Use of a Bayesian maximum-likelihood classifier to generate training data for brain–machine interfaces

2011 J. Neural Eng. 8 046009

(http://iopscience.iop.org/1741-2552/8/4/046009)

IOP PUBLISHING JOURNAL OF NEURAL ENGINEERING

J. Neural Eng. **8** (2011) 049801 (1pp) doi:10.1088/1741-2560/8/4/049801

Corrigendum

Use of a Bayesian maximum-likelihood classifier to generate training data for brain–machine interfaces
Kip A Ludwig *et al* 2011 *J. Neural Eng.* **8** 046009

Received 5 July 2011
Published 20 July 2011

There was an error in the published version of the supplementary data PDF. The correct version is supplied with this corrigendum.

J. Neural Eng. **8** (2011) 046009 (15pp) doi:10.1088/1741-2560/8/4/046009

Use of a Bayesian maximum-likelihood classifier to generate training data for brain–machine interfaces

Kip A Ludwig, Rachel M Miriani, Nicholas B Langhals, Timothy C Marzullo and Daryl R Kipke

Department of Biomedical Engineering, University of Michigan, 1101 Beal Ave., 2247 LBME, Ann Arbor, MI 48109, USA

E-mail: ludwigk@umich.edu and dkipke@umich.edu

Received 2 March 2011
Accepted for publication 11 May 2011
Published 8 June 2011
Online at stacks.iop.org/JNE/8/046009

Abstract

Brain–machine interface decoding algorithms need to be predicated on assumptions that are easily met outside of an experimental setting to enable a practical clinical device. Given present technological limitations, there is a need for decoding algorithms which (a) are not dependent upon a large number of neurons for control, (b) are adaptable to alternative sources of neuronal input such as local field potentials (LFPs), and (c) require only marginal training data for daily calibrations. Moreover, practical algorithms must recognize when the user is not intending to generate a control output and eliminate poor training data. In this paper, we introduce and evaluate a Bayesian maximum-likelihood estimation strategy to address the issues of isolating quality training data and self-paced control. Six animal subjects demonstrate that a multiple state classification task, loosely based on the standard center-out task, can be accomplished with fewer than five engaged neurons while requiring less than ten trials for algorithm training. In addition, untrained animals quickly obtained accurate device control, utilizing LFPs as well as neurons in cingulate cortex, two non-traditional neural inputs.

S Online supplementary data available from stacks.iop.org/JNE/8/046009/mmedia

(Some figures in this article are in colour only in the electronic version)

1. Introduction

The study of the relationship between neural firing rates and associated motor activities in the 1980s and early 1990s led to a number of groundbreaking neural control applications, known more commonly as brain–machine interfaces (BMIs), over the last ten years [1–7]. Typical BMIs use previously measured movements from a subject to map a linear relationship between a subject's recorded neural firing rates and each movement parameter, and then apply this map directly to the control of a device [2, 6, 8–12]. This approach has the advantage of being both simple and intuitive, while still yielding insights into the relationship between cortical neurons and natural movement.

Unfortunately, generating training data from which to derive a linear relationship between neuronal firing rates and movement parameters can be problematic in practical application. BMIs are typically designed for use by individuals with severe motor disabilities and are therefore unable to make movements for calibration purposes. In recent studies, researchers addressed this issue by mapping cortical neural activity to *observed* movements [8, 12]. Similarly, Hochberg *et al* successfully mapped neural activity to *imagined* movements [13]. Although calibrating a BMI device based on observed or imagined movement presents a promising alternative to *a priori* movement measurements, the exact relationship between observed/imagined movements and neuronal firing rates is ill-defined [11]. There is no obvious indicator of when the neural activity related to the observed/imagined movement begins or ends, or if any relationship is present at all on a particular trial [11]. Consequently, neural data unrelated to any movement parameter may unintentionally be used to train the neural decoding algorithm, degrading device performance [11].

Another issue which needs to be addressed before the experimental promise of BMI systems can be translated to practical application is self-paced control [11, 14, 15]. According to Bashashati *et al*, in a self-paced system the user directs the BMI by intentionally changing neural activity [11]. Between periods of intentional control (IC), the subject is not actively attempting to use the BMI device (No-Control, NC) [11]. Current cortical control paradigms, however, operate under the assumption that the subject is always intentionally attempting to control the output device. This methodology presents two problems, (a) the output device responds when the user is not intending to generate a control signal, and (b) neural activity unrelated to observed/imagined movement parameters is intermittently used to train the system [8, 11, 16].

For BMI devices to translate to real clinical applications, decoding algorithms need to be based on assumptions that are easily met outside of an experimental setting [17, 18]. Given present technological limitations, a low number of potentially unstable neuronal units must be assumed from day to day [19], driving a need for decoding algorithms which (a) are not dependent upon a large number of neurons for control, (b) are adaptable to alternative sources of neuronal input such as local field potentials (LFPs), and (c) require only marginal training data for day-to-day calibrations. Moreover, practical decoding algorithms must be able to recognize and eliminate poor training data, and operate without initial movement measurements for calibration.

In this paper we introduce and evaluate a Bayesian maximum-likelihood estimation (bMLE) strategy to identify optimized training data for BMI devices in order to address the issues of isolating quality training data and self-paced control. Instead of correlating small samples of noisy—and potentially unrelated—neural data to previously measured movement parameters in order to train a decoding system, we analyzed neural recordings for separate, distinct neural states created during normal behavior, and then assigned these individual states to control directions.

We implanted six Sprague-Dawley rats and employed our modified bMLE scheme to allow the subjects to build their own neuronal output states to accomplish a multiple state classification task, loosely based on the standard center-out task [20]. Subjects were able to quickly adapt to the bMLE scheme and generate neuronal control sufficient to reliably complete a useful BMI task, despite limited training data and no external movement information. By limiting the initial assumptions in our decoding strategy, we also demonstrate that useful neural control is possible using LFPs and neurons in cingulate cortex, two prospective neuronal information sources without a clearly defined linear relationship to movement parameters. All of this was accomplished using very small training data sets and five or fewer isolated neurons, without requiring any *a priori* movement information.

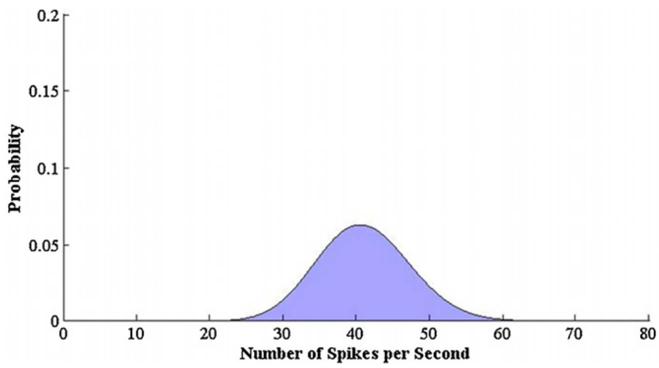

Figure 1. Probability density function of a single neuron with a mean firing rate of 40 Hz based on a Poisson distribution. Note that over a 1 s interval, there is only a 6.3% chance of observing exactly 40 spikes from this neuron.

2. Experimental design and methods

In order to develop a classification strategy aimed at identifying separable neural states from real-time recordings, we chose to base our work on previous nonlinear bMLE schemes [15, 21–23]. A nonlinear decoding scheme was selected as prior studies have indicated that there is significant nonlinear information present in the neural recordings [1, 9, 23–26]. In addition, Santhanam *et al* found that the information transfer rate of BMI devices can be greatly enhanced by eliminating the prediction of unnecessary movement parameters (for example continuous trajectory) using MLE techniques [22].

2.1. Review of Bayesian maximum-likelihood estimation based on a priori measurements

Bayesian maximum-likelihood prediction of stimuli is dependent upon determining the probability distribution of neuronal responses for each delivered stimulus [27–32]. For purposes of this discussion, a stimulus may also include a measured movement parameter or a target position. Given the past history of neuronal responses with respect to different stimuli, observed neural responses can then be used to predict which stimulus is presently being delivered.

As a simple example, assume that the mean firing rate for a specific neuron when a subject's arm remained stationary was previously determined to be 40 Hz. The probability of this neuron firing n times over a short interval Δt while the arm remains stationary is given by the following Poisson statistic (see figure 1):

$$P(n \text{ spikes}|\text{stationary}) = (40\Delta t)^n \, e^{-(40\Delta t)}/n! \qquad (1)$$

A Poisson distribution of spike counts was assumed in this work, as a Poisson distribution has been demonstrated to be a more accurate model of neural firing rates than standard Gaussian distributions [22, 25, 27, 29].

Furthermore, assume that the firing rate of the same neuron during a subject's arm movement to the right was previously determined to be 80 Hz. The probability of this neuron firing n times over a short interval Δt while a subject's

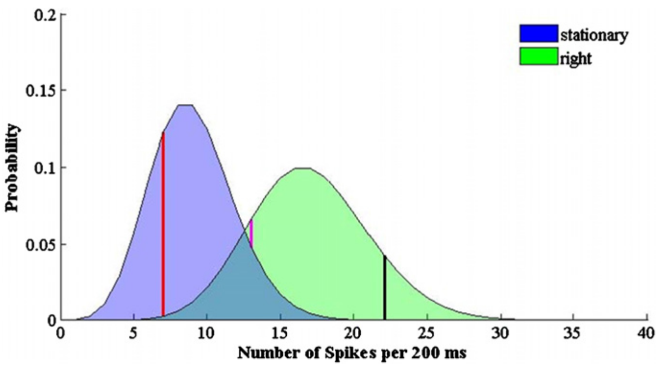

Figure 2. Probability density function of a single neuron's firing rate when the subject's arm remains stationary or moves to the right. The red line is located at 7 spikes over a 200 ms interval, the magenta line is located at 13 spikes, and the black line is located at 22 spikes. Note that the probability of 7 spikes occurring during a 200 ms interval is much greater when the subject's arm remains stationary than when the subject's arm moves right. In terms of a BMI, an observed firing rate of 7 spikes per 200 ms indicates the device should remain stationary. Conversely, an observed firing rate of 22 spikes per 200 ms would indicate the prosthetic device should move to the right. An observed firing rate of 13 spikes per 200 ms would not yield a useable prediction, as this firing rate is approximately equally likely when the subject's arm remains stationary or moves to the right.

arm moves to the right is given by the following Poisson statistic:

$$P(n \text{ spikes}|\text{right}) = (80\Delta t)^n \, \mathrm{e}^{-(80\Delta t)}/n! \tag{2}$$

Using Bayes' law, the observed firing rate of the neuron can then be used to predict whether the subject's arm is presently stationary or moving to the right (figure 2). Given an observed firing rate of 7 spikes over 200 ms, equation (1) becomes

$$\begin{aligned} &P(7 \text{ spikes over 200 ms}|\text{stationary}) \\ &\quad = (40\text{ Hz} \times 0.2\text{ s})^{7\text{ spikes}} \, \mathrm{e}^{-(40\text{ Hz}\times 0.2\text{ s})}/7\text{ spikes}! = 0.1396. \end{aligned}$$

Similarly, equation (2) becomes

$$\begin{aligned} &P(7 \text{ spikes over 200 ms}|\text{right}) \\ &\quad = (80\text{ Hz} \times 0.2\text{ s})^{7\text{ spikes}} \mathrm{e}^{-(80\text{ Hz}\times 0.2\text{ s})}/7\text{ spikes}! = 0.0060. \end{aligned}$$

Assuming that the likelihood of subject's arm remaining stationary or moving to the right is equally likely, Bayes' law states that the probability of the subject's arm remaining stationary given an observed firing rate of 7 spikes over 200 ms can be calculated as follows:

$$\begin{aligned} &P(\text{stationary}|7 \text{ spikes over 200 ms}) \\ &\quad = \frac{P(7 \text{ spikes over 200 ms}|\text{stationary})}{P(7 \text{ spikes over 200 ms}|\text{stationary})+P(7 \text{ spikes over 200 ms}|\text{right})} \\ &\quad = 0.1396/(0.1396 + 0.0060) \\ &\quad = 0.9588, \text{ or approximately } 96\%. \end{aligned}$$

Note that this formula is simply the ratio of the probability of observing 7 spikes while the subject's arm remains stationary over the total probability of observing 7 spikes.

