

Naïve coadaptive cortical control

Gregory J Gage¹, Kip A Ludwig¹, Kevin J Otto², Edward L Ionides³
and Daryl R Kipke^{1,4}

¹ Department of Biomedical Engineering, University of Michigan, Ann Arbor, MI, USA

² Kresge Hearing Research Institute, Department of Otolaryngology, University of Michigan, Ann Arbor, MI, USA

³ Department of Statistics, University of Michigan, Ann Arbor, MI, USA

⁴ Department of Electrical Engineering, University of Michigan, Ann Arbor, MI, USA

E-mail: gagegreg@umich.edu

Received 14 February 2005

Accepted for publication 6 May 2005

Published 31 May 2005

Online at stacks.iop.org/JNE/2/52

Abstract

The ability to control a prosthetic device directly from the neocortex has been demonstrated in rats, monkeys and humans. Here we investigate whether neural control can be accomplished in situations where (1) subjects have not received prior motor training to control the device (naïve user) and (2) the neural encoding of movement parameters in the cortex is unknown to the prosthetic device (naïve controller). By adopting a decoding strategy that identifies and focuses on units whose firing rate properties are best suited for control, we show that naïve subjects mutually adapt to learn control of a neural prosthetic system. Six untrained Long-Evans rats, implanted with silicon micro-electrodes in the motor cortex, learned cortical control of an auditory device without prior motor characterization of the recorded neural ensemble. Single- and multi-unit activities were decoded using a Kalman filter to represent an audio ‘cursor’ (90 ms tone pips ranging from 250 Hz to 16 kHz) which subjects controlled to match a given target frequency. After each trial, a novel adaptive algorithm trained the decoding filter based on correlations of the firing patterns with expected cursor movement. Each behavioral session consisted of 100 trials and began with randomized decoding weights. Within 7 ± 1.4 (mean \pm SD) sessions, all subjects were able to significantly score above chance ($P < 0.05$, randomization method) in a fixed target paradigm. Training lasted 24 sessions in which both the behavioral performance and signal to noise ratio of the peri-event histograms increased significantly ($P < 0.01$, ANOVA). Two rats continued training on a more complex task using a bilateral, two-target control paradigm. Both subjects were able to significantly discriminate the target tones ($P < 0.05$, Z-test), while one subject demonstrated control above chance ($P < 0.05$, Z-test) after 12 sessions and continued improvement with many sessions achieving over 90% correct targets. Dynamic analysis of binary trial responses indicated that early learning for this subject occurred during session 6. This study demonstrates that subjects can learn to generate neural control signals that are well suited for use with external devices without prior experience or training.

1. Introduction

There are over 250 000 cases of spinal cord injuries in the United States of America, with a majority of these injuries resulting in quadriplegia: the loss of movement and sensation in both the arms and legs (Lucas *et al* 2004). Electroencephalographic (Lacourse *et al* 1999) and functional magnetic resonance imaging (Shoham *et al* 2001) studies

have shown that spinal cord injured patients who imagine movements in their paralyzed limbs can still produce activation of the motor cortex, even after an extended period of time post trauma. Early studies demonstrating that single units of the motor cortex could be operantly conditioned (Olds 1965, Fetz 1969, Fetz and Finocchio 1971) led Edward Schmidt to propose in 1980 that unit recordings from the motor cortex may constitute a viable control signal for external devices (Schmidt

1980). Recent technology advances have enabled several groups to begin investigating the possibility of cortically controlled neural prostheses (Kennedy *et al* 2000, Serruya *et al* 2002, Taylor *et al* 2002, Carmena *et al* 2003, Shenoy *et al* 2003, Musallam *et al* 2004, Olson *et al* 2005). Many of the current cortical control paradigms consist of analyzing the relationship between cortical activity and measured motor parameters (Chapin *et al* 1999, Wessberg *et al* 2000, Sanchez *et al* 2004). This known relationship is then used to transform the neuronal population signals into real-time prosthetic device movements. Unfortunately, such motor information cannot be obtained in patients with traumatic spinal lesions or neurological disorders which prevent movement.

Several groups have investigated the possibility of neural prostheses that could adapt to cell tuning properties (Taylor *et al* 2002, Eden *et al* 2004, Musallam *et al* 2004). One study (Musallam *et al* 2004) built a database of cortical responses to motor reaches that were subsequently used to decode brain-controlled tasks. Data collected during brain-controlled reaches to targets were used to continually update the database which eventually contained only brain-controlled trials. They reported that this could be done without leading to a loss in performance. Taylor *et al* (2002) demonstrated that an *a priori* neural database was not required for an adaptive algorithm to allow brain-controlled movements. By starting with random tuning properties and allowing these estimates to be iteratively refined, subjects could make long sequences of three-dimensional movements using a brain-controlled cursor. These algorithms mutually adapted to the learning-induced changes in cell tuning properties, thus creating a *coadaptive* neural prosthetic system.

One common thread in these neural prosthetic research models is that the animals are first trained on a motor task. The experiment is then repeated with a brain-controlled task that represents the trained motor control. It is not clear if this *a priori* motor training is required to learn control of a neural prosthetic device. There may also be devices that do not have inherent correlates with physical motor control (e.g. powered wheelchairs, trolleys, communication boards and various adapted vocational tools). Fine control of such neural prosthetic devices would not come from motor training, but rather through adapting arbitrary control signals from the neocortex.

Our objective in this study was to determine if untrained subjects were able to learn control of an unfamiliar neural prosthetic device. We hypothesized that by coadapting to the subject's neuronal responses, we could allow previously naïve subjects to gain control of a foreign neural prosthesis. Here we show that subjects can learn one-dimensional neural control of a novel auditory device without prior motor training.

2. Methods

2.1. Surgical procedure

Six Long-Evans rats weighing 275–300 g (Charles River Laboratories) were used during this study. Animals were kept on a reversed light schedule and housed within the animal

facility of the University of Michigan. Subjects were handled before surgery; however, training did not occur until the cortical control experiments began. Implantation methods have been discussed previously in more detail (Kipke *et al* 2003, Vetter *et al* 2004) and will only briefly be described here. Prior to surgery, anesthesia was induced through an intraperitoneal injection of a mixture of 50 mg ml⁻¹ ketamine, 5 mg ml⁻¹ xylazine and 1 mg ml⁻¹ acepromazine at an injection volume of 0.125 ml/100 g body weight. Anesthesia was maintained with hourly intraperitoneal injections of 0.1 ml ketamine (50 mg ml⁻¹). Subjects were placed in stereotaxic ear bars (MyNeuroLab.com, St. Louis, MO) and an incision was made down the midline of the head. Tissue was removed to reveal the reference fissures on the skull and three bone screw holes were drilled using a surgical bit. A craniotomy was created over the forelimb area of the primary motor cortex (MI) of the left hemisphere. The dura was resected to allow insertion of the penetrating electrode. One 16-channel chronic silicon-substrate microelectrode array (Kipke *et al* 2003, Vetter *et al* 2004) was implanted by hand using fine forceps into the brain (see table 1 for stereotaxic locations). Each electrode (Center of Neural Communication Technology, Ann Arbor, MI; catalog 4 × 4 4mm200 chronic) had four separate shanks (200 μm inter-shank spacing) with four recording sites spaced evenly along each shank (200 μm inter-site spacing). The craniotomy was filled with a hydrogel polymer (ALGEL™, Neural Intervention Technologies, Ann Arbor, MI) to anchor the electrode, and a silicone polymer (Kwik-Sil™, World Precision Instruments) was applied to protect the electrode ribbon cable. Finally, a protective headcap was created using dental acrylic (Co-Oral-Ite, Dental Mfg. Co., Santa Monica, CA). The animals were allowed 48–72 h to recover from surgery. All surgical and animal care procedures were in accordance with the National Institute of Health guidelines and were approved by the University of Michigan Institutional Animal Care and Use Committee.

