

Primary Motor Cortex Tuning to Intended Movement Kinematics in Humans with Tetraplegia

Wilson Truccolo,^{1,3} Gerhard M. Friehs,^{2,6} John P. Donoghue,^{1,3,7} and Leigh R. Hochberg^{1,3,4,5}

Departments of ¹Neuroscience and ²Clinical Neurosciences (Neurosurgery), and ³Brain Science Program, Brown University, Providence, Rhode Island 02912, ⁴Center for Restorative and Regenerative Medicine, Rehabilitation Research and Development Service, Department of Veterans Affairs, Veterans Health Administration, Providence, Rhode Island 02908, ⁵Department of Neurology, Massachusetts General Hospital, Brigham and Women's Hospital, and Spaulding Rehabilitation Hospital, Harvard Medical School, Boston, Massachusetts 02114, ⁶Department of Neurosurgery, Rhode Island Hospital, Providence, Rhode Island 02905, and ⁷Cyberkinetics Neurotechnology Systems, Foxborough, Massachusetts 02035

The relationship between spiking activities in motor cortex and movement kinematics has been well studied in neurologically intact nonhuman primates. We examined the relationship between spiking activities in primary motor cortex (M1) and intended movement kinematics (position and velocity) using 96-microelectrode arrays chronically implanted in two humans with tetraplegia. Study participants were asked to perform two different tasks: imagined pursuit tracking of a cursor moving on a computer screen and a “neural cursor center-out” task in which cursor position was controlled by the participant’s neural activity. In the pursuit tracking task, the majority of neurons were significantly tuned: 90% were tuned to velocity and 86% were tuned to position in one participant; 95% and 84%, respectively, in the other. Additionally, velocity and position of the tracked cursor could be decoded from the ensemble of neurons. In the neural cursor center-out task, tuning to direction of the intended target was well captured by a log-linear cosine function. Neural spiking soon after target appearance could be used to classify the intended target with an accuracy of 95% in one participant, and 80% in the other. It was also possible to extract information about the direction of the difference vector between the target position and the instantaneous neural cursor position. Our results indicate that correlations between spiking activity and intended movement velocity and position are present in human M1 after the loss of descending motor pathways, and that M1 spiking activities share many kinematic tuning features whether movement is imagined by humans with tetraplegia, or is performed as shown previously in able-bodied nonhuman primates.

Key words: motor cortex; neuromotor prostheses; paralysis; motor control; spinal cord injury; stroke; quadriplegia

Introduction

Neural activity in the primate primary motor cortex (M1) is known to correlate with movement kinematics. Previous studies in nonhuman primates have described M1 tuning to position, velocity and acceleration of the hand during point-to-point reaching, pursuit tracking, and drawing movements (Georgopoulos et al., 1982; Kettner et al., 1988; Schwartz, 1994, 2004; Fu et al., 1995; Moran and Schwartz et al., 1999; Paninski et al., 2004). Although the interpretation and identification of the mechanisms underlying these tuning functions remain contro-

versial (Kalaska et al., 1989; Todorov, 2000; Scott et al., 2001; Scott, 2004), their statistical characterization has been important in informing and constraining models of M1’s participation in motor control. Studies of the relationships between spiking activities and movement control have been, however, largely limited to the study of nonhuman primates because of the need for an invasive sensor to record neuronal action potentials.

The launch of pilot clinical trials to develop neural interface systems based on signals derived from intracortical microelectrode arrays provides a unique opportunity to evaluate the properties of M1 neurons in humans. We have previously reported neural prosthetic results (Hochberg et al., 2006a) for the first participant in a pilot clinical trial of an intracortically based neural interface system [BrainGate Neural Interface System (NIS); Cyberkinetics Neurotechnology Systems, Foxborough, MA]. Among the many unknowns, it was unclear whether and how the presumed arm area in M1 [the “knob” (Yousry et al., 1997)] would be related to intended or attempted movement kinematics (position and velocity), rather than actual movements, in patients with lost or highly impaired voluntary control of limb motion. This participant could successfully control the position of external devices (e.g., a computer cursor) via a linear filter that related M1 spiking to position. Although the type of implemented tasks, training procedures and online neural decoding

Received March 31, 2007; revised Oct. 26, 2007; accepted Nov. 29, 2007.

This work was supported in part by the Office of Research and Development, Rehabilitation Research and Development Service, Department of Veterans Affairs, and National Institutes of Health—National Institute of Neurological Disorders and Stroke Grants 1K01NS057389, NS-25074 (Javits Award), R01NS369744, and ONR N0014-06-0185. We thank our trial participants for their dedication to this research. We also thank Abe Caplan, Misha Serruya, Maryam Saleh, and other employees of Cyberkinetics Neurotechnology Systems for data collection, device engineering, manufacturing, and clinical trial design and management; Jon Mukand for his role as clinical investigator for participant 1; and John Simeral, Sung-Phil Kim, and Michael Black for discussions.

G.M.F. is a consultant and stock holder of Cyberkinetics Neurotechnology Systems (CYKN). J.P.D. is the Chief Scientific Officer and a director of CYKN; he holds stocks and receives compensation. L.R.H. receives clinical trial support from CYKN.

Correspondence should be addressed to Wilson Truccolo, Department of Neuroscience and Brain Science Program, 60 Olive Street, Providence, RI 02912. E-mail: Wilson_Truccolo@Brown.edu.

DOI:10.1523/JNEUROSCI.4415-07.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/281163-16\$15.00/0

algorithms in the pilot trial were primarily designed with neuro-prosthesis goals in mind, analysis of the collected data provides the opportunity to investigate more general encoding properties in human M1, also furthering neural interface system development and intended movement velocity and position in two trial participants

targets, positioned at 0, 90, 180, and 270°, subtending a visual angle of 4.8° (5 cm) from the center of the monitor; the distance from screen center to target center was 9.5 cm (9.2°). The position of this cursor was obtained by regressing spiking activity onto screen position via the linear filter built in the pursuit tracking task. This on-line neurally controlled cursor is henceforth referred to as the neural cursor (NC). The NC subtended a visual angle of 2.43° (2.5 cm diameter). A trial began after the participant held the NC over the center target for 500 ms. At the start of a trial, one of four peripheral targets appeared on the screen. Participants were instructed to move the NC to the target location and to hold it at that location for 500 ms. A trial terminated after 7 s. Trials were considered successful if the target was acquired in 7 s or less; trials were considered failures if the target was not acquired in 7 s. Regardless of whether or not the target was successfully acquired, because the NC was under constant neural control, the participant needed to return the cursor to the center target (and dwell there for 500 ms) before the next trial would begin. In each session, 20 trials were collected for each of the 4 pseudo-randomly presented radial targets. Videos showing examples of these center-out trials are available in the supplemental material (available at www.jneurosci.org).

