

The Problem of Parametric Neural Coding in the Motor System

Jacob Reimer and Nicholas G. Hatsopoulos

Abstract In the early visual and auditory system neurons are sensitive to a variety of parameters including orientation, contrast, and spatial and temporal frequencies, amplitude, timing, and spectral variables. There are theoretical reasons to believe that neural tuning for these particular parameters is fundamental to the information processing in each area. In contrast, we argue on both principled and empirical grounds that the idea of parametric encoding that has been so fruitfully applied to processing in early sensory systems does not have the potential to achieve more than heuristic or operational status in explanations of the motor system. In the motor system, inherent correlations among parameters of motion that occur in natural movements will necessarily make a neuron that is tuned to one variable also appear to be sensitive to other variables at different time lags. Similarly, depending on the nature of the task, neurons that appear to be tuned to parameters in one coordinate frame will often appear to be tuned to correlated variables in other coordinate frames. Finally, we point out that the tuning for any parameter can vary significantly with time lag. For all these reasons, we suggest that it may not be particularly meaningful to ask whether one or another movement parameter is represented in motor cortex. Instead, we propose that the tuning of any movement-sensitive cortical neuron is best envisioned as carving out a specific hyper-volume in a high-dimensional movement space. When one considers the way this tuning space changes over time, the time-varying preferred parameter values of the neuron describe a small segment of movement that we call a “movement fragment”.

Introduction

Since the advent of behavioral electrophysiology in the sixties, a large body of evidence has accumulated that motor cortical (MI) neurons encode almost every imaginable movement parameter. Pioneering work by Evarts and others

N.G. Hatsopoulos (✉)

Department of Organismal Biology and Anatomy, University of Chicago, Chicago, IL 60637, USA

e-mail: nicho@uchicago.edu

showed that activity in motor cortex is correlated with forces generated at the periphery (Evarts 1968). In the early 1980s, Georgopoulos and colleagues found that the activity of many cells in MI was tuned to the direction of hand movement (Georgopoulos et al. 1982). Subsequent studies showed that MI activity correlates with limb position (Georgopoulos et al. 1984; Paninski et al. 2004), velocity (Moran and Schwartz 1999), acceleration (Georgopoulos et al. 1982), movement distance, joint torques, and hand force (Cabel et al. 2001; Cheney and Fetz 1980; Hepp-Reymond et al. 1978; Kalaska et al. 1989; Smith et al. 1975; Taira et al. 1996), or some combinations of those parameters (Ashe and Georgopoulos 1994; Fu et al. 1995; Fu et al. 1993; Kurata 1993). The diversity of these results has led to a lack of consensus as to the functional role of motor cortex. In this chapter, we hope to illustrate the issues that make interpreting parametric tuning difficult, especially with regards to elucidating the functional role of motor cortical neurons. A similar argument has been made previously on different grounds (Fetz 1992).

Correlated Parameters

A fundamental problem in motor physiology is that statistical dependencies among movement parameters ensure that neurons that appear to encode one parameter will also appear to encode other correlated parameters. In sensory physiology, this problem can be overcome because the experimenter has full control of the stimulus and because it is possible to uncouple the different variables that are present in natural stimuli. In motor control experiments, the experimenter can only guide or provide constraints to movements that are under the subject's control. Naturally generated movements involve biological constraints (i.e. empirical constraints peculiar to the neural control of movement) such as minimization of jerk and the two-thirds power law between movement velocity and curvature; these constraints can introduce dependencies between different parameters. However, even if the experimenter were hypothetically able to train a human or animal subject to violate these biological constraints, they would still be unable to eliminate statistical dependencies due to physical constraints (i.e. the presence of interaction torques between the shoulder and elbow, restrictions on hand trajectories imposed by joint geometry, the relationship between torques and joint angles due to passive muscle properties, and Newton's laws of motion relating acceleration to force).

In addition, behavioral paradigms that are used to study the motor system introduce their own correlations among movement parameters. For example, the center-out task introduces a correlation between the position and speed of the hand such that positions near the center and peripheral targets are associated with slow speeds. In an attempt to attenuate these correlations, we and others have employed a "random target pursuit" (RTP) task in which movements are made to sequences of randomly positioned targets. In this task, the monkey generates arm movements in the horizontal plane in order to move a

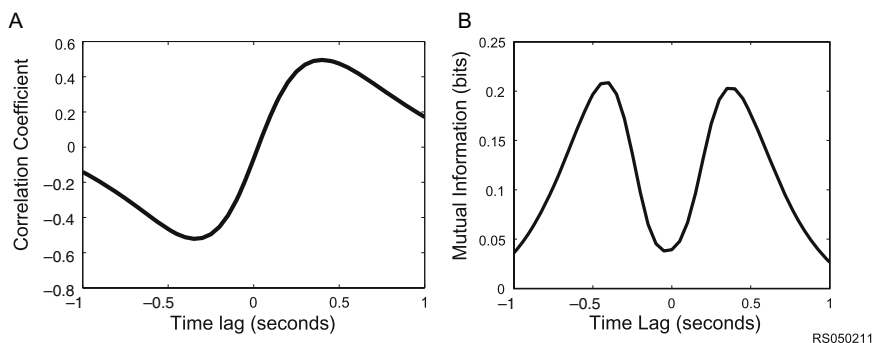


Fig. 1 Correlations among movement parameters (**A**) Despite the lack of correlation between position and velocity at zero lag in the RTP task, there remains a time-lagged correlation between these two variables as measured by the cross-correlation. This is partly due to the limited workspace such that a leftward velocity will result in a leftward position after some time. (**B**) Information theory can also be used to quantify this statistical dependency by computing the mutual information between these two variables which measures how much information in bits one variable provides about the other. It is evident that the variations in the mutual information profile closely match the cross correlation

cursor through a sequence of randomly-appearing targets. There is no planning or “hold” period in the task, and even though the animals are rewarded with juice every seven target acquisitions, they typically generate a continuous trajectory over many trials.

