Published in final edited form as:

J Neural Eng. 2014 April ; 11(2): 026001. doi:10.1088/1741-2560/11/2/026001.

Self-recalibrating classifiers for intracortical brain-computer interfaces

William Bishop^{1,4}, Cynthia C Chestek^{5,11,13}, Vikash Gilja^{6,11}, Paul Nuyujukian^{7,10}, Justin D Foster⁵, Stephen I Ryu^{5,12}, Krishna V Shenoy^{5,7,8,9,11}, and Byron M Yu^{2,3,4}

Byron M Yu: byronyu@cmu.edu

¹Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA 15213

²Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213

³Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213

⁴Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213

⁵Department of Electrical Engineering, Stanford University, Stanford, CA 94305

⁶Department of Computer Science, Stanford University, Stanford, CA 94305

⁷Department of Bioengineering, Stanford University, Stanford, CA 94305

⁸Department of Neurobiology, Stanford University, Stanford, CA 94305

⁹Neurosciences Program, Stanford University, Stanford, CA 94305

¹⁰School of Medicine, Stanford University, Stanford, CA 94305

¹¹Stanford Institute for NeuroInnovation and Translational Neuroscience, Stanford University, Stanford, CA 94305

¹²Palo Alto Medical Foundation, Palo Alto, CA 94301

¹³Department of Biomedical Engineering, University of Michigan, Ann Arbor, MI 48109

Abstract

Objective—Intracortical brain-computer interface (BCI) decoders are typically retrained daily to maintain stable performance. Self-recalibrating decoders aim to remove the burden this may present in the clinic by training themselves autonomously during normal use but have only been developed for continuous control. Here we address the problem for discrete decoding (classifiers).

Approach—We recorded threshold crossings from 96-electrode arrays implanted in the motor cortex of two rhesus macaques performing center-out reaches in 7 directions over 41 and 36 separate days spanning 48 and 58 days in total for offline analysis.

Main results—We show that for the purposes of developing a self-recalibrating classifier, tuning parameters can be considered as fixed within days and that parameters on the same electrode move up and down together between days. Further, drift is constrained across time, which is reflected in

the performance of a standard classifier which does not progressively worsen if it is not retrained daily, though overall performance is reduced by more than 10% compared to a daily retrained classifier. Two novel self-recalibrating classifiers produce a ~15% increase in classification accuracy over that achieved by the non-retrained classifier to nearly recover the performance of the daily retrained classifier.

Significance—We believe that the development of classifiers that require no daily retraining will accelerate the clinical translation of BCI systems. Future work should test these results in a closed loop setting.

1. Introduction

Brain-computer interfaces (BCIs) aim to assist disabled patients by translating recorded neural activity into control signals for assistive devices. Intracortical BCI systems rely upon arrays of chronically implanted microelectrodes which penetrate into the cortex to record the activity of tens to hundreds of neurons. The last decade has witnessed a substantial investment of resources by the community into developing decoders for intracortical BCI which are more accurate, both in offline simulations with prerecorded data and in online, closed-loop settings. However, in almost all cases these decoders are retrained daily in a supervised manner to maintain their accuracy (e.g., Musallam et al. (2004); Santhanam et al. (2006); Velliste et al. (2008); Li et al. (2009); Hochberg et al. (2012); Gilja et al. (2012); Collinger et al. (2012)). This retraining typically requires that the subject perform a series of trials where the goal of each trial, such as desired reach target, is known so labelled data for training can be collected.

While retraining procedures may only require minutes of a subject's time, future users may find daily repetition of a task which stands between them and device use burdensome. Indeed, imagine the annoyance modern cell phone users would experience if they had to calibrate the touch screens of their devices everyday before use. A few minutes of calibration would be a small price to pay for the connectivity these devices provide, but it is difficult to imagine the widespread adoption of modern phones under these conditions. In the same way, we suggest that modern BCI systems should aim to achieve predictable performance each day without undue burden to users.