Likewise, the probability of the subject's arm moving to the right given an observed firing rate of 7 spikes over 200 ms is given by

$$\begin{aligned} &P(\text{right}|7 \text{ spikes over 200 ms}) \\ &\quad = \frac{P(7 \text{ spikes over 200 ms}|\text{right})}{P(7 \text{ spikes over 200 ms}|\text{stationary})+P(7 \text{ spikes over 200 ms}|\text{right})} \\ &\quad = 0.0060/(0.1396 + 0.0060) \\ &\quad = 0.0412, \text{ or approximately } 4\%. \end{aligned}$$

In terms of a BMI, a threshold of 95% probability can be used to determine a significant prediction. If the observed firing rate from a group of neurons predicts a movement to the right with a probability of greater than 95%, the BMI device is then moved to the right. For this one neuron example, a firing rate of 7 spikes per 200 ms indicates that the BMI controlled device should remain stationary (see figure 2). Conversely, an observed firing rate of 22 spikes per 200 ms indicates the BMI device should move to the right (P(stationary | 22 spikes per 200 ms) < 0.001, P(right | 22 spikes per 200 ms) > 0.999). An observed firing rate of 13 spikes per 200 ms does not yield a useable prediction, as this firing rate is likely to occur when the subject's arm remains stationary or moves to the right (P(stationary | 13 spikes per 200 ms) < 0.2668, P(right | 13 spikes per 200 ms) > 0.7332).

This methodology can easily be extended to incorporate multiple neurons and multiple stimulus parameters. Assuming that the previously determined mean firing rate of a group of neurons with respect to a specific stimulus s_j is defined as the vector $\mathbf{M}_j$:

$$\mathbf{M}_j = \{\mu_{j1}, \ \mu_{j2}, \ \ldots, \mu_{jn}\},$$

where n is the total number of recorded neurons and μ_{j1} is the mean firing rate of neuron 1 during stimulus s_j.

Given a vector of observed neuronal responses R from this group of neurons, where r_1 is the observed firing rate from neuron 1 of n neurons

$$R = \{r_1, \ r_2, \ldots, r_n\}.$$

The probability of an observed response vector R during the specific stimulus s_j is

$$P(R|s_j) = \prod_{i=1}^{n} (\mu_{ji}\Delta t)^{r_i} \, \mathrm{e}^{-(\mu_{ji}\Delta t)}/r_i! \tag{3}$$

If s is one of a set of m possible stimuli, then the set of possible stimuli S is given by

$$S = \{s_1, \ s_2, \ldots, s_m\}.$$

Assuming the delivery of all stimuli in set S is equally probable, the probability of any specific stimulus s_j based on observed response vector R can be computed as follows:

$$P(s_j|R) = P(R|s_j) \Big/ \sum_{i=1}^{m} P(R|s_i). \tag{4}$$

As in the one neuron example, a response vector R that predicts a specific movement s_j with a confidence level greater than 95% indicates that the BMI device should also perform movement s_j. This basic Bayesian framework is often extended to include movement kinematics (previous position, velocity, and acceleration, for example) to further refine instantaneous movement predictions [2, 23, 24, 27, 29].

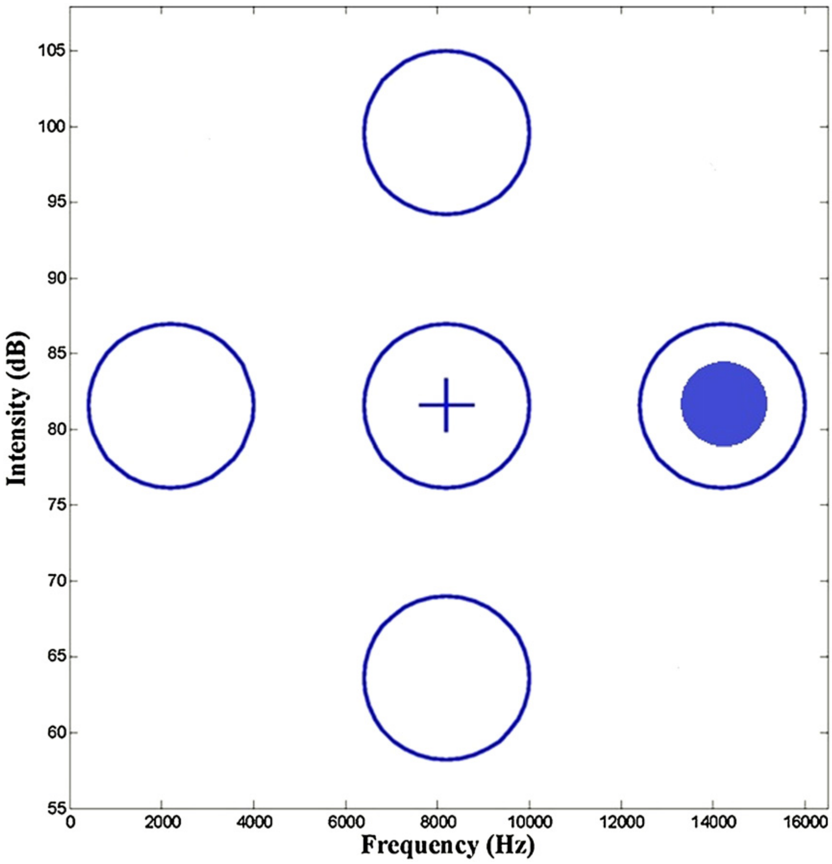

Figure 3. The virtual center-out task (visual representation). Each circle on the periphery represents possible targets. The middle circle represents the center 'start' button. The subject is given neural control of a virtual cursor on the two-dimensional axis, represented by the cross hairs. On a given trial, the subject is required to first move the cursor into the center circle for 1 s. A target circle is then presented, denoted by the filled blue circle. During the response window, the subject is required to move the virtual cursor into the target circle for 1 s. Moving the cursor into a circle other than the target results in an incorrect trial. Note that the axes in this figure generate an auditory analog of the center-out task (see section 2.2.3).

2.2. Bayesian MLE classification scheme

In the standard center-out task, subjects are trained to reach out and interact with indicated targets. A center target is typically placed directly in front of the subject, and additional targets are placed equidistant from the center at regular angular intervals. In a given trial, the center target is indicated first, and the subject is required to move their hand to touch the center target for 1 s. Next, one of the additional targets is randomly indicated. For a successful trial, a subject is required to move their hand from the center to the indicated target [21, 33].

In a typical neural controlled center-out task, first neural recordings are obtained during manual performance of the task. These neural recordings are then analyzed to find neural states that were indicative of either measured kinematic parameters during movement, or a successfully touched target position (end goal) [22]. In this fashion, a map is generated between the subject's neural state and specific target positions or kinematic parameters. This neural map is then applied to determine the position of a virtual cursor during a neural controlled center-out task. Instead of manually moving their hands to touch a center target and then a specified target, subjects are required to manipulate their neural states to move a virtual cursor into a center circle, and then to a specified target position (see figure 3).

Figure 4 depicts a flow chart of our bMLE decoding scheme, which proceeds identically to standard bMLE with one major exception. In our bMLE scheme, *a priori* movement measurements were not available to create a neural map. Similar to the work of Bashashati, neural data for training was identified based on external knowledge of the 'approximate' time of intended control (IC) [11]. During the IC response window, IC may or may not be present, and the timing of IC is unknown. Neural data during the IC response window may belong to either the IC or NC class.

Our task required the subject to first develop a distinct neural state for the center position on the two-dimensional (2D) virtual axis, and subsequently for each of the possible target positions (see figure 3). After developing a neural state for each target, real-time control proceeded according to the standard bMLE scheme. To successfully complete a trial, subjects were required to match their neural state to the neural state associated with the center position on the 2D axis for 1 s (see figure 4, question 1), and then match their neural state