The RTP task produces its own characteristic set of correlations. For example, despite the lack of correlation between position and velocity at zero lag, there remains a time-lagged correlation between these two variables as measured by the cross-correlation (Fig. 1A). This is partly due to the limited workspace such that a leftward velocity will result in a leftward position and then a rightward velocity after some time. Information theory can be used to quantify this statistical dependency by computing the mutual information between these two variables. Mutual information measures how much information in bits one variable provides about the other (Fig. 1B), and has the additional advantage that it can capture non-linear statistical dependencies. It is evident that the variations in the mutual information profile closely match the cross correlation.

Because of the statistical dependencies among movement parameters in the RTP task, a motor cortical neuron that is tuned to velocity will also appear to be tuned to position at a characteristic time lag. To illustrate this phenomenon, we computed the mutual information temporal profile between the firing rate of a neuron and the velocity of the hand (Fig. 2A, light trace). This cell provides the most velocity information in its rates at a time lag of ~ 100 ms, which is typical of many MI neurons, and is consistent with a causal role in generating movement (Moran and Schwartz 1999; Paninski et al. 2004). However, the same neuron also provides information about hand position at different time leads and lags (Fig. 2A, dark trace) which can be explained by the inherent correlations between

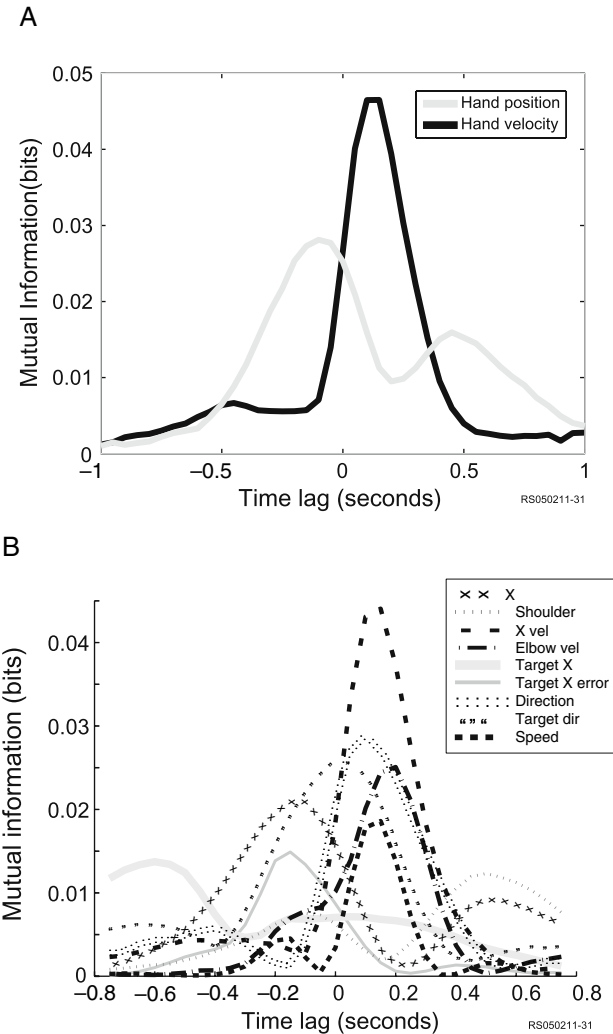


Fig. 2 Characterizing motor cortical tuning using information temporal profiles (A) Mutual information temporal profile computed between the binned firing rate (50 ms bins) of an example neuron and the velocity of the hand (*light trace*). The same neuron also provides information about hand position at different time leads and lags (*dark trace*) which can be explained by the inherent correlations between velocity and position in the random target pursuit task (see Fig. 1). Peaks to the left of zero indicate that the neuron is providing information about the value of the parameter in the past relative to the firing rate of the neuron, while peaks to the right of zero indicate that the neuron is providing information about the value of the parameter in the future (B) Information temporal profiles for the same neuron for a number of parameters

velocity and position in the random target pursuit task (see Fig. 1). By examining information in the temporal profiles for a number of parameters, it is apparent that this neuron is sensitive to a vast number of movement features (Fig. 2B). Many of these sensitivities can be explained by the statistical dependencies among these parameters.

Correlated Coordinate Frames

Parameterizing movement requires one to define a coordinate system that allows one to quantify behavior. In research on upper limb movements in primates, common examples include body, shoulder, or joint coordinate frames. There are ethological and biomechanical justifications for these three coordinate frames, and in theory each captures different dimensions of movement. In practice, however, experimentally-generated reaching movements are often highly stereotyped and traverse a much lower-dimensional space than is theoretically possible. As a consequence of this dimensional redundancy, movement parameters measured in different coordinate system are often highly correlated.

