Cortical Involvement in the Recruitment of Wrist Muscles

Ashvin Shah Andrew H. Fagg Andrew G. Barto

January 23, 2004

Abstract

In executing a voluntary movement, one is faced with the problem of translating a specification of the movement in task space (e.g. a visual goal) into a muscle recruitment pattern. Among many brain regions, the primary motor cortex (MI) plays a prominent role in the specification of movements. In what coordinate frame MI represents movement has been a topic of considerable debate. In a two-dimensional wrist step tracking experiment, Kakei et al. (1999) described some MI cells as encoding movement in a muscle coordinate frame and other cells as encoding movement in an extrinsic coordinate frame. This result was interpreted as evidence for a cascade of transformations within MI from an extrinsic representation of movement to a muscle-like representation. However, we present a model that demonstrates that, given a realistic extrinsic-like representation of movement, a simple linear network is capable of representing the transformation from an extrinsic-space to the muscle recruitment patterns implementing the movements on which Kakei et al. (1999) focused. This suggests that cells exhibiting extrinsic-like qualities can be involved in the direct recruitment of spinal motor neurons. These results call into question models that presume a serial cascade of transformations terminating with MI pyramidal tract neurons that vary their activation exclusively with muscle activity. Further analysis of the model shows that the correlation between the activity of an MI neuron and a muscle does not predict the strength of the connection between the MI neuron and muscle. This result cautions against the use of correlation methods as a measure of cellular connectivity.

Introduction

The primary motor cortex (MI) plays an important role in the control of voluntary movement. Early experiments investigating the control of reaching movements recorded single neuron activity in MI of the monkey as it performed a center-out reaching task with its hand (Georgopoulos et al. 1982; Schwartz et al. 1988). The activity of the MI neurons was described as varying with the cosine between the direction of hand movement in extrinsic space and the neuron's *preferred direction* (PD), also expressed in extrinsic space. A variety of subsequent studies suggest that MI neural activity may capture other aspects of movement. By varying the origin of hand position in a three-dimensional center-out task, Caminiti et al. (1990, 1991) described MI as encoding hand movement in a polar coordinate system centered at the shoulder. Kinematic parameters such as velocity, acceleration, target direction, and target position may also be represented by the activity of an MI neuron — in some cases simultaneously (Moran and Schwartz 1999; Schwartz and Moran 2000; Ashe and Georgopoulos 1994) or sequentially (Fu et al. 1995). Isometric studies suggest that MI may also encode aspects of directional force (Georgopoulos et al. 1992; Sergio and Kalaska 1998).

While an extrinsic representation of movement in MI seems to follow from the results of many of the studies outlined above, representations of movement in *intrinsic space*, such as muscle activity or joint angle deviation, cannot be excluded. If the primate varied its arm configuration while it performed the same hand trajectory, corresponding MI activity was described more accurately by models based on joint kinematics or joint torques than models based on hand movement direction (Scott and Kalaska 1995, 1997). Other models argued that MI activity could be described by equations based on muscle-shortening velocity (Mussa-Ivaldi 1988) or joint angle deviations (Ajemian et al. 2000). Support for a muscle-based representation of movement in MI arose from studies which showed that MI neural activity (and PDs) changed with arm configuration similar to how muscle electromyographic (EMG) activity did in an isometric task (Sergio and Kalaska 1997) and passive arm movement (Scott 1997). Other studies show that the activity of red nucleus neurons, which may encode movement similarly to MI (Houk et al. 1993; Miller and Sinkjaer 1998), correlated better with muscle activity than other variables (Gibson et al. 1985; Houk et al. 1987; Miller et al. 1993; Miller and Houk 1995; Miller and Sinkjaer 1998). Because MI sends projections directly to the spinal cord, among other areas, the idea that MI represents movement in an intrinsic space is an attractive one.

While these studies suggest that MI best represents movement in a particular coordinate frame, many of them indicate that other coordinate frames may be represented as well (though to a lesser degree). The results of more recent studies (Kakei et al. 1999, 2003) make the possible simultaneous representation of multiple coordinate frames in MI neural activity more apparent. These studies investigated the variation of primate wrist muscle EMG and MI neural activity during a two degree-of-freedom (DOF) wrist movement task and provided evidence showing the existence of two subpopulations of MI neurons: neurons with PDs that capture extrinsic properties of movement and neurons with PDs that encode muscle activation patterns. The existence of both types of neurons

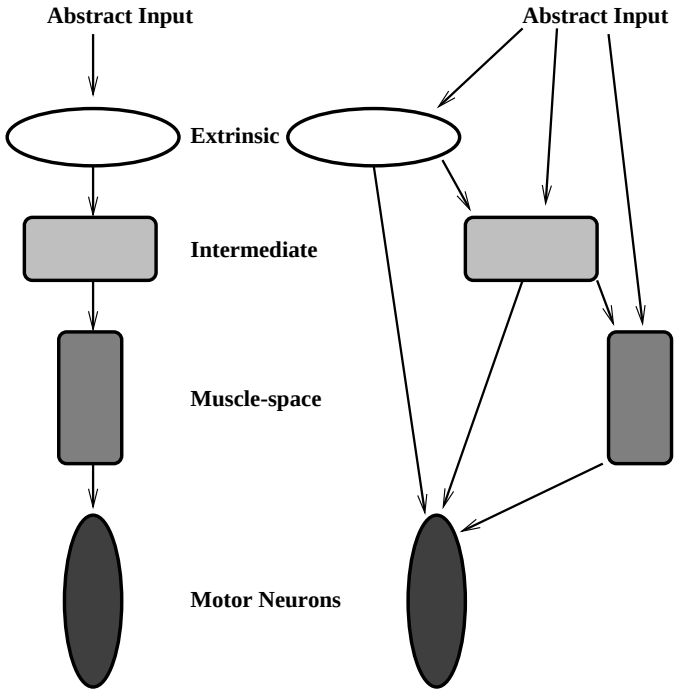

Figure 1: Graphical representation of a serial processing scheme (left) and a parallel processing scheme (right). The different shapes represent populations of neurons whose activities are characterized as encoding movement in different spaces. "Extrinsic," "Intermediate," and "Intrinsic" shapes represent MI neural populations; shapes labeled "Motor Neurons" represent spinal motor neurons.

may support the idea that a serial processing scheme (see figure 1, left side) is implemented within MI (although Scott 2003 offers a different interpretation — see discussion section). Under this model, MI is actively involved in the transformation of an abstract, extrinsic representation of movement into an intrinsic space (Kakei et al. 1999, 2003).

A serial processing scheme is appealing because it offers a simple mechanism for controlling all visually-guided movements. It also implies that the only projections MI sends to the spinal cord to control muscles arise from neurons that explicitly encode movement in a muscle coordinate frame. However, the complex architecture of the central nervous system allows for the possibility of non-serial processing schemes (*cf.* Kalaska and Crammond 1992). Figure 1 illustrates a serial processing scheme (left) versus a series-parallel processing scheme (right, hereafter referred to as "parallel"), in which multiple populations of MI neurons that are characterized as encoding movement in different spaces can each directly command muscles. Can an extrinsic-like representation of movement participate in the direct activation of muscles? Modeling work by Kakei et al. (2003) and Salinas and Abbott (1995) suggests that this is possible. Both show that a linear transformation exists between an extrinsic representation of movement and an intrinsic one under specific assumptions of how movement-related variables are encoded. However, the intrinsic representations used deviate from how muscles are recruited.

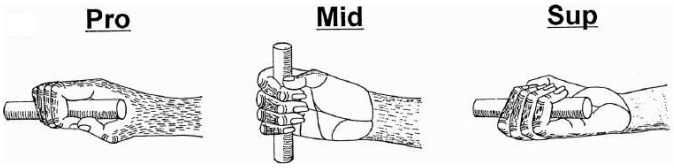

Figure 2: A schematic showing a monkey right hand gripping a handle in pronated (left), midrange (middle), and supinated (right) wrist posture. From Kakei et al. (1999). NEED PERMISSION.

We examine the feasibility of the parallel processing scheme through a neural network model inspired by the work of Kakei et al. (1999). We show that it is computationally possible for a population of neurons with PDs defined in extrinsic space to directly control muscles appropriately. The model produces muscle activation patterns similar to those recorded during the task used in Kakei et al. (1999) (Hoffman and Strick 1999). We train the model with a performance-based optimization procedure and do not impose an *a priori* representation of intrinsic movement such as those used in Kakei et al. (2003) and Salinas and Abbott (1995). We also use the model to examine the use of correlation methods to make inferences of coding schemes and connectivity. Elements of this work have been presented previously in poster form (Shah et al. 2002).

MI and Muscle Involvement in the Production of Wrist Movements

Hoffman and colleagues (Kakei et al. 1999, 2003; Hoffman and Strick 1999) described a two DOF step tracking task in which a human or monkey subject moved a manipulandum with its wrist, fixed in a pronated, supinated, or midrange posture (figure 2), to move a cursor on a computer screen from a central point to one of several targets falling on a circle around the starting location. The kinematics of movement, EMG activity from several muscles, and single neuron activity in MI were recorded. Muscle activity as a function of target direction exhibited a "truncated cosine-like" shape — for values of target direction for which a cosine-like function is negative, the truncated cosine is zero. A cosine of the form $B\cos(\theta - C) + D$ was fit to the muscle activation patterns, with a low weight given to values of muscle activity near zero to account for the truncation. The parameter C defined the muscle's PD for that wrist posture.

The wrist step-tracking task distinguished three coordinate frames: *extrinsic*, represented by the movement of the cursor and unaffected by wrist posture, *joint-centered*, defined by wrist flexion/extension and radial/ulnar deviation, and *muscle*, defined by how the muscle PDs rotated as the wrist posture rotated. As the wrist rotated 180° from pronation to supination, muscle PDs rotated between 40° and 110° (Kakei et al. 1999). MI neurons were labeled *extrinsic-like*, *muscle-like*, or *joint-like*, depending on how their PDs rotated as wrist posture rotated 180°. Because categorization of neurons was based solely on how their PDs rotated, the suffix "-like" was explicitly included. 50% of the MI neurons recorded were labeled "extrinsic-like" because their PDs changed by only a small amount. The depth of modulation of some neurons in this category varied with wrist posture. 32% of the MI neurons were labeled "muscle-like" because their PDs shifted

similarly to muscle PD shifts. None of the neurons were labeled "joint-like," and the rest were not easily classified. The presence of both extrinsic- and muscle-like neurons led Kakei et al. (1999) to suggest that a serial processing scheme was implemented in MI. This implies that only intrinsic neurons (*e.g.* muscle-like MI neurons) can be pyramidal tract neurons (PTNs). If PTNs consisted of only intrinsic neurons, one might expect a spatial and temporal distinction between the intrinsic and extrinsic neurons. However, Kakei et al. (1999) found no such differences.

