallett M. Early consolidation in human primary motor cortex. *Nature* 2002;415:640–644. [PubMed: 11807497]
- Mussa-Ivaldi FA, Hogan N, Bizzi E. Neural, mechanical and geometric factors subserving arm posture in humans. *Journal of Neuroscience* 1985;5:2732–2743. [PubMed: 4045550]
- Nezafat R, Shadmehr R, Holcomb HH. Long-term adaptation to dynamics of reaching movements: a PET study. *Experimental Brain Research* 2001;140:66–76.
- Osu R, Hirai S, Yoshioka T, Kawato M. Random presentation enables subjects to adapt to two opposing forces on the hand. *Nature Neuroscience* 2004;7:111–112.
- Poggio T. A theory of how the brain might work. *Cold Spring Harbor Symposium on Quantitative Biology* 1990;55:899–910.
- Pouget A, Dayan P, Zemel R. Information processing with population codes. *Nature Reviews Neuroscience* 2000;1:125–132.
- Pouget A, Sejnowski TJ. Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience* 1997;9:222–237.
- Rao AK, Shadmehr R. Contextual cues facilitate learning of multiple models of arm dynamics. *Society for Neuroscience Abstracts*. 27[302.4]. 2001
- Schwartz AB. Direct cortical representation of drawing. *Science* 1994;265:540–542. [PubMed: 8036499]
- Scott SH, Kalaska JF. Reaching movements with similar hand paths but different arm orientation: I. Activity of individual cells in motor cortex. *Journal of Neurophysiology* 1997;77:826–852. [PubMed: 9065853]
- Sergio LE, Kalaska JF. Systematic changes in directional tuning of motor cortex cell activity with hand location in the workspace during generation of static isometric forces in constant spatial directions. *Journal of Neurophysiology* 1997;78:1170–1174. [PubMed: 9307146]
- Sergio LE, Kalaska JF. Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *Journal of Neurophysiology* 1998;80:1577–1583. [PubMed: 9744964]
- Seung HS, Sompolinsky H. Simple models for reading neuronal population codes. *Proceedings of the National Academy of Science U S A* 1993;90:10749–10753.
- Shadmehr R, Brashers-Krug T. Functional stages in the formation of human long-term motor memory. *Journal of Neuroscience* 1997d;17:409–419. [PubMed: 8987766]
- Shadmehr R, Holcomb HH. Neural correlates of motor memory consolidation. *Science* 1997;277:821–825. [PubMed: 9242612]
- Shadmehr R, Holcomb HH. Inhibitory control of motor memories: a PET study. *Experimental Brain Research* 1999;126:235–251.
- Shadmehr R, Moussavi ZMK. Spatial generalization from learning dynamics of reaching movements. *Journal of Neuroscience* 2000;20:7807–7815. [PubMed: 11027245]
- Shadmehr R, Mussa-Ivaldi FA. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience* 1994;14:3208–3224. [PubMed: 8182467]
- Steinberg O, Donchin O, Gribova A, Cardosa dO, Bergman H, Vaadia E. Neuronal populations in primary motor cortex encode bimanual arm movements. *European Journal of Neuroscience* 2002;15:1371–1380. [PubMed: 11994131]
- Thoroughman KA, Shadmehr R. Electromyographic correlates of learning internal models of reaching movements. *Journal of Neuroscience* 1999;19:8573–8588. [PubMed: 10493757]
- Thoroughman KA, Shadmehr R. Learning of action through adaptive combination of motor primitives. *Nature* 2000;407:742–747. [PubMed: 11048720]
- Tolhurst DJ, Thompson ID. Organization of neurones preferring similar spatial frequencies in cat striate cortex. *Experimental Brain Research* 1982;48:217–227.
- Turner RS, Anderson ME. Pallidal discharge related to the kinematics of reaching movements in two dimensions. *Journal of Neurophysiology* 1997;77:1051–1074. [PubMed: 9084582]

Figure 1.

Experimental setup and typical data. **(A)** Subjects hold the handle of a robot and reach to a target. The plot shows hand trajectory (dots are 10 ms apart) for typical movements to 8 targets in the null field, i.e., robot motors disengaged. **(B)** Typical force fields produced by the robot. Forces are plotted as a function of hand velocity. **(C)** Average hand trajectories ($\pm$ SD) for movements during the initial trials in the saddle force field. **(D)** Simulation results for reaches in the saddle force field. **(E)** Hand trajectories during catch trials in the saddle field. **(F)** Simulation results for catch trials in saddle field. Redrawn from Shadmehr and Brashers-Krug (1997d).