For example, when the movement of the arm is constrained in two dimensions, hand position and joint angles are highly correlated (Fig. 3A). Positions and velocities in one set of coordinates are correlated with positions and velocities in another set at a zero time lag. Therefore, a neuron that appears to encode hand velocity in a body-centered coordinate system will also appear to be sensitive to shoulder and elbow joint velocity (Figs. 3B and 3C). Increasing the number of available degrees of freedom (by removing experimental constraints) can reduce the dependencies between coordinate frames, but only if the task involves movements that explore the full range of parameter values (unlike a 3-D center-out task, for example). Even unconstrained natural movements explore only a lower-dimensional region of the movement space than is theoretically possible (Graziano et al. 2004). Partly due to these correlations, there is still no clear evidence of a privileged coordinate frame in which arm movements are encoded in motor cortex (Wu and Hatsopoulos 2007).

Multi-Dimensional Sensitivity

Despite the fact that some neurons exhibit multiple parametric sensitivities that fall out of the statistical dependencies among the parameters, we find that a significant proportion of neurons in MI contain information about movement that cannot be explained in terms of the encoding of a single parameter at a single time lag (Fig. 4; each panel is a separate simultaneously-recorded neuron). Note that all six of these example neurons are well-tuned for direction (as well as position and velocity), but that the time-lags of peak directional information vary from neuron to neuron, and some neurons contain multiple

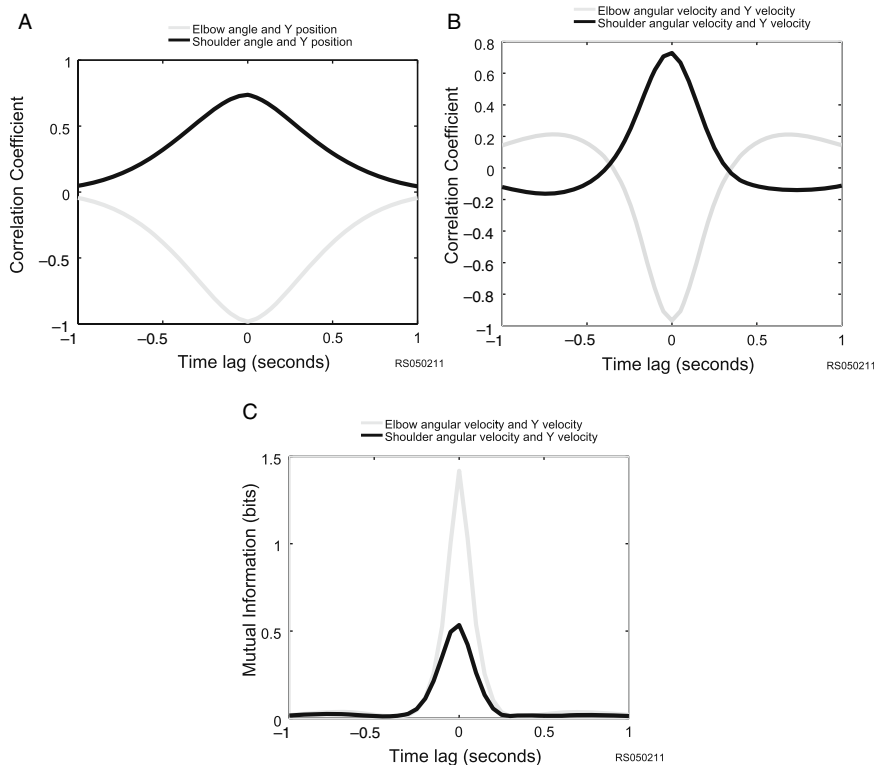


Fig. 3 Coordinate frames and correlated parameters **(A)** When the movement of the arm is constrained in two dimensions, at any moment hand position and joint angles are highly correlated. Positions in one set of coordinates are correlated with positions in another set at a zero time lag. **(B)** Joint velocities are correlated with hand velocities, also at a zero time lag. **(C)** Mutual information between velocities measured in two different coordinate frames mirrors the correlation structure of the task

peaks. Although it is certainly possible to characterize the direction tuning of these cells at a single time lag, each neuron actually displays a more complex sensitivity, in the sense that their temporal information profiles are not explicable by the correlation structure in the task.

For example, the panel in Fig. 4A shows the information profile for a cell that provides maximum information about position and velocity at a similar time lag. If information about position was an epiphenomenon of the information provided about velocity, one would expect the peaks to be offset by the delay in the cross-correlation peaks between the two parameters (as in the simultaneously-recorded cell shown in Fig. 2A). Other neurons exhibit complex multimodal profiles. For example, some profiles contain bimodal peaks in velocity information at both positive and negative time lags (Fig. 4B–E). Figure 4F shows a cell with a trimodal profile for velocity information and a roughly