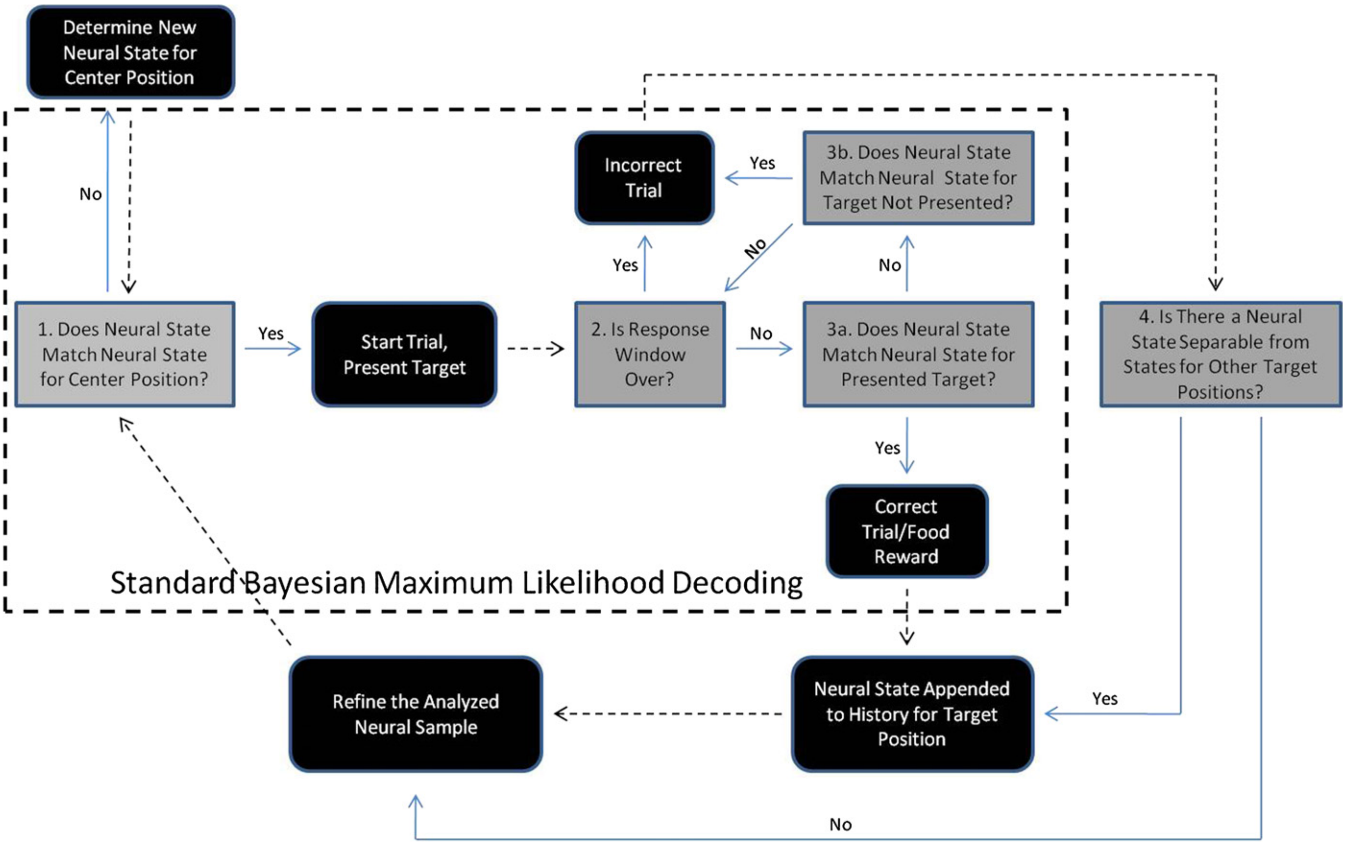

Figure 4. Flowchart of bMLE decoding scheme. A trial begins after question 1 is answered 'Yes', indicating that the observed neural state matched the neural state indicative of the center position for 1 s. Gray boxes indicate questions, whereas black boxes indicate actions. Solid lines denote a path taken depending on the answer to a question. Dotted black lines indicate a path automatically taken as the next in a sequence. The large dotted black box indicates steps that occur in a standard bMLE scheme based on *a priori* measurements.

to the neural state associated with a randomly selected target position for 1 s (see figure 4, question 3a).

At the beginning of an experimental run, baseline neural recordings were taken for 20 s to establish a mean neural state representative of the NC state. The firing rate of each recorded neuron was calculated in 200 ms non-overlapping windows (also known as 'bins'), and then averaged over the 20 s of baseline recordings to determine the mean firing rate of each neuron during the baseline NC window (see figure 5, step 1). This NC neural state was then mapped to the center of the 2D auditory axis. In order to start a trial, subjects were required to match their current neural state to the NC state with a confidence level of greater than 95% for five consecutive 200 ms bins (see section 2.1, equation (4)).

2.2.1. Developing separable neural states for each target. At the start of each trial, one of the targets was randomly selected to be presented to the subject. On the first presentation of a target position, there was no previous neural history from which to generate a unique neural state for the target. As a result, the subject could not match their current neural state to the neural state indicative of the target, and the first presentation of any target always resulted in an incorrect trial. After all incorrect trials, the neural activity over the 10 s response window was inspected for a candidate neural state to be added to the history for the target in subsequent trials (see figure 4, question 4).

Candidate neural states were selected on the basis of 'separability' from the neural state already assigned to the center and prior target positions. 'Separable' neural states can be statistically distinguished from the neural states already associated with the center and other target positions using a bMLE classification scheme, and therefore can be effectively used to predict an additional target position in the future. To identify 'separable' neural states, the 10 s response window was inspected for neural states that were unlikely to have been generated by random fluctuations around the mean neural firing rates already associated with previous target positions. Returning to a simple one neuron example from section 2.1, a neuron with a mean firing rate of 40 Hz during the baseline NC state and a mean firing rate of 45 Hz at target position 1 would have low separability, as this neuron's probability density functions for the NC and target position 1 states largely overlap, providing little information to distinguish between these two states. Conversely, a neuron with a mean firing rate of 20 Hz during the NC state and 80 Hz at target position 1 provides information to distinguish between these two states, as the probability density function for each state is separate and distinct, and therefore has high separability.

Target Presented

Step 1 Neuron Firing Rate (0.2 Second Bins)

	B1	B2	B3	B4	B5	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	R16	R17	R18	R19	R20	R21	R22	R23	R24	R25
Neuron 1	15	8	17	14	11	13	10	13	6	19	11	13	18	16	16	10	10	11	11	15	12	8	16	19	15	12	15	15	15	12
Neuron 2	1	1	0	1	3	1	2	0	1	2	1	0	1	0	0	2	1	0	1	1	0	1	0	1	2	1	1	0	2	2
Neuron 3	15	12	9	8	13	14	11	12	10	8	9	14	11	12	0	1	0	0	0	0	0	1	0	1	0	11	9	13	9	15
Neuron 4	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	1	0	1	0	1
Neuron 5	4	10	8	6	6	6	5	7	6	7	8	9	7	10	7	8	5	8	10	9	8	6	5	11	12	6	10	5	9	7
Neuron 6	9	9	11	8	8	14	7	11	11	11	10	6	8	10	6	8	10	9	8	8	7	12	8	12	9	6	6	8	7	6
Neuron 7	5	4	4	4	4	6	4	2	3	2	4	5	1	6	15	10	15	13	14	12	12	15	12	12	12	4	8	1	4	5
Neuron 8	11	15	11	8	15	18	15	15	15	9	10	10	8	12	1	0	1	3	2	1	1	1	1	2	2	12	12	14	18	16

Center Position/NC Mean Block 1 Mean Block 3 Mean Block 17 Mean

Step 2 Mean Firing Rate for the Neurons in Each Block

	Mbase	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18	M19	M20	M21
Neuron 1	13.0	12.2	11.8	12.4	13.4	15.4	14.8	14.6	14.0	12.6	11.6	11.4	11.8	11.4	12.4	14.0	14.0	14.0	15.4	15.2	14.4	13.8
Neuron 2	1.2	1.2	1.2	0.8	1.0	0.8	0.4	0.6	0.8	0.6	0.8	1.0	0.6	0.6	0.6	0.6	0.8	1.0	1.0	1.0	1.2	1.2
Neuron 3	11.4	11.0	10.0	10.6	10.4	10.8	9.2	7.6	4.8	2.6	0.2	0.2	0.0	0.2	0.2	0.4	0.4	2.6	4.2	6.8	8.4	11.4
Neuron 4	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.4	0.4	0.4	0.4	0.2	0.0	0.0	0.0	0.0	0.2	0.2	0.4	0.4	0.6
Neuron 5	6.8	6.2	6.6	7.4	7.4	8.2	8.2	8.2	7.4	7.6	7.6	8.0	8.0	8.2	7.6	7.8	8.4	8.0	8.8	8.8	8.4	7.4
Neuron 6	9.0	10.8	10.0	9.8	9.2	9.0	8.0	7.6	8.4	8.6	8.2	8.6	8.4	8.8	8.6	9.4	9.6	9.4	8.2	8.2	7.2	6.6
Neuron 7	4.2	3.4	3.0	3.2	3.0	3.6	6.2	7.4	9.4	11.8	13.4	12.8	13.2	13.2	13.0	12.6	12.6	11.0	9.6	7.4	5.8	4.4
Neuron 8	12.0	14.4	12.8	11.8	10.4	9.8	8.2	6.2	4.4	3.4	1.4	1.4	1.6	1.6	1.2	1.2	1.4	3.6	5.8	8.4	11.6	14.4

Step 3 Percentage of 'Variance' Vectors Centered on Mj Predicting Target Position j Versus MBase

	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18	M19	M20	M21
P	3.8	0.4	0.0	0.0	1.2	20.8	36.4	70.8	87.6	99.2	98.4	99.8	99.0	99.0	99.2	98.2	86.8	99.0	99.2	98.2	86.8

Step 4 Percentage of 'Variance' Responses Centered on Mj,k Predicting Target Position j Versus Mbase,k (Neuron by Neuron Analysis)

	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18	M19	M20	M21
~~Neuron 1~~	0.0	0.3	0.0	0.0	20.7	6.2	5.0	0.1	0.0	2.0	3.9	0.5	4.7	0.0	0.3	0.1	0.1	19.2	15.8	2.2	0.0
~~Neuron 2~~	7.7	6.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	8.5	7.4
Neuron 3	0.0	0.0	0.0	0.0	0.0	3.6	40.3	89.4	99.3	100.0	100.0	100.0	100.0	100.0	100.0	100.0	99.3	93.5	59.5	21.5	0.0
~~Neuron 4~~	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
~~Neuron 5~~	0.2	0.0	0.0	0.0	4.5	4.7	5.8	0.0	0.0	0.0	2.3	2.2	5.0	0.0	0.1	11.4	1.9	20.5	18.7	9.0	0.0
~~Neuron 6~~	36.6	15.6	8.8	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.0	3.1	1.7	0.0	0.0	1.7	13.4
Neuron 7	0.0	0.4	0.0	0.5	0.0	58.3	77.3	92.1	97.3	99.0	98.7	99.3	98.8	99.0	98.5	98.3	96.0	92.3	78.1	47.0	1.6
Neuron 8	14.8	0.0	0.0	8.3	20.4	57.8	85.9	97.3	99.1	100.0	100.0	100.0	100.0	100.0	100.0	100.0	98.3	90.2	54.4	0.0	13.1

*** Note: Grey background indicates neurons which have changed their firing rates significantly during the response window.**
White font was used to highlight the change in firing rate of these neurons during separable response vectors.

Figure 5. Analyzing the response window for a new neural state. The response window for each trial began immediately after the target position, j, was presented. The firing rates of individual neurons during the response window were binned every 0.2 s to generate an observed response vector R_i (step 1). The vectors for each bin were then grouped into overlapping blocks of five bins, and a mean vector for each block, $\mathbf{M}_j$, was determined (step 2). Using $\mathbf{M}_j$ as the mean neural state for the presented target position, 1000 variance vectors were generated around $\mathbf{M}_j$ and tested as observed responses (step 3). For a block to be included in the history for the target position j, at least 95% of the variance vectors were required to predict target position j (as compared to the mean neural state for the center and other target positions) with a probability of less than $p < 0.01$. In step 4, steps 1 through 3 were repeated on a neuron-by-neuron basis to determine which neurons (if any) were modulating significantly during the response window. Neurons excluded from future Bayesian MLE analysis have been indicated by a dotted red line.

