Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics

Timothy P. Lillicrap[1,2,*] and Stephen H. Scott[2,3,*]
[1]Department of Pharmacology, University of Oxford, Oxford OX1 3QT, UK
[2]Centre for Neuroscience Studies
[3]Department of Biomedical and Molecular Sciences
Queen's University, Kingston, ON K7L 3N6, Canada
*Correspondence: timothy.lillicrap@pharm.ox.ac.uk (T.P.L.), steve.scott@queensu.ca (S.H.S.)
http://dx.doi.org/10.1016/j.neuron.2012.10.041

SUMMARY

Neurons in monkey primary motor cortex (M1) tend to be most active for certain directions of hand movement and joint-torque loads applied to the limb. The origin and function of these biases in preference distribution are unclear but may be key to understanding the causal role of M1 in limb control. We demonstrate that these distributions arise naturally in a network model that commands muscle activity and is optimized to control movements and counter applied forces. In the model, movement and load preference distributions matching those observed empirically are only evident when particular features of the musculoskeletal system were included: limb geometry, intersegmental dynamics, and the force-length/velocity properties of muscle were dominant factors in dictating movement preferences, and the presence of biarticular muscles dictated load preferences. Our results suggest a general principle: neural activity in M1 commands muscle activity and is optimized for the physics of the motor effector.

INTRODUCTION

Two fundamental aspects of motor control which underlie the diverse set of behaviors displayed by primates are (1) the ability to make directed movements and (2) the ability to stabilize a limb against imposed forces, such as those induced by grasped objects (Figures 1A and 1B). In experiments designed to probe the neural mechanisms underlying these tasks, single neurons in monkey primary motor cortex (M1) were found to be most active for (i.e., prefer) certain movement directions and applied forces. These preferred movement directions were initially thought to be distributed uniformly in space (Georgopoulos et al., 1988; Caminiti et al., 1990a). Subsequent studies have revealed that these movement preferences have tendencies to cluster around *particular* directions (Mitsuda and Onorati, 2002; Naselaris et al., 2006) and that these biases are markedly

stronger when the arm is maintained in the horizontal plane (Figure 1C; Scott and Kalaska, 1997; Scott et al., 2001). Similarly, M1 neurons have been shown to exhibit pronounced preference distribution biases for certain forces applied during postural tasks (Figure 1D; see Experimental Procedures for details; Cabel et al., 2001; Herter et al., 2007; Ajemian et al., 2008). The origin and function of the observed preference biases has remained unclear but may be key to understanding the causal role of M1 activity in the control of posture and movement.

It has been suggested that the nonuniform distribution of preferred movement directions reflects directional hyperacuity and develops as a result of biases in spatial experience (Naselaris et al., 2006), e.g., from having more experience reaching away from and toward the body than left or right, perhaps akin to how orientation biases are thought to emerge in primary visual cortex (Blakemore and Cooper, 1970; Li et al., 2008). A number of studies support this idea, demonstrating that the cortical representation of movements is use-dependent, such that more M1 neurons become involved in frequently performed actions (Nudo et al., 1996; Classen et al., 1998).

Alternatively, it has been proposed that the observed nonuniform distributions in preferred movement and torque directions (PMD and PTD, respectively) are dictated by the mechanical and anatomical properties of the limb (Scott and Kalaska, 1997; Scott et al., 2001; Herter et al., 2007). For example, if motor cortical units directly control muscle activity (Evarts, 1968; Bennett and Lemon, 1996; Todorov, 2000; Holdefer and Miller, 2002), then movements which require more muscle activity might require correspondingly higher levels of neural activation. This hypothesis is supported indirectly by the fact that muscles recorded during reaching and loaded-posture show trends in their distribution of preferences similar to those observed for neurons (Kurtzer et al., 2006a, 2006b; Figures 1E and 1F).

This debate on the nature of preference distributions in M1 parallels the debate over what motor cortex encodes or represents (Evarts, 1968; Mussa-Ivaldi, 1988; Kakei et al., 1999). A long history of studies have documented correlations between neural activity in M1 and both ''intrinsic'' (e.g., muscle activity and joint forces; Phillips and Porter, 1964; Evarts, 1968, 1969; Humphrey, 1972; Conrad et al., 1977; Asanuma et al., 1979; Cheney and Fetz, 1980; Kalaska et al., 1989; Holdefer and Miller,
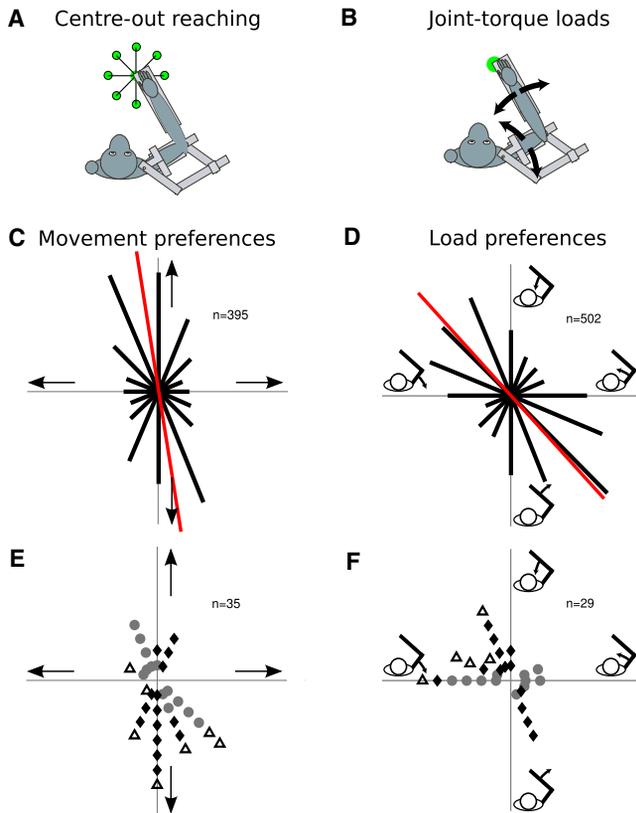
Figure 1. M1 Neurons Tend to Prefer Particular Movements and Loads

(A) In the center-out reaching task a monkey moves its hand from a central target to peripheral targets evenly spaced around a circle (radius = 0.06 m).

(B) In the loaded-posture task the monkey maintains a fixed posture while nine different joint-based loads are applied by the KINARM robot.

(C) Polar histogram of preferred movement directions of M1 neurons ($n = 395$) recorded during the center-out task performed by the contralateral arm; the activity of each neuron was fit by a plane to determine its directional preference and those with significant fits ($p<0.05$) are included here. The length of each line represents the number of neurons which fell into one of 16 equally spaced bins. Neurons tend to prefer movements away from the body and a little to the left, or movement toward the body and a little to the right. Red line denotes a significant bimodal distribution at the given orientation (Bimodal Rayleigh $r = 0.37$, orientation $\theta = 98.4°$, $n = 396$, $p<10^{-3}$ by bootstrap).

(D) Polar histogram of preferred torque directions (plane fit, $p<0.05$) of M1 neurons ($n = 502$) recorded during a postural-load task. Line denotes a significant bimodal distribution. Neurons tend to prefer shoulder flexor torques combined with elbow extensor torque, or else shoulder extensor torques combined with elbow flexor torques (Bimodal Rayleigh $r = 0.31$, orientation $\theta = 132.4°$, $n = 502$, $p<10^{-3}$ by bootstrap).

(E and F) Muscles recorded during reaching (E) and loaded-posture (F) show similar bimodal distributions of preferred directions (Kurtzer et al., 2006a). Grey circles (●) denote monoarticular shoulder muscles; black diamonds (◆), monoarticular elbow muscles; empty triangles (△), biarticular muscles.

2002) and ''extrinsic'' (e.g., hand position and movement in Cartesian space; Kalaska et al., 1989; Georgopoulos et al., 1988; Caminiti et al., 1990a, 1990b; Kakei et al., 2001) variables. Implicit in the hyperacuity hypothesis is that the direction of hand motion or other goal-related features are explicitly encoded by M1. By contrast, the hypothesis that neural preferences have

their origin in the properties of the peripheral biophysics implies that M1 is intimately involved in the development of low-level muscle activity. Under this view, correlations observed with high level variables such as hand direction, velocity, or force are viewed as incidental and are thought to occur because of causal links between muscle activity and limb physics (Mussa-Ivaldi, 1988; Todorov, 2000).

Due to systemic correlations and the interconnected nature of limb physics and neural control (Todorov, 2000; Mussa-Ivaldi, 1988; Kalaska, 2009; Reimer and Hatsopoulos, 2009), it has been difficult to experimentally address these hypotheses. How can we separate the manner in which neural processing in M1 is influenced by factors such as biased spatial experience, limb geometry, intersegmental dynamics, musculoskeletal organization, and muscle mechanics? Optimal control theory, the field of mathematics concerned with finding the best way of acting (Stengel, 1994), is a natural way to formalize and understand the complex interplay inherent in biological motor control and has proven a powerful tool for explaining why animals display particular behaviors (Parker and Smith, 1990; Alexander, 1996; Todorov, 2004; Todorov and Jordan, 2002; Scott, 2012). Here, by applying optimal control theory to a simple network model, we systematically examine how the above-mentioned factors dictate neural processing in simulated neural populations.

Our model is fundamentally different from many of the models which have previously been used to examine M1 function (e.g., Todorov, 2000; Guigon et al., 2007; Shah et al., 2004; Trainin et al., 2007). The network we develop is capable of generating optimal movements online without the need for inputs specifying trajectory kinematics, performing in the presence of noise, and responding to external perturbations. Thus, our model is not simply a phenomenological (i.e., curve fitting) model of preferred direction distributions. Instead the model is mechanistic and functions by generating movements using peripheral feedback—the preference distributions we examine are merely a way to statistically characterize this mechanism and compare it to experimental data.

