

## **On the role of primary motor cortex in arm movement control**

Emanuel Todorov

Department of Cognitive Science  
University of California, San Diego

September 11, 2002

### **Abstract**

The role of primary motor cortex (M1) in arm movement control has been the subject of lasting and often heated debates. Early views of low-level muscle control have been discredited by numerous correlations with higher-level parameters related to hand kinematics. The abstract directional coding hypothesis advanced to explain these observations has also been discredited, this time by numerous correlations with lower-level muscle- or joint-related parameters. Taken at face value, the contradictions accumulated over the years may seem to suggest that no coherent account of M1 function can ever be given, and one has to accept a catalog of disparate empirical findings in place of an explanation. This is in principle possible, but, at the risk of disagreeing with some of the neurophysiologists whose findings we find it necessary to re-interpret, we prefer the following more optimistic approach: assume that many of the observed phenomena are in fact epiphenomena, and search for simple underlying mechanisms that give rise to the seemingly contradicting results. To uncover such mechanisms, we need models.

The present chapter summarizes a series of recent studies as well as unpublished results aimed at building a coherent model of M1. A coherent model requires an answer to the following question: if the M1 population output encodes just one thing, what could that thing be? The model we arrive at is related to earlier views of muscle-based control. It accounts for later results by incorporating the properties of the musculo-skeletal system: M1 neurons are assumed to activate groups of muscles while taking into account the state-dependence of muscle force production and effects of multijoint mechanics. The M1 signal needed to compensate for these peripheral properties turns out to be correlated with behavioral parameters on multiple levels of description, quantitatively explaining many results which appear unrelated to the pattern of muscle activation. The chapter ends with an extended discussion of several important issues (origins of directional tuning, muscle synergies, redundancy of population codes, spinal cord function, and equilibrium-point control) from the point of view of our model.

## 1. Introduction

Primary motor cortex (M1) plays a fundamental role in the control of voluntary arm movements, as evidenced by the profound deficits following M1 lesions. Precisely what that role is has been the subject of lasting debates. These debates are fueled by the numerous correlations found between the activity of M1 neurons and various behavioral parameters.

The earliest experiments in awake behaving monkeys (Evarts, 1968) showed that M1 firing is better correlated with isometric force than limb position. Given the dense projection from M1 to the spinal cord, sometimes directly to motoneurons, it is natural to interpret these early observations as evidence that M1 controls muscle force or activation (Evarts, 1981). It was later observed, however, that in multijoint movement tasks the majority of M1 neurons encode not the acceleration (which is proportional to force), but the direction (Georgopoulos et al., 1982) or velocity (Schwartz, 1994) of hand movement. In addition the M1 directional tuning curves were quite broad, which, in the context of single-joint studies that dominated the field at the time, seemed inconsistent with low-level muscle control. These observations, and the ability of the population vector (PV) to recover movement direction, led to the opposing view that M1 sends a much more abstract signal specifying a three-dimensional vector in extrapersonal space.

To distinguish between the two alternatives, movement velocity and external load were varied separately in the same experiment (Kalaska et al., 1989). The results however supported both views: velocity- and load-related signals were both present and combined additively in the activity of the same neurons. Although the directional coding hypothesis can be extended to incorporate the observed tuning for load (Kalaska et al., 1989; Georgopoulos et al., 1992; Taira et al., 1996) as well as position (Georgopoulos et al., 1984; Kettner et al., 1988), it offers no explanation as to why all these parameters should be encoded simultaneously. Furthermore, many additional correlations have been observed (see (Todorov, 2000a) that do not seem to fit in either view of M1 function. One can use regression analysis to identify the multiple parameters correlated with the activity of each neuron (Ashe and Georgopoulos, 1994; Fu et al., 1995; Taira et al., 1996); the activity of any one neuron is of course correlated with some parameters better than others, but the presence of simultaneous correlations makes it impossible to determine what that neuron really encodes (Fetz, 1992).

Is it then possible to make any coherent statement regarding the role of M1 in arm movement control, at least to a first approximation? Is there *one* parameter of motor behavior which, if encoded in the M1 population, would explain most available results? We argue that of all the parameters previously considered, the only one that can provide the basis for a coherent model of M1 is muscle activation. The present chapter summarizes a series of recent studies (Todorov, 2000a; Todorov et al., 2000; Todorov,

2002) as well as unpublished work in which we have developed a model of direct cortical control of muscle activation.

### **1.1. The case for muscle-based encoding**

If one were to approach the M1 literature without a preexisting bias and search for the single idea that is most likely to lead to a simple yet successful quantitative model, which idea should one choose? We found the idea of muscle-based encoding to be the most promising candidate, for the following reasons:

***Evidence for muscle control.*** M1 neurons and muscles display similar patterns of directional variation in their onset times and activation magnitudes (Scott, 1997). The distributions of preferred directions of both M1 neurons and muscles are roughly uniform in the natural posture, but become substantially elongated (along the same axis) when the arm is constrained to the horizontal plane (Scott and Kalaska, 1997; Scott et al., 2001). Increased muscle cocontraction is accompanied by increased M1 activity (Humphrey and Reed, 1983); see also Section 3.6. M1 and muscle activity are synchronized in certain frequency bands (Conway et al., 1995)

***Evidence against directional coding.*** Several studies have found systematic differences in M1 activity between experimental conditions with identical movement directions in extrapersonal space but different joint configurations (Caminiti et al., 1991; Scott and Kalaska, 1995; Sergio and Kalaska, 1998; Kakei et al., 1999). Such results clearly contradict the hypothesis of directional coding in extrapersonal space. The changes of M1 activity observed in most of these experiments are consistent with both muscle- and joint-based encoding, except for the wrist movement task used by (Kakei et al., 1999) which quite convincingly ruled out a joint-based encoding (see also Section 7.3). Interestingly, neural activity in ventral premotor cortex (PMv) was not affected by arm posture (Kakei et al., 2001), suggesting that the directional coding hypothesis is much more applicable to PMv than M1.

***Direct projections.*** Some M1 neurons (labeled cortico-motoneuronal, or CM, neurons) project monosynaptically to motoneurons (Fetz and Cheney, 1980; Lemon et al., 1986). If the M1 population encoded just one parameter and that parameter was not closely related to muscle activation, motoneurons would not be able to "decode" the descending signal correctly and this direct pathway would have no reason to exist. Note that although monosynaptic projections are a minority, there is an evolutionary trend for their percentage to increase – being most prominent in the hand muscles of humans. Therefore the existence of this pathway is not an "accident" that can be ignored.

***Proprioceptive feedback.*** Most M1 neurons are affected at short latencies by proprioceptive feedback, so that imposed movement in the preferred direction suppresses the activity of the neuron in a stretch-reflex-like manner (Evarts, 1981). The existence of this pathway indicates that the M1 output is compatible with such feedback, i.e. adding the two signals on the level of individual neurons produces a

meaningful quantity. It is hard to imagine how the combination of abstract directional signals encoded in extrapersonal space, and proprioceptive feedback encoded in a muscle- or joint-related frame of reference, could be useful or even interpretable.

***Evolutionary considerations.*** One would expect the general function of the spinal cord in primates, to the extent that it is used, to be related to functions it performed long before any supraspinal systems existed. The supraspinal systems most likely evolved to enhance those functions, by modulating the state of the spinal interneuronal networks in a way that helped them solve whatever problems they were already trying to solve (Loeb, 1999). The interpretation of abstract motor commands was not among these problems. A descending signal which modulates the state of the spinal networks should be compatible with the representation they use, and that representation is likely to be closely related to the pattern of muscle activation.

***Computational considerations.*** While the motor system is clearly hierarchical, the division of labor implicit in the directional coding hypothesis is unlikely. According to that hypothesis the vast supraspinal system is dedicated to the rather trivial calculation of a three-dimensional vector, while all the difficult control problems (e.g. coordinating redundant nonlinear actuators, dealing with complex interaction forces, ensuring stability in the presence of delays and uncertainty) are left to the ancient spinal circuitry.

## **1.2. Model overview**

As explained above there are plenty of reasons to adopt a muscle-based model of M1. The question is, how can such a model account for the correlations with various behavioral parameters that seem unrelated to muscle activation? The main contribution of our work is in answering that question. Here we outline the features of the model and explain intuitively how they help us account for the many disparate results.

***M1 output.*** We focus exclusively on M1 neurons projecting to the spinal cord. Interneurons and output neurons projecting to other areas will not be modeled, since most of the arguments from Section 1.1 do not apply to them. The activity of such neurons can be systematically different, and so by focusing on the M1 output we avoid some of the response diversity that a more complete model would have to explain.

***M1-to-behavior mapping.*** We are only interested in the instantaneous activity of the M1 output neurons, and its causal relationship to the instantaneous motor behavior at some later point in time. All predictions regarding M1 will be derived by "inverting" that causal relationship (see Fig 1A). Motor planning, sensorimotor transformations, and transcortical feedback corrections are all enclosed in the black box labeled "brain" (Fig 1A) and do not concern us here. The "brain" receives online feedback,

takes into account all goals, plans, expectations, preferences, etc. and somehow generates the M1 output. This pathway will not be used to draw any predictions, and so we do not need to model it.

**Population analysis.** Our goal is to explain average population (or subpopulation) responses rather than individual response types (Kalaska et al., 1989). Because of the redundancy in the M1-to-muscle mapping, the same pattern of muscle activity can be generated with an infinite family of M1 activation patterns. How the brain chooses one of them is an interesting question which we address separately in Sections 6 and 7. For the most part however we will not distinguish between M1 activity patterns that result in the same EMG pattern.

**Simplifications.** We introduce a number of simplifications that allow us to avoid curve fitting and adhoc assumptions about poorly understood mechanisms (see Section 2). In particular, we use linear approximations to the spinal circuitry, muscle force production mechanisms, and kinematics of the multijoint arm. The effects of other descending systems are ignored, although there is evidence (Miller and Sinkjaer, 1998) that rubrospinal projections can be modeled similarly. Sensory feedback through spinal pathways is assumed to modulate the muscle activity generated by descending signals (see Section 7). While such feedback modulation is very important for achieving stability, its effect is washed out in the trial-averaged analyses used in the M1 literature, and so we do not model it explicitly here.

**Muscle state dependence.** It is well established that for a fixed activation level, the force generated by a muscle varies with its length and velocity (Joyce et al., 1969; Zajac, 1989; Brown et al., 1999). The force-length-velocity surface at maximum (tetanic) activation is illustrated in Fig 1B, using the state-of-the-art muscle model of (Brown et al., 1999); details in Table 1 and Section 2. Note that for the same activation level and within the physiological range of muscle lengths and velocities, the force output varies from almost 0 to above maximum isometric force (defined as 1). Clearly the neural circuits controlling muscle activation have to take this state dependence into account if the resulting behavior is to resemble the desired one even remotely. The M1 signal compensating for the state dependence illustrated in Fig 1B plays a central role in our model since it correlates with position and velocity – and thus explains the observations that appear most problematic.

## **2. From M1 output to motor behavior**

The relationship between M1 output and motor behavior is no doubt complex. If we incorporate every detail whose existence is currently suspected the resulting model will probably have enough free parameters to fit every available observation. The usefulness of a model, however, comes not from incorporating all proposed mechanisms and fitting all existing data, but from illuminating the underlying principles whose interplay gives rise to, or at least can give rise to the data. Our strategy is to avoid free parameters and the curve fitting associated with them. We do so with the help of simplifying assumptions:

linearizing functions that are mildly nonlinear, assuming quantities that vary slightly to be constant, ignoring terms that tend to average out, linking together parameters that have similar values. Every simplification is justified when first introduced, with the exception of cosine tuning which is addressed separately in Section 6.

It is important to realize that the simplifications themselves do not give rise to the phenomena we are trying to explain. Therefore our ability to reproduce these phenomena indicates that they arise from the underlying principle behind the model: direct cortical control of muscle activation taking into account the basic properties of the musculo-skeletal system.

## 2.1. From M1 output to muscle activation

Let  $t$  denote time,  $c_j(t)$  the mean firing rate of the  $j_{th}$  pyramidal tract neuron relative to its baseline,  $a_i(t)$  the EMG activity of the  $i_{th}$  muscle relative to its baseline,  $w_{ij}$  the net effect that neuron  $j$  has on muscle  $i$ , and  $d$  the average delay between movement-related activity in M1 and movement onset<sup>1</sup>. Baseline is defined as the mean activity during posture in the center of the workspace. To relate M1 output to muscle activation we will use the first-order model

$$a_i(t) = \sum_j w_{ij} c_j(t-d) + \varepsilon_i(t) \quad (1)$$

where  $\varepsilon_i(t)$  denotes the effects of other descending systems and spinal feedback. We will not impose any restrictions on the pattern of weights  $w_{ij}$ , i.e. each neuron can excite or inhibit any combination of muscles. The weights  $w_{ij}$  represent not only monosynaptic projections from pyramidal tract neurons to motoneurons, but a linear approximation to the input-output behavior of the spinal cord. Similarly the delay  $d$  corresponds not only to transmission delays, but approximates the net effect of any low-pass filtering or recurrent processing in the spinal cord.

Note that Eq 1 will be used to describe the relationship between M1 and muscle activity during natural voluntary behaviors such as reaching or pushing. During instructed delay periods, sleep, or after training the subject to control a cursor with signals recorded in M1 (Helms Tillery et al., 2001), the

---

<sup>1</sup> Notation:  $x$  is a scalar,  $\mathbf{x}$  is a column vector,  $\mathbf{x}^T$  is a transposed (row) vector,  $\mathbf{x}^T \mathbf{y}$  is the dot-product of  $\mathbf{x}$  and  $\mathbf{y}$ ,  $X$  is a matrix,  $\dot{\mathbf{x}}$  and  $\ddot{\mathbf{x}}$  are the temporal derivatives of  $\mathbf{x}$ ,  $\|\mathbf{x}\|$  is the vector length,  $|X|$  is the matrix determinant,  $X^{-1}$  is the matrix inverse,  $\lfloor x \rfloor = x$  for  $x > 0$  and 0 otherwise,  $\delta(k,s) = 1$  for  $k = s$  and 0 otherwise.

descending activity may somehow be gated at the spinal level. Anatomically the M1 output is not the only descending source of activation: there are other descending systems, and the pyramidal tract itself originates from multiple cortical areas (Dum and Strick, 1991). It is not yet known, however, whether the pyramidal tract neurons in premotor and supplementary motor areas behave more similarly to the mixed population in their respective area, or to the pyramidal tract neurons in M1 (in which case the present model applies to them as well). Recordings from the red nucleus (Miller and Sinkjaer, 1998) indicate that Eq 1 may be a reasonable approximation to the effects of other descending systems. It should also be noted that lesions in M1 produce more profound deficits than lesions in any other supraspinal motor area (Johnson, 1992), further justifying our emphasis on M1. As discussed above and in Section 7, the spinal feedback contributions are assumed to average out over repeated trials.