Figure 5 depicts a detailed example of the process used to identify separable neural states present in the response window, and proceeds as follows.

Step 1, creating bins: first, the firing rate of each neuron was calculated in 200 ms non-overlapping bins to create the vector R_k, a C $\times$ 1 vector containing the firing rates of C observed neurons for bin k (where $k = 1$ indicates the 200 ms bin immediately following the presentation of the target, $k = 2$ indicates the bin starting at 200 ms after target presentation and ending at 400 ms, etc; see figure 5 step 1 for an example).

Step 2, calculating the mean vector: the vectors for each bin were then grouped into overlapping blocks of five bins (where block 1 consists of vectors R_{1-5}, block 2 consists of R_{2-6}, etc), and a C $\times$ 1 mean vector for each block was determined (see figure 5 step 2 for an example). This was performed to generate a 1 s mean updated every 200 ms for subsequent calculations.

Step 3, testing neural states for separability: next, the mean vector for each block was tentatively assigned as $\mathbf{M}_j$, a C $\times$ 1 vector to be examined as a possible new neural state for target position j. From a practical application perspective, a subject cannot consistently generate a neural response state exactly matching a mean vector; there is an inherent variability in the observed firing rate of a neuron around a given mean [34, 35]. For a new neural state to be useful in the context of a BMI task, expected fluctuations around the mean firing rates for the state must be differentiable from the neural states already assigned to other target positions. To account for this, 1000 artificially generated 'variance' vectors, drawn from the probability distribution resulting from the expected variance centered around $\mathbf{M}_j$ (see supplementary material (B) available at stacks.iop.org/JNE/8/046009/mmedia for details), were tested as observed responses. Equation (4) was then used to calculate the probability of target position j given each observed response, using $\mathbf{M}_j$ as the vector mean for the neural state representing position tone j. For $\mathbf{M}_j$ to be considered 'separable', at least 95% of the 'variance' vectors were required to predict target position

j (as opposed to the neural states for the center and the other target positions) with a probability of $p < 0.01$ (see figure 5 step 3 for an example).

If a 'separable' block was evident in the response window, this block was appended to the history for target position j. If multiple blocks were separable, the block with the greatest number of 'variance' vectors accurately predicting target position j was appended to the history for target position j. If no 'separable' block was evident during the IC response window, no block was appended to the history for the target position. This distinction is important, as adding data to the history for the target when no useful neural modulation was present would significantly alter the mean firing rates for the neural state associated with the target position, and therefore negatively impact performance. The mean of all the individual vectors in the history for target position j was used to generate the mean neural state for target position j, $\mathbf{M}_j$, for future trials.

Step 4, refining the neuronal sample: including neurons that do not change firing rates with respect to at least one target detracts from neural control performance, as these neurons act as additional noise sources [8]. The neuronal sample for a Bayesian MLE classifier is normally pared down to include only the neurons positively contributing to performance by iterating through all possible combinations of neurons as input to the classifier. Due to the time constraints of a real-time application, this strategy was not technologically feasible. Therefore, an alternative methodology to eliminate neurons failing to positively contribute to completing the BMI task was implemented. Using the procedure outlined earlier to locate separable neural states in the response window, neurons were analyzed on an individual basis to locate neurons that changed firing rates significantly during the IC response window. A neuron, *k*, was deemed to generate a significant response to a target if at least 90% of the artificially generated variance responses centered around the neuron's mean for a block in the response window, $\mathbf{M}_{j,k}$, were sufficient to differentiate a new target position from center position/NC with $p < 0.05$ (see figure 5 step 4). For a neuron to be included in the Bayesian MLE analysis, it was required to exhibit one significant response to at least one of the target positions in the trial history. This strategy ensured that every neuron included in the real-time analysis significantly modulated during the presentation of at least one target position.

The size of the history for each target was adjusted based on the percentage of correct trials for the presentation of each target (see supplementary material (C) available at stacks.iop.org/JNE/8/046009/mmedia for details). The prior history for a given target was shortened during periods of poor performance and lengthened during periods of strong performance. This strategy was implemented to allow the subject to quickly generate new neural states for a target position while performing poorly, but minimize the influence of occasional aberrant responses during periods of strong performance.

2.2.2. Subsequent trials. After a history for each tone was built to establish the neural state for each target, prediction of the current position based on the present neural firing rate proceeded identically to Bayesian MLE classification based on *a priori* measurements. Defining the mean firing rate based on the history of n neurons with respect to the target position j as the vector $\mathbf{M}_j$:

$$\mathbf{M}_j = \{\mu_{j1},\ \mu_{j2}, \ldots, \mu_{jn}\}.$$

The real-time observed response R can be used to predict the probability of the center position and each target position by equation (4).

In order to start a trial, the observed response vector R generated by the subject was required to predict the NC state (center position, $p < 0.05$) for five consecutive bins (see figure 4, question 1). If the subject failed to start a trial, a 20 s recording was taken to establish a new mean vector for the baseline NC state. Upon starting a trial, a random target position was presented to the subject, and the subject was required to generate an observed response vector R that predicted the neural state for that target with a $p < 0.01$ for five consecutive bins to receive a food reward. The five 'correct' bins were then appended to the history for the target. Failing to accomplish this task, or generating an observed response vector R that predicted a target position other than the one presented, resulted in an incorrect trial (see figure 4, questions 2 and 3). After all incorrect trials, the response window was inspected for candidate neural states to be appended to the history for the target position (see figure 4, question 4).

2.2.3. Auditory variant of the standard center-out task. As Sprague-Dawley rats were used as test subjects in this study, an auditory BMI classification task was developed that parallels the standard center-out task in order to cater to the rat's heavy reliance on auditory cues. As noted earlier, in the standard center-out task subjects are trained to reach out and touch targets that have been lit. A center target is typically placed directly in front of the subject, and additional targets are placed equidistant from the center at regular angular intervals. In a given trial, first the center target is indicated, and the subject is required to touch the center target for 1 s. Next, one of the other targets is randomly indicated. For a successful trial, a subject is required to move their hand to touch the indicated target [21, 33]. In a virtual representation of this paradigm, the subject is given neural control of a cursor on a computer screen and tasked with moving the cursor on the *x*- and *y*-axis; first into a center circle on the screen and then into a target circle (see figure 3).

In our auditory classification task, the subject was given neural control of a tone which could be manipulated in frequency and intensity, which we refer to as an auditory cursor. The auditory cursor was represented by 200 ms tone pips, corresponding to a specific location on a two-dimensional axis, with frequency and intensity as the individual axes (see figure 3). In order to start a trial, subjects were required to hold the auditory cursor in the center of the two axes for 1 s, matching the cursor to the frequency and intensity selected for the center position. Each trial began with the random presentation of one of four possible target tones. The target tones were selected to be equidistant from the center of both axes, and spaced at regular angular intervals (see figure 3). After the presentation of the target tone, subjects were given a 10 s response window to match the auditory cursor to the presented target tone in order to obtain a food reward. If the subject failed to accomplish this task in 10 s, or matched the audio cursor to one of the other possible target tones, the trial ended and no reward was given.

2.2.4. Local field potentials. The basic MLE methodology presented here can easily be extended to include LFPs as a potential neural input source. After being relayed from the Multichannel Acquisition Processor System (Plexon, Inc., Dallas, TX), LFPs were additionally filtered in MATLAB using a fourth-order Butterworth filter with a passband between 10 and 40 Hz. The RMS voltage of the filtered LFP signal for each channel was calculated over non-overlapping 200 ms bins and used as the input to the MLE filter. The mean (μ) and standard deviation (σ) of the recorded RMS was then calculated for each channel. Similar to other LFP studies, the probability distribution for the RMS for each channel was assumed to be Gaussian [36, 37]. The previously determined RMS mean for each channel with respect to a specific stimulus s_j was defined as the vector $\mathbf{M}_j$:

$$\mathbf{M}_j = \{\mu_{j1},\ \mu_{j2}, \ldots, \mu_{jn}\},$$

where n is the total number of recorded LFP channels and μ_{j1} is the RMS mean for channel 1 during stimulus s_j.

Given a vector of observed RMS responses R from each channel where r_1 is the observed RMS from channel 1 of n channels:

$$R = \{r_1,\ r_2, \ldots, r_n\}.$$

Under the Gaussian assumption for RMS voltages, equation (3) becomes

$$P(R|s_j) = \frac{1}{\sqrt{2\pi}} \sum_{i=1}^{n} \frac{\mathrm{e}^{-(r_i-\mu_{ji})^2} \big/ \left(2\sigma_i^2\right)}{\sigma_i}. \tag{5}$$

Bayesian MLE analysis after this modification proceeded identically to single neuron analysis.

2.3. Behavioral testing

Six Sprague-Dawley rats were implanted with standard Michigan 4 × 4 arrays to test the MLE classifier under practical BMI conditions. Surgical procedures were identical to those outlined previously [38–40]. While five animals were implanted in motor cortex, one animal was implanted in cingulate cortex using coordinates 1.5–2.5 mm anterior to bregma, 0.3–0.7 mm lateral from bregma, and 1.6–2.5 mm deep from the surface of the brain [41]. The cingulate animal was used to assess the suitability of the naïve MLE classifier for use with cortical areas not known to have a linear relationship between neuronal firing rates and movement parameters.

2.3.1. Electrophysiology and behavior system. Units were sorted via a Multichannel Acquisition Processor System (Plexon, Inc., Dallas, TX), and spike times were relayed with nominal delays via TCP/IP to a Dell Dimension Computer (Dell, Inc., Austin, TX) that both analyzed the spike activity using in-house designed software (Mathworks, Inc., Natick, MA) and controlled the behavioral box (Coulborne Instruments, Inc., Allentown, PA). During each experimental session neural electrophysiological data from the 16 electrode channels sampled at 40 kHz were simultaneously amplified and bandpass filtered (450–5000 Hz). Manual spike sorting was conducted prior to each experimental session. Autocorrelograms were generated for each sorted unit and visually inspected for an obvious absolute and relative refractory period. Local field potential recordings from all 16 channels were initially bandpass filtered from 3 to 90 Hz and sampled at 1000 Hz. Auditory stimuli were delivered via a speaker (Yamaha NS-10M Studio, Yamaha Corporation, Buena Park, CA) located 35 cm directly above the test box. The system delivered a near-flat frequency response between 500 Hz and 32 kHz. The system was calibrated to a position at the food delivery tray, although calibration measurements indicated that the test box approximated a free field.

2.3.2. Behavioral training. All rats entering training were deprived to 85% of their free-feeding weight to provide motivation to receive a food reward. For each behavioral session, the rats were plugged into the headstage and commutator cables and placed into the behavioral box.

