Behavioral/Systems/Cognitive

Relationship between Unconstrained Arm Movements and Single-Neuron Firing in the Macaque Motor Cortex

Tyson N. Aflalo and Michael S. A. Graziano

Department of Psychology, Princeton University, Princeton, New Jersey 08544-1010

The activity of single neurons in the monkey motor cortex was studied during semi-naturalistic, unstructured arm movements made spontaneously by the monkey and measured with a high resolution three-dimensional tracking system. We asked how much of the total neuronal variance could be explained by various models of neuronal tuning to movement. On average, tuning to the speed of the hand accounted for 1% of the total variance in neuronal activity, tuning to the direction of the hand in space accounted for 8%, a more complex model of direction tuning, in which the preferred direction of the neuron rotated with the starting position of the arm, accounted for 13%, tuning to the final position of the hand in Cartesian space accounted for 22%, and tuning to the final multijoint posture of the arm accounted for 36%. One interpretation is that motor cortex neurons are significantly tuned to many control parameters important in the animal's repertoire, but that different control parameters are represented in different proportion, perhaps reflecting their prominence in everyday action. The final posture of a movement is an especially prominent control parameter although not the only one. A common mode of action of the monkey arm is to maintain a relatively stable overall posture while making local adjustments in direction during performance of a task. One speculation is that neurons in motor cortex reflect this pattern in which direction tuning predominates in local regions of space and postural tuning predominates over the larger workspace.

Key words: primary motor; reaching; grasping; posture; direction tuning; motor control

Introduction

Neurons in the motor cortex of the monkey brain are active during arm movements. This activity is correlated with a range of movement parameters, including force, muscle activity, the direction, speed, and position of the hand, joint rotation, the multijoint posture of the arm, and other aspects of movement (Evarts, 1968; Georgopoulos et al., 1982, 1984, 1986, 1992; Cheney et al., 1985; Kettner et al., 1988; Caminiti et al., 1990; Hocherman and Wise, 1991; Fu et al., 1993; Ashe and Georgopoulos, 1994; Scott and Kalaska 1995, 1997; Kakei et al., 1999; Moran and Schwartz, 1999; Reina et al., 2001; Holdefer and Miller, 2002; Sergio and Kalaska, 2003). One possibility is that the search for the "correct" motor parameter is an ill-posed experimental question because many different movement parameters are correlated with motor cortex neurons. Recently, it has been suggested that the motor system uses an "optimal control" method, potentially controlling any parameter, whether spatial, joint based, or muscle based, if it is useful for the performance of the current task (Todorov and Jordon, 2002).

One limitation in previous studies is a tendency to use highly constrained movement sets and animals that are extensively trained. We recorded from neurons in the motor cortex of untrained monkeys. The arm was free to move spontaneously and naturalistically while the movements were measured. One goal of

this study was to compare the results of a naturalistic, untrained movement set with the previous results obtained with structured, trained movement sets. Can the previous findings be replicated in the present movement set, or will removing the behavioral control also remove the previously obtained neuronal tuning curves?

A second goal of this study was to assess the proportion of total neuronal variance attributable to different movement parameters. In most previous experiments, the constrained movement sets were designed to focus on one or a small number of movement parameters. Once all other sources of variance have been removed, minimized, or averaged out of the data, then the particular parameter under study may account for most of the remaining neuronal variance, with an R^2 value that may be as high as 0.9. Such studies, although legitimately addressing a variety of questions, do not address a certain fundamental question. If most sources of neuronal variance are left in the data by using a more natural range of movement, then what proportion of the total neuronal variance can be explained by any single parameter? Are neurons primarily tuned to one parameter, with the majority of their total variance explained by that parameter, or instead are neurons partially tuned to a diversity of parameters, with each parameter capturing a small portion of the total variance?

Of the 25 analyses included here, five overlap with a preliminary report (Aflalo and Graziano, 2006). These five analyses were updated and included in the present report because they provide the necessary context for the full set of analyses.

Materials and Methods

All procedures were approved by the Princeton University Institutional Animal Care and Use Committee and the attendant veterinarian and were in accordance with National Institutes of Health and United States
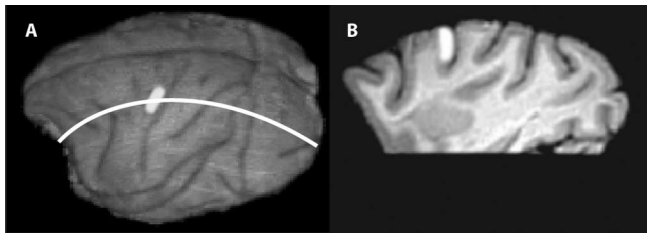
Figure 1. Images of the brain of monkey 1 showing the studied area of cortex. ***A***, MRI reconstruction of the cortical surface showing sulcal pattern. Studied region is in lightened shading. The white line indicates the section shown in ***B***. ***B***, Parasagittal section (0.5 mm thick) through the studied area of cortex. Studied region is in lightened shading.

Department of Agriculture guidelines. We studied the motor cortex in the left hemispheres of two adult male *Macaca fascicularis*.

Surgery. For each monkey, an initial surgical operation was performed under isoflurane anesthesia and strict aseptic conditions, during which an acrylic skullcap was fixed to the skull with bone screws. A steel bolt for holding the head and a 2.5 cm diameter steel chamber for neuronal recording were also imbedded in the acrylic. The recording chamber was positioned for a vertical (dorsoventral) approach to the precentral gyrus. Each animal recovered from the surgery within 1 week but was given 2 additional weeks to allow the skull to grow tightly around the skull screws. In a subsequent procedure, also under deep anesthesia and aseptic conditions, the recording chamber was opened and a hole ~10 mm in diameter was drilled through the layer of acrylic and the bone, exposing the dura.

Neuronal recording. During the daily recording sessions, the monkey sat in a Lexan primate chair with the head restrained by the head bolt. A hydraulic microdrive (Narishige, Tokyo, Japan) was mounted to the top of the recording chamber. A steel guide tube (a 22 gauge syringe needle) was lowered through the hole in the skull and into the granulation tissue that lay over the dura. Then the varnish-coated tungsten microelectrode (impedance of 0.5–2 MΩ; Frederick Haer Company, Bowdoinham, ME) was advanced from the guide tube through the dura and into the brain. Neural signals were amplified (model 1800 amplifier; A-M Systems, Carlsborg, WA), filtered (300–5000 Hz), and recorded at 25,000 Hz. An off-line spike-sorting algorithm was used to assign spikes to individual neurons. Typically one to three neurons could be reliably isolated on the electrode at one time.

Single-neuron activity was sampled at various depths within the motor cortex ranging from the first depth at which any neurons could be isolated to the depth at which the neurons disappeared and the white matter was presumed to begin. A systematic test of different layers of cortex was not attempted in the present experiment. We saw no clear distinction in the tuning properties reported here for shallow or deeper recordings within motor cortex and therefore included all neurons into one analysis population. Neurons were not preselected in any way based on response properties. Instead, all neurons that were encountered by the electrode and that could be held long enough for collection of a full dataset were included.

Location of recording sites. Neither monkey was killed at the termination of this experiment. For monkey 1, after the experiment, the implant was removed and the brain was scanned in a 3-T Siemens (Munich, Germany) Allegra magnetic resonance imaging (MRI) head scanner using a 16 cm birdcage coil (NMSC-023; Nova Medical, Wakefield, MA). A high-resolution, 0.5 × 0.5 × 0.5 mm anatomical scan of the whole brain was taken (magnetization-prepared rapid-acquisition gradient echo sequence; field of view, 128 × 128 mm; matrix, 256 × 256; repetition time, 2500 ms; echo time, 4.4 ms; flip angle, 8°). Figure 1*A* shows a surface reconstruction of the cortex with the studied area in brightened shading. Figure 1*B* shows a sagittal section through the motor cortex with the studied area again in brightened shading.

To find the motor cortex before recording, the central and arcuate sulci were located by shining a bright light on the dura during the initial craniotomy surgery. Both sulci were clearly visible through the dura. The microdrive was then mounted to the recording chamber, and the loca-

tions of the visualized sulci were measured with the tip of the guide tube. In this way, the locations of the sulci were obtained in microdrive coordinates.

After the craniotomy surgery, during the daily experiments, the measured location of the central sulcus was confirmed to within 0.5 mm by recording and stimulating to either side of the sulcus. Just posterior to the sulcus, in primary somatosensory cortex, we observed the expected small tactile receptive fields on the contralateral limb and also the expected rarity of effect of intracortical electrical microstimulation even with currents as high as 200 μA (negative leading biphasic pulses, 0.2 ms phase width, 200 Hz, 100–500 ms train durations, 5–200 μA current). Just anterior to the sulcus, we obtained neuronal responses during hand and arm movement and low microstimulation thresholds, typically <20 μA, sometimes as low as 5 μA, as expected from primary motor cortex. The location of the arcuate sulcus was confirmed by microstimulating just anterior to it and obtaining no skeletomotor movements, but instead obtaining stimulation-evoked saccadic eye movements presumably in the frontal eye fields.

The sites tested were located in the arm representation in motor cortex and were within the anterior bank of the central sulcus or on the cortical surface within 2 mm of the central sulcus. They therefore lay within the boundaries of traditional primary motor cortex. Additional information on the location of sites within the functional map was provided by the analyses of the joint angles to which the neurons were responsive, as described below.

Measurement of joint angles. The positions of points on the limb contralateral to the recording electrode were measured by means of an Optotrak 3020 system (Northern Digital, Waterloo, Ontario, Canada). This system tracks the three-dimensional position of infrared light emitting diodes (LEDs). Each LED could be separately tracked to a spatial resolution of 0.1 mm. The position was measured every 14.3 ms. To create a marker that could be detected by the Optotrak cameras from any angle, we glued five individual LEDs together to produce an omni-directional marker ball. A marker ball was taped to the monkey's forefinger on the dorsal surface so it would not interfere with grasping, on the thumb, again on the dorsal surface so it would not interfere with grasping, on the back of the hand between the knuckles of the third and fourth digits, on the lateral aspect of the elbow, and on the shoulder. In addition, 14 individual markers were taped in a double ring around the monkey's wrist, with seven markers per ring and a 1 cm spacing between the rings. A marker was taped to the side of the primate chair to calibrate the position of the monkey with respect to the chair. For the LEDs attached to the arm and hand, the wires were taped in a bundle to the underside of the arm and draped behind the monkey. The primate chair was open at the front and side, allowing for almost total range of movement of the arm. The monkey's other arm, ipsilateral to the electrode, was not studied with Optotrak markers. To ensure that this hand did not reach for the fruit rewards during trials or tear off the markers taped to the measured hand, the untested hand was fixed to the side of the chair in an arm holder.

The double ring of 14 markers around the wrist was subjected to a rigid body computation to calculate the location and spatial orientation of the wrist. In this computation, for each time point, a three-dimensional rigid model of the double ring of markers was fitted to the measured positions of the currently visible markers, using a least-squares method of optimal fit. The orientation and position of the model could then be used to estimate the orientation and center of the wrist. The center of the wrist was taken to be the mean position of the 14 points in the model.

The position of the shoulder in space was calculated by analyzing the position of the elbow over time. Over many time points, the elbow described a portion of a sphere, the origin of which was located at the shoulder joint. For each 3 min block of data, a shoulder position was calculated by fitting a sphere to the data using a least-squares best-fit algorithm and using the center of the sphere as the shoulder location. Because the shoulder is capable of small translational movements in addition to rotations, this method of estimating shoulder joint location is approximate but was sufficient for the purposes of this study. When the shoulder position was calculated multiple times over different time segments, it varied within <3 cm.

Three shoulder angles were computed: the elevation, the azimuth, and

the "twist" or internal/external rotation of the shoulder joint. We also calculated the flexion of the elbow, the pronation of the forearm, the extension of the wrist, the adduction of the wrist, and the grip aperture. In total, eight degrees of freedom were calculated for the arm. This model of the arm was verified by applying forward kinematics to estimate the position of the hand. This calculated position of the hand matched the actual, measured position of the hand within 1.5 cm.

Description of movements in the dataset. During testing of a neuron, the monkey was allowed to move its contralateral arm freely to touch and explore parts of the primate chair, to reach for small pieces of fruit held out on the end of forceps, to bring food to its mouth, to retrieve food from its mouth, to hold and examine food in central space, and to rotate and explore food items. Occasionally, the monkey also scratched at its skin, scratched rhythmically at a portion of the monkey chair, or attempted to scratch the experimenter with a fast semi-ballistic arm movement. The movement of the arm was recorded through all of these behaviors. Different types of behaviors were not separated in the analysis, partly because one type of behavior tended to grade into another type and the distinction between behaviors could only be made subjectively and partly because the purpose of the study was to include all arm movements in as large and naturalistic a range as possible given the constraints of the primate chair.

For each neuron, the position of the hand in three-dimensional space was measured every 14.3 ms during a continuous time interval ranging from 10 to 30 min. The hand position data were smoothed using a cubic spline with a smoothing coefficient of 0.05. The instantaneous hand speed was then calculated using a three-point running average to obtain a smoothed speed profile. Separate movements were extracted from the dataset on the basis of a speed analysis. Minima in the speed were identified and the intervals between minima were flagged as potential separate movements. To enter the final dataset, the movement had to be at least 150 ms in duration and the peak speed had to be at least 20 cm/s. These parameters seemed to successfully divide the data into discrete segments that matched our subjective impression of separate hand movements. The average speed profile is further described in Results (see Speed tuning I).

The various analyses in this report focus on the neuronal activity associated with periods of arm movement as defined by the velocity analysis described above. We originally attempted to analyze both the periods of movement and the interleaved periods of non-movement but found that neurons were frequently inactive during periods of hand stasis between movements. This lack of activity may be related to the spontaneous nature of the movements. There may have been little planning between movements. Sometimes the arm was relaxed or braced against a part of the chair between movements, and this state of relaxed stasis could not be distinguished from more active stasis in the recorded data. Often the animal performed a mouth movement or a foot movement between arm movements, perhaps drawing attention and motor planning to a different body part. These reasons may have contributed to a reduction in and a variability of neuronal activity between movements. The analyses reported here are therefore focused on the periods of arm movement.

Figure 2*A* shows a typical movement set collected during the testing of a neuron. This set shows 514 separated movement segments that densely sampled the workspace of the hand. Vertically, the movements ranged from 29.7 cm below the mouth to 7.8 cm above the mouth. Horizontally, the movements ranged from 19.4 cm on the contralateral side (same side as the studied arm, opposite side to the electrode) to 16.6 cm on the ipsilateral side. In depth (direction along the monkey's forward line of
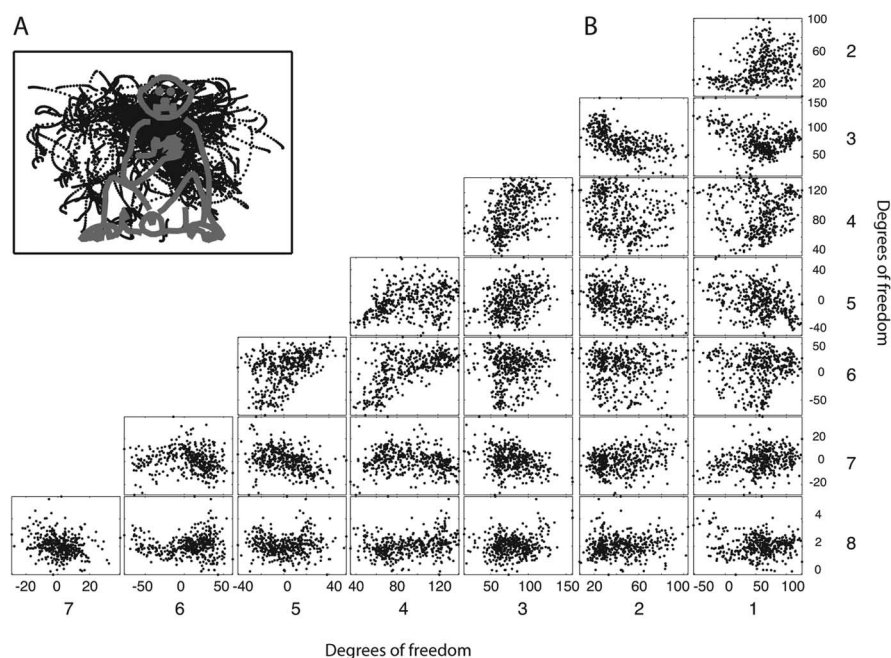

Figure 2. Range of movements in the naturalistic movement set. *A*, Front view of 514 hand movements made during 10 min of testing one neuron. Each trail of dots is equivalent to one movement measured at 14.3 ms intervals. Frame is 45 cm tall. *B*, Correlations among the final positions of the eight degrees of freedom among the recorded movements. Degrees of freedom are as follows: 1, shoulder azimuth; 2, shoulder elevation; 3, shoulder internal/external rotation; 4, elbow extension; 5, forearm pronation; 6, wrist extension; 7, wrist adduction; 8, grip aperture. Units on *x*- and *y*-axes of each plot are degrees of joint angle, except for degree of freedom 8, which is expressed in centimeters.

sight), the movements ranged from 7.5 cm behind the level of the mouth (such as when the monkey reached to its flank or to its ear) to 20.2 cm in front of the mouth (normal for a fully extended reach). The range of movement starting points was not significantly different from the range of ending points.