## 2.2. From muscle force to endpoint force

Let  $f_i(t)$  be the force (tension) that muscle  $i$  generates, and  $\mathbf{e}_i(t)$  the corresponding endpoint force in the two-dimensional workspace of the hand. While the direction of  $\mathbf{e}_i(t)$  depends on joint configuration, that dependence (see Fig 1A) is negligible for the small workspaces and stereotyped postures adopted in M1 experiments, so we will ignore it. Then

$$\mathbf{e}_i(t) = \mathbf{p}_i f_i(t) \quad (2)$$

where  $\mathbf{p}_i$  is a constant 2D vector whose length determines the scaling from muscle force to endpoint force. Note that the present formulation ignores the effects of muscle forces causing redundant joint rotations. If we take this redundant subspace into account as outlined in (Todorov, 2000a), we could derive additional constraints on the M1 population activity (without affecting the constraints derived below). However M1 data has traditionally been analyzed in endpoint space, so these additional constraints could not be compared to existing experimental results.

## 2.3. Arm dynamics in endpoint space

Let  $M$  be the 2x2 matrix expressing arm inertia in endpoint space,  $\mathbf{f}_{\text{ext}}(t)$  the endpoint force that the hand applies against external objects, and  $\mathbf{x}(t)$  the hand position in endpoint space (the origin of the 2D coordinate system is defined at the center of the workspace). We will assume that changes in inertia are negligible for the small movement extents and velocities used in the M1 literature, and approximate the arm dynamics with an anisotropic point mass. Making explicit the dependence of muscle force on activation, length and velocity (the latter two being functions of  $\mathbf{x}$  and  $\dot{\mathbf{x}}$ ), the endpoint dynamics is

$$\sum_i \mathbf{p}_i f_i(a_i(t), \mathbf{x}(t), \dot{\mathbf{x}}(t)) = \mathbf{f}_{\text{ext}}(t) + M\ddot{\mathbf{x}}(t) \quad (3)$$

We now encounter a redundancy problem: there are many more muscle activations  $a_i$  (and even more cell activations  $c_j$ ) than constraints given by Eq 3. We can proceed in two ways: a) linearize the left hand side of Eq 3 to obtain a first-order approximation to the M1-to-behavior mapping; b) introduce further assumptions that resolve redundancy and allow us to use the detailed muscle model illustrated in Fig 1B. We will pursue the first approach (described in Section 2.4A) throughout this chapter. The second approach (described in Section 2.4B) will only be used when conclusions reached on the basis of the linearization might seem sensitive to details of the muscle model.

### 2.4A. Using a linearized muscle model

Assembling all  $a_i$ 's in the vector  $\mathbf{a}$ , all  $c_j$ 's in the vector  $\mathbf{c}$ , all  $f_i$ 's in the vector  $\mathbf{f}$ , all  $w_{ij}$ 's in the matrix  $W$ , and all  $\mathbf{p}_i$ 's in the columns of the matrix  $P$ , Eq 3 can be written in matrix form

$$P\mathbf{f}(\mathbf{a}, \mathbf{x}, \dot{\mathbf{x}}) = \mathbf{f}_{\text{ext}} + M\ddot{\mathbf{x}} \quad (3a)$$

Linearizing  $\mathbf{f}$  around the baseline  $\mathbf{f}(\mathbf{0}, \mathbf{0}, \mathbf{0}) = \mathbf{0}$  yields an approximation to the left hand side of Eq 3a:  $P\mathbf{f}(\mathbf{a}, \mathbf{x}, \dot{\mathbf{x}}) = PS\mathbf{a} - K\mathbf{x} - B\dot{\mathbf{x}}$ . Here  $K$  and  $B$  are the endpoint stiffness and damping by definition, and the scaling matrix  $S = \frac{\partial \mathbf{f}}{\partial \mathbf{a}}$  is diagonal because each muscle is only affected by its own activation. From

Eq 1 we have  $PS\mathbf{a} = PSW\mathbf{c}$ , and so  $PSW\mathbf{c} = \mathbf{f}_{\text{ext}} + M\ddot{\mathbf{x}} + B\dot{\mathbf{x}} + K\mathbf{x}$ .

We will now approximate  $PSW$  with  $FU$ , where the  $j_{th}$  column vector  $\mathbf{u}_j$  of the matrix  $U$  is a 2D unit vector pointing in the same direction as the  $j_{th}$  column vector of  $PSW$ . Thus the  $j_{th}$  neuron activates a linear combination of muscles generating endpoint force in the direction  $\mathbf{u}_j$ . The 2x2 matrix  $F = PSWU^T(UU^T)^{-1}$ , chosen to minimize the norm of the approximation error  $PSW - FU$ , can be thought of as fitting an ellipse to the direction-dependent variation of the lengths of the column vectors in  $PSW$ . Since the scaling from firing rates to forces is arbitrary, we can assume without loss of generality that  $|F| = 1$ . Summarizing the derivation so far

$$FU\mathbf{c}(t-d) = \mathbf{f}_{\text{ext}}(t) + M\ddot{\mathbf{x}}(t) + B\dot{\mathbf{x}}(t) + K\mathbf{x}(t) \quad (4)$$

The 2x2 matrices  $M, B, K, F$  all reflect the geometry of the arm, in roughly the same way. Consider a 4-degree-of-freedom arm, with the elbow pointing down and the hand constrained to a



horizontal plane at about shoulder level. Endpoint inertia is clearly larger along the forward-backward y-axis than the lateral x-axis. Stiffness and damping are known to be larger along the hand-shoulder axis (Tsuji et al., 1995), which in this configuration coincides with the y-axis. Assuming that the weights  $W$  and the activation-to-force scaling terms  $S$  do not vary systematically with direction, the only anisotropy in  $F$  comes from the matrix  $P$  and is due to the fact that generating endpoint force along the y-axis requires joint torques of smaller magnitude<sup>2</sup>. Then the ellipse corresponding to  $F$  is also elongated along the y-axis, and since the aspect ratios of the four ellipses are similar we will assume the four matrices to be proportional to each other:  $M = mF, B = bF, K = kF$ . Since  $|F| = 1$ , the scalars  $m, b, k$  are the direction-averaged inertia, damping and stiffness. Multiplying Eq 4 by  $F^{-1}$  yields

$$U\mathbf{c}(t-d) = F^{-1}\mathbf{f}_{\text{ext}}(t) + m\ddot{\mathbf{x}}(t) + b\dot{\mathbf{x}}(t) + k\mathbf{x}(t) \quad (5)$$

Eq 5 is the relationship between M1 population activity  $\mathbf{c}$  and motor behavior  $\mathbf{f}_{\text{ext}}, \ddot{\mathbf{x}}, \dot{\mathbf{x}}, \mathbf{x}$  that we wanted to derive. The left hand side is the sum of neuronal activities  $c_j$  multiplied by the corresponding vectors  $\mathbf{u}_j$ , or in other words a Population Vector (Georgopoulos et al., 1983). Note that unlike the PV methods used in the M1 literature, the vector  $\mathbf{u}_j$  we assign to neuron  $j$  is a force direction rather than a physiological preferred direction. These two directions however are similar (Fetz and Cheney, 1980). Even if the difference between them is substantial for some neurons, the PV defined in the M1 literature is an unbiased estimate of the quantity  $U\mathbf{c}$  as long as that difference is zero on average. Once we fix the parameters  $m, b, k$ , and the aspect ratio of  $F$ , we can generate predictions regarding population vectors *without* resolving the redundancy in Eq 5. The parameter values we use everywhere:  $m = 1 \text{ kg}$ ,  $b = 10 \text{ Ns/m}$ ,  $k = 50 \text{ N/m}$ ,  $F$  aspect ratio = 2:1, are derived from the human psychophysics literature (see (Todorov, 2000a)). Since the model predictions depend on the relative rather than absolute magnitudes of  $m, b, k$ , scaling these constant down to values appropriate for the monkey arm should not affect our conclusions.

The constraint given by Eq 5 is sufficient to model the PV, but the redundancy mentioned earlier has not yet been resolved (i.e. Eq 5 does not tell us how individual neurons behave). Since several of the phenomena we want to address require such a model, we augment Eq 5 with the assumption of cosine tuning and obtain individual neuron responses

---

<sup>2</sup> In other words, the Jacobian relating endpoint force to joint torque is anisotropic.

$$c_j(t-d) = \frac{\mathbf{u}_j^T}{2} \left( F^{-1} \mathbf{f}_{\text{ext}}(t) + m\ddot{\mathbf{x}}(t) + k\mathbf{x}(t) \right) + b \left[ \mathbf{u}_j^T \dot{\mathbf{x}}(t) \right] \quad (6)$$

This equation is essentially the dot-product of  $\mathbf{u}_j$  with the right hand side of Eq 5 (resulting in cosine tuning). The reason we treat the damping term differently is that muscle damping is known to be asymmetric – predominantly present during shortening (see Fig 1B). The truncated cosine term  $\left[ \mathbf{u}_j^T \dot{\mathbf{x}}(t) \right]$  implies that neuron  $j$  compensates for damping only when the hand is moving within 90 deg of its force direction  $\mathbf{u}_j$  (corresponding to shortening of the muscles it activates).

Note that Eq 6 prescribes an “ideal” response which does not capture the diversity of real neuronal responses (Kalaska et al., 1989), and should be thought of as the average response of all cells whose force directions are close to  $\mathbf{u}_j$ . To obtain an explicit model of response variability, we replace the endpoint impedance parameters  $m, b, k$  in Eq 6 with cell-specific  $m_j, b_j, k_j$ . These constants are sampled independently from uniform distributions between zero and twice the average value:  $m_j \in [0; 2]$ ,  $b_j \in [0; 20]$ ,  $k_j \in [0; 100]$ . Multiplying by 2 and adding a cell-specific baseline  $\bar{c}_j$  sampled uniformly from  $\bar{c}_j \in [0; 34]$  yields an average response of 45 Hz in the preferred direction and 5Hz in the opposite direction, which is close to the range observed experimentally (Crammond and Kalaska, 1996).

Eq 6 can also be used to analyse 1-degree-of-freedom movements, or to compute the average firing of all neurons in their preferred direction ( $\mathbf{u}_A$ ) and the direction opposite to it ( $\mathbf{u}_N$ ). In either case there are only two directions of interest ( $\mathbf{u}_A$  and  $\mathbf{u}_N$ ), which simplifies Eq 6 to

$$\begin{aligned} c_A(t-d) &= \frac{1}{2} \left( f_{\text{ext}}(t) + m\ddot{x}(t) + kx(t) \right) + b\dot{x}(t) \\ c_N(t-d) &= -\frac{1}{2} \left( f_{\text{ext}}(t) + m\ddot{x}(t) + kx(t) \right) \end{aligned} \quad (7)$$

## 2.4B. Using a detailed muscle model

We now return to Eq 3 and rederive the M1-to-behavior mapping using the detailed muscle model illustrated in Fig 1B. That model is identical to (Brown et al., 1999), except for the effects of sag, yielding, low-pass filtering and tendon elasticity which we ignore. The muscle force  $f(a, l, v)$  given in Table 1 is a function of activation  $a$ , length  $l$ , and velocity  $v$ . All parameters are set to the average values for fast-twitch and slow-twitch muscles, which account for data from a wide range of experiments (Brown

et al., 1999). The activation  $a$  is normalized so that  $a = 0$  corresponds to the recruitment threshold and  $a = 1$  to the frequency at which the force output saturates. Length  $l$  and velocity  $v$  are expressed in units of  $l_0$  and  $l_0/\text{sec}$ , where  $l_0$  is the length at which maximum isometric force is generated. Force output  $f$  is in units of maximum isometric force.

The results from this section will only be applied to movement tasks where the directional asymmetries cancel each other (see Eq 5), so Eq 3 becomes

$$e_{\max} \sum_i \mathbf{p}_i f_i(a_i(t) + \bar{a}_i, l_i(\mathbf{x}(t)), v_i(\dot{\mathbf{x}}(t))) = m\ddot{\mathbf{x}}(t) \quad (8)$$

The vectors  $\mathbf{p}_i$  now have unit length,  $l_i$  and  $v_i$  denote the length and velocity of muscle  $i$ ,  $\bar{a}_i$  is the baseline activity corresponding to posture at the center of the workspace, and  $e_{\max}$  is the maximum isometric force each muscle can contribute in endpoint space. The baseline  $\bar{a}_i$  is needed here because the muscle model is nonlinear and requires the total muscle activation.

The direction of hand movement for which muscle  $i$  shortens most rapidly is very close to the direction  $\mathbf{p}_i$  in which it produces endpoint force, and its orientation varies little over a small workspace (see Fig 1A). We will approximate<sup>3</sup> this direction as being constant and equal to the (constant) line of action  $\mathbf{p}_i$ . Then the length  $l_i$  is proportional to the projection  $\mathbf{p}_i^T \mathbf{x}$  of the endpoint position  $\mathbf{x}$  on  $\mathbf{p}_i$ . The relationship between endpoint and muscle kinematics becomes

$$\begin{aligned} l_i &= r_i - h\mathbf{p}_i^T \mathbf{x} \\ v_i &= -h\mathbf{p}_i^T \dot{\mathbf{x}} \end{aligned} \quad (9)$$

The negative sign is due to the fact that muscles shorten when the hand moves in the direction in which they produce force. The constants  $r_i$  specify the muscle lengths when the hand is at the center  $\mathbf{x} = (0,0)$  of the workspace, and  $h$  is a scaling constant that relates endpoint distance to muscle length.