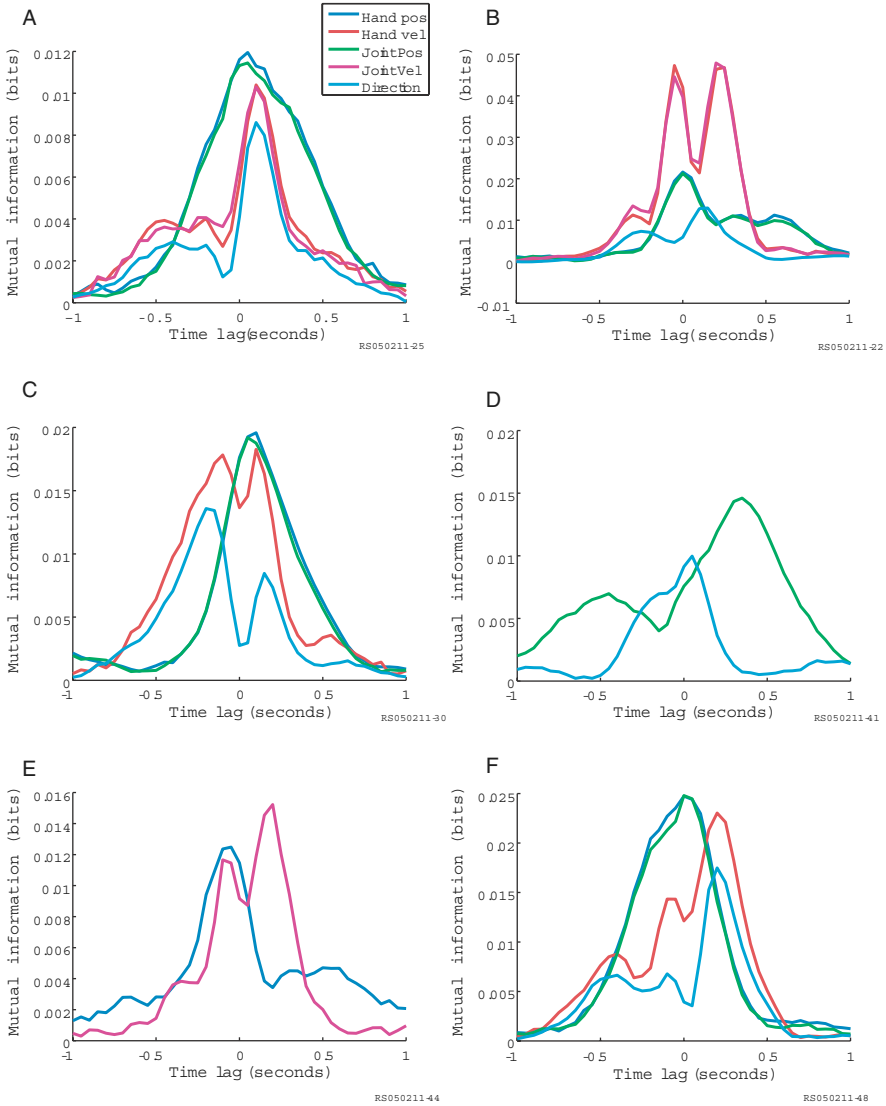


Fig. 4 Multiple information temporal profiles unexplained by inherent correlations Mutual information between neural firing and movement at multiple time lags is computed in the same way as in Fig. 2 (see Methods) for hand position (blue), joint angles (green), hand velocity (red), joint velocities (magenta), and direction (cyan)

unimodal profile for direction information, another relationship that would not be predicted by the correlation structure in the task (since direction and velocity are well-correlated at a zero time lag). Note that because of the extremely high correlations between joint angles and hand position in this task, information




Fig. 5 Multidimensional sensitivity Scatter plot of mutual information peaks for 20 different MI neurons (*black circles*). Peak values of mutual information for either direction or elbow speed (whichever was larger) are plotted along the X axis. The peak information value for each cell when the mutual information was calculated using the joint distribution of the two parameters is plotted along the Y axis

profiles for variables in the two coordinate frames are very similar. However, the similarity in mutual information for joint angles and hand position does not necessarily imply that individual neurons are actually equally sensitive to both sets of variables in different contexts.

If neurons are sensitive to multiple dimensions, the information provided by one movement parameter about the firing rate of the cell should not be redundant with the information provided by a second parameter. Consider, for example, two parameters such as movement direction (measured in a body-centered coordinate system) and elbow speed. The peak mutual information provided by both movement parameters jointly about neuronal activity is larger than the peak mutual information provided by either parameter alone (Fig. 5). If the two parameters were completely redundant, the peak information provided by the joint distribution of the variables would equal the maximum provided by either of the single parameters (i.e. the points would fall on the diagonal line in Fig. 5). We find that in the RTP task, the information about neural activity calculated from combinations of parameters is generally larger than the information provided by either parameter individually.

Time-Varying Tuning and Movement Fragments

The temporal information profiles displayed in Figs. 2 and 4 indicate that the degree of tuning for any specific movement parameter is not static but varies substantially at different time leads and lags. The actual parameter values that

the neuron is tuned to can also vary dramatically with time. When one considers the way the preferred values and tuning strength for multiple parameters evolve over time, the result is a probabilistic trajectory through space that we refer to as a “movement fragment”.