Analysis of preference distributions in our network model favors the hypothesis that biases are dictated by the biomechanics of the limb. Our results demonstrate that the prominent biases in M1 preference distributions emerge robustly in artificial networks optimized to control movements and counter loads under two crucial assumptions: (1) network units command muscle activity via a simple linear filter and (2) that muscle and neuron activity are kept small. In addition, we use our model to demonstrate how variations in arm position and changes in the anatomical organization of biarticular muscles influence network preference distributions.

RESULTS

Dynamic Network Model

We built a dynamic network model that controlled a model of the primate upper limb (Figure 2). The model was optimized to make reaches and maintain postures under static loads while keeping neural and muscle activities and synaptic weights small (Fagg et al., 2002; Shah et al., 2004; Trainin et al., 2007).

The inclusion of a penalty term which encourages small synaptic weights is not strictly necessary—our results are qualitatively the same without this penalty. However, we found that the generalization capabilities of the network were improved by incorporation of this regularization term (Krogh and Hertz, 1993; Hinton and van Camp, 1993). The network itself was feedforward, but was connected in closed loop with the model of the limb, and consisted of a vector of standard sigmoidal units, $z(t)$, which sent their weighted activity to a vector of six lumped muscle actuators, $u(t)$. The units $z(t)$—which may be thought of as motor cortical neurons—received limb state feedback, $x(t)$, and goal information, $y^*(t)$, passed through a preceding layer of sigmoidal units acting as an input filter. The corresponding muscle activity at time t is given by $u(t) = \sigma_u(W^{out} \cdot z(t))$, where $\sigma_u(\cdot)$ is a smooth ramp function and W^{out} is a matrix of synaptic weights which dictate how activity in $z(t)$ leads to changes in muscle activity. In the simulations presented here W^{out} was random and fixed with elements of the matrix drawn from a Normal distribution (Todorov, 2000; Shah et al., 2004; Trainin et al., 2007). Using a version of backpropagation through time (Rumelhart et al., 1986; Werbos, 1990; Stroeve, 1998) modified for our model, we computed the partial derivatives of the cost function with respect to the adjustable parameters in the model and used a conjugate gradient descent algorithm to find a local minimum (see Experimental Procedures, and see Supplemental Information available online).

Following optimization, we instructed the network to perform the same *center-out reaching* and *loaded-posture* tasks performed by our monkeys, with the limb at a central position in the workspace (Scott et al., 2001; Figure 2B). The activity of the network units, $z(t)$, were "recorded" as it performed each task and analyzed using planar regression to determine each unit's movement and load preferences (i.e., analogous to the analysis of M1 neurons). The network produces relatively straight handpaths with bell-shaped velocity profiles (Figures 2C and 2D). As well, the muscle activity and neural activity profiles qualitatively agree with that reported in the literature (Figures 2E and 2F). Units tended to be broadly tuned to hand movement direction and exhibited a mixture of phasic and tonic responses (Scott and Kalaska, 1997; Scott et al., 2001; Cheney and Fetz, 1980; Georgopoulos et al., 1988). The behavior during posture is not pictured but is comparatively simple, and easy to describe: the network successfully stabilized the hand against loads at the central target and neural and muscle activity were relatively constant during these trials. As observed empirically during loaded-posture, modeled units and muscles were broadly tuned to load combinations and were well fit by a linear regression (Kurtzer et al., 2006a; Herter et al., 2007).

Figure 2. The Model Performing Center-Out Reaching Task

(A) Schematic of the dynamic model. The network computes commands internally in a feedforward fashion, but receives on-line feedback from the plant.
(B) Example reaches made by the network over the trained workspace. Thin arcs delimit the reachable workspace, and the dashed box shows the area of the workspace over which the network was trained. Green circles denote start (no border) and end (black border) targets. Eight example reaches made by the network in the center-out task are highlighted. Centre target placed shoulder and elbow joints at 32.6° and 84.2°, respectively.

(C) Close-up of center-out reaches with direction of reaches labeled.
(D) Velocity of hand ($\dot{x}$ and $\dot{y}$) for the rightward, 0°, reach direction (reach is blue in C).
(E and F) Modeled muscle (E, 6 lumped muscles) and unit (F, 16 random examples out of total 1,000 units) activity during center-out reaching movements in 64 equally spaced directions (eight of the 64 shown above). Horizontal axis is time. Vertical axis is movement direction. Color indicates activation level in arbitrary units (au): red high/blue low. Each lumped muscle and unit have been normalized to their maximum and minimum activation levels.
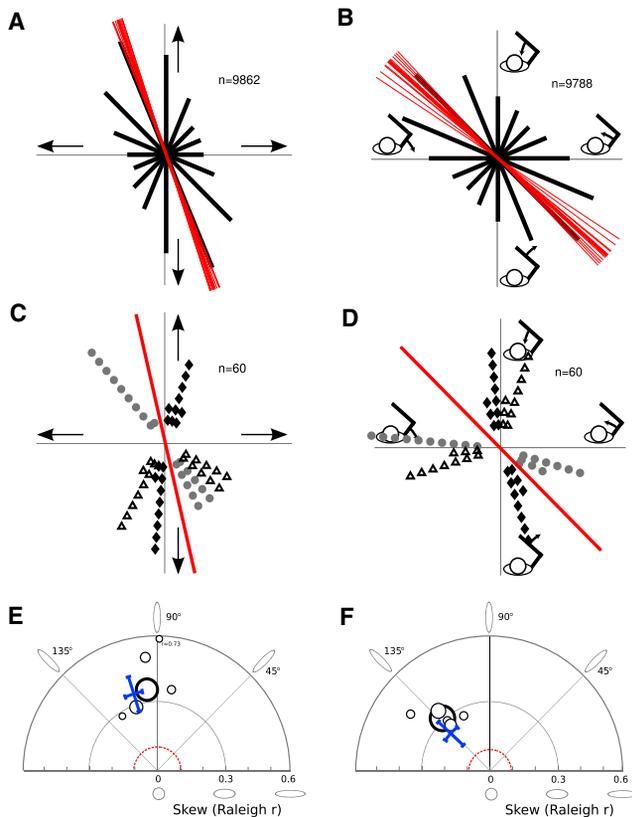
Figure 3. Modeled Unit and Muscle Preference Distributions

(A) Polar histogram of PMDs of simulated units pooled across ten networks (units $n = 10 \times 1000 = 10,000$) recorded during the center-out task. The activity of each unit was fit by a plane to determine its directional preference; only those which had significant fits ($p < 0.05$) are shown here. The length of each line represents the number of units which fell into one of 16 equally spaced bins. As a population, most units prefer hand movements away and a little left, or toward and a little right. Thin red line denotes the orientation of a significant bimodal distribution for each network. Thick red line denotes the average bimodal orientation (Bimodal Rayleigh averaged orientation $\theta = 108.9°$, averaged skew $r = 0.35$, $p < 10^{-3}$ by bootstrap).

(B) Polar histogram of PTDs of units recorded during the loaded-posture task. Units tend to prefer shoulder flexor torques combined with elbow extensor torque or shoulder extensor torques combined with elbow flexor torques ($\theta = 137.7°$, $r = 0.24$, $p < 10^{-3}$); notation similar to (A).

(C and D) Polar histograms of PMDs (C) and PTDs (D) of modeled muscles; 32 equally spaced bins. PMD distribution: $\theta = 102.2°$, $r = 0.46$, $p < 10^{-3}$; PTD distribution: $\theta = 135.4°$, $r = 0.224$, $p < 10^{-3}$ (thin red lines suppressed; thick red lines indicates mean bimodal orientation). Circles denote monoarticular shoulder muscles; diamonds, monoarticular elbow muscles; triangles, biarticular muscles.

(E) Empirical and modeled neuron PMD distributions characterized by primary axis orientation and skew. In this polar plot, the primary axis orientation is represented by the polar coordinate, and the skew by the radial coordinate, of a given point. Each monkey's PMD distribution is plotted as a white circle whose size indicates the number of neurons which make up the distribution; bolded circle indicates the mean distribution (shown in Figure 1; $n = 395$). Dashed red line indicates region into which 99% of 100,000 PMD distributions fell when 395 PMDs are drawn randomly from a uniform distribution. Blue crosshair indicates the PMD distribution found using the model, where the crosshairs delimit the skew and orientation within which 99% of 100,000 random samples of 395 units from the larger collection of optimized units fell.

For each optimized network, the distribution of PMDs during the center-out reaching task was found to be strongly biased for movements forward and left, and backward and right (Figure 3A; Rayleigh test for bimodality, see Supplemental Information; dominant orientation $\theta = 109.9°$, distribution skew $r = 0.30$, $p < 10^{-3}$ by bootstrap). The distribution of PTDs for the loaded-posture task were also significantly bimodally distributed for each network (Figure 3B; $\theta = 137.7°$, $r = 0.24$, $p < 10^{-3}$). The muscle PMD and PTD distributions show similar trends, also matching empirical observations (Figures 3C and 3D). Figures 3E and 3F illustrate that the model exhibited bimodal preference distributions which closely parallel those observed for monkeys.