The average ± SD length of a movement was 8.0 ± 4.9 cm. The average ± SD duration was 351 ± 133 ms. The average ± SD hand speed was 23.74 ± 9.9 cm/s. Each movement had a peak speed, and the average ± SD peak speed among all movements was 40.7 ± 20.4 cm/s. The average time between movements was 4.06 s. This pause between movements ranged from <1 s (40% of pauses, representing brief periods of stasis during ongoing arm movement) to longer than 10 s (15% of pauses, representing periods when the monkey had stopped making arm actions and was stationary or engaged in movement of other body parts).

For each movement, we calculated a standard curvature metric as follows. The straight-line distance between the start and end of the movement was found. The total path length of the movement was found. The ratio of these two quantities provided a curvature metric in which 1.0 corresponds to no curvature and smaller numbers correspond to increasingly curved movements. The average ± SD curvature was 0.91 ± 0.09, indicating that the movements tended to be straight. In Figure 2*A*, some movements appear to be highly curved. This appearance is a result of collapsing a three-dimensional movement into a two-dimensional depiction in which the long axis of the movement is not fully shown.

For each movement, we calculated a direction by connecting the start point to the end point and obtaining the azimuth and elevation angles. The distribution of these directions was then examined. The directions appeared to be relatively evenly distributed. The sphere of all possible directions was divided into 20 equal sectors, and the movement directions were distributed over these 20 sectors with all sectors represented. Thus, the unstructured movement set was well behaved, followed expected characteristics of normal movement, and included a diverse range in terms of workspace, length, speed, and direction.

Inevitably, the movement set included some non-uniformities in the distribution of hand positions and joint angles. For example, hand posi-

tions near the face were overrepresented. Hand positions in lower space were less well represented. Because the movement sets were spontaneous, we were not able to enforce an even distribution of all types of movement. As a result, there is some concern whether the dataset is diverse enough, and the movement parameters independent enough, to allow for meaningful regression analyses. One approach to this issue is addressed in Results (see Simulated neurons). We simulated noisy neurons that were tuned in a variety of ways, calculated the expected response of these neurons during the actual recorded movement sets, and then applied the regression models to the results. We found that the movement sets were sufficiently diverse to allow the regression analyses to uncover the correct tuning curves.

Non-uniformities in the movement set present a particular problem for separating the effect of direction tuning from the effect of hand-position tuning. For example, if a neuron is tuned to upward motion of the hand, it may appear to prefer an upper hand location simply because that location in space can only be reached by upward motion. This possible confound was addressed partly through the use of the simulated neurons (simulated direction-tuned neurons did not evidence a high degree of tuning to hand position) and partly by means of the analysis of neuronal data in Results, showing that the preferred direction does not easily account for the hand-position tuning (see End-point tuning).

Non-uniformities in the movement set may also result in correlations among the different joints of the arm. Such interjoint correlations would be problematical particularly in the analysis for neuronal tuning to final posture. A neuron may actually be related to one joint but, because of interjoint correlations, appear to be related to other joints as well. It was therefore important that the movement set contain enough independence among joints that the contributions of each joint could be separated through regression analysis. Figure 2B shows the correlations among the joint angles reached at the end of the movements. The joints were not tightly coupled. Some showed correlations with each other, as expected given the known covariance of joints during movement, but the scatter was easily sufficient to allow a regression analysis to separate the different contributions. This lack of any strong coupling among joints was probably caused by the great diversity of movements in the set and is presumably the reason for the success of the simulated neuron analysis in which simulated tuning functions could be successfully recovered using our regression analyses. This issue is further addressed in Results.

Preliminary analysis to specify somatotopic location. To further confirm the somatotopic portion of motor cortex that was studied, we performed a preliminary analysis on each neuron. Using a stepwise regression, we obtained the degree of correlation between the neuronal activity and the velocity of each of the eight measured joints. If we were recording primarily in a distal representation, we would expect to find significant regressions with distal joints including hand aperture, wrist flexion, wrist adduction, and forearm pronation. If we were recording primarily in a proximal representation, we would expect to find significant regressions with proximal joints, including elbow flexion and the three degrees of shoulder rotation. Given the known overlap in motor cortex somatotopy, we expected to find neurons related to both proximal and distal joints. The results indicated that 89% of the neurons were significantly related to the proximal joints and 67% were significantly related to the distal joints. These results indicate that the studied neurons were in the forelimb representation in a region that emphasized the proximal joints over the distal joints but represented both.

Regression models. Each neuron was tested for a range of possible tuning functions. Each model of a tuning function was compared with the behavior of the neuron by means of a regression analysis. Some models, such as tuning to speed, involved one free parameter. Other models, such as tuning to posture, involved several free parameters. To avoid inflating the R^2 value with the addition of more parameters, we used the standard adjusted R^2 metric that takes into account the number of regressors

(Cohen et al., 2003). Also following standard practice, we performed a regression only if the number of data points was 20 or more times the number of regressors, thus minimizing the risk of overfitting the data. Finally, as a validity check, we performed an analysis on simulated neurons (see Results, Simulated neurons). For example, we tested a simulated neuron that was direction tuned. When tested on an actual, recorded movement set and with our regression models, the neuron yielded a high R^2 value for direction tuning. It yielded a low R^2 value for postural tuning, despite the larger number of regressors in the postural analysis. Likewise, a simulated neuron with no tuning signal yielded a low R^2 value on all regression models regardless of the number of regressors in the models.

For all regression models described below, we followed a standard procedure in which we offset the movement data from the neuronal spike data by a latency. In this procedure, the kinematic data were taken from the movement period defined by the velocity analysis and the neuronal spike data were taken from an equivalent time period offset by a latency. We tested latencies between −286 and 286 ms at 14.3 ms increments (the temporal resolution of the movement data) and used the result of the regression analysis to choose the latency that optimized the R^2 between neuronal activity and movement. The reason for this approach is that, in general, a latency will exist between the activity of a cortical neuron and the movement of the arm, and different neurons may have different latencies. This analysis differs from the analysis we described previously (Aflalo and Graziano, 2006) in which we used a fixed latency based on the results of electrical stimulation of each site in cortex. In the present paper, because of the optimization of latency, the regression analyses resulted in slightly higher R^2 values than in our previous publication.

In most analyses, the neuronal firing rate was averaged over the duration of each movement, and therefore no additional smoothing of the neuronal data was required. However, for the analyses of hand speed, the instantaneous neuronal firing rate was required. For these analyses, we smoothed the firing rate data using a 10 Hz upper cutoff. The reason for using this frequency cutoff was to ensure that our results were directly comparable with previous studies that used a similar technique (Moran and Schwartz, 1999).

The four main regression models used in this paper are described below, and other models are described in the relevant subsections of Results.

Direction tuning. Each neuron was tested for direction tuning in the following manner. For each movement, we calculated a mean firing rate of the neuron (spikes per second during the movement, offset by a latency as described above). Each hand movement was assigned a direction in Cartesian space based on the vector connecting the beginning and end point of the movement. Firing rate was modeled as a function of the angular deviation ($\Delta\theta$) between this movement vector and a preferred direction: firing rate = $A \cos(\Delta\theta) + B$. A regression analysis was used to find the optimal preferred direction and coefficients, following the method of Georgopoulos et al. (1986). The regression analysis provided an R^2 value indicating how much of the variance in neuronal activity could be attributed to the direction-tuning model.

Direction tuning that rotates with the start position of the hand. Each neuron was also tested with a second direction-tuning model. In this model, the preferred direction vector (PD) was obtained using a cosine tuning model of neuronal firing rate, as above. However, PD was not assumed to be fixed in all regions of the workspace. Instead, it was assumed to vary depending on the start position of the arm. The start vector (ST) was defined as the vector pointing from the shoulder to the position of the hand at the start of the movement. The direction of ST was described by two parameters: the start azimuth θ and the start elevation ϕ. PD was assumed to vary linearly with ST. In this model, the azimuth of PD was linearly related to the azimuth of ST with a slope defined by the parameter AZ, and the elevation of PD was linearly related to the elevation of ST with a slope defined by the parameter EL. Thus, if both AZ and EL are 0, then PD is indeed independent of ST and always points in the same direction regardless of starting position. If AZ and EL are both 1,

then PD rotates exactly in tandem with ST, remaining at a fixed spatial relationship with respect to the arm. In this model, the equation for PD as a function of AZ, θ, EL, and ϕ, is as follows:

$$\text{PD} = \begin{pmatrix} \cos(\text{AZ}\theta)\cos(\text{EL}\phi)) & \sin(\text{AZ}\theta) & \cos(\text{AZ}\theta)\sin(\text{EL}\phi) \\ \sin(\text{AZ}\theta)\cos(\text{EL}\phi) & -\cos(\text{AZ}\theta) & \sin(\text{AZ}\theta)\sin(\text{EL}\phi) \\ \sin(\text{EL}\phi) & 0 & \cos(\text{EL}\phi) \end{pmatrix} \begin{pmatrix} \text{PD}_x^0 \\ \text{PD}_y^0 \\ \text{PD}_z^0 \end{pmatrix},$$

where PD^0 indicates the preferred direction when the hand is located at $\theta = 0$ and $\phi = 0$.

End-point tuning. For this model of neuronal tuning, all data concerning the direction or trajectory of the movement was discarded and only the end point of the movement was considered. Firing rate was modeled as a Gaussian function of these end points in Cartesian space. In the following equation, x_1, x_2, and x_3 refer to the three Cartesian coordinates of the end point of the movement, P1, P2, and P3 refer to the coordinates of the peak of the Gaussian, the SDs of the Gaussian around that peak are indicated by $\sigma 1$, $\sigma 2$, and $\sigma 3$, the height of the Gaussian is given by A, and the floor of the Gaussian is B. A nonlinear regression technique (Bates and Watts, 1988) was used to fit this equation to the data for each neuron:

$$\text{Firing rate} = Ae^{\frac{(x_1-P_1)^2}{2\sigma_1^2}+\frac{(x_2-P_2)^2}{2\sigma_2^2}+\frac{(x_3-P_3)^2}{2\sigma_3^2}} + B.$$

End-posture tuning. This model followed the same general equation as the end-point model except that it involved the eight dimensions of arm-posture space (x_1 through x_8) rather than the three dimensions of Cartesian space. Firing rate was modeled as a Gaussian function that had a peak at a specific, preferred posture. Again, a nonlinear regression technique was used to fit the model to the data for each neuron:

$$\text{Firing rate} = Ae^{\sum_{i=1:8}\frac{(x_i-P_i)^2}{2\sigma_i^2}} + B.$$

Results
Speed tuning I
We first tested whether previous standard methods of obtaining speed tuning from motor cortex neurons would also reveal speed tuning in the present, unconstrained movement set. In the first test of speed tuning, we used a procedure based partly on that of Moran and Schwartz (1999). In this method, for each neuron, many movements are averaged together to produce a mean speed profile and a mean firing-rate profile. The two profiles are then compared with each other using a regression analysis to determine whether the average hand speed tends to track the average firing rate throughout the movement.

Figure 3 shows the results for one example neuron. The data on the changing position of the hand through time was segmented into separate movements (for details, see Materials and Methods). Because of the unconstrained nature of the animal's behavior, different movements were of different durations. To average across movements, we first normalized the length of movements by dividing each movement into 20 time bins. For each time bin, the instantaneous firing rate and the hand speed were calculated. These numbers were then averaged across the 321 movements tested for this neuron. The thin line in Figure 3A shows the average hand speed rising and falling during the movement in a smooth, bell-shaped velocity profile typical of normal movement. The thick line shows the average neuronal activity, also rising and falling during movement.

We performed a regression analysis to compare the average hand speed with the average firing rate. Following a standard method, we used a fixed time window of kinematic data that corresponded to the hand movement and used a time window of equal duration but adjustable start time for the neuronal data. The temporal offset between the neuronal and kinematic data were then optimized to yield the maximum R^2 value. A similar
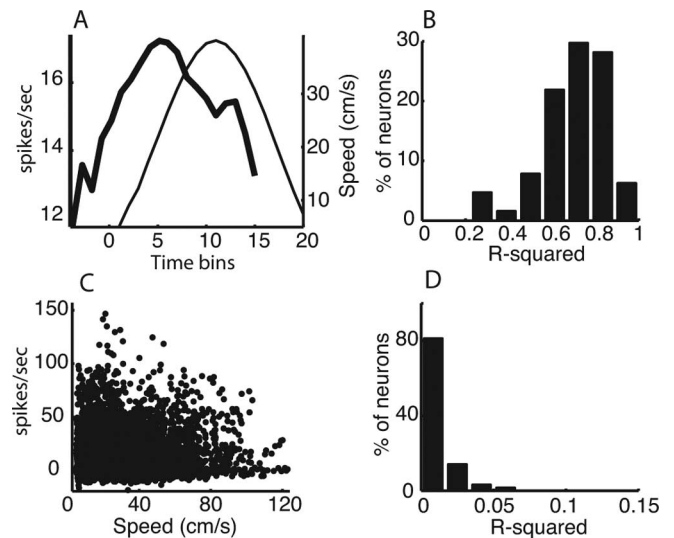
Figure 3. Speed tuning of motor cortex neurons. ***A***, Firing rate and hand speed averaged over 321 movements for one example neuron. When neuronal data were shifted forward by 71 ms (the optimal time lag for this neuron), the two curves matched closely with an R^2 of 0.87. ***B***, Frequency histogram of R^2 values for all 64 neurons tested as in ***A***. ***C***, Firing rate versus hand speed for 9018 time bins during hand movement studied for one example neuron. The two variables showed a small but highly significant correlation ($R^2 = 0.03$; $p = 1.5 \times 10^{-17}$). ***D***, Frequency histogram of R^2 values for all 64 neurons tested as in ***C***.

offset optimization was used for all analyses (see Materials and Methods). For this neuron, the regression analysis returned the largest R^2 value at a temporal offset of -71 ms with neuronal activity preceding movement. The R^2 value for this example neuron was 0.87, and the regression was highly significant ($p = 3.5 \times 10^{-9}$).

Figure 3B shows a histogram of R^2 values for all 64 cells tested in this manner. The mean R^2 value was 0.71, and 93% of the neurons had a statistically significant regression against speed (0.05 significance level, corrected for the number of neurons using the Bonferroni method). These findings approximately match the findings of Moran and Schwartz (1999) who also reported a high degree of speed tuning. The results indicate that the same type of neuronal tuning to speed found in previous experiments can be obtained in a freely moving condition.