Methods

Overview and Motivation

The results of Kakei et al. (1999) can be interpreted to support the serial processing scheme. However, is it computationally feasible for *each* population of MI neurons recruited for a task — including extrinsic-like neurons — to *directly* command muscles? One could approach this question through the implementation of a model in which both intrinsic-like and extrinsic-like neurons command muscles. However, in doing so, one would have to make assumptions about the origin of the intrinsic responses. In addition, it would not be clear what the resulting contribution of the extrinsic population would be in the movement generation process. In order to avoid these difficulties, we instead choose to ask a stronger question: can extrinsic-like neurons *alone* produce the appropriate muscle responses? To answer this question, we have created a neural network model in which a population of extrinsic-like neurons directly activate muscles. The model is based on the task and experimental findings of Kakei et al. (1999). The model presented in this paper is an extension of the model presented in Fagg et al. (2002), which directly finds muscle activation patterns that satisfy optimization criteria; the resulting muscle activation patterns are similar to those reported in Hoffman and Strick (1999) for the wrist movement task.

Model Architecture

The architecture of this network is shown in figure 4. In the following paragraphs, vectors and matrices are in bold type while scalers and the elements of the vectors and matrices are in italics. For example, the array of MI neurons is denoted by $\mathbf{m}$, a column vector of which element m_i refers to either $\mathbf{m}$ neuron i or the activation level of $\mathbf{m}$ neuron i. (For simplicity, the term "neuron" will be used to refer to a unit from the $\mathbf{m}$ array hereafter.) By design, the extrinsic-like neurons behave like those recorded in Kakei et al. (1999): their PDs are expressed in extrinsic space, but their activation levels are modulated by wrist posture. Some neurons are more active when the wrist is pronated than when it is supinated and *vice versa*. Each neuron sends projections ($\mathbf{K}$) directly to each of five muscles ($\mathbf{a}$), which correspond to the muscles from which Hoffman and Strick (1999) and Kakei et al. (1999) recorded.

As described in detail in Fagg et al. (2002), each modeled muscle contributes to the endpoint of wrist movement through its activation level and pulling direction, which depends on wrist posture. Hoffman and Strick (1999) determined the pulling direction of a muscle by individually stimulating

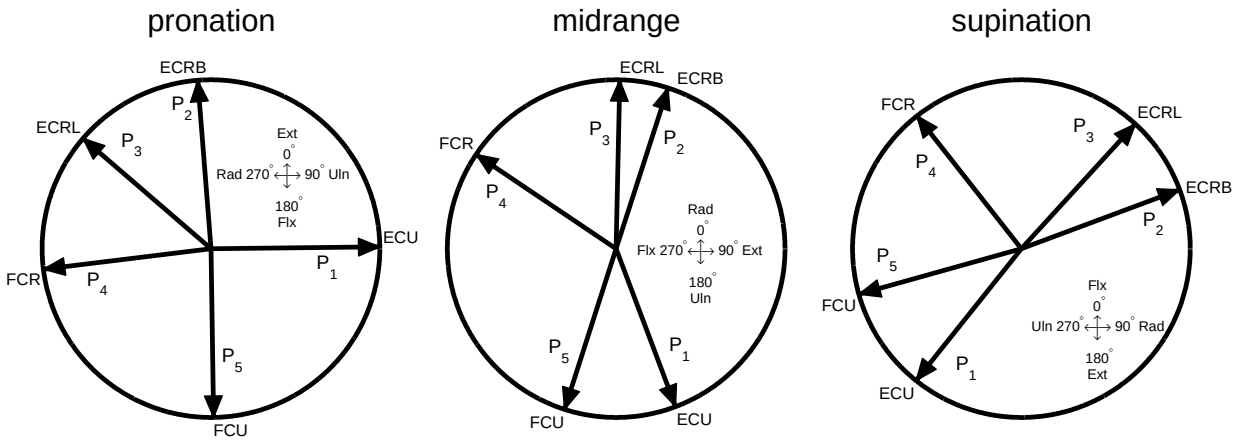

Figure 3: Pulling directions of the five muscles primarily responsible for wrist actuation in the pronated (left), midrange (center), and supinated (right) wrist postures. Pulling direction was defined to be the immediate direction of wrist movement after stimulation of the muscle (Hoffman and Strick 1999). Each vector represents the average pulling direction of the wrist muscle as derived from two monkey subjects. The legend in each circle indicates extrinsic direction (degrees) and joint movement (Rad, radial; Uln, ulnar; Flx, flexion; Ext, extension). Data from DS Hoffman, personal communication. 1-ECU, Extensor Carpi Ulnaris; 2-ECRB, Extensor Carpi Radialis Brevis; 3-ECRL, Extensor Carpi Radialis Longus; 4-FCR, Flexor Carpi Radialis; 5-FCU, Flexor Carpi Ulnaris.

the muscle and observing immediate wrist movement. Average pulling directions (in extrinsic space) for five muscles and three postures are shown in figure 3 (DS Hoffman, personal communication). The activation of the muscles determine extrinsic movement of the wrist ($\mathbf{x}$), analogous to cursor movement in the task.

The neurons in this model are non-spiking; their activation levels are expressed as a real value between zero and one, which can be thought of as the firing-rate of the neuron. $\mathbf{m}$ has $2N$ neurons (we choose $2N = 96$). Within a given wrist posture, each neuron is most active when the target direction is the same as its PD in extrinsic space. For $i = 1, \ldots, N$, the PD of neuron m_i, PD_{m_i}, is defined to be $(\frac{i}{N})360°$, while $PD_{m_{i+N}} = (\frac{i-N}{N})360°$ (thus, $PD_{m_i} = PD_{m_{i+N}}$). In order to create the extrinsic-like behavior found in Kakei et al. (1999), we defined the activity of m_i to take on a Gaussian-like shape around PD_{m_i} and included a wrist posture term (w_i) to modulate its depth:

$$m_i = \left[\exp\left[-\left(\frac{PD_{m_i} - \theta}{\sigma} \right)^2 \right] - w_i(\rho) \right]^+ ,$$

where θ is the target direction and ρ is wrist posture. For $i = 1, \ldots, N$, $w_i(\rho) = 0$, 1/4 or 1/2 for $\rho =$ pronated, midrange, or supinated, respectively, while $w_{i+N}(\rho) = 1/2$, 1/4, or 0. Thus, m_1 through m_N are more active in the pronated wrist posture while m_{N+1} through m_{2N} are more active in the supinated wrist posture. σ defines the width of the Gaussian; this model uses $\sigma = 74.5°$, which produces a width at half maximum of 120°. While a narrow tuning function has some advantages in coding one-dimensional features (Zhang and Sejnowski 1999; Amirikian and Georgopoulos 2000),