The results shown in Figure 3 are robust to a variety of manipulations. We varied the size (e.g., 100–1,000 neurons) and structure (e.g., the addition of recurrent connections) of the network, the unit activation function, the muscle activation function $\sigma_u(\cdot)$, the regularization scalars α and β which weighted the importance of keeping neural and muscle activity small, the distribution from which elements of W^{out} were drawn, and the movement duration and the integration timestep (see Variations on Simulation Setup in Supplemental Information). As well, while the model whose behavior is shown in Figures 3A–3F was optimized to make reaches of a predetermined movement duration chosen to coincide with the average reach duration of monkeys, this is not a necessary assumption. If we instead use an instantaneous cost which penalizes both neural/muscle activity and distance-to-target at each time-step, optimization finds solutions which trade off arriving at the target quickly with keeping neural/muscle activity small. In all of these cases, as long as optimization successfully found a solution with kinematic and gross muscle behavior matching empirical data (i.e., relatively straight handpaths, bell-shaped velocity profiles, and engagement of all six lumped muscles), the optimal distribution of PMDs and PTDs exhibited significant bimodal skew with similar orientations (within ± 12°). Importantly, the set of reach target directions and applied loads which were trained on during optimization were roughly uniformly distributed in Cartesian and joint-torque space, respectively. Therefore, in our model, the bimodal distributions reflect optimization of neural activity for the biomechanical properties of the limb without the need to introduce a bias in spatial (or force) experience.

How Do Biomechanics Influence Optimal Preference Distributions?

How should we unravel these results and begin to understand which facets of the model underlie the observed bimodal distributions? The dynamic model allowed us to study a network that performs real-time control, permitting direct comparison to empirically observed behavior. To more easily link the various biomechanical features of the modeled limb to optimal preference distributions we developed a simplified version of the mechanistic model described above (Figure 4; see Experimental Procedures for details).

(F) Follows the same format as (E) except that neuron PTD distributions are plotted and the mean empirical distribution contained $n = 502$ units (all samples for 99% regions used $n = 502$ as well).

A

$$\dot{\mathbf{y}} = G \cdot F \cdot M \cdot (H \bullet \boldsymbol{\sigma}_u(W^{\mathrm{out}} \cdot \mathbf{z}))$$

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} \; \tfrac{1}{2}||\dot{\mathbf{y}} - \dot{\mathbf{y}}_{\mathrm{target}}||^2 + \tfrac{1}{2}\alpha||\mathbf{u}||^2 + \tfrac{1}{2}\beta||\mathbf{z}||^2$$

B

$$\boldsymbol{\tau} = M \cdot (H \bullet \boldsymbol{\sigma}_u(W^{\mathrm{out}} \cdot \mathbf{z}))$$

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} \; \tfrac{1}{2}||\boldsymbol{\tau} - \boldsymbol{\tau}_{\mathrm{target}}||^2 + \tfrac{1}{2}\alpha||\mathbf{u}||^2 + \tfrac{1}{2}\beta||\mathbf{z}||^2$$
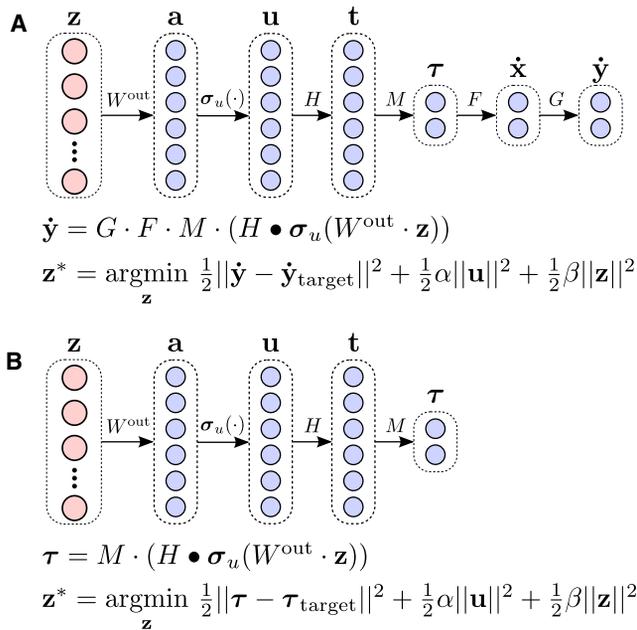
Figure 4. Static Approximation of the Dynamic Model

The static model optimizes initiating hand velocity or steady state joint torques. The network computes output for only a single time step and thus receives no feedback from the periphery.

(A) Schematic of the pathway from unit activity, $\mathbf{z}$, to hand velocity, $\dot{\mathbf{y}}$. Optimal neural activity in our model, $\mathbf{z}^*$, is found by minimizing the error between target and actual hand velocity, $\mathbf{e} = \dot{\mathbf{y}}_{\mathrm{target}} - \dot{\mathbf{y}}$, while keeping unit and muscle activity, $\mathbf{u}$, small. Unit activity causes muscle activity via, $\mathbf{u} = \boldsymbol{\sigma}_u(W^{\mathrm{out}} \cdot \mathbf{z})$, where W^{out} is the matrix of connection strengths between units and muscles and $\boldsymbol{\sigma}_u(\cdot)$ is the standard sigmoid function which keeps muscle activity positive, ensuring that muscles can only "pull." Muscle tension forces, $\mathbf{t}$, are obtained by element-wise multiplication of muscle activity with F-L/V scaling factors appropriate for the movement direction, i.e., $\mathbf{t} = H \cdot \mathbf{u}$. Finally, hand velocity is determined by the linear transformation, $\dot{\mathbf{y}} = G \cdot F \cdot M \cdot \mathbf{t}$, where M is a fixed moment arm matrix and F and G are local linear approximations to limb dynamics and the mapping between joint and hand velocity, respectively.

(B) Schematic of the pathway from unit activity, $\mathbf{z}$, to joint torques, $\boldsymbol{\tau}$. To maintain posture while loads are applied to the limb, equal and opposite forces must be generated by the nervous system. Thus, in the posture task, optimal neural activity $\mathbf{z}^*$ is found by minimizing the error between target and actual joint torques, $\mathbf{e} = \boldsymbol{\tau}_{\mathrm{target}} - \boldsymbol{\tau}$, while keeping unit and muscle activity small.

Static Network Model

The simplified model is temporally and spatially local. The model optimizes only a single time-step for each reach or loaded-posture (i.e., the initiation of the movement in the center-out task, or the steady state performance in the loaded posture task) and thus does not use feedback from the periphery. We optimized this model to generate instantaneous target hand velocities, $\dot{\boldsymbol{y}}_{\mathrm{target}}$, and target joint-torques, $\boldsymbol{\tau}_{\mathrm{target}}$, equally spaced around the unit circle in order to emulate the center-out reaching (Figure 4A) and loaded posture (Figure 4B) tasks, respectively.

We examined a highly simplified model of the limb constrained to the plane—a 2D point-mass model. We then reinitialized, reoptimized, and reanalyzed (see Static Model Analysis in Supplemental Information) the static network model for a series of increasingly realistic abstractions of limb physics, allowing us

to probe the relationship between various features of the limb and optimal preference distributions. Figures 5A and 5B illustrate the progressive change in optimal movement and torque related distributions as a function of limb abstraction. As one might expect, the predicted distributions for the full static model were found to closely parallel those predicted by the dynamic model (shown in Figure 3) for both units and muscles (unit PMD distribution: $\theta = 99.7°$, $r = 0.44$; muscle PMD distribution: $\theta = 103.7°$, $r = 0.67$; unit PTD distribution: $\theta = 137.0°$, $r = 0.23$; muscle PTD distribution: $\theta = 135.9°$, $r = 0.23$).

Key changes in the distribution for movement direction become evident when we added limb geometry, intersegmental dynamics, and F-L/V muscle mechanics to the model. Perhaps unsurprisingly, modeling the limb as a 2D point-mass (model [1]) resulted in a uniform distribution of preferences. Addition of limb geometry (model [2]) creates a large bias in preferences away from the body and slightly to the left, and toward the body and slightly to the right. Inclusion of intersegmental dynamics (model [3]) rotates the distribution counterclockwise, following the directions of movement corresponding to maximum joint-torque (Graham et al., 2003). Finally, inclusion of F-L/V mechanics (model [6]) shifts the distribution clockwise toward the direction of maximum peak joint velocities (Graham et al., 2003). Thus, the orientation of the PMD distribution follows the directions of movement which require the greatest amount of muscle (or torque actuator) activity.

For the torque-related distributions (Figure 5A) there is only one major change evident—the distribution becomes uniform when the biarticular muscles are removed from the model. Removal of F-L/V mechanics had little or no effect due to the fact that the arm is assumed to be at rest and so the force-velocity contribution is the same for every direction of torque production.

The results of this systematic manipulation of limb physics seem to indicate that the optimal PMD distribution is determined by multiple features of limb mechanics: limb geometry, intersegmental dynamics, and F-L/V mechanics all appear to play a part in determining the optimal distribution of PMDs observed empirically. The best fit between model and empirical observations is seen when all of these features are included in the biomechanics. On the other hand, the PTD distribution is determined entirely by the presence/absence of biarticular muscles.