An alternative to retraining in a supervised manner is to employ self-recalibrating neural decoders which can simultaneously estimate a user's intent along with changes in neural tuning parameters. These decoders, generally referred to as adaptive decoders, continuously train themselves during normal BCI use without requiring knowledge of a user's true intended actions. In this way, the retraining period at the beginning of each day can be eliminated. There have been relatively few studies on self-recalibrating decoders for intracortical BCI, and the work that does exist has focused on BCI decoders which decode a continuously valued control signal, such as the position of a computer cursor or prosthetic limb (Eden et al., 2004a,b; Srinivasan et al., 2007; Li et al., 2011).

While continuous decoders have many important applications, classifiers of discrete actions are an important part of existing real-world BCI systems (Ohnishi et al., 2007). Potential clinical applications for BCI classifiers include making selections from a menu interface or

virtual keyboard (e.g., Scherer et al. (2004); Santhanam et al. (2006); Brunner et al. (2011)), selecting among grasp types for a prosthetic hand (e.g., Pistohl et al. (2012); Chestek et al. (2013)), classifying finger movements (e.g., Aggarwal et al. (2008); Wang et al. (2009)), identifying periods of planning and movement (e.g., Shenoy et al. (2003); Achtman et al. (2007); Kemere et al. (2008); Wang et al. (2012)) and providing discrete control signals for a wheelchair (e.g., Leeb et al. (2007); Galán et al. (2008)). In some settings it may be possible to interface with a device using either a continuous or discrete control signal. For example, when interfacing with a menu, a cursor can be moved over items to be selected or a discrete decoder can be used to make a direct menu selection. In certain cases, the use of a classifier may provide a substantial performance improvement to an end user. Indeed, we have demonstrated a discrete decoder which transmits the user's intent at a rate of up to 6.5 bits per second (~15 words per minute), which is faster than continuously guiding a cursor to make discrete selections (Santhanam et al., 2006).

Despite the potential clinical benefit that intracortical BCI classifiers may provide, the subject of self-recalibrating classifiers has received little attention in the intracortical community. In contrast, there is previous work in the EEG community which characterizes signal drift (e.g., Shenoy et al. (2006); Pascual et al. (2011)), and proposes self-recalibrating classifiers. Example algorithms include approaches based on variational Kalman filtering (Sykacek and Roberts, 2002; Sykacek et al., 2004), expectation maximization (Blumberg et al., 2007) and a variety of other methods to update the parameters of an LDA classifier (Vidaurre et al., 2008; Vidaurre and Blankertz, 2010). However, many of these algorithms consider only binary classification, and it is not clear the complexity of some (e.g., Sykacek and Roberts (2002); Sykacek et al. (2004); Blumberg et al. (2007)) is required for BCI application.

In the present work, we seek to develop self-recalibrating classifiers for intracortical BCI. We note this differs from the previous work of Santhanam et al. (Santhanam et al., 2009). Santhanam et al. develop classifiers with improved within-day accuracy by accounting for unobserved influences on neural activity, while we focus on achieving stable across-day performance through self-recalibration. We chose to forgo spike sorting in favor of the use of threshold crossings. That is we mark any voltage event which crosses a set threshold on each electrode as a spike. In one monkey, we also enforce that threshold events must pass a shape heuristic. Previous work has shown only minor reductions in performance when threshold crossings are used in place of spike sorting (Fraser et al., 2009; Chestek et al., 2011), and this removes the need to track and identify units across days.

Having made this design decision, we study how firing rates drift within and between recording sessions on different electrodes. This enables us to then develop self-recalibrating classifiers which take advantage of structure we find in this drift. Along the way, we also examine the performance of a standard BCI classifier that undergoes no daily retraining. In addition to being a novel result in itself, this provides a baseline when assessing the performance of the self-recalibrating classifiers presented in this work.

