

Learning Dynamics of Reaching

**Reza Shadmehr, Opher Donchin, Eun-Jung Hwang, Sarah E. Hemminger,
and Ashwini Rao**

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

Contact information: Reza Shadmehr phone 410-418-5508, fax: 410-614-2458, email:
reza@bme.jhu.edu

Abstract

When one moves their hand from one point to another, the brain guides the arm by relying on neural structures that estimate physical dynamics of the task. For example, if one is about to lift a bottle of milk that appears full rather than empty, the brain takes into account the subtle changes in the dynamics of the task and this is reflected in the altered motor commands. The neural structures that compute the task's dynamics are "internal models" that transform the desired motion into motor commands. Internal models are learned with practice and are a fundamental part of voluntary motor control. What do internal models compute, and which neural structures perform that computation? We approach these problems by considering a task where physical dynamics of reaching movements are altered by force fields that act on the hand. Experiments by a number of laboratories on this paradigm suggest that internal models are sensorimotor transformations that map a desired sensory state of the arm into an estimate of forces, i.e., a model of the inverse dynamics of the task. If this computation is represented as a population code via a flexible combination of basis functions, then one can infer activity fields of the bases from the patterns of generalization. We provide a mathematical technique that facilitates this inference by analyzing trial-to-trial changes in performance. Results suggest that internal models are computed with bases that are directionally tuned to limb motion in intrinsic coordinates of joints and muscles, and this tuning is modulated multiplicatively as a function of static position of the limb. That is, limb position acts as a gain field on directional tuning. Some of these properties are consistent with activity fields of neurons in the motor cortex and the cerebellum. We suggest that activity fields of these cells are reflected in human behavior in the way that we learn and generalize patterns of dynamics in reaching movements.

Introduction

Our arm has inertial dynamics that dictate a complex relationship between motion of the joints and torques. In order to reliably produce even the most simple 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 [1]. Yet the principle is the same for control of biological limbs, as has been confirmed in EMG recordings from the human arm [2]. This has led to the idea that contrary to earlier hypotheses [3], passive properties of muscles are not enough to compensate for the complex physics of our limbs. Rather, the brain must *predict* the specific force requirements of the task in generating the motor commands that eventually reach the muscles.

To illustrate this idea, consider picking up an opaque carton of milk that appears full but has been drained empty. The brain overestimates the mass of the carton by only a couple of pounds (the weight of the missing milk) yet the error is sufficient so that the resulting motor commands produce a jerky motion of the hand. The visual appearance of the bottle apparently retrieves a motor memory in a neural system that predicts the forces that are necessary to move the bottle. Motor commands are constructed based on this prediction and the predicted forces must be accurate if we are to produce smooth movements.

The accuracy of force prediction is particularly important for control of our arm because our hands evolved in large part to support manipulation. For example, a trip to your local natural history museum will confirm that the hand of a chimpanzee has a much longer palm length as compared to a human hand. This means that while we can easily touch our index finger to our thumb and hold an object, say a string that is attached to a yo-yo, a chimpanzee's hand is poorly suited for this. Holding different objects can dramatically change mechanical dynamics of our arm. The neural system that predicts force properties would have to be able to accommodate this variability and adapt. To study the properties of the neural system with which the brain learns to predict forces, we have used a paradigm (Fig. 1) where arm dynamics are systematically changed through forces on the hand [4]. The subject is provided with a target and asked to reach while holding the handle of a robot. When the robot's motors are off (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 once again become smooth and nearly straight. The brain's ability to modify motor commands and predict the novel forces is revealed as sudden removal of force in *catch trials*. Very early in training, the hand's trajectory in the catch trials is a straight path to the target. With further training, trajectories in field trials become straight. More importantly, the 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 called after-effects.

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 can be sufficient to resist perturbations imposed by the robot. However, in a catch trial, 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 direction, predicts the forces that will be imposed on the hand by the robot. One way to do this is to imagine a tape that is played out as a function of time for each movement direction. 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. [5] trained subjects to reach to a small number of targets in a force field and then suddenly 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, brief 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, in performing the reaching movements the neural system had learned to associate the sensory states of the limb—specially limb position and velocity—to forces. The particular order in which those states were visited and the trajectory at the time they were visited (e.g., in a straight line trajectory or in a curved movement) was immaterial. What was important was the region of the state space—the limb’s velocity at a given position—that the reaching movements had visited. 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 [6] tested this by asking whether subjects could adapt to force fields that explicitly depended on time. Remarkably, the experimental results indicated that they could not. When a predictable, time-dependent pattern of force was imposed on reaching movements, generalization trials (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, that prediction depended on an input that described the desired state of the arm.

These experiments suggested that with practice, participants learned a sensory to motor transformation where a velocity-like input signal was transformed into a force-like output signal. This is an *internal model* of the force field.

Neural correlates of learning internal models of dynamics

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 long period of training, it has adapted in the sense that it correctly predicts forces that are produced for typical states visited in reaching movements. 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 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 spinal reflexes, we add to the equations a simple low-gain spring-damper element that stabilizes the limb about the desired trajectory (the straight line). 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. These equations have been detailed in Shadmehr and Mussa-Ivaldi [4].

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 trajectory that is similar to what we have recorded in our participants, that is, a peculiar hooking pattern [4,7]. Now we change the internal model so that it completely takes into account the added dynamics of the force field. That is, we assume that the internal model is fully trained. If we now simulate 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 are consistent with learning an internal model that transformed desired sensory states into forces.

It is an easy next step to extend the mathematical formulation and predict not just the limb's trajectory before and after the internal model adapts to imposed forces, but also a correlate of that adaptation at the level of the neural commands to the muscles. The equations of motion that produced trajectories in Figure 1 included torque generators but not specific muscles. In Thoroughman and Shadmehr [8] we added to these equations constraints regarding function of muscles. The most important constraint was that opposing torques in two antagonistic muscles should be inversely proportional. That is, the more one muscle was activated, the less the antagonist was activated. This assumption allowed us to translate a pattern of expected forces on the hand onto changes in muscle activations. To visualize the changes, we plotted the average magnitude of activation for each simulated muscle as a function of movement direction in hand-centered coordinates and computed a preferred direction (PD) for each muscle [9,10]. The preferred direction for a muscle was the direction of movement in Cartesian coordinates centered on the hand for which the modeled muscle was most active. For example, the simulations predicted that adaptation of the internal model to a clockwise curl field should accompany a clockwise rotation by $\sim 27^\circ$ for elbow muscles and $\sim 18^\circ$ for shoulder muscles. A curl field is a particular force field where the forces are always pushing the hand perpendicular to its current direction of motion. In the clockwise version of this, the force vectors are pointing in the clockwise direction perpendicular to the direction of hand motion. EMG in biceps, triceps, anterior and posterior deltoids in a group of participants confirmed this prediction [8]. This confirmation of the model's prediction illustrated that one way to represent the change in motor commands due to adaptation of the internal model was as rotations in preferred direction of muscles.

The result is not surprising because of course the brain would have to change the commands to the muscles if forces are to be produced to counter the effects of the robot imposed field. Rather, the results are useful for the following reasons. First, because the model 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 by which 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 [11], one can suggest that the rotation in EMG that the model predicts 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 [12], Bizzi and colleagues reported that task related cells in M1 underwent a median clockwise shift in PD of 16° . However, whereas we had found that the EMG patterns returned to baseline conditions once the field was turned off (i.e., a washed out phase), many cells in the motor cortex kept the change in their PD. Bizzi and colleagues labeled these "memory" cells. Remarkably, whereas the memory cells kept their clockwise change in PD when the force field was turned off, another population of cells that had not changed their PD in the force field now changed their PD in a counter clockwise direction when the field was removed. Therefore, after completion of training and return to a null field condition, distinct population of cells maintained the effect of the field in terms of rotations in their PD. As demonstrated many times previously, the cells in M1 are not simply upper motoneurons. Rather, in this case there are likely involved in representing the memory of the internal model.