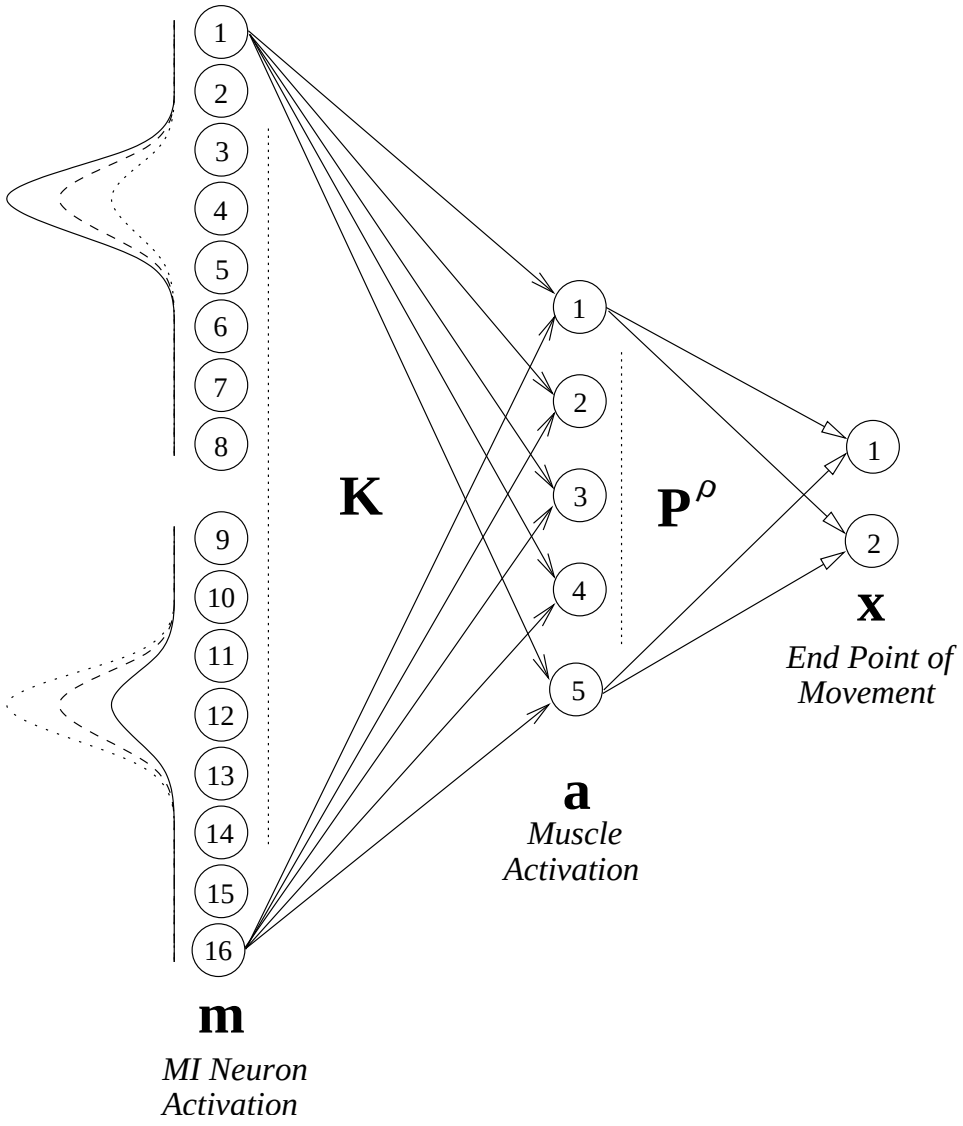

Figure 4: Architecture of this model with $2N = 16$ for clarity. The **m** array represents MI neurons. The curve to the left of the **m** array approximates the activation of the **m** neurons for a target direction near $\theta = 180^\circ$ when the wrist is in the pronated (solid line), midrange (dashed), and supinated (dotted) wrist postures. The **a** array represents the muscles. Each m_i connects to each a_j through the connection matrix **K**. The **x** array represents the endpoint of wrist movement. Each a_i is connected to each x_j through the connection matrix $\mathbf{P}^\rho$, which depends on wrist posture ρ and represents muscle pulling directions for that wrist posture. Full connections are not illustrated for clarity; connections that are not illustrated are represented by vertical dotted lines, which indicate that the connections follow the same pattern as the surrounding illustrated connections. Closed unfilled arrows indicate excitatory connections while open arrows indicate mixed connections. All units have linear activation functions except **m**, which has a lower threshold of zero.

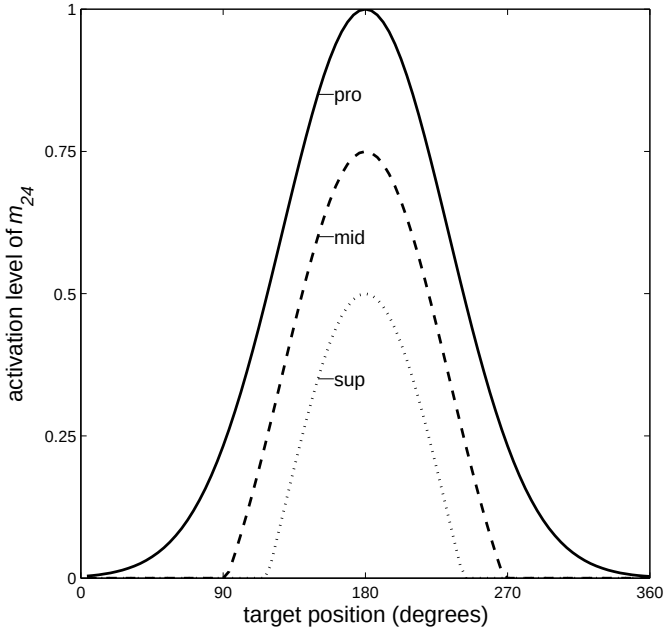

Figure 5: Activity of m_{24} ($PD_{m_{24}} = 180°$) as a function of extrinsic target position for wrist posture pronation (pro: solid line), midrange (mid: dashed line), and supination (sup: dotted line). Note that m_{24+N} has the same PD as m_{24}, but its activation is higher in supination and lower in pronation.

we use a relatively broad tuning function in the interest of generalization. The function $[\cdot]^+$ returns zero if its argument is less than zero — this sets the minimum possible activation level of any m_i to zero.

The hard-wired pattern of activity of the neurons behave similarly to the extrinsic-like neurons modulated by wrist posture recorded in MI by Kakei et al. (1999). Figure 5 shows how the activity level of one neuron varies with extrinsic target position. Although the activity of this neuron is modulated by wrist posture, its PD does not change with wrist posture. Figure 6 shows the activity of the array of neurons in $\mathbf{m}$ when the target is at $\theta = 180°$; the activation pattern is different for the three different wrist postures. Both extrinsic target direction and wrist posture are represented in the activation of the neuron array. While the modulation by wrist posture captures some intrinsic information, a muscle coordinate frame is *not* represented in the activity of the $\mathbf{m}$ array.

The $2N$ $\mathbf{m}$ neurons directly project to an array, $\mathbf{a}$, of five muscles. The connection matrix from $\mathbf{m}$ to $\mathbf{a}$ is referred to as $\mathbf{K}$. There are no imposed constraints on the elements of $\mathbf{K}$; they can take on any real value and each neuron has a direct connection to each muscle (McKiernan et al. 1998). Negative connections can be accounted for with inhibitory interneurons; however, since the abstraction level of this model is high, we choose instead to allow $\mathbf{K}$ to contain both positive and negative elements. Muscle activity is a linear function of $\mathbf{m}$:

$$\mathbf{a} = \sum_{i \in M} \mathbf{K}_i m_i,$$

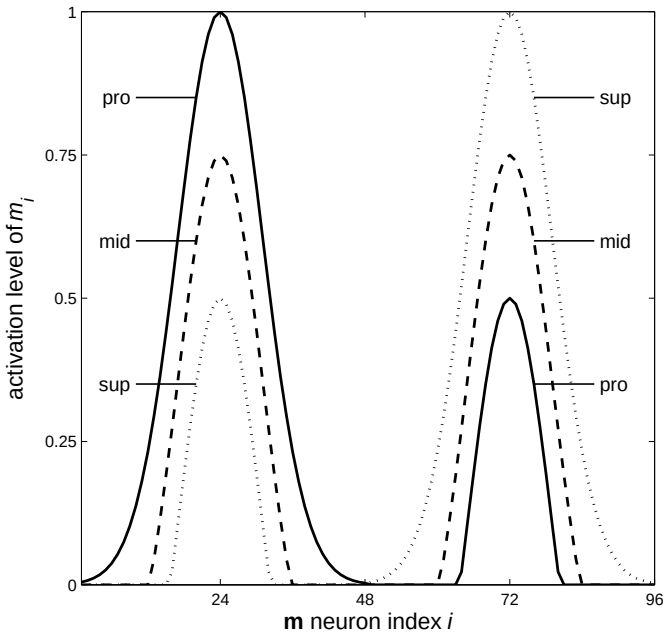

Figure 6: Activity of the array of MI neurons ($\mathbf{m}$, $2N = 96$ neurons) for pronation (pro: solid line), midrange (mid: dashed line), and supination (sup: dotted line) wrist postures when the target direction is $\theta = 180^\circ$. For m_1 through m_{48}, the activity in pronated position is greater than the activity in midrange, which is greater than the activity in supinated position. For m_{49} through m_{96}, the activity in supinated position is greater than the activity in midrange, which is greater than the activity in pronated position. Both extrinsic target direction and intrinsic wrist posture are represented by the activity of this array of neurons. Muscle activity is not represented.

where $\mathbf{K}_i$ denotes the weights of the projections from neuron m_i to each of the five muscles and M is the set of all neurons in $\mathbf{m}$.

Each muscle is represented as a unit vector in the direction of its pulling direction, indicated in figure 4 as the matrix $\mathbf{P}^\rho$, which changes with wrist posture, but is not plastic. We assume that muscles contribute to the endpoint of movement along their vector of action, where the length of the vector is proportional to the muscle's activation level, a_i (analogous to the muscle EMG level). The muscles are also assumed to contribute to the movement endpoint independently of one-another. Thus, the endpoint of movement, $\mathbf{x}$, is determined by the weighted sum of the muscle pulling directions as follows:

$$\mathbf{x} = \sum_{i \in A} \mathbf{P}_i^\rho a_i,$$

where $\mathbf{P}_i^\rho$ denotes the pulling direction of muscle i when in posture ρ and A is the set of all muscles. Note that only the endpoint of wrist movement, not the path of movement, is simulated.

Selecting the MI → Muscle Connections

We wish to show that extrinsic-like neurons can directly generate realistic muscle activation patterns. We do this in the context of the model by showing that a $\mathbf{K}$ exists such that for each target position and wrist posture, *i.*) the wrist reaches the target and *ii.*) muscles only pull (and never push). In Fagg et al. (2002), we discuss the redundancy in muscle activation patterns and show that the introduction of a minimum effort criterion applied to the muscles results in a cosine-like recruitment of muscles similar to that observed experimentally (Hoffman and Strick 1999). Thus, we include a third criterion: *iii.*) a minimal degree of effort is used in making the movement.

We employ a gradient-descent method to identify an appropriate $\mathbf{K}$ matrix. This is accomplished as follows: the initial elements of $\mathbf{K}$ are randomly chosen from a uniform distribution between $-1/2$ and $1/2$. For a given target and wrist posture, a random $\mathbf{K}$ yields a random muscle activation pattern produced by $\mathbf{m}$ and hence a random (and inaccurate) endpoint of wrist movement. As in Fagg et al. (2002), we define the error as follows:

$$E(\mathbf{x_{targ}}, \mathbf{a}, \rho) = \frac{1}{2} \left\| \mathbf{x_{targ}} - \sum_{i \in A} \mathbf{P}_i^\rho a_i \right\|^2 + \frac{\lambda}{2} \|\mathbf{a}\|^2 ,$$

where $\mathbf{x_{targ}}$ is the target location, $\| \cdot \|$ returns the magnitude of a vector, and λ is a regularization parameter set to 0.02. λ represents a trade-off between target error and muscle activation. Since all movements are of unit magnitude, and $\mathbf{x}$ is a linear summation of muscle activity, $\|\mathbf{a}\|^2$ is on the order of 1 for most movements. Thus, $\lambda = 0.02$ represents allowable errors on the order of 2% of movement magnitude. The first term in $E(\mathbf{x_{targ}}, \mathbf{a}, \rho)$ represents movement accuracy (criterion *i*); the second term represents total muscle activation (effort, criterion *iii*). The error gradient with respect to a_j is:

$$\frac{\partial E}{\partial a_j} = - \left(\mathbf{x_{targ}} - \sum_{i \in A} \mathbf{P}_i^\rho a_i \right)^T \mathbf{P}_j^\rho + \lambda a_j,$$

where T denotes the transpose of a vector. The error for muscle j is defined as:

$$e_j = \begin{cases} -\frac{\partial E}{\partial a_j} & \text{if } a_j \geq 0, \\ -a_j & \text{otherwise.} \end{cases}$$

To constrain muscle activation to non-negative levels (criterion *ii.*), e_j is set to minimize $E(\mathbf{x_{targ}}, \mathbf{a}, \rho)$ only if $a_j \geq 0$; otherwise, e_j is set to bring a_j toward zero.

