

Large-Scale Organization of Preferred Directions in the Motor Cortex.

I. Motor Cortical Hyperacuity for Forward Reaching

Thomas Naselaris,¹ Hugo Merchant,^{1,2,6} Bagrat Amirikian,^{1,2} and Apostolos P. Georgopoulos^{1,2,3,4,5}

¹Brain Sciences Center, Veterans Affairs Medical Center; ²Department of Neuroscience, ³Department of Neurology, and ⁴Department of Psychiatry, University of Minnesota Medical School; ⁵Cognitive Sciences Center, University of Minnesota, Minneapolis, Minnesota; and ⁶Instituto de Neurobiología, Universidad Nacional Autónoma de México, Campus Juriquilla, Querétaro, Mexico

Submitted 8 May 2006; accepted in final form 6 September 2006

Naselaris, Thomas, Hugo Merchant, Bagrat Amirikian, and Apostolos P. Georgopoulos. Large-scale organization of preferred directions in the motor cortex. I. Motor cortical hyperacuity for forward reaching. *J Neurophysiol* 96: 3231–3236, 2006. First published September 13, 2006; doi:10.1152/jn.00487.2006. We used statistical methods for spherical density estimation to evaluate the distribution of preferred directions of motor cortical cells recorded from monkeys making reaching movements in 3D space. We found that this distribution, although broad enough to represent the entire 3D continuum of reaching directions, exhibited an enrichment for reaching forward from the body and, to a lesser degree, for reaching backward toward the body. The distribution of preferred directions of cells in the motor cortex may have important implications for motor cortical function and for the decoding of arm trajectories from population activity.

INTRODUCTION

Directional tuning in reaching is a prominent feature of motor cortical activity (Caminiti et al. 1990; Georgopoulos et al. 1982, 1986; Schwartz 1994). The tuning of a motor cortical cell is characterized by its preferred direction (PD), which is a unit vector that specifies the direction of reaching at which the cell's discharge rate peaks. Although PDs have been calculated for thousands of cells, the detailed distribution of PDs (DPD) across the directional continuum has not yet been investigated.

A rigorous characterization of this distribution could influence our understanding of how the directional signal generated by populations of motor cortical neurons is decoded, or “read-out.” Such knowledge is important in the context of understanding how natural reaching behaviors are controlled by the motor system and in the context of achieving control of neuroprosthetic devices (Schwartz 2004; Taylor et al. 2002). The structure of the DPD may also have important implications for our understanding of motor cortical plasticity, as discussed in the following text.

It is important to investigate the DPD in the context of unconstrained reaching movements in three-dimensional (3D) space. Such movements occur naturally and most frequently in the everyday life of primates and thus it is reasonable to suppose that the neural mechanisms of reaching would relate to them. In contrast, two-dimensional (2D) movements (Georgopoulos et al. 1982) can only partially relate to those neural mechanisms because they are, by definition, constrained. Furthermore, correct inferences concerning the DPD can be obtained only from the 3D case. Although the basic functional

property of directional tuning is present for both 3D and 2D movements, a particular PD obtained from 2D movements represents a projection into the movement plane of the cell's true, 3D PD; therefore the 2D PD is consistent with an infinite number of 3D PDs (Amirikian and Georgopoulos 2003). In this work, we considered only PDs that were estimated in the context of a 3D reaching task.

Current knowledge about the DPD for unconstrained, 3D movements derives from three studies (Caminiti et al. 1990; Mitsuda and Onorati 2002; Schwartz et al. 1988). In the first investigation of the DPD (Schwartz et al. 1988), it was observed that a sample of roughly 500 PDs extracted from two monkeys was dispersed across the 3D directional continuum. It was also observed that the DPD exhibited an enrichment for reaches made forward, away from the body; however, when applied to this sample, a Rayleigh test for uniformity against a unimodal alternative did not reject the null hypothesis of uniformity. In the study of Caminiti et al. (1990), a sample of nearly 200 PDs extracted from three monkeys again revealed that the DPD is sufficiently broad to represent all portions of the 3D continuum, although the null hypothesis of uniformity was not explicitly tested against any other alternatives. Most recently, it was reported in Mitsuda and Onorati (2002) that the hypothesis of uniformity was rejected by a Rayleigh test ($P < 0.05$) for a sample of 126 PDs obtained from a single monkey, although explicit measures of the location and strength of the implied bias in the DPD were not provided.