If motor cortex was involved in representing the internal model, what might be a testable behavioral consequence of this hypothesis? One of the consistent properties of task related cells in M1 is that if a cell is active for reaching movement to a group of targets at one arm configuration, it is also likely to be active when the configuration of the arm is changed and the targets are moved to the new workspace. However, the change in the workspace often results in a

change in the PD of cells. The PD of a typical M1 cell will rotate approximately with the shoulder angle [13,14]. The reason for this rotation is possibly the observation that many M1 cells are sensitive to force requirements of the task. Sergio and Kalaska [11] 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 memory cells that Bizzi and colleagues found [12] 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 will change as the workspace in which the reaching movements are done changes. We would expect that their PDs will rotate with the shoulder in a way similar to rotation of PDs in arm muscles in the same task. Imagine that the force field related changes in PD and the posture related changes in PD 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 training elsewhere should be observed, i.e., we should observe generalization. This is a behavioral prediction of the model.

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 [15,16]. Indeed, in their force field learning experiment, Bizzi and colleagues [12] 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 assumes 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, would be predicted to be the ones that contribute most to the representation of internal models for dynamics of reaching movements.

Generalization as a function of arm’s position

Because M1 cells that have muscle-like properties in their tuning tend to rotate their PD with the shoulder angle, learning an internal model with these cells should result in a specific pattern of generalization. For example, consider adaptation to a force field described by $\mathbf{f} = B_1\dot{\mathbf{x}}$, where \mathbf{f} is a force vector acting on the hand, $\dot{\mathbf{x}}$ is a hand velocity vector, and $B_1 = [-11, -11; -11, 11]$ N.sec/m. This is a “saddle field” (Fig. 1b) where movements toward 120° and 300° encounter a resistive force and movements toward 210° and 30° encounter an assistive force. If the right arm is near the horizontal plane and the shoulder is flexed so that the hand is at a “left” workspace (meaning that reaching movements are performed in a flexed posture for the shoulder), preferred direction of triceps is about 90°. When a subject trains in the field, one observes a 30° clockwise rotation in the PD of triceps. Now imagine that there are cells in the motor cortex that also 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 a right workspace, we would expect that 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., $180^\circ - 30^\circ$), and movement of the hand to the right workspace should bring the PD to 60° (i.e., $180^\circ - 30^\circ - 90^\circ$). If the subject had not practiced movements in the field, 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 “memory cells” that rotated their PD at the left workspace maintain 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 the shoulder angle.

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 [17]. For our saddle field, this procedure produces the surprising result that a 90° rotation in the shoulder results in a 180° rotation of the matrix B_1 . This theoretical result means that the force field described by B_1 should be generalized to $-B_1$ at the right workspace. 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 [18]. 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 [17]. The results suggested the intriguing theory that not only the motor cortex might take part in representing the memory of the internal model, but that the properties of activity fields (or tuning) of cells in M1 might be related to the behavioral patterns of generalization in force fields. The property of activity fields that is relevant in this case is the change in PD as a function of shoulder angle.

Computing an internal model with a population code

How does one test the idea that activity fields of certain cells influence patterns of generalization during learning of reaching movements? Alternatively, how does one infer the shape of the activity fields from the patterns of behavioral generalization? We need to advance beyond a description of the input-output variables that are encoded by internal models (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 internal models.

One of the most widely used models of neural computation is population coding. While the idea of using populations of neurons to code variables of interest is old [19], 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 [20]. To motivate our approach, let us put aside for now the problem of predicting force as a function of velocity and consider the simpler problem of representing direction of movement of the hand. Georgopoulos et al. [20] 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, a vector of unit length \mathbf{w}_i . The movements were in a plane. Therefore, \mathbf{w} is a 2D vector that might point along any direction about a unit circle. In a given trial, imagine that the movement direction is θ , and each cell i discharges by amount r_i . This discharge can be decomposed into two terms. The first term is an average response $g_i(\theta)$ 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 :

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

In this equation, the first term is the tuning curve of the cell and the second term is noise. 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° [21]. 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 [22]. If cells did not have this noise and we could record from a large number of cells at the same time, we would note that cell j happened to fire most during some movement and estimate the movement direction $\hat{\theta}$ to be the preferred direction of that cell:

$$\hat{\theta} = \mathbf{w}_j$$

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 [20]. 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:

$$\hat{\theta} = \sum_i \mathbf{w}_i r_i = \sum_i \mathbf{w}_i g_i(\theta) + \mathbf{w}_i n_i$$

This approach is better in the sense that it produces a smaller variance in its estimate from trial to trial (when the movement direction is fixed) than in 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 [23]. 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 [20,24].

The example of population coding above is an instance of neural computation of an *identity map*, i.e., a map where the output is an estimate of the input variable. In general, a population code could also be used to map an input variable x into any other variable y [25,26]. In this case, the tuning curves of the neurons that participate in this computation become the basis functions with which the output is approximated. Basis functions are a set of functions such that when they are linearly combined, they can approximate almost any linear or non-linear function. For example, Pouget and Sejnowski [27] suggested that neurons in the parietal cortex might serve as basis functions with which the brain could compute 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 [28]. 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:

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

The appropriate weighting \mathbf{w}_i would have to be learned to form this map. However, Pouget and Sejnowski [27] 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 more general problem of computing nonlinear maps. Another point is that 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 \mathbf{x} (i.e., the position and velocity of the arm) into force \mathbf{f} . Let us assume this neural computation is performed via a population code. Each neuron that participates in this computation has a tuning curve g_i that describes the average discharge of that cell as a function of hand position and velocity. Each cell also has a *preferred force vector* \mathbf{w}_i . The population vector response of the network is:

$$\hat{\mathbf{f}} = \sum_i \mathbf{w}_i g_i(\mathbf{x}) + \mathbf{w}_i n_i \quad (1)$$

We now have a framework to relate tuning properties with behavioral generalization. Consider the following experiment: participants are initially trained with force field \mathbf{f}_1 for movement along arm state \mathbf{x}_1 . The error that they experience in a movement changes the preferred force vector \mathbf{w} . Assuming Hebbian learning rules, the weight change will be maximum for those neurons that happened to be most active about state \mathbf{x}_1 . The subject is then asked to make a movement with the arm along state \mathbf{x}_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 \mathbf{x}_1 must have been broad enough to not only be active for \mathbf{x}_1 , but also \mathbf{x}_2 . Therefore, if the internal model is represented via a population code, then generalization is described by the shape of the tuning curves of the elements of computation.

From the experiment in which we observed generalization in one arm from one workspace to another workspace [17], we can now conclude that the tuning functions could not have had “preferred positions” of the hand, in the sense that this implies sharply tuned activity functions around that position. If they did, then experiencing force at one hand position could not be generalized to another hand position far away.

Inferring coding of limb position and velocity from patterns of generalization

The 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 discharge varies with arm position and velocity, we note that cell activity in M1 is modulated globally and often linearly as a function of limb position [29], and cells have preferred directions of movement that often change as a function of the shoulder angle [13]. To capture these observations, Hwang et al. [30] hypothesized that cells that are involved in computing the internal model have tuning functions that are described as follows:

$$\begin{aligned} g(\underline{q}, \underline{\dot{q}}) &= g_{\text{position},i}(\underline{q}) \cdot g_{\text{velocity},i}(\underline{\dot{q}}) \\ g_{\text{position},i}(\underline{q}) &= \mathbf{k}^T \cdot \underline{q} + b \\ g_{\text{velocity},i}(\underline{\dot{q}}) &= \exp\left(-\frac{\|\underline{\dot{q}} - \underline{\dot{q}}_i\|^2}{2\sigma^2}\right) \end{aligned} \quad (2)$$

The above function represents output of a basis function. The position-dependent term is a linear function that encodes joint angles, $\underline{q} = (\theta_{\text{shoulder}}, \theta_{\text{elbow}})$, while the velocity-dependent term encodes joint velocities. Fig. 2C plots this function for reaching movements to various directions at various starting positions. The basis is sensitive to both the static position of the limb and its velocity. It combines the two via a gain field, i.e., directional tuning is modulated multiplicatively as a function of limb position. As a result, both the preferred direction of the tuning and the depth of modulation vary with the starting position of the reach. The gradient vector \mathbf{k} reflects sensitivities for the shoulder and elbow displacement, and b

is a constant. The velocity-dependent term is a Gaussian that encodes joint velocity \dot{q} centered on the preferred velocity \dot{q}_i .