Data analysis 1: off-line analyses of tuning functions

Modeling of spiking activity as a function of kinematic covariates during pursuit tracking. We represented the spiking activity (recorded spike times) of each sorted unit as a discrete time neural point process (Truccolo et al., 2005), i.e., as a binary sequence (spike train) obtained by determining whether a spike occurred or not in each consecutive time bin of width $\Delta t = 1$ ms. The value of this sequence at a particular time t is here denoted by $\Delta N_t \in [0,1]$. To investigate the relationship between a neuron's spiking activity and specific kinematic covariates, we modeled, in the log domain, the instantaneous spiking rate as a function of the covariates. Four models were examined: position, velocity, direction, and speed models. For the model of spiking activity as a function of TC position (position model), we used the following:

$$\log \lambda_t^c(x_{t+\tau}, y_{t+\tau}) = \mu + \beta_1 x_{t+\tau} + \beta_2 y_{t+\tau} \quad (1)$$

where λ_t^c is the instantaneous spiking rate function (in spikes/s) at time t for the c th neuron (out of C recorded neurons), x and y are the horizontal and vertical positions, respectively, τ is a time lag, and $[\mu, \beta_1, \beta_2]$ are model parameters to be estimated. A related position model in three dimensions was presented by Kettner et al. (1988). The chosen form of the model of spiking activity as a function of TC velocity (velocity model) was a variation of the model proposed previously by Moran and Schwartz (1999):

$$\log \lambda_t^c(\dot{x}_{t+\tau}, \dot{y}_{t+\tau}) = \mu + \alpha |\dot{\mathbf{x}}_{t+\tau}| \cos(\theta_{t+\tau} - \theta_0) = \mu + \beta_1 \dot{x}_{t+\tau} + \beta_2 \dot{y}_{t+\tau} \quad (2)$$

where $\dot{\mathbf{x}}_t = [\dot{x}_t, \dot{y}_t]^T$ is the velocity vector, θ_t is the movement direction of TC and θ_0 is the preferred movement direction (PD) of the neuron. Although the same letters μ and β appear in different models, they relate to different parameter values depending on the model. This particular velocity model has been successfully used before in approximations of M1 velocity tuning functions in monkeys performing pursuit tracking and center out tasks with their hand (Truccolo et al., 2005). Also, our model is closer in form to the exponential form based on the von Mises distribution, which has been shown to capture well the tuning to direction in spiking activity in monkeys performing center-out tasks (Amirikian and Georgopoulos, 2000). To separately investigate M1 tuning to direction and speed of the TC, we also considered the direction model:

$$\log \lambda_t^c(\theta_{t+\tau}) = \mu + \beta_1 \cos(\theta_{t+\tau}) + \beta_2 \sin(\theta_{t+\tau}) \quad (3)$$

and speed model

$$\log \lambda_t^c(|\dot{\mathbf{x}}_{t+\tau}|) = \mu + \beta |\dot{\mathbf{x}}_{t+\tau}| \quad (4)$$

A unit was considered tuned to a covariate whenever the following two conditions were satisfied (McCullagh and Nelder, 1989): (1) at least one

of the coefficients in the model was statistically significant (Wald confidence intervals, p value < 0.05), and (2) a significant reduction in deviance was obtained from a kinematic model when compared with a model including only the estimated mean spiking rate (χ^2 test, p value < 0.05). The deviance was computed under the Poisson distribution. (For intuition purposes, one can think of the deviance under the Gaussian distribution; in this case the deviance is simply the residual sum of squares.) Optimal time lags τ were determined by finding the time lag that maximized the log-likelihood function of the point process. All of the above models were fitted using standard functions for generalized linear models in Matlab (MathWorks, Natick, MA). More details about the motivation and maximum likelihood fitting of these nonlinear neural point process models can be found in the study by Truccolo et al. (2005).

Assessment of changes in preferred direction across experimental sessions. We examined whether the preferred direction of a tuned unit during the pursuit tracking task changed significantly from one session to the other. A bootstrap resampling procedure was used to assess the statistical significance of preferred direction changes (Chestek et al., 2007). In brief, the bootstrap procedure consisted of the following steps. First, for a chosen unit and session, a distribution of preferred directions was generated by bootstrap sampling with replacement of the observed unit's spiking activity and the related velocity covariate at the previously estimated optimal time lag. One thousand bootstrap samples were used. For each of these bootstrap samples, a velocity model (Eq. 2) was fitted and the corresponding preferred direction was obtained as $\tan^{-1}(\beta_2/\beta_1)$ resolved to the proper quadrant. Second, once a distribution of PDs for a particular unit and session was generated, the circular mean (Zar, 1999) of the distribution was subtracted from each of the sampled PDs. Third, for a chosen unit, we sampled one PD from the session 1 zero mean distribution and another PD from the session 2 zero mean distribution, and computed their absolute angle difference. By repeating this procedure (1000 samples) the distribution of changes in PD corresponding to the null hypothesis (no change in PD) was obtained. This distribution was used to compute the probability that the actually observed change in PD was statistically zero. Units whose PD difference had a p value < 0.05 [false discovery rate (FDR) corrected for multiple comparisons] (Benjamini and Hochberg, 1995; Benjamini and Yekutieli, 2001) were considered to have a significant change in PD between the two sessions. To attenuate the effect of the spiking activity's temporal dependencies in the bootstrap sampling, we dropped all of the spike train samples that occurred within the spike train autocorrelation length.

Assessment of the power of a model to predict spiking activity during pursuit tracking. We assessed how much of a neuron's spiking activity variation was explained by the above described models. Here, this assessment is performed by measuring the power of a model to predict the occurrence or not of single spikes. A measure of predictive power can be derived from receiver operating characteristic (ROC) curve analysis (Fawcett, 2006). ROC curves were generated as follows. A spike is predicted to happen at a particular time t whenever the spike probability conditioned on a particular model, in other words:

$$P(\Delta N_t = 1 | \lambda_t(\cdot)) \approx \lambda_t(\cdot) \Delta t \quad (5)$$

is greater than a specified threshold. A ROC curve can then be generated from the true positive (hit probability) and false positive rates of predicted spikes obtained at different thresholds values. The area under the ROC curve (AUC) gives the probability that, when two different samples (one containing a spike, $\Delta N = 1$, and the other not, $\Delta N = 0$) are randomly drawn from the data, the model will give a higher probability for the sample with a spike. A measure of predictive power can then be derived from the AUC as follows:

$$\text{predictive power} = 2 \times \text{AUC} - 1, \quad (6)$$

with values ranging from 0 to 1. If a model does not improve spike prediction beyond chance level, the AUC will be close to 0.5 (i.e., chance level) and the predictive power will be zero as expected for a model without any predictive power. However, if the model always assigns a higher probability to a sample containing a spike, both the AUC and the predictive power will be one, again as expected for a model with maxi-

imum predictive power. In other words, there exists in this case a specific threshold such that the prediction would be perfect: the model would have then explained all of the variation of the spiking activity in the recorded data.

Relative importance of different kinematic covariates to M1 spiking during pursuit tracking: model independent assessment. Additionally, a model independent comparison of the relative importance of a kinematic covariate to the prediction of single neuron spiking activity was attempted by computing the mutual information between the covariate and the spiking activity. However, given the small sample and low spiking rates, the variance of mutual information estimates was high. An alternative, less powerful but still informative measure based on linear second-order correlation was taken by computing the cross-correlation function between sequences of spike counts (50 ms time bins) and position, velocity, direction and speed at multiple time lags. Cross-correlation functions for position and velocity were computed separately for each coordinate. For direction, a circular-linear correlation measure was used (Zar, 1999). We report, for each neuron and covariate, the maximum absolute value of the cross-correlation function. Confidence intervals (95%) for the null hypothesis were obtained via random permutation tests that used phase randomization of the covariate time series (for details, see supplemental text C, available at www.jneurosci.org as supplemental material).

Modeling of spiking activity as a function of target direction during the neural cursor center-out task. We examined the tuning of M1 neurons to target direction during the initial period after target onset. We used the same direction model (Eq. 3) with the exception that, because the target direction was constant during a trial, the spike trains were transformed into counts over a 400 ms time window. Confidence intervals for the direction tuning functions in the neural cursor center-out task were estimated via bootstrap resampling (5000 bootstrap samples).