2. METHODS

Reach task and neural recordings

Animal protocols were approved by the Stanford University Institutional Animal Care and Use Committee. We trained two adult male monkeys (*Macaca mulatta*, monkeys I and L) to perform center-out reaches for liquid rewards. As illustrated in Fig. 1, visual targets were back-projected onto a vertical screen in front of the monkey. A trial began when the monkey touched a central yellow square. Following a randomized (300 – 500 ms) hold period, a visual reach target appeared. After a brief (35 ms) delay, the animal was permitted to reach. Reach targets could appear at one of twenty-eight locations. Each of the twenty-eight targets was typically located at a distance of 4.5 cm, 7.0 cm, 9.5 cm, or 12 cm from the central target and at an angle of 40°, 85°, 130°, 175°, 220°, 310°, or 355° from horizontal so that the targets were arranged in 4 concentric rings of 7 targets each. For all analyses in this paper, we collapse across reach distance and only consider the direction of reach on each trial. After the reach target was held for 200 ms for monkey L and 300 ms for monkey I, a liquid reward was delivered.

Monkeys sat in a custom chair (Crist Instruments, Hagerstown, MD) with the head braced and nonreaching arm comfortably restrained. A 96-channel silicon electrode array (Blackrock Microsystems, Salt Lake City, UT) was implanted in motor cortex contralateral to the reaching arm. When identifying neural spikes on each electrode, threshold crossings (Chestek et al., 2009; Fraser et al., 2009) were used in place of spike sorting. For monkey L, we also required voltage waveforms to pass a shape heuristic. Threshold levels for each electrode were set at the beginning of each recording day with an automated method and then, with possible rare exception, held fixed for the remainder of data collection. The automated method of setting threshold levels consisted of collecting 2 seconds of filtered voltage data sampled at 30-kHz from each electrode and calculating a root-mean-square (RMS) like value from this data. Thresholds levels were then set at three times this value for each electrode.

Recordings were made on a total of 41 (monkey L) and 36 (monkey I) days spanning a period of 48 (monkey L) and 58 (monkey I) days in total. In all of the work that follows, we limit ourselves to only analyzing successful trials. In the data analyzed, there were (mean $\pm$ s.d.) 1737 $\pm$ 482 (monkey L) and 1688 $\pm$ 368 (monkey I) successful trials per day spanning 53 $\pm$ 16 (monkey L) and 88 $\pm$ 16 (monkey I) minutes per day.

2.1. Mathematical Notation

We aim to develop BCI classifiers which maintain stable day-to-day performance without any daily supervised retraining. For clarity, we will speak of developing classifiers within the context of the datasets used in this study, but the techniques presented here should generalize to other intracortical BCI classification scenarios.

Throughout this work, we rely on probabilistic methods to infer the intended reach direction from neural data. For trial t on day d , we form a feature vector, $\mathbf{x}_{d,t} \in \mathbb{N}^{E \times 1}$, from our neural data by counting the number of threshold crossings on each of E electrodes of the recording array. The notation $x_{d,e,t}$ will refer to the specific spike count for electrode $e^\dagger$. Similarly, $y_{d,t}$

$\in \{1, \dots, J\}$ indicates which one of J directions a reach is made in for a specific trial and day. Throughout this work we follow the convention of using bold script exclusively to denote vectors. The symbol P will be used to indicate probability distributions.

On any given trial there is a probabilistic model relating observed spike counts to reach direction. We model this with the distribution $P(\mathbf{x}_{d,t} | \mathbf{y}_{d,t}; \Theta_{d,t})$, where $\Theta_{d,t}$ is a set of parameters characterizing the distribution. In this work for electrode e , these parameters consist of the means $\mu_{e,j}$ and variances $\nu_{e,j}$ of spike counts when reaches are made in each direction j . We will refer to this set of parameters as *tuning parameters*. While both means and variances can drift in principle, we focus on achieving stable classification through tracking changes in mean spike counts. We term the set of $\mu_{e,j}$ parameters for a fixed reach direction j as class means, as they are mean counts conditioned on a class of movement. We can index class means across days and trials as $\mu_{d,e,j,t}$, where d indexes day and t indexes a trial within a day. In some of the work that follows, we will speak of class means across all trials of a day, which we will denote as $\mu_{d,e,j}$.