The multiplicative nature of this encoding is one of the requirements of basis functions [27]. However, we should note that the properties that we assigned to the tuning function are not unique to cells in M1. For example, in addition to M1, linear modulation of discharge with respect to limb position has been observed in the spinocerebellar tract [31] and the somatosensory cortex [32]. Tuning functions that have preferred directions or velocities of movement have been reported in the cerebellar cortex [33]. Indeed, it appears that a linear modulation of discharge with respect to limb position and a Gaussian tuning with respect to arm velocity may be a fundamental property of many cells in the motor system.

Consider a situation in which the internal model is constructed as a linear combination of these non-linear bases. How would their activity fields be reflected in behavioral generalization? With training at a given arm configuration, the preferred force vector of some of these bases will change. The change will occur in those bases that happen to be most active at this arm configuration. The way that these active bases change their static discharge with arm position dictates how far in position space the learning will generalize.

Next, note that because the bases encode joint velocity and not hand velocity, the preferred direction of movement (which is represented in Cartesian coordinates) will rotate for some of the cells as the shoulder joint changes position. The way that the elements change their PDs with arm position dictates the coordinate system of the generalization.

Let us first examine how adaptation with these elements results in generalization in terms of spatial displacement of the hand. Suppose one trains subjects in a small workspace (of reaching) with the arm at one configuration, and then tests them at another workspace. 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 spatial direction of movement in two different spatial locations.

To quantify how people generalize forces as a function of hand position, Hwang et al. [30] performed an experiment in which participants made reaching movements from different starting locations to targets in the same direction (i.e., the direction of movement was the same in joint velocity space). The arm was covered by a semitransparent screen, upon which an overhead projector painted targets. The handle of the robot, held by the subject, housed an LED that was visible on the screen. Therefore, subjects had visual and proprioceptive cues regarding hand position. In these parallel movements, opposite curl fields acted on the hand in the *left* and *right* targets (Fig. 2a). A null field was always present in the *middle* target. The robot brought the hand to a random start position after completion of each movement. When the left and right targets were 24 cm apart, subjects could learn these movements and little after-effects were present in the middle target (Fig. 2b). As targets were brought closer, learning became more difficult and interference became apparent in the middle target.

Fig. 2d plots the dependence of interference on the spatial distance. The shape of this interference pattern constrains the gain k , i.e., the slope of the gain field. We simulated learning with Eqs. 1 and 2 and kept k the same for all bases. We found the gain k that produced an interference pattern similar to what we had seen in our subjects. Now, we took these same bases and asked whether they could also explain the global generalization pattern that we had seen earlier: that is, when subjects learned a field at the left workspace and generalized to 80 cm away

at the right workspace [17]. Remarkably, we found that the same k also explained the amount of generalization that we had seen in that experiment (Fig. 2e). Furthermore, the forces were generalized in terms of torques on the joints, rather than forces on the hand. Therefore, bases that are linearly modulated by arm position (Fig. 2c) and encode joint velocity rather than hand velocity appear to explain the pattern of interference in Fig. 2d, the intrinsic coordinate system of generalization, along with its large spatial generalization in Fig. 2e.

The reason why we assumed that the bases linearly coded hand position space was because discharge of cells in the spinocerebellar tract [31], somatosensory cortex [32], and M1 [29] 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 [34,35]. Therefore, generalization as a function of static position of the arm seems consistent with bases that encode limb velocity and position in intrinsic, joint-like coordinates. The bases appear to be tuned to direction of movement and that tuning is multiplicatively modulated as a linear function of limb position.

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, one of us (O.D.) recently observed 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 [36]. Kalaska and colleagues also observed similar properties of tuning functions for cells in the premotor cortex [37]. An important technical point in both of these experiments is that both arms performed reaching movements directly in front the animal in the same workspace.

Let us now imagine that this invariance of PD with respect to the right and left arms is also a property of the cells that we hypothesized were involved in computing an internal model. Consider a subject that trained with her right hand in a curl force field, resulting in a rotation of PDs in some 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 arm to the left arm. Furthermore, it makes the surprising prediction that the coordinate system of generalization from one arm to another should be in an extrinsic, Cartesian-like coordinates!

The prediction is surprising 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.

Criscimagna-Hemminger et al. [38] tested this using 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. Was there generalization from one arm to the other, and if so, was the transfer in extrinsic coordinates?

We first quantified generalization in right-handed individuals from right to left. In comparing performance of the extrinsic and intrinsic groups to a control group (Fig. 3), we found significant inter-limb generalization in extrinsic coordinates only. In the transfer trials, the extrinsic group's performance with the left hand was significantly better than controls (Fig 3c), whereas in the intrinsic group performance with the left hand was significantly worse than controls (Fig 3c).

We next quantified generalization in right-handed individuals from left to right. Subjects trained in a procedure similar to Fig. 3a, except that they trained with the left arm in a curl field and were then tested with the right arm on either the extrinsic or the intrinsic representation of the same field. We found that performance during test of generalization was not significantly different from controls on either the extrinsic or the intrinsic representation (data not shown). This suggested that in humans, generalization of arm dynamics in right-handed individuals occurred only from the dominant right to the left arm, and its coordinate system was extrinsic in the workspace that we tested.

Interestingly, the same pattern of generalization was observed in a right-handed callosotomy patient. In callosotomy patients, when visual information is restricted to one hemisphere, that hemisphere can produce a reaching movement with the ipsilateral arm [39]. This is because a small but significant number of corticospinal projections to the proximal muscles of the arm are from the ipsilateral hemisphere [40]. However, converging evidence [41] indicates that the dominant hemisphere may have a significant role in controlling the non-dominant arm, but not vice versa. The fact that we observed generalization only from the dominant to the non-dominant arm suggests that the cells in the non-dominant hemisphere that participate in learning in this task are not tuned to movements with the ipsilateral arm. In contrast, cells in the dominant hemisphere (left) are tuned to movements of both arms and they maintain their PD when the task is performed with one or the other arm. Therefore, the fact that some M1 cells maintain their PD irrespective of the arm that is used for reaching is consistent with the coordinate system of inter-limb generalization that we observed in learning of force fields.

Activity fields with respect to color of the target

The main claim of the hypothesis is that tuning properties of 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 proprioceptive state of the arm. This made sense because cells in the motor cortex are sensitive to these states and we wished to infer how this sensitivity influences learning. Let us now consider a task where the forces do not depend on state of the arm. For example, imagine a reaching movement where position or velocity (or any other kinematic variable of the arm) 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 on the color of the target.

If the cells 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. [42] 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 [43]. On any given trial, the color of the target was randomly chosen as either red or green. Because the movement was always in the same direction, the pattern of forces on that movement depended exclusively on this cue. We trained subjects ($n=3$) extensively on this task, providing them with

over 3000 trials, spread over three days. Remarkably, in catch trials we consistently found no evidence of after-effects (Fig. 4) and performance showed no suggestion of adaptation. However, with longer training [44] or with explicit instruction about the nature of the forces [45], it is possible to associate color to force fields. The remarkable difficulty in learning of this apparently simple task makes the prediction that the activity fields of the bases are typically only weakly modulated by color of the target.

Inferring shape of the tuning curves from patterns of generalization

The assumption about formation of an internal model via a population code means that when one measures generalization, one might be able to infer activity fields 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 second problem is that in motor control, we have to consider coordinate systems. Generalization depends not only on the distance between training and test locations, but also on the coordinate system in which that space is measured. 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 this section, we show that all four questions can be approached.

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 [46,47]. 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 preferred force vector associated with a basis is likely to change most for those bases that are most active. 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 bases are wide enough to encompass both the movements, or, alternatively, whether they are so narrow that they cannot span the gap. The generalization function depends on how the tuning curves encode movements and one can acquire critical information regarding generalization from the trial-to-trial variations in behavior.

The importance of this idea is that 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 movement-to-movement 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.

Begin with the assumption that the internal model is composed of a linear combination of an unknown set of bases (Eq. 1). These bases encode the state of the arm (in the current case, only velocity because we limit movement to a small spatial workspace). Assume that the purpose of learning is to minimize the difference between expected force in a movement and the actual force, and that adaptation is through gradient descent that results in modification of the preferred force vector associated with each basis. How does the shape of the bases affect the pattern of trial-to-trial errors?

Donchin et al. [48] demonstrated a method to quantify generalization from the trial-to-trial measures of behavior. The idea was to represent adaptation with basis functions as a hidden-state dynamical system. Once that system was expressed mathematically, fitting it to the data would provide an estimate of the generalization function.