3. Results

3.1. In vivo performance

All six animals in this study were first started on a one-state detection task, where the same target tone was presented each trial. This task required the animal to generate a second neural state on command—different than the center position/NC state—after the target tone was presented. Figure 6 depicts a representative example of one animal progressively learning to generate a second distinct neural state in response to the presentation of a target. All six animals were able to use the bMLE classifier to complete this task at a 90% or greater level within three days of the first experimental run (see table 1). To complete this task, the animals developed individualized repeated motor behaviors (e.g. lifting of the front paws into the air, turning the head in a particular direction, etc) in response to the presentation of target tone 1. These stereotyped motor behaviors in turn generated a repeatable, separable, neural state sufficient to drive the auditory cursor.

After learning the one-state task, four of the six animals—three implanted in motor cortex, the other implanted in cingulate cortex—were advanced to a two-state discrimination task where one of two possible targets was randomly presented to the animal. The other two animals no longer had obvious unit activity on at least two electrode sites, and consequently were moved to the LFP task (section 3.3). Again, animals were typically able to complete the two-state task above forced

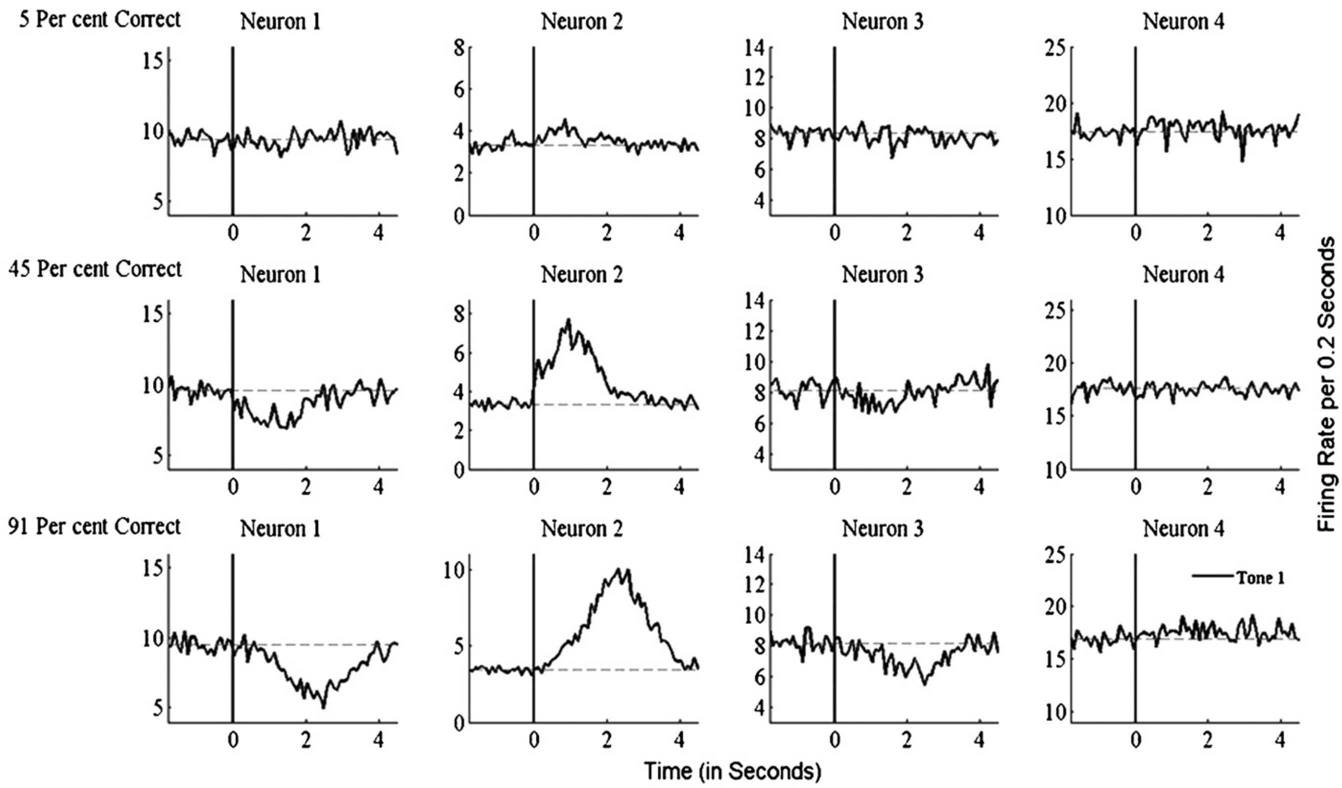

Figure 6. Building a distinct neural state for a target. Each row depicts the same four neurons recorded from one animal, across the same recording day. The dotted line denotes the mean firing rate for each neuron during the baseline window. The dark black line centered at 0 s indicates the presentation of a target tone. Row 1 depicts the mean firing rate averaged over the first 200 trials for the animal, where the subject successfully completed the one-state center-task on only 5% of the trials. Row 2 depicts the mean firing rate averaged over the second set of 200 trials for this animal, where 45% of the trials were successfully completed. Row 3 depicts the mean firing rate averaged over the third set of 200 trials for this animal, where 91% of the trials were successfully completed. Note that as performance improves, the animal progressively learns to decrease the firing rates of neurons 1 and 3, while increasing the firing rate of neuron 2, in response to the presentation of the target.

Table 1. Summary of *in vivo* results.

Neural input source	Number of target tones	Number of animals	Average day 1 performance	Average day 3 performance	Average day 7 performance	Forced choice random chance
Units	1	6, 5 Motor 1 cingulate	16.8 ± 7%	92.8 ± 5%[a]	N.A.	N.A.
	2	4, 3 Motor 1 cingulate	32.5 ± 11%	64 ± 10%[a,b]	73.3 ± 9%[a,b]	50%
	4	2, 1 Motor 1 cingulate	15 ± 8%	26 ± 8%	53 ± 13%[a,b]	25%
LEPs	1	6, 5 Motor 1 cingulate	67.4 ± 13%	75.2 ± 10%	83 ± 14%[a]	N.A.
	2	2, 1 Motor 1 cingulate	37.5 ± 9%	51 ± 9%[a]	65 ± 6%[a,b]	50%

[a] Significant improvement over day 1 performance ($p < 0.01$).
[b] Significant improvement over forced choice random chance ($p < 0.01$).

choice chance within three days of training (see table 1). Note that as this task was not a forced choice paradigm, but instead resulted in a failed trial when no choice was made, chance based on a forced choice paradigm should be considered a very conservative reference for comparison. Figure 7 depicts a peri-stimulus time histogram (PSTH) of an experimental run from a subject successfully completing this task well above forced choice chance. When first exposed to the task, the animal generated only one neural response state different than the center position/NC state, presumably a hold-over response from training with the one-state detection task. Hold-over behavior was evident every time an animal was switched to a more difficult/new task (see figures 7 and 8). After only three days of training (100 trial sessions, seven sessions per day), subjects were able to consistently perform the two-state discrimination above chance.

After reaching a two-state task performance level of greater than 80%, two of the animals were anesthetized for a 300 trial experimental run. In both cases, the anesthetized animal's performance level was under 2%. Under anesthetized conditions, the bMLE classifier rarely generated training data, and almost never deviated from the center position/NC condition.

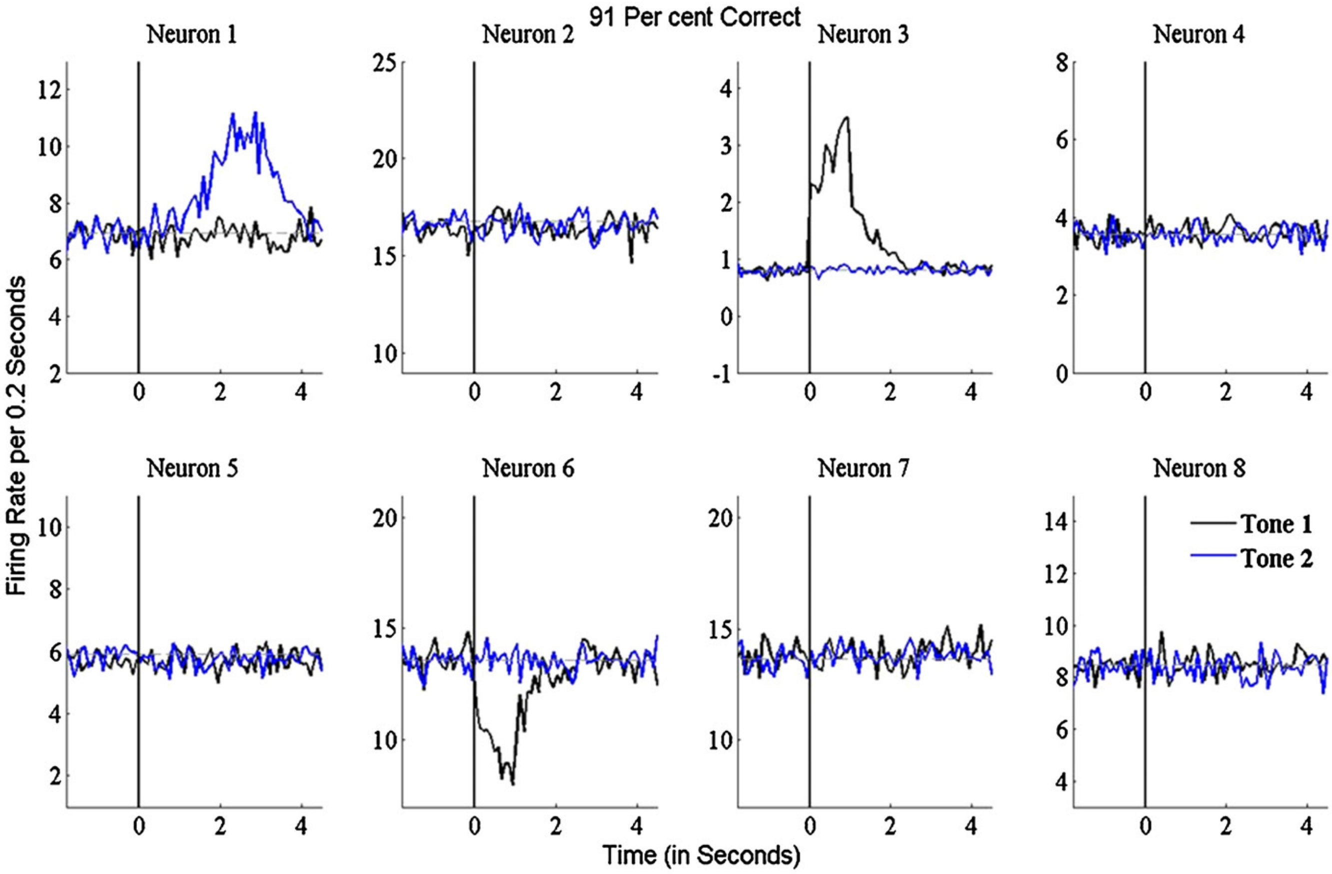

Figure 7. Representative example, two-state classification task. The dotted line denotes the mean firing rate for each neuron during the baseline (center position/NC) window. The dark black line centered at 0 s indicates the presentation of a target. Data set depicts the mean firing rate averaged over a 200 trial interval in which one subject (KMI 5) successfully completed 91% of trials (significant, $p < 0.001$). Note that for tone 1, the subject increases the firing rate of neuron 3 while decreasing the firing rate of neuron 6. For tone 2, the subject only increases the firing rate of neuron 1. In a forced choice paradigm, chance would be 50%.