To resolve the redundancy in Eq 8 we will again use the assumption of cosine tuning. The total force  $m\ddot{\mathbf{x}}$  that needs to be generated will be distributed among individual muscles as

$$dis_i(\ddot{\mathbf{x}}) = \frac{1}{n e_{\max}} \left[ co + 2m\mathbf{p}_i^T \ddot{\mathbf{x}} \right] \quad (10)$$

where  $n$  is the number of muscles, the cocontraction level  $co$  is the scalar sum of the magnitudes of endpoint forces contributed by individual muscles, and the dot-product  $\mathbf{p}_i^T \ddot{\mathbf{x}}$  corresponds to cosine

---

<sup>3</sup> This approximation is very similar to the linearization used in (Mussa-Ivaldi, 1988).

tuning. For a uniform distribution of  $\mathbf{p}_i$ 's and  $co$  large enough to ensure that the functions  $dis_i(\ddot{\mathbf{x}})$  are full cosines, it can be verified that setting  $f_i = dis_i(\ddot{\mathbf{x}})$  satisfies Eq 8.

Given motor behavior  $\ddot{\mathbf{x}}, \dot{\mathbf{x}}, \mathbf{x}$ , we can use Eq 9, 10 to compute muscle forces  $f_i = dis_i(\ddot{\mathbf{x}})$ , lengths  $l_i(\mathbf{x})$ , and velocities  $v_i(\dot{\mathbf{x}})$ . Let  $act(f_i, l_i, v_i)$  be the activation level for which a muscle at length  $l_i$  and velocity  $v_i$  generates force  $f_i$ . This function is uniquely defined because muscle force increases monotonically with activation for fixed length and velocity. The function  $act$  is found by inverting the muscle model  $f$  given in Table 1, i.e. solving numerically the equation  $f(act, l_i, v_i) = f_i$  with respect to  $act$ . The baselines, and time-varying activations relative to the baselines can now be computed

$$\begin{aligned}\bar{a}_i &= act\left(\frac{co}{ne_{\max}}, 0, 0\right) \\ a_i(t) &= act(dis_i(\ddot{\mathbf{x}}(t)), l_i(\mathbf{x}(t)), v_i(\dot{\mathbf{x}}(t))) - \bar{a}_i\end{aligned}\tag{11}$$

Eq 11 is the relationship between muscle activation and motor behavior that we wanted to derive. The population vector computed from muscle activations is  $\sum_i \mathbf{p}_i a_i(t)$ . Under the assumption that weight magnitudes  $|w_{ij}|$  do not vary systematically with direction (see Section 2.4A), it can be verified that the population vector computed from cell firing rates is identical (apart from the delay  $d$ ).

The parameters of the model were set as follows: mass  $m = 1\text{ kg}$ , number of muscles  $n = 10$ , maximum endpoint isometric force per muscle  $e_{\max} = 30\text{ N}$ , scaling constant  $h = 0.6$  mapping a 1m range of motion in  $\mathbf{x}$  to a  $[0.7; 1.3]$  physiological range of normalized muscle lengths. The cocontraction level was adjusted to  $co = 30\text{ N}$ , which resulted in empirical stiffness  $k = 76\text{ N/m}$  and damping  $b = 9\text{ Ns/m}$ . The empirical stiffness and damping were identified via perturbation experiments applied to the model. Note that unlike the linearized model in Section 2.4A, stiffness and damping cannot be set independently here. Instead they are implicitly determined by the time-varying muscle activations computed in Eq 11. The muscle lengths  $r_i$  at the center of the workspace were chosen randomly in the interval  $[0.9; 1.1]$ . Ten sets of random  $r_i$ 's were generated and the results averaged over the 10 simulation runs.

### 3. Basic results

Many experiments have been designed to identify what parameter of motor behavior is being represented in M1. Instead of a clear-cut answer, however, this effort has produced a long list of correlations. Our

model predicts that such a list should indeed accumulate, because muscle activation is correlated with all the candidate representations that have been considered.

### 3.1. Position

The postural activity of individual neurons predicted by the model (Eq 6) is illustrated in Fig 2A, for 8 hand positions arranged in a circle around the center of the workspace. The gradient of the response surface is oriented along the neuron's preferred direction, and the slope is determined by the amount of stiffness that needs to be compensated. Such monotonically increasing, and roughly linear postural responses have been well documented (Georgopoulos et al., 1984;Kettner et al., 1988).

### 3.2. Load

The predicted M1 activity in isometric force conditions (Fig 2B) is a linear function of the force being generated. In the 1-dof case (Eq 7) the response increases linearly with force magnitude for neurons whose force vector  $\mathbf{u}$  is within 90 deg of the isometric force  $\mathbf{f}_{\text{ext}}$ , and decreases otherwise. In the 2-dof case (Eq 6) the response varies linearly with the cosine of the angle between  $\mathbf{u}$  and  $\mathbf{f}_{\text{ext}}$ . Such monotonically increasing, and roughly linear load responses have been well documented (Fetz and Cheney, 1980;Kalaska et al., 1989).

### 3.3. Movement

The average activity of all M1 neurons in their preferred and opposite directions (Eq 7) is illustrated in Fig 2C, for a typical 10 cm movement with 500 msec duration and bell-shaped speed profile. The M1 signal predicted from the linearized muscle model (left) is the sum position, velocity, and acceleration terms weighted by the direction-averaged stiffness, damping, and inertia. When the detailed muscle model is used instead (Eq 11) the individual contributions of these terms can no longer be separated. The composite M1 signal however remains similar (right).

The predicted movement response closely resembles the speed profile and varies with direction – explaining why the activity of M1 neurons has often been interpreted as an encoding of movement direction (Georgopoulos et al., 1982) or velocity (Schwartz, 1994). In our model this phenomenon occurs simply because the damping that needs to be compensated dominates the other terms making up the M1 response. Note that the time-varying fluctuations relative to baseline are different in the preferred and opposite directions, which is a direct consequence of the fact that muscle damping is asymmetric. The asymmetry seen in Fig 2C has been observed experimentally (Kalaska et al., 1989;Crammond and Kalaska, 1996;Moran and Schwartz, 1999b). In comparison to the predicted response during an isometric

force ramp, the signal in Fig 2C has a more elaborate temporal structure – in agreement with experiments comparing movement and isometric force conditions (Sergio and Kalaska, 1998).

One might wonder why the predictions in Fig 2C do not resemble the classic triphasic burst pattern of muscle activity. The reason is that this pattern is typical for large ballistic movements (Brooks, 1986), which are much faster than the movements studied in the M1 literature. If we apply our model (Eq 7) to a movement which is  $x$  times faster, the magnitude of the position component remains the same, the velocity component scales by  $x$ , and the acceleration component scales by  $x^2$ . Thus for faster movements the model predicts that the acceleration component of the M1 signal will become more pronounced, causing PV reversals consistent with the triphasic burst pattern.

One might ask whether maintaining a position away from the center always requires a substantial increase of agonist activity (as shown in Fig 2C). This is indeed a side-effect of the linearized model where muscle stiffness is constant, but is no longer true in the detailed muscle model. As Fig 1B shows, the stiffness (i.e. partial derivative in the “length” direction) of realistic muscles is activation-dependent. Therefore the peripheral position can be maintained by substantially lower muscle activations. In the model, reducing muscle cocontraction is accomplished by reducing the cocontraction parameter  $co$  in Eq 10 and 11, and the prediction is that the overall activity in M1 will also decrease. The relation between overall M1 activity and muscle cocontraction is further discussed in Section 3.6.

### 3.4. Movement with external load

If a constant load is added during movement, the model prediction (Eq 6) is the sum of the corresponding movement and load responses – as observed experimentally (Kalaska et al., 1989). The predicted response as a function of movement and load direction is shown in Fig 2D.

If an isotropic inertial load  $m_{\text{ext}}$  is added during movement, the model prediction (Eq 6) is the sum of the movement response and the term  $F^{-1}m_{\text{ext}}\ddot{\mathbf{x}}(t)$ . This added term accentuates the acceleration dependent component of the signal shown in Fig 2C, especially in lateral directions where  $F^{-1}$  is larger. Thus our model predicts that when monkeys move with inertial loads of increasing magnitude, reversals of PV direction during the movement will begin to be observed – first in lateral movement directions. This is indeed the case (Kalaska et al., 1992).

### 3.5. Population vectors

Fig 2E shows the population vectors (Eq 5) for different postures, movements, loads, and movements in the presence of a constant load. In movement conditions the PV prediction has been averaged over the

movement time. For posture and movement, the PV successfully recovers the target direction as it does in experimental data (Georgopoulos et al., 1983).

This is no longer the case during load compensation and movement with load. The load compensation PV is systematically distorted, so that vectors pointing in lateral directions are elongated and vectors pointing at 45 deg are rotated towards the lateral axis. These distortions result from the anisotropy of the matrix  $F^{-1}$  in Eq 5 (which captures the anisotropy of the Jacobian transformation relating endpoint forces and joint torques). Very similar distortions in load PVs have been observed experimentally (Kalaska et al., 1989). It has also been observed that PVs computed in a mixed movement and load condition (Kalaska and Crammond, 1992) completely fail to reconstruct the movement direction, and show the same pattern of distortions as Fig 2E. This failure of the PV in the model is of course due to the added load-related signal.

### **3.6. Muscle cocontraction**

A number of studies have found systematic differences in M1 activity recorded during movements with identical directions made under varying experimental conditions (Caminiti et al., 1991; Scott and Kalaska, 1995; Kakei et al., 1999; Li et al., 2001). While these studies contradict the hypothesis of abstract directional encoding and are generally consistent with our model (see Sections 1 and 7), they do not point directly to a muscle-based representation in M1. To provide direct evidence for a muscle-based representation, one would have to compare experimental conditions in which both kinematic and kinetic parameters are identical but muscle activity is different. It is possible to design such an experiment by instructing, training, or otherwise provoking a subject to modulate the level of muscle cocontraction. Apart from some interesting early results (Humphrey and Reed, 1983), the control of muscle cocontraction has received surprisingly little attention in the M1 literature.

We analysed the changes in M1 activity associated with different levels of muscle cocontraction, in both monkeys and humans (Todorov et al., 2000). Monkeys trained in force field adaptation experiments (Li et al., 2001) show increased muscle cocontraction during the movement time, even in the baseline condition before any forces are applied. In that condition their movements are kinematically and kinetically indistinguishable from the typical movements that monkeys make in center-out reaching tasks, yet the M1 activity is markedly different. Both the directionally tuned and non-tuned populations of neurons increase their firing rates after movement onset, for all movement directions. In contrast, previous studies (Crammond and Kalaska, 1996; Moran and Schwartz, 1999b) have found that after movement onset, responses increase in the preferred direction and *decrease* in the anti-preferred direction – as our

model predicts in the absence of cocontraction. Thus increased muscle cocontraction is accompanied by a non-specific increase of activity in M1.

We also conducted an fMRI experiment with human subjects asked to cocontract their forearm and wrist muscles while maintaining posture. Preliminary results were described in (Todorov et al., 2000). In agreement with the monkey data, activation over M1 was significantly elevated in the cocontraction condition compared to rest in the same posture. Since subjects were not producing movements or isometric forces in any direction, the observed change in M1 activity cannot correspond to any directional command. This converging evidence from monkeys and humans provides strong evidence against the directional coding hypothesis, and fits naturally in our framework.

## 4. Trajectory reconstruction

Given the correlations between M1 activity and various behavioral parameters illustrated in Fig 2, it should be possible to reconstruct the movement trajectory more or less accurately from a population of M1 responses. The potential for neuroprosthetic applications has recently generated substantial interest in this topic (Wessberg et al., 2000; Helms Tillery et al., 2001). Here we propose a simple reconstruction method that takes advantage of the multiple correlations. We also explain in the framework of our model why trajectory reconstructions based on the Population Vector hypothesis (Schwartz, 1994; Moran and Schwartz, 1999a) are systematically distorted.

### 4.1. Reconstruction methods

Since the M1 signal most closely resembles movement velocity (Fig 2C), the simplest viable method is to integrate the PV over time – yielding reasonably accurate trajectory reconstruction (Schwartz, 1994; Moran and Schwartz, 1999a). We applied this method to the PV predicted by our model (Eq 5) in the spiral movements used experimentally (1.5-7.5 cm spiral radius, traced in 2.5 sec according to the  $2/3$  power law). The reconstruction we obtained (Fig 3A) is very similar to (Schwartz, 1994; Moran and Schwartz, 1999a), and displays the same distortion near the center of the spiral (the reason for the distortion is discussed later). An alternative method that has been studied is fitting a multilinear regression model (or a neural network) that maps cell firing rates to hand trajectories (Wessberg et al., 2000).

These methods implicitly assume that cell firing contains information about a *single* movement parameter: velocity in (Schwartz, 1994; Moran and Schwartz, 1999a) and position in (Wessberg et al., 2000). Empirically this is not the case, and therefore it is possible to develop better methods taking advantage of the multiple correlations that have been documented (as well as the regularities of hand kinematics). Optimal reconstruction can be achieved by a Bayesian method which involves: i) a



generative model of how neural activity depends on arm state; ii) a dynamic model predicting the current arm state from the previous one; iii) a recursive estimator that computes the posterior distribution of the arm state given the distribution inferred previously and the current neural activity. Although steps towards building such a method are already being taken (Gao et al., 2001), this is a difficult estimation problem and the optimal solution is not likely to be found in the immediate future.

Here we propose a simple reconstruction method which is related to PV integration (Schwartz, 1994) but also takes advantage of the correlations with hand position and acceleration (or load). According to our model (Eq 5) the population vector  $Uc(t-d)$  can be interpreted as the force acting on a point mass that moves in a visco-elastic environment. Therefore we can reconstruct the hand trajectory by simulating such a point mass and driving it with the (instantaneously computed) population vector. The state of the simulation is our prediction regarding the state of the hand. As in the PV integration method, the initial state has to be known. Endpoint impedance and external loads also have to be known (or fitted to optimize performance).

To assess the performance of our method, we generated a synthetic database containing 200 minimum-jerk trajectories (100 straight, 100 curved) with 300-600 msec duration and 5-10 cm extent. Each trajectory was executed in 4 dynamic conditions: normal, external load, external friction, external spring. For each execution we simulated the responses of 100 Poisson-spiking cells with mean firing rates given by Eq 6 (using the model of response variability), binned the spike data in 20 msec bins, and smoothed it with a 4th order Butterworth filter whose parameters were optimized separately for each reconstruction method. Fig 3B shows that the reconstruction error of our method (Dynamic Simulation) is lower than multilinear regression. As expected, prediction error decreases with the number of cells in the dataset. Note that our method generates simultaneous predictions about position and velocity, while the regression model is different in each case. We do not show an explicit comparison to the PV integration method because it is a special case of multilinear regression applied to velocity prediction.