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 and would like to understand how that generalization depends on the shape of the basis functions g . To simplify matters, let us ignore the noise in Eq. 1, assume that the bases are only a function of velocity because our reaching movements will all be performed in a small spatial workspace, and rewrite that equation in terms of vector quantities:

$$\hat{\mathbf{f}} = W\mathbf{g}(\dot{\mathbf{x}})$$

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

$$\mathbf{g}(\dot{\mathbf{x}}) = [g_1(\dot{\mathbf{x}}) \quad \cdots \quad g_m(\dot{\mathbf{x}})]^T$$

Here $\hat{\mathbf{f}}$ is a 2x1 vector. It is an estimate of actual force \mathbf{f} . The error in our estimate is:

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

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

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

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

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

After performing trial n , the error in that trial $\tilde{f}^{(n)}$ 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 α :

$$w_{ij}^{(n+1)} = w_{ij}^{(n)} + \alpha \tilde{f}_i^{(n)} g_j(\dot{\mathbf{x}}^{(n)})$$

Writing this in vector form we have:

$$W^{(n+1)} = W^{(n)} + \alpha \tilde{\mathbf{f}}^{(n)} \mathbf{g}(\dot{\mathbf{x}}^{(n)})^T$$

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

$$W^{(n+1)} \mathbf{g}(\dot{\mathbf{x}}^{(n+1)}) = W^{(n)} \mathbf{g}(\dot{\mathbf{x}}^{(n+1)}) + \alpha \tilde{\mathbf{f}}^{(n)} \mathbf{g}(\dot{\mathbf{x}}^{(n)})^T \mathbf{g}(\dot{\mathbf{x}}^{(n+1)})$$

This is equivalent to:

$$\hat{\mathbf{f}}^{(n+1)}(\dot{\mathbf{x}}^{(n+1)}) = \hat{\mathbf{f}}^{(n)}(\dot{\mathbf{x}}^{(n+1)}) + \underbrace{\alpha \mathbf{g}(\dot{\mathbf{x}}^{(n)})^T \mathbf{g}(\dot{\mathbf{x}}^{(n+1)})}_{\text{generalization function}} \hat{\mathbf{f}}^{(n)} \quad (3)$$

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 projection of the bases in trial n upon the bases evaluated at trial $n+1$. Intuitively, we see that the projection will be largest when the two consecutive trials are along the same velocity. The shape of the tuning functions will determine the generalization as the distance between the two movements changes in the state space.

In an artificial system, a “trial” would be an example data point. That is, the internal model would make a guess about the force at a particular velocity, and a “teacher” would provide the actual force, and the error would be used to modify the weights. The shape of the bases would then dictate the generalization to the next trial. However, movements are not a single point in velocity space, but a trajectory. Because of delays in sensorimotor feedback, we could reasonably assume that feedback about the actual forces might not be available to update the internal model until the movement is completed. After completion of movement, the internal model would be updated along the entire trajectory of the desired velocity. Because this desired trajectory is along a straight line to the target, we will represent it simply as a direction. For example, if we have 8 directions of movement, the “generalization function” is a matrix of size 8×8 . Element (i, j) of this matrix describes the fraction of error in movement direction i that is generalized to direction j . To simplify things, we can assume that what matters is not the specific directions of the two consecutive movements, but rather their angular distance. In this way, the generalization function becomes a vector of size 8×1 where each element indicates generalization between two consecutive movements that are separated by an angular distance of $0, 45, \dots, 270^\circ$. Let us call this generalization function b . With this approximation, a trial becomes a single reaching movement.

Now the important thing to notice is that despite the fact that the error in movement n potentially affects the internal model for all possible directions of movement, we can observe that effect for only one direction, the actual direction for which movement was made in trial $n+1$. Therefore, while the effects of the generalization to the other 7 possible directions are hidden to us, they nevertheless exist and we must account for them in order to accurately represent the trial-to-trial changes in the internal model. To do so, let k be an integer variable that can take a value from 1 to 8. It represents the possible directions of movement that could occur in trial $n+1$. After movement n is completed, the internal model is updated in all these directions, and we have:

$$\hat{\mathbf{f}}_k^{(n+1)} = \hat{\mathbf{f}}_k^{(n)} + b(k - \dot{x}^{(n)})\tilde{\mathbf{f}}^{(n)} \quad k = 1, \dots, 8$$

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 $\hat{\mathbf{f}}$. 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 with respect to where it “should be” if the internal model was perfect. Let us call that position error vector \mathbf{y} . Let us further assume that it will be related to the force error $\tilde{\mathbf{f}}$ in the estimate of the internal model via a compliance matrix D . This matrix relates how force error produces a displacement from the intended trajectory. We now have the following:

$$\begin{aligned} \mathbf{y}^{(n)} &= D\tilde{\mathbf{f}}^{(n)} \\ \hat{\mathbf{f}}_k^{(n+1)} &= \hat{\mathbf{f}}_k^{(n)} + b(k - \dot{x}^{(n)})\tilde{\mathbf{f}}^{(n)} \quad k = 1, \dots, 8 \end{aligned} \quad (4)$$

Now let us introduce a new variable z , and define it as follows:

$$\mathbf{z}_k^{(n)} \equiv D\hat{\mathbf{f}}_k^{(n)}$$

With substitution of the above equation into Eq. (4), we arrive at a coupled dynamical system:

$$\begin{aligned} \mathbf{y}^{(n)} &= D\mathbf{f}^{(n)} - \mathbf{z}_{k^{(n)}}^{(n)} \\ \mathbf{z}_k^{(n+1)} &= \mathbf{z}_k^{(n)} + b(k - \dot{x}^{(n)})\mathbf{y}^{(n)} \quad k = 1, \dots, 8 \end{aligned} \quad (5)$$

$\mathbf{y}^{(n)}$ is the error vector on n^{th} movement, made in direction k ; $\mathbf{f}^{(n)}$ is the force experienced in that movement, and is scaled by a compliance-like matrix D . Compliance is the inverse of stiffness. Whereas stiffness describes force produced when a body is displaced, compliance describes displacement produced when a body experiences force. When an error occurs in a movement, the internal model is updated (reflected in the 8 equations). b is the generalization function that characterizes the effect of error that was experienced in a given state on all other states. We measure a sequence of movement errors $\mathbf{y}^{(n)}$ and fit them to the system in Eq. (5) in order to find the best fit for matrices D and vector b . There are 12 unknown parameters in these two variables. The procedure for fitting these equations to a sequence of movements is provided in Donchin et al. [48]. 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.

We begin by considering the fit of these equations to human data (Fig. 5a). Data from a large group of subjects ($n=75$) was 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 fit the trial-to-trial variations in performance remarkably well (Fig. 5a); and 2) the generalization function B and compliance matrix D remained consistent across repeated measures (Fig. 5b). 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 our out-and-back target sequence. We tested a new group of subjects ($n=8$) 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 (Fig. 5b).

We next tested another group of subjects ($n=11$) 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, or counter clockwise curl. As the field was random, we did not expect any adaptation. Remarkably, analysis of the trial-to-trial changes in performance produced a generalization function similar to that which we had estimated from trials in which subjects learned 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 our behavioral data is shown in Fig. 5c.

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 may exist a single basis function that encodes movement kinematics and explains learning in all of these tasks. Our best guess today is that this function encodes hand position of the contralateral arm linearly and hand velocity with a bimodal activation function:

$$g_{velocity,i}(\dot{q}) = \exp\left(-\frac{\|\dot{q} - \dot{q}_i\|^2}{2\sigma^2}\right) + \frac{1}{s} \exp\left(-\frac{\|\dot{q} + \dot{q}_i\|^2}{2\sigma^2}\right) \quad (6)$$

Relating the inferred activity fields to the neurophysiology of the motor system

Are the bases that we inferred with this abstract model interpretable in terms of the neurophysiology of the motor system? From the patterns of generalization, we inferred that:

- 1) The bases encode hand velocity with a function that has a preferred direction and is modulated broadly but is bimodal.
- 2) The bases encode position of the arm linearly in the horizontal plane. This position coding multiplicatively modulates tuning with respect to direction.
- 3) The preferred direction of the bases rotates 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 only weakly modulated by color of the target.