The idea of “movement fragment” encoding can be illustrated for just two dimensions by generating averaged spike-triggered fragments (STFs) of hand trajectories from different neurons while the animal performs the RTP task (Fig. 6A,B,C). Imagine that each time a neuron spikes, we take a “snapshot” of the actual movement trajectory that the monkey produced from one second prior to the spike to one second after it. The spike-triggered movement fragment for this cell is the average (taken independently at each time lead or lag) of all these snapshots taken over many such spikes. In Fig. 6, one half of the STF (blue) represents the average x and y positions up to one second prior to the occurrence of a neuronal spike, and the other half (red) represents the average of all the position values subsequent to the spike. Darker colors represent larger time lags between the mean position and the cell’s spiking. The green segment represents the average position simultaneous with and immediately subsequent to the spike occurrence. The STF shown in Fig. 6D is for the same cell as in Fig. 6C; the difference is that elbow and shoulder joint angles are plotted instead of Cartesian coordinates.

The time-varying preferred directional tuning of a neuron can also be used to describe the movement fragment which it encodes. Figure 6 E, F, and G show the instantaneous directional tuning of the same cells as in Fig. 6 A, B, and C. Each arrow points in the preferred direction of a cell at a given time lead or lag (± 1 second; calculated in 50 ms bins) with respect to the neural activity. Arrows are plotted with their quivers positioned at the corresponding time lead/lag along the STF for each cell for ease of comparison, and the color of the arrows represents the lag with respect to the neural activity, as in the STF. For the first two cells, the temporal shifts in preferred direction from past to future (E,F) are roughly consistent with the path traced out by the STF (A,B). We argue that these two views of the movement sensitivity of the cell are equivalent – that the trajectory described by changes in a neuron’s multidimensional tuning over time in essence describes a fragment of movement.

The specific equivalence between directional tuning and Cartesian position illustrated in Fig. 6 A, B and E, F does not necessarily hold in general for MI neurons. As we’ve argued, it is extremely unlikely that any small number of parameters could adequately capture the tuning properties of every movement-sensitive cell. Of course, if every neuron was exclusively tuned to hand position and direction (in a body-centered coordinate system), the preferred direction would always line up with the STF. That this is not always true is illustrated in Fig. 6C and G. For this cell, the preferred directions in Cartesian coordinates (6G) trace out a path that is almost exactly the opposite of the spike-triggered fragment (6C). However, when the STF is generated in a joint-angle coordinate system, the directional tuning (in joint space) and the spike-triggered fragment are much more consistent (Fig. 6D,H). One way to interpret this result is to

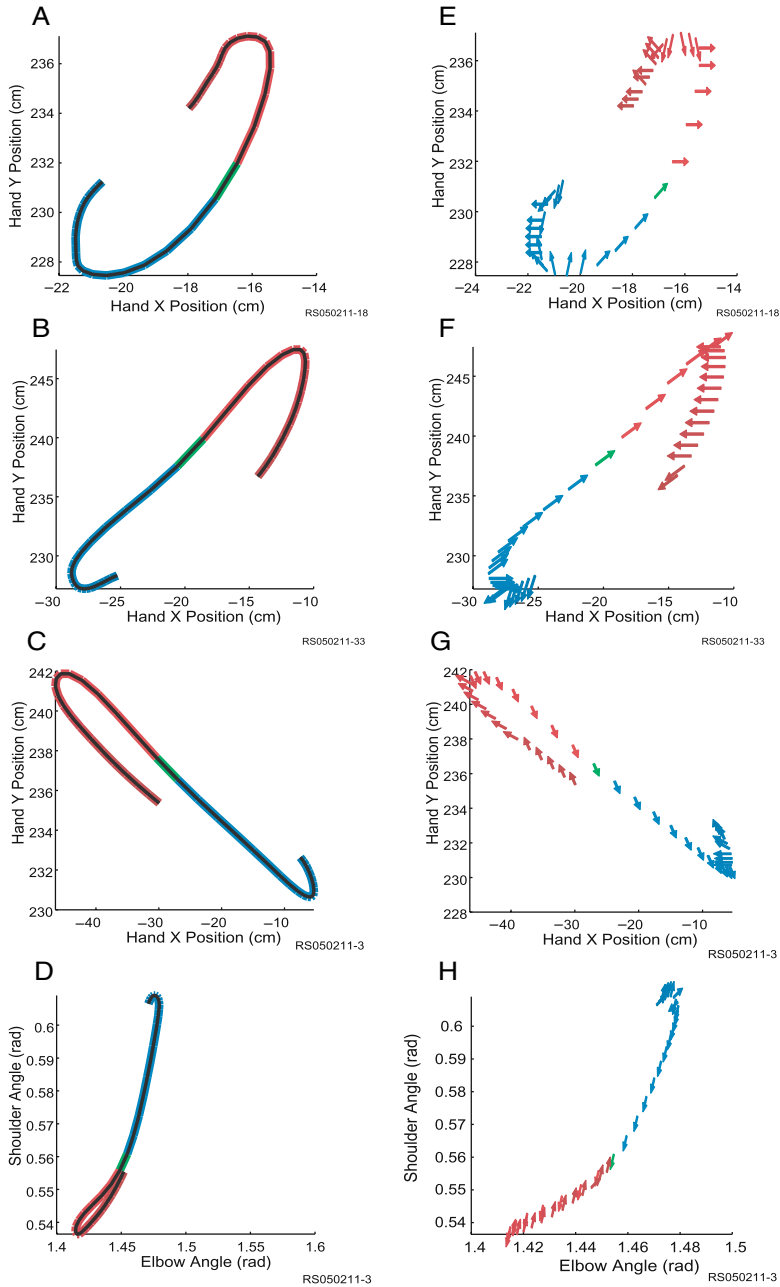


Fig. 6 Movement fragments (A,B,C) Spike-triggered fragments for three neurons. Each 50 ms bin with one or more spikes was used to select a 2 second “window” of movement centered on the spiking bin. The average (computed independently for x and y position at each lead/lag) of these movement segments is the STF. The time lag with respect to the neural activity is indicated by color. The blue segment represents the average movement made before

conclude that neurons in MI encode movements in two different coordinate frames. An alternative approach is to go no further than to say that the two parameters we happened to choose captured different aspects of the multi-dimensional movement fragment to which the cell was selective.