The above method of testing for speed tuning has a limitation. The method involves averaging across many movements to obtain relatively smooth, well behaved curves from which the between-movement variance has been removed. If the averaging is not performed and therefore the full range of neuronal variance is left in the analysis, what percentage of the total variance will be attributable to speed tuning? This question is addressed in the next section.

Speed tuning II
To examine speed tuning in the raw data without using averaging to remove variance, we used a technique based partly on that of Ashe and Georgopoulos (1994). Figure 3C shows the result for one example neuron. For this neuron, 339 movements were studied. These movements, placed end to end with periods of non-movement removed, spanned 129 s. During this span of time, did the instantaneous firing rate of the neuron correlate with the instantaneous hand speed? We treated each 14.3 ms time bin within each movement as a distinct data point in the regression. For this neuron, 9018 time bins were used. The regression between neuronal activity and hand speed (offset by an optimized
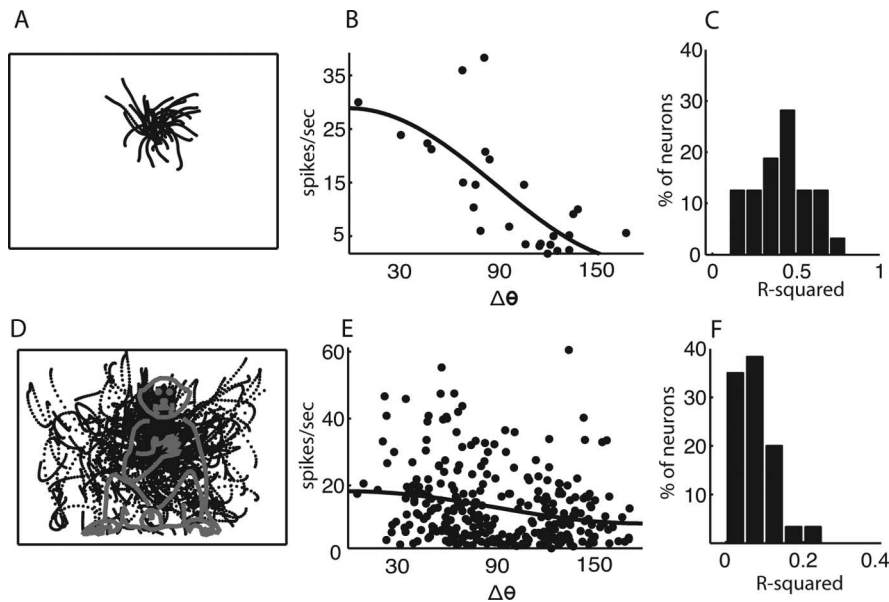
Figure 4. Direction tuning of motor cortex neurons. **A**, Front view of 26 selected hand movements made during 10 min of testing one neuron. Each trail of dots is equivalent to one movement measured at 14.3 ms intervals. Frame is 45 cm tall. Each movement shown originated within a 5 cm radius sphere of central space and was between 6 and 15 cm in length. **B**, Tuning of an example neuron to direction, based on selected movement set. x-Axis shows angular difference between the direction of each movement and the preferred direction; y-axis shows mean firing rate during each movement; for cosine tuning to direction, $R^2 = 0.43$, $p = 0.0001$. **C**, Frequency histogram of R^2 values for all neurons tested as in **B**. **D**, Front view of full set of 320 hand movements made during testing of one neuron. **E**, Direction tuning of an example neuron (same neuron as in **B**), based on full movement set. $R^2 = 0.05$, $p = 0.00008$. Note that a new preferred direction was obtained by regression, and therefore the data points shown in **B** do not plot to the same location on the x-axis as in **E**. **F**, Frequency histogram of R^2 values for all neurons tested as in **E**.

latency; see Materials and Methods) was highly significant ($p = 1.5 \times 10^{-17}$), yet the R^2 value was 0.03. Of the variance in neuronal activity during hand movement, only 3% could be attributed to speed tuning.

Figure 3D shows the R^2 values for all 64 cells tested. The mean R^2 was 0.01, and 98% of the neurons had a statistically significant regression (0.05 significance level, corrected for the number of neurons using the Bonferroni method). These results show that almost all neurons are tuned to hand speed and that this tuning accounts for a small proportion of the total neuronal variance.

Many different analytic methods can be used to test for speed tuning. We tested two additional regression analyses with similar results. We compared the peak speed of each movement (hand speed during the 14.3 ms time bin around the peak speed) with the firing rate associated with that peak speed (during the 14.3 ms time bin offset from the peak speed by a latency optimized for each neuron). This regression returned a mean R^2 value of 0.03. We also compared the mean speed during each movement with the mean firing rate associated with the movement (during a time window of the same duration as the movement and offset by a latency optimized for each neuron). This regression returned a mean R^2 value of 0.02. All of these methods converged on a similar result: the neuronal firing rate varies from movement to movement, and 1–3% of this variance is attributable to speed tuning. Only when the between-movement variance is removed by averaging (as above, Speed tuning I) will the remaining variance be well correlated with hand speed. These results do not show that neurons have no real speed tuning. Rather, they are genuinely speed tuned, accounting for a small proportion of the total neuronal variance. Previous studies report that motor cortex neurons are tuned to both speed and direction. Can a large

proportion of the remaining variance be attributed to direction tuning? The following sections examine tuning to direction and also to velocity.

Direction tuning I: local
Figure 4A shows 26 movements selected from the 320 spontaneous movements performed by the monkey during the testing of one example neuron. These 26 movements were selected from the full set of movements on the basis of start location (within a central ball of space with radius 5 cm) and length (between 6 and 15 cm). These movements therefore roughly approximated the center-out movement set commonly used to test direction tuning in motor cortex neurons (Georgopoulos et al., 1986). For each of these movements, we defined the direction to be the vector pointing from the start location to the end location of the movement. Also for each movement, a mean neuronal firing rate was calculated. The mean firing rate was based on a time interval equal to the movement interval but offset from the movement interval by an optimized temporal lag (see Materials and Methods). We performed a regression analysis to determine the preferred direction of the neuron. We used a standard cosine tuning model (Georgopoulos et al., 1986) in which the firing rate of the neuron is proportional to the cosine of $\Delta\theta$, the angle between the preferred direction and the actual direction of movement.

Figure 4B shows the result of this analysis for an example neuron. On average, the firing rate was high during movements near the preferred direction (thus with low $\Delta\theta$) and low during movements far from the preferred direction (with high $\Delta\theta$). The regression fitting to the cosine model returned an R^2 value of 0.43 and was highly significant ($p = 0.0001$).

Figure 4C shows the R^2 values for all cells tested. The mean R^2 value was 0.42, and 68% of the cells showed a significant fit to the cosine model of directional tuning (0.05 significance level, corrected for the number of neurons using the Bonferroni method). The present results are therefore similar to previous results using a center-out task, in that most neurons showed a significant fit to a cosine tuning function.

Just as in the case for speed tuning, the R^2 values obtained for the present analysis showing strong direction tuning must be interpreted cautiously. The movement set is limited such that the direction of movement represents a main source of variance and thus results in a relatively high R^2. Other sources of variance, including differing starting positions, starting postures, and movement distances, although present in this selected dataset, are minimized. If the entire range of movements were used and, thus, all sources of variance were admitted to the data, how much of the total variance would be attributable to direction tuning? This question is addressed in the next section.

Direction tuning II: global
Figure 4D shows the full set of 320 movements performed by the monkey during the testing of one example neuron. We per-

formed a regression analysis to determine how well the neuronal firing during the full range of movements fit a cosine model of direction tuning. Figure 4*E* shows the result for the same example neuron shown in Figure 4*B*. The regression fit to the cosine model was highly significant ($p = 0.00008$). The R^2 value for the global movement set, however, was 0.05. Thus, direction tuning accounted for only ~5% of the total variance in the behavior of this neuron. Figure 4*F* shows the R^2 values for all 64 cells tested. The mean R^2 value was 0.08, and 63% of the cells showed a significant fit to the cosine model of directional tuning (0.05 significance level, corrected for the number of neurons using the Bonferroni method).

These results show that most neurons in motor cortex are indeed direction tuned, but by itself direction tuning accounts for only ~8% of the global neuronal variance. Other parameters must account for the remaining variance. Numerous studies have shown that hand position and arm posture have a large effect on the activity of motor cortex neurons (Kettner et al., 1988; Caminiti et al., 1990; Scott and Kalaska, 1995, 1997; Sergio and Kalaska, 2003). These effects of posture may be even greater in the present dataset because a greater range of postures was included.

Direction tuning III: rotation of preferred direction with changes in starting hand position

The previous section examined whether each neuron was tuned to a single preferred direction. One possibility, however, is that a single preferred direction might not capture the full extent of direction tuning for a typical neuron. Perhaps direction tuning would capture a greater share of the total neuronal variance if we used a more sophisticated tuning model in which the preferred direction rotates systematically as the starting position of the hand changes.

Caminiti et al. (1990) tested direction tuning in motor cortex neurons when the hand was placed at three different starting positions separated horizontally. Averaged across neurons, the preferred direction rotated in the same direction that the starting position of the arm rotated (for a similar finding of shoulder-centered directional tuning, see Wu and Hatsopoulos, 2006). We asked whether a similar result could be confirmed in the present dataset and whether this model of a rotating preferred direction could capture more of the total neuronal variance.

We analyzed each neuron as follows (for equations, see Materials and Methods). First, for each movement, we defined the start vector to be the straight-line vector from the monkey's shoulder to the position of the hand at the start of the movement. The direction of this start vector was defined by a specific start azimuth angle and start elevation angle. We then fit the neuronal data to a regression model that included a cosine tuning to a preferred direction but in which the preferred direction rotated in the azimuth in a manner linearly related to the start azimuth angle of the arm and also rotated in elevation in a manner linearly related to the start elevation angle of the arm. The ratio between the rotation of the preferred vector of the neuron and the start vector of the arm was parameterized by two numbers, AZ for the azimuth ratio and EL for the elevation ratio. For example, if the regression settled on an AZ = 0, this result would indicate that the preferred direction for the neuron remained constant in different regions of space and did not rotate in the azimuth regardless of the start vector. If the regression returned an AZ > 0, this result would indicate that the preferred direction of the neuron tended to rotate in the same direction as the starting vector of the arm. An AZ = 1 would indicate that the preferred vector rotated exactly in tandem with the start vector of the arm. An AZ < 0 would
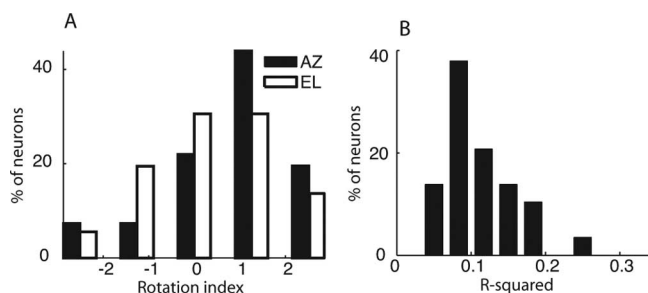

Figure 5. Rotation of preferred direction with starting hand position. *A*, A preferred direction model was tested in which the preferred direction of a cell was not fixed but instead could rotate as the start position of the hand rotated. Frequency histogram shows AZ and EL values for all cells tested. The AZ rotation index indicates the ratio between the starting azimuth angle of the arm and the azimuth angle of the preferred direction of the neuron. The peak in AZ near 1 indicates that the preferred direction of most neurons tended to rotate in the same direction and by a similar amount as the starting angle of the arm. Similarly, the EL rotation index indicates the ratio between the starting elevation angle of the arm and the elevation angle of the preferred direction of the neuron. *B*, Frequency histogram of R^2 values for all cells tested with this model of a rotating preferred direction.

indicate that the preferred vector rotated in the opposite direction as the start vector of the arm. Given the noisy nature of neuronal signals, a range of AZ values was expected, and a null result would emerge as a distribution of AZ values centered around 0. Similar relationships would pertain to the elevation of the preferred vector and the elevation of the start vector, characterized by the parameter EL.

Figure 5*A* shows the results of this regression for all neurons. The black bars show the results for the azimuth. The AZ values peaked near 1. The mean AZ value was 0.88, and the distribution was significantly more than 0 ($t = 3.87; p = 0.00039$). This result replicates the finding of Caminiti et al. (1990) in three ways. First, the preferred direction rotated when the start position of the hand was changed. Second, the extent and direction of rotation of the preferred direction varied considerably from neuron to neuron, as shown by the wide distribution of AZ values. Third, on average across neurons, the preferred direction rotated in the same direction and by approximately the same extent (approximately a ratio of 1) as the start vector. On this final point, our results are even more clear than the results of Caminiti et al., possibly because Caminiti et al. used only three start positions of the hand whereas we used a continuous range of start positions over the entire workspace, improving our sampling.

The experiment of Caminiti et al. examined changes in start position in the azimuth only. Here we were able to examine changes in both azimuth and elevation, and the open bars in Figure 5*A* show the result for elevation. The mean EL value was 0.54, and the distribution was significantly more than 0 ($t = 2.28; p = 0.029$). This result indicates that the preferred direction tended to rotate in elevation in the same direction that the start vector rotated. However, the distribution was not as clearly shifted toward 1 for elevation as it was for azimuth.

These results show that the preferred direction that best fits a neuron, as extracted by regression analysis, tends to be one that changes depending on the start position of the hand, rotating approximately in tandem with the arm. This trend is more clear and highly significant for rotations in the azimuth and is less clear although still significant for elevation. Figure 5*B* shows the distribution of R^2 values obtained with this model of a rotating preferred direction. The mean R^2 value is 0.13. This mean R^2 value is significantly greater than the mean R^2 of 0.08 obtained with the previous, simpler model of a fixed preferred direction
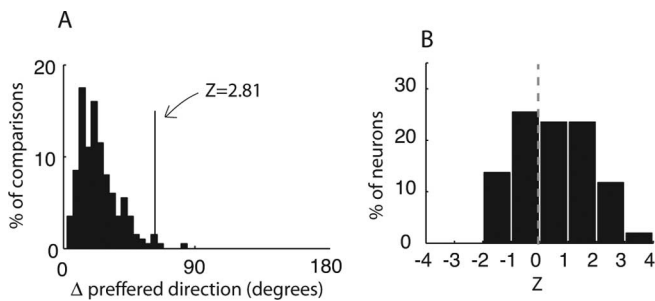
Figure 6. Change in preferred direction with starting arm posture. **A**, Results for one neuron. All movements were divided on a median split according to a raised elbow parameter (see Results, Direction tuning IV: changes of preferred direction with changes in initial posture). The preferred direction was calculated separately for the two sets of movements, and the difference in preferred direction ($\Delta\theta$) is shown as the vertical black line. To assess the reliability of this result, the movements were also randomly divided into two groups and a $\Delta\theta$ was calculated. This random division was performed 200 times, and the results are shown as a frequency histogram. On this randomized distribution, the $\Delta\theta$ of the nonrandom, posture-based split was significantly above the mean ($Z = 2.81$; $p = 0.0051$). **B**, Z scores for all cells tested as in **A**, This distribution of Z scores was significantly >0 (mean of 0.51; $t = 2.89$; $p = 0.0057$).

and shown in Figure 4F (comparison of the two R^2 distributions, $t = 3.5$; $p = 0.00072$).

Direction tuning IV: changes of preferred direction with changes in initial posture

Scott and Kalaska (1995, 1997) showed that the preferred direction of a neuron tends to change when the initial posture of the arm is changed. In their experiment, the start posture of the arm was varied by altering the height of the elbow, essentially increasing the elevation angle of the shoulder and the internal rotation angle of the shoulder in tandem, to keep the hand in approximately the same spatial location. Here we examine whether a similar result applies to the present naturalistic movement set.