The connection matrix $\mathbf{K}$ is modified as follows:

$$\mathbf{K_{new}} \leftarrow \mathbf{K_{old}} + \eta \mathbf{e}\mathbf{m}^T,$$

where η is a learning rate set to 0.02 in this model and $\mathbf{e}$ is a column vector representing the error terms for all five muscles. Note that this is not intended as a biologically based learning approach. A $\mathbf{K}$ is chosen such that all a_js are non-negative (criterion *ii.*); within this constraint, a $\mathbf{K}$ is selected such that it minimizes $E(\mathbf{x_{targ}}, \mathbf{a}, \rho)$, which includes criteria *i.* and *iii.* The gradient descent method was applied for each of the three wrist postures and 12 targets equally spaced along a unit circle surrounding the central starting position, yielding a total of 36 distinct tasks; the 36 tasks constitute one training epoch. Note that each task, by construction, is represented by a unique $\mathbf{m}$ activation pattern (figure 6). $\mathbf{K}$ was updated iteratively on each of the 36 tasks until the mean target error over all 36 tasks was less than 0.05.

30 independent training runs were executed, each with a unique initial $\mathbf{K}$. Over the 30 runs, the iterative procedure described above converged to a solution within an average ($\pm$ standard deviation) of $210,150$ ($\pm 7,137$) training epochs. For each run, the standard deviation in target error (over the 36 tasks) was computed. The mean standard deviation in target error over the 30 runs was 0.041 ($\pm 4.13 \times 10^{-4}$). The mean maximum change in any single muscle activation on the last epoch over all 36 tasks was 0.028 ($\pm 3.19 \times 10^{-4}$), while the mean maximum change in any single element of $\mathbf{K}$ was 0.0014 ($\pm 1.95 \times 10^{-5}$). Thus, termination of a run was accompanied by convergence of the $\mathbf{K}$ matrix. For an individual task, the muscle activation pattern ($\mathbf{a}$, the five-element vector describing the activation levels of the five muscles) at convergence deviated from the average over all 30 runs by 0.011 ($\pm 7.4 \times 10^{-3}$), regardless of initial conditions (the initial $\mathbf{K}$ matrix). The $\mathbf{K}$ matrix describing the weights of the five projections from each $\mathbf{m}$ neuron, on the other hand, deviated from the average over all 30 runs by 0.52 ($\pm$ 0.15). However, there was a general pattern across the 30 runs (see figures 10B and 10C). While there is no unique solution for $\mathbf{K}$, the different $\mathbf{K}$s do yield the same muscle activation patterns for each of the 36 tasks. The existence of unique muscle activation patterns for each of the 36 distinct tasks is discussed in detail in Fagg et al. (2002). The length of the muscle activation vector, averaged over all 36 tasks, was $\|\mathbf{a}\| = 1.20$ (± 0.38).

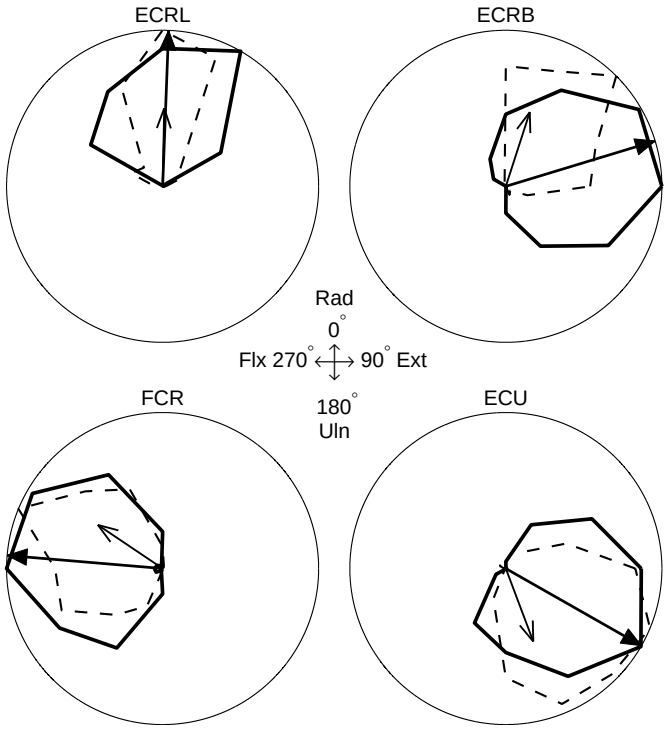

Figure 7: Normalized polar plots of muscle activation as a function of target direction for muscles ECRL, ECRB, FCR, and ECU in the midrange wrist posture. Solid lines indicate muscle activation as produced by the model; dashed lines indicate muscle activation as produced by a monkey subject (Hoffman and Strick 1999; DS Hoffman, personal communication). Short open arrows are muscle pulling directions. Long closed arrows are model muscle PDs. The legend in the middle of the figure denotes joint angle deviation and corresponding extrinsic direction of movement in degrees for the midrange wrist posture: Rad, radial; Uln, ulnar; Ext, extension; Flx, flexion.

Results

Muscle Activation

The following analysis uses the results from one of the 30 runs. All other runs have very similar results. Figure 7 shows normalized model muscle activation versus target direction as polar plots for muscles ECRL, ECRB, FCR, and ECU in the midrange wrist posture. Hoffman and Strick (1999) analyzed these four muscles from a monkey subject; their data (Hoffman and Strick 1999; DS Hoffman, personal communication) are included in figure 7 as dashed lines for comparison. The function $B\cos(\theta - C) + D$, where the PD of a muscle is defined to be the parameter C, was fitted to each muscle activation pattern. Because muscle activation patterns followed a "truncated cosine" (due to the constraint that all $a_i \geq 0$), low muscle activation levels (< 0.05 in our analyses) were given a zero weight in the fitting process. The muscle activation patterns as produced by this model are similar to those reported in Fagg et al. (2002), which used a gradient descent method to directly find muscle activation patterns that satisfy the same three criteria as this model. The model presented in this paper, in contrast, finds a connection matrix $\mathbf{K}$ to transform the activity of the $\mathbf{m}$ array into the appropriate muscle activation pattern.

Figure 7 also shows the pulling directions (short open arrow) and calculated PDs (long closed arrow) of the model muscles. In all four cases, the PD deviates from the pulling direction. The discrepancy between pulling directions and preferred directions is also seen in the EMG data from Hoffman and Strick (1999). This difference is the result of the uneven distribution of pulling directions of the muscles (Fagg et al. 2002). If there is a large gap between the pulling directions of two muscles, then the muscles have to devote additional effort to pulling against each other in order to reach a target located within the gap. Thus, the PD of a muscle will tend toward this gap.

The similarity between the muscle activation patterns as produced by this model and those recorded experimentally demonstrate that it is possible for an extrinsic-like representation to be translated (via a linear transformation) into an intrinsic representation of movement (muscle activation patterns).

Correlation Analysis

Inspired by studies that used correlation analysis to suggest that neurons in the red nucleus and MI directly encode muscle activity (Miller and Sinkjaer 1998; Miller and Houk 1995; Miller et al. 1993; Houk et al. 1993; Gibson et al. 1985; Houk et al. 1987; Gibson et al. 1985), we have calculated the correlation coefficient for each neuron-muscle pair over all 36 tasks in our model. The formula used to calculate the correlation coefficient ($corr_{ij}$) between $\mathbf{m}$ neuron i (m_i) and muscle j (a_j) is:

$$corr_{ij} = \frac{\sum_{\vartheta \in \Theta} (m_i(\vartheta) - \mu_{m_i})(a_j(\vartheta) - \mu_{a_j})}{\sqrt{\sum_{\vartheta \in \Theta} (m_i(\vartheta) - \mu_{m_i})^2} \sqrt{\sum_{\vartheta \in \Theta} (a_j(\vartheta) - \mu_{a_j})^2}},$$

where ϑ is the index of task in Θ (the set of 36 tasks), $m_i(\vartheta)$ is the activation of m_i for task ϑ, $a_j(\vartheta)$ is the activation of a_j for task ϑ, μ_{m_i} is the average activation of m_i over all 36 tasks, and μ_{a_j} is the average activation of a_j.

The correlations between each neuron and each muscle ranged between -0.68 to 0.88. While a high positive correlation might be interpreted as the neuron coding in a muscle coordinate frame, in this model it merely indicates that the neuron and the muscle happen to have similar PDs — they are highly active over the same range of tasks. Figure 8 plots, as a function of target direction and for each of the three wrist postures, the activities of the neuron and muscle that have the highest $corr_{ij}$ over all neuron-muscle pairs and the activities of the neuron and muscle that have the lowest $corr_{ij}$. The highest correlation ($corr_{ij} = 0.88$) was found between neuron number 49 (out of $2N = 96$) and muscle ECRL. m_{49} has a PD of 7.5°, while ECRL has a PD of 342.7° in the pronated wrist posture, 2.2° in the midrange posture, and 17.6° in supinated posture. Because the neurons and muscles are active over a wide range of target directions, the PDs of m_{49} and ECRL are considered to be similar (figure 8A). The lowest correlation ($corr_{ij} = -0.68$) was found between m_{39} and muscle ECU. Neuron m_{39} has a PD of 292.5°, while ECU has a PD of 89.3° in the pronated wrist posture, 119.7° in the midrange posture, and 155.4° in the supinated posture. m_{39} and ECU were active over opposite target directions (figure 8B).