2.2. Tracking Drifting Tuning Parameters Across Time

Movements of the recording array relative to the brain (Santhanam et al., 2007), scar tissue buildup (Polikov et al., 2005), behavioral changes (Chestek et al., 2007) and neural plasticity (Taylor et al., 2002; Jarosiewicz et al., 2008; Ganguly and Carmena, 2009; Chase et al., 2012) may all cause tuning parameters to drift. In this section, we track drift in class means across days and time by convolving spike counts for all reach directions in a particular direction with a Gaussian kernel. This is a simple way of smoothing spike counts to estimate tuning parameters when true reach direction is known on each trial. This is an appropriate assumption for this analysis which aims to establish basic properties of tuning parameter drift. However, the techniques introduced in this section should not be confused with the methods that are introduced later to track drifting tuning parameters by our self-recalibrating classifiers when knowledge of true reach direction is not assumed known. Mathematically, estimates of tuning parameters are formed through convolution with a kernel as

$$\hat{\mu}_{d,e,j,t} = \frac{1}{C} \sum_{s: y_{d,s}=j} K(t-s; w) x_{d,e,s}, \quad (1)$$

where $K(t-s; w)$ is Gaussian density function with 0 mean and standard deviation w and $C = \sum_{s: y_{d,s}=j} K(t-s; w)$ is a normalizing constant. Mathematically, the function K indexes over trial number, and therefore, $t-s$ in the argument is the distance in *number of trials* between two trials and not time. There remains the challenge of selecting the proper value of the width, w , of the Gaussian density function. For the purposes of this paper, we will present results for a single value of w and then show in Supplementary Results that the qualitative results we obtain are robust as the value of w is varied.

We are interested in comparing the amount of drift between and within days. If we smooth only within days, we enforce smoothing within but not across days and we may bias our

[‡]Throughout this paper we will use the convention of listing subscripts in alphabetical order.

results to show larger between than within day drift. To account for this, we ignore day boundaries and smooth across all trials for a given electrode and reach direction as if they were on the same day. This may tend to “over-smooth” between days, but to foreshadow what is to come, we would like to argue that within-day variation in tuning parameters is small relative to between-day variation. Therefore, this procedure is conservative for our purposes.

2.3. A Standard BCI Classifier

We desire to examine the impact of tuning parameter drift on a standard BCI classifier. We select a Gaussian, naïve Bayes decoder as our standard classifier, which has been used for BCI classification before (Santhanam et al., 2006; Maynard et al., 1999). Briefly, this classifier models spike counts on electrodes as independent Gaussian distributions conditioned on reach direction. The notation $\mu_{e,j}$ and $v_{e,j}$ refer to the mean and variance of the distribution for electrode e and reach direction j . For notational convenience, we will group all the means and variances for a single electrode e into the vectors $\boldsymbol{\mu}_e \in \mathbb{R}^{J \times 1}$ and $\boldsymbol{v}_e \in \mathbb{R}^{J \times 1}$, respectively. The joint probability of observing spike counts for all electrodes can be computed as the product of probabilities for individual electrodes. Formally, we write this as

$$P(\mathbf{x}_{d,t} | y_{d,t}) = \prod_{e=1}^E P(x_{d,e,t} | y_{d,t}). \quad (2)$$

A uniform prior probability of reaching in any direction, $P(y_{d,t}) = 1/J$, is employed and Bayes’ rule can then be used to find the posterior probability of reaching in direction j given neural data, $P(y_{d,t} = j | \mathbf{x}_{d,t})$. This posterior is calculated for each reach direction, and the direction with the highest probability is selected as the decoded direction for each trial.