The reconstruction method described here may appear to be tailored to our model of M1 responses, but this is not quite true. It is a very generic method, applicable to any population that (for whatever reason) happens to correlate with position, velocity, and force. Such simultaneous correlations in M1 are an empirical fact, and therefore they can be taken advantage of by using similar reconstruction methods.

## 4.2. Variable timelag

The distortion seen in Fig 3A and (Schwartz, 1994; Moran and Schwartz, 1999a) is due to an interesting phenomenon described by (Schwartz, 1994). The timelag between the M1 population vector and hand

velocity is not constant, but increases with the curvature of the hand path. The increase observed experimentally is quantitatively similar to the model predictions (Fig 3C and 3D) in the spiral movements from Fig 3A. In the linearized muscle model (Eq 5) this increase can be computed analytically (Todorov, 2000a). The answer is  $\text{atan}((m\omega - k/\omega)/b)/\omega$ , where the angular velocity  $\omega$  is related to the curvature  $\kappa$  by the 2/3 power law  $\omega = 12\kappa^{2/3}$ , and the scaling constant 12 rad cm/sec is extracted from the data in (Schwartz, 1994).

An intuitive explanation of why the model produces this effect is given in Fig 3C (top). In the outer part of the spiral the curvature is small (thus the acceleration component  $m\ddot{\mathbf{x}}$  is small), while the distance from the center of the workspace is large (thus the position component  $k\mathbf{x}$  is large). As a result the PV direction is biased in a way that matches the direction of hand movement at some earlier point in time (i.e. the PV appears time-advanced). In the inner part of the spiral the interplay between the position and acceleration components causes the opposite bias, making the PV appear time-delayed. Note that the instantaneous speed in the inner and outer parts of the spiral is also different; this is not visible in the intuitive plot but is taken into account in the analytical calculation.

Our explanation of the variable timelag phenomenon (Todorov, 2000a) has been challenged (Moran and Schwartz, 2000) on the basis that the model is too simple. It is not at all clear how the simplifications we introduced could cause a phenomenon so similar to the one observed experimentally, but to rule out that possibility we have repeated the analysis with the detailed muscle model (Todorov, 2000b). Fig 3D shows the timelag between the PV and hand movement direction computed from Eq 11. The four parameters of the detailed model were each varied by a factor of 2, demonstrating the robustness of the phenomenon. In one simulation we added gravity compensation, which caused fluctuations of the curvature-timelag function similar to those observed experimentally (Moran and Schwartz, 1999a) but did not alter its underlying shape. It is not surprising that the detailed and linearized muscle models produce almost identical results: the variable timelag phenomenon in the detailed model is due to the compensation for visco-elasticity, and the linearized model captures that visco-elasticity to first order.

## 5. Evidence for direction encoding?

According to the hypothesis of directional control M1 encodes predominantly the direction, rather than the magnitude, of movement velocity (Georgopoulos et al., 1982) and isometric force (Taira et al., 1996; Ashe and Georgopoulos, 1994). Apart from the damping asymmetry in Eq 5, 6, 7 (which makes the encoding of speed more difficult to detect) our model treats velocity and force as vectors whose direction and magnitude are nowhere separated. Thus it is important to examine in detail the results that have been

interpreted as encoding of direction irrespective of magnitude. After examining 6 such results (see below), we conclude that none of them is convincing evidence for direction encoding and in fact they are all consistent with our model.

Note that the magnitude of both velocity (Schwartz, 1994) and force (Fetz and Cheney, 1980) is clearly being encoded in M1. So the question we are asking here is whether the encoding of direction dominates the encoding of magnitude, in a way inconsistent with treating either force or velocity as a vector.

### 5.1. Force magnitude versus direction

Let  $k$  denote the trial number,  $r(k)$  the magnitude of the 3D isometric force vector  $\mathbf{f}_{\text{ext}}(k)$  produced by the monkey,  $x(k), y(k), z(k)$  the coordinates of the unit vector which points in the same direction as  $\mathbf{f}_{\text{ext}}(k)$ , and  $mfr(k)$  the rate at which a given neuron fires in trial  $k$ . To assess the relative contributions of force magnitude and direction to the firing rate of that neuron, it might seem reasonable to compare the two regression models

$$\begin{aligned} \text{(M)} \quad mfr(k) &= b_0 + b_r r(k) \\ \text{(D)} \quad mfr(k) &= b_0 + b_x x(k) + b_y y(k) + b_z z(k) \end{aligned}$$

For 79% of the M1 neurons studied by (Taira et al., 1996) only regression model (D) provided a significant fit. Can we then conclude that 79% of the neurons in M1 do not encode force magnitude, and if we reach such a conclusion how can we reconcile it with earlier studies (Fetz and Cheney, 1980) showing the opposite?

All we can conclude from the results of (Taira et al., 1996) is that the majority of M1 neurons do not encode force magnitude in the way prescribed by regression model (M). Note however that model (M) prescribes a rather unnatural encoding of magnitude: the response of the neuron is supposed to increase when the force magnitude increases, even when the force direction is opposite to the neuron's preferred direction. Consider instead a more natural encoding (Eq 6) where the response of the neuron is simply the dot-product of its preferred direction and the force vector. Geometrically such a response is a plane (Fig 4A left). The regression models used by (Taira et al., 1996) attempt to fit the surfaces shown in Fig 4A (right) to M1 data. If M1 neurons encode force as we predict, regression model (D) would clearly provide a significant (although not perfect) fit, while regression model (M) would not fit at all. We have confirmed this using synthetic responses generated from Eq 6 (Todorov, 2000a).

## 5.2. Bias versus dynamic force

To assess the relative contributions of static and dynamic isometric force, monkeys were trained to maintain a static (bias) force in a specified direction and then produce an additional (dynamic) force pulse in the direction of a randomly chosen target (Georgopoulos et al., 1992). The finding that the preferred direction for dynamic force does not depend on the direction of bias force (and vice versa) was taken as evidence that bias and dynamic forces are somehow encoded separately in M1. This finding was illustrated with a plot almost identical to Fig 4B.

Our model predicts the same effect (Fig 4B) without explicitly separating bias and dynamic force. Consider a neuron whose firing  $mfr$  is given by the dot-product of its preferred direction  $\mathbf{u}$  and the total isometric force  $\mathbf{f}_{\text{total}}$ . Since  $\mathbf{f}_{\text{total}}$  is the sum of  $\mathbf{f}_{\text{dynamic}}$  and  $\mathbf{f}_{\text{bias}}$ , we have

$$mfr = \mathbf{u}^T \mathbf{f}_{\text{total}} = \mathbf{u}^T (\mathbf{f}_{\text{dynamic}} + \mathbf{f}_{\text{bias}}) = \mathbf{u}^T \mathbf{f}_{\text{dynamic}} + \mathbf{u}^T \mathbf{f}_{\text{bias}}$$

If we now fix  $\mathbf{f}_{\text{bias}}$  and determine the preferred direction by varying  $\mathbf{f}_{\text{dynamic}}$ , the term  $\mathbf{u}^T \mathbf{f}_{\text{bias}}$  will simply act as a baseline and the preferred direction will be  $\mathbf{u}$ . If we change the direction or magnitude of  $\mathbf{f}_{\text{bias}}$  (or remove it completely) and repeat the same procedure, the only difference will be in the baseline and so the preferred direction will remain unchanged. The same holds for fixing  $\mathbf{f}_{\text{dynamic}}$  and computing the preferred direction with respect to  $\mathbf{f}_{\text{bias}}$ .

## 5.3. Cell classification

Although it is clear that the responses of individual neurons are correlated with multiple movement parameters, one could classify neurons according to the movement parameter showing the strongest correlation. This was done in a 2D movement task (Ashe and Georgopoulos, 1994) where the time-varying mean firing rate of each neuron was correlated with target direction (which does not vary within a trial), hand position, velocity, and acceleration. The responses classified as direction-related were most common (47%) while the acceleration-related responses were least common (6%). Does this finding contradict our model which contains no explicit representation of movement direction? The answer turns out to be negative: the same classification scheme applied to Eq 6 (with response variability) yields up to 30% direction-related, and only 2% acceleration-related responses.

How can a seemingly reasonable classification scheme find a large percentage of direction-related responses in a model without any explicit encoding of target direction? This is possible because a weighted sum of position, velocity and acceleration (Eq 6) does not necessarily resemble any one of these signals – as illustrated in Fig 4C. The family of synthetic response types given by Eq 6 is parameterized

by the cell-specific constants  $m_j, b_j, k_j, \bar{c}_j$ . The baseline and overall scale do not affect the outcome of the classification, so it is sufficient to study the family of responses parameterized by  $k_j/m_j$  and  $b_j/m_j$ . Each point in that two-parameter family corresponds to a response type, one of which is shown as an inset in Fig 4C. The classification boundaries induced by the procedure of (Ashe and Georgopoulos, 1994) are marked with thick lines, while the correct classification boundaries (correct classification is computed in the model by identifying the component signal with largest contribution) are marked with thin lines. The correlation-based classification boundaries are substantially different from the correct ones. In particular, a direction region (black lines) which should not exist has been created. This region is centered around the point where the correct boundaries meet, which corresponds to the synthetic response type most difficult to classify as either position, velocity, or acceleration-related. The correlation-based procedure is misclassifying the responses in this region as “directional” simply because they do not fit well in any of the alternative categories. For example, the response in the middle of the misclassified region resembles the velocity profile in the preferred direction and the acceleration profile in the opposite direction.

The results in Fig 4C are obtained using the mean firing rates given by Eq 6, and do not reflect any classification biases due to data transformations. Smoothing single-trial spike trains can cause further biases as demonstrated on synthetic data (Todorov, 2000a). It was shown (Georgopoulos and Ashe, 2000) that smoothing does not affect the results of (Ashe and Georgopoulos, 1994), but that may happen in other datasets.

A deeper problem with classifying cells into discrete categories is that even if we had access to the correct classification, it is likely to change between task conditions. That problem is illustrated in Fig 4C where the dashed lines show the correct classification boundaries for a two times faster movement. As explained in Section 3.3, executing the movement faster increases the relative contribution of the acceleration component, shifting the classification boundaries substantially. What does it mean then to label a cell as “velocity-related”, if it can become “acceleration-related” in a faster movement, “position-related” in a slower movement, and “direction-related” in a condition where the response is too complex to fit in any of the alternative categories?

#### 5.4. Phasic responses

The simple relationship between M1 and EMG activity given by Eq 1 does not strictly hold, as (Fetz et al., 1989) showed by averaging the activity of all corticomotoneuronal cells and muscles they recorded (Fig 4D). In particular, the M1 response at the onset of an isometric force ramp is more phasic than muscle activation. Is this evidence for abstract directional encoding? The answer is obviously negative.

As Fig 4D illustrates, EMG activity is more phasic than joint torque, just like M1 activity is more phasic than EMG activity. If we agree to interpret such differences as evidence for directional encoding, we must conclude that muscle activation carries an abstract directional component – which is quite absurd. The EMG-force difference is to be expected given that muscles act as low-pass filters (Zajac, 1989; Brown et al., 1999). This suggests a similar explanation for the M1-EMG difference: if the processing in the spinal cord (which we approximated with a fixed delay in Eq 1) has the effect of a low-pass filter, the descending M1 signal has to be more phasic in order to generate the necessary EMG activity. The black lines in Fig 4D show that the M1 signal given in (Fetz et al., 1989) can be filtered and thresholded to yield a close match to the EMG signal, and filtered again (using the nonlinear muscle filter from (Zajac, 1989)) to match the force trace.

An alternative explanation is that the phasic M1 signal serves a “priming” role such as bringing near threshold the spinal networks that activate the group of agonist muscles. The difference between a signal designed to compensate the effects of low-pass filtering and a signal designed to overcome thresholds is that the former should always be present, while the latter should only be present around the onset of EMG activity. The fact that phasic responses have not been reported in drawing tasks of longer duration (Moran and Schwartz, 1999a; Schwartz, 1994) makes the priming explanation more likely.

Note that phasic responses found in movement tasks (Kalaska et al., 1989) are most common in superficial layer neurons which we are not concerned with here (since they do not project to the spinal cord). This may explain why recordings from a mixed M1 population show an unusually strong phasic component, leading to a large population vector for dynamic force and an almost nonexistent population vector for bias force (Georgopoulos et al., 1992). The latter result however is difficult to interpret: it is not clear how the population vector method could fail to reconstruct the bias force given that the majority of M1 neurons were significantly tuned for bias force (Georgopoulos et al., 1992). One possible explanation for this puzzling result is that the preferred directions for dynamic force (which were used in computing the bias force population vector) were not sufficiently close to the preferred directions for bias force (which should ideally be used). Even if the difference between dynamic and bias preferred directions averaged over the population is zero, the bias population vector computed in this way scales down by an amount related to the circular variance of the distribution of differences.

## **5.5. Cell versus muscle tuning modulation**

To assess how directly M1 neurons control muscles, one can compute the ratio of maximal-to-minimal activation (which is a measure of directional modulation) and compare these ratios for M1 neurons and muscles. It has been shown that M1 tuning curves are less modulated by the direction of bias force than muscle tuning curves are (Georgopoulos et al., 1992). Is that finding inconsistent with Eq 1? Not at all, in

fact such a difference is to be expected if M1 output controls the activation of muscle groups directly. One possible mechanism is illustrated in Fig 4E: if we have just two neurons driving two muscles with both excitatory and inhibitory weights, an M1 tuning modulation of 2:1 can be transformed into a muscle tuning modulation of 8:1. Suppression of muscles pulling opposite to the neurons' preferred direction – which is essential for this mechanism – has been documented (Fetz and Cheney, 1980). Another mechanism that can change the tuning modulation is simply the subtraction of some “baseline”, which physiologically could correspond to descending activity that results in subthreshold spinal activation. If we modify the setup in Fig 4E to have only excitatory projections and subtract a baseline of 4, an M1 tuning modulation of 2:1 can be transformed into a muscle tuning modulation of 6:1. Therefore differences in tuning modulation are consistent with our model.