2.2. Data acquisition

During each experimental session neural electrophysiological data from the 16 electrode channels were sampled at 40 kHz. These signals were simultaneously amplified and bandpass filtered (450–5000 Hz) on a Multichannel Neuronal Acquisition Processor (MNAP; Plexon Inc., Dallas, TX). At the beginning of each recording session, units on each electrode channel were separated and identified using online thresholding, template matching and principal component analysis. Spike times were transmitted with nominal delays over a local TCP/IP connection to a second computer running custom software (MATLAB, Mathworks Inc., Natick, MA) for neural decoding and the environmental hardware control (Tucker-Davis Technologies, Gainesville, FL). Spike times and waveshapes were stored to disk for offline analysis. Event timings for target tone onset and food delivery were captured and stored by pulse signals that synchronized the hardware events with spike timing.

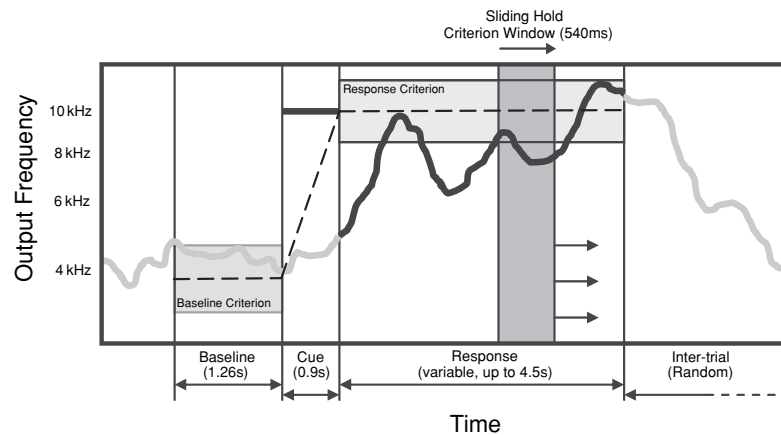


Figure 1. Behavioral paradigm. Dark line indicates auditory tone frequency played back to the subjects (initial cue followed by 90 ms feedback pips). Horizontal shaded regions represent the criterion windows for the baseline and response frequencies. Responses were determined correct and were rewarded if the feedback frequency was maintained within the response criterion of the target frequency, for the duration of a 540 ms sliding window, incrementing in 90 ms steps. The unobserved ideal response or ‘intended’ response used during the training of the adaptive filter is indicated by a thin dotted line. The result of the illustrated trial would be determined as a ‘correct target’ trial.

2.3. Behavioral paradigm

The cortical control system used in these experiments is a one-dimensional auditory analog of the center-out reaching task (Georgopoulos *et al* 1986, Schwartz *et al* 1988). In center-out reaching experiments, the hand (or cursor, in the case of brain-controlled tasks) is held at the center of a circle until a target cue is placed in one of a fixed number of points on the perimeter of the circle. The subject’s task is to move the hand (cursor) into the target position and hold it for a fixed amount of time. Target acquisition must be completed within an allowable response period. In our auditory version, an audio cursor is represented by 90 ms sound pips representing the one-dimensional location within the logarithmically spaced 250 Hz to 16 kHz frequency spectrum. Baseline firing rates were mapped to the center of the frequency space, and trials began with the presentation of a target tone at a given frequency. Subjects had a fixed amount of time to match the target frequency using the auditory cursor. The movement of the auditory cursor was dependent on the real-time decoding of the cortical firing rates as described below. As with the center-out paradigm, trials are marked either as correct (held at correct target frequency for the hold period), wrong (held at an incorrect target frequency for the hold period) or late (no answer within the response period).

Subjects were kept at 85% of their free feeding weight and were tested using either a fixed target task (10 kHz tone, $N = 6$) or a target discrimination task (1.5 kHz or 10 kHz tone, $N = 2$). The fixed target task was run for two to three sessions a day for 8 days. Each session consisted of 100 trials. Figure 1 illustrates trial timing and sequence for this task. Trials began when subjects held the auditory cursor within the criterion window of the baseline level (4 kHz) for 540 ms. The criterion window was set as $\pm 17\%$ of the logarithmic workspace (1.5–10 kHz). Pilot studies determined that this criterion window allowed naïve users to acquire the target in 15–20% of trials by chance. An auditory cue (10 kHz) was then presented for 900 ms to indicate the target frequency. Auditory feedback

of the predicted cursor position was presented to the subjects in 90 ms pips. Subjects had a 4.5 s response window to maintain the cursor within the criterion window of the 10 kHz target for 540 ms. Correct responses were reinforced with a food pellet (45 mg; BioServe #F0021, Laurel, MD). A random inter-trial interval (5–15 s) separated each trial.

The behavioral trial for the target discrimination task was similar to the fixed target, with the exception of an additional target tone. The two target tones were equally spaced (log scale) from the baseline tone. Baseline, target 1, and target 2 frequencies were set to 4 kHz, 10 kHz, and 1.5 kHz respectively. Subjects of the discrimination task were initially trained on a fixed target before being presented with two targets. Experimental runs were constructed as either training or testing sessions. Training sessions repeated missed targets up to four times, while testing sessions pseudo-randomly presented an equal number of the 1.5 kHz or 10 kHz targets. All sessions began with randomized weights and consisted of 200–300 trials. Subjects ran one training session and one to two testing sessions per day. Only testing sessions were used for analysis in this study. At the end of each day of training, supplemental dry food was provided (if necessary) to maintain body weight near 85% *ad lib*.

Our behavioral paradigm provided a contingency between the stimulus (target tone), the desired response (reaching the correct target) and the presentation of a reinforcer (food pellet) which allowed subjects to associatively learn the cortical control task. Initially, the probability of a correct response was low, but through this operant conditioning paradigm, the number of correct trials occurred with higher frequency.