Abandoning the Essentialist Stance

The difficulties we have discussed so far suggest two possible directions for future motor control research. The first is to continue to devise experimental frameworks or analytic techniques that reduce the dependencies between a few chosen parameters of interest. A number of clever behavioral paradigms have been developed to attenuate or eliminate the correlations among certain movement parameters. For example, researchers have attempted to use loading experiments to uncouple kinematic from dynamic variables and have shown that primary motor cortical neurons are sensitive to force direction and not simply kinematic direction (Kalaska et al. 1989). In other experiments, researchers have attempted to uncouple movement direction from target direction (Alexander and Crutcher 1990) or Cartesian coordinates from joint coordinates (Kakei et al. 1999) and have shown that some motor cortical neurons appear to encode one or the other movement parameter or coordinate frame. A recent analytic technique has suggested a way to tease apart the contribution of different parameters of motion but is limited to considering only linear correlations among a few parameters (Stark et al. 2006). Despite the successes of these approaches, they do not address the larger issue that there are no definitive a priori grounds for choosing one set of parameters over another, and no way to eliminate statistical dependencies between all parameters simultaneously.

A second possible direction for research in motor control is to accept that it may not be especially meaningful to ask if any particular movement parameter is or is not encoded in motor cortex. In this view, motor cortical neurons encode idiosyncratic “movement fragments” characterized by multiple movement parameters that evolve in time, with the consequence that in any suitably high dimensional motor task, neurons that modulate their activity selectively for



Fig. 6 (continued) the neuron spiked, and the red segment represents the average movement made after the neuron spiked. The green segment represents the average positions in the two 50 ms bins simultaneous with and immediately subsequent to the spike occurrence. (D) The STF for the same cell as in 6C, but instead of Cartesian coordinates, average elbow and shoulder trajectories are plotted. (E,F,G) The instantaneous preferred direction was calculated by binning the direction every 50 ms into one of ten bins, and then calculating which of these ten directions tended to elicit the highest firing rate at each time lag ± 1 seconds with respect to the neural activity. The *arrow points* in the preferred direction of the cell at each time lag, and is plotted along the trajectory of the spike-triggered fragment for visualization purposes. Colors represent time lag as in A,B,C. (H) The instantaneous preferred direction for the same cell as in 6G, calculated in joint space (elbow and shoulder angles)

different movements will appear to be tuned along one or more axes. A similar point of view was originally proposed by Sherrington and colleagues but was later abandoned in favor of analytic theories that postulated that the motor system built movements out of rudimentary variables much like the visual system appears to do in visual perception.

Of course, we accept that there are plenty of reasons to think that characterizing the movement-related tuning of a cell along an arbitrary axis may be *useful* even if there is no reason to believe that tuning for a particular parameter is necessarily relevant or informative with regard to the neuron's functional role in generating movement. In fact, it is almost always necessary to restrict one's analysis to a few choice parameters. Most efforts to decode movements from cortical activity rely on tuning exclusively for position, velocity, or direction of movement in two or three-dimensional space (Serruya et al. 2002; Taylor et al. 2002; Wessberg et al. 2000; Wolpaw and McFarland 2004). In other experimental contexts it is often desirable to characterize a cell's tuning for a single parameter (such as direction) that evenly partitions the full space of possible movements. If appropriately chosen, a single parameter can provide a useful measure of movement similarity. Of course not all measures of similarity give the same results. For example, movements that are similar in direction are not necessarily equivalently similar in joint power (Scott et al. 2001). In fact, it is worth taking a closer look at directional tuning in motor cortex, because its significance is often taken for granted.

Directional Tuning Reconsidered

It would clearly be foolish to argue against the demonstrated utility of classifying cells according to their preferred direction. However, several authors have taken a more extreme position, arguing that directional tuning is itself more than a useful heuristic and that the robustness of directional tuning and the success of population vectors in decoding hand movements indicate that the directional tuning of MI neurons is intimately related to their function in generating movements. In this view, directionally-tuned motor cortical neurons are involved in generating a directional command to downstream areas. Again, the intuition driving this view may rely on an analogy with direction or orientation-tuned cells in the visual system, where such tuning appears to play a functional role in processing (Liu and Newsome 2005; Salzman et al. 1992). However, the observation that directional tuning is common and robust in MI is no evidence at all for the view that the directional tuning of MI neurons is intimately related to their function in generating movements. On the contrary – any behavioral paradigm which varies a parameter in such a way that it evenly spans the bulk of the movement space will cause that parameter to appear to be encoded in movement-sensitive neurons but will fail to capture other multi-dimensional sensitivities present in many motor cortical neurons (see Fig. 4).