After learning the two-state task, two of the four animals—one implanted in motor cortex, the other implanted in cingulate cortex—were moved to a four-state discrimination task, where four target tones were randomly presented to the animal. At this point, the other two animals had obvious unit activity on one or fewer channels, and consequently were moved to the LFP task (section 3.2). Both remaining subjects were able to successfully complete the four-state task significantly above forced choice chance after one week (see table 1). A full eight-state task was never attempted, due to the limitations of the rat model.

Figure 8 depicts PSTHs of an experimental run from a subject (KMI 5) successfully completing this task well above forced choice chance. Note that KMI 5 retained the same neural states for tones 1 and 2 as generated for the two-state task (see figure 7). However, for tone 3, KMI 5 increased the firing neuron 3 *without* decreasing the firing rate of neuron 6, therefore creating a neural state distinct from tone 2. For tone 4, the subject decreased the firing rate of neuron 7 slightly without changing the firing pattern for the other neurons. The performance of both animals when presented with tone 1 or tone 2 was superior to the subject performance for tones 3 and 4, presumably because of previous training.

3.2. Performance using local field potentials

As noted in other studies, recorded unit activity for animals in this study typically diminished substantially over the three weeks following surgery [2, 38, 39, 42]. When fewer than two well-isolated neuronal units were evident, animals were transitioned to a single-state classification task using their recorded LFPs as the input to the classifier. All six animals were able to generate a second neural state sufficient to complete the one-state task on the first day of LFP training (see table 1). Subjects typically retained the same stereotyped motor responses noted when performing the one-state unit task, which may explain how quickly the animals were able to generate viable LFP responses (see figure 10). Subjects tended to increase the RMS voltage across the entire array in response to a presented tone. Subject performance increased significantly ($p < 0.01$) over the first week of LFP training.

After learning the one-state LFP task, two animals—one implanted in motor cortex, the other implanted in cingulate cortex—were transitioned to a two-state LFP discrimination task. Both animals were able to successfully complete the two-state task significantly above chance for a two-state forced choice paradigm ($p< 0.01$, see figure 10). Again, as this task

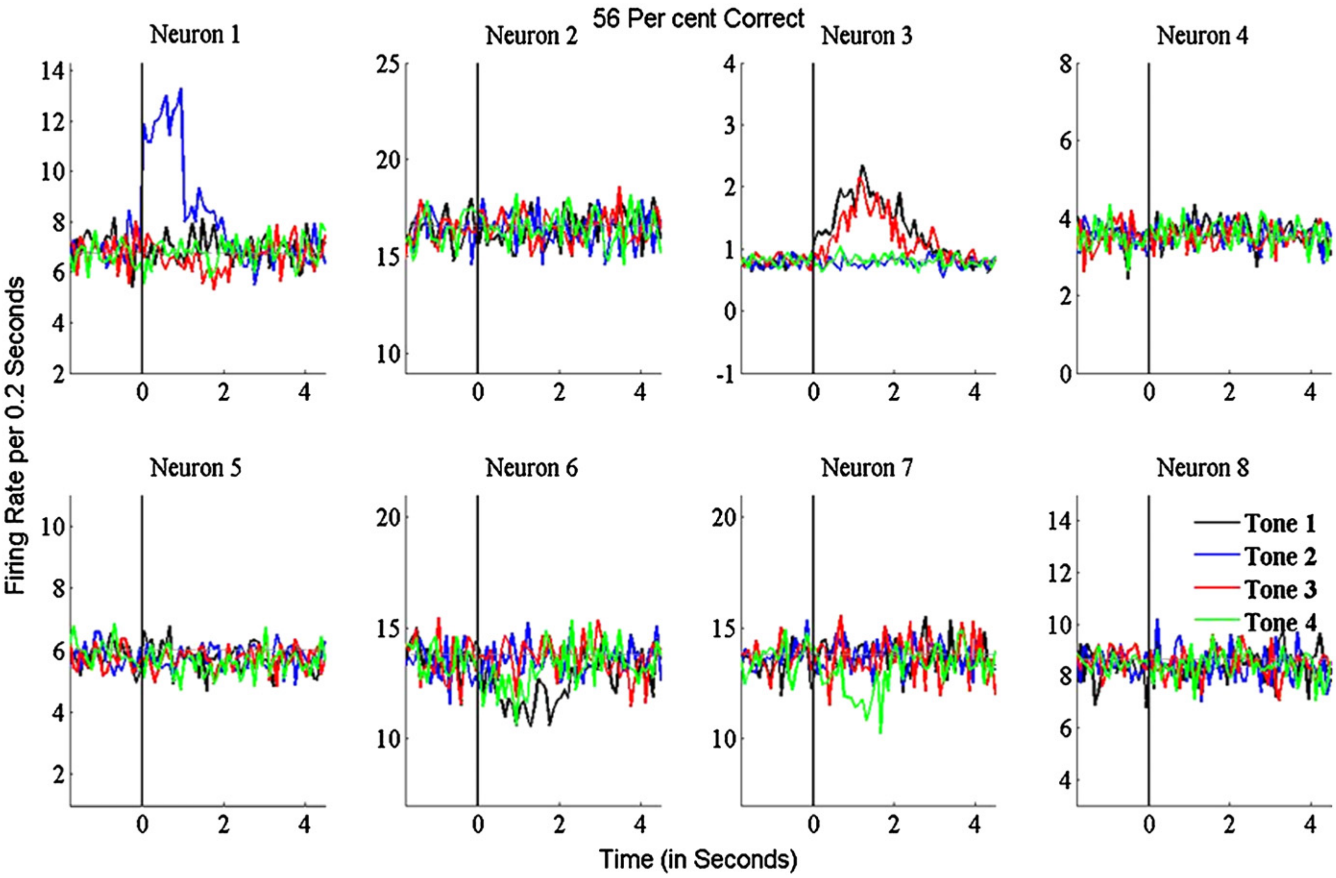

Figure 8. Representative example, four-state classification task. The dotted line denotes the mean firing rate for each neuron during the baseline window. The dark black line centered at 0 s indicates the presentation of a target tone. Data set depicts the mean firing rate averaged over a 200 trial interval in which one subject (KMI 5) successfully completed 56% of trials (significant, $p < 0.001$). Note that for tones 1 and 2, KMI 5 has retained approximately the same neural states as those adopted for the two-state center-out task (figure 7). However, for tone 3, the subject modulates neuron 3 *without* demodulating neuron 6, therefore creating a neural state distinct from tone 2. For tone 4, the subject demodulated neuron 7 slightly without changing the firing pattern for the other neurons. As subjects progressed from the one-state task to the four-state task, they typically retained neural states associated with previous tones, and added new neural states for additional tones. Per cent correct for each tone: tone 1 = 85%, tone 2 = 74%, tone 3 = 38%, tone 4 = 28%. In a forced choice paradigm, chance would be 25%.

was not a forced choice paradigm, but instead resulted in a failed trial when no choice was made within the time limits, forced choice chance should be considered a very conservative reference for comparison. As with the one-state LFP task, both subjects tended to increase the RMS voltage in the 10–40 Hz frequency band across the entire array in response to tone 1. In response to tone 2 presentation, however, both subjects increased the RMS voltage across only the upper half of the array.

4. Discussion

4.1. Animal behavior during task performance

In the process of developing new neural states, animals implanted in motor and cingulate cortices often developed specific stereotyped movements for each of the target tones. This result is not surprising, as the animal is probably learning a new movement/behavior for each target tone, which incidentally changes the firing rates of recorded neurons, as opposed to deliberately modulating neural activity. Behaviors associated with a specific target would often carry over as the number of targets was increased. In cases where the neurons recorded across the array changed dramatically from one day to the next, animals would often learn a new behavior for each target in order to generate the separate neural states necessary to complete the task.

4.2. Algorithm performance

Results from implanted animals utilizing our bMLE algorithm in real time to develop neural control indicate solid performance using very few neurons and very little training data. Animals well trained in a particular task typically required only one trial for each target to establish a useable neural state. Moreover, animals performing the two-state and four-state tasks used fewer than five total neurons to generate the neural states necessary to perform these tasks[1]. Finally, the speed at which animals were able to successfully perform

[1] The number of neurons actively engaged in the task, out of the total number of neurons isolated on a given day, was determined as described in section 2.2.1 step 4, 'refining the neuronal sample'.