2.4. Ensemble decoding

To enable real-time neural control of the auditory cursor, we used a Kalman filter to infer the cursor frequency from the neural recording data (Wu *et al* 2004). A detailed description of the Kalman filter can be found in Maybeck (1979). Briefly, the Kalman filter is a mathematical procedure that provides

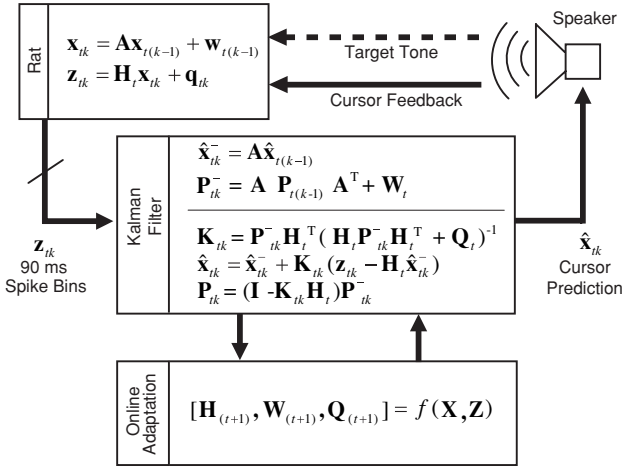


Figure 2. Closed loop cortical control schematic. Spike bins (\mathbf{z}_{ik}) from the motor cortex were decoded using a Kalman filter to predict the cursor frequency ($\hat{\mathbf{x}}_{ik}$). The predicted frequency was fed back to subjects via a speaker every 90 ms of the response window.

an efficient computational means to estimate the state of a process based on noisy Gaussian observations. Here, the observations are the binned spike times, which can be assumed to be Gaussian distributed provided that unit firing rates are sufficiently high. In our paradigm, the cursor frequency was modeled as a system state variable x_{tk} , where t is the trial index ($t = 1, 2, \dots, T$; where T is the total number of trials in the session), and k is an index of time ($k = 1, 2, \dots, K$; where K is the number of 90 ms time steps in trial t). For each trial, the cursor x_{tk} was assumed to propagate in time according to an unobserved difference equation

$$x_{tk} = Ax_{t(k-1)} + w_{t(k-1)}, \quad (1)$$

where A relates the prior cursor position to the current position ($A = 1$ in our experiments) and w_{tk} is a white noise term that was assumed to have a normal probability distribution, $w_{tk} \sim N(0, W_t)$. Note that x_{tk} is a scalar in our one-dimensional paradigm. However, similar equations to those that follow can be written when x_{tk} is a vector of multiple dimensions.

Unit firing times were collected in 90 ms bins and modeled as an observed noisy response to the unobserved state process (1). We define our measurement difference equation that describes the relationship between the cursor frequency (x_{tk}) and recorded spike bins (\mathbf{z}_{tk}) as

$$\mathbf{z}_{tk} = \mathbf{H}_t x_{tk} + \mathbf{q}_{tk} \quad (2)$$

where \mathbf{z}_{tk} is a $C \times 1$ vector of spike bins from C cells, \mathbf{H}_t is a $C \times 1$ vector that linearly relates the frequency state to the neural firing. Again, we assume the noise in the observations has zero mean and is normally distributed, i.e. $\mathbf{q}_{tk} \sim N(0, \mathbf{Q}_t)$.

As figure 2 illustrates, the Kalman filter uses the above state equations (1) and (2) as a model to infer, or predict, the cursor position, $\hat{\mathbf{x}}_{ik}$, given only the observed spike bins, \mathbf{z}_{ik} , of a trial. In order to use this model, we must first define three parameters: \mathbf{H}_t , W_t and \mathbf{Q}_t . It is with these parameters that we adapt the Kalman filter to the neural response of our subjects.

2.5. Filter adaptation

The challenge presented to a ‘naïve’ prosthetic controller is to predict the cursor position, while simultaneously estimating the weights that should be assigned to an evolving neural ensemble experiencing learning-induced changes. For assistance we turn to stochastic control theory, a well-established engineering discipline for tracking non-stationary control signals (Åström 1970, Maybeck 1979). Stochastic control theory offers many tools for dealing with parameter estimation of dynamical state-space systems. Algorithms such as recursive least squares (RLS) or recursive Newton–Raphson provide a way for filter coefficients and cursor predictions to be simultaneously calculated on a bin-by-bin basis (Davis and Vinter 1984). Additionally, block estimation allows for standard system identification techniques to be used on non-stationary signals (Haykin 1996). In block estimation, the available data are divided into individual blocks which are small enough to assume pseudo-stationarity. The filter coefficients are then computed and updated on a block-by-block basis. We selected the block estimation approach as it allowed the filter coefficients to be trained on select regions within the block that maximized our selection criteria (see equation (3)). This selective training technique provided an opportunity to identify and focus on units whose firing rate properties were best suited for control.

The blocks used in this study consisted of data obtained from the past ten trials, and were updated on a trial-by-trial basis. Data from new trials were added to a block by selecting the appropriate time lag in the response window where the unit firing rates had the largest correlation to the expected cursor movement. The time lag at which this occurred was determined by calculating the correlation coefficient of a sliding window of the recorded unit responses with a window of the ‘intended’ target frequency movement for the given trial. The lag chosen for training (l) was calculated across C cells via the formula

$$l = \underset{j}{\operatorname{argmax}} \left(\max_{c \in \{1, 2, \dots, C\}} \operatorname{corr}([\mathbf{x}_{t(1:n)} \quad \mathbf{x}_{t(j:j+n)}], [\mathbf{z}_{t(1:n),c} \quad \mathbf{z}_{t(j:j+n),c}]) \right), \quad (3)$$

where $z_{tk,c}$ is the spike bin count of cell c at time k , and j is the index of a sliding window of length n across R response bins, $j \in [1, 2, \dots, R - n]$. In our experiments, $R = 50$ and $n = 6$. The colon operator indicates concatenation; for example, a vector containing the values of x for the first five time steps of trial t is written as $\mathbf{x}_{t(1:5)} = [x_{t1} \ x_{t2} \ x_{t3} \ x_{t4} \ x_{t5}]$. The function corr is defined as

$$\operatorname{corr}(\mathbf{x}, \mathbf{z}) = \frac{\sum_i (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{z}_i - \bar{\mathbf{z}})}{\sqrt{\sum_i (\mathbf{x}_i - \bar{\mathbf{x}})^2 \sum_i (\mathbf{z}_i - \bar{\mathbf{z}})^2}}, \quad \bar{\mathbf{x}} = \frac{1}{n} \sum_{i=0}^n \mathbf{x}_i. \quad (4)$$

Note that the vector \mathbf{x} of equation (3) does not refer to the predicted cursor position of trial t , but rather to the ideal unobserved cursor movement (see the dashed line of figure 1). The vectors used in the correlation function are a concatenation

of the baseline ($1 : n$) and response ($j : j + n$) windows of the trial, thus allowing for one-dimensional movements. By correlating vectors of the ideal movements with vectors of the observed firing rates \mathbf{z}_t , we can adapt the decoding filter to focus on unit responses that are potentially suited for control.

The lag calculated from (3) over the past $M + 1$ trials was used to train the parameters of the decoding filter for the immediately upcoming trial ($M = 9$ in our experiments). The transformation matrix \mathbf{H} was estimated online by the regression equation

$$\hat{\mathbf{H}}_{(t+1)} = (\mathbf{Z}\mathbf{X}^T)(\mathbf{X}\mathbf{X}^T)^{-1}, \quad (5)$$

where \mathbf{X} is the *intended cursor block* that concatenates \mathbf{X}_t of the past $M + 1$ trials

$$\mathbf{X} = [\mathbf{X}_t \ \mathbf{X}_{(t-1)} \ \mathbf{X}_{(t-2)} \ \cdots \ \mathbf{X}_{(t-M)}], \quad (6)$$

and \mathbf{X}_t is defined as the ideal cursor movement from baseline to response

$$\mathbf{X}_t = [\mathbf{x}_{t(1:n)} \ \mathbf{x}_{t(l:l+n)}]. \quad (7)$$

Similarly, the matrix \mathbf{Z} of (5) is the *observation block* which is defined as the concatenation of the ensemble spike bins from the past $M + 1$ trials at their appropriate lag, l , that best correlated to the ideal movement

$$\mathbf{Z} = [\mathbf{Z}_t \ \mathbf{Z}_{(t-1)} \ \mathbf{Z}_{(t-2)} \ \cdots \ \mathbf{Z}_{(t-M)}] \quad (8)$$

$$\mathbf{Z}_t = [\mathbf{z}_{t(1:n)} \ \mathbf{z}_{t(l:l+n)}]. \quad (9)$$

By using a trial-by-trial sliding block to obtain the least-squared estimate of the relationship between spike rates and cursor frequency ($\hat{\mathbf{H}}$) and by operantly rewarding correct trials, we hypothesized that subjects would learn the behavioral task and develop a strategy to control the auditory cursor.