Figure 9 summarizes this same relationship between all neuron-muscle pairs. The $corr_{ij}$ between each muscle and neuron as a function of the PD of the neuron is plotted. Note that by construction, neurons i and $i + N$ (for $1 \leq i \leq N$) have the same PD. The scale bars on the top left of figure 9 illustrate the relationship between the PDs of the neurons and the index of the neurons (also see methods section for definition of PD_{m_i}). Also shown for each muscle are their PDs in each of the three wrist postures. The $corr_{ij}$s between each muscle and the neurons vary sinusoidally with the PDs of the neurons. For each muscle, there is some neuron with which it has a high $corr_{ij}$ (ranging from 0.77 to 0.88, depending on the muscle). In each case, the PD of that neuron falls within the range of the PDs of the muscles as the wrist rotates from pronation to supination.

Figure 9 also shows that the $corr_{ij}$ between a muscle and a neuron is not determined purely by the neuron's PD. As shown in figure 6, neurons 1 through N are more active when the wrist is in the pronated posture, while neurons $N + 1$ through $2N$ are more active when the wrist is in the supinated posture. If the muscle has the highest $corr_{ij}$ with a neuron of index $1 \leq i \leq N$, then that muscle's PD in pronation is closer to PD_{m_i} than its PD in supination (for example, muscle FCU in figure 9). Similarly, if the muscle has the highest $corr_{ij}$ with a neuron of index $N + 1 \leq i \leq 2N$, then that muscle's PD in supination is closer to PD_{m_i} than its PD in pronation (for example, muscle ECRB in figure 9). This is the case because the higher activation level of the neuron in one wrist posture over the others contributes more strongly to the correlation measure. Therefore, the $corr_{ij}$ between a muscle and a neuron is determined by their PDs and how active they are in each of the three wrist postures.

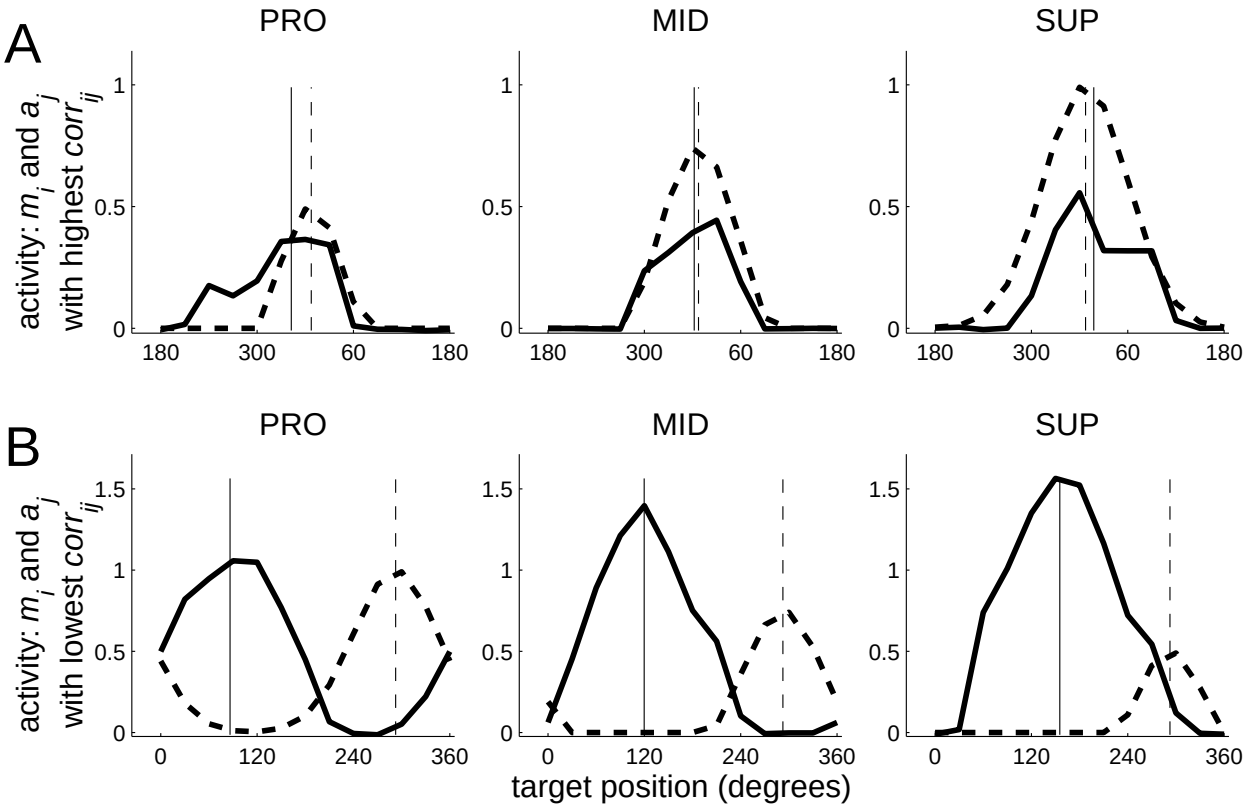

Figure 8: Activity of an **m** neuron (dashed curve) and activity of a muscle (solid curve) as a function of target direction for each of the three wrist postures. *A, top three graphs*: m_{49} and muscle ECRL, which have a high correlation (= 0.88). Note that these graphs are centered around a target direction of 7.5°, the PD of m_{49}, for clarity. *B, bottom three graphs*: m_{39} and muscle ECU, which have a highly negative correlation (= −0.68). Vertical lines indicate the PDs of the **m** neurons (dashed lines) and muscles (solid lines). PRO: pronated, MID: midrange, SUP: supinated. Note that the activation function of the **m** neurons are only plotted for the 12 targets, hence the curves are not as smooth as those in figure 5.

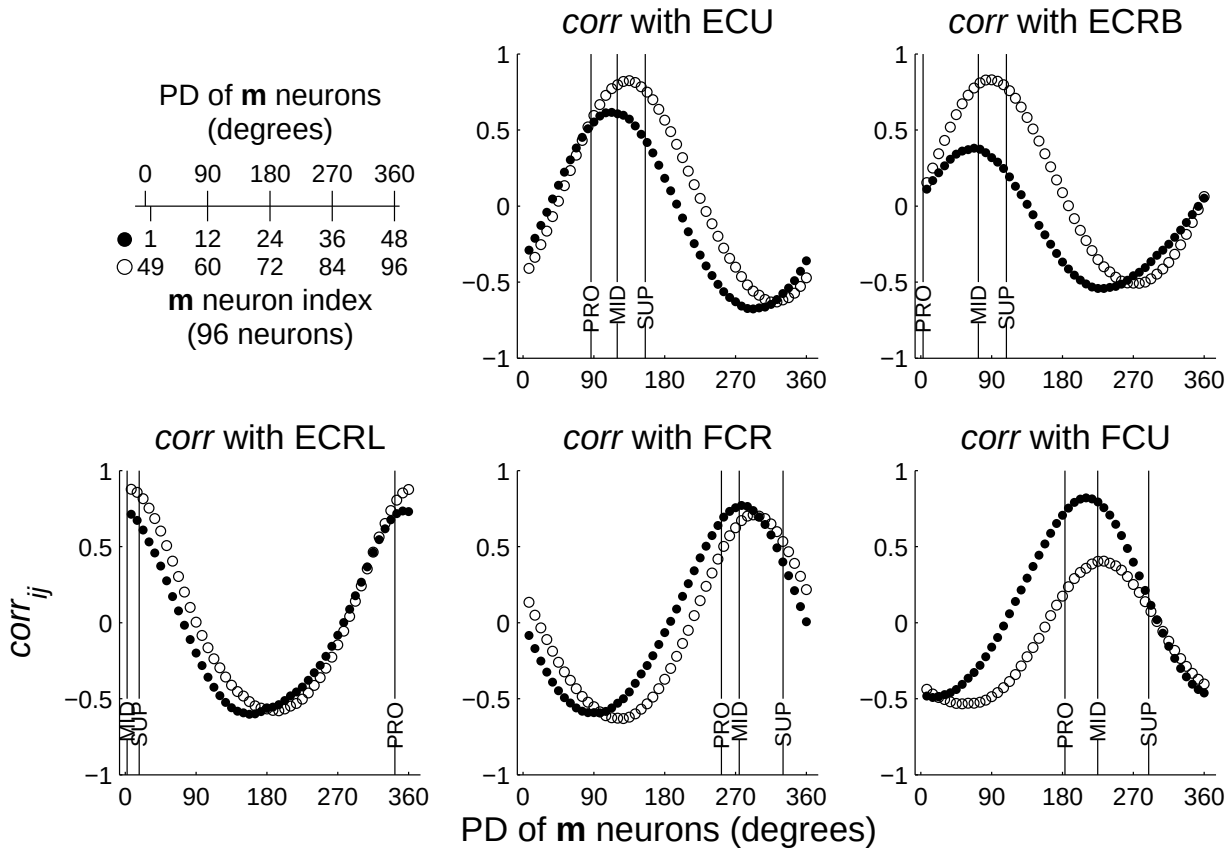

Figure 9: Correlation coefficient ($corr_{ij}$) of each muscle with each **m** neuron, versus the PDs of the **m** neurons. Note that by construction, **m** neurons i and $i + N$ (for $1 \leq i \leq N$) have the same PD. Closed circles mark $corr_{ij}$s between the muscle and **m** neurons 1 through N, while open circles mark $corr_{ij}$s between the muscle and **m** neurons $N + 1$ through $2N$. $N = 48$. The PDs of each muscle when the wrist is pronated (PRO), midrange (MID), or supinated (SUP) postures are shown as vertical lines. *Top left*: scale bars illustrating the relationship between **m** neuron index and **m** neuron PD. The $corr_{ij}$ depends on how similar the PD of the **m** neuron is to the PDs of the muscle.