For each neuron, we performed the following analysis. We analyzed the posture of the arm at the start of each movement. We focused on two aspects of arm posture: the elevation angle of the shoulder and the internal rotation angle of the shoulder. We constructed a single parameter, the "raised elbow" parameter, in which both of these shoulder angles varied together. A high value of raised elbow indicates that the shoulder elevation was high and the shoulder was internally rotated, essentially bringing the arm to a "chicken wing" posture with the elbow in upper space. Likewise a low value of raised elbow indicates that the shoulder elevation was low and the shoulder was externally rotated, essentially bringing the elbow to a low position. All movements were divided into two equal groups along a median split: those with a high raised elbow value at the start of the movement, and those with a low raised elbow value at the start of the movement. For each group of movements, we separately calculated a preferred direction using a regression analysis as above (see Direction tuning I). We then compared the two preferred directions by calculating the angular difference between them.

Figure 6A shows the result for one example neuron. The preferred direction was different for the trials with a high value of raised elbow than for the trials with a low value of raised elbow, with a change in preferred direction of 62° (Fig. 6A, thin vertical line). Is this change in preferred direction the result of the change in arm posture, or is it merely the result of noise in measuring an unreliable preferred direction? To address this question, we took the same total movement set and divided it in half randomly, calculating a preferred direction for each half and then finding

the difference between the two preferred directions. We performed this random division of the data 200 times. The results are shown in the histogram in Figure 6A. When the movements were randomly assigned to two groups, the difference in preferred direction tended to be small, with a mean difference of 24°. When the movements were assigned to two groups on the basis of arm posture, the difference obtained was far from the randomized mean, with a Z score of 2.81. For this particular cell, the start posture of the arm had a significant effect on the preferred direction, beyond that expected by chance ($p = 0.0051$). This example neuron was not typical; it showed an unusually clear effect.

Figure 6B shows the distribution of Z scores obtained for all neurons. If the posture of the arm had no significant effect on direction tuning, then the Z scores should be equally distributed about 0. The distribution is broad and extensively overlaps 0 but is significantly more than 0 (mean of 0.51; $t = 2.89$; $p = 0.0057$). Thus, the effect of starting posture on the preferred direction is variable from neuron to neuron but across the population of neurons is statistically significant. These results replicate the finding of Scott and Kalaska (1995, 1997), indicating that the same effect obtained in a constrained movement set can also be obtained in a free-movement paradigm. The results of this section and the previous section (see Direction tuning III: rotation of preferred direction with changes in starting hand position) together confirm, in the free-moving paradigm, one of the basic properties of motor cortex neurons that had been described previously: direction tuning is not consistent across all parts of space or all configurations of the arm. Instead, it is most consistent locally and changes with global changes of hand position or arm posture (Caminiti et al., 1990; Scott and Kalaska 1995, 1997; Sergio and Kalaska, 2003).

Direction tuning V: velocity tuning

The above sections explore speed tuning and direction tuning separately. It has also been shown that motor cortex neurons are tuned to velocity or to the combination of speed and direction (Moran and Schwartz, 1999; Reina et al., 2001; Paninski et al., 2004). Because speed accounted for almost none of the overall variance, we expected a velocity model to account for approximately the same amount of variance as direction tuning alone. We tested two velocity models.

First, we performed the same regression as above (see Direction tuning II: global), but the direction vectors in the regression equation were replaced by velocity vectors (the velocity vector for each movement was defined as the mean hand velocity during the movement, whose components were the mean x, y, and z hand velocity). The regression model returned a preferred velocity vector for each neuron with a cosine tuning to that preferred velocity. This model of velocity tuning resulted in a mean R^2 value of 0.09, not significantly different from the value obtained for regression against direction alone (0.08).

Second, we performed the same regression as above (see Direction tuning III: rotation of preferred direction with changes in starting hand position) but again with the direction vectors in the regression equation replaced by hand velocity vectors. This regression model returned a preferred velocity vector that could rotate depending on the start position of the hand. This model of velocity tuning resulted in a mean R^2 value of 0.11, also not significantly different from the value obtained for regression against direction alone (0.13).

Direction tuning VI: modulation by a hand-position signal

The activity of motor cortex neurons is modulated not only by the direction of a reach but also by the position of the hand (Kettner et al., 1988; Caminiti et al., 1990; Scott and Kalaska, 1995, 1997; Sergio and Kalaska, 2003). This hand-position signal is often included in the direction-tuning regression model as a linear modulation of the baseline firing rate of the neuron or as a linear modulation of the gain of the direction-tuning response.

We examined a model of this type that took into account both direction and hand position. We began with the regression equation from above (see Direction tuning III), in which the preferred direction was allowed to rotate depending on the starting position of the hand. This model, as described above, resulted in a mean R^2 value of 0.13. It was the best of the direction tuning models in the sense that it captured the greatest percentage of the variance. We then modified the regression in two ways. The A term in the model (see Materials and Methods) was replaced by a term that depended linearly on the x, y, and z starting position of the hand. In this manner, the gain of the direction tuning response could be modulated by hand position. Likewise, the B term in the model was replaced by a term that depended linearly on the x, y, and z starting position of the hand, thereby allowing the baseline of the direction tuning response to be also modulated by hand position. This model therefore examined the amount of neuronal variance that could be attributed to a combination of direction tuning and a linear effect of starting hand position. Because the model took into account more than direction tuning, it was expected to capture more of the neuronal variance than did direction tuning alone.

The mean R^2 value for this combined regression was 0.16. One interpretation is that direction tuning has improved by taking linear hand-position trends into account. This interpretation, however, is not correct. The model is not a direction-tuning model; it is a mixed model. The 16% of the variance is not attributable to direction tuning alone. Instead, the results indicate that direction tuning captures ∼13% of the variance and that adding a linear dependence of firing rate on starting hand position can add an additional ∼3% of the variance.

End-point tuning I: Gaussian tuning

In a standard test of direction tuning, the hand begins in a central location and reaches to a variety of peripheral targets. The movements therefore differ in both direction and final hand position. Many studies have examined whether the tuning obtained in these tasks is truly a tuning to direction or instead a tuning to the final position of the hand (Georgopoulos et al., 1985; Kettner et al., 1988; Caminiti et al., 1990; Fu et al., 1993). The studies test a variety of starting and ending positions, such that both hand position and direction can be independently varied. These studies find that direction tuning is genuine and persists independently of hand position, although there is also some effect of final hand position. These experiments however were inadvertently designed to favor the effect of direction tuning and to diminish the effect of end-point tuning, in the following manner. Most of these studies sampled the full range of possible directions, thereby maximizing the chance of detecting direction tuning, even if that tuning is broad. The studies, however, typically explored only a small range of positions near central space. For example, a central cube of space that is 10 cm in width comprises only ∼5% of the volume of the workspace used by the monkey in the present experiment. In one of the most extensive previous tests of the workspace (Caminiti et al., 1990), three cubes of space, each 10 cm across, were tested, totaling ∼15% of the total workspace. If a
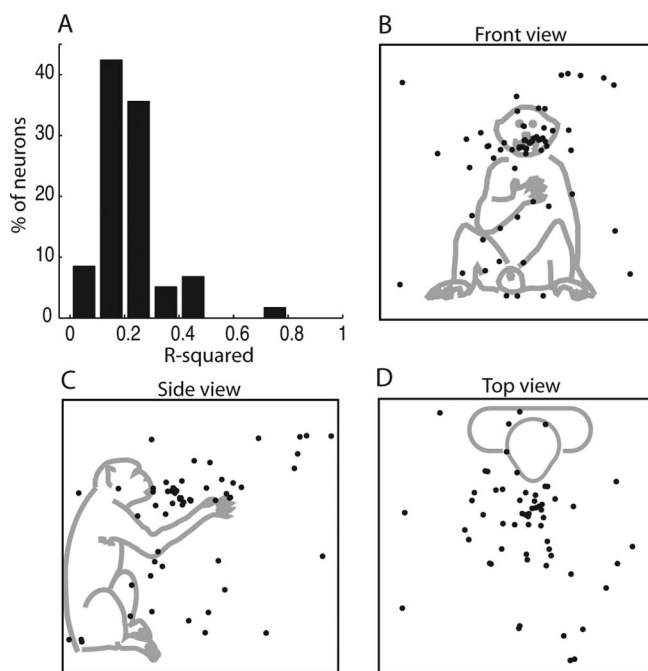

Figure 7. Tuning of neurons to hand end point. ***A***, Each neuron was tested with a tuning model in which the neuron fired most during movements that terminated with the hand at or near a specific location in space and fired progressively less during movements for which the hand terminated progressively farther from the preferred location. The graph shows a frequency histogram of R^2 values for all cells tested with this end-point model. ***B–D***, Preferred hand positions as determined by the end-point tuning model, displayed from three perspectives. The schematic monkey drawing indicates approximate scale and orientation.

neuron is broadly tuned to hand end point, such that the tuning curve can be seen only over the full workspace and varies little over the limited range tested, then end-point tuning would be minimized in these previous experiments. Here we examine whether neurons are significantly end-point tuned over the global movement set. It is important to note that we do not treat direction tuning and end-point tuning as rival hypotheses. Neurons may in principle be significantly tuned to direction in a manner that emerges over local regions of space and in an independent manner significantly tuned to end point in a manner that emerges over a more global expanse of the workspace.

To test for end-point tuning, we modeled the firing rate of the neuron as a Gaussian function of end point in which the Gaussian was peaked at a preferred spatial location of the hand (see Materials and Methods). Movements that ended near that preferred end point should be associated with high firing rates, and movements that ended far from that preferred end point should be associated with low firing rates. For each neuron, we obtained an R^2 value indicating how well this end-point model fit the neuronal data. Figure 7*A* shows the distribution of R^2 values for the population of neurons. On average, the end-point model accounted for 22% of the total neuronal variance.

Figure 7*B–D* shows the distribution of preferred hand positions for the tested neurons, as obtained with the end-point tuning model. These preferred hand positions were dispersed through the space around the animal. For some neurons, the regression analysis found a preferred hand position at the edge of the normal range of the workspace. In such a case, the neuron essentially had a monotonic tuning to hand position, preferring one extreme side of the workspace over the other. Most neurons, however, had a preferred hand position inside the workspace and
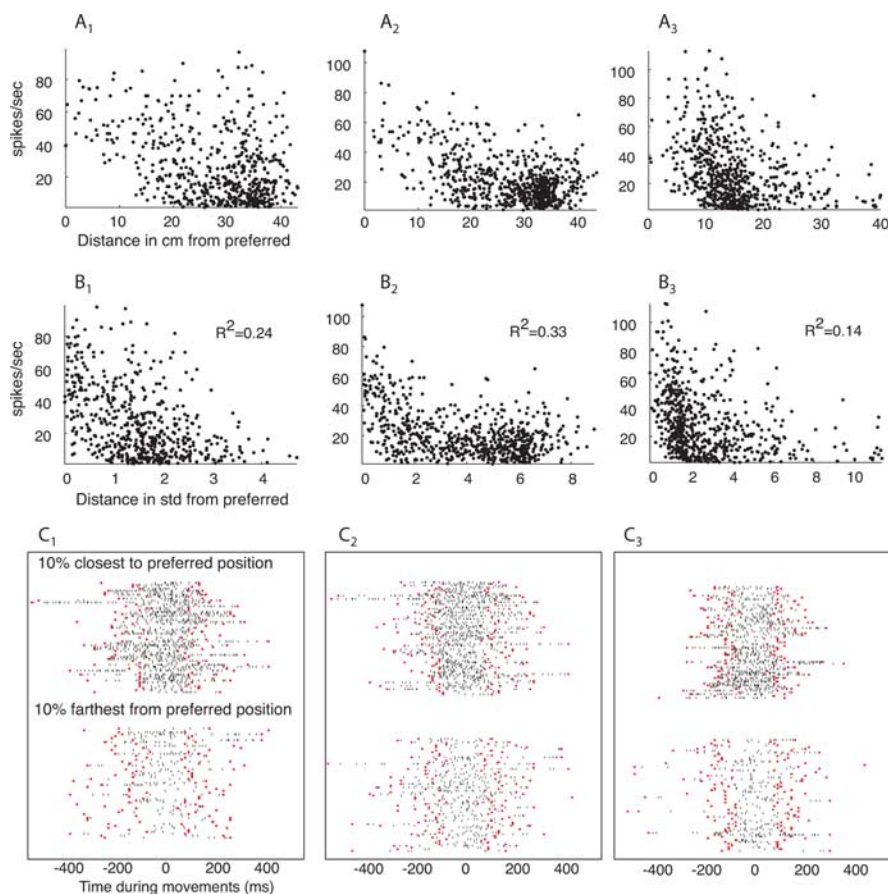
Figure 8. End-point tuning of three example neurons. **A1–A3**, Mean firing rate of example neurons 1–3 during each movement as a function of the distance (centimeters) between the end of the movement and the preferred end point determined by regression analysis. The neurons fired more during movements that terminated closer to the preferred end point. **B1–B3**, Same data as in **A** but with the x-axis replotted in units of the SDs of the Gaussian tuning curve, displaying the tuning to end point more clearly. **C1–C3**, Rasters showing high neuronal activity of example neurons 1–3 during the 10% of movements that terminated nearest to the preferred end point and low neuronal activity during the 10% of movements that terminated farthest from the preferred end point. Red tic marks indicate start and end of movement.

movements that terminated progressively farther from the preferred end point had progressively lower firing rates. This trend is noisy for two reasons. First, only ~24% of the variance of this neuron was attributable to the Gaussian tuning to the preferred end point. Second, the SD of the Gaussian tuning was not equal in all spatial directions. This cell was more sharply tuned in the x dimension that in the y and z dimensions. In terms of the Gaussian tuning, 1 cm in the x direction is not equivalent to 1 cm in the y or z direction. To better show the Gaussian fit to the data, Figure 8B1 shows a replotting of the same graph. In this plot, the x-axis represents the distance between the end point of each movement and the preferred end point of the neuron, but this distance now is expressed in SDs of the Gaussian tuning function. Again, on average, movements that terminated near the preferred end point (within an SD of the peak of the Gaussian) had higher firing rates, and movements that terminated progressively farther from the preferred end point had progressively lower firing rates. This graph shows in a more direct way the relationship between firing rate and the proximity of the movement end point to the preferred end point.

Figure 8C1 shows rasters of neuronal activity for the same example neuron. Each line in the raster display shows spike data during the analysis window for a single movement. Because different movements were of different durations in the naturalistic movement set, these lines in the raster display are of different lengths. The first raster display shows data from the 10% of movements that terminated nearest to the preferred position; the second raster display shows data from the 10% of movements that terminated farthest from the preferred position. These rasters show that the firing rate was variable from movement to movement but that the neuron tended to fire more during movements that terminated near the preferred hand position (mean ± SD of 49.8 ± 20.6 spikes per second) and tended to fire less during movements that terminated far from the preferred hand position (mean ± SD of 13.9 ± 12.4 spikes per second) with a significant different between these two firing rates ($t = 11.3$; $p = 4 \times 10^{-20}$).