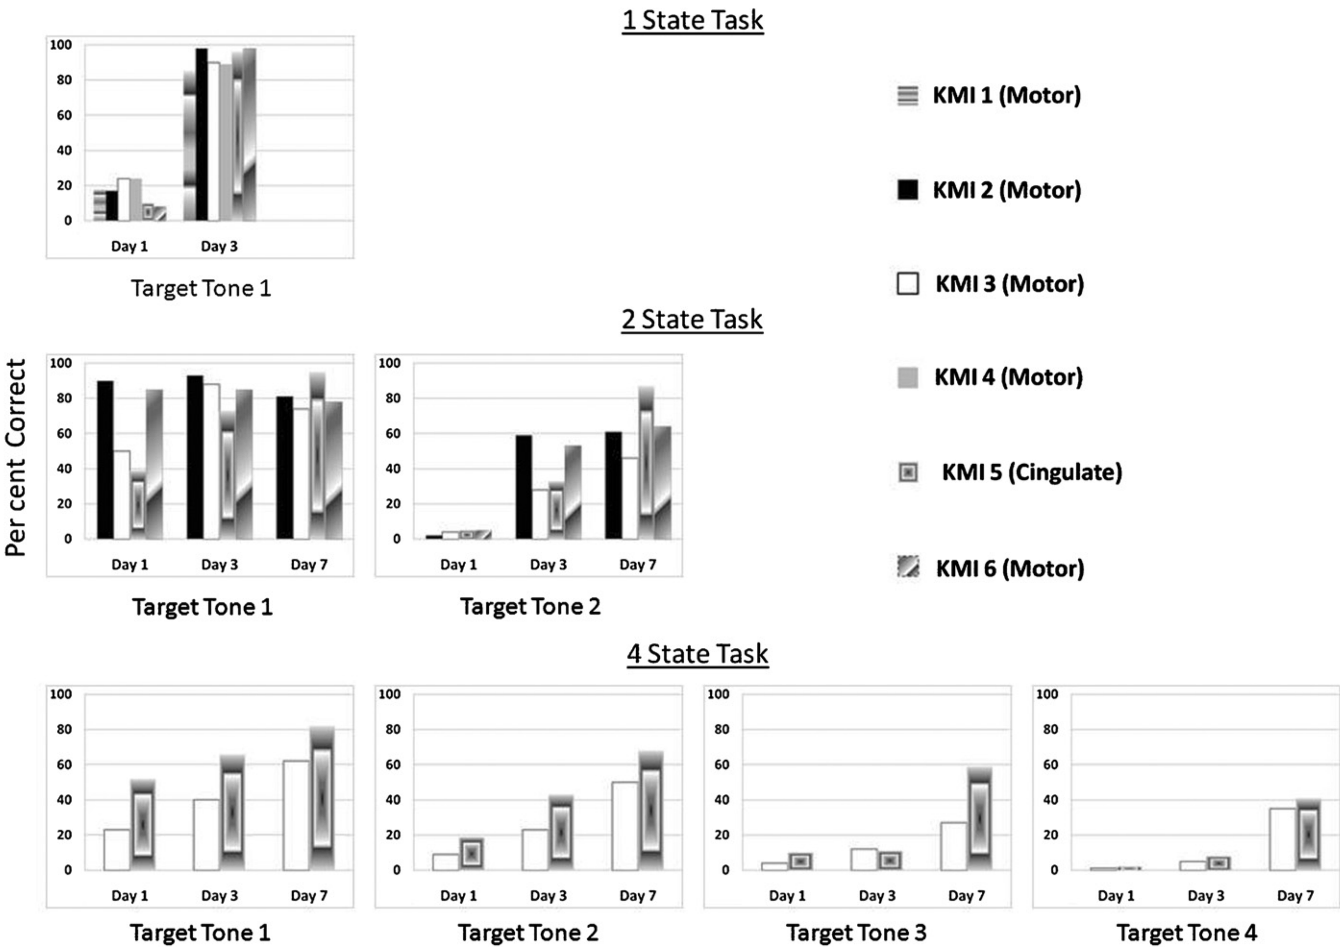

Figure 9. Performance improvement over time for each task. The y-axis for all subplots indicates the percentage of correct trials over 200 trials. The first row depicts each animal's performance improvement from day 1 to day 3 when performing the one target task ($n = 6$). Note that on day 1, the percentage of correct trials is only 16.8%. Over the next two days, all animals were able to greatly improve their ability to match the auditory cursor to a single target tone via neural control. The second row depicts the performance improvement of each animal from day 1 to day 7 when performing a two-state target identification task ($n = 4$). All animals performed well when presented with target tone 1 on the first day of the two-state task, presumably because of prior training in the one-state task. Over seven days, an obvious improvement in correctly matching the auditory cursor to target tone 2 is evident across all four animals. The third row depicts the trend in performance over the first week of the four-state classification task ($n = 2$). Both animals again exhibited some residual ability to match the neurally controlled auditory cursor to target tones presented in previous tasks.

new tasks above chance—typically within a week—suggests the self-paced bMLE decoding strategy presented in this study is easy to use.

4.3. Algorithm parameters

Algorithm parameters—window size, number of consecutive 'correct' bins, required degree of statistical separability of neural states, etc—were initially determined from the relevant literature, and then modified through systematic trial and error to optimize animal performance. Depending on specific requirements for a particular BMI, these parameters can be modified to tailor device performance. For example, the required p value to specify a particular target from a NC state can be decreased to reduce accidental movements. Similarly, the required percentage of variance vectors around a prospective neural class indicating a new target can be varied to decrease the probability of training on poor data. Window size can be decreased to increase information transfer rates, but at the expense of accuracy. In clinical applications, these parameters would need to be determined based off the preferences of the subject.

4.4. Limiting decoding assumptions

Only two assumptions were necessary for the bMLE classifier presented in this study to work: (1) the approximate time of IC, and (2) the probability distribution of the neural input source (Poisson or Gaussian for individual neurons and LFPs, respectively). In contrast, the linear mapping of neural firing rates to prosthetic control is dependent on a number of assumptions [43] that can be difficult to achieve in practice [8, 10, 13, 16, 22, 23, 37]. In order to generate an effective neural map, the recorded sample of neurons must be sufficient to fully define the movement parameter space [2, 10, 22, 23]. However, often only a small number of stable neuronal units can be recorded under chronic conditions due to technological limitations and the chronic immune response

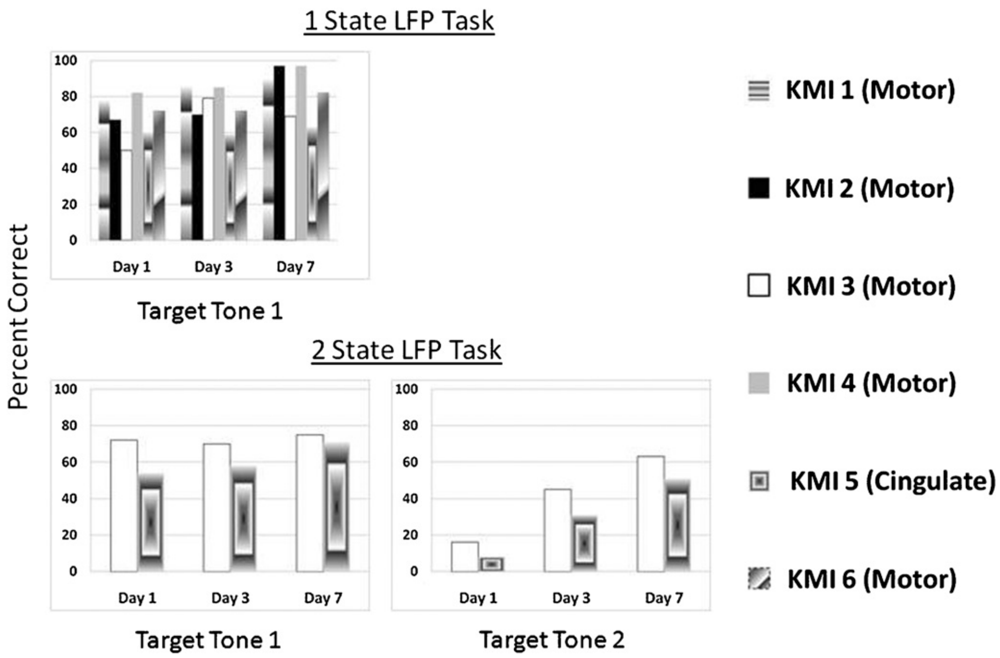

Figure 10. Performance improvement over time for each LFP task. The y-axis for all subplots indicates the percentage of correct trials over 200 trials. The first row depicts each animal's performance over the first week when performing the one-state LFP task ($n = 6$). All animals were immediately able to match the auditory cursor to a single target tone using the LFPs as a control signal. This result is not surprising; through the unit driven task the prior week the animals developed specific motor behaviors (e.g. lifting of the front paws into the air, turning the head in a particular direction, etc) in response to the presentation of target tone 1. These stereotyped motor behaviors in turn generated a repeatable, separable neural state sufficient to drive the auditory cursor. As the frequency and intensity of a target tone never changed from task to task, the animals tended to retain stereotyped motor behaviors with respect to familiar target tones—also generating a differentiable neuronal signal in the LFP recordings. The second row depicts the performance improvement of each animal from day 1 to day 7 when performing a two-state LFP task ($n = 2$). On the first day of the two-state task, both animals were able to match the auditory cursor to target tone 1 using LFPs as the control signal, as they had been trained on the one-state LFP task the week immediately prior. Conversely, target tone 2 had not been presented to the animals the previous week. Consequently, the animals did not perform well when target tone 2 was presented on day 1, but instead demonstrated a marked improvement in performance over the course of the first week.

[2, 10, 13, 22, 39, 44–50]. Moreover, the firing rates of many neurons recorded in motor cortex have no identifiable relationship to movement parameters [8, 16, 20], and the firing rates of neurons that are related to movement are only approximately linear during precisely defined stereotypical movements [10, 16, 23]. In addition, accurately isolating individual neurons can be difficult given low signal-to-noise recordings [10, 51, 52]. Consequently, multi-unit clusters are often used as the input for neural prosthetic applications, which can negatively affect the performance of linear decoding algorithms [8, 16].

Calibrating a neural control system based on previously determined relationships between neural firing rates and measured movement parameters may also ignore the subject's ability to adapt neural firing patterns in order to generate a greater number of usable output states. Studies which use decoding algorithms that can incorporate transient changes in neural ensemble firing patterns have indicated that animal subjects can learn to adapt their firing patterns in order to improve performance in neural control tasks [4, 6, 53–55]. In these studies, neurons which had little or no previous relationship to a known movement parameter often adapted over time to provide information relevant to completing the neural control task [4, 6, 53–55].

By limiting the initial assumptions used in our decoding algorithm, subjects were allowed to utilize neuronal sources without a known linear relationship to a movement parameter to create an output signal for neural control. Specifically, neurons from cingulate cortex and LFPs were effectively used as neuronal input sources to the bMLE algorithm in real time. In a previous study, rats were able to modulate single neurons from cingulate cortex on command [41]. By successfully completing a four-state classification task significantly above chance levels, the cingulate animal in this study demonstrated an even more refined neural control signal, demonstrating that cortical areas outside of motor cortex can be used to generate effective neural control. As an additional benefit, by limiting the initial decoding assumptions, the bMLE classification scheme presented in this study should be less dependent upon experimental environment and set-up for practical clinical use.

4.5. Local field potential recordings

Although prior studies have investigated the utility of the recorded LFPs to predict observed movements off-line [36, 37], real-time control of a neuronal output signal utilizing LFPs has only been demonstrated in one study [56, 57]. LFPs could represent the optimal input signal for BMI systems, as LFPs exhibit greater spatial and temporal resolution than electroencephalograms (EEGs) [5, 36, 37], and are presumed to be easier to obtain and more stable than single unit recordings for chronic implantations [5, 37]. Unlike action potentials, which can be differentiated from noise sources due to their characteristic waveforms [51], there is no universally accepted method to distinguish viable LFP signal from confounding noise sources.

In this study, animals were able to modulate the RMS voltage of LFPs in the 10–40 Hz band to generate a real-time control signal sufficient to complete the one-state and two-state center-out tasks above forced choice chance (see table 1). However, there is no way to determine if the animals were manipulating their LFPs through neurological mechanisms, or by inducing artifact into the recordings. As noted in many other studies [36, 37, 39, 42], obvious artifact was intermittently apparent during animal motion, which would on occasion create a perturbation in the RMS voltage of the LFP between 10 and 40 Hz. Not surprisingly, despite the self-paced architecture of the bMLE classifier, these noise perturbations could generate occasional unwanted control outputs. However, the specificity demonstrated by completing the two-state LFP task suggests that accidentally induced recording artifact was unlikely to be responsible for animal performance. Unlike other studies which removed periods of extensive noise during *post hoc* analysis [36, 37, 58], the constraints of real-time decoding preclude this option. In order to generate more refined real-time control with LFPs in a self-paced context, techniques drawn from EEG (electroencephalography) control tasks to isolate signal from noise should be applied in future experiments [3, 59].