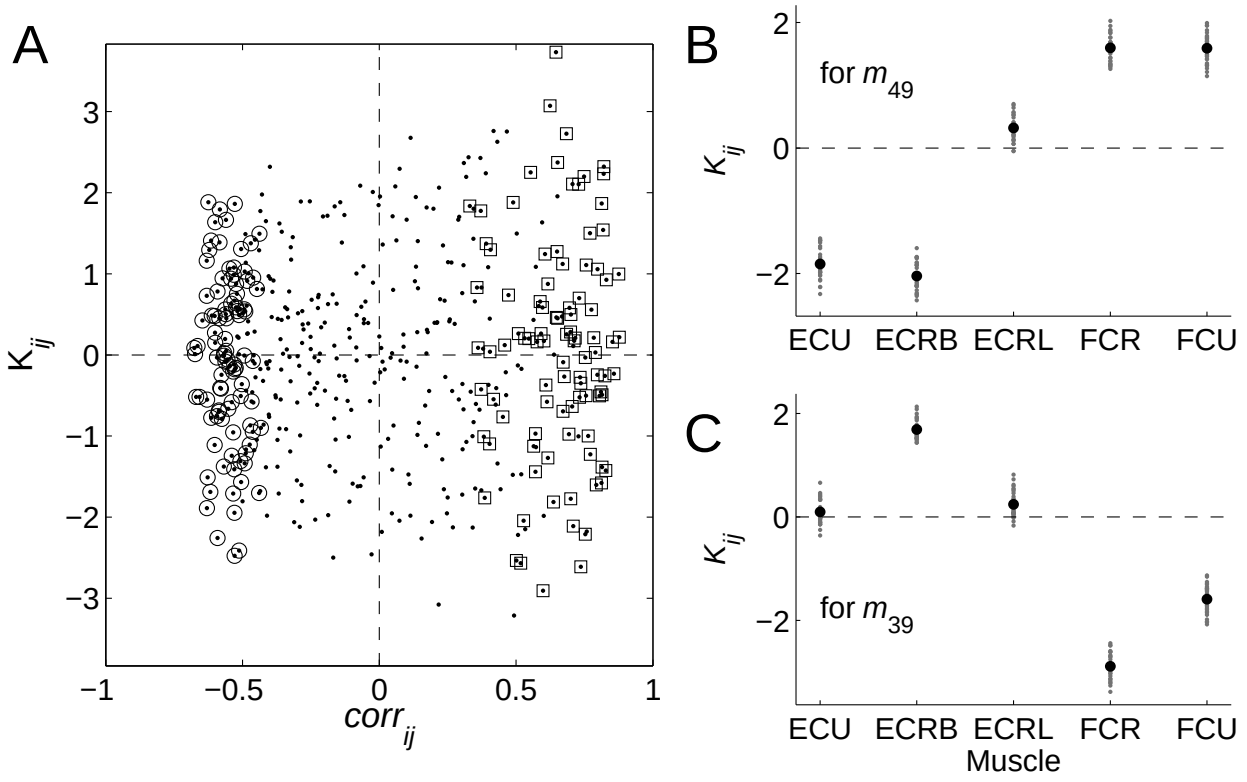

Figure 10: *A*: Scatter plot of the connection weight between m_i and a_j (K_{ij}) as a function of $corr_{ij}$.The lowest $corr_{ij}$ associated with each neuron across the muscle set is indicated by a circle, while the highest $corr_{ij}$ is indicated by a square. Dashed lines show $corr_{ij} = 0$ and $K_{ij} = 0$. *B*: The strength of the connections emanating from m_{49} to each of the five muscles for all 30 runs (gray dots) and the mean of the 30 runs (black dots). Dashed line shows $K_{ij} = 0$. *C*: Same as *B* but for m_{39}. This figure shows that $corr_{ij}$ does not predict K_{ij}.

Connection Strength and Correlation

High correlations between the activity of a neuron and the activity of a muscle may be used as evidence that there is a positive connection between that neuron and muscle. However, this is not the case in this model. Figure 10A plots the weight, K_{ij}, of the connection between m_i and a_j as a function of $corr_{ij}$. The highest correlation associated with each neuron (across the muscles) is highlighted as a square, while the lowest correlation associated with each neuron is highlighted as a circle. The corresponding connection weights for both highest and lowest $corr_{ij}$s are normally distributed: the mean K_{ij} ($\pm$ standard deviation) for the highest $corr_{ij}$s is 0.096 ($\pm$1.36), while the mean K_{ij} for the lowest $corr_{ij}$s is -0.049 ($\pm$1.04). Thus, $corr_{ij}$ is not related to K_{ij} — for some neurons, the connection weight between it and the muscle with which it has a high $corr_{ij}$ is highly negative. The variability in the distribution of lowest $corr_{ij}$s (circles) is less than that of the highest $corr_{ij}$s (squares). This is due to the requirement that the linear transform specified by $\mathbf{K}$ must produce non-negative muscle activation levels.

Figures 10B and 10C show the weight vector (the strength of the connections emanating from a single neuron to each of the five muscles) for neurons 49 and 39, respectively, for all 30 runs (gray dots) and the mean of the 30 runs (black dots). While there is considerable variability between the 30 runs, the weight vector exhibits a general pattern: although the set of all K_{ij}s ranges from -3.2 to 3.7 for the single run focused on in these results, an individual K_{ij} will only vary on the order of 1 across the 30 runs. Figure 10B shows that the K_{ij} between m_{49} and ECRL (which have the highest $corr_{ij}$ for the single run) is near zero, while the K_{ij} between m_{39} and FCU (which have the lowest $corr_{ij}$) is moderately negative. In addition, the K_{ij} between m_{49} and FCU is fairly high, but their $corr_{ij}$ is -0.44 for the single run. The K_{ij} between m_{39} and FCR is highly negative, but their $corr_{ij}$ is 0.74. Thus, even given the pattern across the 30 runs, K_{ij} and $corr_{ij}$ may be conflicting.

Figure 10 shows that $corr_{ij}$ does not predict the connection strength between m_i and a_j (*cf.* König and Engel 1995; Munk et al. 1995; Engel et al. 1991). By construction, one neuron will exhibit correlated activity with other neurons. All that is required to produce the observed neuron-muscle correlation is that the net connection strength to the muscle be positive. This effect is possible because there are many more neurons than muscles or controlled degrees-of-freedom. Figures 9 and 10 also show that each neuron exhibits the full range of $corr_{ij}$s from highly positive to highly negative with some muscle.

Sensitivity to Noise

To determine the robustness of the model in the presence of noise, signal dependent noise (Harris and Wolpert 1998) was added to neuron activity. For each neuron, noise was selected from a normal distribution with a mean of zero and a standard deviation of $m_i\varepsilon$. Simulations were run 30 times each with $\varepsilon = 0.01$ (low noise) and $\varepsilon = 0.1$ (high noise), which represent a range of biologically plausible magnitudes of noise (Todorov 2002; van Beers et al. 2003). For both cases, the model converged to a unique $\mathbf{K}$ (in contrast to the highly variable $\mathbf{K}$s found by the zero-noise model). Like the $\mathbf{K}$s trained without noise, the elements of the $\mathbf{K}$s trained with a low noise level had a

wide range (-4.4 to 4.7). However, for high noise, the elements of $\mathbf{K}$ had a narrow range (-0.19 to 0.34). In addition, for low noise, $corr_{ij}$ did not predict K_{ij}, but for high noise, there was a weak predictive relationship. The mean K_{ij} for the highest $corr_{ij}$s was 0.06 (± 0.09) while the mean K_{ij} for the lowest $corr_{ij}$s was -0.02 (± 0.04). While this tendency was significant (t-test, $p < 0.001$), some K_{ij}s between a neuron and muscle with a positive $corr_{ij}$ were negative (and *vice versa*). We examine the implications of these results in the discussion section.

In neither case did the model reach the termination condition (average target error < 0.05) within one million training iterations. However, the resulting $\mathbf{K}$ produced an endpoint of movement near the target. Model strategy was tested by removing noise and examining movement endpoint and muscle activation. The average target error (over the 36 tasks) was 0.07 ($\pm$ 0.06) for low noise and 0.12 ($\pm$ 0.08) for high noise. For both levels of noise, a unique (across the 30 runs) muscle activation vector was found for each of the 36 tasks. To determine if neural noise resulted in a change in muscle recruitment, we computed $\|\mathbf{a}_\circ(\vartheta) - \mathbf{a}_\varepsilon(\vartheta)\|$, where $\mathbf{a}_\circ(\vartheta)$ is the muscle activation vector for task ϑ derived from zero-noise training and $\mathbf{a}_\varepsilon(\vartheta)$ is derived from noisy training. Averaged over the 36 tasks, the deviation in muscle recruitment was 0.15 ($\pm$ 0.10) for low noise and 0.37 ($\pm$ 0.18) for high noise. The deviation in movement endpoint, computed in the same manner, was 0.05 (± 0.03) for low noise and 0.09 ($\pm$ 0.05) for high noise. Thus, even though the introduction of noise may result in different muscle activation patterns and movement endpoints for some tasks, the strategy was similar.

We examined the model's ability to reach targets on which it was not explicitly trained. A model with $\mathbf{K}$ trained on the original 12 targets and no noise was presented with 144 random targets (taken from a uniform distribution between 0° and 360°). While movement endpoints were in the vicinity of the targets, accuracy was compromised — mean target error increased to 0.22 ($\pm$ 0.14) over the 432 tasks (144 targets and 3 wrist postures). If the model was trained with random targets for each training epoch (for one million epochs), and thus sampled a greater proportion of possible target directions, the mean error over 144 random targets was 0.11 ($\pm$ 0.08). We also examined the model's sensitivity to the width of the target encoding. When the width of the $\mathbf{m}$ neuron tuning function was increased to $\sigma = 114.5^\circ$, a model trained on the original 12 targets and tested with 144 random targets had a mean target error of 0.14 ($\pm$ 0.10). Thus, when trained over a greater variety of targets or with a wider neuron tuning function, the model's ability to generalize increased. Finally, we tested a model with random PDs for the $\mathbf{m}$ neurons (as opposed to evenly distributed PDs) for 30 runs. The only notable difference between the two models was that the model with randomly distributed MI PDs took an average of 343,330 ($\pm$ 200,060) iterations to reach the termination condition.