The noise matrices $\hat{\mathbf{Q}}_{(t+1)}$ and $\hat{\mathbf{W}}_{(t+1)}$ were then estimated using the equations

$$\begin{aligned} \hat{\mathbf{W}}_{(t+1)} &= (\mathbf{X}_{1:(N-1)} - \mathbf{A}\mathbf{X}_{1:(N-1)}) \\ &\times (\mathbf{X}_{2:N} - \mathbf{A}\mathbf{X}_{1:(N-1)})^T / (N - 1), \end{aligned} \quad (10)$$

and

$$\hat{\mathbf{Q}}_{(t+1)} = (\mathbf{Z} - \hat{\mathbf{H}}_{(t+1)}\mathbf{X})(\mathbf{Z} - \hat{\mathbf{H}}_{(t+1)}\mathbf{X})^T / N, \quad (11)$$

where N is the block length, $N = (M + 1)2n$.

The updated filter parameters were used to decode the cursor frequency for the immediately following trial. The end of one trial was the beginning of the next, so $x_{t0} = x_{(t-1)K}$. On the initial trial of each session \mathbf{H}_1 , \mathbf{Q}_1 and \mathbf{W}_1 were randomized.

2.6. Behavioral performance

The behavioral performance for each subject was investigated by comparing the percentage of correct targets with the percentage that would be expected by chance. For the target discrimination task, late trials were discarded and each session was treated as a two-state forced choice paradigm in which chance was 50%. However, the fixed target paradigm consisted of trials that could not end in a ‘wrong target’ state. Therefore, a stimulus randomization method was employed to determine the amount of correct targets that could have been selected by chance during the response period.

The stimulus randomization method is described as follows. If the predicted frequency (\hat{x}) was not related to the stimulus (target) tones, then the tone times can be randomized without affecting the chance that the rat meets the reward criterion. For each training session, 300 sequences of random tone times were drawn. Each random sequence had the same number of tone times and the same tone time distribution as the experimental session. The number of times that the subject would have been rewarded for each of the 300 sequences was used to determine the distribution of chance. For each behavioral session, the *percent correct above chance* was calculated.

For the fixed target task, we defined the *early learning session* as the first session on which there was reasonable certainty ($P < 0.05$) that the subject performed better than chance. For the target discrimination task, we also performed a dynamic analysis of learning using a state-space framework that analyzes binary observations (Smith *et al* 2005). Correct and wrong trials from all testing sessions were arranged to form a time series of binary trial responses and were used to compute the learning curve and its confidence intervals using the state-space smoothing algorithm described in Smith *et al* (2004). We defined the early learning trial as the first trial on which there is reasonable certainty ($P < 0.05$) that a subject performs better than chance for the following 500 trials.

An additional measure of chance was estimated through the use of catch trials. Catch trials consisted of stimuli in which the intensity of the tone and feedback were set to 0 dB to determine the number of trials which reached the correct target by chance.

2.7. Unit analysis

To monitor the evolution of the tuning properties over sessions, we measured the signal strength of the peri-event histograms (PEH) relative to background noise for each unit for every session. The PEH, or signal histogram, was centered on target tone onset and had the width of a full trial window (6.66 s, 90 ms bins). Noise histograms were generated by selecting random times from the recording session, and centering the trial window at this random location. We calculated the root mean square (RMS) of both the signal and noise histograms as

$$\text{RMS}_k = \sqrt{\frac{1}{N} \sum_{n=1}^N (\mathbf{s}_{k,n} - \mu)^2}, \quad (12)$$

where k is a label, $k \in \{\text{signal}, \text{noise}\}$, N is the total number of bins for the histogram, \mathbf{s} is the firing rate value of the n th bin and μ is the mean firing rate. Using the RMS calculations of both signal and noise histograms, we then calculated the signal to noise ratio (SNR) as

$$\text{SNR}(\text{dB}) = 20 \log_{10} \left[\frac{\text{RMS}_{\text{signal}}}{\text{RMS}_{\text{noise}}} \right]. \quad (13)$$

To determine if a recorded unit showed a significant inhibitory or excitatory response, confidence intervals were calculated for each target-onset-centered PEH. The

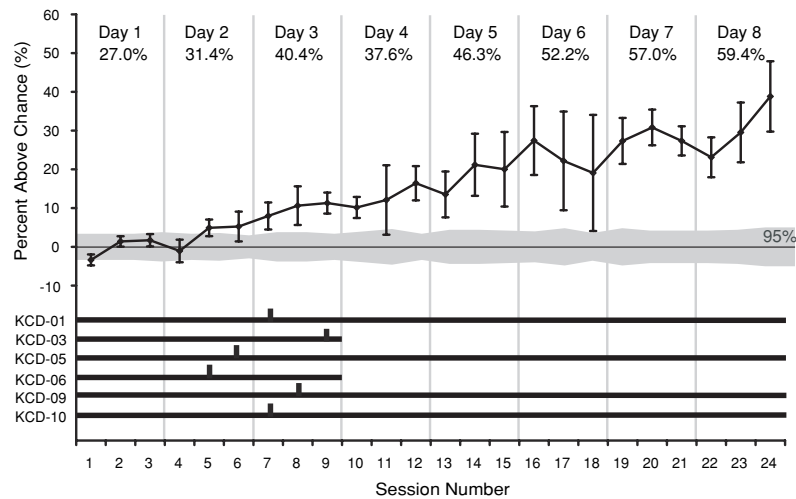


Figure 3. Performance in the fixed target task for six rats across 8 days. The top line shows the mean of the group's performance above chance. Error bars represent standard error. Gray band indicates the 95% confidence interval. Bottom lines indicate the length of time each subject was performing the task. Tick marks indicate early learning sessions. Days of training and mean percent correct for all subjects for each day are indicated across the top.

computations for these intervals were based on the null hypothesis that spike trains are the realization of independent Poisson-point processes as described in the literature (Abeles 1982). Unit responses that crossed the upper 99% confidence interval during the response window were labeled as *excitatory* while those that crossed the lower 99% confidence interval during the response window were labeled *inhibitory*. Responses crossing both confidence intervals were excluded from this analysis.

To test if discrimination had occurred in our subjects, we looked for distinct changes in the unit response pattern for each target. Mean firing rates during the response window of all trials were loaded into MATLAB (Mathworks Inc., Natick, MA) and were separated and labeled according to the target for each trial. Trial by trial analysis was then performed using a Support Vector Machine (SVM) classification toolbox (Ma *et al* 2002). The SVM filter was used to predict the target frequency given the test data from each trial using training data (mean firing rates and target labels) from all other trials. This analysis was repeated for each trial in the session ('*leave-one-out*' method) and the total percent correct was noted. Sessions where the classification percent correct was significantly greater than the calculated binomial distribution ($P < 0.05$, Z-test) were labeled as demonstrating *target discrimination*.