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 [14,14,29,49,50,50]. 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 [36], premotor cortex [37], 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 [33], whereas no such bimodality is reported in the same task in the primary motor cortex [50].

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 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 elements 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 particular encoding of velocity. In Eq. (6), we have chosen to represent 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, where as cells in M1 generally increase their discharge with increased movement speed [15,50]. Only one study has considered how the internal model generalizes in terms of speed of movement [51]. 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.

Consolidation

Thus far we have been describing learning of internal models using a mathematical framework where acquisition of information is one and the same as memory. In this framework, preferred force vectors associated with the bases change to minimize error in the task. Once the task is over, presumably these changes are maintained and that forms the basis of long-term memory.

The scope of our naiveté was plainly demonstrated when we found that acquisition of memory of an internal model is merely the first step in a sequence of events that eventually results 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 [7]. 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. 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 [52]. 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 [53]. 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 [54,55], vs. at 2 or 4 weeks after initial practice [56].

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.

Major shortcomings of the theory

Both our measures of performance and the construction of the internal model focused on the early component of the reach (typically up to 250 ms), a period when one expects little influence from feedback. Therefore, even if the theory is successful, it only addresses adaptation associated with the motor commands that initiate the movement. However, our recent work [57,58] and those of our colleagues [59,60] has found that with training, the brain also learns to respond to feedback during a movement by producing appropriate motor responses. We currently have no model to account for trial-by-trial adaptation or generalization of this form of adaptation.

The alert reader would also note that while we started our story with the problem of using visual appearance of objects to estimate their dynamics, the theory that we developed made little mention of these cues but rather focused on proprioceptive measures of limb state. How do non-proprioceptive cues like color, spatial cues about the pattern of forces, sequential cues regarding movement order, or cognitive cues affect computation of internal models? In our daily interaction with the environment, it is these cues that must dominate selection and adaptation of internal models. This important question remains poorly understood.

Summary

The specific coding of movement parameters in the neurons that compose internal models have a significant, measurable influence on behavior. That influence can be observed in how our brain learns to predict forces in control of reaching movements. Training to make reaching movements in a force field results in a specific, highly reproducible pattern of force 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 neurons that participate in this computation are the bases with which the force field is approximated. From the patterns of generalization one can infer the shape of these bases: 1) The bases are modulated as a function of hand velocity with a broad function that has a preferred direction and but is bimodal. 2) The bases are modulated linearly with arm position in the horizontal plane and this position coding multiplicatively modulates directional tuning, resulting in a gain field. 3) The preferred direction of the bases rotates 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 weakly modulated by color of the target. These are also some of the properties of cells in the primary motor cortex and the cerebellum.

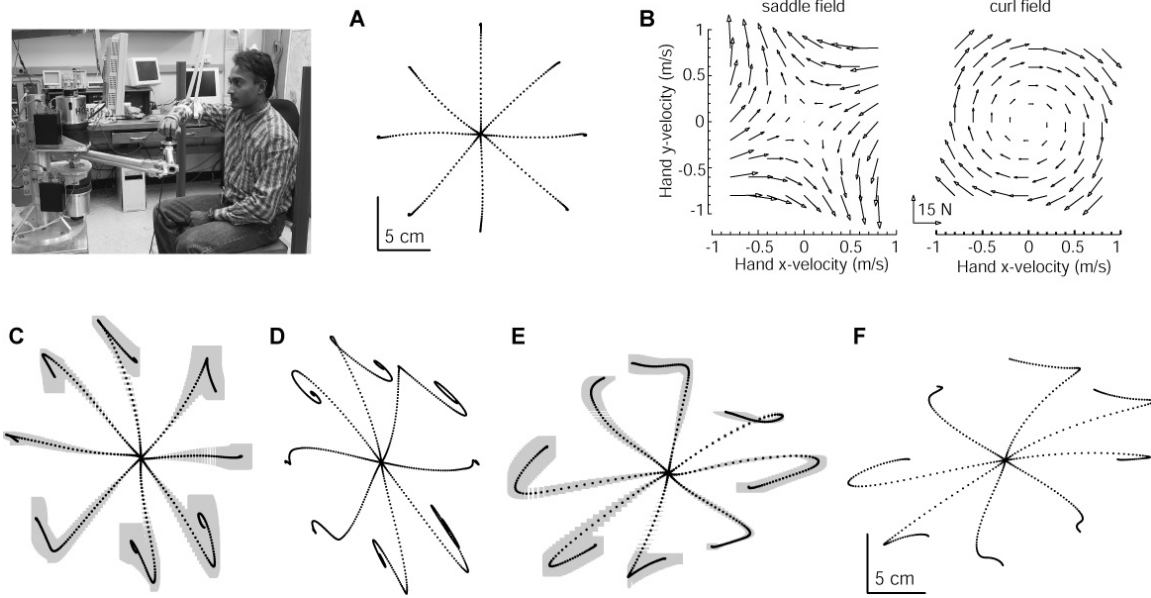


Figure 1. Experimental setup and typical data. **(a)** Subjects hold the handle of the robot and reach to a target. The plot shows hand trajectory (dots are 10 ms apart) for typical movements to 8 targets. **(b)** Examples of two force fields produced by the robot. **(c)** Average hand trajectories (+/- SD) for movements during the initial trials in the saddle field. **(d)** Simulation results for movements in the saddle field. **(e)** Hand trajectories during catch trials. **(f)** Simulation results during catch trials. The controller in this simulation had fully adapted to the field and was expecting the field to be present in these movements. Adapted from Shadmehr, R. and Mussa-Ivaldi, F. A., *J. Neurosci.*, 14, 3208, 1994. Used with permission.