Is the tuning to final hand position merely an artifact of direction tuning? The region of final hand positions that the neuron appears to prefer may have been approached predominantly from the preferred direction. For example, a neuron that prefers upward movements might appear to prefer an upper location in space. This explanation, however, is unlikely for several reasons. First, it predicts that the preferred hand locations should cluster systematically around the edges of the workspace, but, as shown in Figure 7, the neurons preferred a range of final positions including many within the workspace. Second, as described below (see Simulated neurons), we simulated a direction-tuned neuron and tested it on the actual movements measured from the monkey. The simulated neuron performed well on a direction-tuning
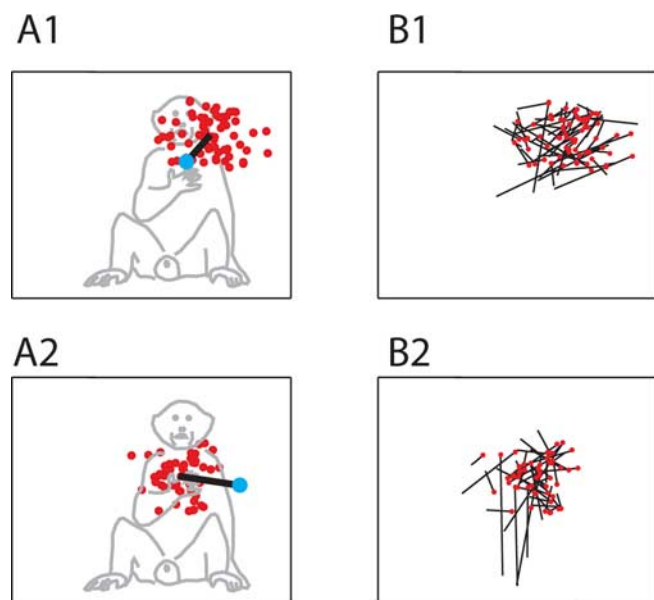
therefore had a nonlinear response function that fell off to either side of the preferred location. The tuning functions were generally broad: the mean width of the Gaussian tuning curve at half-height was 18 cm, approximately half the range of the workspace. Because of the broadness of these tuning curves, the end-point tuning is unlikely to be a major source of signal in experiments that test a small part of the workspace in front of the monkey. If only a small volume of the workspace is tested, most neurons will have a preferred hand position outside that tested region, and the tuning to end point will either be missed (if the tested part of space happens to overlap a relatively flat part of the tuning curve) or will be measured as a graded, monotonic preference for one side of the tested space over the other (if the tested part of space overlaps with the rising slope of the tuning curve).

Figure 8A1 shows data from an example neuron. First, the end-point tuning was found using the regression analysis described above. This analysis obtained a Gaussian surface in three-dimensional space for which the peak of the Gaussian corresponded to the preferred end point of the neuron. In this graph, the x-axis represents the distance in centimeters between the end point of each movement and the preferred end point of the neuron. The y-axis represents the firing rate of the neuron during the movement. On average, movements that terminated near the preferred end point (distance of 0) had higher firing rates, and

Figure 9. Direction tuning does not explain end point tuning. *A1*, *A2*, Data from two example neurons. The dots show the final hand position for the 10% of movements that terminated closest to the calculated preferred position. The line shows the preferred direction of movement as calculated using the same local set of movements. The blue dot shows the end of the movement vector. *B1*, *B2*, Data from the same neurons shown in *A*. Each line depicts a movement (straight line connecting the start and end point of the movement), and each red dot shows the end point of the movement, for the same movements whose end positions are shown in *A*.

model and showed little tuning to final hand position. A neuronal signal that is known to be direction tuned, therefore, does not produce an artifactual hand-position tuning.

Finally, Figure 9 shows a third analysis to address this question. Figure 9*A1* contains data from one example neuron. The dots show the final hand positions for the 10% of movements that terminated closest to the preferred hand position. Figure 9*B1* shows the directions of the movements that terminated in those locations. This movement set included a range of directions and did not approach the end points from predominantly one direction. We used this limited movement set to obtain a preferred direction (using a standard cosine tuning model as above, see Direction tuning I). The direction-tuning R^2 value was 0.18. (The relatively low R^2 value, below the mean of 0.42 obtained for local direction tuning in Direction tuning I, probably reflects the fact that these movements, although more local than a full movement set, are nonetheless more dispersed than in our test of local direction tuning.) The preferred direction among these movements, plotted as a thick line in Figure 9*A1*, was aimed down and to the left, whereas the preferred hand location was up and to the right. For this neuron, the preferred direction was pointed the wrong way to account for the preferred hand location. Figure 9 also shows a second example cell. In general, the preferred direction had no obvious relationship to the preferred position. Thus, although neurons are to some degree direction tuned (as outlined in the previous sections) and are to some degree tuned to the end point of the movement, the preference for a particular final position of the hand apparently cannot be explained as a consequence of a preference for a particular direction of motion. This question is further addressed in the next section.

End-point tuning II: directional convergence
Previous experiments suggest that neurons rarely show a preference for movements directed toward a specific location in space

(Caminiti et al., 1990; Sergio and Kalaska, 2003). The preferred direction of a neuron does not change systematically across space in such a way that it aims toward a point of convergence. Instead, when the preferred direction is assessed in different regions of space, it tends to change somewhat idiosyncratically. On average across neurons, it tends to rotate in the same direction as the rotation of the shoulder. Similar results were confirmed in the present experiment (see Direction tuning III: rotation of preferred direction with changes in starting hand position). How can this result, a failure of the preferred direction to converge to a point, be reconciled with the result of the previous section, that neurons fire most during movements that end near a preferred location in space? On first sight, the two results appear to be contradictory.

There are, however, at least two ways to define end-point tuning. One is a proximity model, and the other is a directional model. In the proximity model, tested in the previous section, a neuron will fire best during movements that terminate near a preferred position, regardless of the direction of the movement. For example, a movement that starts at the preferred position and ends 1 cm away might result in a high firing rate, although the movement is directed away from the preferred position. Likewise, a movement that starts 20 cm from the preferred position and ends 15 cm from it might result in a low firing rate, although the movement is directed toward the preferred position. This model, in which trajectory information is discarded and only the final position enters the equation, accounted for 22% of the neuronal variance on average, as described above.

In a directional model of end-point tuning, in contrast, the neuron fires most during movements directed toward the preferred position and fires least during movements directed away from the preferred position, regardless of whether the movement terminates near or far from the preferred position. For example, a movement that starts at the preferred position and ends 1 cm away might result in a low firing rate, although the movement ends near the preferred position. Likewise, a movement that starts 20 cm from the preferred position and ends 15 cm from it might result in a high firing rate, although the movement ends far from the preferred position. This model, in which the final position of the hand at the end of each movement is discarded and only the direction enters the equation, has generally been found to account for little of the neuronal behavior (Caminiti et al., 1990; Sergio and Kalaska, 2003).

To determine whether our dataset was consistent with previous findings, we also tested a directional model of end-point tuning. In this model, for each movement, two vectors were defined. The first vector connected the start point of the movement to the end point of the movement. The second vector connected the start point of the movement to a preferred end point of the neuron. The angular deviation between these vectors was defined as $\Delta\theta$. In the regression equation, the firing rate was modeled as a cosine function of $\Delta\theta$ [firing rate = $A \cos (\Delta\theta) + B$]. For each neuron, the regression obtained the preferred end point that best fit this model. For a neuron that follows this model, the firing rate should be highest when the movement is directed toward the preferred end point and lowest when the movement is directed away from the preferred end point. This model resulted in a mean R^2 value across neurons of 0.07, less than the mean R^2 value of 0.22 obtained for the Gaussian tuning to end point described in the previous section.

It is instructive to compare the three general classes of direction model tested in the present experiment. First, a standard direction model, in which the preferred direction is the same

everywhere in space (see above, Direction tuning II: global), accounted for 8% of the variance on average. Second, a direction model in which the preferred direction could rotate in different regions of space (see above, Direction tuning III: rotation of preferred direction with changes in starting hand position) accounted for 13% of the variance. This rotating model tended to select a preferred direction that rotated approximately with the shoulder angle. This rotation of the preferred direction results in a diverging pattern, in which preferred directions in adjacent parts of space point slightly away from each other. It is not surprising therefore that the third model (of the present section), in which the preferred directions in different regions of space were constrained to converge toward a common location, performed somewhat less well than either of the previous ones, accounting for 7% of the variance.

These results together with the results of the previous sections suggest that neurons are to some extent tuned to end point (mean $R^2 = 0.22$). This end-point tuning, however, should not be construed as a set of preferred directions aimed at a common location in space. The neurons are not well tuned to a converging pattern of preferred directions. Instead, the neurons are tuned to any movements that terminate near a preferred location. The spatial tuning curve is broad, the preferred region is large, the natural hand movements are often relatively short (mean length of 8 cm), and therefore a variety of movements aimed in all directions can terminate within the preferred region yielding a high firing rate. In addition, superimposed on this spatial tuning for end point, the neurons are tuned to some extent to direction (mean R^2 of $\sim0.08 – 0.13$ depending on the model used), in which the preferred direction has little obvious relationship to the preferred hand location.

End-point tuning III: start-point or end-point tuning?
As described above, the firing rate of the neurons tended to be larger for movements that terminated near a preferred hand location, and this correlation accounted for $\sim22\%$ of the neuronal variance. However, because the monkey's spontaneous movements were relatively short (mean of 8 cm), the movements tended to start and end in a similar region of space. Start point and end point, therefore, were correlated across the global movement set. Were the neurons tuned specifically to the end point of the movement, or was neuronal activity equally determined by the proximity of the start point to the preferred location?

To approach this issue, we performed the following analysis. For each neuron, we selected the 10% of movements that terminated nearest to the preferred position. Figure 9B1 shows this set of movements for one example neuron. These movements contain some variance in ending and starting position. Because this local region of space around the preferred location had a diameter that approximated the average length of a movement, therefore, within this movement set, no significant correlation existed between start point and end point. We then asked, within this set of movements near the preferred location, whether the distance between the start point of a movement and the preferred point was significantly related to firing rate (using a linear regression between firing rate and distance from start point to preferred point). This was not typically the case. Across neurons, the mean R^2 value for this linear regression was 0.02, and 7% of the neurons showed a significant relationship. In contrast, we asked whether the distance between the end point of a movement and the preferred point was significantly correlated with firing rate. A relationship was obtained in which the firing tended to be higher for movements that terminated closer to the preferred point. The
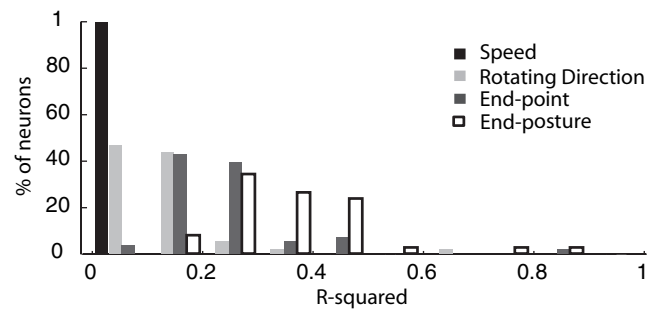

Figure 10. Comparison of four regression models. Each neuron was tested with a tuning model in which the neuron fired most during movements that terminated with the eight joints of the arm at or near a specific postural configuration. This end-posture tuning is shown as a frequency histogram of R^2 values. Also shown is a frequency histogram of speed tuning (from Fig. 3D), direction tuning in which the preferred direction was allowed to rotate depending on starting hand position (from Fig. 5B), and tuning to the three-dimensional end point of the movement (from Fig. 7A).

mean R^2 value for this linear regression was 0.20, and 65% of the neurons showed a significant relationship. These results suggest that the proximity of the end points to a preferred location significantly affected the firing rate of neurons and that the proximity of the start points to the preferred location was at least a much weaker influence. This question of whether the neurons preferentially encode the end state of the movement or whether the firing rate equally reflects the initial state is further addressed below in two sections (see Neurons are more end-posture than start-posture tuned I: analysis of temporal offsets; and Neurons are more end-posture than start-posture tuned II: analysis of movements near the preferred posture).

End-posture tuning
Electrical stimulation of motor cortex on a behaviorally relevant timescale can evoke movements that terminate in specific postures (Graziano et al., 2002a, 2004a,b, 2005; Cooke and Graziano, 2004a,b). For example, stimulation of one location in cortex might bring the hand to the mouth with the elbow in lower space, the forearm supinated such that the palm faces the mouth, the wrist slightly flexed, the fingers in a grip posture, and the mouth open. The effect of stimulation is therefore not merely to bring the hand to a specific end point but to bring the many joints of the arm to a specific final configuration or posture. Observational studies of the normal spontaneous behavior of monkeys indicate that a high proportion of time is spent using the arm as a postural device (Graziano et al., 2002b, 2004a). Psychophysical studies in humans suggest some degree of encoding of the final posture of a movement (Feldman, 1986; Rosenbaum et al., 1995; Desmurget and Prablanc, 1997). We therefore asked to what degree neurons in motor cortex are tuned to the posture reached by the arm at the end of each movement.

Eight degrees of freedom of the arm were monitored, including grip aperture and seven joint angles. These degrees of freedom define an eight-dimensional posture space. We modeled neuronal firing rate as a Gaussian surface in eight-dimensional space whose peak corresponds to the preferred end posture. Movements that terminate at a posture near the peak of the Gaussian should be associated with a high neuronal firing rate, and movements that terminate at a posture far from the peak of the Gaussian should be associated with a low neuronal firing rate. For each neuron, we obtained an R^2 value indicating how well this model fit the neuronal data. The distribution of R^2 values across the population of neurons is shown in Figure 10. The mean R^2 value

was 0.36. For comparison, Figure 10 also shows the distribution of R^2 values for speed tuning, for the complex direction tuning in which the preferred direction could rotate with the starting position of the hand, and for end-point tuning in three-dimensional Cartesian space. Of the models of neuronal tuning tested here, the one that performed best involved so called intrinsic coordinates (in this case joint angles) rather than extrinsic coordinates (direction or Cartesian hand location), in agreement with previous suggestions that motor cortex emphasizes intrinsic coordinates (Scott and Kalaska, 1997). These tuning curves to end posture are examined in greater detail in the following six sections.

Examples of end-posture tuning

Figure 11A1 shows data from an example neuron. First, the end-posture tuning was found using the regression analysis described above. This analysis obtained a Gaussian surface in eight-dimensional posture space for which the peak of the Gaussian corresponded to the preferred posture of the neuron. In this graph, the x-axis represents the distance in posture space between the final posture of each movement and the preferred posture of the neuron. It is important to note that this distance is not the distance of the hand from a preferred location in space, but rather the distance of the arm from a preferred configuration in posture space. Thus, the units cannot be in Cartesian centimeters. Instead, we expressed the units in SDs of the Gaussian fitting function. In this manner, all eight dimensions of posture space can be expressed in the same units, and a one-dimensional graph can be presented. The advantage of the one-dimensional graph is that it shows the Gaussian fit to the data. On average, movements that terminated near the preferred posture (within an SD of the peak of the Gaussian) had higher firing rates, and movements that terminated progressively farther from the preferred posture had progressively lower firing rates. Approximately 33% of the variance of this neuron was attributable to the Gaussian tuning to the preferred posture ($R^2 = 0.33$; $p < 0.0001$).