Data analysis 2: off-line decoding of instantaneous velocity, position, difference vector direction and intended target

Off-line decoding of instantaneous velocity and position during pursuit tracking. The observed spike trains from all of the tuned neurons together with the corresponding velocity or position models (fitted to training data sets) were used to decode off-line the velocity or position of the TC cursor in test data sets. We adopted a state space formulation to decode the TC velocity from the spike trains in the recorded neural ensemble (Eden et al., 2004a; Truccolo et al., 2005). Following Bayes' rule, the posterior probability distribution for velocity at time $t + \tau$ can be expressed as follows:

$$p(\dot{\mathbf{x}}_{t+\tau} | \Delta N_{0:t}) = \frac{\Pr(\Delta N_t | \Delta N_{0:t-\Delta t}, \dot{\mathbf{x}}_{t+\tau}) p(\dot{\mathbf{x}}_{t+\tau} | \Delta N_{0:t-\Delta t})}{\Pr(\Delta N_t | \Delta N_{0:t-\Delta t})}, \quad (7)$$

where $\Delta N_{0:t} = \{\Delta N_{0:t}^1, \Delta N_{0:t}^2, \dots, \Delta N_{0:t}^C\}$ corresponds to the collection of spike trains from C observed neurons during an interval from time 0 up to time t . A stochastic state space point process filter can be derived as a recursive decoding algorithm (see supplemental text A, available at www.jneurosci.org as supplemental material). In this algorithm, the de-




Figure 1. Spiking activity. **A**, Average spike waveforms of sorted units in the corresponding microelectrode (labeled C1, C2, etc., of 96 recording channels) for S1. Examples are from sessions 1 and 2. In some cases, more than one unit was identified in the same electrode and session (for example, C10). The vertical bar on the left of each of the waveforms corresponds to 100 μV (i.e., -50 to $50 \mu V$). Waveforms span 1.6 ms. **B**, Average spike waveforms for S3. Examples from session 1 are shown. Similar examples were recorded in session 2. **C, D**, Two examples of recorded spike waveforms and the corresponding ISI distribution are shown in **C** and **D** for S1 and S3, respectively. The thick darker curve corresponds to the average waveform in the session. Three hundred randomly chosen waveforms are shown.

coded velocity is taken to be the estimated mode of the posterior distribution (Eq. 7). The initial value for the decoded velocity at time $t = 0$ was sampled from a zero mean Gaussian distribution with variance equal to the variance of the observed TC velocity. Note that the true observed velocities never entered the recursive algorithm; the input to the decoding algorithm consisted only of the observed spike trains. For visualization purposes, we plotted the decoded velocity in polar coordinates, i.e., decoded direction and speed, in the Results section. The position model (Eq. 1) and this stochastic state point process filter were used for off-line position decoding.

Correlation coefficients between true and decoded covariate values, as well as specific decoding examples (see Fig. 4), were computed on test data sets. Training and test data sets were obtained as follows. Each 1 min block was selected once to be the testing data set whereas the other remaining three blocks constituted the training data set. For example, if




Figure 2. Pursuit tracking task: kinematics of the TC. **A**, Examples of the paths of the cursor tracked by S1 and S3. Each plot corresponds to a 1 min tracking block. A technician moved the cursor (TC) to targets (data not shown) that appeared sequentially in random locations on the screen. **B**, Joint distribution of direction (in radians) and magnitude (speed) of the velocity vectors from the pursuit tracking blocks (two sessions) for S1 and S3. Speed is given in visual angle per second (degree/second). Speeds of the presented TC were generally higher for S1 than S3. The two plots on the right show the marginal distributions for direction and speed, respectively. **C**, Correlation functions. Autocorrelation functions, $C(\cdot, \tau)$, for position and velocity in both horizontal (x) and vertical (y) coordinates are shown in the top plots. The bottom plots show the cross-correlation functions, denoted by $C(\cdot, \dot{\cdot}, \tau)$. Blue and red curves correspond to correlation functions for S1 and S3, respectively.

the first three tracking blocks provided 3 min of training data to build the neural encoding models, the test data consisted of the fourth pursuit tracking block, and so on until each block had been selected once for test data. The reported correlation coefficients correspond to the average of the correlation values obtained for each test data set separately.

Off-line decoding of the direction of the instantaneous difference vector during the neural cursor center-out task. The same direction model (Eq. 3) and stochastic state point process filter described above for velocity were used to decode the direction of the instantaneous difference vector between the NC position vector and the target position vector during the neural cursor center-out task. To avoid confusion with the TC direction denoted above by θ , the direction of the difference vector will be denoted by ϕ . For details regarding the definition of the difference vector and its direction, see Figure 10A.

Off-line decoding of target direction during neural cursor center-out task. We were interested in the ability to decode (predict) the location of a newly appearing target in the right, left, down, or up radial position, from the ongoing spiking activity. A naive Bayes classifier was used to decode the target direction. The conditional probabilities $p(N^1, N^2, \dots, N^C)$, where N^c is the spike count in a given time window for the c th neuron and θ the direction (0, 90, 180, or 270°) of the target in a given center-out presentation, were estimated nonparametrically. The representation of the spiking activity in terms of spike counts was chosen for simplicity given that the covariate of interest, target direction, was constant during the trial. Leave-one-out cross-validation was used to assess the classification performance as function of window length and time with respect to target onset. As an example using a 100 ms long window, we computed the classification performance based on spike counts observed in this time window at several different times with respect to target onset, starting at time 0 and moving forward in steps of 25 ms. In other words, for time 0 ms the classification was based on spiking counts observed in the interval $[-100 \text{ ms}, 0 \text{ ms}]$, for time 25 ms the classification was based on $[-75 \text{ ms}, 25 \text{ ms}]$, and so on. The length of the explored windows was varied from 100 to 700 ms in steps of 100 ms.

Results

A total of 29 sorted units (S1) and 121 sorted units (S3) per recording sessions were included in the analyses presented here, resulting in 58 (S1) and 242 (S3) neuron recordings. Figure 1 shows the average extra-cellular recorded action potentials for the sorted neurons from each participant. The mean of the peak-to-peak amplitudes of average waveforms across neurons was $73 \pm 66 \mu\text{V}$ (mean ± 2 SD) for S1 and $173 \pm 312 \mu\text{V}$ for S3. Larger waveform amplitudes were observed in S3, perhaps reflecting closer proximity of electrode tips to larger neurons in deep layers given that longer electrodes were used in this participant. In a few channels, two or three units were separated.

M1 spiking activity and intended kinematics during pursuit tracking task

Before investigating relationships between neural spiking and kinematics, we briefly describe the main statistical properties of the trajectories of the TC. Figure 2A shows examples of paths of the TC during one of the 1 min tracking blocks. These paths were generated by randomly placing, in sequence, a new target on the screen. Once the target was hit by the TC, the target disappeared and a new one appeared on the screen (see Materials and Methods). The distribution of movement direction and speed of the TC in the four 1 min tracking blocks used to fit velocity models are shown in Figure 2B. Motion of the TC in S1 sessions was generally faster than in S3 sessions. Figure 2C shows that both position and velocity had temporal correlations spanning ± 5 s, and tended to show negative correlations at time




Figure 3. M1 velocity tuning during pursuit tracking. **A, B**, Optimal time lag analysis. Examples of velocity tuning functions at different time lags for a neuron (C10a, S1, session 1) are shown in **A**. Velocity is represented in polar coordinates. The color coding represents spiking rates with units of spikes per second. The corresponding log-likelihood under the velocity model as a function of different time lags for these examples is shown in **B**. The stars on the curve indicate that the velocity model was statistically significant for that time lag. This neuron (*Figure legend continues*.)