The static and dynamic model variants mutually reinforce these conclusions. The static model is simpler, demonstrating that the central results are not dependent on various design choices made for the dynamic model. For example, in the static model there is no feedback and neuron activity is optimized in a nonparametric fashion (i.e., the neurons are able to take on any real value) on a trial-to-trial basis. Since the results hold across models, this helps to mitigate concern that the choice of non-linearity for the dynamic model is key to observing our results. The results from the dynamic model demonstrate that peripheral feedback and full nonlinear biomechanics do not substantially alter the results observed in the static linearized case. As well, a parametric model such as our dynamic variant is required to explore the question of biased experience (as discussed below).
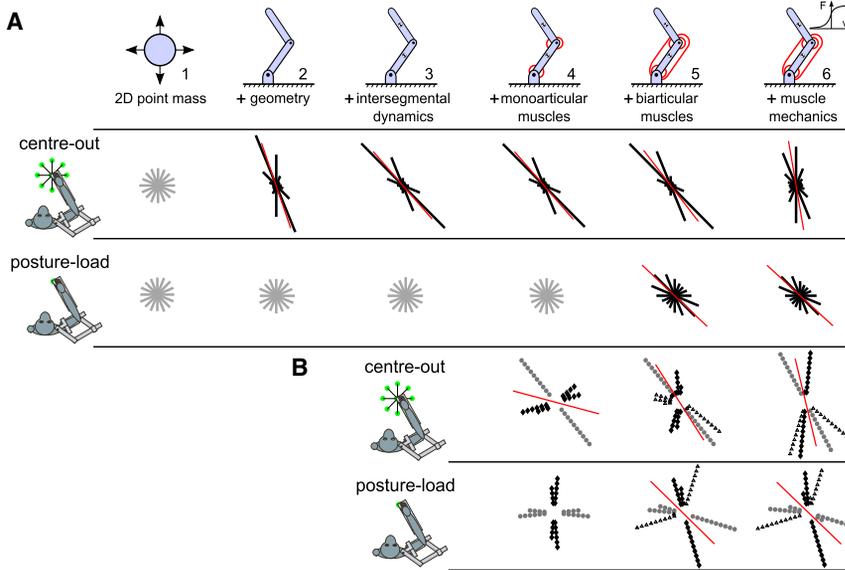
Figure 5. In Simulations, Unit and Muscle Preference Distributions Reflect the Biomechanics of the Limb

(A) Mimicking the center-out reach and loaded-posture tasks, the static model was optimized to produce velocities and torques evenly spaced around the unit circle. Rows one and two show unit PMD and PTD distributions reflecting optimal solutions for a series of abstractions of the limb. Distributions represent averages across ten instantiations of the model, each with 1,000 units, and each with the connectivity matrix, W^{out}, initialized randomly. Red lines indicate a significant bimodal distribution with the dominant axis given by the orientation of the line. Polar histograms are colored black to denote a bimodal distribution—those which failed to pass significance are shown in gray. Note that the histograms show the real data from our static model; each of these is composed of PDs from 10,000 units and thus will tend to be a good approximation to the underlying distribution. Distribution statistics [1] 2-D point-mass (PMD distribution: $\theta = $ n/a°, $r = 0.02$, $p > 0.5$; PTD distribution: $\theta = $ n/a°, $r = 0.002$, $p > 0.05$), [2] addition of geometry (PMD distribution: $\theta = 108.8°$, $r = 0.61$, $p < 10^{-3}$; PTD distribution: $\theta = $ n/a°, $r = 0.03$, $p > 0.05$), [3] addition of intersegmental dynamics (PMD distribution: $\theta = 131.4°$, $r = 0.67$, $p < 10^{-3}$; PTD distribution: $\theta = $ n/a°, $r = 0.02$, $p > 0.05$), [4] addition of monoarticular muscles (PMD distribution: $\theta = 131.4°$, $r = 0.67$, $p < 10^{-3}$; PTD distribution: $\theta = $ n/a°, $r = 0.03$, $p > 0.05$), [5] addition of biarticular muscles (PMD distribution: $\theta = 127.9°$, $r = 0.65$, $p < 10^{-3}$; PTD distribution: $\theta = 136.4°$, $r = 0.22$, $p < 10^{-3}$), [6] addition of F-L/V properties (PMD distribution: $\theta = 99.7°$, $r = 0.44$, $p < 10^{-3}$; PTD distribution: $\theta = 137.0°$, $r = 0.22$, $p < 10^{-3}$), (B), Rows three and four show static model optimized muscle PMD and PTD distributions for those limb models which included muscles. Distribution statistics: [4] monoarticular muscles (PMD distribution: $\theta = 165.0°$, $r = 0.27$, $p < 10^{-3}$; PTD distribution: $\theta = $ n/a°, $r = 0.001$, $p > 0.5$), [5] additions of biarticular muscles (PMD distribution: $\theta = 122.9°$, $r = 0.58$, $p < 10^{-3}$; PTD distribution: $\theta = 135.9°$, $r = 0.219$, $p < 10^{-3}$), [6] addition of F-L/V properties (PMD distribution: $\theta = 103.8°$, $r = 0.71$, $p < 10^{-3}$; PTD distribution: $\theta = 135.9°$, $r = 0.23$, $p < 10^{-3}$). Grey circles (●) denote monoarticular shoulder muscles; black diamonds (◆), monoarticular elbow muscles; empty triangles (△), biarticular muscles.

Optimal Preference Distributions Change as a Function of Limb Posture and Musculoskeletal Organization

Beyond determining the causes of the observed bimodal preference distributions, our model also makes concrete predictions about how PMD and PTD distributions ought to change as a function of limb posture and anatomy. Figure 6 shows two such predictions. If center-out reaching is performed in the right half of the workspace, with the shoulder and elbow joints placed at −5.7° and 74.5°, respectively (Figure 2A shows 0° orientation for both joints), then the optimal PMD distribution is found to rotate substantially in the clockwise direction (Figure 6A; $\theta = 54.3°$, $r = 0.47$). This result agrees, at least qualitatively, with previous empirical work which has shown that M1 neural populations systematically rotate their preferred directions in different parts of the workspace (Caminiti et al., 1990a, 1990b; Sergio and Kalaska, 2003) and in different wrist orientations (Kakei et al., 1999).

Figure 6B shows how the anatomical organization of the musculoskeletal system influences the distribution of PTDs in the loaded-posture task. If the biarticular muscles were reattached so that each one had a flexion/extension or extension/flexion action rather than flexion/flexion and extension/extension, then the optimal PTD distribution would be mirrored into the first and third quadrants ($\theta = 44.3°$, $r = 0.22$). While this prediction is not easily testable, it clearly demonstrates the role of musculoskeletal structure in determining observed neural activity during a task.

How Does Biased Spatial Experience Affect Preference Distributions in Our Model?

In our static model, since the optimal neural activity for a given target velocity or torque is determined independently of other trials, the statistics of movements and loads have, by definition, no effect on the optimized solutions. This is also true of any other model, e.g., those of Guigon et al. (2007), and Todorov (2000), where the neural activity on one trial shares no underlying parameters with other trials. As such, these models cannot be used to evaluate how biases in spatial experience might affect preference distributions. This is not the case in our dynamic model where parameters are shared across tasks and movements (see Experimental Procedures). Thus, at least in principle, spatial biases may alter the optimal solutions found by our model and thus influence preference distributions.

To examine the affect of biased experience in our network, we trained our dynamic model to move the 2D point-mass abstraction of the limb to targets distributed unevenly in space. We used the 2D point-mass because we know that a uniform distribution of target directions leads to a uniform distribution of PMDs (see the shaded area #1 in Figure 6A). It follows that if we see a change from uniform in the PMD distribution, we know that this effect results from training on a biased set of reaching directions. In this test, the network was trained to reach from a central location to targets placed around a circle, but with many more targets presented forward and backward than left and right (Figures 7A and 7B; see Biased Experience in Supplemental Information).
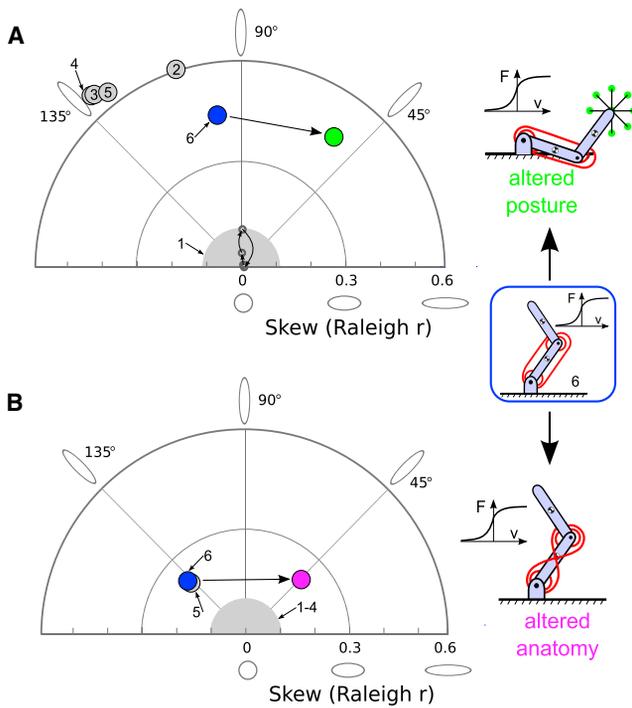
Figure 6. PMD and PTD Distributions for Various Limb Models and Predicted Changes as a Function of Posture and Anatomy

(A) PMD distribution for full limb model plotted as blue circle, whereas gray circles denote distributions for simplified limb models (see Figure 5 legend for specific numbers). Shaded gray half circle indicates location of PMDs distributions associated with the two-dimensional point mass. Green circle indicates the PMD distribution found using the full limb model but with the hand positioned far to the right in the workspace (mean $\theta = 54.3°$, $r = 0.47$, $p<10^{-3}$). Small dark gray circles connected with arrows illustrate the bimodal PMD distribution parameters for the four distributions illustrated in Figure 7C.

(B) Follows the same format as panel a except that PTD distributions are plotted and the mean empirical distribution contained $n = 502$ units. Pink circle is the predicted orientation/skew of the PTD distribution if the moment arms of the biarticular muscles are altered so that each muscle simultaneously flexes one joint and extends the other, rather than having flexion/flexion and extension/extension actions (mean $\theta = 44.3°$, $r = 0.22$, $p<10^{-3}$).