In this work we will examine the performance of two versions of this standard classifier, which differ only in the way they are trained. The first version, referred to as the standard retrained classifier, learns the $\boldsymbol{\mu}_e$ and $\boldsymbol{v}_e$ parameters fresh with trials from the beginning of each day on which classification will be performed. This emulates standard practice in the BCI field. The second version, referred to as the standard non-retrained classifier, is trained initially on all trials collected on a small number of days in which neural data and desired reach direction is known. After this training, the parameters of the standard non-retrained classifier are held fixed, and it is then used to predict reach direction for trials on subsequent days without any type of retraining or recalibration.

2.4. A Probabilistic Self-Recalibrating Classifier

We now describe the intuition and probabilistic model behind our two self-retraining classifiers. The first version, which we refer to as the *SR* classifier, takes a principled probabilistic approach which provides insight into the simplified self-recalibrating (*SR_S*) classifier we subsequently present. Readers interested only in the *SR_S* classifier may skip to section 2.5 after reading the part of this section outlining the probabilistic model common to both classifiers without loss of understanding.

The *SR* classifier can be considered an augmented standard classifier. It continues to assume tuning parameters are fixed within a day, but it no longer assumes class means are known at the beginning of each decoding day. Mathematically, it achieves this by introducing a distribution over class mean values that a decoding day begins with. As a decoding day progresses, the *SR* classifier produces refined distributions over class means while simultaneously estimating reach direction on a trial-by-trial basis.

A formal description of probabilistic model for the *SR* classifier follows. When explicitly referring to the class mean for electrode e , reach direction j on day d , we will use the notation $\mu_{d,e,j}$. In results we show that the $\mu_{d,e,j}$ parameters for the same electrode drift from day to day in a highly correlated manner. Intuitively, this means that the average spike counts observed for two different reach directions on the same electrode tend to move up and down together across days. We capture this intuition by referencing all class means for electrode e to the same base value $b_{d,e}$ for day d , and we allow this base value to be drawn from a Gaussian distribution independently on each day. Mathematically,

$$b_{d,e} \sim \mathcal{N}(m_e, s_e), \quad (3)$$

where m_e is the mean and s_e is the variance of the distribution from which base values are drawn. For notational convenience we will group all the base values for all electrodes for day d into a vector $\mathbf{b}_d \in \mathbb{R}^{E \times 1}$. We model the base values for different electrodes as independent, and the prior probability for the set of base values for all electrodes is therefore given by

$$P(\mathbf{b}_d) = \prod_{e=1}^E P(b_{d,e}). \quad (4)$$

We relate the class mean for reach direction j for an electrode to its base value through an offset, $o_{e,j}$, so that

$$\mu_{d,e,j} = o_{e,j} + b_{d,e}. \quad (5)$$

The offset, $o_{e,j}$, is fixed across days. Class means for a single electrode can then vary from day to day but they must move up and down together. Having defined a probabilistic model for $b_{d,e}$, we specify an observation model for a single electrode by defining the probability of a spike count on electrode e for trial t on day d given $b_{d,e}$ and reach direction j as

$$x_{d,e,t} | b_{d,e}, y_{d,t} = j \sim \mathcal{N}(o_{e,j} + b_{d,e}, v_{e,j}). \quad (6)$$

This observation model is almost identical to that of the standard classifier with one important exception: spike counts now depend on the base value, which is a random variable that can vary from day to day. To specify the complete observation model, we again model electrodes as conditionally independent given reaching direction so that

$$P(\mathbf{x}_{d,t}|\mathbf{b}_d, y_{d,t}) = \prod_{e=1}^E P(x_{d,e,t}|\mathbf{b}_{d,e}, y_{d,t}). \quad (7)$$

To complete the specification of the probabilistic model for the self-recalibrating classifier, we define the prior probability of reaching in any direction as uniform, exactly as specified for the standard model.

2.4.1. Decoding with the Self-Recalibrating Classifier—We now show how to decode reach direction while simultaneously refining knowledge of class means with the *SR* classifier. For each trial, we seek to find two distributions. The first is $P(y_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^t)$, the posterior distribution over reach directions. The most likely direction from this distribution is selected as the decoded reach direction on each trial. The second distribution is $P(\mathbf{b}_d|\{\mathbf{x}_{d,s}\}_{s=1}^t)$, which embodies the decoder's knowledge of the base values for a day given the neural activity that it has seen.