Figure 11B1 shows rasters of neuronal activity for the same example neuron. Each line in the raster display shows spike data during the analysis window for a single movement. Because different movements were of different durations in the naturalistic movement set, these lines in the raster display are of different lengths. The first raster display shows data from the 10% of movements that terminated nearest to the preferred posture; the second raster display shows data from the 10% of movements that terminated farthest from the preferred posture. These rasters show that the firing rate was variable from movement to movement, but that the neuron clearly fired more during movements
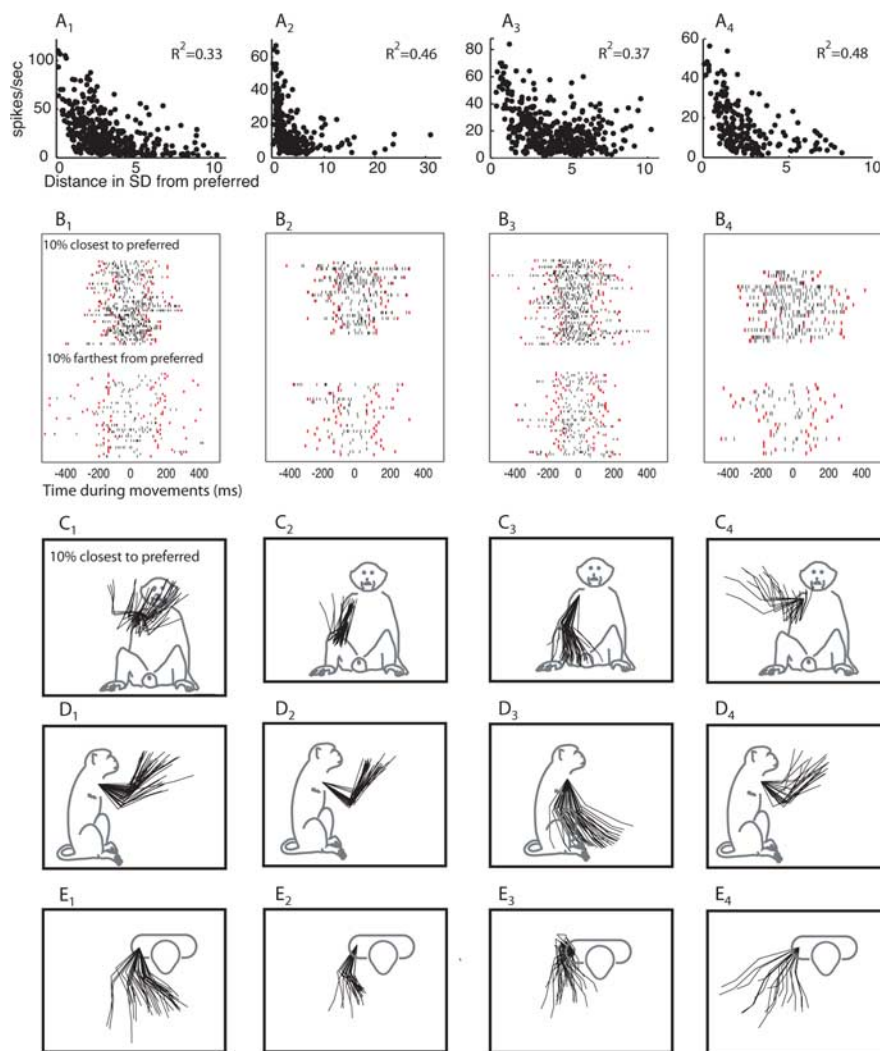

Figure 11. End-posture tuning of four example neurons. **A1–A4**, Data from example neurons 1–4. For each neuron, the preferred multijoint posture was determined by regression analysis. The final posture of each movement was compared with the preferred posture. The distance between them was calculated in eight-dimensional posture space. Distance was measured in units of SDs of the Gaussian tuning function, to express all eight dimensions in posture space in equivalent units. This distance is plotted on the x-axis, and firing rate during the movement is plotted on the y-axis. **B1–B4**, Rasters showing high neuronal activity during the 10% of movements that terminated nearest to the preferred end posture and low neuronal activity during the 10% of movements that terminated farthest from the preferred end posture. Red tic marks indicate start and end of movement. **C1–C4**, Stick-figure drawings showing front view of the arm for the 10% of movements that terminated nearest to the preferred end posture. Three joints are shown: shoulder, elbow, and wrist. The schematic monkey drawing indicates approximate scale and orientation. **D1–D4**, Same as **C** but side view. **E1–E4**, Same as **C** but top view.

that terminated near the preferred posture. Although only 33% of the neuronal variance was attributable to end-posture tuning, the effect is clearly visible in the rasters.

Figure 11C1–E1 shows more data from the same example neuron. Each stick figure shows the configuration of the arm at the termination of a movement. The figure shows the 10% of movements that terminated closest to the preferred posture. This display does not capture the full eight-dimensional posture space. Several degrees of freedom, especially the internal/external rotation of the shoulder and the pronation/supination of the forearm, are not easily visible. However, the graph does give some sense of the range of postures preferred by this neuron. The neuron is more narrowly tuned to some joints (such as the shoulder elevation) than to others (such as the shoulder azimuth). As a result, this range of preferred postures does not correspond neatly to a small cluster of hand positions in space. Instead, a large and ir-
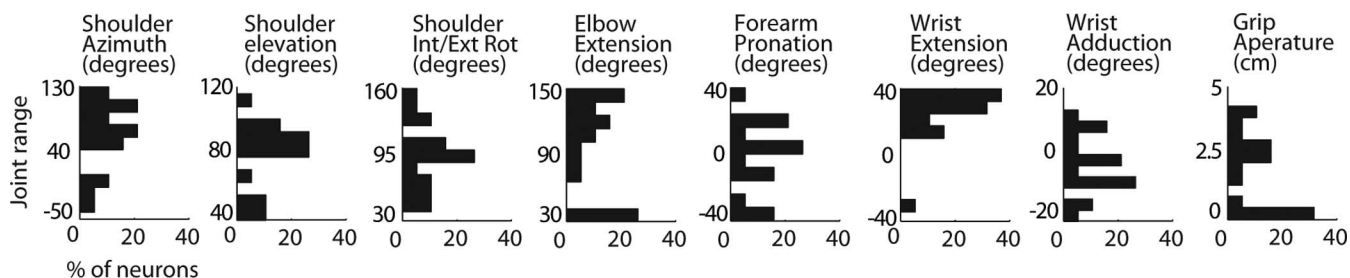
Figure 12. Distribution of joint angles preferred by neurons. For each joint, a frequency histogram shows the proportion of neurons tuned to a particular final joint angle. Only the 50% most sharply tuned neurons were plotted for each graph. Most neurons were tuned to more than one joint and are therefore represented on more than one histogram.

regular region of hand space is represented by these postures. This finding that postural tuning is more sharply peaked for some degrees of freedom than for others is discussed in greater detail in below (see Posture subspaces). For this neuron, the preferred posture involved a raised elevation of the shoulder lifting the hand mainly into upper space, a range of azimuth angles that placed the hand in a band of frontal space, a partially extended elbow that placed the hand away from the body, an extended wrist angle, and a grip aperture (data not shown) that was on average 2 cm. The posture is similar to a naturalistic reach to grasp or manipulate a small object in upper frontal space.

Figure 11 also shows the end-posture tuning for three other example neurons.

Range of preferred joint angles
Some experiments have examined linear models of tuning to joint rotation (Reina et al., 2001). In a linear model, a neuron will prefer rotation of a joint to one or another extreme. The present model of end-posture tuning is nonlinear; it assumes that neurons fire best during rotation from any direction toward a preferred joint angle. However, the preferred angle could in principle lie at or beyond the edge of the range of movement. In this case, the nonlinear model used here would approximate the linear models used elsewhere. An important question, therefore, is whether the preferred joint angles as determined by the regression model tend to fall at the extreme of the range of motion or within the range of motion. Figure 12 shows the result. To examine those neurons that had clearly preferred postures, we ranked neurons by the sharpness of the Gaussian tuning function and arbitrarily chose the 50% of neurons that were most sharply peaked, thus the neurons for which the end posture was most clearly specified. These neurons are represented in Figure 12. Each frequency histogram shows data for one joint. For the most proximal joints, especially the three shoulder joints, neurons were tuned to a range of different preferred angles including those in the middle of the range. For the more distal joints, neurons were more likely to be tuned to an extreme. For example, neurons tended to prefer a closed-grip aperture. Note that the same neuron might be tuned to an intermediate shoulder angle and a closed grip. It was not the case that some neurons were tuned to proximal joints and others to distal joints; rather, neurons tended to be tuned to a combination of many joints, an issue discussed in greater detail below (see Posture subspaces).

The plot in Figure 12 provides some indication of the postures most commonly preferred by neurons. Many neurons preferred a shoulder azimuth at ~90°, orienting the arm into the front portion of space rather than to the side (0°), a shoulder elevation of ~80° and an elbow extension of ~130°, which places the arm in a raised and extended posture, the wrist extended, and the grip

closed. These most common joint angles correspond to a typical arm posture of the monkey, as if it had reached out to grip an object in central space. A range of other postures was also represented, as is evident from the range of joint angles in Figure 12. Because electrical stimulation in different regions of the precentral gyrus tended to evoke different common postures of the arm (Graziano et al., 2002a), it is possible that neuronal recordings from other cortical regions in the precentral gyrus might have revealed a different distribution of preferred postures.

To compare directly with a linear model of neuronal tuning in joint space, we also tested neurons using a simple eight-dimensional linear regression comparing mean firing rate during a movement with mean joint angular speed. Such a regression model assumes that a neuron prefers movement of a joint in one direction only, toward one extreme angle or the other. The mean R^2 for this linear model was 0.14, significantly less than the mean R^2 of 0.36 for the Gaussian end-posture model ($t = 7.0$; $p = 3.2 \times 10^{-9}$). This result confirms that neurons are tuned to movement through joint space, but that the linear model misses a common feature of the tuning, a tuning to joint rotations that terminate at an intermediate posture.

Posture subspaces
As described above, for each neuron a regression model was used to obtain the best-fitting Gaussian tuning function in eight-dimensional posture space. On examining the resultant Gaussian tuning functions, we found that neurons were not equally well tuned to all eight dimensions of posture space. We examined the sharpness of tuning to each of the eight degrees of freedom. Figure 13A shows the result for one example neuron. For each joint, we calculated the width of the tuning curve at half height, in degrees of joint angle. We then divided this width by the range of movement measured during the animal's spontaneous behavior, resulting in a tuning width index that could be compared across joints. A tuning width index of 1 indicates that the width of the tuning curve was approximately the same as the full range of movement. A tuning width index <1 indicates a narrower tuning curve, and a tuning width index >1 indicates a tuning curve so broad that it essentially was not relevant over the normal range of movement of the joint. If the index was 1 or less, we classified the neuron as well tuned to that joint. On this criterion, this example neuron was well tuned to four joints of the eight (four joints show a tuning width index <1). This neuron was tuned to shoulder azimuth, internal/external shoulder rotation, wrist extension, and grip aperture. It was therefore tuned to a mixture of proximal and distal joints. It was essentially untuned to the other joints. Figure 13B–D shows several more example neurons tuned to different combinations of joints. Figure 13E shows the mean across neurons. The joints are arranged in this graph in order
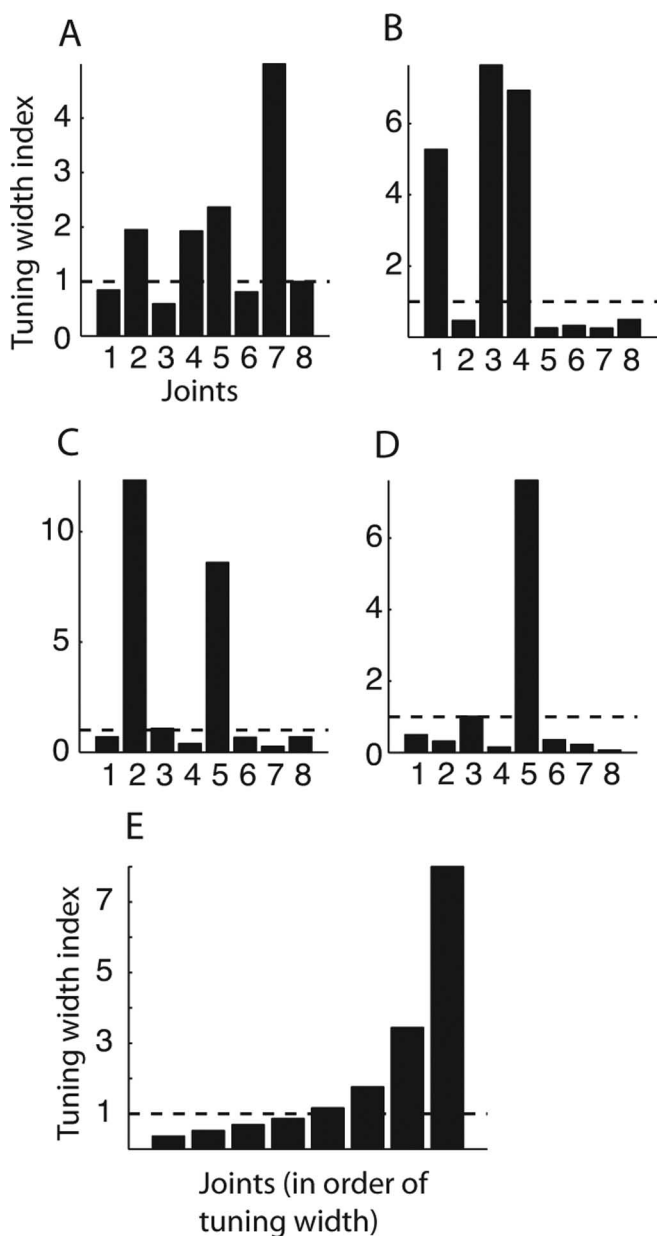
Figure 13. Tuning of neurons to subsets of joints. **A**, Data from one example neuron. The end-posture regression analysis returned a set of preferred joint angles to which the cell was tuned in a Gaussian manner. For each joint, a tuning width index was calculated. Tuning width index is the width of Gaussian tuning curve at half height in degrees of joint angle/natural range of joint motion measured during movement. Tuning index of 1 indicates that the width of the tuning curve at half height approximated the normal range of joint motion. Lower tuning index corresponds to sharper tuning. The graph shows that the neuron was sharply tuned to four joints and poorly tuned to the remaining four joints. The joints are (in order) shoulder azimuth, shoulder elevation, shoulder external/internal rotation, elbow extension, forearm pronation, wrist extension, wrist adduction, and grip aperture. **B**, An example neuron sharply tuned to five of the eight joints. **C**, An example neuron sharply tuned to six of the eight joints. **D**, An example neuron sharply tuned to seven of the eight joints. **E**, For each neuron, the joints were arranged from most sharply tuned to least sharply tuned; the results were then averaged across neurons. On average, neurons were well tuned to four of the eight joints.

from most sharply tuned to least sharply tuned. On average, neurons were well tuned to four joints (tuning width index <1) and relatively poorly tuned to the remaining joints (tuning width index >1). These results show that it is not quite correct to view the neurons as tuned to a preferred final posture in an eight-dimensional posture space. Rather, the neurons are tuned within

posture subspaces, each neuron primarily tuned to a subset of joints and less well tuned to the remaining joints.

One caveat of the above analysis is that the joints may be correlated with each other; hence, a neuron that is genuinely tuned to one joint may appear to be tuned to other, correlated joints. To address this caveat, we performed an alternative, conservative analysis that removed the covariance among joints. For each neuron, we first tested the end-posture model with a regression analysis as described above and arrived at an R^2 value indicating how much of the total variance could be explained by that model. We then performed an $n - 1$ regression to determine how much of that explained variance could be attributed solely to each joint, once the variance explained by the remaining seven joints had been removed. In this manner, the correlation among joints was explicitly taken into account and removed from the analysis. We then asked, for each joint, whether it made a statistically significant unique contribution to the R^2. This method is conservative, because it eliminates signal that overlaps among joints and considers only the component of signal that is unique to each joint. However, using this regression method, we obtained a similar result. Each neuron on average was significantly tuned to four joints of the eight.

Comparison of end-posture tuning and end-point tuning

The previous sections examined two types of end-state tuning: tuning to end point and tuning to end posture. Do the hand positions associated with the preferred postures tend to match the hand positions associated with the preferred end points? Such a match would provide a validity check, showing that two different analysis methods arrived at a similar result.

This question is complicated by two caveats. First, as described above, the neurons are not typically narrowly tuned to all eight angles that make up the posture of the arm. Each neuron tends to be narrowly tuned to an average of four joints of the eight. A set of postures near the preferred posture therefore corresponds to a narrow range of angles for some joints and a wide range of angles for other joints. This in turn can correspond to a large, irregular set of hand positions (for examples, see Fig. 11).

Second, a hand position associated with a preferred posture can also be reached by postures that are not preferred by the neuron. With eight degrees of freedom in the arm, any hand position can be achieved through a large range of postures.

For these two reasons, there is no clean mapping between hand position and arm posture. Therefore, a strict correspondence between the hand positions preferred by the neurons and the arm postures preferred by the neurons should not be expected. With these two caveats in mind, however, it may still be possible to find some degree of correspondence between the cloud of hand positions associated with the preferred posture and the cloud of hand positions associated with the preferred end point.