The progress of training is shown in Figure 7C. Behavior during learning was quite variable, in part because gradient updates depended on which targets were drawn on a given optimization run. However, when aligned by cost (i.e., average performance), clear trends emerged. The networks tended to learn to move more quickly to targets forward and backward—the directions where targets appeared most often. As a consequence, midterm evaluations of PMD distributions show significant bimodal trends. This is because the network moves the point-mass further in some directions than others and thus requires more "muscle" activation for these trials. This trend persists even into later learning when kinematics in all directions look relatively similar, since there are still differences at the level of neural/muscular activity. However, the optimal PMD distributions found when optimization terminated were indistinguishable from uniform ($\theta = n/a°$, $r = 0.01$). Thus, even though there were patterned deviations from uniformity during optimization (the

small gray circles connected with arrows in Figure 6A illustrate the bimodal PMD distribution parameters for the 4 distributions shown in Figure 7C), the final distributions exhibit almost no bimodal trend. Even though our model shares parameters between reaches, the optima show little trace of the biased movement statistics in the distributions. This experiment suggests a complex interplay between spatial experience and neural activity, an exploration of which is beyond the scope of this work. The basic conclusion is clear though—for well practiced/optimized movements, our model predicts that spatial biases are likely to have little effect on the PMD distributions.

DISCUSSION

Biological motor control reflects a complex interplay between brain, behavior, and biomechanics that makes it difficult to experimentally examine the causal roles of the constituent factors (Kalaska, 2009; Reimer and Hatsopoulos, 2009). The application of optimization theory has been successfully applied to understand many biological phenomena, including visual cortex receptive fields (Olshausen and Field, 1996; Karklin and Lewicki, 2009), motor behaviors (Alexander, 1996; Todorov and Jordan, 2002; Todorov, 2004), muscle activations (Kurtzer et al., 2006a; Fagg et al., 2002), and some features of motor cortical neuron activity (e.g., Todorov, 2000; Trainin et al., 2007). We used this approach to examine the characteristics of a simple network model trained to control mechanical plants with various abstractions of the primate musculoskeletal system. Several key issues were demonstrated. First, our model illustrates that the patterns of activity of its units tended to display preferences for certain movement and torque directions and that such biases depend critically on specific features of the musculoskeletal system. Second, these distributions closely parallel those observed in M1 neurons of nonhuman primates performing similar tasks. Third, biased experience had little effect on optimized unit preferences in our network, though it does have some effect when movements are still poorly practiced (i.e., when behavioral performance is unrefined). Thus, our model motivates a shift back toward the view espoused by early work on primary motor cortex: that M1 is intimately involved in the generation of low level muscle commands (Phillips and Porter, 1964; Evarts, 1968, 1969; Humphrey, 1972; Conrad et al., 1977; Asanuma et al., 1979; Cheney and Fetz, 1980).

By varying the level of abstraction of the limb physics controlled by our network model, we were able to tease apart which factors dictate preference distributions. In the center-out reaching task we found that limb geometry, intersegmental dynamics, and muscle F-L/V mechanics all play a role in determining the optimal distribution of PMDs. For the loaded-posture task we found that the biarticular muscles are the dominant factor in shaping the optimal distribution of PTDs.

Why is the PMD distribution dominated by limb geometry/dynamics and relatively insensitive to the switch from mono- to biarticular muscles, while the PTD distribution exhibits the opposite trend? In the postural task, there is by definition little or no movement at the optimal solution. Since loads are applied at the joints, any contributions from geometry or intersegmental
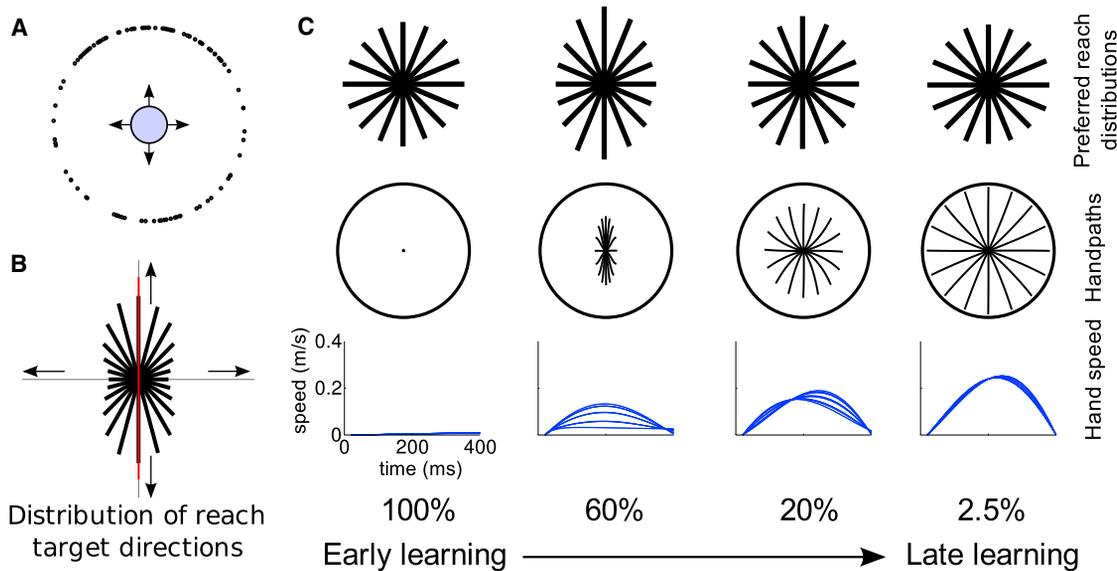
Figure 7. The Effect of Biased Experience on Preference Distributions
(A) The dynamic model is trained to move the 2D point-mass from a central location to targets distributed unevenly around a circle. Many more targets were presented forward and backward than left or right; green circles show an example training batch of 100 targets.
(B) The bimodal Vonmises distribution from which target directions were drawn (Vonmises parameters were $\mu = \pi/2$, $\kappa = 1/2$; see Supplemental Information).
(C) Reaching behavior on 16 equally spaced targets was used to asses unit PMDs at 4 times during learning (aligned by percent remaining cost). Top row: as optimization progressed, the average PMD distribution became significantly bimodal (red lines indicates significant skew orientations), and then returned to a uniform distribution during late stages of learning ($\theta = $ n/a$°$, $r = 0.008 \rightarrow \theta = 87.8°$, $r = 0.04 \rightarrow \theta = 8.4°$, $r = 0.009 \rightarrow \theta = $ n/a$°$, $r = 0.01$). Middle/bottom rows: mean handpaths and speeds for 16 targets equally spaced about a circle. Networks initially learn to move in the direction of targets which are more prevalent during training.

terms are virtually nonexistent. Thus, the postural distribution is unaffected by limb dynamics. In the reaching task, the strong contributions from geometry and intersegmental terms simply limit the visible influence of muscle configuration on the PMD distribution even though, when present, the biarticulars are substantially involved in generating movement torques (see Figure 2E). This second observation warns that there may be significant information present in the M1 neural population which cannot be read by merely characterizing neuron PDs and their distributions.

In most cases, the directional preferences of neurons and muscles were similar for the various models tested. This is not surprising due to the fact that muscle activation for reaching or postural control depends on limb properties, and in turn, neural activity controlled muscle activity. However, PMD and PTD distributions of muscles and neurons are by no means mirror images of each other. For example, the addition of biarticular muscles resulted in substantive changes in the PMD distribution of muscles that were not observed for the network units (Figure 5). Further, although the orientation of the distributions is relatively similar, the muscle distributions show clear clustering of preferences for each muscle group, whereas the distributions for neurons are smooth and continuous.

Some of our results are prefigured by the simulations described by Scott and Kalaska (1997). They constructed populations of units which explicitly encoded either hand direction, angular direction of shoulder/elbow, or required torque at the shoulder/elbow. Each model unit discharged according to the dot product of its random receptive field with a given movement variable. Populations which encoded intrinsic variables (i.e., joint angles and torques) showed bimodal distributions similar to those observed empirically, while those which encoded extrinsic variables (i.e., hand-direction) did not. In contrast, units in our model do not "encode" anything per se, they merely discharge with activities optimized for the current task goal. In our model, even if feedback is given in terms of hand position/ direction, neurons will exhibit the same bimodal distributions because optimized behavior must still account for the biophysics of the limb (see Variations on Simulation Setup in the Supplemental Information). Though Scott and Kalaska (1997) saw qualitatively similar trends in the population tuning of model units neurons, these trends arise for fundamentally different reasons.

Several other models have studied M1 function by including biomechanics and solving the associated optimal control problems (e.g., Guigon et al., 2007; Trainin et al., 2007). The results of these models are firmly consistent with the idea that limb biomechanics plays a role in shaping M1 neural activity. Our work moves beyond this fundamental notion, offering a paradigm for teasing apart the specific contributions from different parts of the periphery. Additionally, the fact that our model is parametrized by synaptic weights which are shared across solutions, rather than non-parametric as in the case of Guigon et al. (2007) and Trainin et al. (2007) has allowed us to begin to examine the contributions of biased experience in determining M1 neural activity.

How Should We Understand and Generalize These Results?