In this study we analyzed the structure of the DPD using a variety of statistical tests and methods for density estimation. We applied these methods to a sample that included the original data described in Schwartz et al. (1988), as well as a more recent data set that includes roughly 800 PDs extracted from widely distributed recording sites in the motor cortex of two monkeys (Naselaris et al. 2005). Our findings are consistent with previous observations that PDs are distributed broadly across the 3D directional continuum. However, we found that the DPD is not *strictly* uniform, as the result of an enrichment of the representation for forward reaching directions (as observed originally; Schwartz et al. 1988), as well as for backward reaches directed toward the body. This enrichment induces a relative increase in accuracy of the motor-cortical representation for forward and backward reaches, as measured by the neuronal population vector; however, when averaged across all directions, the effect on accuracy is quite

Address for reprint requests and other correspondence: A. Georgopoulos, Brain Sciences Center, Veterans Affairs Medical Center, University of Minnesota, Minneapolis, MN 55455 (E-mail: omega@umn.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

small. Our results are consistent with previous demonstrations (Georgopoulos et al. 1986, 1988; Schwartz 1994) that the distribution of PDs in motor cortex provides an excellent substrate for a population-level code of reaching direction.

METHODS

Behavioral task

The neural data used in these studies were collected from four rhesus monkeys (NI, IO, A, and B) engaged in the 3D center-out task (Georgopoulos et al. 1986). Light-cued reaches were made from a fixed, central starting position toward eight targets located near the corners of a cube. Data from all four monkeys refer to reaches of the left arm. The starting position was level with the shoulder of the monkey's reaching arm. The monkey was instructed to hold the starting position for a variable amount of time (0.5–1.5 s), after which a button located on a randomly selected target was illuminated. The monkeys reached toward the target and depressed the button for 0.08 s to obtain a juice reward. For each recording site, monkeys performed 40 reaches toward the targets (five blocks of eight reaches; targets were presented in random order within each block).

Neural recordings

All recordings were obtained from the right hemisphere of the motor cortex. A detailed account of the recording protocols for monkeys NI and IO was previously given elsewhere (Schwartz et al. 1988). For monkeys A and B, recording sites were located at various depths beneath a patch of cortex that extended 3–4 mm along the central sulcus and 7–12 mm in the direction perpendicular to the central sulcus. This region was centered about 15 mm from the midline along the medial–lateral axis. Microstimulation at 5–20 μ A within the boundaries of this region evoked contractions of distal and proximal arm muscles. Recording sites along each electrode penetration were spaced 150 μ m apart. At each site, raw extracellular membrane potentials were sampled at 60 kHz. Single-unit activity was extracted from these records using the Plexon off-line sorter (Plexon, Dallas, TX).

Calculation of PDs

PDs were calculated using a multiple linear regression analysis (Georgopoulos et al. 1986). For each reach, an average firing rate was obtained for the interval spanning from 0 to 60 ms after the onset of the target to the end of the movement. The average firing rate r was treated as the dependent variable in a regression relating the direction of reach to the firing rate

$$r = b_0 + b_x d_x + b_y d_y + b_z d_z$$

where d_x , d_y , and d_z are the direction cosines of the target and b_0 , b_x , b_y , and b_z are regression coefficients. The components of the cell's preferred direction (PD) vector $\mathbf{p}$ are given by

$$p_x = b_x/k \quad p_y = b_y/k \quad p_z = b_z/k$$

where

$$k = (b_x^2 + b_y^2 + b_z^2)^{1/2}$$

All cells included in these analyses demonstrated a significant directional tuning effect ($P < 0.05$), as indicated by a bootstrapping procedure (Lurito et al. 1991).