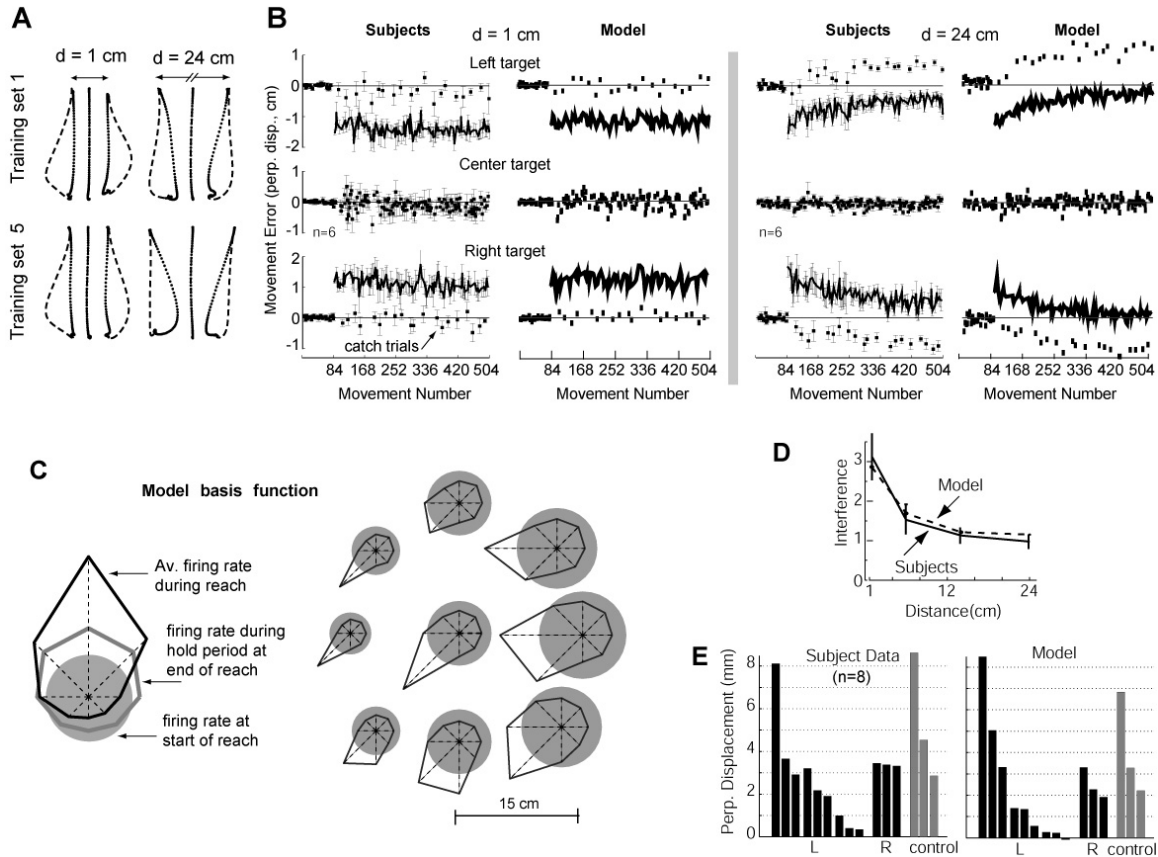


Figure 2. A gain field coding of limb position and velocity in the internal model of arm dynamics. **(a)** Subjects made parallel, 10 cm movements at left, center, and right toward target at -90° . At the left and right targets a curl field was present (clockwise at left, counter-clockwise at right). The field at the center target was always null. Subjects were divided into four groups based on the distance between the left and right targets: $d=1$, $d=6$, $d=14$, and $d=24$ cm. Typical movements for single subjects in the $d=1$ and $d=24$ cm groups are shown (field trials are dash, catch-trials are dotted). The 24 cm group could learn the task: they had clear after-effects and movements in field improved. The 1 cm group did not show after-effects and movements did not improve. **(b)** Group data (mean \pm SD) for the 1 and 24 cm groups ($n=6$ in each group). Field trials are connected with a solid line, catch-trials are dots. Note the increased interference in the center target for the 1 cm group as field trials start. **(c)** We simulated learning with bases that encoded static hand position and movement velocity multiplicatively via a linear function of position and a Gaussian function of velocity. The firing rate of one basis function is plotted for 8 directions of movement during movement time, and at the hold time at start and end of the movement. Firing rate is also plotted for another basis for movements that start from 9 different start positions. Note that the preferred direction of the basis rotates with start position. Firing at start of reach varies with start of reach. Because of multiplicative interaction of position sensitivity with directional tuning, depth of modulation varies with start position. **(d)** An interference measure that quantifies how the left and right movements affected the movements at the center. This is the ratio of standard deviations (SD) for the center movements 420-504 vs. 1-84. As distance between the targets decreases, interference increases. The bases appear to account for the generalization pattern of our subjects over this small displacement of the hand. **(e)** We asked whether the bases could also account for the generalization pattern that we had observed in Shadmehr and Moussavi (2000). In that report, we observed that after subjects were

trained in a small workspace at a “left” workspace, performance was significantly better than naïve when they were tested at a “right” workspace (80 cm away) in a rotated version of the same field. The bases produced similar patterns of generalization as our subjects over this large displacement. Adapted from Hwang, E. J., Donchin, O., Smith, M. A., and Shadmehr, R., *PLoS Biol.*, 1, 209, 2003. Used with permission.

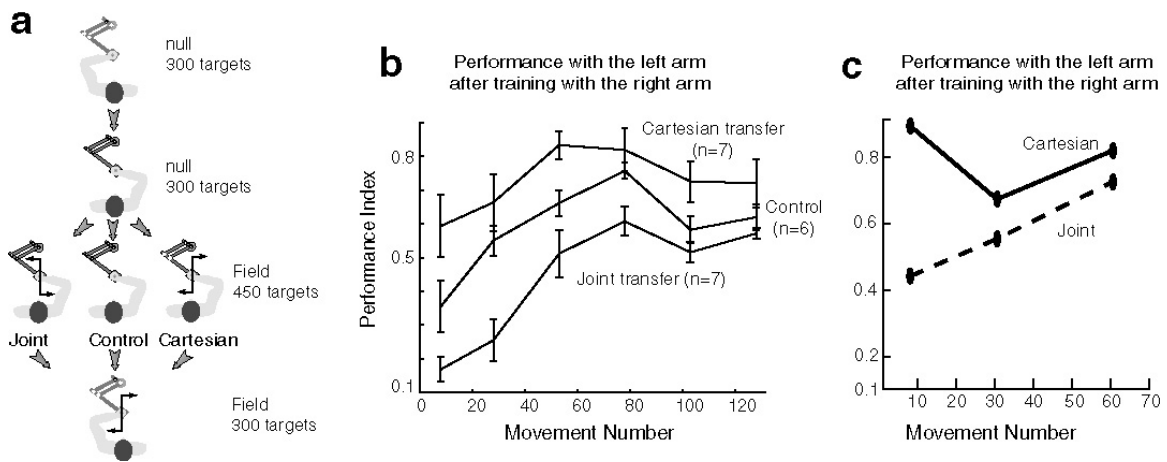


Figure 3. Inter-limb generalization. **(a)** Right handed subjects ($n=20$) trained with the left and then the right arm in a null field. They were then assigned to one of 3 groups: training with the right arm in a clockwise curl, counter clockwise curl, or a control group that received further training in the null field. All of these groups were then tested with their left arm on a clockwise curl field. For subjects who trained with their right arm on a clockwise curl field, this was a test of inter-limb transfer in extrinsic (Cartesian) coordinates. For subjects who trained with their right arm in a counter clockwise curl field, this was a test of inter-limb transfer in intrinsic (joint) coordinates. **(b)** Performance index in the test trials (mean \pm SEM). Cartesian coordinate group shows transfer. Joint coordinate group shows interference. **(c)** Performance of subject J.W. (a split-brain patient) on the Cartesian and joint transfer of the field from right arm to left. This subject also showed generalization from right arm to left. Adapted from Criscimagna-Hemminger, S. E., Donchin, O., Gazzaniga, M. S., and Shadmehr, R., *J. Neurophysiol.*, 89, 168, 2003. Used with permission.

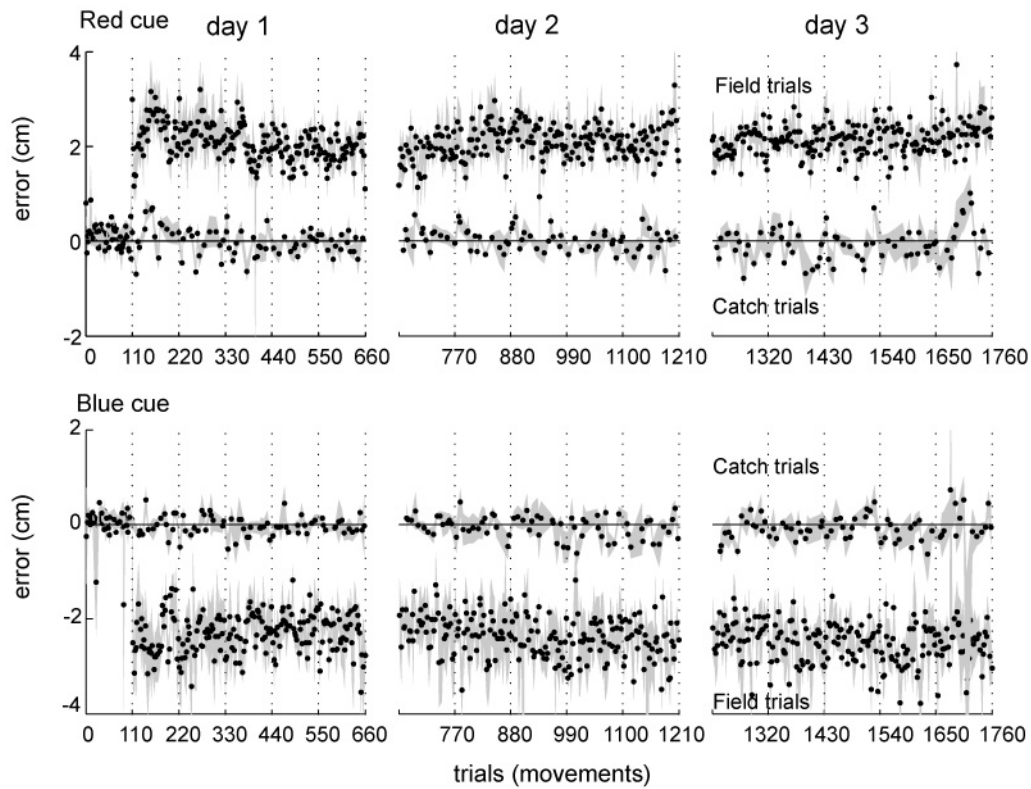


Figure 4. Naïve subjects ($n=3$) were trained to associate color cues (blue or red colored square appearing in the direction of the primary target) to force field B (a clockwise curl field) or $-B$ (a clockwise curl field). Error in each trial was measured as perpendicular displacement from a straight-line trajectory. Data are means \pm SD. Performance remained poor and no field specific after-effects developed in catch trials despite 3 days of training