Our static model clarifies precisely what determines the optimal preference distributions in a general sense: the optimal unit activity is a function of all the components of the controlled effector between network units, $z(t)$, and the desired output, whether the desired output is a force, position, or velocity. For example, if s_{target} is a downstream target state (e.g., velocity or torque), then optimal neural activity might be determined by $z^* = A^+ \cdot s_{target}$, where A is a matrix (playing the role of the controlled effector/plant dynamics) which maps neural activity to states, s, and A^+ is the pseudoinverse of A. Solving for z^* finds the smallest z (in terms of the standard vector norm) which minimizes the squared difference, $1/2\|A \cdot z - s_{target}\|^2$. This approach works nicely if the motor effector downstream of z can be effectively approximated by a linear mapping such as A. However, this cannot always be done effectively (e.g., the fact that muscle can only ''pull'' is an important fact which does not admit a simple linear approximation), and so it is often important to keep the general form.

What does this mean in straightforward terms? How can we explain the fact that both M1 neurons and the units in our model develop prominent biases in their preferred directions? To begin, we note simply that some movements require more force (Mussa-Ivaldi et al., 1985; Graham and Scott, 2003) and thus more muscle activity than others. Under the hypothesis that M1 activity leads directly to muscle activation via a relatively simple linear filter, these movements require correspondingly more neural activity. While there are theoretically an infinite number of ways that this neural activity might be chosen, with the simple assumption that redundancy is resolved by penalizing the squared magnitude of z, the result is that more neurons have their largest response in the directions of movement which require more muscle activity.

The view of M1 presented above is admittedly simplistic. While convenient for the purposes of exposition, it must be acknowledged that there are many facets of M1 activity which cannot be straightforwardly accounted for by this paradigm. For example, many M1 neurons are known to correlate more strongly with features of movement kinematics (Kakei et al., 1999) than force or EMG, or exhibit complex or context-dependent responses which cannot be easily accounted for by their connection to downstream muscles (Fetz and Cheney, 1987; Bennett and Lemon, 1996; Churchland and Shenoy, 2007). These facts serve as a compelling reminder that we are a long way from a synoptic view of M1 function. So, how should we begin to think about the multitude of correlations observed in M1? Our approach takes a fundamentally neutral stance on the debates surrounding correlations in M1. Though many of the neurons in our model exhibit relatively simple ''muscle-like'' responses, we do not explicitly require them to do so, and we expect that, as the complexity of the modeled spinal cord, physics, and task increase, the predicted optimal activity will show a corresponding increase in complexity.

Implications for Spinal Cord

The vertebrate spinal cord supports a range of sophisticated computational processes to control body and limb movements.

In our model, $z(t)$ is connected directly to motoneurons via a simple linear filter, W^{out}, and thus ignores any unique computational processes provided by the spinal cord. In monkeys, some pyramidal tract neurons in M1 project directly to motor neurons (Bennett and Lemon, 1996; Cheney and Fetz, 1980; Rathelot and Strick, 2009) but many more project indirectly via interneurons in the spinal cord (Porter and Lemon, 1993). It would be more accurate to include a mapping between M1 units and spinal interneurons and some rudimentary feedback process at the spinal level. Implementation of these features was beyond the scope of the present study. Nonetheless, the similarity in the distributions between modeled units and M1 neurons suggests that the influence of spinal computations on the observed distributions of PMDs and PTDs is relatively modest for the voluntary behaviors examined in the present study.

How Does Spatial Experience Affect Preference Distributions?

Our work also addressed the hypothesis that the bimodal distribution of PMDs observed in M1 reflects directional hyperacuity and is caused by biases in experience (Naselaris et al., 2006). Although some work has examined the statistics of primate limb movement experience (Graziano et al., 2004; Howard et al., 2009), these studies do not quantitatively address whether fore/aft movements are more common than side-to-side movements. However, we have shown that bimodal preference distributions emerge in a straightforward mechanistic model without the introduction of biases in spatial experience. We also demonstrated that biases in spatial training do not cause any significant nonuniformity in the optimal distribution of PMDs for practiced movements (Figure 7). Taken together, these observations motivate the conclusion that the bimodal preference distributions seen in M1 are predominantly the product of the structure of the downstream biomechanics. This is not to say that the statistics of training does not have effects on motor processing. There are almost certainly combinations of movements of the limb or body that are not extensively experienced, and our model points to a subtle interplay between training, behavioral performance, and preference distributions.

Contrast with Neural Coding in Primary Visual Cortex

At first glance, it appears our results are at odds with computational models of primary visual cortex (V1), which demonstrate that the tuning properties of units are determined by: a cost function (e.g., find a coding which minimizes information loss while maximizing sparseness), the upstream sensory processing (i.e., filtering done by the retina and LGN) and the statistics of natural images (Blakemore and Cooper, 1970; Li et al., 2008; Olshausen and Field, 1996; Karklin and Lewicki, 2009). These models have tended to emphasize the importance of the statistics of natural scenes in determining the tuning properties of neurons. Our work with artificial neural networks for control places the emphasis elsewhere. We found that the biomechanics of the downstream controlled system, rather than the statistics of environmental experience, appears to be the dominant factor determining unit direction preferences for our planar movement task. In general, both factors—the statistics of images (or

behavior) and the properties of the interceding sensory or motor plant—influence neural processing to varying degrees. The contrasts between these two systems may hinge on this fact: in our model of the motor system the number of neurons far exceeds the number of degrees of freedom of the effector and thus we face an overdetermined problem where redundancy must be resolved, whereas models of visual processing must deal with data compression problems since the visual world presents much more data than it is possible to store in visual cortex synapses.

EXPERIMENTAL PROCEDURES

Empirical Results

While the focus of this study is on modeled results, we also present neural data which augments findings from previous studies (Scott et al., 2001; Herter et al., 2007) by the inclusion of additional data of the same form. The tasks, neural recording procedures, and analyses are identical to those previously published. Briefly, we used the KINARM robotic system to monitor planar limb movements involving flexion and extension movements of the shoulder and elbow joints and apply joint-based loads (BKIN Technologies, Kingston, Canada). An augmented reality system using a semitransparent mirror coplanar with the limb allowed presentation of spatial targets in the workspace while permitting the monkeys to view their entire limb. Four monkeys (*Macaca mulatta*) were trained to perform *center-out reaching* and *loaded-posture* tasks. All experiments were approved by the Queen's University Research Ethics Board. In the reaching task, the monkey was required to move its hand from a central target to either 8 or 16 peripheral targets located uniformly around a circle with radius 6 cm (Figure 1A). In the posture task, the monkey was required to maintain its hand at the central target while 9 static torque load combinations were applied to the shoulder and elbow joints (flexor, extensor and null at each joint; Figure 1B). See Figure 1 and Supplemental Information for further details.

Dynamic Model Specification

We constructed a simple network model which was optimized to make reaches and maintain postures under static loads while keeping neural and muscular activities small. The network controlled a 2-degree of freedom (shoulder and elbow) arm model which was constrained to move in a plane and included arm geometry (Cheng and Scott, 2000; Trainin et al., 2007), intersegmental dynamics (Cheng and Scott, 2000), mono- and biarticular muscle groups with fixed moment arms (Graham and Scott, 2003; Trainin et al., 2007), and force-length and force-velocity curves (Brown et al., 1999; see Musculoskeletal Model Specification in Supplemental Information). The four-dimensional arm state at time t, $\boldsymbol{x}(t)$, is a column vector of joint angles and velocities and evolved according to, $\Delta\boldsymbol{x}(t+1) = \boldsymbol{f}(\boldsymbol{x}(t), \boldsymbol{\tau}(t))$, where $\boldsymbol{f}(\cdot, \cdot)$ are the system dynamics and $\boldsymbol{\tau}(t)$ was a 2 dimensional column vector containing the torque applied at the shoulder and elbow joints. The current state of the hand is given by $\boldsymbol{y}(t) = \boldsymbol{g}(\boldsymbol{x}(t))$, where $\boldsymbol{g}(\cdot)$ is the mapping from arm state to hand state in Cartesian coordinates. The arm dynamics were integrated forward with simple Euler integration and the time step, Δt, was set to 20 ms.

The network was feedforward, and consisted of two layers of standard sigmoidal units (i.e., $\sigma(x) = 1/(1 + e^{-x})$). The output layer of neurons, $\boldsymbol{z}(t)$, sent their weighted activity to a vector of 6 lumped muscle actuators, $\boldsymbol{u}(t)$. The corresponding muscle activity at time t is given by, $\boldsymbol{u}(t) = \boldsymbol{\sigma_u}(W^{\text{out}} \cdot \boldsymbol{z}(t))$, where $\boldsymbol{\sigma_u}(\cdot)$ is smooth ramp function, and W^{out} is a matrix of synaptic weights which dictate how activity in $\boldsymbol{z}(t)$ leads to changes in muscle activity. For muscle activation in our dynamic model, $\boldsymbol{\sigma_u}(\cdot)$, we used a smooth version of a ramp function which is 0 for values less than or equal to zero and linear for values greater than zero (see Muscle Activation in Supplemental Information).

In simulations, W^{out} was random and fixed, with elements of the matrix drawn from a Normal distribution (Todorov, 2000; Shah et al., 2004; Trainin et al., 2007; we examined several variations; see Discussion and Variations on Simulation Setup in Supplemental Information). The two joint torques

were given by $\boldsymbol{\tau}(t) = M \cdot \boldsymbol{h}(\boldsymbol{x}(t), \boldsymbol{u}(t))$, where M is the matrix of muscle moment arms, and $\boldsymbol{h}(\cdot, \cdot)$ is the force-length/velocity (F-L/V) function which computes the tension force produced by each muscle given its activity, $\boldsymbol{u}(t)$, and the length and velocity of the muscle (computed from the joint angles and velocities, $\boldsymbol{x}(t)$).