Nonparametric density estimation

We used a kernel-based method to obtain a nonparametric estimate of the DPD as the probability density $g(\mathbf{x})$, of reaching direction $\mathbf{x}$, given a set of n experimentally determined PDs

$$g(\mathbf{x}) = \frac{1}{n} \sum_i f(\mathbf{x}|\mathbf{c}, \mathbf{p}_i)$$

where $\mathbf{p}_i$ is the PD of cell i . The kernel function f is the von Mises–Fisher (vMF) probability density function (Fisher et al. 1987)

$$f(\mathbf{x}|\kappa, \boldsymbol{\mu}) = \frac{\kappa}{4\pi \sinh \kappa} \exp(\kappa \mathbf{x}^T \boldsymbol{\mu})$$

In this context, $\kappa > 0$ acts as a smoothing parameter.

The value of κ was determined using a maximum likelihood with cross-validation approach suggested in Fisher et al. (1987). In this approach, we estimate the density at $\mathbf{p}_j$ using a sample of $n - 1$ PDs created by omitting $\mathbf{p}_j$

$$g_j(\mathbf{p}_j) = \frac{1}{n-1} \sum_{i \neq j} f(\mathbf{p}_j|\kappa, \mathbf{p}_i)$$

We then maximized the log-likelihood function $L(\kappa) = \sum_i \log [g_i(\mathbf{p}_i)]$, with respect to κ over all n PDs in the sample.

Parametric density estimation

The parametric model of the DPD was specified as

$$g(\mathbf{x}|\boldsymbol{\alpha}, \boldsymbol{\kappa}, \mathbf{M}) = \sum_{i=1}^m \alpha_i f(\mathbf{x}|\kappa_i, \boldsymbol{\mu}_i)$$

where $\alpha_i \in \boldsymbol{\alpha}$ is a mixing coefficient, $\kappa_i \in \boldsymbol{\kappa}$ is the concentration parameter for the i th component of the model, and $m = 2$ is the number of components. The density associated with the i th component peaks at $\boldsymbol{\mu}_i \in \mathbf{M}$, the mode location. The size of the peak at $\boldsymbol{\mu}_i$ is determined by $\kappa_i > 0$ and $\alpha_i > 0$; peak size increases as κ_i and/or α_i increases. The mixing coefficients satisfy $\sum \alpha_i = 1$ and are interpreted as the probability that a given PD belongs to component i . $\boldsymbol{\alpha}$, $\boldsymbol{\kappa}$, and $\mathbf{M}$ were estimated using the expectation maximization (EM) procedure (Banerjee et al. 2003). To avoid solutions corresponding to local maxima, the EM algorithm was run 100 times from different, randomly chosen initial conditions ($0 < \kappa < 1$ for each initial condition). Solutions obtained from these randomly varying initial conditions were highly consistent. The solution that generated the highest likelihood under the model is reported in Table 1.

Test for significance of density peaks in the DPD

We used a bootstrapping procedure to test the significance of the deviation from uniformity of individual peaks in the nonparametric distribution. In this procedure, 10^3 samples of n PDs were drawn from a uniform distribution on the sphere. For each such sample, a nonparametric density estimate was constructed using the procedure described above and the maximum density values were stored. The resulting distribution was used to test the significance of local peaks in the density values obtained from the nonparametric model of the DPD.

Decoding of simulated neural activity

The firing rate $r(\mathbf{x})$ for a simulated cell when a reach is made in direction $\mathbf{x}$ was defined as $r(\mathbf{x}) = N(\mathbf{x})/\Delta t$. The temporal interval Δt was set to 700 ms, corresponding to the average trial duration. The spike count $N(\mathbf{x})$ was drawn from a Poisson distribution with mean given by $(b_0 + k\mathbf{p}^T \mathbf{x})\Delta t$, where b_0 and k were drawn at random from the population of real motor-cortical cells and $\mathbf{p}$ was drawn either from this same population, a uniform distribution on the sphere, or a two-component mixture of vMF distributions. Movement direction was decoded from a population of such cells using the neuronal population vector (NPV; Georgopoulos et al. 1986). The normalized

TABLE 1. Summary of the distribution of preferred directions for separate data sets