2.8. Histology

Upon completion of training, rats were transcardially perfused with 4% formaldehyde. The brains were removed, sectioned into 40 μm coronal slices and stained with a conventional cresyl violet Nissl stain. The sections were then analyzed using a Leica MZFLIII light microscope (Leica Microsystems, Inc., Germany) to determine probe placement.

3. Results

We summarize two sets of experiments. The first involves data from six naïve subjects in which the task was to move an auditory cursor to a fixed target. Next, we considered the case in which multiple targets were presented and the task was to discriminate and move to the correct target. Two subjects from the initial study were used during the later stage.

3.1. Fixed target task

After recovery from surgery, populations of 5–23 (mean, 11.5) single and multi-unit clusters were discriminated from the recording electrodes and used for the fixed target control task. Subjects were able to control the feedback cursor significantly above chance (early learning session) within 7 ± 1.4 (mean \pm SD) behavioral sessions ($P < 0.05$, stimulus randomization method, $N = 6$). Figure 3 shows the group performance across 24 successive sessions. The 95% confidence intervals indicate two standard deviations above and below the mean chance distribution (randomization method). Horizontal lines indicate the number of days that each subject participated in the study. Tick marks indicate the early learning sessions for each of the subjects. Two subjects were not included after session 9 due to illness (KCD-06) and the loss of all recording units (KCD-03). The average percent of trials on which the target was successfully acquired is shown for each day. This percentage increased from the first session ($21.7\% \pm 6.5\%$, mean \pm SD) to the last ($69.6\% \pm 17.6\%$, mean \pm SD; $P < 0.01$, one-way ANOVA).

Regression coefficients of the subjects' performance over chance as a function of session number are given in table 1. These coefficients were all positive and in five out of six subjects (83.3%) significant ($P < 0.05$, one-way ANOVA), indicating that performance increased with training. The coefficient for the entire group of subjects was 1.53 percentage




Figure 4. Example of trial output aligned with neural ensemble response. (a) Raster plot of spike trains from 14 units over two trials. (b) Behavioral output over two trials. The estimated cursor frequency (\hat{x}) is shown as a thick line, while the expected frequency (x) is shown as a thin line. The criterion windows for the target and baseline frequencies are indicated as horizontal dashed lines. The beginning of each target tone is marked with a black vertical line, while gray vertical lines indicate when the subject produced a correct response and was positively reinforced. (c) Relatively weighted translation decoding \hat{H} matrix. (d)–(g) Peri-event histograms over 100 trials, centered on target onset (indicated by arrow). Bottom bar (R) indicates mean \pm SE of the reward distribution for the PEH plots. Time scale for (a) and (b) is indicated by the 1 s bar. Time scale for (d)–(g) is shown in seconds.

Table 1. Electrode placement and performance for all subjects in the fixed target task. Total number of sessions and the regression coefficients of subject’s performance as a function of session number are given. P values < 0.05 indicate that this coefficient was significant.

Subject	Coordinates (mm)	Sessions	Regression coefficients	P
KCD-01	AP:+3.8 ML:2.0	20	1.6	< 0.001
KCD-03	AP:+2.5 ML:3.0	9	1.3	< 0.05
KCD-05	AP:+2.6 ML:2.5	24	2.4	< 0.02
KCD-06	AP:+3.0 ML:2.5	9	0.24	> 0.58
KCD-09	AP:+1.2 ML:2.9	23	1.1	< 0.02
KCD-10	AP:+1.5 ML:1.6	24	0.96	< 0.01

points per session and was significant ($P < 0.01$, one-way ANOVA).

Figure 4 shows a typical output of the fixed target experiments. Two trials (20–21) from subject KCD-01 are shown from session 24. A raster of the spike train outputs (a) from the 14 sorted single- and multi-unit clusters are aligned with the behavioral output (b). After each trial, the coadaptive algorithm updated the weights (c) of the decoding matrix (\hat{H}) to minimize the output frequency error based on the neural ensemble responses from the past ten trials. The processing time for this updating algorithm to run is indicated in the figure (typically between 2 and 8 s). PEHs averaged

across all 100 trials are shown in (d)–(g). Each PEH is centered at the time that the 10 kHz tone was presented (indicated by arrows). Excitatory and inhibitory responses can be observed in the response window in both the raster plot (a) and the averaged PEH (d)–(g), which allowed for control of the auditory cursor along a one-dimensional frequency axis. Upon tuning, units that responded in an excitatory fashion to the tone during the response window received high positive decoding weights, while inhibitory responses during the response window received larger negative weights. Units that did not respond, or only responded outside the trial window (g), received low weights as they were not useful for device control. During this session, the estimated auditory cursor (\hat{x}) was strongly correlated ($\rho = 0.707$, $P < 0.01$) to the ideal target and baseline frequencies (x).

To visualize the cursor across an entire session, we examined the distribution of the cursor predictions during all of the baseline and response windows. Figure 5 shows a histogram of \hat{x} predictions in 100 logarithmically spaced bins for both the baseline window (4 kHz, gray) and target response window (10 kHz, black) from all trials ($N = 100$) of three sessions of KCD-01. During the first session (a), the distributions for both the baseline and response are centered on the baseline frequency. The response distribution begins to spread into the target window during the early learning session (b) as the subject learns to control the cursor towards the target.




Figure 5. Histogram of auditory cursor frequency estimation (\hat{x}) during two trial windows: baseline (gray) and response (black), for three sessions: naïve, early learning and late learning of KCD-01. Baseline period was fixed at 1.2 s, while the response window was variable (up to 4.5 s) depending on response time. Criterion windows for the baseline and response frequencies are indicated by vertical bars on the \hat{x} axis.

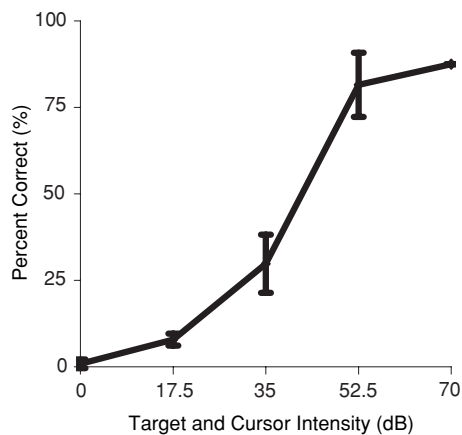


Figure 6. Intensity randomization test. Behavioral results of a well-trained subject (KCD-01) from 341 trials across two sessions where the intensity of both target and feedback tones was randomized between 0 and 70 dB. Error bars indicate standard error.

By the late learning session (c), the distribution becomes bimodal indicating that the subject was able to hold the cursor at the baseline target (4 kHz) during the baseline window, then quickly acquire the target (10 kHz) and hold it for the required criterion (540 ms) during the response period. The change in scaling of the response distribution is an indication of the increased correct responses (thus shorter response windows) during learning. The mean response window length was 4.02 s, 2.82 s and 2.05 s for the naïve session, the early learning session and the late learning session respectively.