The units $\boldsymbol{z}(t)$ received a version of state feedback, $\boldsymbol{x}(t)$, goal information, $\boldsymbol{y}^*(t)$, and load context information, $\boldsymbol{c}(t)$, filtered by the first layer of neurons, $\boldsymbol{v}(t)$ (see Sensory Feedback Filtering in Supplemental Information). The goal information was specified in Cartesian coordinates and the load context information was described in terms of the loads applied at the shoulder and elbow joints. The network was optimized to perform both reaching and loaded postures over a central portion of the workspace (Figure 2A). For an individual trial, i, of reaching or posture, the network is optimized to minimize the cost function,

$$J_i = \frac{1}{2} \sum_{t=0}^{t=t_f} \ell_i(\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{z}(t), \boldsymbol{y}_i^*(t), t) \qquad \text{(Equation 1)}$$

where, $\ell_i(\cdot, \cdot, \cdot, \cdot, \cdot)$ is the instantaneous cost, $\boldsymbol{y}_i^*(t)$ is the desired hand state for trial i, and t_f is the final time. This is the general form of the total cost for either a posture or reach trial; it sums the cost from the limb state ($\boldsymbol{x}(t)$) and muscle and neural activities ($\boldsymbol{u}(t)$ and $\boldsymbol{z}(t)$, respectively) at each time step, t. During posture and at the beginning and end of reaches the instantaneous cost was

$$\ell_i(\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{z}(t), \boldsymbol{y}_i^*(t), t) = \|\boldsymbol{g}(\boldsymbol{x}(t)) - \boldsymbol{y}_i^*(t)\|^2 + \alpha\|\boldsymbol{u}(t)\|^2 + \beta\|\boldsymbol{z}(t)\|^2$$
$$\text{(Equation 2)}$$

where $\|\cdot\|$ is the standard vector norm and α and β are constants specifying the importance of keeping neural and muscle activity small. The kinematic error term, $\|\boldsymbol{g}(\boldsymbol{x}(t)) - \boldsymbol{y}_i^*(t)\|^2$, penalized the deviation of the hand from the desired state. The instantaneous cost was similar for reaching except that during the reach the kinematic error term was dropped from all but the final time-step, t_f, and a term which penalized deviations from a straight hand path was introduced in its place (i.e., no explicit reference trajectory was provided to the network as input, but handpaths were encouraged to be roughly straight). Equations 1 and 2 mean that in the reaching task, the network is required to minimize the distance between the hand and target at the final time and move along a roughly straight path to get there; in the posture task, these equations require that the network minimize the distance between hand and target at every time step; in both tasks, the network must keep muscle and neural activities small at all times.

Each network had 1,000 units in $\boldsymbol{z}(t)$, and the synaptic weights of each network, W^{out}, were initialized randomly (from a Normal distribution with a mean of 0 and variance 0.001) prior to the optimization. Network activity was computed as

$$\boldsymbol{z}(t) = \boldsymbol{\sigma_z}\big(W^{\text{filt}} \cdot \boldsymbol{\sigma_v}\big(W^{\text{in}} \cdot [\boldsymbol{x}(t); \boldsymbol{c}(t); \boldsymbol{y}^*(t)] + \boldsymbol{b_v}\big) + \boldsymbol{b_z}\big) \qquad \text{(Equation 3)}$$

where $\boldsymbol{\sigma_z}(\cdot)$ and $\boldsymbol{\sigma_v}(\cdot)$ are the M1 and filter activation functions which were standard sigmoids, W^{filt}, $\boldsymbol{b_z}$ and W^{in}, $\boldsymbol{b_v}$, are the weight matrices and bias vectors for the M1 and filters, respectively, and ; is the vertical concatenation operation. In most of our simulations, W^{out} is fixed and the aggregate parameter vector, $\boldsymbol{w} = [\text{vec}(W^{\text{filt}}); \boldsymbol{b_z}; \text{vec}(W^{\text{in}}); \boldsymbol{b_v}]$, is optimized (here, $\text{vec}(\cdot)$ returns a vector version of the matrix given as input). In the simulations where $\boldsymbol{w} = [\text{vec}(W^{\text{out}}); \text{vec}(W^{\text{filt}}); \boldsymbol{b_z}; \text{vec}(W^{\text{in}}); \boldsymbol{b_v}]$ is also optimized then. In either case, we adjoined to the total cost function with the regularization term, $1/2 \, \gamma \|\boldsymbol{w}\|^2$ (γ was a scalar and was set to 10^{-5} in our simulations). This term is often referred to as a weight decay term and is a principled choice for regularization, tending to confer good generalization properties; note that in the static model there are no weight decay terms since there are no upstream parameters—neural activity is optimized directly. $\boldsymbol{w}$ was initialized randomly using a Normal distribution centered on 0 with variance 0.001. We computed the derivatives of the total cost function with respect to the aggregate parameter vector, $\partial J_{\text{total}}/\partial \boldsymbol{w}$, and used this to perform gradient descent.

To optimize the network to perform reaches and loaded-postures over the central workspace, we optimized an aggregate cost function, $J_{\text{total}} = \sum_i J_i$, composed of the sum of the costs for many random reach and loaded-posture trials (i.e., the J_i's). Using a version of backpropagation through time

(Rumelhart et al., 1986; Werbos, 1990; Stroeve, 1998) modified for our model, we found the partial derivatives of J_{total} with respect to the adjustable parameters in the model and used a conjugate gradient descent algorithm to find a minimum (see Optimization and Analysis and Computational Details in Supplemental Information).

Static Model Specification

For the center-out reaching task, we optimized this model to generate 16 target initial hand velocities equally spaced around the unit circle. These targets are here specified in the 2 × 16 matrix, $\dot{y}_{target}$. Similarly, to emulate the loaded-posture task, we optimized the model to generate 16 joint torque combinations equally spaced around the unit circle and specified in a 2 × 16 matrix, τ_{target}.

As before, we used a network with 1,000 units and neural activity for the 16 targets was thus specified by the matrix, z. To ensure that muscles could only ''pull,'' we employed the standard sigmoidal function for the muscle activation, $u = \sigma_u(W^{out} \cdot z)$, and W^{out} was again the matrix of connectivity strengths between units and muscles. The elements of W^{out} were drawn from a Normal distribution with a mean of 0 and variance 0.001.

To approximate the effect of muscle F-L/V mechanics, we multiplied each muscle's activity by the scaling factor found by averaging (from movement onset to 100 ms after peak velocity) the F-L/V effects associated with the optimal solution found with the dynamic model for a given movement direction. Thus, muscle tensions were found via, $t = H \cdot u$, where, H, is the 6 × 16 matrix of F-L/V scaling factors appropriate for each of the movement directions and · is the element-wise product. Joint torques were found via the linear mapping, $\tau = M \cdot t$, where M is the 2 × 6 matrix of fixed moment arms. We employed a linear approximation, the 2 × 2 matrix F, of the mapping from joint torques to joint velocities, $\dot{x}$, so that, $\dot{x} = F \cdot \tau$. Similarly, we used a linear approximation, the 2×2 matrix G, of the mapping from joint velocities to hand velocities in Cartesian coordinates, so, $\dot{y} = G \cdot \dot{x}$.

For the center-out task, we optimized the cost function given by

$$J(z) = \frac{1}{2}\|G \cdot F \cdot M \cdot (H \cdot \sigma_u(W^{out} \cdot z)) - \dot{y}_{target}\|_F^2 + \frac{1}{2}\alpha\|u\|_F^2 + \frac{1}{2}\beta\|z\|_F^2 \quad \text{(Equation 4)}$$

where, $\|\cdot\|_F$ is the Frobenius norm of a matrix and is equal to the square root of the sum of the squared elements of the matrix and α and β are scalars which weight the importance of keeping muscle and neural activity small. In practice, α and β were both set to 10^{-5}, though any value low enough to allow the optimization to find solutions where the summed target errors were small (e.g., $1/2 \|\dot{y} - \dot{y}_{target}\|_F^2 < 10^{-1}$) produced results similar to those reported here. Under this model, the optimal unit activity is given by, $z^* = \arg\min_z J(z)$. We computed the partials of J with respect to the elements of z, used conjugate gradient descent to find z^* and then fit a plane to the activity of each unit to find its PMD. Thus, in the static model the neural activities are sought directly—in this sense the optimal activities are nonparametric, which has the added benefit that the solutions sought in the static case are not the result of choosing a specific nonlinear activation function. Optimizations were terminated when cost changes remained negligible for 1,000 consecutive updates. For the loaded-posture task, we optimized the cost function given by

$$J(z) = \frac{1}{2}\|M \cdot (H \cdot \sigma_u(W^{out} z)) - \tau_{target}\|_F^2 + \frac{1}{2}\alpha\|u\|_F^2 + \frac{1}{2}\beta\|z\|_F^2 \quad \text{(Equation 5)}$$

and found PTDs for each unit. Equations 4 and 5 require the network to minimize the squared distance between appropriate initiating velocities and torques for the reaching and postural tasks, respectively, while simultaneously minimizing the neural and muscle activities. As with the dynamic model, we repeated these optimizations 10 times with W^{out} initialized randomly each time. See Supplemental Information for details.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2012.10.041.

REFERENCES

Ajemian, R., Green, A., Bullock, D., Sergio, L., Kalaska, J., and Grossberg, S. (2008). Assessing the function of motor cortex: single-neuron models of how neural response is modulated by limb biomechanics. Neuron 58, 414–428.

Alexander, R.M. (1996). Optima for Animals (Princeton, NJ: Princeton University Press).

Asanuma, H., Zarzecki, P., Jankowska, E., Hongo, T., and Marcus, S. (1979). Projection of individual pyramidal tract neurons to lumbar motor nuclei of the monkey. Exp. Brain Res. 34, 73–89.