Discussion

A range of experimental and modeling studies have been used to argue that the primary motor cortex encodes movement in one or a small number of distinct coordinate frames (*e.g.*, Georgopoulos et al. 1982; Caminiti et al. 1991; Georgopoulos et al. 1992; Mussa-Ivaldi 1988; Scott and Kalaska

1997; Ajemian et al. 2000; Miller and Houk 1995; Kakei et al. 1999). These studies rely on evidence of MI cell activity correlating with some aspect (or aspects) of an executed movement. Emerging from some of these studies, and from analogies with the robot control domain, is a serial processing scheme (*cf.* Dum and Strick 2002; Scott 2000; Loeb et al. 1999) in which an extrinsic representation of movement (such as a visual cue) is transformed through multiple stages and brain regions and culminates within MI, where movement may be encoded as an explicit representation of muscle recruitment levels. However, the results of Kakei et al. (1999) argue against such a specialized role for MI. In particular, a significant subset of the cells observed in this study exhibited what Kakei et al. (1999) termed an "extrinsic-like" behavior, in which cell preferred direction (PD) did not change substantially despite changes in the muscle recruitment pattern. It should be noted that a large subset of these cells did exhibit a change in the depth of modulation as a function of wrist configuration, and hence encoded some intrinsic information. In addition, a subset of cells exhibited "muscle-like" behavior, in which cell PD changed in a fashion similar to muscle PD. The existence of both extrinsic- and intrinsic-like populations of cells within MI led Kakei et al. (1999) to support the serial processing scheme and to suggest that a portion of the transformation might take place within MI itself.

The modeling results presented in this paper suggest that populations of MI neurons that encode movement in different coordinate spaces can each be directly involved in the recruitment of spinal motorneurons (we refer to this approach as a *parallel processing scheme*; figure 1). This perspective challenges the pure serial processing scheme in which intrinsically-behaving (specifically, muscle-like) MI cells are the only cortical source of motorneuron input. In addition, experimental evidence suggests that premotor areas, which encode movement in a more abstract coordinate frame and likely take part in an earlier stage of the sensorimotor transformation (Kakei et al. 2001, 2003), have direct connections to spinal motorneurons (Dum and Strick 2002; Luppino et al. 1994). Through the use of a model of the Kakei et al. (1999) task, we ask about the computational feasibility of the parallel processing scheme. Rather than implementing a model in which both intrinsic- and extrinsic-like cells activate muscles, we chose instead to include only the extrinsic-like cells. Such an approach asks a more interesting question than the former (since intrinsic-like cells should be capable of driving muscles without the addition of extrinsic-like cells), and also asks a harder computational question. One benefit of making such a simplification is that it avoids the need for making specific assumptions about the origins of the muscle-like cell responses. Nevertheless, this approach is not meant to argue against the existence or utility of intrinsically behaving cells.

Our model shows that a linear transformation exists from a population of extrinsic-like cells to a muscle activation pattern capable of producing movements to a specified target for the task described in this paper. In addition to generating appropriate movements, the model also produces muscle activation patterns similar to the EMG patterns recorded by Hoffman and Strick (1999). Taken together, these results demonstrate that it is computationally feasible for extrinsic-like MI neurons to directly command muscle activation patterns, which in turn supports the plausibility of the parallel processing scheme.

Interpreting MI Neural Activity

Why do the results of different experimental studies lead to different interpretations of what is being encoded by MI? Todorov (2000) suggests that because muscle activation patterns (and hence forces) are affected not only by the descending motor command, but also by the muscle's length, change in length, and acceleration of length, the form of the experimental question can lead to different interpretations about what is being coded by cells of MI. Thus, in some tasks MI activity appears to reflect extrinsic variables, while in other tasks MI activity appears to reflect intrinsic variables. In his paper, the issue of multiple representations being simultaneously exhibited by MI is not addressed (although the issue is also not excluded). An explanation for simultaneous representations is given by Scott (2003). In his review of the role of MI in goal-directed movements, Scott (2003) suggests that the different populations of MI neurons, each seemingly representing different facets of movement, control different elements of the task. For example, in order to move its arm in a certain way, a primate must also control its posture, different types of spinal neurons, and a multitude of other variables often ignored in data analysis and modeling work. The different coordinate frames simultaneously represented in MI neural activity may be due to its control of the myriad variables necessary to accomplish a task. Although this may be the case, we show that it is computationally possible for units that behave in one coordinate frame to directly control units that behave in a different coordinate frame — neurons do not have to correlate their activity with that of the units they command. Therefore, neurons that behave differently (*i.e.* extrinsic-like and muscle-like neurons) can control the *same* variables (*i.e.* muscle activation) in parallel. In addition, Loeb et al. (1999) discusses how having different populations of neurons commanding the same variables may increase performance and robustness.

Related Modeling Work

Several other modeling studies used error-driven backpropogation to show that a sensorimotor transformation can be computed by a simple network. Anastasio and Robinson (1989) and Anastasio and Robinson (1990) developed three-layer neural network models of the vestibulo-oculomotor reflex in which the transformation from the vestibular signals to the appropriate oculomotor responses were performed by a hidden layer. The authors found that for a network with an equal number of units in the input, hidden, and output layers, the hidden layer units behaved as tensor theory would predict: the units were maximally active for a specific direction of eye or head rotation that coincided with the direction of motor response of an output neuron. However, when the number of hidden layer units was greater than the number of output units, the PDs and activation magnitudes of the hidden layer units were distributed and did not coincide with the output units. The latter model produced hidden layer units that behaved much more like the vestibular nucleus neurons of the cat, which are thought to perform the sensorimotor transformation (Anastasio and Robinson 1990).

In another study, Xing and Andersen (2000) developed several three-layer network models of varying complexity; the task for all models was to produce the correct motor output in response

to the sensory input. The simplest of these models required no coordinate transformations while the most complicated of these models required both integration of multiple coordinate frames and coordinate transformations. In models where coordinate transformations were required, the coordinate spaces of the hidden layer neurons were distributed and did not correspond to any particular coordinate frame represented in the output array. In addition, the hidden units of the more complicated models activated output neurons of different coordinate frames simultaneously.

With a large number of neurons, neurons do not have to behave in a stringent manner in order to perform a coordinate transformation. A simple linear mapping may be sufficient for the necessary transformation (Sanger 1994). As shown by these studies (Anastasio and Robinson 1989, 1990; Xing and Andersen 2000) and our model, the characteristics of the activities of neurons (*i.e.* how activity varies with aspects of movement) does not necessarily lead to a causal relationship with the units the neurons command.

Two other studies use models to show that a linear transformation can accomplish a transformation similar to the one examined in our model. Kakei et al. (2003) present a simple model which shares some fundamental architectural features with our model. The model shows that the weighted sum of two cosines, each of the form $B\cos(\theta - PD)$ and with different PDs, produces a third cosine, the PD of which depends on the relative weighting of the first two cosines. The first two cosines represent two neurons that can be described as extrinsic-like modulated by wrist posture, while the third cosine represents a neuron that can be described as muscle-like. Kakei et al. (2003) suggest that this linear summation can account for the muscle-like neurons found in MI. In another modeling study, Salinas and Abbott (1995) hard-wired an array of sensory neurons (which behave similar to the extrinsic-like MI neurons in Kakei et al. 1999) and an array of motor neurons (which behave similar to the muscle-like MI neurons in Kakei et al. 1999) and used a Hebbian learning rule to modify the connections between the two. After training, the sensory array could directly produce the appropriate activation patterns in the motor array. We view the results of these studies as lending support to the parallel processing scheme argued for in this paper.

While our model investigates a similar transformation as those presented in Salinas and Abbott (1995) and Kakei et al. (2003), our approach allows for the use of an optimization procedure to select the motor output (muscle activation patterns in our model) as opposed to committing to a predefined representation. Salinas and Abbott (1995) and Kakei et al. (2003) assume that the output array follows a Gaussian or cosine form; our method allows for significant flexibility in selecting an arbitrary output pattern. However, in all cases, the simple, linear transformation is made possible by the sparse representation exhibited by the input and output representations. The sparse representation in our model is the result of the optimization procedure and was not explicitly included. In addition, we show how the muscle activation output of our model leads to endpoint of movement; this creates a plausible motor plant which allows us to use a performance-dependent optimization procedure.

The gradient descent method used in our model has certain advantages. While Kakei et al. (2003) show how the linear summation of two extrinsic-like neuron can produce an arbitrary muscle-like neuron, our model specifically finds how to weigh the contributions of the extrinsic-like neurons to

produce the appropriate muscle activation patterns. Salinas and Abbott (1995) employ a Hebbian-style learning rule which, in the limit, tends toward a linear transformation in which connection strength is predictable from input/output correlation. Unlike the learning rule employed in our model, this rule reduces the number of effective degrees-of-freedom (DOFs) that are exploitable in selecting a connection matrix. Although their learning rule does represent one that is more biologically plausible (in cortex) than the error-driven rule used in our model, it must presume the existence of an external module that is able to independently drive the sensory inputs and motor outputs correctly.

Effects of Noise

The introduction of signal dependent noise in our model leads to a harder problem and thus a reduction in the DOFs afforded by neural redundancy. While a transformation exists from the extrinsic-like neurons to muscle activation, there are more restrictions on this transformation. For example, because the total amount of noise to a muscle is minimized when averaged over many neurons, it is advantageous to recruit many rather than fewer neurons to provide the same input (*cf.* Todorov 2002). Thus, in the model with high neural noise, neurons with similar PDs will tend to activate a muscle the same way (correlation and connection strength will be related to some degree). However, in a high noise simulation with more MI neurons (192), and thus greater DOFs, the predictive relationship between $corr_{ij}$ and K_{ij} was weaker. In addition, to keep muscle activity from growing unreasonably (because neurons are cooperating more), the elements of $\mathbf{K}$ must be close to zero, also a property we observe in the model with high neural noise. While some of the noisy model behavior can be explained by a reduction in the DOFs, some aspects can also be explained by similarities with stochastic gradient search methods (Hassoun 1995; Hoptroff and Hall 1989), which intentionally introduce noise in gradient descent optimization. One of the advantages of this method is that it aids in finding a global minimum (Schoen 1991), seen in the unique $\mathbf{K}$s found in both of our models with noisy neurons.

Future Work

In order to focus on the issue of coordinate transformations, we kept the level of abstraction of our model very high. In the model presented in this paper and in Fagg et al. (2002), we make the assumption that all muscles pull with equal strength in their pulling directions. Fagg et al. (2002) discusses how muscle behavior changes when relative pulling strengths are changed. In short, the qualitative nature of muscle recruitment does not change. We also use an abstract representation of muscle, in which we do not include time-related aspects of muscle action, in our models. However, the temporal behavior of muscle recruitment is also important (Hoffman and Strick 1999). One area of future research with this model is to incorporate the temporal aspects of movement generation and the dependence on spinal and muscular dynamics (Houk et al. 2002).