To ensure that the subject's behavior was based on the target and feedback tones, we conducted an experiment using randomized intensities of these tones from a set of five levels (0, 17.5, 35, 52.5 and 70 dB). Figure 6 shows the mean behavioral results from 341 random intensity trials across two sessions (KCD-01). The results demonstrate that the subject's behavior was a function of the stimulus intensity, indicating that the target stimuli (and not random fluctuations in cursor predictions) were driving the observed behavior. The sigmoidal shape of figure 6 is consistent with psychophysical measurements of tone detection tasks (Thomas and Setzer 1972). Furthermore, the 0 dB catch trials provided an alternative online determination of chance that we could measure against the randomization method. During these

sessions, 0 dB trials received almost no correct responses ($5.1\% \pm 1.1\%$, mean \pm SD). Chance for this task using the randomization method was calculated to be $19.9\% \pm 2.2\%$ (mean \pm SD), indicating that the randomization method provides a conservative method for determining chance.

The adaptive algorithm selected units that showed potential for control based on the criterion that they were able to modulate from their baseline firing rate during the response window. This selection criterion resulted in changes to the signal to noise ratio of the PEHs across behavioral sessions. The mean SNR increased significantly for the unit PEHs of all rats ($P < 0.01$, ANOVA), indicating that at least one unit in the ensemble had developed the ability to control the system via changes in the firing rates. Moreover, the median SNR also increased significantly ($P < 0.01$, ANOVA) across all subjects indicating that many units in the ensemble were active in control, a desirable property for establishing a robust control signal. Figure 7 shows the distribution of SNR calculations for the PEHs of all units of KCD-05 across 24 training sessions. Median and mean are indicated, and increased significantly with session number ($P < 0.001$, ANOVA). Representative PEHs and the SNR values are shown from sessions 2 and 23. Across all subjects, units from sessions in which the subject did not show significant performance above chance had a mean SNR of 1.7 ± 2.1 dB (mean \pm SD, $N = 377$ units) with median 1.2 dB, while units from sessions where control was significantly above chance measured 3.7 ± 4.0 dB (mean \pm SD, $N = 509$ units) with a median of 4.4 dB.

3.2. Target discrimination task

Two subjects (KCD-09 and KCD-10) were trained on a two target discrimination control task for 30 and 8 testing sessions, respectively. Figure 8 shows the behavioral results for both subjects. The gray band indicates the 95% confidence intervals of chance as calculated from a binomial distribution based on the number of trials for each session. KCD-09 showed a positive learning trend (0.8 percentage points per session, $P < 0.01$, ANOVA) which allowed for several sessions where $>90\%$ of acquired targets were correct. KCD-10 remained at chance for most sessions; however, offline SVM classification (Ma *et al* 2002) of the recorded unit activity was able to correctly classify trials above chance ($P < 0.05$,

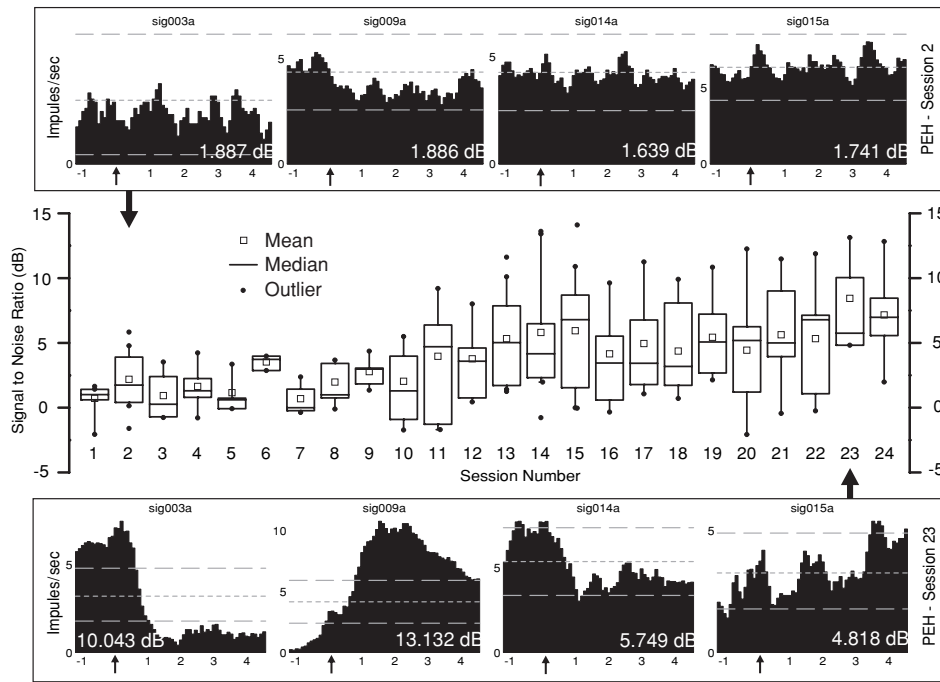


Figure 7. Signal to noise ratio (SNR) results for KCD-05 and examples of eight units from two sessions. Boxes indicate the upper and lower quartile of the distribution, and whiskers indicate the 10th and 90th percentile. The regression coefficients of both the mean and median SNR across sessions were positive and significant (mean, 0.26 dB/session, $P < 0.001$, ANOVA; median, 0.27 dB/session, $P < 0.001$, ANOVA). Representative PEHs and their 99% confidence intervals are shown from sessions 2 and 23. The SNR calculation for each PEH is shown in white.

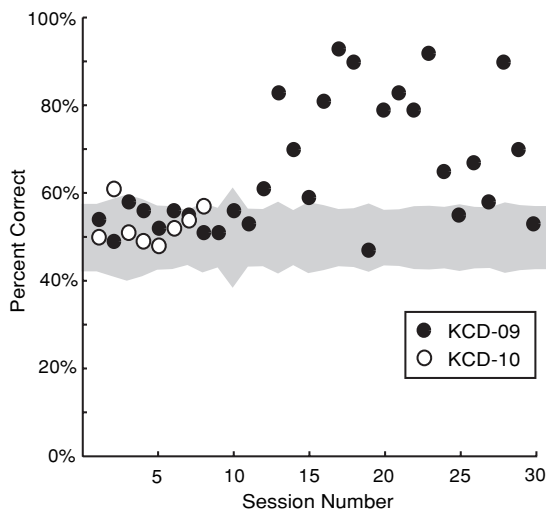


Figure 8. Performance of two subjects in discrimination control task. Gray band indicates the 95% confidence interval of chance as calculated from a binomial distribution based on the total number of trials. Changes in band width are due to slight variations in the number of trials per session. In sessions where both subjects are present, the gray band indicates the larger of the two chance intervals.

Z-test) for both subjects indicating that KCD-10 had learned to discriminate the targets, but had not yet learned control of the cursor to select both targets. The behavioral performance of KCD-09 rises and stays above the 95% confidence interval of chance for several sessions after session 12. However,

dynamic analysis using the method of Smith *et al* (2004) in which correct and wrong trials ($N = 3708$) from all sessions were used as binary observations of the unobservable learning state process indicated that early learning occurred much earlier, on trial 715 (session 6). Dynamic analysis of KCD-10 trials indicated that early learning had not occurred.

Three sessions from KCD-09 fell to chance after the early learning session. During session 19, the units previously used for control could not be sorted from the noise. Offline unit analysis showed that the subject was able to adopt an alternative strategy involving units from other channels in the array on session 20, leading to a recovery in performance.