Figure 4. Off-line decoding of TC velocity during pursuit tracking. Spike trains of velocity 2-tuned neurons were used, together with the fitted velocity models (tuning functions as shown in Fig. 3A), to generate a neural ensemble prediction of the velocity of the tracked cursor. In the examples shown, the models were fitted first to data from three 1 min pursuit tracking blocks (training data) and then used to decode velocity on test data consisting of a different 1 min block. A stochastic state point process filter was used for decoding (Eq. 7) (Eqs. S1–S7 in the supplemental material, available at www.jneurosci.org). At time 0, a random velocity vector was assigned to the initial state and then at every millisecond, during the 60 s of tracking, the TC velocity vector was updated according to the stochastic state point process filter output. For plotting purposes, the decoded velocity vector is shown in polar coordinates, with speed given in visual angle (degree) per second. **A, C**, The neural decoded movement direction (red) and the actual direction (black) are shown for S1 (**A**) and S3 (**C**). The dashed portion of the black curve represents segments where we considered the cursor to be moving too slowly to have a clearly defined or perceivable movement direction (speed, $<2.5^\circ/\text{s}$). The decoded speed of the tracked cursor is shown in **B** for S1 and **D** for S3. For computational simplicity, a single time lag was used for all of the neurons in the ensemble.

lags that differed for S1 and S3. Further, position and velocity in the same coordinate were also positively and negatively correlated depending on the time lag. Although the pursuit task was designed to minimize all of these temporal correlations in and among kinematic covariates, the observed correlations are likely to have resulted from edge effects imposed by the workspace and

←

(Figure legend continued.) seems to be maximally tuned to velocity at $\tau = 0.5$ s (i.e., to future values of the TC's velocity), suggesting anticipation effects. **C**, Distribution of the optimal time lags over the population of tuned neurons in S1. **D**, Distribution of the optimal time lags over the population of tuned neurons in S3. Positive time lags suggest anticipation effects. Such effects could happen because a given target was present on the screen until it was acquired by the TC. **E, F**, Summary of PDs and spike rate modulation for the population of velocity tuned neurons in S1 (**E**) and S3 (**F**). The direction of each vector gives the preferred direction (in degrees) and the length of the vector denotes the tuning depth or spike rate modulation caused by variations in velocity. Distribution for the same preferred directions and tuning depths over the tuned population in S1 and S3 are shown in **G** and **H**, respectively. Summaries and histograms include neurons from the two recording sessions.

properties of natural motion (the TC was controlled by a technician).

Tuning to velocity during pursuit tracking

M1 neurons were tuned to TC velocity in both participants. We examined the relationship between spiking activity and past, simultaneous and future velocities of the TC (Fig. 3A). The fact that the TC motion was relatively slow and that the TC's target was present on the screen during the entire TC motion could have induced both delay and anticipation effects (see below). For this reason, we fitted the velocity model (Eq. 2) for a wide range of different time lags: $\tau \in [-4, -3.9, \dots, 0, \dots, 3.9, 4]$ s. For each neuron, we found the optimal time lag (OTL), i.e., the time lag at which the log-likelihood under the velocity model achieved its maximum value, and checked whether the model satisfied the two significance criteria (see Materials and Methods). A neuron was considered tuned to velocity if both significance criteria were met. For further analyses, we used only the velocity tuning functions at the optimal time lag. An example of optimal time lag analysis based on the log-likelihood function is shown in Figure 3B. The distributions for the optimal time lag over the population of tuned neurons are shown in Figure 3, C and D. In S1, optimal time lags over the population were clustered ~ 100 –200 ms. This could be a reflection of anticipation effects, given the structure of the task: the participant knew the target position at the start of the TC's movement and this movement toward the target was almost rectilinear. Optimal time lags for neurons in S3 were bimodally distributed over the interval, with clustering around positive and negative time lags (see Discussion). A summary of preferred direction and tuning depth at the OTLs over all of the tuned neurons is shown in Figure 3, G and

H. The tuning depth was defined as the total rate modulation in the tuning function accounted for by the covariate, in this case velocity. Of the studied neurons across both sessions, 90% and 95% were tuned to velocity in S1 and S3, respectively. The velocity model can be seen as the introduction of speed effects to the simpler direction model (Eq. 3). For this reason we assessed whether the reduction in deviance obtained by the velocity model with respect to the direction model was statistically significant (second criterion for significance) (see Materials and Methods). For all of the recorded units, this reduction in deviance was statistically significant (χ^2 test, $p < 0.05$), indicating that the speed component in the velocity model added to the explanation of spiking activity. A similar analysis comparing the direction and the speed models (Eq. 4) revealed also that, for all of the tuned neurons, the direction model achieved a larger reduction in deviance (χ^2 test, $p < 0.05$), suggesting that the direction model explained more of the spiking activity than the speed model.

Additional evidence supporting velocity tuning in this pursuit

tracking task was provided by the off-line neural decoding of TC velocity in cross-validated blocks. We set the time lag of the velocity covariate for all of the neurons to the same value given by the mean of the optimal time lag over the population. Obviously, this choice of a single time lag might not have been optimal, but it was taken, nevertheless, for simplifying the computations. Correlation coefficients between true and decoded velocity computed over all of the four 1 min block test datasets (see Materials and Methods) were relatively small (0.28 and 0.35 in S1, 0.37 and 0.35 in S3, for velocities in x and y coordinates, respectively). Despite these relatively small correlations, reasonably accurate off-line decoding of the direction of the velocity vector was achieved (Fig. 4A, C). Figure 4 shows decoding examples for both participants.

We also examined whether the preferred direction of a tuned unit changed significantly from one session to the other. The two recording sessions for each participant were separated by four and one day(s) in S1 and S3, respectively. As in Chestek et al. (2007), we visually inspected waveforms and interspike distributions to determine whether a recorded unit in a given electrode seemed similar or not in both sessions. Only the recorded units that were tuned to direction in both sessions were included in the analysis. In the case that more than one unit was sorted from a single electrode, only the largest unit entered the analysis. A bootstrap sampling procedure was used to assess the statistical significance of changes in preferred directions (see Materials and Methods). In S1, 5 of 12 examined units (41%) showed a significant change in PD ($p < 0.05$, FDR correct for multiple comparisons) (Benjamini and Hochberg, 1995). Of these five units, changes of 38.5, 41.5, 59.6, 114.0 and 130.9° were observed. In S3, 33 of 76 (43%) units showed a significant change in PD. The changes in PD ranged from 30.9 to 176.2°, with a median of 125°. Possible interpretations of these results are given in the Discussion.

Tuning to position during pursuit tracking

Our previous work (Hochberg et al., 2006a) demonstrated the availability of position information in the recorded neuronal ensembles in S1, by showing that he was able to successfully control a computer cursor and other devices via a position linear filter decoder. Here, we report more extensive analyses for him and an additional participant (S3) with a different neurological condition (see Materials and Methods). As was done for the analysis of velocity tuning, we fitted the position model (Eq. 1) for a range of different time lags $\tau \in [-4, -3.9, \dots, 0, \dots, 3.9, 4]$ s, and we report analyses based on the position tuning function at the optimal time lag. Figure 5A shows an example of an empirically




Figure 5. M1 position tuning during pursuit tracking. **A, B**, An example of a position-tuned neuron (C18a, S1, session 1) is shown. The empirical mean spike rate conditioned on x, y position of the tracked cursor at $\tau = 1$ s is shown in **A**. To compute the conditional mean rate, we used spike counts in 100 ms bins. The conditional mean rate is given in spikes per second (pseudocolor coding). Blank bins refer to states in the position space where we considered there were not enough samples to estimate the conditional mean (i.e., bins with < 10 samples during the four 1 min pursuit tracking blocks). Position is given in centimeters. The corresponding position tuning function based on position model (Eq. 1) fitted to the actual point process data (1 ms time resolution spike trains, not the conditional mean) is shown in **B**. Summaries of the PD in position space and of the tuning depth modulation over the population of tuned neurons are shown in **C** for S1 and **E** for S3. These summaries include neurons from the two recording sessions. The corresponding distributions of optimal time lags (OTL) are shown in **D** for S1 and **F** for S3. Optimal time lags were determined based on the same log-likelihood analysis used for the velocity model (see Fig. 3A, B).