## 5.6. Temporal multiplexing

An interesting possibility suggested by (Fu et al., 1995) is that the nature of the M1 encoding changes in different phases of the movement (although it is not clear what “phase” means in continuous movements of longer duration). In the case of reaching it was shown that M1 activity is best correlated with target direction around movement onset, and correlations with other parameters (such as target location and distance) peak later in the movement (Fu et al., 1995). Does that imply an encoding of target direction in extrapersonal space around movement onset? It does not, for two reasons.

The first explanation (already mentioned in Section 5.4) is that M1 may have to send a signal that brings spinal neurons to their thresholds. This signal will depend on which muscle groups need to be activated (and that depends on the movement direction) but will not be affected by how far the target is. Note that varying the arm posture (Scott and Kalaska, 1997) should affect such a compensation signal, so it can be distinguished experimentally from an abstract direction encoding.

The second explanation is even simpler, in the sense that it does not require additional assumptions. In the study of (Fu et al., 1995) monkeys made movements of different extent in each direction. As expected from Fitt's law, both the duration and speed of the movement scaled with distance. Such scaling (Fig 4F) tends to make the initial portions of the trajectories indistinguishable, i.e. it is difficult to tell how far the indented target is given the arm kinematics around movement onset. Therefore our model (Eq 6) predicts a weak correlation with target distance or location at the beginning of the movement, and a stronger correlation with target direction – see Fig 4F. The details of the partial  $R^2$  calculation can be found in (Todorov, 2000a).

## 6. Cosine tuning

Cosine tuning of both M1 neurons (Georgopoulos et al., 1982;Kettner et al., 1988;Kalaska et al., 1989;Caminiti et al., 1991) and muscles (Turner et al., 1995;Herrmann and Flanders, 1998;Hoffman and Strick, 1999) has been observed repeatedly. The robustness of this phenomenon suggests that cosine tuning curves must be “appropriate” in a rather fundamental way, yet a satisfactory explanation is lacking. In a historical context dominated by single-joint studies one could have interpreted the broad directional tuning of M1 neurons as a principle of cortical computation. But the fact that muscles (as well as muscle spindles) have the same tuning curves as cortical neurons suggests that the explanation is more likely to involve known peripheral properties rather than unknown computational principles (see also Section 7.1).

Muscle length and velocity vary with the cosine of hand displacement and velocity (Eq 9) as first pointed out by (Mussa-Ivaldi, 1988). That would of course explain cosine tuning of spindle afferents, but is it an adequate explanation of cosine tuning in the motor output? We do not believe so, for the following reasons: i) cosine tuning is present in isometric force tasks and before movement onset; ii) muscle preferred directions can be systematically different from their lines of action (Cisek and Scott, 1998;Hoffman and Strick, 1999); iii) many tuning curves are sharper than cosine (Hoffman and Strick, 1999;Amirikian and Georgopoulos, 2000); iv) the arm musculature is redundant and so any set of net joint torques can be achieved with an infinite variety of EMG patterns (i.e. each muscle does not have to compensate for its own visco-elasticity or mimic the activity of its own spindle afferents).

Why is it then that the motor system tends to choose one out of the infinitely many activation patterns which would generate the same behavior? We have recently argued that the apparent redundancy is only due to the traditional emphasis on average behavior (see also (Loeb, 2000)). If factors such as robustness to noise and fatigue are taken into account, it can be shown that cosine tuning is the unique muscle activation pattern that optimizes expected performance. Observed deviations from cosine tuning also fall out of this minimization (Todorov, 2002). This argument is summarized below.

### 6.1. Consequences of multiplicative noise

It is well established that neuromotor noise is multiplicative, in the sense that the standard deviation of the net force increases linearly with its mean. Such scaling has been demonstrated for tonic isometric force (Sutton and Sykes, 1967;McAuley et al., 1997) as well as brief isometric force pulses (Schmidt et al., 1979). These empirical results lead to a neuromotor noise model where each muscle contributes endpoint force with standard deviation  $\sigma$  linear in the mean  $\mu$ :  $\sigma = \lambda\mu$ . Incorporating this scaling law into a minimum variance model of trajectory planning provides an explanation of Fitt's law and the smoothness of arm trajectories (Harris and Wolpert, 1998).



The above scaling law has another interesting consequence which is the basis for our explanation of cosine tuning. Suppose we have two identical muscles pulling in the same direction, and we want them to produce average net force  $\mu$ . If we activate only one muscle so that it generates average force  $\mu$ , the variance of the net force will be  $\lambda^2 \mu^2$ . If instead we activate both muscles so that each one of them generates average force  $\mu/2$ , the force output of each muscle will have variance  $\lambda^2 \mu^2 / 4$ . Assuming uncorrelated noise the variance of the net force will be  $\lambda^2 \mu^2 / 2$ , which is two times smaller! We have confirmed experimentally both the scaling law and the predicted noise reduction due to redundancy (Todorov, 2002). Fig 5A shows that the standard deviation of isometric grip force is linear in the mean. As predicted, the noise magnitude decreases when subjects use both hands (instead of the dominant hand alone) to generate the same instructed force.

Thus it is advantageous to activate all muscles pulling in the direction of desired net force. What about muscles pulling in slightly different directions? If all of them are recruited simultaneously the noise in the net force direction will still decrease, but at the same time extra noise will be generated in orthogonal directions. In other words the advantage of activating redundant muscles decreases with the angle away from the net force direction. Because of this systematic decrease, there exists a unique tuning curve for which the average net force is generated with minimum expected error. To identify that optimal tuning curve we need a formal model.

## 6.2. A model of noisy force production

As in Section 2.4B, muscle  $i$  generates endpoint force along a fixed line of action given by the unit vector  $\mathbf{p}_i$ . The muscle force (or tension)  $f_i$  is now divided in two components:  $f_i = \mu_i + z_i$ . The mean force  $\mu_i$  is assumed to be under the control of the nervous system, while the zero-mean term  $z_i$  combines all sources of noise affecting the force output of muscle  $i$ . The net endpoint force is then  $\sum_i (\mu_i + z_i) \mathbf{p}_i$ . The cost of using the vector of mean muscle forces  $\boldsymbol{\mu} = [\mu_1; \dots; \mu_n]$  to generate the desired net endpoint force  $\mathbf{f}^*$  will be defined as

$$\text{Cost}(\boldsymbol{\mu}, \mathbf{f}^*) = E_{z|\boldsymbol{\mu}} \left\| \mathbf{f}^* - \sum_i (\mu_i + z_i) \mathbf{p}_i \right\|^2 + \eta \sum_i \mu_i^2 \quad (12)$$

The first term in Eq 12 penalizes the expected difference between the desired and realized endpoint force, the second term is an effort cost, and the weighting parameter  $\eta$  sets the tradeoff between accuracy and effort. We use a supra-linear (i.e. quadratic) function to approximate effort because larger muscle forces are produced by larger and more rapidly fatiguing motor units (according to Henneman's size principle).

The expectation is taken over the probability distribution of the noise terms  $z_i$  conditional on the mean forces  $\mu_i$ . Since the quantity whose expectation we have to compute is quadratic in the random variables  $z_i$ , all we need to specify is the mean and covariance of the conditional distribution  $p(\mathbf{z} | \boldsymbol{\mu})$ . The mean is zero by definition, and the conditional covariance will be modelled as

$$\text{Cov}(z_k, z_s | \mu_k, \mu_s) = (\lambda_1 \mu_k \mu_s + \lambda_2 (\mu_k + \mu_s)) (\delta(k, s) + \lambda_3 + \lambda_4 \mathbf{p}_k^T \mathbf{p}_s) \quad (13)$$

The quadratic function parameterized by  $\lambda_1$  and  $\lambda_2$  generalizes the scaling law discussed above. The linear term  $\lambda_2 (\mu_k + \mu_s)$  is needed because of the non-zero intercept in Fig 5A (a constant term is also needed to model the non-zero intercept, but it does not affect the minimization with respect to  $\boldsymbol{\mu}$  so we ignore it). The second multiplicative term in Eq 13 does not depend on  $\boldsymbol{\mu}$ , but instead models the effects of the angle between the lines of action: the delta function  $\delta(k, s)$  corresponds to independent noise in the output of each muscle;  $\lambda_3$  is a noise term common to all muscles;  $\lambda_4 \mathbf{p}_k^T \mathbf{p}_s$  is a noise correlation whose strength varies with the angle between the two lines of action. The latter term is needed to capture the observation that muscles with similar lines of action show correlated EMG fluctuations (Stephens et al., 1999).

The constraints we impose on the mean forces  $\mu_i$  are

$$\begin{aligned} \mu_i &\geq 0 \text{ for all } i \\ \sum_i \mu_i &= co \quad (\text{optional}) \end{aligned} \quad (14)$$

The non-negativity constraint is needed because muscles can only pull. The cocontraction constraint is optional, and allows us to specify the desired arm impedance (see Section 2.4B). Note that generating muscle forces that cancel each other is always suboptimal according to the cost function in Eq 12, because it increases both the noise and energy consumption without affecting the net force. Ideally the cost function would take into account the advantages of higher impedance (allowing us to derive the optimal impedance and tuning function simultaneously), but that leads to technical complications which we leave for future work. For the purposes of deriving the optimal tuning curve it is sufficient to minimize Eq 12 subject to Eq 13 and 14, for every possible setting of the parameters  $\mathbf{f}^*, \eta, co, \lambda_1, \dots, \lambda_4$ .

### 6.3. Cosine tuning is optimal

The general minimization problem defined above was solved for many cases of interest in (Todorov, 2002). The main results are summarized here.

When the distribution of lines of action  $\mathbf{p}_i$  is uniform, it can be shown analytically that the optimal tuning curve is either a cosine or a truncated cosine. If the optional cocontraction constraint is specified, the tuning curve is a cosine for  $2\|\mathbf{f}^*\| \leq c_0$  and a truncated cosine otherwise (Fig 5B,  $p = 2$ ). If cocontraction is not specified, this condition becomes  $2\|\mathbf{f}^*\| \leq -\lambda_2 / 2\lambda_1$ . While analytical results could only be derived for a square cost (Eq 12), we also solved the minimization problem numerically for a variety of cost functions in which the force error term was raised to the power  $p \in [0.5; 4]$ . As Fig 5B shows, the optimal tuning curves were very similar to those derived analytically for  $p = 2$ . Both cosine and truncated cosine tuning curves have been observed experimentally (Hoffman and Strick, 1999). Our model predicts that the tuning curves will become broader when the cocontraction level is increased. Indeed we have obtained indirect evidence supporting this prediction (Todorov et al., 2000).

For a nonuniform distribution of lines of action  $\mathbf{p}_i$  the optimal tuning curves are not exactly cosine, but very close (Fig 5C). The direction of net force for which a muscle is maximally active (i.e. its preferred direction) is no longer aligned with its line of action. Instead the optimal preferred direction is rotated to fill in gaps in the distribution of lines of action. This effect is illustrated for a bimodal distribution in Fig 5C. Such systematic misalignment has been observed experimentally for both arm (Cisek and Scott, 1998) and wrist (Hoffman and Strick, 1999) muscles.

Another prediction of the model is a negative force bias (or undershoot) which increases with force level. This is because both the error and effort terms in Eq 12 are minimized when the net force is smaller than desired. The predicted magnitude of the effect depends on the error-effort tradeoff specified by  $\eta$ . In our isometric force experiment (Todorov, 2002) we observed such a negative bias. As predicted, it increased with the force level and was larger in the one hand condition.

## 7. The model in perspective

There is an interesting parallel (Fig 6) between the M1 model presented here and the classical model of orientation selectivity in primary visual cortex (V1): both models involve a primary cortical area (M1 or V1), a corresponding peripheral organ (muscle or retina), and an intermediate processing stage (spinal cord or lateral geniculate nucleus) approximated with a linear mapping. Making this parallel explicit can help us place the issues discussed above within the broader context of computational neuroscience, and identify directions for future research.

## 7.1. The origin of directional tuning

The defining feature of a muscle is its line of action (determined by the fixed attachment to the skeleton), while the defining feature of a photoreceptor is its location on the retina. In any given posture, a fixed line of action implies a preferred direction, just like a fixed retinal location implies a spatially localized receptive field. Therefore directional tuning in the motor system is no more surprising than spatial tuning in the visual system, in the sense that both types of tuning reflect fundamental peripheral properties. Of course, directional tuning of muscles does not necessitate such tuning in M1. Similarly, V1 neurons do not have to display spatial tuning – one can imagine for example a spatial Fourier transform<sup>4</sup> in the retina or LGN that completely abolishes the localized receptive fields of photoreceptors. But perhaps the nervous system avoids such drastic changes in representation, and so the tuning properties that arise at one level propagate to neighbouring levels, regardless of the direction of connectivity. The same argument may apply with regard to topographic organization: 2D topography arises in the retina and is found in V1, whereas muscle lines of action are not topographically organized (at least not in a sufficiently low-dimensional space) and correspondingly no detailed topography has been found in M1 (Schieber, 2001).

Proponents of the directional coding hypothesis have criticized the idea of muscle control on the basis that M1 neurons cannot possibly be "just higher-level motoneurons". The parallel with V1 illustrates how unreasonable such arguments are. Consider criticizing the related V1 model on the basis that V1 neurons cannot possibly be "just higher-level retinal ganglion neurons". Indeed they are not: V1 neurons display robust orientation tuning (Hubel and Wiesel, 1962) which emerges at the cortical level and does not reflect any peripheral properties. But that phenomenon is perfectly consistent with the simple linear model in Fig 6B, where V1 layer 4 neurons are driven by sets of retinal ganglion neurons with coaligned center-surround receptive fields. This parallel emphasizes the potential of structured high-dimensional mappings to generate interesting physiological responses, and the importance of elaborating the structure of the M1-to-muscle weight matrix  $W$  in future research. The latter would be much easier if a purely cortical phenomenon, comparable in robustness to V1 orientation selectivity, was discovered in M1. In contrast, all the phenomena addressed in this chapter reflect the basic properties of the musculo-skeletal system – which is why we could account for them without specific assumptions regarding the structure of  $W$ .

---

<sup>4</sup> A Fourier transform is a linear mapping, and therefore can be implemented by the connectivity illustrated in Fig 6.

## 7.2. The organization of muscle synergies

Each column of the M1-to-muscle weight matrix  $W$  corresponds to the “muscle field” of a pyramidal tract neuron. Are these muscle fields composed of arbitrary subsets of muscles, or is there some organizing principle, as is the case with V1?