Monkey	N	Forward Component				Backward Component				Angle Between Components, deg	Bingham (P <)
		α	κ	μ_1 , deg (θ, ϕ)	Bootstrap (P <)	α	κ	μ_2 , deg (θ, ϕ)	Bootstrap (P <)		
A, B, NI, IO	1,139	0.42	2.6	(66, 184)	10^{-3}	0.58	1.2	(133, 354)	10^{-3}	160	10^{-6}
A, B	848	0.43	2.76	(63, 185)	0.01	0.57	1.15	(138, 346)	0.01	154	10^{-6}
NI, IO	291	0.74	0.81	(78, 167)	0.05	0.26	4.35	(121, 350)	0.01	158	10^{-6}

firing rate for each cell, in the form $(r - b_0)/k$, was used as the weighting factor for calculating the NPV. Decoding error was measured as the angle between the actual reaching direction and the direction estimated by the decoding algorithm.

RESULTS

We analyzed a sample of 1,139 PDs recorded from four monkeys performing a 3D center-out task. We will refer to this as the *combined sample*. Recordings were obtained from sites across the arm area of the primary motor cortex, including sites in the anterior bank of the central sulcus, and the exposed surface of the precentral gyrus. An equal-area projection plot of the PDs in the combined sample (Fig. 1D; see Fig. 1, A–C for an explanation of the projected coordinate system) illus-

trates the two most salient features of the DPD. First, PDs cover the entire 3D directional continuum. Second, there is a local increase in PD density for reaches that are almost straight forward ($\phi \approx 180^\circ$) and slightly above the equator ($\theta < 90^\circ$) and another, smaller increase in PD density for reaches that are nearly backward ($\phi \approx 360^\circ$) and below the equator ($\theta > 90^\circ$).

We estimated PD density directly by constructing a nonparametric model of the DPD (Figs. 1, E and F and 2). The nonparametric model revealed local peaks in PD density for forward and backward reaches; it also confirmed that the magnitude of the backward peak was clearly smaller than that of the forward peak. A two-component mixture model was then used to capture locations of the peaks (black dots, Figs. 1E and 2; see Table 1). The ratio of PD densities at the forward




FIG. 1. Distribution of preferred directions (DPDs). A: a PD (black arrow) is a unit vector that defines a point on the surface of the sphere (outlined in black) whose center (red dot) rests at the starting position for the reaching movements. In this consideration, a PD is specified by a pair of angles (θ, ϕ). θ gives the angular deviation of the PD from the north pole ($\theta = 0^\circ$); ϕ gives the angular deviation in the plane from the axis that points directly toward the monkey ($\phi = 0^\circ$). Thus a reach in the horizontal plane directly toward the monkey is given by the pair $(90^\circ, 0^\circ)$, whereas a reach directly away is given by $(90^\circ, 180^\circ)$. B: surfaces of the backward (left) and forward (right) hemispheres of the sphere mapped into separate planes using an equal area projection. Blue numbers represent values of θ , which varies along the dashed blue meridian. Red numbers represent values of ϕ , which varies along the dashed red equator. Backward and forward poles are located at the intersection of the meridian and the equator. C: similar to B, but with spherical coordinates replaced by reach directions, specified from the perspective of the monkey. D: an equal-area projection plot of the combined sample of 1,139 PDs. E: to obtain a representation of the DPD, the surface of the sphere was partitioned into a grid of 10^3 equal-area patches. A nonparametric density estimate was obtained at each node on the grid. Grid node was then mapped into the plane under the equal-area projection; the corresponding density value was plotted in grayscale. Black dots indicate the mode locations estimated from the mixture model. On the grayscale bar, the left arrow indicates $1/4\pi \approx 0.08$, the density value predicted by a uniform distribution on the sphere. Right arrow indicates a significance bound of $P = 10^{-3}$, as determined by a bootstrapping test. F: 3D representation of the DPD. Sphere is rotated so that the forward pole appears in the foreground. Monkey is located near the red X. Perspective is that of an observer facing the monkey; thus the hemisphere to the viewer's right of the blue meridian corresponds to reaching directions to the monkey's left. Gray cylinders extend from the up/down and left/right poles. Meridians correspond by color with those seen in B and C. Green arrow extending from the surface shows the direction corresponding to the peak of the DPD. In this figure, density is given by a color scale.

and backward peak locations identified by the mixture model is 1.3, indicating that the enrichment of the representation for backward reaches is indeed weaker than that for forward reaches.