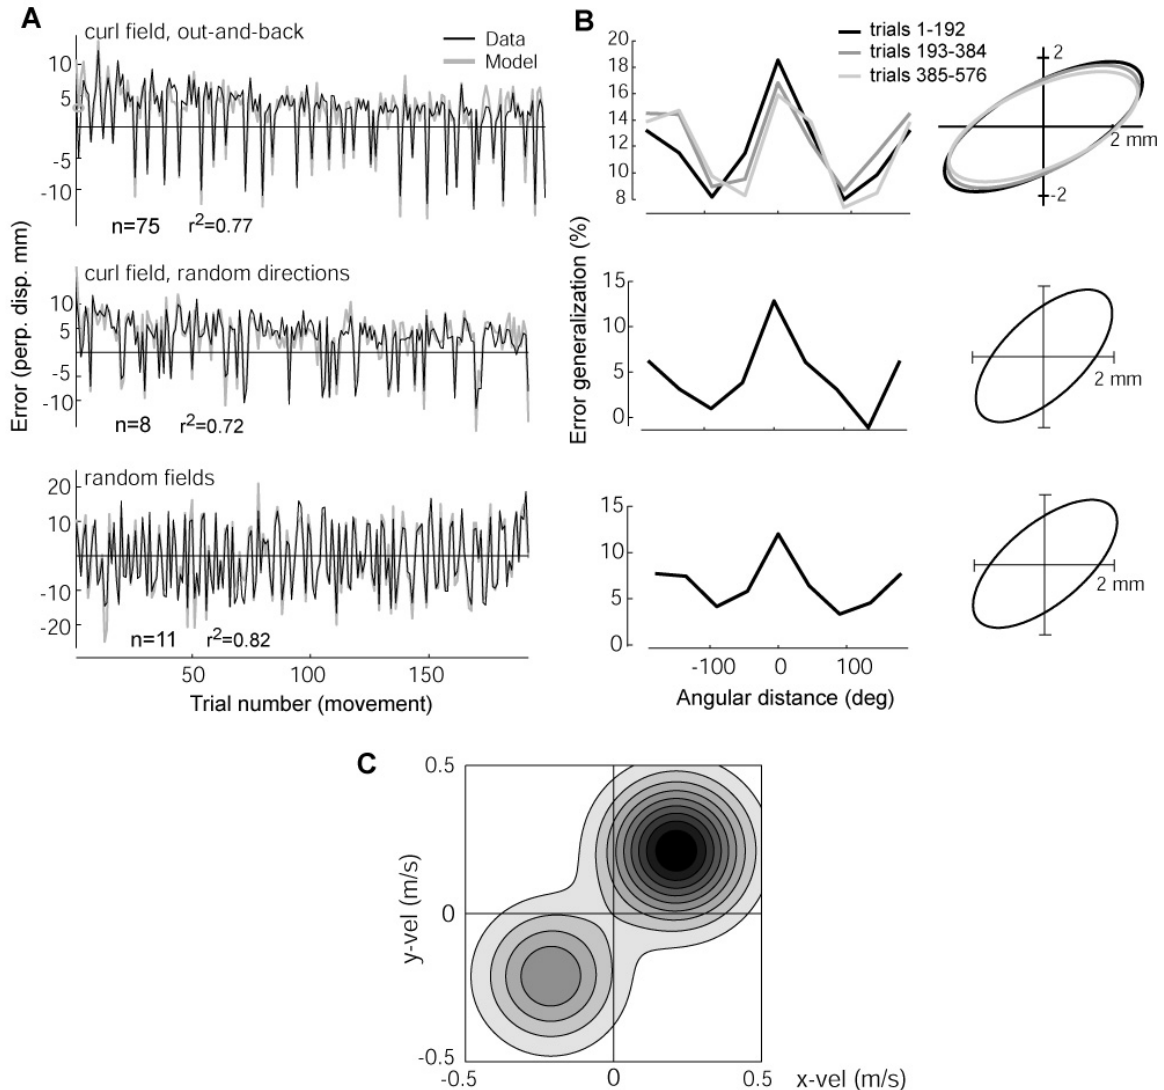


Figure 5. Generalization as a function of direction. **Top row: (a)** Black lines are movement errors during 192 movements (out-and-back pattern) in a standard curl field paradigm to 8 directions of targets ($n=72$ subjects). Sharp negative spikes are catch trials. Gray lines are $y^{(n)}$ as fit to Eq. (5), where the influence of error in any given movement on subsequent movements is estimated. Subjects performed 3x192 movements (3 target sets), but data for only one set is shown. **(b)** The estimated generalization function (b in Eq. 5) and estimated compliance matrix D for each target set. 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. The matrix D is plotted as a transformation of a circle. The estimate changes little in repeated measures and its orientation and shape are consistent with previous estimates of the stiffness of the arm (Mussa-Ivaldi et al., 1985). **Second row:** In this experiment, a group of subjects ($n=8$) 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, a group of subjects ($n=11$)

trained in a force field that randomly changed from movement to movement. Despite no obvious learning of this field, the generalization function is similar to other “learnable” tasks. (c) The shape of the basis function implied by 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. Velocity dependent component parameter values in Eq. (6): $\sigma = 0.15$, $s = 2$. Adapted from Donchin, O., Francis, J. T., and Shadmehr, R. *J. Neurosci.*, 23, 9032, 2003. Used with permission.

Bibliography

- [1] Hollerbach, J. M. and Flash, T., Dynamic interactions between limb segments during planar arm movement, *Biol.Cybern.*, 44, 67, 1982.
- [2] Gribble, P. L. and Ostry, D. J., Compensation for interaction torques during single- and multijoint limb movement, *J.Neurophysiol.*, 82, 2310, 1999.
- [3] Flash, T., The control of hand equilibrium trajectories in multi-joint arm movements, *Biol.Cybern.*, 57, 257, 1987.
- [4] Shadmehr, R. and Mussa-Ivaldi, F. A., Adaptive representation of dynamics during learning of a motor task, *J.Neurosci.*, 14, 3208, 1994.
- [5] Conditt, M. A., Gandolfo, F., and Mussa-Ivaldi, F. A., The motor system does not learn the dynamics of the arm by rote memorization of past experience, *J.Neurophysiol.*, 78, 554, 1997.
- [6] Conditt, M. A. and Mussa-Ivaldi, F. A., Central representation of time during motor learning, *Proc.Natl.Acad.Sci.USA*, 96, 11625, 1999.
- [7] Shadmehr, R. and Brashers-Krug, T., Functional stages in the formation of human long-term motor memory, *J.Neurosci.*, 17, 409, 1997.
- [8] Thoroughman, K. A. and Shadmehr, R., Electromyographic correlates of learning internal models of reaching movements, *J.Neurosci.*, 19, 8573, 1999.
- [9] Flanders, M. and Soechting, J. F., Arm muscle activation for static forces in three-dimensional space, *J.Neurophysiol.*, 64, 1818, 1990.
- [10] Sergio, L. E. and Kalaska, J. F., Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task, *J.Neurophysiol.*, 80, 1577, 1998.
- [11] Sergio, L. E. and Kalaska, J. F., 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, *J.Neurophysiol.*, 78, 1170, 1997.
- [12] Li, C. S. R., Padoa-Schioppa, C., and Bizzi, E., Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field, *Neuron*, 30, 593, 2001.
- [13] Caminiti, R., Johnson, P. B., and Urbano, A., Making arm movements within different parts of space: dynamic aspects in the primate motor cortex, *J.Neurosci.*, 10, 2039, 1990.
- [14] Caminiti, R., Johnson, P. B., Galli, C., Ferraina, S., and 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, *J.Neurosci.*, 11, 1182, 1991.
- [15] Moran, D. W. and Schwartz, A. B., Motor cortical representation of speed and direction during reaching, *J.Neurophys.*, 82, 2676, 1999.
- [16] Kakei, S., Hoffman, D. S., and Strick, P. L., Muscle and movement representations in the primary motor cortex, *Science*, 285, 2136, 1999.
- [17] Shadmehr, R. and Moussavi, Z. M. K., Spatial generalization from learning dynamics of reaching movements, *J.Neurosci.*, 20, 7807, 2000.