In order for subjects to control the cursor to reach targets in opposite directions, ideally the unit responses to each target should also be tuned to deviate from the baseline firing rate in opposite directions. Across the 39 behavioral sessions of both subjects, 438 units were manually sorted and used for control (mean \pm SD, 11.1 ± 2.4 units per session). Of these units, 235 (53.7%) showed significant responses immediately following the 10 kHz target tone. These responses were further classified as either excitatory (135 units, 57.5%) or inhibitory (100 units, 42.6%). For the 1.5 kHz tone, 228 units (52.1%) showed significant modulation during the response window, where 176 (77.2%) of these were excitatory and 52 (22.8%) were inhibitory. Of the units that responded significantly to either of the targets (371 units, 87.9%), 92 (24.8%) did not show a discrimination between the tones. These units were either excitatory (71 units, 77.2%) for both of the presented targets, or inhibitory for both (21 units, 22.8%). However, 232 units (62.5%) selectively showed a significant response for one




Figure 9. (a) Raster and PEH plots of three units (selected from ten units) from KCD-09 on session 17 of the two-target task. All plots are centered on target onset (indicated by arrows). The first 15 trials for both 1.5 kHz target (gray) and 10 kHz target (black) are shown in the raster plots; the PEH includes all 200 trials of the session. Thickness of PEH indicates standard error. Dashed lines indicate the upper and lower 99% confidence intervals. (b) Trajectory plot for 10 kHz target (black) and 1.5 kHz target (gray). Late trials were ignored. Horizontal lines indicate the criterion window for the respective target. Thickness indicates standard error.

but not the other target tone. Moreover, 47 units (12.7% of all behaviorally responding units) selectively had significant excitatory responses for one target, while having significantly inhibitory responses for the other.

These bilateral target discriminating units, combined with units that responded for only one of the targets, allowed one subject (KCD-09) to control the cursor to both targets significantly above chance ($P < 0.05$, randomization method) after 12 testing sessions. Figure 9 shows the response of three units (selected from ten) in the target discrimination task from session 17. Raster plots of the first 15 pseudo-random trials for each target have been sorted into 10 kHz (black) and 1.5 kHz (gray). The PEHs for all trials of both tones are plotted directly beneath. The arrow indicates when the target was presented. One unit shows an excitatory response for the 10 kHz tone and inhibitory response for the 1.5 kHz tone, while the other units show the opposite response. The trajectory plot of all trials is shown in (b). The subject in this session was able to reach the correct target in 93% of the trials.

4. Discussion

This study demonstrates that naïve subjects can learn closed-loop, real-time control of one-dimensional cursor movements using single-unit and multi-unit activity of the neocortex. Several investigators have demonstrated similar degrees of adaptation and robustness (Musallam *et al* 2004, Wolpaw and McFarland 2004, Taylor *et al* 2002). Previous studies have demonstrated that one-dimensional control signals could be inferred from simultaneously recorded neurons in rats that were first trained on a lever task (Chapin *et al* 1999, Olson *et al* 2005). Our results suggest that this initial training period may not be required for neural prosthetic control.

We did observe stereotyped motor behaviors in some animals performing the cortical control task. We are unable to determine whether these stereotyped behaviors were related to the recorded units' firing rate modulation, or whether the motor behaviors were merely superstitious behaviors. We can state, however, that all feedback and rewards were based solely on the unit activity of the motor cortex. Other neural prosthetic studies which have looked at EMG signals from the muscle

groups of the cortical region used for control, report that EMG modulations were eventually absent in brain-controlled control tasks (Taylor *et al* 2002, Carmena *et al* 2003). These and other studies (Fetz and Finocchio 1971, Chapin *et al* 1999) indicate that cortical control is possible without direct motor movements.

Our model contains assumption violations that do not affect our conclusions but should be addressed. First, the Kalman filter implicitly assumes a linear Gaussian relationship between cursor movement and unit firing rate. While this assertion is not exact, a linear assumption has been shown to be a reasonable approximation by several investigators for real-time control (Taylor *et al* 2002, Serruya *et al* 2002). The Gaussian assumption only holds for binned data when spike rates are sufficiently high. An alternative approach to avoid this violation would be to use a filter based not on binned data, but on the point process observations of spike times (Eden *et al* 2004). Additionally, the variance of the noise used in state equation (1) was not an accurate representation of the feedback cursor, as this was determined by regressing intended movements, which were specified as ideal (noiseless) movements. Nonetheless, the resulting equations allowed our subjects to quickly learn control of the modeled feedback cursor. Further work needs to be done to establish whether decoding filters that remove these assumptions will assist subjects to learn faster or achieve higher performance.

One limitation of our adaptive algorithm is that it is based on a supervised learning paradigm in which both the intended target and neural activity were used to estimate each neuron's receptive field. This paradigm works well for goal-directed tasks but has disadvantages for longer, free-ranging tasks where receptive fields may change or new neurons become available after the training stage. Our work could be extended to allow the adaptive algorithm to simultaneously predict the cursor estimates while tracking the evolution of the receptive field parameters ($\hat{\mathbf{H}}$) on a bin-by-bin basis. This technique would allow newly added neurons to contribute to the decoding, even when they were not present during the encoding (training) stage. Eden *et al* (2004) provide an example of receptive field tracking in their development of an adaptive stochastic state point process filter for neural decoding.

Subjects were able to control the one-dimensional cursor in the fixed target task using a few (11.5 ± 3.5 , mean \pm SD) MI units per session. This is comparable to similar experiments in monkeys where 7–30 MI neurons were used for two-dimensional control (Serruya *et al* 2002) and 18 units for three-dimensional control (Taylor *et al* 2002). While the number of units used per session did not increase or decrease significantly, there were changes in the number of units sorted and used for control from session to session (coefficient of variance = 30.1%). The regression coefficient of the subjects' performance (percent over chance) as a function of the number of units used for control was calculated to be an increase of 2.4% per unit, and was significant ($P < 0.03$, ANOVA). While this analysis indicates that performance increased with the number of units, it does not mean subjects could not perform the task with fewer units. There were sessions where subjects reached >90% of targets (>50% above chance) using as few as five units.

Our control algorithm requires no prior knowledge of the recorded neuron's tuning properties for adequate system performance. For the initial trial of each session, random weights were assigned to each recorded unit. Subsequently, the system adapted based on stereotyped unit responses across the microelectrode array to enable control. This coadaptive process allowed the rats in our paradigm to learn neural control within 3 days of training. These findings suggest that the described adaptive decoding filter may be a means of training paralyzed human patients where motor tuning properties of neurons are unobtainable. Several groups are investigating brain computer interfaces in which human subjects must learn to derive control signals using electroencephalographic recordings (Wolpaw and McFarland 2004, Fabiani *et al* 2004), electrocorticographic recordings (Leuthardt *et al* 2004), cortical local field potentials (Kennedy *et al* 2004), and single-unit activity (Serruya *et al* 2004, Kennedy *et al* 2000). As each of these systems must derive tuning properties from an initial naïve state, a coadaptive decoding filter (Taylor *et al* 2002) may decrease the time to learn brain control or may maximize control performance by allowing the brain to explore all possible neural responses and to adopt the strategy that is the easiest for control.