For each neuron, we obtained the 10% of movements that terminated closest to the preferred posture of the neuron, as determined by the eight-dimensional regression against end posture (Fig. 11C1). We obtained the mean x, y, and z location of the hand at the end of these 10% of movements. For the same neuron, we obtained the 10% of movements that terminated closest to the preferred hand position as determined by the three-dimensional regression against end point. We obtained the mean x, y, and z location of the hand at the end of these 10% of movements. For each neuron, we therefore obtained a mean hand position associated with the preferred posture (posture-defined position) and a mean hand position associated with the preferred

end point (end-point-defined position). We then asked whether the posture-defined position and the end-point-defined position were correlated across neurons. In the vertical dimension, the posture-defined position and the end-point-defined hand position were highly significantly correlated ($R = 0.62$; $p = 0.00003$). In the mediolateral dimension, the correlation was also highly significant ($R = 0.58$; $p = 0.0001$). Finally, in the front-to-back dimension, the correlation was also highly significant ($R = 0.43$; $p = 0.007$). These results indicate that the preferred hand positions obtained through the end-point model tended to approximately match the hand positions associated with the preferred postures obtained through the end-posture model. One possibility is that the tuning to end point, accounting for a mean of 22% of the neuronal variance, was partly a noisy reflection of a more robust tuning to arm posture, accounting for 36% of the neuronal variance.

Neurons are more end-posture than start-posture tuned I: analysis of temporal offsets

Numerous studies have reported that neuronal activity in motor cortex is modulated by posture (Kettner et al., 1988; Caminiti et al., 1990; Fu et al., 1993; Scott and Kalaska, 1995, 1997; Paninski et al., 2004). The present study supports the same finding, showing in addition that over the larger workspace the postural tuning curves can take the form of a broad, nonlinear tuning to a preferred posture. One question, however, is whether this postural tuning reflects the active control of movement, such as movement planning and execution, or whether it reflects feedback signals that inform the motor cortex about the state of the limb. One way to address this question is to ask whether the neurons are tuned specifically to the final posture of a movement, as if in preparation of that end state, or whether the neurons are equally tuned to the starting posture of the movement. We therefore performed the regression analysis described above (see End-posture tuning) but regressing against the start posture of each movement rather than the end posture to determine whether neurons show a similar Gaussian tuning, and a similar degree of fit, to a start-posture model.

The mean R^2 for the start-posture tuning model was 0.33, nearly as high as the R^2 for end-posture tuning (0.36). At first glance, this high value would seem to suggest that the neurons are modulated by current or recent posture rather than preferentially encoding the anticipated posture. An analysis of the temporal offsets, however, revealed a different reason for this high R^2 value. In this section, we first describe the temporal offsets for several regression models and then return to the question of start-posture coding.

Figure 14 shows the optimized temporal offsets obtained for five of the above regression analyses. In this graph, negative values represent offsets in which the window for neuronal activity was shifted earlier than the movement window. For example, Figure 14A shows the offsets obtained for direction tuning (following the method used in Direction tuning II: global). The mean ± SD temporal offset for direction tuning was −36 ± 117 ms and was not significantly different from 0, indicating that on average the regression against direction was optimized when using neuronal activity that was approximately coincident with the movement. Figure 14 also shows the result for the more complex direction tuning model in which the preferred direction could rotate depending on starting hand position (Fig. 14B), speed tuning following the method used in Speed tuning 1 (Fig. 14C), and tuning to end posture (Fig. 14D). In all of these cases, the
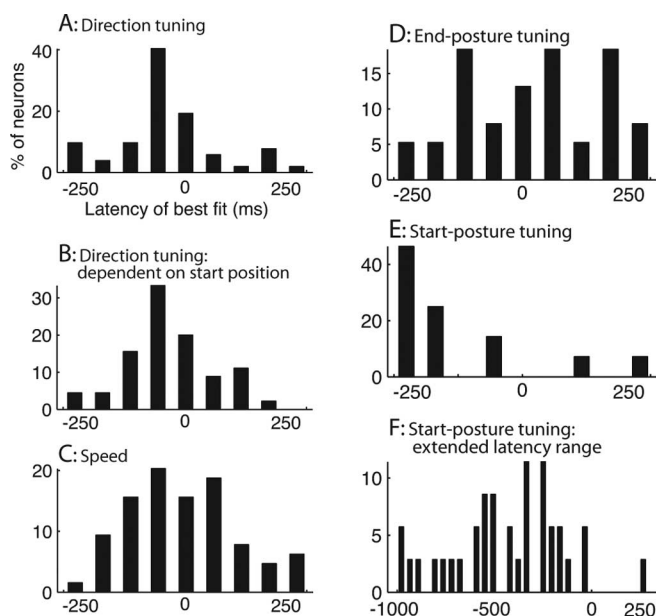

Figure 14. Temporal offsets for six regression models. For each regression model, the kinematic data were selected from a fixed time period during the movement, and the neuronal data were selected from a time period equal in duration but offset by a temporal lag. For each neuron, the temporal offset that optimized the regression was selected. These optimal temporal offsets are shown in the frequency histogram in which the x-axis represents temporal offset. Negative values indicate neuronal activity leading movement. **A**, Temporal offsets for direction tuning following the method in Results (Direction tuning II: global). **B**, Temporal offsets for direction tuning following the method in Results (Direction tuning III: rotation of preferred direction with changes in starting hand position). **C**, Temporal offsets for speed tuning following the method in Results (Speed tuning I). **D**, Temporal offsets for end-posture tuning. **E**, Temporal offsets for start-posture tuning. **F**, Temporal offsets for start-posture tuning in which the range of possible offsets was extended downward to −1000.

offsets were distributed around and not significantly different from 0.

Previous studies reported offsets in which neuronal activity significantly preceded movement (Moran and Schwartz, 1999; Reina et al., 2001; Paninski et al., 2004). In our data, the temporal offsets for direction tuning trended in that direction but were not significantly different from 0. One possible reason is that, in a trained movement task in which the animal responds to task cues, the movement involves a relatively consistent reaction time and planning period, whereas in the present study in the spontaneous movement set, the animal may be guiding the movements more on the fly and with less preplanning, resulting in shorter and more variable temporal offsets. As a result, the lag between neuronal activity and movement no longer reaches significance.

Figure 14 also shows that the temporal offsets for speed and direction models were more tightly distributed than for the arm-posture model. This pattern resembles previous findings (Paninski et al., 2004). One interpretation is that the control of direction and speed can be exercised only in association with the movement itself because these parameters are not defined when the arm is stationary, whereas the active control of posture must be exercised at all times. Therefore, one might expect the neuronal control of posture to extend over a greater range of times than the neuronal control of direction, resulting in a spread-out distribution of temporal offsets. A second speculation (proposed by Paninski et al., 2004) is that the modulation of neuronal activity by posture reflects both feedforward and feedback signals and therefore has a greater range of temporal offsets. As described next, however, other aspects of the data suggest that the postural tun-

ing reflected relatively more of a feedforward signal than a feedback signal.

Figure 14*E* shows the temporal offsets for start-posture tuning. These offsets were shifted with respect to the others and were highly significantly below 0 ($t = 4.94$; $p = 0.00004$). A comparison between Figure 14, *D* and *E*, shows the difference between end-posture tuning, for which the temporal lags were distributed broadly around the movement period and start-posture tuning, for which the temporal lags mainly preceded the movement. To further explore this effect, we extended the range of possible offsets downward from -286 to -1000 ms and reran the start-posture analysis. Figure 14*F* shows the resulting distribution of temporal offsets. The mean is -413 ms.

These results help to answer the question raised at the beginning of this section. Are the neurons tuned to the anticipated, upcoming posture, as though contributing to the control of the movement, or are they equally tuned to the previous, starting posture, as though reflecting a feedback signal? The activity during a movement is indeed tuned to the final posture of the movement, but it is not really tuned to the start posture. Rather, to obtain a high R^2 value for the start posture, it is necessary to select neuronal activity from well before the start of the movement. In effect, neuronal activity during the previous movement is correlated with the end posture of that previous movement. This result suggests that the neuronal tuning to posture reflects relatively more an anticipation of the future state of the limb rather than a feedback signal about a recent state of the limb. Both influences probably contribute, but the planned future posture appears to be the larger influence according to this analysis. The next section provides a second, independent approach to the same question and arrives at the same answer.

Neurons are more end-posture than start-posture tuned II: analysis of movements near the preferred posture

With relatively short movements (mean length of 8 cm), the end posture of a movement might often be similar to the start posture. Movements that end in the preferred region may therefore also tend to start in the preferred region and remain in the preferred region throughout the movement. Although the firing rate is typically high during these movements near the preferred posture, is it because the start posture was near the preferred posture or because the end posture was near the preferred posture?

To approach this issue, we performed the following analysis. For each neuron, we selected the 10% of movements that terminated nearest to the preferred posture. These movements were typically associated with a high firing rate, because they terminated near the peak of the postural tuning curve (for examples, see Fig. 11*B1*–*B4*). The purpose of the analysis was to examine this set of movements clustered around the preferred posture and to determine whether the firing rate was more related to the proximity between the start posture and the preferred posture or to the proximity between the end posture and the preferred posture.

We performed two linear regressions. First, we asked whether the distance between the start posture of a movement and the preferred posture was significantly related to firing rate. This was not the case. Across neurons, the mean R^2 value for this linear regression was 0.02, and none of the neurons showed a significant relationship. The firing rate during movement tended to be high because the movements were near the preferred posture, but the firing rate tended to be uniformly high regardless of the proximity of the start posture of the movement to the preferred posture.

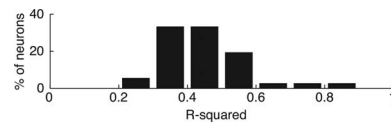In contrast, using the same subset of movements, we asked

Figure 15. Frequency histogram of R^2 values for hierarchical regression. This regression model included tuning to a preferred final posture, tuning to a preferred final hand position, tuning to a preferred direction that can rotate with changes in the starting hand position, and a linear dependence on mean speed.

whether the distance between the end posture of a movement and the preferred posture was significantly correlated with firing rate. A relationship was obtained in which the firing tended to be higher for movements that terminated closer to the preferred posture. The mean R^2 value for this linear regression was 0.4, and 71% of the neurons showed a significant relationship.

Thus, among the subset of movements near the preferred posture, the firing rate was dependant on the end posture and independent of the start posture. This result, like the result of the previous section, strongly suggests that the postural tuning obtained here reflects primarily an anticipation of the planned state of the limb rather than a feedback signal about the recent state of the limb.

Combined tuning: speed, direction, end point, and end posture

We combined speed tuning, direction tuning, end-point tuning, and end-posture tuning in a hierarchical regression to determine how much of the total variance can be explained by these variables. We first regressed against end posture, then regressed the residual against end point, then regressed the second residual against the direction-tuning model in which the preferred direction could rotate systematically with changes in the initial hand position, and finally regressed the third residual against the mean speed of the hand during each movement. Because of the possible partial overlap of the different tuning models, the combined R^2 value is expected to be less than the sum of the separate R^2 values given in the previous sections. Figure 15 shows the R^2 values obtained by this sequential regression. On average, 44% of the variance was explained by the combination of parameters. The remainder of the variance may be attributable to the many other parameters critical to movement but not tested here, to the changing motivational or attentional state of the animal, and to noise in the measurement of arm movement and neuronal activity.

Simulated neurons

To test the validity of the above regression models, we modeled artificial neurons. One artificial neuron was direction tuned. The data were generated using the cosine direction-tuning model with Gaussian noise added to create a neuron that was noisily tuned to a preferred direction. We simulated the firing rate that this artificial neuron should have during a set of actual, recorded hand movements. These firing rates were then subjected to the regression analyses described above. Figure 16*A* shows the result. When tested with a direction-tuned model, the artificial direction tuned neuron showed a midrange R^2 value consistent with its noisy tuning. When tested with the other regression models, it showed near zero R^2 values. Thus, when testing a neuron that is truly direction tuned, the analyses extracted a directional tuning signal and not a spurious hand-position-tuning or posture-tuning signal.

In a similar manner, we generated an artificial neuron that was

tuned to final hand location in space and an artificial neuron that was tuned to final posture of the arm. As shown in Figure 16, B and C, the R^2 value for each type of neuron was highest for the matching type of regression model and near zero for the non-matching regression models. Finally, we generated an artificial neuron that was not tuned. Its firing rate was randomized. Figure 16 D shows that all regression models returned a low R^2 value for this artificial neuron, with no tendency for the models that had more regressors to extract a higher R^2 value.

Muscle activity

We are currently performing an experiment examining muscle activity during untrained movement of the arm, and we report here some preliminary data to compare with the neuronal data. In these experiments, we measure electromyographic (EMG) activity. One technical limitation that we have not yet solved is that our method of tracking the arm requires powered LEDs fixed to the arm that generate electrical noise in the EMG signal. Thus far, recordings from the forearm and hand muscles are not reliable. Recording EMG from the upper arm has been more successful, partly because the muscles are bigger and produce more reliable signals and partly because fewer powered LEDs are fixed to the upper arm. Here we report on preliminary data from the biceps.

Fine insulated stainless steel wires were threaded into a 22 gauge syringe needle and inserted into the long biceps. The wires had an exposed tip of 1–2 mm. Three wires spaced ~5 mm apart were inserted to provide input to a differential amplifier and its ground (single-neuron amplifier model 1800; A-M Systems). The amplifier filters were set with a low cutoff at 300 Hz and a high cutoff at 5000 Hz. Although some studies use a lower level for the low cutoff (Cheney et al., 1985), we found that a 300 Hz cutoff was helpful to remove low-frequency noise. The EMG signal was digitized at a rate of 20,000 Hz, rectified, and integrated in 14.3 ms time bins.

We predicted that the EMG signal from the biceps would show good tuning to final posture, but that the tuning would differ from the typical tuning of motor cortex neurons in two ways. First, the neurons often were related to movement of more than one joint, whereas the biceps EMG should be primarily related to the elbow. Second, the neurons often preferred movement toward an intermediate joint angle, whereas the biceps EMG was expected to prefer a flexed elbow angle. In these ways, the muscle signal was expected to be simpler than the neuronal signals.