estimated mean spiking rate conditioned on the TC position, and the corresponding fitted model is shown in Figure 5B. Figure 5, C and E, shows the distribution of PD in the position space for the population of tuned units. This position PD corresponds to the direction, in the two-dimensional position space, in which the spiking rate changes the most with changes in position. The distribution of optimal time lags for both participants is shown in Figure 5, D and F. Optimal time lags for S1 and S3 tended to cluster around $\tau = 1$ s, with wider spread for S3. Summaries for both participants include tuned neurons from both sessions 1 and 2. Of the studied units, 86% and 84% were tuned to position in S1 and S3, respectively. Position decoding has been previously examined (Hochberg et al., 2006a); here we add that off-line position decoding using the point process filter algorithm (see Materials and Methods) (supplemental text A, available at




Figure 6. Velocity and position spike prediction power during pursuit tracking: ROC analysis. **A**, ROC curve example for one neuron. The curve was constructed by computing the true positive rate (hit probability) and false positive rate in spike prediction based on a specific model (position model in this example) and different threshold values. The probability of a spike at a given discrete time is derived from the instantaneous spiking intensity given a model. If this probability is higher than a specified threshold, a spike is predicted. The shaded region corresponds to the AUC. **B**, The position and velocity models' power to predict a spike are shown for S1. Each point value corresponds to predictive power for a particular tuned neuron from either session 1 or 2. The diagonal line corresponds to equal prediction power for position and velocity. The prediction power was defined as $2 \times \text{AUC} - 1$, ranging from 0 to 1 (maximum power). For details, see Materials and Methods (Eq. 6). **C**, Same as in **B**, but for S3. Tuned neurons from session 1 and 2 were included. Over all, differences between predictive power of position and velocity models were small in this pursuit tracking task.

www.jneurosci.org as supplemental material) yielded correlation coefficients between true (TC) and decoded position of 0.41 and 0.29 for S1, and 0.42 and 0.45 for S3 in x and y coordinates, respectively.

Comparison between velocity and position models

Because differences in decoding performance can also be a reflection of the nature and assumptions in the decoding algorithm itself, we used a different approach to compare position and velocity models. We compared the velocity and position models in terms of their power to predict the occurrence of single spikes (1 ms time resolution). As stated in the Materials and Methods section, the predictive power (Eq. 6) is a value between 0 (minimum predictive power) and 1 (maximum predictive power) and it is derived from the area under the ROC curve. A predictive power of 1 would mean that the model explained all of the variation in the spike train.

Position and velocity models were fitted to data from the four 1 min pursuit tracking blocks. ROC curves were computed on the same data. Because the number of parameters in the velocity and position model was the same and this number was very small

compared with the number of samples, data overfitting was not an issue and a cross-validation scheme was not necessary. Figure 6A provides an example of ROC curve analysis for a single neuron. Summaries of the predictive power of velocity and position over the population of recorded neurons from both participants are shown in Figure 6, B and C. Differences between the predictive power of position and velocity were minor and did not motivate further detailed statistical testing.

We also performed a model-independent comparison of the strength of the relation between different kinematic covariates (position, velocity, direction, and speed) and M1 spiking activity. Low spike rates and small sample data prevented us from applying standard mutual information analyses. Instead, we adopted a less powerful but still informative measure: cross-correlation functions between the 1 min long spike trains (transformed into 50 ms counts) and the corresponding time series of the kinematic covariate. These functions were computed at multiple lags and averaged across the four pursuit tracking blocks. Cross-correlations for position and velocity were computed separately for x and y coordinates, and the circular-linear correlation (Zar, 1999) was computed for the TC movement direction. Neurons whose cross-correlation function values (at any time lag) achieved statistical significance were included in the analyses. 95% confidence intervals were obtained from distributions generated via phase randomization tests (see Materials and Methods) (supplemental text C, available at www.jneurosci.org as supplemental material). Figure 7 summarizes this analysis for both participants. A sign test performed separately for each session and participant revealed no mean differences between correlation values for position and velocity in each coordinate over the recorded neuron population (p value > 0.05). Highly significant differences were obtained, however, when comparing direction and speed (sign test, p value < 0.0001) in each session and participant. Correlation values were higher for direction in 93 and 96% of the neurons for S1 (sessions 1 and 2, respectively), and 100% and 96% for S3. In summary, both predictive power and model-independent comparisons did not reveal major differences in the relative importance of velocity and position for explanation of M1 spiking activity in this task. Significant differences were detected only for the relative importance of direction and speed.

participant revealed no mean differences between correlation values for position and velocity in each coordinate over the recorded neuron population (p value > 0.05). Highly significant differences were obtained, however, when comparing direction and speed (sign test, p value < 0.0001) in each session and participant. Correlation values were higher for direction in 93 and 96% of the neurons for S1 (sessions 1 and 2, respectively), and 100% and 96% for S3. In summary, both predictive power and model-independent comparisons did not reveal major differences in the relative importance of velocity and position for explanation of M1 spiking activity in this task. Significant differences were detected only for the relative importance of direction and speed.

Relation of M1 neural activity to intended kinematics in a neural cursor center-out task

As shown above, during the imagined pursuit tracking task, M1 neural activity in both participants with tetraplegia was tuned to both TC position and velocity. We were also interested in further investigating neuronal spiking activity when the participant was using a neurally controlled cursor to acquire a series of targets.




Figure 7. M1 spiking activity and kinematics during pursuit tracking: correlation analysis. **A**, **B**, A model independent assessment of the relative importance of different kinematic covariates [position (pos), velocity (vel), speed, and direction] to M1 spiking activity was obtained by computing cross-correlation functions between each covariate and the sequence of spike counts for each neuron (see Materials and Methods for details). The plots summarize the distribution of the maximum absolute cross-correlation function values between spike count sequences and the time series of a specific kinematic covariate over the recorded population of neurons, including data from both sessions 1 and 2. The bottom and top boundary of each box corresponds to the 2.5 and 97.5 percentiles, respectively; the interior line shows the median and the crosses correspond to the computed maximum absolute correlation values for each neuron. A sign test, applied to data from each session separately, revealed no mean differences ($p > 0.05$) between cross-correlation values for position and velocity in each coordinate. Highly significant differences were detected when comparing correlation means for direction and speed (sign test, $p < 0.0001$) in both S1 and S3.

Spiking activity during a neural cursor control task

We start by investigating the time course of spiking activities and the properties of their tuning functions to intended target direction during the neural cursor center-out task. Target onset was preceded by a holding phase in which the participant kept the NC at the center of the workspace for at least 500 ms. When this hold period was satisfied, a target appeared at one of four radial positions (right, up, left, and down) and the participant proceeded, following prior instruction, to move the NC toward the target. Most of the analyses described below focus on the initial period [0, 1.5 s] after target onset. Neurons recorded in both participants showed clear modulation of their spiking rates after target onset. Figure 8 shows examples of perievent time histograms (PETHs), centered at target appearance (time 0). Modulation in the PETH tended to occur around 300 ms for S1. In S3, the modulation onset showed a broader range. In some cases, the temporal profile of spike rate modulation varied according to the target.