To determine whether the forward and backward peaks represent true departures from uniformity, we tested their significance using a bootstrap procedure. This test revealed that both the forward and backward peaks are significantly ($P < 10^{-3}$) larger than the random fluctuations in density expected from a uniform DPD (left and right arrows in Fig. 1E show expected density and bound at which $P = 10^{-3}$, respectively). Consistent with this result, the Bingham test of uniformity against a bimodal alternative rejected the null hypothesis of uniformity ($P < 0.01$; see Table 1).

To test the consistency of the properties of the DPD, we applied the same statistical tests and density estimation procedures to two data sets separately (Fig. 2). Enrichments of forward and backward reaching directions are evident in both data sets. The results of the tests of significance for the separate data sets, shown in Table 1, leave little doubt that the properties of the DPD are robust.

The nonuniformity of the DPD could affect the accuracy with which movement direction can be decoded from neuronal population activity, enhancing decoding accuracy for the forward and backward reaching directions, while introducing a global bias that increases the average level of decoding error. A widely used method for decoding arm directions is the neuronal population vector (NPV) (Caminiti et al. 1991; Georgopoulos et al. 1986, 1988; Schwartz 1994; Steinberg et al. 2002; Taylor et al. 2002). The NPV is a weighted sum of PDs and is thus sensitive to first-order features of the underlying PD distribution (Georgopoulos et al. 1988; Mussa-Ivaldi 1988; Salinas and Abbott 1994;

Sanger 1994; Scott et al. 2001; Steinberg et al. 2002). To assess the effect of the DPD on global decoding accuracy, we calculated population vectors from the activity of simulated cells whose PDs were drawn at random from our experimental sample. The asymptotic level of error, averaged over 100 uniformly and randomly distributed reaching directions, is roughly 12° (Fig. 3A). Thus the statistically significant bias in the DPD has a fairly small effect on average decoding accuracy. However, the error of the NPV is smallest for reaching directions that are close (or opposite) to the forward peak in the DPD (Fig. 3B). Furthermore, the length of the NPV—which is an indication of its reliability as an estimate of reaching direction—is greatest for reaching directions near (or opposite to) the forward peak. Thus the quality of the directional command produced by the NPV is greatest for reaching directions near the forward/backward peaks in the DPD (Fig. 3C).

As shown in Fig. 3B, there is a local minimum in the error for reaches at exactly 90° to the forward peak. For reaches in this direction, cells with preferred directions near the forward peak do not fire at rates significantly different from baseline and thus make no contribution to the population vector. As a consequence, the population vector length reaches a global minimum for these directions.

DISCUSSION

Our findings are consistent with previous observations that preferred directions in the motor cortex are sufficiently broadly dispersed to cover the full continuum of 3D reaching directions (Caminiti et al. 1990; Schwartz et al. 1988). We have also shown that embedded within this broad dispersal of PDs is an enrichment for reaches directed forward (Schwartz et al. 1988) and an enrichment for reaches directed backward. These enrichments are manifested as local peaks in the DPD, and represent a significant departure from strict uniformity. The structure of the DPD permits a highly accurate representation of the direction of reach, as measured by the application of the NPV. The accuracy of this representation is greatest for forward and backward reaching directions.

What is the explanation for the enrichment of the motor cortical representation for forward and backward reaches? One possible explanation has to do with the motor cortex's capacity for massive, behaviorally dependent reorganization (Sanes and Donoghue 2000). Motor-cortical plasticity has been demonstrated using a variety of techniques, including intracortical microstimulation (Nudo et al. 1996), slice preparations (Hess and Donoghue 1994), functional magnetic resonance imaging (Karni et al. 1995), and transcranial magnetic stimulation (TCMS) (Classen et al. 1998; Pascual-Leone et al. 1995). Collectively, work on motor-cortical plasticity has demonstrated that the cortical representation of movements and movement sequences associated with the learning and production of a stereotyped motor task are enlarged as a result of the learning. This well-replicated finding has been interpreted as evidence that "motor practice induces the recruitment of additional M1 units into a network specifically representing the trained motor sequence" (Karni et al. 1995).