Our model includes a well defined representation of movement based on neurons recorded in Kakei et al. (1999). Why do we see structure in neural activity at all? In continuing work, we have

adapted this model to include MI neurons that are not hard-wired to behave a certain way. We are examining the effects of local interactions (not included in the model presented in this paper) and optimization criteria (such as minimization of metabolic energy, Balasubramanian et al. 2001; Levy and Baxter 1996; Schreiber et al. 2002) in the formation of neural behavior. Preliminary results show that the local interactions and optimization criteria produce both extrinsic-like and muscle-like MI neurons which are active simultaneously during the wrist task.

Acknowledgments

The authors had very helpful discussions with Dr. Donna Hoffman, Dr. Peter Strick, Dr. Lee Miller, and Dr. Thomas Anastasio. Preparation for this report was supported by NIH Grant #NIH MH 48185-10, NSF Grant #EIA 9703217, and NASA Grant NAG9-1445.

References

Ajemian, R., Bullock, D., and Grossberg, S. (2000). Kinematic coordinates in which motor cortical cells encode movement direction. *Journal of Neurophysiology*, 84:2191–2203.

Amirikian, B. and Georgopoulos, A. P. (2000). Directional tuning profiles of motor cortical cells. *Neuroscience Research*, 36:73–79.

Anastasio, T. J. and Robinson, D. A. (1989). The distributed representation of vestibulo-oculomotor signals by brain-stem neurons. *Biological Cybernetics*, 61:79–88.

Anastasio, T. J. and Robinson, D. A. (1990). Distributed parallel processing in the vertical vestibulo-ocular reflex: Learning networks compared to tensor theory. *Biological Cybernetics*, 63:161–167.

Ashe, J. and Georgopoulos, A. P. (1994). Movement parameters and neural activity in motor cortex and area 5. *Cerebral Cortex*, 6:590–600.

Balasubramanian, V., Kimber, D., and Berry, M. J. I. (2001). Metabolically efficient information processing. *Neural Computation*, 13:799–815.

Caminiti, R., Johnson, P. B., Galli, C., Ferraina, S., and Burnod, Y. (1991). Making arm movements within different parts of space: The premotor and motor cortical representations of a coordinate system for reaching to visual targets. *The Journal of Neuroscience*, 11:1182–1197.

Caminiti, R., Johnson, P. B., and Urbano, A. (1990). Making arm movements within different parts of space: Dynamic aspects in the primate motor cortex. *The Journal of Neuroscience*, 10:2039–2058.

Dum, R. P. and Strick, P. L. (2002). Motor areas in the frontal lobe of the primate. *Physiology & Behavior*, 77:677–682.

Engel, A. K., Kónig, P., Kreiter, A. K., and Singer, W. (1991). Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science*, 252:1177–1179.

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

Fu, Q.-G., Flament, D., Coltz, J. D., and Ebner, T. J. (1995). Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *Journal of Neurophysiology*, 73:836–854.

Georgopoulos, A. P., Ashe, J., Smyrnis, N., and Taira, M. (1992). The motor cortex and the coding of force. *Science*, 256:1692–1695.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience*, 2:1527–1537.

Gibson, A. R., Houk, J. C., and Kohlerman, N. J. (1985). Relation between red nucleus discharge and movement parameters in trained macaque monkeys. *Journal of Physiology (London)*, 358:551–570.

Harris, C. M. and Wolpert, D. M. (1998). Signal dependent noise determines motor planning. *Nature*, 394:780–784.

Hassoun, M. H. (1995). *Fundementals of Artificial Neural Networks*, chapter 8. The MIT Press, Cambridge, MA.

Hoffman, D. S. and Strick, P. L. (1999). Step-tracking movements of the wrist. iv. muscle activity associated with movements in different directions. *Journal of Neurophysiology*, 81:319–333.

Hoptroff, R. G. and Hall, T. J. (1989). Learning by diffusion for multilayer perceptron. *Electronic Letters*, 25:531–533.

Houk, J. C., Dessem, D. A., Miller, L. E., and Sybirska, E. H. (1987). Correlation and spectral analysis of relations between single unit discharge and muscle activities. *Journal of Neuroscience Methods*, 21:201–224.

Houk, J. C., Fagg, A. H., and Barto, A. G. (2002). Fractional power damping model of joint motion. In Latash, M., editor, *Progress in Motor Control: Structure-Function Relations in Voluntary Movements*, pages 147–178. Human Kinetics, Champaign, IL.

Houk, J. C., Keifer, J., and Barto, A. G. (1993). Distributed motor commands in the limb premotor network. *Trends in Neuroscience*, 16:27–33.

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 area. *Nature Neuroscience*, 4:1020–1025.

Kakei, S., Hoffman, D. S., and Strick, P. L. (2003). Sensorimotor transformations in cortical motor areas. *Neuroscience Research*, 46:1–10.

Kalaska, J. F. and Crammond, D. J. (1992). Cerebral crotical mechanisms of reaching movements. *Science*, 255:1517–1523.

König, P. and Engel, A. K. (1995). Correlated firing in sensory-motor systems. *Current Opinion in Neurobiology*, 5:511–519.

Levy, W. B. and Baxter, R. A. (1996). Energy efficient neural codes. *Neural Computation*, 8:531–543.

Loeb, G. E., Brown, I. E., and Chen, E. J. (1999). A hierarchical foundation for models of sensorimotor control. *Experimental Brain Research*, 126:1–18.

Luppino, G., Matelli, M., Camarda, R., and Rizzolatti, G. (1994). Corticospinal projections from mesial fronal and cingulate areas in the monkey. *Neuroreport*, 5:2545–2548.

McKiernan, B. J., Marcario, J. K., Karrer, J. H., and Cheney, P. D. (1998). Corticomotoneuronal postspike effects in shoulder, elbow, wrist, digit, and intrinsic hand muscles during a reach and prehension task. *Journal of Neurophysiology*, 80:1961–1980.

Miller, L. E. and Houk, J. C. (1995). Motor co-ordinates in primate red nucleus: Preferential relation to muscle activation versus kinematic variables. *Journal of Physiology*, 488:533–548.

Miller, L. E. and Sinkjaer, T. (1998). Primate red nucleus discharge encodes the dynamics of limb muscle activity. *Journal of Neurophysiology*, 80:59–70.

Miller, L. E., van Kan, P. L. E., Sinkjaer, T., Andersen, T., Harris, G. D., and Houk, J. C. (1993). Correlation of primate red nucleus discharge with muscle activity during free-form arm movements. *Journal of Physiology*, 469:213–243.

Moran, D. W. and Schwartz, A. B. (1999). Motor cortical represenation of speed and direction during reaching. *Journal of Neurophysiology*, 82:2676–2692.

Munk, M. H. J., Nowak, L. G., Nelson, J. I., and Bullier, J. (1995). Structural basis of cortical synchronization ii. effects of cortical lesions. *Journal of Neurophysiology*, 74:2401–2414.

Mussa-Ivaldi, F. A. (1988). Do neurons in the motor cortex encode movement direction? an alternative hypothesis. *Neuroscience Letters*, 91:106–111.

Salinas, E. and Abbott, L. F. (1995). Transfer of coded information from sensory to motor networks. *The Journal of Neuroscience*, 15:6461–6474.

Sanger, T. D. (1994). Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation*, 6:29–37.

Schoen, F. (1991). Stochastic techniques for global optimization: A survey of recent advances. *Journal of Global Optimization*, 1:207–228.

Schreiber, S., Machens, C. K., Herz, A. V. M., and Laughlin, S. B. (2002). Energy-efficient coding with discrete stochastic movements. *Neural Computation*, 14:1323–1346.

Schwartz, A. B., Kettner, R. E., and Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual target in three-dimensional space. i. relations between single cell discharge and direction of movement. *The Journal of Neuroscience*, 8:2913–2927.

Schwartz, A. B. and Moran, D. W. (2000). Arm trajectory and representation of movement processing in motor cortical activity. *European Journal of Neuroscience*, 12:1851–1856.

Scott, S. H. (1997). Comparison of onset time and magnitude of activity for proximal arm muscles and motor cortical cells before reaching movements. *Journal of Neurophysiology*, 77:1016–1022.

Scott, S. H. (2003). The role of primary motor cortex in goal-directed movements: Insights from neurophysiological studies on non-human primates. *Current Opinion in Neurobiology*, 13:671–677.

Scott, S. H. and Kalaska, J. F. (1995). Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of Neurophysiology*, 73:2563–2567.

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. *Journal of Neurophysiology*, 77:826–852.

Scott, S. J. (2000). Role of motor cortex in coordinating multi-joint movements: Is it time for a new paradigm? *Canadian Journal of Physiology and Pharmacology*, 78:923–933.

Sergio, L. E. and Kalaska, J. F. (1997). Systematic changes in directional tuning of motor cortex cell activity with hand location in the workspace during generation of static isometric forces in constant spatial directions. *Journal of Neurophysiology*, 78:1170–1174.

Sergio, L. E. and Kalaska, J. F. (1998). Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *Journal of Neurophysiology*, 80:1577–1583.

Shah, A., Fagg, A. H., and Barto, A. G. (2002). Cortical involvement in the recruitment of wrist muscles. Naples, FL. Neural Control of Movement Conference.

Todorov, E. (2000). Direct cortical control of muscle activations in voluntary arm movements: a model. *Nature Neuroscience*, 3:391–398.

Todorov, E. (2002). Cosine tuning minimized motor errors. *Neural Computation*, 14:1233–1260.

van Beers, R. J., Haggard, P., and Wolpert, D. M. (2003). The role of execution noise in movement variability. *Journal of Neurophysiology*, (in press).

Xing, J. and Andersen, R. A. (2000). Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *Journal of Cognitive Neuroscience*, 12:601–614.

Zhang, K. and Sejnowski, T. J. (1999). Neuronal tuning: To sharpen or broaden? *Neural Computation*, 11:75–84.