M1 tuning to intended movement direction during neural cursor center-out task

We studied the above described dependence of spiking rates on different targets in terms of tuning to target direction. It is not obviously expected that M1 spiking activity should be related to intended movement direction in this center-out task. That is because the participants used a position based control interface to perform this task (for more details, see Materials and Methods): the position of the neural cursor was given by the output of a position filter which explicitly extracted position, not velocity, information from the recorded M1 neuronal ensemble via a linear mapping. As exemplified in Figure 9, **A** and **C**, tuning to target direction was well captured by the direction model (Eq. 3, Materials and Methods). The percentage of neurons that showed statistically significant tuning to target direction in specified time windows after target onset was 66 and 64% for S1 and S3, respectively. Figure 9, **B** and **D**, summarizes the preferred directions in the tuned population.

Because there were only four radial targets, intended target (possibly also location) and intended direction could be confounded during the initial phase (0–1.5 s) of “reaching.” Therefore, the detected direction tuning could actually be a reflection of tuning to intended target identity or position. For this reason, we extended our analyses of direction tuning to later phases of the center-out reaching. As reported previously, the on-line position linear filter decoder used to animate the neural cursor was far from optimal. Although the participants were able to navigate the cursor from the center to the target and back, often in a few seconds, the cursor motion commonly strayed from a straight or a simple curvilinear path linking the center and the target. As a result, there were many more than just four directions in which to move the neural cursor toward the target during the reaching phase. This provided us, therefore, with a much wider range of intended directions for the examination of M1 direction tuning. This examination was performed as follows.

Given target and instantaneous NC position vectors during a center-out trial, the direction of the difference vector between these two vectors can provide a reasonable approximation to the intended movement direction in which to move the NC (for a schematic description of these vectors, see Fig. 10A). In this scenario, the direction of the difference vector would then approximate the direction of the intended velocity vector at each moment in time during the center-out task. We fitted direction models (Eq. 3, with the covariate now representing the direction of the difference vector) to the neural point process data (1 ms time bins) and used these models together with the stochastic state point process filter (see Materials and Methods) to decode the direction of the difference vector. The instantaneous angle error (i.e., angle difference between true and decoded direction of the difference vector) obtained by using this decoding algorithm was relatively small. It can be seen in Figure 10B that the distribution of the angle error concentrated below 45°, demonstrating that decoding of this direction clearly departed from chance levels. Correlation coefficients (circular correlation) (Zar 1999) between the true and decoded difference vector’s direction were $r = 0.21$ and 0.4 for S1 and S3, respectively. An assessment of the power of this covariate to predict spike activity is provided in Figure 10C, where we have used the same ROC approach as done in the analysis of velocity and position during the pursuit tracking task. Both the decoding and predictive power analyses indicated, therefore, that information about the direction of the difference vector, or intended movement direction, was available in the M1 spiking activity during the neural cursor center-out task. Issues




Figure 8. *A, B*, Target onset PETHs during neural cursor center-out task. Examples of PETHs for nine different neurons from S1 (*A*) and S3 (*B*) from either session 1 or 2 are shown. Each row in the figure corresponds to a particular neuron (e.g., C18a), and the four columns correspond to the four target directions (i.e., right, up, left, and down). The PETH for the top neuron includes the corresponding raster plot. Target onset is at time 0. A time bin width of 50 ms was used to compute the spike counts and the corresponding spike rate. The red curve gives the smoothed histogram. Each PETH was computed from 20 center-out trials.

regarding neural adaptation, including changes in direction tuning properties attributable to the type of the task (e.g., pursuit tracking vs neural cursor center-out), are addressed in the Discussion.

Off-line decoding of intended target during neural cursor center-out task

We examined how well intended target could be decoded from M1 activity during the initial phases of the reaching center-out trials. We assessed the decoding performance of intended target during the interval [0, 1.5 s] of the center-out task for trials that lasted at least 2 s. Figure 11, *A* and *B*, shows the distributions of the times taken by the participants to move the neural cursor from the center of the screen to the radial target and to hold at the target for 0.5 s. By limiting our off-line analysis to the first 1.5 s, we ensured that the neural spikes entering the decoding analyses had occurred before target acquisition in all of the examined trials. Figure 11, *C* and *D*, shows the temporal evolution of the classification performance based on spike counts in a given moving time window (see Materials and Methods). A naive Bayes classifier was used for classification. Models were fitted on training data and target classification was performed on test data using a leave-one-out cross-validation scheme. In other words, a trial was selected as the test trial, and all of the remaining

others were used to fit the model. We varied the window lengths from 100 to 700 ms, in 100 ms steps. For each time and window length, the conditional probability of a neuron's spike count in the window was estimated nonparametrically. In S1, classification performance departed from chance classification (0.25) at ~200 ms after target onset. In S3, this departure happened later at ~400 ms. Maximum performance levels and peak time depended on the length of the time window used for the spike counts. The maximum correct classification performance and performance peak time increased monotonically with the window length (Fig. 11 *E–H*). For each participant, the time course in classification performance was about the same in the two analyzed sessions.

Discussion

In two humans, M1 tuning to position and velocity of intended movement shared many of the features previously observed in able-bodied nonhuman primates. In particular, we have demonstrated that M1 neural spiking in these two humans is strongly tuned to intended movement direction in both pursuit tracking and point-to-point (center-out) 'reaching' tasks. It is notable that these features were observed years after severe damage to descending motor pathways alone (S3) or combined damage to




Figure 9. M1 tuning to target direction during neural cursor center-out task. Examples of direction tuning functions are shown in **A** and **C** for S1 and S3, respectively. Direction tuning functions were fitted to single neuron spike count data collected over a 400 ms time window covering the time interval 200–600 ms after target onset for S1 and either 150–550 ms or 400–800 ms, depending on the neuron, for S3. The four red stars represent the observed mean spiking rate conditioned on the four intended movement directions. Models were fitted to single trial data, not to the computed conditional mean rates. The fitted model (black curve) allowed interpolation to directions in $[0, 360]$ degrees. The red curves give the 95% confidence interval estimated via bootstrap resampling (5000 samples). Different neurons from either session 1 or 2 were used in these examples. The polar plots in **B** and **D** summarize the estimated preferred directions for the tuned neurons in both sessions for S1 and S3, respectively. The length of the vector represents the depth in modulation in the tuning function.

both descending and ascending pathways (S1), although it cannot be ruled out that these features emerged as a result of the neurologic injury or adaptation to specific task constraints (e.g., “massless” cursor and pure kinematics control) in the given experimental tasks. Although the analyses do not point to a particular mechanism underlying the nature of M1 representations and their role in motor control in the intact nervous system, the existence of these tuning features in humans with tetraplegia has implications for both models of motor control and for the development of neuroprostheses for persons with paralysis or limb loss. Aspects of these results warrant further elaboration.