- [18] Brashers-Krug, T., Shadmehr, R., and Bizzi, E., Consolidation in human motor memory, *Nature*, 382, 252, 1996.
- [19] Humphrey, D. R., Schmidt, E. M., and Thompson, W. D., Predicting measures of motor performance from multiple cortical spike trains, *Science*, 170, 758, 1970.
- [20] Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E., Neural population coding of movement direction, *Science*, 233, 1416, 1986.
- [21] Amirikian, B. and Georgopoulos, A. P., Directional tuning profiles of motor cortical cells, *Neurosci.Res.*, 36, 73, 2000.
- [22] Tolhurst, D. J. and Thompson, I. D., Organization of neurones preferring similar spatial frequencies in cat striate cortex, *Exp.Brain Res.*, 48, 217, 1982.
- [23] Seung, H. S. and Sompolinsky, H., Simple models for reading neuronal population codes, *Proc.Natl.Acad.Sci.U.S.A*, 90, 10749, 1993.
- [24] Schwartz, A. B., Direct cortical representation of drawing, *Science*, 265, 540, 1994.
- [25] Poggio, T., A theory of how the brain might work, *Cold Spring Harbor Symp.Quant.Biol.*, 55, 899, 1990.
- [26] Pouget, A., Dayan, P., and Zemel, R., Information processing with population codes, *Nat.Rev.Neurosci.*, 1, 125, 2000.
- [27] Pouget, A. and Sejnowski, T. J., Spatial transformations in the parietal cortex using basis functions, *J.Cog.Neurosci.*, 9, 222, 1997.
- [28] Andersen, R. A., Essick, G. K., and Siegel, R. M., Encoding of spatial location by posterior parietal neurons, *Science*, 230, 456, 1985.
- [29] Georgopoulos, A. P., Caminiti, R., and Kalaska, J. F., Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space, *Exp.Brain Res.*, 54, 446, 1984.
- [30] Hwang, E. J., Donchin, O., Smith, M. A., and Shadmehr, R., A Gain-Field Encoding of Limb Position and Velocity in the Internal Model of Arm Dynamics, *PLoS.Biol.*, 1, 209, 2003.
- [31] Bosco, G., Rankin, A., and Poppele, R. E., Representation of passive hindlimb postures in cat spinocerebellar activity, *J.Neurophysiol.*, 76, 715, 1996.
- [32] Helms Tillery, S. I., Soechting, J. F., and Ebner, T. J., Somatosensory cortical activity in relation to arm posture: nonuniform spatial tuning, *J.Neurophysiol.*, 7x, 2423, 1996.
- [33] Coltz, J. D., Johnson, M. T. V., and Ebner, T. J., Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking, *J.Neurosci.*, 19, 1782, 1999.
- [34] Scott, S. H. and Kalaska, J. F., Reaching movements with similar hand paths but different arm orientation: I. Activity of individual cells in motor cortex, *J.Neurophysiol.*, 77, 826, 1997.
- [35] Ajemian, R., Bullock, D., and Grossberg, S., A model of movement coordinates in the motor cortex: posture-dependent changes in the gain and direction of single cell tuning curves, *Cereb.Cortex*, 11, 1124, 2001.

- [36] Steinberg, O., Donchin, O., Gribova, A., Cardoso, d. O., Bergman, H., and Vaadia, E., Neuronal populations in primary motor cortex encode bimanual arm movements, *Eur.J.Neurosci.*, 15, 1371, 2002.
- [37] Kalaska, J. F., Cisek, P., and Crammond, D. J. Effector-independent activity in primate dorsal premotor cortex (PMd) during instructed delay tasks. *Soc.Neurosci.Abs.* 359.2. 2000.
- [38] Criscimagna-Hemminger, S. E., Donchin, O., Gazzaniga, M. S., and Shadmehr, R., Learned dynamics of reaching movements generalize from dominant to nondominant arm, *J.Neurophysiol.*, 89, 168, 2003.
- [39] Gazzaniga, M. S., Cerebral specialization and interhemispheric communication: does the corpus callosum enable the human condition, *Brain*, 123, 1293, 2000.
- [40] Galea, M. P. and Darian-Smith, I., Corticospinal projection patterns following unilateral section of the cervical spinal cord in the newborn and juvenile macaque monkey, *J.Comp Neurology*, 381, 282, 1997.
- [41] Haaland, K. Y. and Harrington, D. L., Hemispheric asymmetry of movement, *Curr.Opin.Neurobiol.*, 6, 796, 1996.
- [42] Gandolfo, F., Mussa-Ivaldi, F. A., and Bizzi, E., Motor learning by field approximation, *Proc.Natl.Acad.Sci.USA*, 93, 3843, 1996.
- [43] Rao, A. K. and Shadmehr, R. Contextual cues facilitate learning of multiple models of arm dynamics. *Soc.Neurosci.Abs.* 27[302.4]. 2001.
- [44] Krouchev, N. I. and Kalaska, J. F., Context-dependent anticipation of different task dynamics: rapid recall of appropriate motor skills using visual cues, *J.Neurophysiol.*, 89, 1165, 2003.
- [45] Wada, Y., Kawabata, Y., Kotosaka, S., Yamamoto, K., Kitazawa, S., and Kawato, M., Acquisition and contextual switching of multiple internal models for different viscous force fields, *Neurosci.Res.*, 46, 319, 2003.
- [46] Thoroughman, K. A. and Shadmehr, R., Learning of action through adaptive combination of motor primitives, *Nature*, 407, 742, 2000.
- [47] Donchin, O. and Shadmehr, R. Linking motor learning to function approximation: Learning in an unlearnable force field. 14, 197-203. 2002. *Advances in Neural Information Processing Systems*.
- [48] Donchin, O., Francis, J. T., and Shadmehr, R., Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control, *J.Neurosci.*, 23, 9032, 2003.
- [49] Turner, R. S. and Anderson, M. E., Pallidal discharge related to the kinematics of reaching movements in two dimensions, *J.Neurophysiol.*, 77, 1051, 1997.
- [50] Johnson, M. T. V. and Ebner, T. J., Processing of multiple kinematic signals in the cerebellum and motor cortices, *Brain Res.Rev.*, 33, 155, 2000.
- [51] Goodbody, S. J. and Wolpert, D. M., Temporal and amplitude generalization in motor learning, *J.Neurophysiol.*, 79, 1825, 1998.
- [52] Krakauer, J. W., Pine, Z. M., and Ghez, C., Visuomotor transformations for reaching are learned in extrinsic coordinates, *Soc.Neurosci.Abs.*, 25, 2177, 1999.

- [53] Muellbacher, W., Ziemann, U., Wissel, J., Dang, N., Kofler, M., Facchini, S., Boroojerdi, B., Poewe, W., and Hallett, M., Early consolidation in human primary motor cortex, *Nature*, 415, 640, 2002.
- [54] Shadmehr, R. and Holcomb, H. H., Neural correlates of motor memory consolidation, *Science*, 277, 821, 1997.
- [55] Shadmehr, R. and Holcomb, H. H., Inhibitory control of motor memories: a PET study, *Exp.Brain Res.*, 126, 235, 1999.
- [56] Nezafat, R., Shadmehr, R., and Holcomb, H. H., Long-term adaptation to dynamics of reaching movements: a PET study, *Exp.Brain Res.*, 140, 66, 2001.
- [57] Bhushan, N. and Shadmehr, R., Computational architecture of human adaptive control during learning of reaching movements in force fields, *Biol.Cybern.*, 81, 39, 1999.
- [58] Wang, T., Dordevic, G. S., and Shadmehr, R., Learning the dynamics of reaching movements results in the modification of arm impedance and long-latency perturbation responses, *Biol.Cybern.*, 85, 437, 2001.
- [59] Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M., The central nervous system stabilizes unstable dynamics by learning optimal impedance, *Nature*, 414, 446, 2001.
- [60] Franklin, D. W., Osu, R., Burdet, E., Kawato, M., and Milner, T. E., Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model, *J.Neurophysiol.*, 90, 3270, 2003.