The short latency contributions of CM neurons to muscle activity, as measured by spike-triggered averaging, clearly have some structure: it is known that individual neurons tend to facilitate groups of muscles pulling in similar directions (i.e. agonists) and to suppress their antagonists (Fetz and Cheney, 1980). This evidence, however, comes from single joint studies; in a multijoint setting the notions of agonists and antagonists are much more elusive, and a quantitative description of the biomechanical features that co-facilitated muscles share is yet to come. An interesting step towards a multijoint description of co-facilitated muscle groups is the recent finding that some CM neurons facilitate proximal muscles, others facilitate both proximal and distal muscles, but very few CM neurons facilitate distal muscles alone (McKiernan et al., 1998). This observation makes sense biomechanically, because the hand can only exert forces against external objects if those forces are transmitted via the proximal joints. Thus, muscle fields appear to be organized along some principles, but such principles are not nearly as clear as those underlying V1 organization. Furthermore, the CM neurons facilitating hand muscles appear to have rather uninterpretable muscle fields (Schieber, 2002); this is somewhat disturbing given that the CM pathway is most prominent in hand control.

The above discussion focused on muscle fields as defined by spike-triggered averaging. But M1 neurons exert most of their effects through spinal interneurons, and therefore a more proper definition of a muscle field would include both direct and indirect contributions. Unfortunately the indirect contributions are hard to quantify. One approach is to rely on long-term correlations between neuronal and EMG activity, keeping in mind that the results are likely to be affected by correlations with other neurons. Using that approach, it was recently shown that the so defined muscle fields appear to be grouped in a small number of clusters (Holdefer and Miller, 2002). But again, a (somewhat different) analysis of hand-related neurons showed no evidence of clustering or any other easily interpretable form of organization (Poliakov and Schieber, 1999). Note also that the correspondence between a neuron’s direct and indirect contributions, as measured by spike-triggered averages and long-term correlations respectively, is somewhat loose (McKiernan et al., 2000).

In summary, the experimental literature on muscle synergies is far from conclusive, and the issues could probably be clarified if the empirical work was complemented by formal models. Is there a modeling approach that could yield specific predictions regarding the structure of muscle synergies? The parallel with V1 again turns out to be illuminating. The discovery of orientation tuning and its simple mechanistic explanation have generated great theoretical interest into why the mapping illustrated in Fig

6B should have that particular structure, and why V1 neurons should be orientation tuned. Several models have led to the same general answer, which is that the V1 encoding is optimized to capture the statistical regularities of natural images (Olshausen and Field, 1996). In other words, orientation tuning cannot be predicted from properties of the visual periphery alone, but only makes sense when the properties of the visual world are also taken into account. By analogy, theoretical predictions regarding the structure of muscle synergies might come from analysis of the dynamic interactions between the musculo-skeletal system and the natural environments in which it operates. We have begun to develop a model of sensorimotor synergies (Todorov and Gharahmani, 2001) based on the same general idea: using unsupervised learning algorithms to capture the statistical regularities in the sensorimotor flow. Such an approach may also be useful in modeling the experience-dependent reorganizations in M1 (Nudo et al., 1996). This work will be presented elsewhere.

### **7.3. Using the redundancy of population codes**

The field of motor control has long suffered from the misconception that redundancy is a problem. It certainly is from the point of view of the researcher trying to understand why a specific motor pattern is being chosen. But from the point of view of the motor system redundancy is a solution rather than a problem, because the existence of multiple acceptable patterns can only make the discovery of one of them easier. Our recent theory of motor coordination (Todorov and Jordan, 2002) clarifies why musculo-skeletal redundancy is useful, and summarizes a number of psychophysical results demonstrating that such redundancy is being utilized by the motor system.

Redundancy is also present, to a much greater extent, on the level of neurophysiology. Any muscle activation pattern can be generated by a large set of cortical activation patterns<sup>5</sup>, and therefore one might expect this flexibility to be somehow taken advantage of. Indeed, a number of studies have shown that individual neuronal responses, and the way they change between experimental conditions, can be hard to interpret; and yet, the population of such responses generally behaves in accordance with our model. The change of preferred directions, when averaged over the M1 population, matches: a) the amount of shoulder rotation in 3D reaching (Caminiti et al., 1991); b) the average rotation of muscle preferred directions in 2D force field adaptation (Li et al., 2001). In both cases, however, the changes displayed by individual neurons span a wide spectrum and are generally uninterpretable. Similarly,

---

<sup>5</sup> In our linear model that redundant set is a linear subspace. The population vector is a projection that maps the entire redundant subspace into a point; this is why in the linear case we could predict population vectors without having to resolve redundancy.

adding the muscle fields of a number of hand-related cortical neurons provides reasonable reconstructions of the pattern of observed muscle activity (Schieber, 2002) – despite the fact that the relationship between the activity of individual M1 neurons and the muscles they directly facilitate or suppress is rather complex (McKiernan et al., 2000) and probably task-specific (Kilner et al., 2000). Finally, linear regression models that reconstruct movement trajectories from cortical population activity can work quite well in the face of substantial trial-to-trial fluctuations of single neuron activity (Wessberg et al., 2000).

The above examples indicate that motor cortex may take advantage of the redundancy of population codes, but do not clarify the nature of that advantage. A rare case where the latter is possible is the center-out wrist movement task of (Kakei et al., 1999), who measured the changes of M1 responses caused by changing the forearm posture. While the preferred directions of one M1 subpopulation rotated as much as muscle preferred directions (and lines of action), another subpopulation showed smaller rotations accompanied by changes in overall firing rate (or gain). Such changes are clearly inconsistent with an abstract encoding of direction, which predicts no effect of posture on any aspect of neural activity. The observed changes can be consistent with our model, as follows. Consider a muscle whose activation  $a$  is generated by two cosine-tuned neurons, with gains  $b_{1,2}$  and preferred directions  $\mathbf{u}_{1,2}$ . Under isometric force  $\mathbf{f}$  we have  $a = b_1 \mathbf{f}^T \mathbf{u}_1 + b_2 \mathbf{f}^T \mathbf{u}_2 = \mathbf{f}^T (b_1 \mathbf{u}_1 + b_2 \mathbf{u}_2)$ , i.e. the muscle activity is cosine-tuned with preferred direction  $b_1 \mathbf{u}_1 + b_2 \mathbf{u}_2$ . The latter vector can be rotated by rotating  $\mathbf{u}_{1,2}$ , but also by keeping  $\mathbf{u}_{1,2}$  fixed and only varying  $b_{1,2}$ . Therefore, the same rotation in muscle space (required to perform the task in a different forearm posture) can be accomplished by two different changes of cortical output: a rotation of preferred directions, or a change in gains. Why would M1 employ a mix of the two mechanisms? We offer the following intuitive answer. It is reasonable to assume that, in the wrist movement task of interest, there are at least two relevant inputs to M1: one specifying the target location in extrapersonal space and probably originating in ventral premotor area (Kakei et al., 2001), and the other one specifying the current arm posture irrespective of the task. If these two input signals are combined additively, they will produce gain changes without preferred direction changes – resembling the “extrinsic-like” neurons of (Kakei et al., 1999). How should M1 transform such an input pattern into an output pattern appropriate for the task, given that many output patterns are task-appropriate? One might expect the transformation to be accomplished with the “least amount of work”, which means that the output should resemble the input as much as possible, subject to the constraint that the necessary rotation in muscle space is implemented. We have performed a number of neural network simulations, in which recurrent networks were trained to transform the input described above into any task-appropriate output. The learning algorithm always converged to networks whose output patterns were a mix

“intrinsic” and “extrinsic” responses, as observed by (Kakei et al., 1999). These network models will be described in detail elsewhere.

#### **7.4. On the role of the spinal cord**

Our model does not clarify how the spinal cord works, but clarifies how it does not work. The fact that so many M1 phenomena can be explained using the simple feedforward transformation in Eq 1 indicates that the spinal cord does not isolate or hide the basic properties of the musculo-skeletal system from the rest of the brain. Indeed that would be a poor strategy, because primate arms (and especially hands) are used in a wide range of tasks with different dynamic requirements, and so a peripheral “detail” that is a nuisance in one task may become a crucial feature in another. The spinal cord could, however, augment musculo-skeletal properties in a way consistent with the model. For example, short latency stretch reflexes can compensate for muscle yielding, and thus maintain spring-like behavior on longer time scales (Nicols and Houk, 1976). To incorporate such a mechanism we simply have to modify the stiffness and damping terms in our model (in fact the psychophysical estimates we relied on probably reflect both muscle and reflex impedance).

More generally, one can think of the spinal cord as integrating two sources of activity: descending and afferent. The effects of the (pyramidal-tract) descending signal on muscle activity are modeled to first order by our linearization. Our results imply that the spinal cord generates the bulk of the muscle activity as a feedforward transformation of the descending signal, and that transformation is sufficiently smooth to appear linear over the small range of displacements, velocities, and loads typically used in the M1 literature. In this scheme, spinal feedback is used not to drive the movement or to translate abstract directional commands, but to counteract undesirable deviations by modulating muscle activity on short time scales.

This is in contrast with views where afferent signals play the leading role in generating muscle activity, while descending signals play the modulatory role of setting thresholds or gains for the spinal circuits. We disagree with such views for two general reasons. 1) Following deafferentations, both monkeys (Bizzi et al., 1984) and humans (Ghez and Sainburg, 1995) can produce substantial and quasi-normal muscle activity, while a leading role of afferent signals would seem to predict a lack of muscle activity in that condition. 2) The large network of supraspinal areas has all the information needed to compute the average muscle activations that propel the arm – especially in simple and overtrained tasks where all the internal models underlying this computation have presumably been formed. The only thing that the supraspinal system is known to be incapable of is generating very rapid corrections for undesirable deviations. Such corrections are crucial for ensuring stability, and since the spinal cord is



ideally suited to generate them, we do not see why it should be burdened with other computations that are best performed elsewhere.

To make this discussion more concrete, we introduce some ideas from stochastic optimal feedback control (Davis and R.B., 1985) that are relevant here. Consider a system with (unobservable) state  $\mathbf{x}_t$ , control  $\mathbf{u}_t$ , noisy feedback  $\mathbf{y}_t = H\mathbf{x}_t + \xi_t$ , and stochastic dynamics  $\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \omega_t$ . Such systems are controlled optimally by computing an internal state estimate  $\hat{\mathbf{x}}_t$  using the Kalman filter  $\hat{\mathbf{x}}_{t+1} = A\hat{\mathbf{x}}_t + B\mathbf{u}_t + K_t(\mathbf{y}_t - H\hat{\mathbf{x}}_t)$ , and mapping this estimate (rather than the raw feedback) into a control signal  $\mathbf{u}_t = -L_t\hat{\mathbf{x}}_t$ . It is not presently clear how this should be implemented in a sensorimotor hierarchy containing multiple feedback loops that operate simultaneously. Nevertheless, the above equations emphasize two important points. 1) Estimation and control are computationally separate processes, and involve separate sets of time-varying parameters ( $K$  and  $L$  respectively). The sensory gains  $K$  reflect the reliability of the afferent signal. When that signal becomes infinitely unreliable (or absent), the optimal  $K$  become 0, which automatically transforms the system into an open-loop controller with similar average behavior – in agreement with deafferentation studies. 2) The short-term contribution of the afferent signal  $\mathbf{y}$  to the control  $\mathbf{u}$  is given by the term  $L_{t+1}K_t(H\hat{\mathbf{x}}_t - \mathbf{y}_t)$ , which is proportional to the difference between expected and observed feedback. Since the difference is 0 on average, that contribution will be washed out in analyses that average data over trials – even in tasks where rapid single-trial feedback corrections are large and crucial for achieving stability and overall performance.

## 7.5. Equilibrium-Point control?

The Equilibrium-Point (EP) hypothesis (Feldman and Levin, 1995) has not been extensively applied to the M1 literature, but we include here a brief discussion of it because EP control is a rare example of a coherent theoretical framework – of the kind that we argued for in the abstract. Unfortunately, we do not see how it can be reconciled with the M1 literature. While the presence of positional gradients and cocontraction-related signals in M1 correspond naturally to the R and C commands postulated by the EP hypothesis, a number of other well established facts appear inconsistent with it:

1. The most salient finding from center-out reaching tasks is that movement-time activity is dominated by a velocity related signal. In particular, most neurons (during movement in their preferred direction) show an increase of firing followed by a decrease. In contrast, neural activity that encodes positional shifts should display monotonic changes in the course of reaching movements. The only way around that problem seems to be a proportional-derivative modification along the lines of (McIntyre and Bizzi, 1993), who proposed that both positional and velocity reference signals may be needed. This

still does not explain why the M1 encoding of velocity is strongly asymmetric – in the same way that muscle damping is asymmetric. Note also that including a transient velocity command results in a mode of control that is neither “equilibrium” nor “point”.

2. Straight reaching movements produced under curl-viscous forces (Li et al., 2001) are encoded differently from kinematically identical movements produced in the absence of such forces. As mentioned above, the average rotation of muscle preferred directions coincides with the average rotation of M1 preferred directions. The only way to reconcile EP control with such observations is to admit that, after all, the control signal takes dynamics into account (Gribble and Ostry, 2000). Then one would have to also admit that the reference trajectory for reaching movements can exhibit reversals, because such reversals have been observed in M1 (Kalaska et al., 1992). Note that, in general, dynamic additions to EP control eliminate one of the main advantages of such control: the avoidance of (either explicit or implicit) dynamic calculations. At the same time, such additions bring EP control closer to being an irrefutable hypothesis.
3. Finally, the ability of deafferented subjects (Bizzi et al., 1984; Ghez and Sainburg, 1995) to generate quasi-normal muscle activity seems fundamentally inconsistent with any view in which afferent signals play the leading role in activating muscles. One could postulate that, following deafferentation, both the spinal cord and M1 are reorganized so profoundly that they begin to function according to our model. But that would be merely a way to dismiss these important observations, rather than deal with them.

All that being said, we readily admit that the apparent difficulties and lack of parsimony may be due to our limited experience with building EP control models. Indeed, Anatol Feldman (personal communication) has suggested that since the command signals do not directly correspond to any one behavioral parameter, they may in principle be correlated with all parameters in the ways observed. We look forward to a specific neural implementation of EP control that addresses the M1 literature in detail. In the meantime, however, we remain skeptical – for the reasons outlined above.