Tuning to kinematics during pursuit tracking

The ability to decode movement direction continuously during 1 min long pursuit tracking blocks was notable: in contrast to previous studies in able-bodied monkeys performing pursuit tracking with actual hand movements (Paninski et al., 2004; Truccolo et al., 2005), here the participants were unable to move their hands or arms. Although we propose that the underlying source of tracking-related

velocity signals in M1 was an ‘imagined’ action or intention to move, a few alternative explanations should be also considered. First, although the two participants sustained damage to descending motor pathways, voluntary motion of the head, neck and eyes remained. These motions could have accompanied the imagined tracking. Our previous analysis (Hochberg et al. 2006a) showed, however, that head motion and trajectories of the neural cursor were poorly correlated, indicating that this remaining motion played a minor role, if any. Furthermore, open loop experiments where the participant (S1) was asked to simply imagine different arm and hand movements yielded diverse and imagined-movement-specific modulation in M1 neuronal activity that could not have resulted from visuomotor cues [Hochberg et al. (2006a), their Fig. 3 and supplemental Fig. 1]. Future studies involving participants with advanced ALS (i.e., persons with almost no remaining head, trunk, or appendicular movements) (Hochberg et al., 2006b) may help to answer this question directly. Other alternative explanations would be that signals about the direction of the tracked cursor could have been made available in M1 via indirect visual inputs originating, for ex-




Figure 10. Decoding of difference vector direction during neural cursor center-out task. **A**, The difference vector is illustrated. The red and blue circles represent the target and the neural cursor, respectively, and the red and blue arrows represent the corresponding target and instantaneous neural cursor position vectors. The green arrow represents the difference vector between the target and neural cursor position vectors, and $\phi(t)$ denotes the instantaneous angle or direction of this difference vector. The dashed curve represents the ongoing trajectory of the neural cursor during a center-out trial. **B**, The normalized histogram of the error between the true direction and the decoded direction, $|\phi(t) - \hat{\phi}(t)|$ (in degrees), obtained by decoding the difference vector direction during the two center-out sessions for both participants. The horizontal line represents the uniform distribution expected from a “random” decoder. The fact that error concentrated at angle errors below 45° suggests that information about the direction of the difference vector was available in the M1 neural ensemble during the neural cursor center-out task. **C**, Estimated predictive power of the difference vector direction based on the AUC measure (for details, see Materials and Methods, Eq. 6, Fig. 6). The crosses correspond to the estimated predictive power values for each neuron; each box corresponds to the 95% data interval (over the neuronal population), and the interior line represents the median predictive power value.

ample, from area MT, or that M1 tuning reflected gaze direction effects (Baker et al., 1999). Additional experimental controls, beyond those adopted under the priorities and constraints of this pilot clinical trial, would have been required to sort out these alternative explanations. Although the precise origin of these kinematic signals is unclear, our findings have nevertheless demonstrated their robust persistence in M1 neuronal spiking activities in humans with tetraplegia.

Optimal time lags for the tracked cursor velocity clustered around 400 ms for S1, suggesting anticipatory effects. These effects might have easily taken place in this task because the target was present at all times during cursor motion, and that the cursor’s motion toward the target was almost rectilinear. Similar clustering was observed for position in S1. A different pattern was seen in S3 for whom velocity and position optimal time lags were more spread out and a bimodal distribution with clusters around -1 and 2 s was observed for velocity time lags. For comparison, note that in able-bodied monkeys actually performing a hand tracking task, optimal time lags for hand velocity tend to cluster around 100 ms, as expected for neurons controlling the actual hand movement, whereas for hand position, time lags tend to distribute more uniformly in a broad interval from $(-1, 2)$ s (Paninski et al.,

2004) (Fig. 11D,E). It should be also noted that the observed distribution of optimal time lags might reflect more the properties of the tasks, different cognitive strategies used by the participants, or their neurological condition and associated compensations, than general M1 functional properties.

The direction of the cursor’s velocity vector during the tracking task contributed more to the explanation of neuronal spiking activity than did speed. A similar observation has been made in M1 spiking activity in monkeys manually executing pursuit tracking (Truccolo et al., 2005) and reaching tasks (Moran and Schwartz, 1999). It should be noted, however, that this weaker contribution of speed in the human data reported here could also reflect the small range of speeds of the tracked cursor in the pursuit tracking task. Our analyses suggest also a balanced relative contribution of position and velocity to the explanation of M1 spiking activity. However, because these analyses relied on model based (ROC) and linear correlation measures, the issue of the relative explanatory or predictive power of these two covariates needs further exploration.

Tuning to intended movement direction and target during neural cursor center-out task

Because there were only four radial targets, their respective identity and direction were confounded at the beginning of each target presentation. Therefore, the observed tuning to direction (Fig. 9) could have resulted from tuning to the intended target, or conversely, the classification of intended target (Fig. 11) could have been based on intended direction signals. Future experiments that prevent such confounding (e.g., by using two layers of radial targets, where a straight line between the center and outer target would intercept an inner target) will be required to fully resolve this ambiguity. Nevertheless, the decoding analysis of the difference vector’s direction provided additional evidence for the existence of information about intended movement direction in M1 spiking during this (position control) neural cursor center-out task. The direction of the difference vector, computed throughout the center-out trials, was largely independent of target identity or target position in our data: the neural cursor did not follow, in general, a straight path from the center to the target; the paths were usually very wiggly. In this way, the difference vector’s direction provided an approximation to the intended movement direction at any instant during the center-out task. (We say an approximation because the intended direction might have also depended on other factors such as velocity and acceleration of the neural cursor.) We also note that tuning to intended direction should not be expected a priori in this task, because participants controlled the neural cursor via an on-line linear position filter. This filter required explicit information about intended position, not direction or velocity, from M1 spiking activity. The fact that many successful trials lasted longer than the length of the on-line

position linear filter (i.e., 1.4 s) argues against a filter that was actually outputting position by integrating velocity information available in the neuronal ensemble activity.

Differences between participants

Although most of our results regarding tuning to kinematics were qualitatively similar between the two participants, there were a few important differences. As mentioned, the distribution of optimal time lags contrasted between the two participants. Additionally, off-line neural decoding of direction during pursuit tracking and of the intended target during the center-out task was better in S1 than in S3. Also, S3 required longer times for target acquisition using the on-line position linear filter decoder. This occurred although this participant's recordings had many more significantly tuned neurons than S1. We can only speculate about a series of factors that may have contributed to these differences: recording from different cortical layers (electrodes were 1.0 mm long in S1 and 1.5 mm in S3), possible differences in the cognitive and control strategies adopted by the participants when handling the position linear filter interface and their different medical conditions, i.e., spinal cord transection in S1 with complete deafference and deafference of the limb, and pontine stroke with intact sensory pathways and intermittent bilateral upper extremity flexor spasms in S3.

Stability of recorded neuronal signals: changes in recorded neuronal populations and changes in tuning properties

Neurons were recorded stably across single sessions: tracking spike waveforms continuously across a session (1–2 h) revealed that there was no observed change in the number or identity of recorded neurons. However, because recordings occurred for only a few hours each day, and thus continuous tracking of waveforms across sessions could not be performed, we do not know whether waveforms discriminated from same electrode across days were from the same or a different neuron. Additionally, we also observed that, in some cases, a unit well isolated in one session could not be discriminated in the other session. (This was one of the motivations for building new decoding filters at the beginning of each session.) We believe that small vertical motions of the microelectrode array, on the order of a few tens of micrometers with respect to the cortical tissue, could be a source of these fluctuations across sessions. Further, there has been interest in assessing the stability of single neuron tuning properties over hours or




Figure 11. Off-line classification of intended target during neural cursor center-out task. *A, B*, The distribution of times required for S1 and S3, respectively, to acquire the target with the neural cursor in all of the successful trials (both sessions). *C* and *D* plot the proportion of correct classification as a function of time, based on counts computed from windows of different lengths (100, 500, and 700 ms) for S1 and S3, respectively. The intended target for each of the center-out trials was decoded based on spike counts. Each window was shifted by 25 ms to obtain the time course of the classification performance. Black dots represent the correct classification rate computed separately for each of the two sessions. The red curves represent a smooth fit to these points. The black line at 0.25 represents the chance level of correct classification given four targets. Time 0 corresponds to target onset time. Only successful trials that lasted >2 s from target appearance to target acquisition were used in this analysis (for details, see Materials and Methods). In *E* and *G*, the maximum achieved correct classification and the corresponding peak time based on all of the explored window lengths (100, ..., 700 ms) are shown for S1. *F, H*, Similarly for S3. The two points for each window length correspond to the values for sessions 1 and 2.