One study that supports this interpretation is particularly important in the context of our findings. In the study of Classen et al. (1998), focal TCMS of the motor cortex of




FIG. 2. NONPARAMETRIC DENSITY ESTIMATES FOR 2 SEPARATE DATA SETS.




FIG. 3. Effect of DPD structure on decoding accuracy. *A*: decoding error as a function of population size. Neuronal population vector (NPV) was applied to simulated populations of varying size. Error associated with the NPV (solid curve) asymptotes at about 12°. If the underlying population had uniformly distributed PDs, error would asymptote at 0° (dashed curve). Error bars show SDs across 50 populations of a given size; they are suppressed on the dashed curve for clarity. *B*: decoding error as a function of reaching direction. Data represent an average of >500 separate populations of 100 simulated cells whose PDs were drawn from the 2-component mixture model (see Table 1, top row, for parameters). For each population, decoding errors and NPV lengths were evaluated for 500 different reaching directions, chosen from a uniform distribution on the sphere. Abscissa gives the angle between each reaching direction and the forward peak of the DPD (μ_1 , top row of Table 1). Black curve gives decoding error for the NPV (left axis). Gray curve gives the length of the NPV (right axis). Error bars show SDs across the 500 separate cell populations. NPV length is normalized to fall within the range [0, 1]. *C*: a signal-to-noise ratio was constructed by dividing the length of the NPV by the error (gray and black curves in *B*, respectively).

humans consistently evoked thumb movements in a single direction. Subjects were then instructed to repeatedly produce movements of the thumb in the opposite of the evoked direction for 30 min. Subsequently, the direction of the TCMS-evoked thumb movement changed to match the “trained” direction. This finding suggests that the repetition

of a movement in a single direction can bias the motor-cortical representation of direction toward the repeated one.

Given these results on motor-cortical plasticity, we conjecture that the enrichment of the motor cortical representation for forward and backward reaching directions may simply reflect an increase in the incidence of reaches in these directions, relative to other directions, in the everyday life of the monkey. A direct test of this conjecture would require a characterization of the statistics of the natural reaching behavior of the monkey and a thorough sampling of PDs taken before and after intentional manipulation of these statistics. To our knowledge, such an experiment has not yet been attempted.

Whether the enrichment for specific reaching directions in the DPD has any functional effect on arm movement depends on the manner in which motor cortical commands are “read out,” or translated into action by circuits in the spinal cord. If read-out is implemented by a weighted sum of cortical activities (an NPV-like mechanism), the nonuniformity of the DPD can induce a relative increase in accuracy (hyperacuity) for forward and backward reaches (see Fig. 3C). However, if read-out is implemented by an optimal method (Salinas and Abbott 1994), the direction-dependent errors will be eliminated. Our results thus provide a basis for testing whether an NPV-like read-out mechanism is implemented by the motor system. Although some studies have shown directionally dependent variations in reaching accuracy (Lamotte and Acuna 1978; Smyrnis et al. 2000), no measurements of reaching accuracy have been made for 3D movements performed by monkeys with a known DPD.

A final important issue related to the observed enrichment of forward and backward directions in the DPD is the spatial scale at which this enrichment exists. PDs in this study were extracted from cells dispersed across a large (about 30 mm²) cortical region. The large sampling area underlying our database thus admits of at least two possibilities for the spatial organization of PDs. One possibility is that spatially separate motor-cortical regions may generate the distinct peaks apparent in the DPD. Alternatively, the structure of the DPD may be replicated locally in all regions of the motor cortex. We address this issue directly in the companion paper.