To understand this assertion, it may be helpful to consider an analogy. Imagine a hypothetical case where neurons are sensitive to combinations of different food ingredients. For example, one neuron may respond to garlic and ginger, and another may respond to butter and pastry flour. Now consider a large population of neurons, where many combinations of ingredients are represented. Despite the fact that the actual sensitivities of neurons in this example are quite specific, when cells are arrayed according to their preferences for a general variable such as ethnic cuisine (i.e. the analogue of movement direction) that spans the space of possible foods, they will appear to be selective for different ethnic cuisines (such as Chinese and French cuisines respectively), because of the distribution of ingredients among foods of different ethnicities. For the same reason, these two cells would also appear to be differentially tuned to another variable such as spiciness that spans and partitions the food space. Concluding that either the “ginger-garlic” or “butter-flour” neurons were ethnically tuned (with respectively Chinese and French preferred ethnicities) – while legitimate in some respects – fails completely to capture the actual functional sensitivities of our hypothetical cells. In the case of tuning for motor parameters, the same may be said for a variable like direction.

Conclusion

When a novel domain of inquiry is in the process of developing, the form of the questions that are initially posed can shape the conceptual structure of the mature field. The early work on directional tuning was field-defining not just because it involved an important discovery, but because it immediately suggested a number of obvious questions of the form: “Is kinematic parameter X represented in the activity of motor area Y?” These questions are presumed to be important because their answers might allow us to ask the subsequent question: “How is parameter X operated on by area Y to produce motor output Z?” This sort of explanatory framework is central to the goals of computational neuroscience more generally, and as we have already indicated, has been central to the success of computational models of visual perception. Admittedly, a significant amount of progress has been made in the motor system along these lines as well. These successes notwithstanding, there are good reasons to think that our understanding of motor control may require a different interpretive framework – in part because of the unique problems associated with parameterizing movement that we have discussed. Although we do not touch on the subject here, it may be that the recent spate of research on optimal feedback control will provide the outlines of such an alternative framework.

One might argue that we will eventually develop experiments and techniques that will allow us to tease apart the contributions of specific movement parameters

to assess which variables are “really” being operated on in the motor system. After all, this approach has been highly successful in helping us understand how information is processed by different sensory areas. However, as we have tried to point out, motor parameters are importantly different from sensory parameters, and as a consequence we argue that motor physiologists should not try to take the kind of “essentialist” stance about the encoding of movement parameters that may be possible with regard to sensory variables. It may be more productive to analyze movements in a framework that references intentions, goals, and coordinated movements, rather than treating the movement of the arm as a problem in control theory and the motor cortex as a parametric controller (Graziano 2006).

We have argued that constructing parametric tuning curves may be practically indispensable, but that neurons are better viewed as participating in the generation of “movement fragments” than as controllers of particular parameters. We have focused on functional quantities like tuning and mutual information, but our argument also rests on an assumption about anatomy and neural wiring. One potential objection to our account might be that the very same correlations exist for robotic arms with many degrees of freedom as for monkey arms with many degrees of freedom, but we can know perfectly well which variables are important in controlling the robotic arm. By analogy then, why could we not eventually figure out which variables are being operated on by motor cortical neurons? The difference is that it is possible to design a robot such that each controller receives a distinct collection of input variables and produces a distinct motor command as output (for example with respect to movement around one joint). In making our argument, we assume that a motor cortical neuron’s inputs and outputs cannot be totally segregated according to the kind of information they transmit in this way. Cortical wiring can be tuned, no doubt, so that neurons with similar functional roles connect to similar downstream targets (Jackson et al. 2003), but the interconnected wiring of the brain makes it unlikely that any motor cortical neuron carries information or commands solely about a single aspect of movement. It is more likely that the neurons that are active during a movement are those with idiosyncratic multi-dimensional tuning properties that nonetheless contain specific similarities consistent with the observed behavior.

Methods

Electrophysiological recording. Silicon microelectrode arrays containing 100 platinized-tip electrodes (1.0 mm electrode length; 400 μ m inter-electrode separation; Cyberkinetics Inc., Salt Lake City, UT) were implanted in the arm area of primary motor cortex (MI) in two juvenile male macaque monkeys (*Macaca mulatta*). Signals were filtered, amplified (gain, 5000) and recorded digitally (14-bit) at 30 kHz per channel using a Cerebus acquisition system

(Cyberkinetics Inc.). Only waveforms (1.6 ms in duration) that crossed a threshold were stored and spike-sorted using Offline Sorter (Plexon Inc., Dallas, TX). Single units were manually extracted by the Contours and Templates methods. Inter-spike interval histograms were computed to verify single-unit isolation by ensuring that less than 0.05% of waveforms possessed an inter-spike interval less than 1.6 ms. The number of units in each electrode varied from one to five. To ensure good single unit isolation, we only studied the single units whose signal-to-noise ratio was larger than 3.

Random-Sequence task. The monkeys were operantly trained to perform a random-target pursuit task (RTP) by moving a cursor to targets via contralateral arm movements. The cursor and a sequence of seven targets (target size: 1 cm x 1 cm) appeared on a horizontal projection surface. At any one time, a single target appeared at a random location in the workspace, and the monkey was required to reach it within 2 seconds. As soon as the cursor reached the target, the target disappeared and a new target appeared in a new, pseudo-random location. After reaching the seventh target, the monkey was rewarded with a drop of water or juice. A new set of seven random targets was presented on each trial. The data used here included some trials where a viscous load was applied to the elbow and/or shoulder joints.