Figure 2.

Generalization from one configuration of the arm to another. **(A)** Subjects trained with the hand at the “left” workspace in the force field shown and were then tested at a “right” workspace in a different field. **(B)** Directional tuning for the biceps muscle. For each direction of movement, muscle activity was averaged and is plotted during the null field (solid line) and after adaptation to the force field (dashed line). The gray wedge indicates preferred direction of the muscle. The field at left produces 38° clockwise rotation in the PD of biceps. At right, PD of biceps in the null field rotated 90° with respect to null at left. However, the field at right also produced a clockwise rotation of biceps PD. **(C)** Performance measure (perpendicular displacement of a reach with respect to a straight line) is plotted for bin size of 16 movements. Training at left generalizes to the field at right. Redrawn from Shadmehr and Moussavi (2000).

Figure 3.

Activation of bases that encode limb position and velocity as a gain field. The figure is a polar plot of activation pattern for a typical basis function in the model. The polar plot at the center represents activation for an eight-direction center-out reaching task (targets at 10 cm). Starting point of each movement is the center of the polar plot. The shaded circle represents the activation during a center-hold period and the polygon represents average activation during the movement period. The eight polar plots on the periphery represent activation for eight different starting positions. Each starting position corresponds to the location of the center of each polar plot. The preferred positional gradient of this particular basis function has a

rightward direction. The preferred velocity is an elbow flexion at $62^\circ/\text{s}$. Redrawn from Hwang et al. (2003).

Figure 4.

Consequences of learning internal models with bases that encode static limb position and movement direction as a gain field. Generalization patterns that are produced by these bases predict that certain tasks will be very difficult to learn. **(A)** Because the bases encode state of the limb in intrinsic coordinates (e.g., joint position and velocity), and associate this state to joint torques, a field that is translation invariant in Cartesian coordinates will be hard to learn. **(B)** Because the bases linearly encode static position of the limb, it is difficult to learn to associate movements that are in the same direction with different forces. However, when the movements are sufficiently far apart (about 14 cm), the task becomes learnable. When the center movement is in a null field, the linear encoding of static limb positions makes it so that

the generalization pattern from left and right movements cancel at center, making this an easy task to learn. When the null field is placed to the right, the middle movement in field F2 generalizes to the right movement, making this task hard to learn. (C) A field where forces depend on hand position. Reaching targets are drawn as small circles. In the easy task, movements that are in the same direction (for example, from bottom target to the center target) have the same force pattern. In the hard task, movements that are in the same direction have opposite forces. Part B redrawn from Hwang et al. (2003).

Figure 5.

Estimation of a generalization function from trial-by-trial patterns of error. **(A) Top row:** Black lines are movement errors during 192 movements (out-and-back pattern) in a standard curl field paradigm to 8 directions of targets. Sharp negative spikes are catch trials. Black lines are measured data and gray lines are fit to Eq. (4). Subjects performed 3×192 movements (3 target sets), but data for only one set is shown. **Second row:** In this experiment, subjects practiced in a target set that was not out-and back, but random directions. The shape of the generalization function and compliance are similar to that obtained in the first row. **Third row:** In this experiment, subjects trained in a force field that randomly changed from movement to movement. Despite no obvious learning trends, the generalization function is similar to other “learnable” tasks. **(B)** The estimated generalization function (b in Eq. 4). The generalization function implies that $\sim 18\%$ of the error that was recorded for a movement toward any given direction updated the internal model for that same direction. About 12% of error was generalized to neighboring directions at 135° and 180° . The same subjects were again tested on the same field a second and a third time (2^{nd} and 3^{rd} target sets, each set 192 movements). The generalization functions for all three sets of targets are shown in (B). Little change is seen in these repeated measures. **(C)** The estimated compliance matrix D for each target set. Compliance matrix is plotted by multiplying D by a unit force vector that goes about a circle. The estimates change little with repeated measures. The orientation of the ellipse is consistent with previous estimates of arm stiffness (Mussa-Ivaldi et al., 1985). **(D)** A basis function

consistent with the generalization functions. This particular basis has a preferred velocity at $[0.21, 0.21]$ m/s, corresponding to the peak velocity for a 10 cm movement toward 45° . Dark regions indicate higher activation. Redrawn from Donchin et al. (2003).