5. Conclusions

In this study, we developed and tested a nonlinear bMLE classification scheme to identify statistically separable neural states evident in real-time neural recordings, and then mapped these states to generate self-paced neural control of a BMI suitable for accomplishing a useful task. By limiting initial decoding assumptions and training only on relevant neural data, accurate neural control was possible with five or fewer neurons, using minimal training data and no *a priori* movement measurements for calibration. In addition, we expanded our pool of possible neuronal input sources to include neurons in motor cortex without a detectable linear relationship to movement parameters, neurons from non-motor cortical areas, multi-unit clusters, and LFPs. This decoding scheme was sufficient to identify periods of NC in the data set, and neither generated training data nor control output during these periods. Finally, as the methodology proposed here is both simple and adaptive, this framework should be more resistant to the inherent instability of the neuronal input source resulting from the chronic immune response.

Acknowledgments

The authors of this paper would like to thank all of the present and past members of the Neural Engineering Laboratory for their help with editing this document and programming assistance, especially Dr Erin Purcell, Dr Greg Gage, Matt Gibson, Mike Joseph, Dr Takashi Kozai, Dr John Seymour, Dr Hirak Parikh, Dr Kevin Otto, and Dr Rio Vetter. This work was supported by the Center for Wireless Integrated Microsystems NSF EEC-9986866, the NIH P41 Center for Neural Communication Technology (EB002030), and the Whitaker Foundation.

References

[1] Chapin J K *et al* 1999 Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex *Nat. Neurosci.* **2** 664–70
[2] Schwartz A B 2004 Cortical neural prosthetics *Annu. Rev. Neurosci.* **27** 487–507
[3] Wolpaw J R *et al* 2002 Brain–computer interfaces for communication and control *Clin. Neurophysiol.* **113** 767–91
[4] Carmena J M *et al* 2003 Learning to control a brain–machine interface for reaching and grasping by primates *PLoS Biol.* **1** e42
[5] Anderson R A *et al* 2004 Cognitive neural prosthetics *Trends Cogn. Sci.* **8** 486–93
[6] Taylor D M, Helms-Tillery S I and Schwartz A B 2002 Direct cortical control of 3D neuroprosthetic devices *Science* **296** 1829–32
[7] Gage G J *et al* 2005 Naive coadaptive cortical control *J. Neural Eng.* **2** 52
[8] Wahnoun R, He J and Tillery S I H 2006 Selection and parameterization of cortical neurons for neuroprosthetic control *J. Neural Eng.* **3** 162
[9] Wessberg J *et al* 2000 Real-time prediction of hand trajectory by ensembles of cortical neurons in primates *Nature* **408** 361–5
[10] Wu W *et al* 2003 A switching Kalman filter model for the motor cortical coding of hand motion *Proc. IEEE Engineering in Medicine and Biology Society, 2003* pp 2083–6
[11] Bashashati A *et al* 2007 A survey of signal processing algorithms in brain computer interfaces based on electrical brain signals *J. Neural Eng.* **4** R32
[12] Velliste M *et al* 2008 Cortical control of a prosthetic arm for self-feeding *Nature* **453** 1098–101
[13] Hochberg L R *et al* 2006 Neuronal ensemble control of prosthetic devices by a human with tetraplegia *Nature* **442** 164
[14] Mason S G and Birch G E 2000 A brain-controlled switch for asynchronous control applications *IEEE Trans. Biomed. Eng.* **47** 1297–307
[15] Achtman N *et al* 2007 Free-paced high-performance brain–computer interfaces *J. Neural Eng.* **4** 336
[16] Lebedev M A and Nicolelis M A L 2006 Brain–machine interfaces: past, present and future *Trends Neurosci.* **29** 536–46
[17] DiGiovanna J *et al* 2009 Coadaptive brain–machine interface via reinforcement learning *IEEE Trans. Biomed. Eng.* **56** 54–64
[18] Digiovanna J *et al* 2010 Cyber-workstation for computational neuroscience *Front Neuroeng.* **2** 17
[19] Ward M P *et al* 2009 Toward a comparison of microelectrodes for acute and chronic recordings *Brain Res.* **1282** 183–200
[20] Georgopoulos A P, Schwartz A B and Kettner R E 1986 Neuronal population coding of movement direction *Science* **233** 1416–9
[21] Shenoy K V *et al* 2003 Neural prosthetic control signals from plan activity *Neuroreport* **14** 591–6
[22] Santhanam G *et al* 2006 A high-performance brain computer interface *Nature* **442** 195

[23] Kemere C, Shenoy K V and Meng T H 2004 Model-based neural decoding of reaching movements: a maximum likelihood approach *IEEE Trans. Biomed. Eng.* **51** 925–32
[24] Brockwell A E, Rojas A L and Kass R E 2004 Recursive Bayesian decoding of motor cortical signals by particle filtering *J. Neurophysiol.* **91** 1899–907
[25] Gao Y *et al* 2003 A quantitative comparison of linear and non-linear models of motor cortical activity for the encoding and decoding of arm motions *1st Int. IEEE/EMBS Conf. on Neural Engineering (Capri, Italy, 2003)* pp 189–92
[26] Kim K H, Kim S S and Kim S J 2006 Superiority of nonlinear mapping in decoding multiple single-unit neuronal spike trains: a simulation study *J. Neurosci. Methods* **150** 202
[27] Barbieri R *et al* 2005 An analysis of hippocampal spatio-temporal representations using a Bayesian algorithm for neural spike train decoding *IEEE Trans. Neural Syst. Rehabil. Eng.* **13** 131–6
[28] Geisler W S and Albrecht D G 1995 Bayesian analysis of identification performance in monkey visual cortex: nonlinear mechanisms and stimulus certainty *Vis. Res.* **35** 2723–30
[29] Lehky S R 2004 Bayesian estimation of stimulus responses in Poisson spike trains *Neural Comput.* **16** 1325–43
[30] Lehmkuhle M J, Normann R A and Maynard E M 2005 Trial-by-trial discrimination of three enantiomer pairs by neural ensembles in mammalian olfactory bulb *J. Neurophysiol.* **95** 1367–79
[31] Oram M W *et al* 1998 The 'Ideal Homunculus': decoding neural population signals *Trends Neurosci.* **21** 259–65
[32] Sanger T D 1996 Probability density estimation for the interpretation of neural population codes *J. Neurophysiol.* **76** 2790–3
[33] Georgopoulos A P *et al* 1982 On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex *J. Neurosci.* **2** 1527–37
[34] Lee D *et al* 1998 Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex *J. Neurosci.* **18** 1161–70
[35] Maynard E M *et al* 1999 Neuronal interactions improve cortical population coding of movement direction *J. Neurosci.* **19** 8083–93
[36] Mehring C *et al* 2003 Inference of hand movements from local field potentials in monkey motor cortex *Nat. Neurosci.* **6** 1253–4
[37] Rickert J *et al* 2005 Encoding of movement direction in different frequency ranges of motor cortical local field potentials *J. Neurosci.* **25** 8815–24
[38] Ludwig K A *et al* 2009 Using a common average reference to improve cortical neuron recordings from microelectrode arrays *J. Neurophysiol.* **101** 1679–89
[39] Ludwig K A *et al* 2006 Chronic neural recordings using silicon microelectrode arrays electrochemically deposited with a poly(3,4-ethylenedioxythiophene) (PEDOT) film *J. Neural Eng.* **3** 59
[40] Vetter R J *et al* 2004 Chronic neural recording using silicon-substrate microelectrode arrays implanted in cerebral cortex *IEEE Trans. Biomed. Eng.* **51** 896–904
[41] Marzullo T C, Miller C R and Kipke D R 2006 Suitability of the cingulate cortex for neural control *IEEE Trans. Neural Syst. Rehabil. Eng.* **14** 401–9
[42] Santhanam G *et al* 2007 HermesB: a continuous neural recording system for freely behaving primates *IEEE Trans. Biomed. Eng.* **54** 2037–50
[43] Kettner R E, Schwartz A B and Georgopoulos A P 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** 2938–47
[44] Nicolelis M A L *et al* 2003 Chronic, multisite, multielectrode recordings in macaque monkeys *Proc. Natl Acad. Sci. USA* **100** 11041–6
[45] Suner S, Fellows M R, Vargas-Irwin C, Nakata K and Donoghue J P 2005 Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex *IEEE Trans. Neural Syst. Rehabil. Eng.* **13** 524–41
[46] Vetter R J *et al* 2004 Spike recording performance of implanted chronic silicon-substrate microelectrode arrays in cerebral cortex *IEEE Trans. Neural Syst. Rehabil. Eng.* **52** 896–904
[47] Williams J, Rennaker R and Kipke D 1999 Long-term neural recording characteristics of wire microelectrode arrays implanted in cerebral cortex *Brain Res. Brain Res. Protoc.* **4** 303–13
[48] Rennaker R L *et al* 2007 Minocycline increases quality and longevity of chronic neural recordings *J. Neural Eng.* **4** L1
[49] Seymour J P and Kipke D R 2007 Neural probe design for reduced tissue encapsulation in CNS *Biomaterials* **28** 3594–607
[50] Ludwig K A *et al* 2011 Poly(3,4-ethylenedioxythiophene) (PEDOT) polymer coatings facilitate smaller neural recording electrodes *J. Neural Eng.* **8** 014001
[51] Lewicki M S 1998 A review of methods for spike sorting: the detection and classification of neural action potentials *Netw. Comput. Neural Syst.* **9** R53–78
[52] Wood F *et al* 2004 On the variability of manual spike sorting *IEEE Trans. Biomed. Eng.* **51** 912–8
[53] Kennedy P R *et al* 2000 Direct control of a computer from the human central nervous system *IEEE Trans. Rehabil. Eng.* **8** 198–202
[54] Serruya M D *et al* 2002 Instant neural control of a movement signal *Nature* **416** 141–2
[55] Moritz C T, Perlmutter S I and Fetz E E 2008 Direct control of paralysed muscles by cortical neurons *Nature* **456** 639–42
[56] Kennedy P *et al* 2004 Using human extra-cortical local field potentials to control a switch *J. Neural Eng.* **1** 72–7
[57] Kennedy P R *et al* 2004 Computer control using human intracortical local field potentials *IEEE Trans. Neural Syst. Rehabil. Eng.* **12** 339–44
[58] Magill P J *et al* 2004 Synchronous unit activity and local field potentials evoked in the subthalamic nucleus by cortical stimulation *J. Neurophysiol.* **92** 700–14
[59] Cooper R *et al* 2003 EEG, paediatric neurophysiology, special techniques and applications *Clinical Neurophysiology* vol 2 ed R Cooper *et al* (Amsterdam: Elsevier) pp 8–103