## **Summary**

This chapter summarized a series of recent studies in which we have developed a simple mechanistic model of M1 control. The model accounts for a large number of experimental results, indicating that in the absence of major unpredictable events the trial-averaged M1 signal contains sufficient information to specify the trial-averaged pattern of muscle activations. This descending signal, which is the major output of the large supraspinal network of motor areas, takes into account peripheral properties such as muscle state dependence and perhaps spinal reflex gains. Therefore the essential dynamic properties of the motor periphery are not “hidden” by the spinal cord, as the directional coding hypothesis has suggested. In

retrospect this is hardly surprising: it would be a waste of resources if the entire supraspinal system was devoted to computationally trivial tasks such as calculating three-dimensional vectors, and the much more difficult problem of controlling the complex musculo-skeletal dynamics in real time was left to the ancient spinal circuits.

### **Acknowledgements**

We would like to thank Anatol Feldman and Gerald Loeb for their extensive comments on the manuscript. This work was supported by the Gatsby Charitable Foundation and the Alfred Mann Institute for Biomedical Engineering. Figures 1A, 2C (left), 2E, 3A, 3C, and 4F are reprinted with permission from (Todorov, 2000a). Figures 1B and 4C are reprinted with permission from (Todorov, 2000b). Figures 5A, 5B, and 5C are reprinted with permission from (Todorov, 2002).

## References

- Amirikian B, Georgopoulos A (2000) Directional Tuning Profiles of Motor Cortical Cells. *Neuroscience Research* 36: 73-79.
- Ashe J, Georgopoulos AP (1994) Movement Parameters and Neural Activity in Motor Cortex and Area 5. *Cereb Cortex* 6: 590-600.
- Bizzi E, Accornero N, Chapple W, Hogan N (1984) Posture control and trajectory formation during arm movement. *J Neurosci* 4: 2738-44.
- Brooks V (1986) *The neural basis of motor control*. New York: Oxford University Press.
- Brown IE, Cheng EJ, Loeb GE (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.
- Caminiti R, Johnson PB, Galli C, Ferraina S, Burnod Y (1991) Making Arm Movements Within Different Parts of Space: The Premotor and Motor Cortical Representation of a Coordinate System for Reaching to Visual Targets. *J Neurosci* 11: 1182-97.
- Cisek P, Scott SH (1998) Cooperative Action of Mono- and Bi-Articular Arm Muscles During Multi-Joint Posture and Movement Tasks in Monkeys. *Soc Neurosci Abst* 164.4.
- Conway BA, Halliday DM, Farmer SF, Shahani U, Maas P, Weir AI, Rosenberg JR (1995) Synchronization Between Motor Cortex and Spinal Motoneuronal Pool During the Performance of a Maintained Motor Task in Man. *J Physiol (Lond)* 489: 917-924.
- Crammond DJ, Kalaska JF (1996) Differential Relation of Discharge In Primary Motor Cortex and Premotor Cortex to Movements versus Actively Maintained Postures During a Reaching Task. *Exp Brain Res* 108: 45-61.
- Davis MHA, R.B. V (1985) *Stochastic Modelling and Control*. Chapman and Hall.
- Dum RP, Strick PL (1991) The origin of corticospinal projections from the premotor areas in the frontal lobe. *J Neurosci* 11: 667-669.
- Evarts EV (1968) Relations of pyramidal tract activity to force exerted during voluntary movements. *J Neurophysiology* 31: 14-27.
- Evarts EV (1981) Role of Motor Cortex in Voluntary Movements in Primates. In: *Handbook of Physiology* (Brooks VB, ed), pp 1083-1120. Baltimore: Williams and Wilkins.
- Feldman AG, Levin MF (1995) The Origin and Use of Positional Frames of Reference in Motor Control. *Behavioral and Brain Sciences* 18: 723-744.
- Fetz EE (1992) Are Movement Parameters Recognizably Coded in the Activity of Single Neurons? *Behav Brain Sci* 15: 679-690.

- Fetz EE, Cheney PD (1980) Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. *J Neurophysiol* 44: 751-772.
- Fetz EE, Cheney PD, Mewes K, Palmer S (1989) Control of Forelimb Muscle Activity by Populations of Corticomotoneuronal and Rubromotorneuronal Cells, chap. 36. *Progress in Brain Research* 80: 437-449.
- Fu QG, Flament D, Coltz JD, Ebner TJ (1995) Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J Neurophysiol* 73: 836-854.
- Gao Y, Bienenstock E, Black M, Shoham S, Serruya M, Donoghue J (2001) Encoding/decoding of arm kinematics from simultaneously recorded MI neurons. *Soc Neurosci Abst* 63.3.
- Georgopoulos AP, Ashe J (2000) One motor cortex, two different views. *Nat Neurosci* 3: 963-965.
- Georgopoulos AP, Ashe J, Smyrnis N, Taira M (1992) The motor cortex and the coding of force. *Science* 256: 1692-1695.
- Georgopoulos AP, Caminiti R, Kalaska JF (1984) Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. *Experimental Brain Research* 54: 446-454.
- Georgopoulos AP, Caminiti R, Kalaska JF, Massey JT (1983) Spatial coding of movement: a hypothesis concerning the coding of movement direction by motor cortical populations. *Experimental Brain Research Supplement* 7: 327-336.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2: 1527-1537.
- Ghez C, Sainburg R (1995) Proprioceptive control of interjoint coordination. *Can J Physiol Pharmacol* 73: 273-84.
- Gribble PL, Ostry DJ (2000) Compensation for loads during arm movements using equilibrium-point control. *Experimental Brain Research* 135: 474-482.
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394: 780-784.
- Helms Tillery SI, Lin WS, Schwartz AB (2001) Training non-human primates to use a neuroprosthetic device. *Soc Neurosci Abst* 63.4.
- Herrmann U, Flanders M (1998) Directional Tuning of Single Motor Units. *Journal of Neuroscience* 18: 8402-8416.
- Hoffman DS, Strick PL (1999) Step-Tracking Movements of the Wrist. {IV}. Muscle Activity Associated with Movements in Different Directions. *Journal of Neurophysiology* 81: 319-333.
- Holdefer R, Miller LE (2002) Primary motor cortical neurons encode functional muscle synergies. *Exp Brain Res*.

- Hubel DH, Wiesel TN (1962) Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex. *J Physiol (Lond)* 160: 106-154.
- Humphrey DR, Reed DJ (1983) Separate Cortical Systems for Control of Joint Movement and Joint Stiffness: Reciprocal Activation and Coactivation of Antagonist Muscles. In: *Advances in Neurology: Motor Control Mechanisms in Health and Disease*. (Desmedt JE, ed), pp 347-372. New York: Raven Press.
- Johnson PB (1992) Toward an Understanding of the Cerebral Cortex and Reaching Movements: A Review of Recent Approaches. In: *Control of Arm Movement in Space: Neurophysiological and Computational Approaches* (Caminiti R, Johnson PB, Burnod Y, eds), Berlin: Springer-Verlag.
- Joyce GC, Rack PMH, Westbury DR (1969) The Mechanical Properties of Cat Soleus Muscle During Controlled Lengthening and Shortening Movements. *J Physiol (Lond)* 204: 461-474.
- Kakei S, Hoffman D, Strick P (1999) Muscle and Movement Representations in the Primary Motor Cortex. *Science* 285: 2136-2139.
- Kakei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. *Nat Neurosci* 4: 1020-1025.
- Kalaska JF, Cohen DAD, Hyde ML, 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.
- Kalaska JF, Crammond DJ (1992) Cerebral cortical mechanisms of reaching movements. *Science* 255: 1517-23.
- Kalaska JF, Crammond DJ, Cohen DAD, Prud'homme M, Hyde ML (1992) Comparison of Cell Discharge in Motor, Premotor, and Parietal Cortex During Reaching. In: *Control of Arm Movement in Space: Neurophysiological and Computational Approaches* (Caminiti R, Johnson PB, Burnod Y, eds), Berlin: Springer-Verlag.
- Kettner RE, Schwartz AB, Georgopoulos AP (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space III. Positional gradients and population coding of movement direction from various movement origins. *J Neurosci* 8: 2938-2947.
- Kilner JM, Baker SN, Salenius S, Hari R, Lemon RN (2000) Human cortical muscle coherence is directly related to specific motor parameters. *J Neurosci* 20: 8838-8845.
- Lemon RN, Mantel GW, Muir RB (1986) Corticospinal facilitation of hand muscles during voluntary movement in the conscious monkey. *J Physiol* 381: 497-527.
- Li CS, Padoa-Schioppa C, Bizzi E (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 30: 593-607.
- Loeb G (1999) What might the brain know about muscles, limbs and spinal circuits? In: *Progress in Brain Research 123: Peripheral and Spinal Mechanisms in the Neural Control of Movement* (Binder MD, ed), pp 405-409. Elsevier.
- Loeb GE (2000) Overcomplete Musculature or underspecified Tasks? *Motor Control* 4: 81-83.

- McAuley JH, Rothwell JC, Marsden CD (1997) Frequency Peaks of Tremor, Muscle Vibration and Electromyographic Activity at 10 Hz, 20 Hz and 40 Hz During Human Finger Muscle Contraction May Reflect Rythmicities of Central Neural Firing. *Exp Brain Res* 114: 525-541.
- McIntyre J, Bizzi E (1993) Servo Hypotheses for the Biological Control of Movement. *Journal of Motor Behavior* 25: 193-202.
- McKiernan BJ, Marcario JK, Karrer JH, Cheney PD (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.
- McKiernan BJ, Marcario JK, Karrer JH, Cheney PD (2000) Correlations between corticomotoneuronal (CM) cell postspike effects and cell-target muscle covariation. *J Neurophysiol* 83: 99-115.
- Miller LE, Sinkjaer T (1998) Primate Red Nucleus Discharge Encodes the Dynamics of Limb Muscle Activity. *J Neurophysiol* 80: 59-70.
- Moran DW, Schwartz AB (1999a) Motor cortical activity during drawing movements: population representation during spiral tracing. *J Neurophysiol* 82: 2693-2704.
- Moran DW, Schwartz AB (1999b) Motor Cortical Representation of Speed and Direction During Reaching. *J Neurophysiol* 82: 2676.
- Moran DW, Schwartz AB (2000) One motor cortex, two different views. *Nat Neurosci* 3: 963-965.
- Mussa-Ivaldi FA (1988) Do Neurons in the Motor Cortex Encode Movement Direction? An Alternative Hypothesis. *Neurosci Lett* 91: 106-111.
- Nicols TR, Houk JC (1976) Improvement in Linearity and Regulations of Stiffness That Result from Actions of Stretch Reflex. *Journal of Neurophysiology* 39: 119-142.
- Nudo RJ, Milliken GW, Jenkins WM, Merzenich MM (1996) Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *J Neurosci* 16: 785-807.
- Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381: 607-609.
- Poliakov AV, Schieber MH (1999) Limited functional grouping of neurons in the motor cortex hand area during individuated finger movements: A cluster analysis. *Journal of Neurophysiology* 82: 3488-3505.
- Schieber MH (2001) Constraints on somatotopic organization in the primary motor cortex. *Journal of Neurophysiology* 86: 2125-2143.
- Schieber, M. H. Controlling the hand: A complex journey from motor cortex to finger movements. (2002). In *Neural Control of Movement*.
- Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT Jr (1979) Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol Rev* 86: 415-451.
- Schwartz AB (1994) Direct cortical representation of drawing. *Science* 265: 540-542.

- Scott S, Kalaska J (1995) Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *J Neurophysiol* 73: 2563-2567.
- Scott S, Kalaska J (1997) Reaching Movements with Similar Hand Paths But Different Arm Orientation. I. Activity of Individual Cells in Motor Cortex. *Journal of Neurophysiology* 77: 826-852.
- Scott SH (1997) Comparison of onset time and magnitude of activity for proximal arm muscles and motor cortical cells before reaching movements. *J Neurophysiol* 77: 1016-1022.
- Scott SH, Gribble PL, Graham KM, Cabel DW (2001) Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature* 413: 161-165.
- Sergio LE, Kalaska JF (1998) Changes in the Temporal Pattern of Primary Motor Cortex Activity in a Directional Isometric Force Versus Limb Movement Task. *J Neurophysiol* 80: 1577-1583.
- Stephens JA, Harrison LM, Mayston MJ, Carr LJ, Gibbs J (1999) The sharing principle. In: *Peripheral and spinal mechanisms in the neural control of movement* (Binder MD, ed), Oxford: Elsevier.
- Sutton GG, Sykes K (1967) The variation of hand tremor with force in healthy subjects. *Journal of Physiology* 191(3): 699-711.
- Taira M, Boline J, Smyrnis N, Georgopoulos A, Ashe J (1996) On the Relations Between Single Cell Activity in the Motor Cortex and the Direction and Magnitude of Three-Dimensional Static Isometric Force. *Exp Brain Res* 109: 367-376.
- Todorov E (2000a) Direct Cortical Control of Muscle Activation in Voluntary Arm Movements: a Model. *Nature Neuroscience* 3: 391-398.
- Todorov E (2000b) Reply to 'One motor cortex, two different views'. *Nat Neurosci* 3: 963-964.
- Todorov, E. and Gharahmani, Z. (2001) A theory of optimal sensory-motor primitives. In *Neural Control of Movement*.
- Todorov E (2002) Cosine tuning minimizes motor errors. *Neural Computation* 14: 1233-1260.
- Todorov, E. and Jordan, M. (2002) Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, in press.
- Todorov, E., Li, R., Gandolfo, F., Benda, B., DiLorenzo, D., Padoa-Schioppa, C., and Bizzi, E. (2001) Cosine Tuning Minimizes Motor Errors: Theoretical Results and Experimental Confirmation. *Soc Neurosci Abs* 785.6.
- Tsuji T, Morasso PG, Goto K, Ito K (1995) Human Hand Impedance Characteristics During Maintained Posture. *Biol Cybern* 72: 475-485.
- Turner RS, Owens J, Anderson ME (1995) Directional Variation of Spatial and Temporal Characteristics of Limb Movements Made by Monkeys in a Two-Dimensional Work Space. *Journal of Neurophysiology* 74: 684-697.



Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ, Srinivasan MA, Nicolelis MA (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408: 361-365.

Zajac FE (1989) Muscle and Tendon: Properties, Models, Scaling, and Application to Biomechanics And Motor Control. *Crit Rev Biomed Eng* 17: 359-411.

## Legends

### Table 1: Detailed muscle model.

The first equation shows the muscle force as a function of activation, length, and velocity. The remaining equations define the different terms in the first equation, and the corresponding parameter values.

### Figure 1: Model overview.

**A)** Solid arrows denote pathways explicitly taken into account, dashed arrows correspond to pathways that we argue can be neglected in a first-order model. The force field illustrates the static position-dependent endpoint forces produced by an elbow flexor, when the endpoint of a 2-link arm with 30cm links is moved over a 20x20 cm workspace (not drawn to scale). Note that the forces are roughly parallel, and become smaller when the endpoint is displaced in the force direction. **B)** Force output of the nonlinear muscle model at maximal activation, shown as a function of physiological length and velocity.

### Figure 2: Basic results.

**A)** The predicted postural response for targets arranged in a circle is a planar surface. **B)** The predicted response for increasing load magnitude in a constant direction is linear. A linear response is also predicted when the direction of a constant-magnitude load is varied; in this case the response is linear in the cosine of the angle between preferred and actual direction. **C)** Predicted movement-time response. Note that the compound response (left) is a sum of positional, velocity, and acceleration terms weighted by system-level stiffness, damping, and inertia respectively. In the nonlinear model (right) these components can no longer be separated, but the overall response is similar. **D)** Predicted response when both the movement direction and imposed load direction are varied independently. **E)** Population vector predictions for displacements or movements in different directions, loads in different directions, and movements in different directions in the presence of a constant load.

### Figure 3: Trajectory reconstruction.

**A)** Comparison of simulated hand kinematics (top) and the sum of population vectors predicted by the model (bottom). **B)** Comparison of the reconstruction errors (Y axis) of different methods. **C)** Analysis of the variable time-lag effect in the linearized model. **D)** Demonstration that the conclusions of the simple analysis hold for the detailed nonlinear muscle model.

**Figure 4: Apparent evidence for direction encoding.**

**A)** The predicted response for 2D isometric force is planar (left). Attempts to fit such responses by separating the contributions of force magnitude and direction are expected to fail, because the corresponding surfaces (right) are inappropriate for fitting planar responses. **B)** Model predictions of firing rate modulation due to independently varying the directions of dynamic and bias force. **C)** Illustration of the problems associated with classifying cells as position-, velocity-, acceleration-, or direction-related. Thin solid lines mark correct classification boundaries; thick solid lines mark the classification boundaries generated by linear regression classifiers; thin dashed lines mark the correct classification boundaries when the movement is made two times faster. **D)** Both cortical and muscle activation are known to be more phasic than the resulting torque pattern, which can be due to low-pass filtering (filter parameters shown in the figure). **E)** Illustration of how simple summation of cortical signals on the muscle level can lead to very different modulation ratios of cortical and muscle tuning (2:1 vs. 8:1). **E)** When targets at different distances are used, speed profiles scale both in magnitude and duration (left). As a result, the initial portion of the trajectory is roughly invariant of target distance, and so the predicted cortical signal is best correlated with direction early in movement (right).

**Figure 5: Cosine tuning.**

**A)** Measured isometric force fluctuations (Y axis) are a linear function of instructed grip force. Each subject participated in two conditions: gripping the force sensor with one hand or both hands. **B)** Optimal tuning curves can be full cosines or truncated cosines, depending on the ratio between specified cocontraction and force magnitude.  $p$  is the exponent of the cost function. The similarity of the curves for the three values of  $p$  indicates that our analytical results are not sensitive to the square cost ( $p=2$ ) used to derive them. **C)** Optimal tuning curves when the distribution of pulling directions (illustrated on the bottom) is non-uniform. The preferred directions are misaligned relative to the lines of action, so that the gap in the distribution is filled.

**Figure 6: Mechanistic models of M1 and V1.**


Schematic comparison of the model proposed here (**A**) and the classic model of orientation selectivity in primary visual cortex (**B**). The structure of the weight matrix  $W$  in the case of V1 is known, and explains how orientation selectivity can emerge on the cortical level. In contrast, the M1-to-muscle weight matrix has unknown structure; as a result, our model can presently address only those phenomena that reflect basic properties of the musculo-skeletal system. Directional tuning is such a phenomenon, and closely corresponds to the spatial localization of receptive fields in the visual system.

**Table 1**

Functions	Parameters
$f(a, l, v) = A F_L F_V + F_{P1} + A F_{P2}$	
$A(a, l) = 1 - \exp\left(-\left(\frac{a}{a_f N_f}\right)^{N_f}\right)$	$a_f = 0.56$
$N_f(l) = n_{f0} + n_{f1} \left(\frac{1}{l} - 1\right)$	$n_{f0} = 2.11; n_{f1} = 4.16$
$F_L(l) = \exp\left(-\left \frac{l^\beta - 1}{\omega}\right ^\rho\right)$	$\beta = 1.93; \rho = 1.87; \omega = 1.03$
$F_V(l, v) = \begin{cases} \frac{v_{\max} - v}{v_{\max} + (c_{v0} + c_{v1}l)v}, & v \leq 0 \\ \frac{b_v - (a_{v0} + a_{v1}l + a_{v2}l^2)v}{b_v + v}, & v > 0 \end{cases}$	$v_{\max} = -5.72; b_v = 0.62$ $c_{v0} = 1.38; c_{v1} = 2.09$ $a_{v0} = -3.12; a_{v1} = 4.21; a_{v2} = -2.67$
$F_{P1}(l) = c_1 k_1 \ln\left(\exp\left(\frac{l - l_{r1}}{k_1}\right) + 1\right)$	$c_1 = 104.25; k_1 = 0.052; l_{r1} = 1.42$
$F_{P2}(l) = c_2 \exp(k_2(l - l_{r2}) - 1)$	$c_2 = -0.02; k_2 = -18.7; l_{r2} = 0.79$

**Figure 1**

**A) Schematic Illustration**



**B) Detailed Muscle Model**



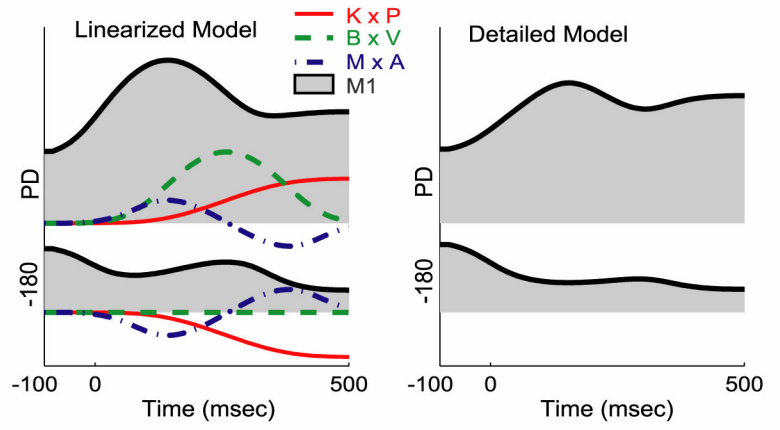


Figure 2

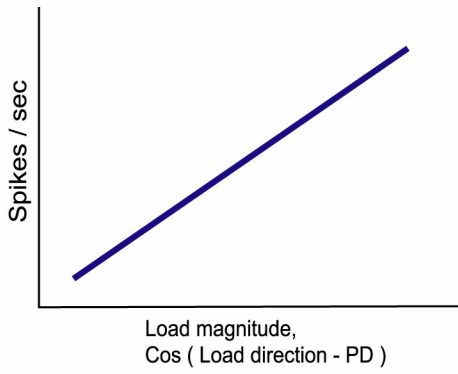
A) Position



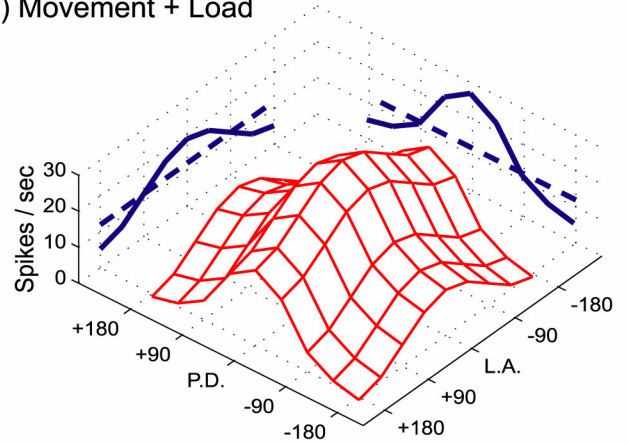
C) Movement




B) Load



D) Movement + Load




E) Population Vectors




**Figure 3**


**A) PV Reconstruction**



**B) Comparison of Reconstruction Methods**



**C) Variable Lag - Linearized Model**



**D) Variable Lag - Detailed Model**


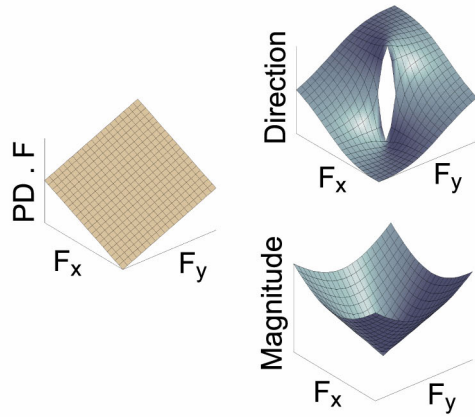
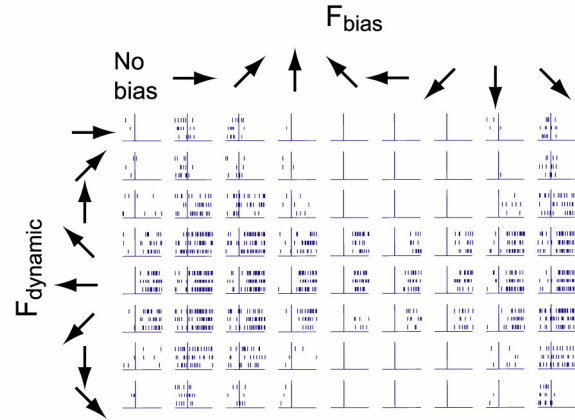


Figure 4


A) Force Magnitude vs. Direction



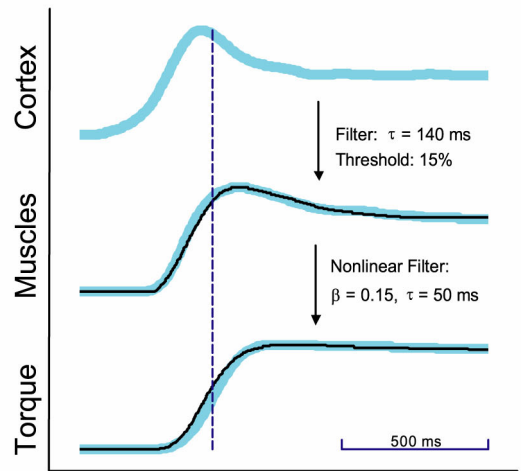
B) Bias vs. Dynamic Force




C) Cell Classification



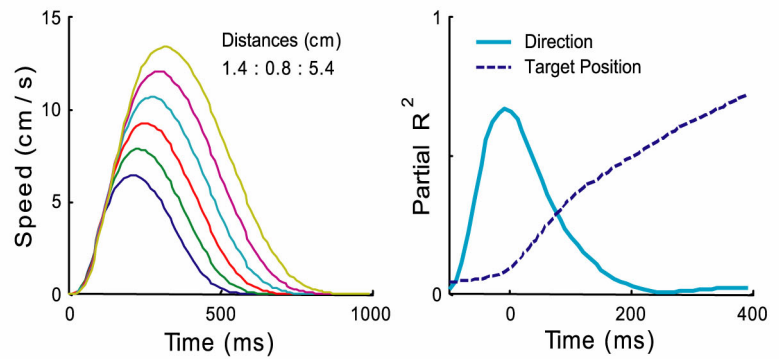
D) Phasic Responses



E) Cell vs. Muscle Tuning

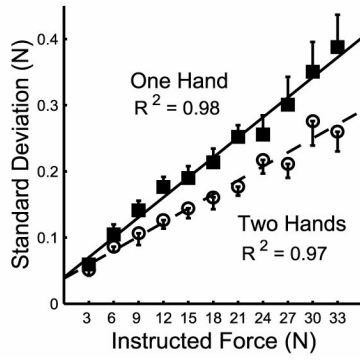


F) Temporal Multiplexing

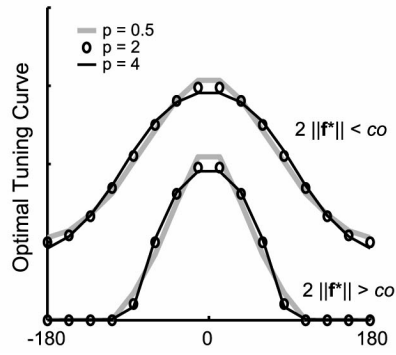


**Figure 5**

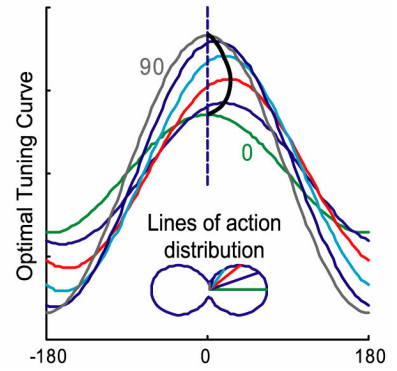
**A) Multiplicative Noise**



**B) Uniform Distribution**

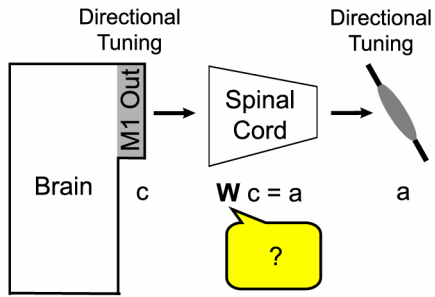


**C) Nonuniform Distribution**



**Figure 6**

**A) Primary Motor Cortex**



**B) Primary Visual Cortex**