REFERENCES

- Amirikian B and Georgopoulos AP. Modular organization of directionally tuned cells in the motor cortex: is there a short-range order? *Proc Natl Acad Sci USA* 100: 12474–12479, 2003.
- Banerjee A, Inderjit D, Joydeep G, and Suvrit S. Generative model-based clustering of directional data. In: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Washington, DC: ACM Press, 2003, p. 19–28.
- Caminiti R, Johnson PB, Galli C, Ferraina S, and Burnod Y. Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. *J Neurosci* 11: 1182–1197, 1991.
- Caminiti R, Johnson PB, and Urbano A. Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *J Neurosci* 10: 2039–2058, 1990.
- Classen J, Liepert J, Wise SP, Hallett M, and Cohen LG. Rapid plasticity of human cortical movement representation induced by practice. *J Neurophysiol* 79: 1117–1123, 1998.
- Fisher NI, Lewis T, and Embleton BJJ. *Statistical Analysis of Spherical Data*. Cambridge, UK: Cambridge Univ. Press, 1987.
- Georgopoulos AP, Kalaska JF, Caminiti R, and Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2: 1527–1537, 1982.

- Georgopoulos AP, Kettner RE, and Schwartz AB.** Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J Neurosci* 8: 2928–2937, 1988.
- Georgopoulos AP, Schwartz AB, and Kettner RE.** Neuronal population coding of movement direction. *Science* 233: 1416–1419, 1986.
- Hess G and Donoghue JP.** Long-term potentiation of horizontal connections provides a mechanism to reorganize cortical motor maps. *J Neurophysiol* 71: 2543–2547, 1994.
- Karni A, Meyer G, Jezard P, Adams MM, Turner R, and Ungerleider LG.** Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377: 155–158, 1995.
- Lamotte RH and Acuna C.** Defects in accuracy of reaching after removal of posterior parietal cortex in monkeys. *Brain Res* 139: 309–326, 1978.
- Lurito JT, Georgakopoulos T, and Georgopoulos AP.** Cognitive spatial-motor processes. 7. The making of movements at an angle from a stimulus direction: studies of motor cortical activity at the single cell and population levels. *Exp Brain Res* 87: 562–580, 1991.
- Mitsuda T and Onorati P.** Three-dimensional tuning profile of motor cortical activity during arm movements. *Neuroreport* 13: 1477–1480, 2002.
- Mussa-Ivaldi FA.** Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci Lett* 91: 106–111, 1988.
- Naselaris T, Merchant H, Amirikian B, and Georgopoulos AP.** Spatial reconstruction of trajectories of an array of recording microelectrodes. *J Neurophysiol* 93: 2318–2330, 2005.
- Nudo RJ, Milliken GW, Jenkins WM, and Merzenich MM.** Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *J Neurosci* 16: 785–807, 1996.
- Pascual-Leone A, Nguyet D, Cohen LG, Brasil-Neto JP, Cammarota A, and Hallett M.** Modulation of muscle responses evoked by transcranial magnetic stimulation during the acquisition of new fine motor skills. *J Neurophysiol* 74: 1037–1045, 1995.
- Salinas E and Abbott LF.** Vector reconstruction from firing rates. *J Comput Neurosci* 1: 89–107, 1994.
- Sanes JN and Donoghue JP.** Plasticity and primary motor cortex. *Annu Rev Neurosci* 23: 393–415, 2000.
- Sanger TD.** Theoretical considerations for the analysis of population coding in motor cortex. *Neural Comput* 6: 29–37, 1994.
- Schwartz AB.** Direct cortical representation of drawing. *Science* 265: 540–542, 1994.
- Schwartz AB.** Cortical neural prosthetics. *Annu Rev Neurosci* 27: 487–507, 2004.
- Schwartz AB, Kettner RE, and Georgopoulos AP.** Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *J Neurosci* 8: 2913–2927, 1988.
- Scott SH, Gribble PL, Graham KM, and Cabel DW.** Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature* 413: 161–165, 2001.
- Smyrnis N, Gourtzelidis P, and Evdokimidis I.** A systematic directional error in 2-D arm movements increases with increasing delay between visual target presentation and movement execution. *Exp Brain Res* 131: 111–120, 2000.
- Steinberg O, Donchin O, Gribova A, Cardosa de Oliveira S, Bergman H, and Vaadia E.** Neuronal populations in primary motor cortex encode bimanual arm movements. *Eur J Neurosci* 15: 1371–1380, 2002.
- Taylor DM, Tillery SI, and Schwartz AB.** Direct cortical control of 3D neuroprosthetic devices. *Science* 296: 1829–1832, 2002.