days. Several previous studies have reported changes in these properties and related them to learning (Li et al., 2001; Padoa-Schioppa et al., 2004) or to background intrinsic fluctuations (drift) of tuning functions on a hypothesized solution manifold (Rokni et al. 2007). Others have found little or no significant variability in tuning (Chestek et al., 2007). We attempted a simple assessment of the stability of tuning properties in our data by examining changes in preferred direction across the two experimental sessions. Our results seem to indicate that whereas ~60% of the examined units showed no significant change in preferred direction, ~40% of the examined units in both participants showed a statistically significant change. These results should be interpreted with caution, however. First, our identification of unit identity across different sessions was based on visual inspection of spike waveforms and interspike distributions. A more rigorous identification would require, as mentioned above, the tracking of waveforms continuously over the time period in between sessions (Santhanam et al., 2007). We hope to systematically investigate tuning function stability and the effects of learning with more controlled tasks and larger data sets in future studies. Second, it should also be noted that nonstationarity of tuning properties, as well as changes in the recorded neuronal population, can be properly handled in neuroprosthetic applications via intermittent recalibration of decoding filters (e.g., as performed here with filter rebuilding at the beginning of a session) or via continuous tracking of the neuronal population and their tuning properties (Eden et al., 2004a,b; Srinivasan et al., 2007). In addition, as pointed out by Rokni et al. (2007), one can also conjecture that an optimal solution manifold might similarly emerge for the neuronal population recorded at the prosthetic interface. If this conjecture is correct, spontaneous background drift of tuning functions on this optimal solution manifold would not significantly affect decoded motor output.

Implications for neural interface systems

The results from these off-line encoding-decoding analyses suggest that intended velocity could be used as a control signal for direct corticomotor prosthetic devices by persons with paralysis. A preliminary evaluation in one participant showed that closed-loop velocity based control resulted in significantly improved performance compared with position based control (Kim et al., 2007). Additionally, the fact that both intended position and velocity signals are available in M1 activity in these two humans with tetraplegia suggests that the combined use of these two variables in decoding algorithms (Kemere et al., 2004; Srinivasan et al., 2006, 2007; Wang et al., 2007; Yu et al., 2007) may enhance performance.

References

- Amirikian B, Georgopoulos AP (2000) Directional tuning profiles of motor cortical neurons. *Neurosci Res* 36:73–79.
- Baker JT, Donoghue JP, Sanes JN (1999) Gaze direction modulates finger movement activation patterns in human cerebral cortex. *J Neurosci* 19:10044–10052.
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J Royal Stat Soc B* 57:289–300.
- Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. *Ann of Statistics* 29:1165–1188.
- Chestek CA, Batista AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* 27:10742–10750.
- Eden UT, Frank LM, Barbieri R, Solo V, Brown EN (2004a) Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Comput* 16:971–998.
- Eden UT, Truccolo W, Fellows MR, Donoghue JP, Brown EN (2004b) Reconstruction of hand movement trajectories from a dynamic ensemble of spiking motor cortical neurons. *Proc IEEE Eng Med Biol Soc Annu Conf* 2:4017–4020.
- Fawcett T (2006) An introduction to ROC analysis. *Pattern Recognition Letters* 27:861–874.
- Fu QG, Flament D, Coltz JD, Ebner TJ (1995) Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J Neurophysiol* 73:836–854.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and neuron discharge in primate motor cortex. *J Neurosci* 2:1527–1537.
- Guillory KS, Normann RA (1999) A 100-channel system for real time detection and storage of extracellular spike waveforms. *J Neurosci Methods* 91:21–29.
- Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006a) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:164–171.
- Hochberg LR, Friehs GM, Brown Jr RH, Cudkowicz ME, Krivickas LS, Donoghue JP (2006b) Voluntary modulation of motor cortical activity by a person with amyotrophic lateral sclerosis: initial BrainGate experience. *Soc Neurosci Abstr* 32:13.1.
- Kalaska JF, Cohen DAD, Hyde ML, Prud'homme M (1989) A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *J Neurosci* 9:2080–2102.
- Kemere C, Shenoy KV, Meng TH (2004) Model-based neural decoding of reaching movements: a maximum likelihood approach. *IEEE Trans Biomed Eng* 51:925–932.
- Kettner RE, Schwartz AB, Georgopoulos (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.
- Kim SP, Simeral JD, Hochberg LR, Donoghue JP, Friehs GM, Black MJ (2007) Multi-state decoding of point-and-click control signals from motor cortical activity in a human with tetraplegia. In: *Proceedings of the Third International IEEE EMBS Conference on Neural Engineering*, pp 486–489. Kohala Coast, HI: IEEE Engineering in Medicine and Biology Society.
- Li CS, Padoa-Schioppa C, Bizzi E (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 30:593–607.
- McCullagh P, Nelder JA (1989) *Generalized linear models*, Ed 2. Boca Raton, FL: Chapman and Hall/CRC.
- Moran DW, Schwartz AB (1999) Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82:2676–2692.
- Padoa-Schioppa C, Li CS, Bizzi E (2004) Neuronal activity in the supplementary motor area of monkeys adapting to a new dynamic environment. *J Neurophysiol* 91:449–473.
- Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP (2004) Spatiotemporal tuning of motor neurons for hand position and velocity. *J Neurophysiol* 91:515–532.
- Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representations. *Neuron* 54:653–666.
- Rousche PJ, Normann RA (1992) A method for pneumatically inserting an array of penetrating electrodes into cortical tissue. *Ann Biomed Eng* 20:413–422.
- Santhanam G, Linderman MD, Gilja V, Afshar A, Ryu SI, Meng TH, Shenoy KV (2007) HermesB: a continuous neural recording system for freely behaving primates. *IEEE Trans Biomed Eng* 54:2037–2050.
- Schwartz AB (1994) Direct cortical representation of drawing. *Science* 265:540–542.
- Schwartz AB (2004) Differential representation of perception and action in the frontal cortex. *Science* 303:380–383.
- Scott SH (2004) Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* 5:534–546.
- Scott SH, Gribble PL, Graham KM, Cabel DW (2001) Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature* 413:161–165.

- Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP (2002) Instant neural control of a movement signal. *Nature* 416:141–142.
- Srinivasan L, Eden UT, Willsky AS, Brown EN (2006) A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural Comp* 18:2465–2494.
- Srinivasan L, Eden UT, Mitter SK, Brown EN (2007) General-purpose filter design for neural prosthetic devices. *J Neurophysiol* 98:2456–2475.
- Suner S, Fellows MR, Vargas-Irwin C, Nakata K, Donoghue JP (2005) Reliability of signals from chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Trans Neural Syst Rehabil Eng* 13:524–541.
- Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci* 3:391–398.
- Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN (2005) A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 93:1074–1089.
- Wang W, Chan SS, Heldman DA, Moran DW (2007) Motor cortical representation of position and velocity during reaching. *J Neurophysiol* 97:4258–4270.
- Yousry TA, Schmidt UD, Alkadhi H, Schmidt D, Peraud A, Buettner A, Winkler P (1997) Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. *Brain* 120:141–157.
- Yu BM, Kemere C, Santhanam G, Afshar A, Ryu SI, Meng TH, Sahani M, Shenoy KV (2007) Mixture of trajectory models for neural decoding of goal-directed movements. *J Neurophysiol* 97:3763–3780.
- Zar JH (1999) *Biostatistical analysis*. Upper Saddle River, NJ: Prentice Hall.