As expected, the biceps EMG was largest during flexion of the elbow joint. Figure 17A shows EMG traces during rotations of the elbow. Flexions of the elbow were associated with higher EMG activity than extensions of the elbow. Just as for the neuronal analysis, we performed a linear regression of the biceps EMG against joint angular velocity with eight degrees of joint motion. The total R^2 value for this regression was 0.41. Figure 17B shows the individual contributions of each joint when the covariance with the remaining seven joints was regressed out. As expected, the elbow joint shows the dominant effect. The three degrees of shoulder rotation show a small effect, probably because the movement of the shoulder is to a
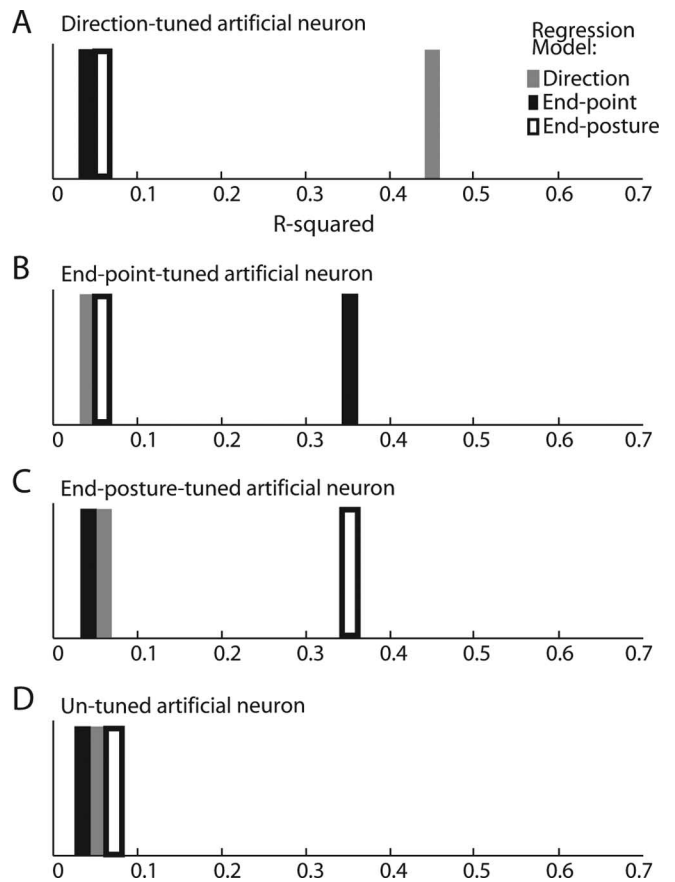

Figure 16. Analysis of simulated noisy neurons. *A*, Simulated direction-tuned neuron tested on direction, end-point, and end-posture regression models. *B*, Simulated end-point tuned neuron tested on the same three models. *C*, Simulated end-posture tuned neuron tested on the same three models. *D*, Simulated neuron with no tuning on the same three models.
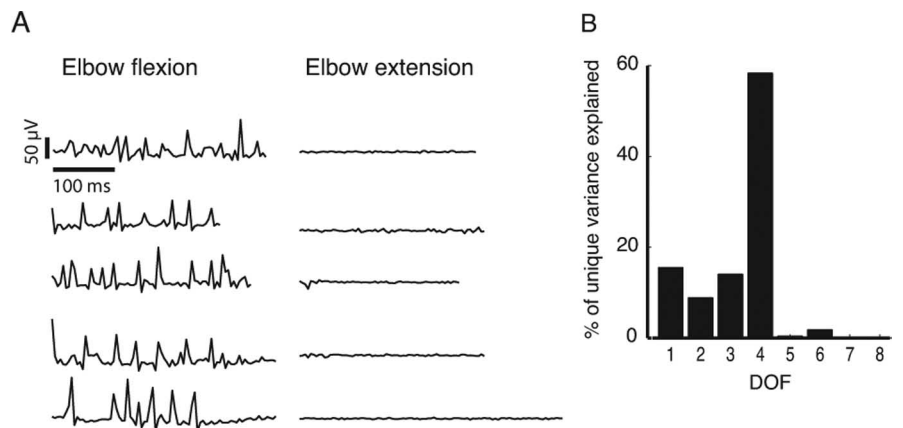

Figure 17. EMG results for biceps during naturalistic movements. *A*, Right, EMG traces during movements involving elbow flexion. Left, EMG traces during movements involving elbow extension. *B*, An eight-dimensional regression against joint velocity was performed and returned an $R^2 = 0.41$. The unique contribution for each joint, when the contributions of the other seven joints were regressed out, was calculated. The figure shows the percentage of contribution to the R^2 for each joint. As expected, the elbow joint showed the largest effect. Joints on *x*-axis are (in order) shoulder azimuth, shoulder elevation, shoulder external/internal rotation, elbow extension, forearm pronation, wrist extension, wrist adduction, and grip aperture. DOF, Degrees of freedom.

small degree mechanically coupled to the elbow. The four distal degrees of freedom show essentially no effect. The overall pattern shows that the activity of the biceps is related primarily to rotation of the elbow joint.

We then analyzed the biceps EMG using the model of a tuning

to a preferred final posture. Because the biceps EMG is largest during elbow movements toward flexion, the signal was expected to be relatively well tuned to final posture, in which the preferred final posture was a flexed elbow. The R^2 value for this tuning to final posture was 0.49. The EMG signal was indeed tuned to a flexed elbow posture. The preferred elbow angle obtained by the regression fit was 2°, outside the normal mechanical range of the elbow (~30–180°). That is, the biceps EMG was on average larger for movements that terminated in more flexed postures, resulting in an extrapolated preferred posture more flexed than the arm could actually go. In contrast, for most neurons, the regression analysis obtained a preferred final posture within the normal mechanical range of the elbow. Only one neuron had a preferred elbow angle outside of that range. The preferred final elbow angle for the biceps EMG was so extreme that it was highly significantly smaller than that obtained for the neurons ($p = 10^{-8}$).

These results show, as expected, that the biceps EMG activity was related primarily to the elbow and was tuned primarily to flexion, not to movement toward a posture in the normal range of the arm. In these ways, the pattern of results for the muscle was simpler and much clearer than the pattern described above for the neurons in motor cortex. It cannot be ruled out that some of the neuronal signals studied here are correlated primarily with muscle activity. Indeed, considerable direct control of muscle activity seems likely. However, to the extent that each neuron was related to many joints and to the extent that neurons often preferred movement of joints to intermediate angles rather than to extremes, the pattern of results was too complex to be entirely explained by a simple linear relationship to muscle activity.

Discussion

The present results suggest that the standard tuning properties of motor cortex neurons described previously in trained tasks are also present during semi-naturalistic, unstructured arm movements. However, the amount of neuronal variance explained by standard directional tuning may be less than suspected previously, at least in monkeys that have not received extensive daily training on a directional reaching task, whereas the variance explained by postural tuning may be more. On average, speed tuning accounted for ~1% of the total variance, simple direction tuning for 8%, a more complex version of direction tuning for 13%, tuning to the final position of the hand for 22%, and tuning to the final posture of the arm for 36%.

Interpretation of R^2

All R^2 values in the present study are below 50%. Do these seemingly low values indicate a weak or invalid result? A high R^2 value for a particular source of variance will be obtained only if all other sources of variance are minimized. In the case of speed tuning, it is possible, by averaging out the variance from other sources, to obtain an R^2 value that surpasses 90% (see above, Speed tuning I). However, when all sources of variance are admitted, although speed tuning remains significant, it explains only ~1% of the variance. In this case, the R^2 value can change for some neurons from 90 to 1% depending on the manner of the test but still reflect the same, valid speed-tuning signal.

Tuning to a final multijoint posture of the arm accounted for the largest share of the neuronal variance: 36%. An incorrect interpretation might suggest that this model cannot be valid because it explained far less than 100% of the variance and even less than 50%. However, the fact that the postural tuning captures one-third of the variance, in a grab bag with all other sources of

variance present, suggests that it is a major contributor to the signal.

Presumably the use of the arm for complex, normal behavior requires the control of numerous variables including hand speed, joint angle, muscle force, final posture, path curvature, grip size, the relative timing of arm and grip movement, and others. These variables are relevant to different extents in different movements. They are all presumably reflected in motor cortex neuronal activity. The variance in neuronal activity should be related in some proportion to each of these control variables, each one capturing a small part of the total variance. The present findings are consistent with the view that neurons are tuned to many control variables important in the animal's repertoire, and that posture is an especially prominent control variable capturing approximately one-third of the total variance.

One major caveat is that many of the movement variables are likely to be correlated with each other. If a neuron contributes to the control of variable *A* but variable *A* is correlated with variable *B*, then the neuron may appear on analysis to be tuned to *B*. Therefore, our results showing that neurons are tuned to many parameters must be taken cautiously. It cannot be excluded that some of the tuning to different movement variables reported here is a result of correlations among the variables. This difficulty of correlations among movement variables has been noted before and has led to some alternative interpretations to direction tuning (Scott and Kalaska, 1995, 1997). Arguably, the correlative nature of single-neuron experiments has been the main obstacle to interpreting single-neuron results in motor cortex.

Local direction tuning and global posture tuning

In the now classical paradigm of Georgopoulos et al. (1982, 1986), monkeys perform a center-out task, reaching from a central location to surrounding locations. Most neurons are broadly direction tuned, preferring one direction of reach and firing less during neighboring directions. Furthermore, when different starting and ending positions are tested, direction tuning still predominates (Georgopoulos et al., 1985; Kettner et al., 1988; Caminiti et al., 1990; Fu et al., 1993). However, the present experiment suggests that, in a more global range of movements, the neuronal variance is only ~8% attributable to this type of direction tuning. Other types of tuning, such as end-point and end-posture tuning, explain more of the variance.

These seemingly disparate results are not, however, in contradiction. In a center-out task and its variants, the direction of the hand is typically explored through a full range of angles, yet the posture of the arm is tested over a narrow range. In most experiments, the hand is limited to a central region of space comprising <5% of the volume tested here. Because the tuning curves for posture and hand position are broad, they will drive little of the neuronal variance if tested over this small range. As a result, when testing local regions of space with a full range of hand directions, direction tuning will tend to predominate.

Our results replicate previous findings that the neurons are locally but not globally direction tuned. The directional tuning changes when the hand is placed in a new starting position or when the arm is given a new starting posture (Caminiti et al., 1990; Scott and Kalaska, 1995, 1997; Sergio and Kalaska, 2003). The present findings further suggest that the behavior of the neurons over the global movement set may be better described as a broad tuning to a preferred final posture.

Many of the actions performed by monkeys (and humans) require placing the arm, including both proximal and distal joints, in a canonical posture and making smaller adjustments to

that posture. To a first approximation, the behavior of the arm is best described as a set of useful postures rather than as a set of useful directional vectors. For example, manipulating an object between the hand and the teeth involves a canonical arm posture including the shoulder in a low elevation, the elbow flexed, the forearm supinated, the wrist slightly flexed, and the grip closed. Within this canonical posture, smaller adjustments are made to each of these joints to manipulate the object at the mouth. An example from the human motor repertoire is typing: the arm remains in a canonical posture with the upper arm vertically downward, the elbow flexed to ~90°, the forearm pronated, and the wrist straight. Superimposed on this relatively stable posture is a set of smaller directional adjustments. Arguably, in this type of behavior, the proximal joints require more stabilization than the distal joints, but all joints whether proximal or distal participate in the canonical posture and also participate in the finer directional adjustments around that posture needed to perform the action. We speculate that so many actions in the primate repertoire have this architecture of a global posture and local directional adjustments to that posture that neurons in motor cortex may have come to reflect that structure in their movement tuning.

A recent study suggests that neurons compensate for a force load on the hand in a manner that emphasizes either postural coding or directional coding (Kurtzer et al., 2005). Other recent work suggests that, in the human, postural coding may be emphasized in the right hemisphere in motor cortex and directional coding may be emphasized in the left hemisphere in motor cortex, although both types of processing are of course represented in both hemispheres (Bagesteiro and Sainburg, 2002). Exactly how these results relate to the present study is not yet clear. Most normal actions require both a postural component and a directional component, as well as control of other aspects of movement, yet it may be that some actions emphasize postural control whereas others emphasize directional control.

Much of the literature in motor control asks how the hand is moved from point A to point B. Is it moved via a postural controller (Rosenbaum et al., 1995; Ghafouri and Feldman, 2001) or a directional controller? We suggest, however, that almost all normal behavior requires a much greater range of objectives than merely a cursor-like movement of the hand from point to point. Control of posture is not merely a means of moving the hand around but can be a prominent part of the behavioral task. The arm configuration provides the correct hand orientation, stability along the correct axis required for the task, a maximum of comfort, and prevention of different parts of the arm from colliding with the body or with nearby objects. For example, bringing the hand to the mouth is not merely a matter of translating the hand, cursor-like, to the mouth; it involves a characteristic whole posture of the arm. Moreover, in some types of behavior, such as nudging something with the elbow or flinging up the arm into a protective posture to block a threat, the position of the hand in space is relatively unimportant and the posture of the arm becomes the primary task requirement. For these reasons, it does not seem surprising that motor cortex should contain such robust signals concerning the configuration of the arm or that postural signals and directional signals should be combined as they are in the behavior.

References

Aflalo T, Graziano MSA (2006) Partial tuning of motor cortex neurons to final posture in a free-moving paradigm. Proc Natl Acad Sci USA 103:2909–2914.

Ashe J, Georgopoulos AP (1994) Movement parameters and neural activity in motor cortex and area 5. Cereb Cortex 4:590–600.

Bagesteiro LB, Sainburg RL (2002) Handedness: dominant arm advantages in control of limb dynamics. J Neurophysiol 88:2408–2421.

Bates DM, Watts DG (1988) Nonlinear regression analysis and its applications. New York: Wiley.

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

Cheney PD, Fetz EE, Palmer SS (1985) Patterns of facilitation and suppression of antagonist forelimb muscles from motor cortex sites in the awake monkey. J Neurophysiol 53:805–820.

Cohen J, Cohen P, West SG, Aiken LS (2003) Applied multiple regression/correlation analysis for the behavioral sciences, Ed 3. Mahwah, NJ: Erlbaum.

Cooke DF, Graziano MSA (2004a) Sensorimotor integration in the precentral gyrus: polysensory neurons and defensive movements. J Neurophysiol 91:1648–1660.

Cooke DF, Graziano MSA (2004b) Super-flinchers and nerves of steel: defensive movements altered by chemical manipulation of a cortical motor area. Neuron 43:585–593.

Desmurget M, Prablanc C (1997) Postural control of three-dimensional prehension movements. J Neurophysiol 77:452–464.

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

Feldman AG (1986) Once more on the equilibrium-point hypothesis (lambda model) for motor control. J Mot Behav 18:17–54.

Fu QG, Suarez JI, Ebner TJ (1993) 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.

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.

Georgopoulos AP, Caminiti R, Kalaska JF (1984) Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. Exp Brain Res 54:446–454.

Georgopoulos AP, Kalaska JF, Caminiti R (1985) Relations between two-dimensional arm movements and single-cell discharge in motor cortex and area 5: movement direction vs. movement endpoint. In: Hand function and the neocortex (Goodwin A, Darian-Smith I, eds). Exp Brain Res [Suppl] 10:175–183.

Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233:1416–1419.

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

Ghafouri M, Feldman AG (2001) The timing of control signals underlying fast point-to-point arm movements. Exp Brain Res 137:411–423.

Graziano MSA, Taylor CSR, Moore T (2002a) Complex movements evoked by microstimulation of precentral cortex. Neuron 34:841–851.

Graziano MSA, Taylor CSR, Moore T, Cooke DF (2002b) The cortical control of movement revisited. Neuron 36:349–362.

Graziano MSA, Cooke DF, Taylor CSR, Moore T (2004a) Distribution of hand location in monkeys during spontaneous behavior. Exp Brain Res 155:30–36.

Graziano MSA, Patel KT, Taylor CSR (2004b) Mapping from motor cortex to biceps and triceps altered by elbow angle. J Neurophysiol 92:395–407.

Graziano MSA, Aflalo T, Cooke DF (2005) Arm movements evoked by electrical stimulation in the motor cortex of monkeys. J Neurophysiol 94:4209–4223.

Hocherman S, Wise SP (1991) Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. Exp Brain Res 83:285–302.

Holdefer RN, Miller LE (2002) Primary motor cortical neurons encode functional muscle synergies. Exp Brain Res 146:233–243.

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

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.

Kurtzer I, Herter TM, Scott SH (2005) Random change in cortical load representation suggests distinct control of posture and movement. Nat Neurosci 8:498–504.

Moran DW, Schwartz AB (1999) Motor cortical representation of speed and direction during reaching. J Neurophysiol 82:2676–2692.

Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP (2004) Spatiotemporal tuning of motor cortical neurons for hand position and velocity. J Neurophysiol 91:515–532.

Reina GA, Moran DW, Schwartz AB (2001) On the relationship between joint angular velocity and motor cortical discharge during reaching. J Neurophysiol 85:2576–2589.

Rosenbaum DA, Loukopoulos LD, Meulenbroek RG, Vaughan J, Engelbrecht SE (1995) Planning reaches by evaluating stored postures. Psychol Rev 102:28–67.

Scott SH, Kalaska JF (1995) Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures J Neurophysiol 73:2563–2567.

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

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

Todorov E, Jordon MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226–1235.

Wu W, Hatsopoulos N (2006) Evidence against a single coordinate system representation in the motor cortex. Exp Brain Res 175:197–210.