Acknowledgments

This work was supported by DARPA contract number N66001-02-C-8059. Probes were provided by the Center for Neural Communication Technology sponsored by NIH NIBIB grant P41-EB002030-11. Aspects of this work were previously presented at the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, San Francisco, USA, Sept. 2004 and at the Society for Neuroscience Annual Meeting, San Diego, CA, Oct. 2004. The authors would like to thank Rio Vetter for surgical assistance, Kristina Shalizi for data analysis, Tim Marzullo for helpful suggestions on the behavioral paradigm, Eugene Daneshvar and Luis Salas for their assistance with animal experiments and Matt Johnson for histology.

References

- Abeles M 1982 Quantification, smoothing, and confidence limits for single-units' histograms *J. Neurosci. Methods* **5** 317–25
- Åström K J 1970 *Introduction to Stochastic Control Theory (Mathematics in Science and Engineering vol 70)* (New York: Academic)
- Carmena J M, Lebedev M A, Crist R E, O'Doherty J E, Santucci D M, Dimitrov D F, Patil P G, Henriquez Craig S and Nicolelis M A L 2003 Learning to control a brain-machine interface for reaching and grasping by primates *PLoS Biol.* **1** e2
- Chapin J K, Moxon K A, Markowitz R S and Nicolelis M A 1999 Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex *Nat. Neurosci.* **2** 664–70
- Davis M H A and Vinter R B 1984 *Stochastic Modelling and Control (Monographs on Statistics and Applied Probability)* (New York: Chapman and Hall)
- Eden U T, Frank L M, Barbieri R, Solo V and Brown E N 2004 Dynamic analysis of neural encoding by point process adaptive filtering *Neural. Comput.* **16** 971–98
- Fabiani G E, McFarland D J, Wolpaw J R and Pfurtscheller G 2004 Conversion of EEG activity into cursor movement by a brain-computer interface (BCI) *IEEE Trans. Neural Syst. Rehabil. Eng.* **12** 331–8
- Fetz E E 1969 Operant conditioning of cortical unit activity *Science* **163** 955–8
- Fetz E E and Finocchio D V 1971 Operant conditioning of specific patterns of neural and muscular activity *Science* **174** 431–5
- Georgopoulos A P, Schwartz A B and Kettner R E 1986 Neuronal population coding of movement direction *Science* **233** 1416–9
- Haykin S 1996 *Adaptive Filter Theory* 3rd edn (Upper Saddle River, NJ: Prentice-Hall)
- Kennedy P R, Andreasen D, Ehirim P, King B, Kirby T, Mao H and Moore M 2004 Using human extra-cortical local field potentials to control a switch *J. Neural. Eng.* **1** 72–7
- Kennedy P R, Bakay R A, Moore M M, Adams K and Goldwithe J 2000 Direct control of a computer from the human central nervous system *IEEE Trans. Rehabil. Eng.* **8** 198–202
- Kipke D R, Vetter R J, Williams J C and Hetke J F 2003 Siliconsubstrate intracortical microelectrode arrays for long-term recording of neuronal spike activity in cerebral cortex *IEEE Trans. Neural Syst. Rehabil. Eng.* **11** 151–5
- Lacourse M G, Cohen M J, Lawrence K E and Romero D H 1999 Cortical potentials during imagined movements in individuals with chronic spinal cord injuries *Behav. Brain Res.* **104** 73–88
- Leuthardt E C, Schalk G, Wolpaw J R, Ojemann J G and Moran D W 2004 A brain-computer interface using electrocorticographic signals in humans *J. Neural Eng.* **1** 63–71
- Lucas J W, Schiller J S and Benson V 2004 Summary health statistics for U.S. adults: national health interview survey, 2001 *Vital Health Stat.* **10** 60–2
- Ma J, Zhao Y and Ahalt S 2002 *OSU SVM Classifier MATLAB Toolbox (v3.00)* (Columbus, OH: Ohio Supercomputer Center (OSC), The Ohio State University) (<http://www.ece.osu.edu/maj/osu.svm>)
- Maybeck P S 1979 *Stochastic Models, Estimation, and Control (Mathematics in Science and Engineering vol 141)* (New York: Academic)
- Musallam S, Corneil B D, Greger B, Scherberger H and Andersen R A 2004 Cognitive control signals for neural prosthetics *Science* **305** 258–62
- Olds J 1965 Operant conditioning of single unit responses *Proc. XXIII Int. Congress Physiol. Sci.*
- Olson B P, Si J, Hu J and He J 2005 Closed-loop cortical control of direction using support vector machines *IEEE Trans. Neural Syst. Rehabil. Eng.* **13** 72–80

- Sanchez J C, Carmena J M, Lebedev M A, Nicolelis M A, Harris J G and Principe J C 2004 Ascertaining the importance of neurons to develop better brain-machine interfaces *IEEE Trans. Biomed. Eng.* **51** 943–53
- Schmidt E M 1980 Single neuron recording from motor cortex as a possible source of signals for control of external devices *Ann. Biomed. Eng.* **8** 339–49
- Schwartz A B, Kettner R E and Georgopoulos A P 1988 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** 2913–27
- Serruya M D, Caplan A H, Saleh M, Morris D S and Donoghue J P 2004 The braingate pilot trial: building and testing a novel direct neural output for patients with severe motor impairments. Program No. 190.22, 2004 Abstract Viewer/Itinerary Planner, San Diego, CA, Society for Neuroscience
- Serruya M D, Hatsopoulos N G, Paninski L, Fellows M R and Donoghue J P 2002 Instant neural control of a movement signal *Nature* **416** 141–2
- Shenoy K V, Meeker D, Cao S, Kureshi S A, Pesaran B, Buneo C A, Batista A P, Mitra P P, Burdick J W and Andersen R A 2003 Neural prosthetic control signals from plan activity *Neuroreport* **14** 591–6
- Shoham S, Halgren E, Maynard E M and Normann R A 2001 Motorcortical activity in tetraplegics *Nature* **413** 793
- Smith A C, Frank L M, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel A M, Suzuki W A and Brown E N 2004 Dynamic analysis of learning in behavioral experiments *J. Neurosci.* **24** 447–61
- Smith A C, Stefani M R, Moghaddam B and Brown E N 2005 Analysis and design of behavioral experiments to characterize population learning *J. Neurophysiol.* **93** 1776–92
- Taylor D M, Helms Tillery S I and Schwartz A B 2002 Direct cortical control of 3d neuroprosthetic devices *Science* **296** 1829–32
- Thomas D R and Setzer J 1972 Stimulus generalization gradients for auditory intensity in rats and guinea pigs *Psychon. Sci.* **28** 22–4
- Vetter R J, Williams J C, Hetke J F, Nunamaker E A and Kipke D R 2004 Spike recording performance of implanted chronic silicon-substrate microelectrode arrays in cerebral cortex *IEEE Trans. Neural Syst. Rehabil. Eng.* **52** 896–904
- Wessberg J, Stambaugh C R, Kralik J D, Beck P D, Laubach M, Chapin J K, Kim J, Biggs S J, Srinivasan M A and Nicolelis M A 2000 Real-time prediction of hand trajectory by ensembles of cortical neurons in primates *Nature* **408** 361–5
- Wolpaw J R and McFarland D J 2004 Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans *Proc. Natl. Acad. Sci.* **101** 17849–54
- Wu W, Black M J, Mumford D, Gao Y, Bienenstock E and Donoghue J P 2004 Modeling and decoding motor cortical activity using a switching Kalman filter *IEEE Trans. Biomed. Eng.* **51** 933–42