Mutual information estimates. Mutual information gives a quantitative measure of the reduction in uncertainty in one random variable provided by knowledge of another. The mutual information between the binned spike rate S and behavioral kinematics K can be defined as follows:

$$I(S; K) = H(K) - H(K|S), \text{ where } H(X) = -\sum_x p(x) \log(p(x))$$

The details of this calculation are described in (Paninski et al. 2004). We used a non-parametric binning approach to calculate the mutual information from the empirical distributions of spike rates and parameter values. To eliminate any bias due to errors in our approximation of the actual probability distributions, we calculated the information obtained from shuffled data and subtracted the result from our original estimate. All of the results shown here were calculated from data that was binned at 50 ms resolution. Mutual information with respect to the neural activity at time t was calculated individually for each parameter at each time lag from $t-1000$ ms to $t+1000$ ms.

References

- Alexander GE and Crutcher MD. Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *J. Neurophysiol.* 64: 164–178, 1990.
- Ashe J and Georgopoulos AP. Movement parameters and neural activity in motor cortex and area 5. *Cerebral Cortex* 6: 590–600, 1994.

- Cabel DW, Cisek P and Scott SH. Neural activity in primary motor cortex related to mechanical loads applied to the shoulder and elbow during a postural task. *J Neurophysiol* 86: 2102–2108, 2001.
- Cheney PD and Fetz EE. Functional classes of primate corticomotoneuronal cells and their relation to active force. *J Neurophysiol* 44: 773–791, 1980.
- Evarts EV. Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology* 31: 14–27, 1968.
- Fetz EE. Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences* 15: 679–690, 1992.
- Fu Q-G, Flament D, Coltz JD and Ebner TJ. Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J. Neurophysiol.* 73: 836–854, 1995.
- Fu Q-G, Suarez JI and Ebner TJ. Neuronal specification of direction and distance during reaching movements in the superior precentral premotor area and primary motor cortex of monkeys. *J. Neurophysiol.* 70: 2097–2116, 1993.
- Georgopoulos AP, Caminiti R and Kalaska JF. Static spatial effects in motor cortex and area 5: Quantitative relations in a two-dimensional space. *Experimental Brain Research* 54: 446–454, 1984.
- Georgopoulos AP, Kalaska JF, Caminiti R and Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. 2: 1527–1537, 1982.
- Graziano M. The Organization of Behavioral Repertoire in Motor Cortex. *Annu Rev Neurosci*, 2006.
- Graziano MS, Cooke DF, Taylor CS and Moore T. Distribution of hand location in monkeys during spontaneous behavior. *Exp Brain Res* 155: 30–36, 2004.
- Hepp-Reymond M-C, Wyss UR and Anner R. Neuronal coding of static force in the primate motor cortex. 74: 287–291, 1978.
- Jackson A, Gee VJ, Baker SN and Lemon RN. Synchrony between neurons with similar muscle fields in monkey motor cortex. *Neuron* 38: 115–125, 2003.
- Kakei S, Hoffman DS and Strick PL. Muscle and Movement Representations in the Primary Motor Cortex. *Science* 285: 2136–2139, 1999.
- Kalaska JF, Cohen DAD, Hyde ML and Prud'homme M. A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. 9: 2080–2102, 1989.
- Kurata K. Premotor cortex of monkeys: Set- and movement-related activity reflecting amplitude and direction of wrist movements. *Journal of Neurophysiology* 77: 1195–1212, 1993.
- Liu J and Newsome WT. Correlation between speed perception and neural activity in the middle temporal visual area. *J Neurosci* 25: 711–722, 2005.
- Moran DW and Schwartz AB. Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82: 2676–2692, 1999.
- Paninski L, Fellows MR, Hatsopoulos NG and Donoghue JP. Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology* 91: 515–532, 2004.
- Salzman CD, Murasugi CM, Britten KH and Newsome WT. Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12: 2331–2355, 1992.
- Scott SH, Gribble PL, Graham KM and Cabel DW. Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature* 413: 161–165, 2001.
- Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR and Donoghue JP. Instant neural control of a movement signal. *Nature* 416: 141–142, 2002.
- Smith AM, Hepp-Reymond MC and Wyss UR. Relation of activity in precentral cortical neurons to force and rate of force change during isometric contractions of finger muscles. *Exp Brain Res* 23: 315–332, 1975.

- Stark E, Drori R and Abeles M. Partial cross-correlation analysis resolves ambiguity in the encoding of multiple movement features. *J Neurophysiol* 95: 1966–1975, 2006.
- Taira M, Boline J, Smyrnis N, Georgopoulos AP and Ashe J. 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, 1996.
- Taylor DM, Tillery SI and Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. *Science* 296: 1829–1832, 2002.
- Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ, Srinivasan MA and Nicolelis MA. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408: 361–365, 2000.
- Wolpaw JR and McFarland DJ. Control of a two-dimensional movement signal by a non-invasive brain-computer interface in humans. *Proc Natl Acad Sci U S A* 101: 17849–17854, 2004.
- Wu W and Hatsopoulos NG. Evidence against a single coordinate system representation in the motor cortex. *Journal of Neurophysiology*, 2007.