Bennett, K.M., and Lemon, R.N. (1996). Corticomotoneuronal contribution to the fractionation of muscle activity during precision grip in the monkey. J. Neurophysiol. 75, 1826–1842.

Blakemore, C., and Cooper, G.F. (1970). Development of the brain depends on the visual environment. Nature 228, 477–478.

Brown, I.E., Cheng, E.J., and Loeb, G.E. (1999). Measured and modeled properties of mammalian skeletal muscle. II. The effects of stimulus frequency on force-length and force-velocity relationships. J. Muscle Res. Cell Motil. 20, 627–643.

Cabel, D.W., Cisek, P., and Scott, S.H. (2001). Neural activity in primary motor cortex related to mechanical loads applied to the shoulder and elbow during a postural task. J. Neurophysiol. 86, 2102–2108.

Caminiti, R., Johnson, P.B., and Urbano, A. (1990a). Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. J. Neurosci. 10, 2039–2058.

Caminiti, R., Johnson, P.B., Burnod, Y., Galli, C., and Ferraina, S. (1990b). Shift of preferred directions of premotor cortical cells with arm movements performed across the workspace. Exp. Brain Res. 83, 228–232.

Cheney, P.D., and Fetz, E.E. (1980). Functional classes of primate corticomotoneuronal cells and their relation to active force. J. Neurophysiol. 44, 773–791.

Cheng, E.J., and Scott, S.H. (2000). Morphometry of Macaca mulatta forelimb. I. Shoulder and elbow muscles and segment inertial parameters. J. Morphol. 245, 206–224.

Churchland, M.M., and Shenoy, K.V. (2007). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J. Neurophysiol. 97, 4235–4257.

Classen, J., Liepert, J., Wise, S.P., Hallett, M., and Cohen, L.G. (1998). Rapid plasticity of human cortical movement representation induced by practice. J. Neurophysiol. 79, 1117–1123.

Conrad, B., Wiesendanger, M., Matsunami, K., and Brooks, V.B. (1977). Precentral unit activity related to control of arm movements. Exp. Brain Res. 29, 85–95.

Evarts, E.V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. J. Neurophysiol. 31, 14–27.

Evarts, E.V. (1969). Activity of pyramidal tract neurons during postural fixation. J. Neurophysiol. 32, 375–385.

Fagg, A.H., Shah, A., and Barto, A.G. (2002). A computational model of muscle recruitment for wrist movements. J. Neurophysiol. 88, 3348–3358.

Fetz, E.E., and Cheney, P.D. (1987). Functional relations between primate motor cortex cells and muscles: fixed and flexible. Ciba Found. Symp. 132, 98–117.

Georgopoulos, A.P., Kettner, R.E., and Schwartz, A.B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space.

II. Coding of the direction of movement by a neuronal population. J. Neurosci. 8, 2928–2937.

Graham, K.M., and Scott, S.H. (2003). Morphometry of Macaca mulatta forelimb. III. Moment arm of shoulder and elbow muscles. J. Morphol. 255, 301–314.

Graham, K.M., Moore, K.D., Cabel, D.W., Gribble, P.L., Cisek, P., and Scott, S.H. (2003). Kinematics and kinetics of multijoint reaching in nonhuman primates. J. Neurophysiol. 89, 2667–2677.

Graziano, M.S.A., Cooke, D.F., Taylor, C.S.R., and Moore, T. (2004). Distribution of hand location in monkeys during spontaneous behavior. Exp. Brain Res. 155, 30–36.

Guigon, E., Baraduc, P., and Desmurget, M. (2007). Coding of movement- and force-related information in primate primary motor cortex: a computational approach. Eur. J. Neurosci. 26, 250–260.

Herter, T.M., Kurtzer, I., Cabel, D.W., Haunts, K.A., and Scott, S.H. (2007). Characterization of torque-related activity in primary motor cortex during a multijoint postural task. J. Neurophysiol. 97, 2887–2899.

Hinton, G.E., and van Camp, D. (1993) Keeping neural networks simple by minimizing the description length of the weights. 6th ACM Conference on Computational Learning Theory.

Holdefer, R.N., and Miller, L.E. (2002). Primary motor cortical neurons encode functional muscle synergies. Exp. Brain Res. 146, 233–243.

Howard, I.S., Ingram, J.N., Körding, K.P., and Wolpert, D.M. (2009). Statistics of natural movements are reflected in motor errors. J. Neurophysiol. 102, 1902–1910.

Humphrey, D.R. (1972). Relating motor cortex spike trains to measures of motor performance. Brain Res. 40, 7–18.

Kakei, S., Hoffman, D.S., and Strick, P.L. (1999). Muscle and movement representations in the primary motor cortex. Science 285, 2136–2139.

Kakei, S., Hoffman, D.S., and Strick, P.L. (2001). Direction of action is represented in the ventral premotor cortex. Nat. Neurosci. 4, 1020–1025.

Kalaska, J.F. (2009). From intention to action: motor cortex and the control of reaching movements. Adv. Exp. Med. Biol. 629, 139–178.

Kalaska, J.F., Cohen, D.A.D., Hyde, M.L., and Prud'homme, M. (1989). A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. J. Neurosci. 9, 2080–2102.

Karklin, Y., and Lewicki, M.S. (2009). Emergence of complex cell properties by learning to generalize in natural scenes. Nature 457, 83–86.

Krogh, A., and Hertz, J.A. (1993). A simple weight decay can improve generalization. Adv. Neural Inf. Process. Syst. 4, 950–957.

Kurtzer, I., Pruszynski, J.A., Herter, T.M., and Scott, S.H. (2006a). Primate upper limb muscles exhibit activity patterns that differ from their anatomical action during a postural task. J. Neurophysiol. 95, 493–504.

Kurtzer, I., Herter, T.M., and Scott, S.H. (2006b). Nonuniform distribution of reach-related and torque-related activity in upper arm muscles and neurons of primary motor cortex. J. Neurophysiol. 96, 3220–3230.

Li, Y., Van Hooser, S.D., Mazurek, M., White, L.E., and Fitzpatrick, D. (2008). Experience with moving visual stimuli drives the early development of cortical direction selectivity. Nature 456, 952–956.

Mitsuda, T., and Onorati, P. (2002). Three-dimensional tuning profile of motor cortical activity during arm movements. Neuroreport 13, 1477–1480.

Mussa-Ivaldi, F.A. (1988). Do neurons in the motor cortex encode movement direction? An alternative hypothesis. Neurosci. Lett. 91, 106–111.

Mussa-Ivaldi, F.A., Hogan, N., and Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. J. Neurosci. 5, 2732–2743.

Naselaris, T., Merchant, H., Amirikian, B., and Georgopoulos, A.P. (2006). Large-scale organization of preferred directions in the motor cortex. I. Motor cortical hyperacuity for forward reaching. J. Neurophysiol. 96, 3231–3236.

Nudo, R.J., Milliken, G.W., Jenkins, W.M., and Merzenich, M.M. (1996). Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. J. Neurosci. 16, 785–807.

Olshausen, B.A., and Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature 381, 607–609.

Phillips, C.G., and Porter, R. (1964). The pyramidal projection to motoneurons of some muscle groups of the baboons forelimb. Prog. Brain Res. 12, 222–245.

Porter, R., and Lemon, R. (1993). Corticospinal Function and Voluntary Movement (Oxford: Claredon Press).

Rathelot, J.A., and Strick, P.L. (2009). Subdivisions of primary motor cortex based on cortico-motoneuronal cells. Proc. Natl. Acad. Sci. USA 106, 918–923.

Reimer, J., and Hatsopoulos, N.G. (2009). The problem of parametric neural coding in the motor system. Adv. Exp. Med. Biol. 629, 243–259.

Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986). Learning representations by back-propagating errors. Nature 323, 533–536.

Scott, S.H. (2012). The computational and neural basis of voluntary motor control and planning. Trends Cogn. Sci. 16, 541–549.

Scott, S.H., and Kalaska, J.F. (1997). Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. J. Neurophysiol. 77, 826–852.

Scott, S.H., Gribble, P.L., Graham, K.M., and Cabel, D.W. (2001). Dissociation between hand motion and population vectors from neural activity in motor cortex. Nature 413, 161–165.

Sergio, L.E., and Kalaska, J.F. (2003). Systematic changes in motor cortex cell activity with arm posture during directional isometric force generation. J. Neurophysiol. 89, 212–228.

Shah, A., Fagg, A.H., and Barto, A.G. (2004). Cortical involvement in the recruitment of wrist muscles. J. Neurophysiol. 91, 2445–2456.

Stengel, R.F. (1994). Optimal Control and Estimation (New York: Dover Publications, Inc.).

Parker, G.A., and Smith, J.M. (1990). Optimality theory in evolutionary biology. Nature 348, 27–33.

Stroeve, S. (1998). An analysis of learning control by backpropagation through time. Neural Netw. 11, 709–721.

Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: a model. Nat. Neurosci. 3, 391–398.

Todorov, E. (2004). Optimality principles in sensorimotor control. Nat. Neurosci. 7, 907–915.

Todorov, E., and Jordan, M.I. (2002). Optimal feedback control as a theory of motor coordination. Nat. Neurosci. 5, 1226–1235.

Trainin, E., Meir, R., and Karniel, A. (2007). Explaining patterns of neural activity in the primary motor cortex using spinal cord and limb biomechanics models. J. Neurophysiol. 97, 3736–3750.

Werbos, P.J. (1990). Backpropagation through time: what it is and how to do it. Proc. IEEE 78, 1550–1560.