As a decoding day progresses, the decoder's knowledge of the true base values for a day, as represented in the distribution $P(\mathbf{b}_d|\{\mathbf{x}_{d,s}\}_{s=1}^t)$, becomes more refined. This is possible even though the decoder is never provided with knowledge of the user's true reach directions. The key intuition is that spike counts will tend to cluster around class means for a day, and this structure can be exploited by our classifier, as illustrated for a one electrode and two reach direction decoding scenario in Fig. 2.

We now describe the formal decoding procedure for arriving at the distributions $P(y_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^t)$ and $P(\mathbf{b}_d|\{\mathbf{x}_{d,s}\}_{s=1}^t)$ for each trial. The algorithm is recursive and will use the results from the previous trial, the distributions $P(y_{d,t-1}|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ and $P(\mathbf{b}_d|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$, as its starting point in its calculations for trial t . When decoding the first trial, the prior distribution on reach directions, $P(y_{d,t})$, is used in place of $P(y_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ and the prior distribution on base values, $P(\mathbf{b}_d)$, is used in place of $P(\mathbf{b}_d|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$.

Calculating $P(y_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^t)$. For trial t , we calculate $P(y_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^t)$ using Bayes' rule as

$$\begin{aligned} P(y_{d,t}=j|\{\mathbf{x}_{d,s}\}_{s=1}^t) &= \frac{P(\mathbf{x}_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1}, y_{d,t}=j)P(y_{d,t}=j|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1})}{P(\mathbf{x}_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1})} \\ &= \frac{c_{t|j}P(y_{d,t})}{c_t}, \end{aligned} \quad (8)$$

where in moving to the second line we have defined $c_{t|j} = P(\mathbf{x}_{d,t}|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1}, y_{d,t}=j)$, made use of the fact that reach direction is independent of neural activity so that

$P(y_{d,t}=j|\{\mathbf{x}_{d,s}\}_{s=1}^{t-1}) = P(y_{d,t})$, and recognized that the denominator of the first line is a normalizing constant calculated as $c_t = \sum_{j=1}^J c_{t|j}P(y_{d,t})$.

The scalar $c_{t|j}$ is the probability of the neural activity for trial t given all past neural activity seen during the day if the subject reaches in direction j . We will soon describe a means of calculating its value, but first note that it appears in the calculations for another distribution, $P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t, y_{d,t}=j)$, which we will need shortly. This distribution gives the probability of base values conditioned on all observed neural activity if the subject were to reach in direction j for the current trial. To calculate it, we again use Bayes' rule to write

$$P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t, y_{d,t}=j) = \frac{P(\mathbf{x}_{d,t} | \mathbf{b}_d, y_{d,t}=j) P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})}{c_{t|j}}, \quad (9)$$

where $P(\mathbf{x}_{d,t} | \mathbf{b}_d, y_{d,t}=j)$ is given by Eq. 7 and $P(\mathbf{b}_d | \{\mathbf{x}_s\}_{s=1}^{t-1})$ is the posterior over base values obtained when decoding the previous trial. In forming the numerator of Eq. 9, we have used the fact that $\mathbf{x}_{d,t}$ conditioned on $\mathbf{b}_d$ is independent of all previous neural observations and that base value is independent of reach direction. We see that the numerator on the right hand side of Eq. 9 is in the form of the product of two Gaussians.

$P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ is defined to be a Gaussian distribution with mean $\boldsymbol{\mu}_{t-1}$ and covariance matrix $\sum_{t-1}^{-1}$. Viewed as a function of $\mathbf{b}_d$, $P(\mathbf{x}_{d,t} | \mathbf{b}_d, y_{d,t}=j)$ is Gaussian with mean $\mathbf{x}_{d,t} - \mathbf{o}_j$, where $\mathbf{o}_j \in \mathbb{R}^{E \times 1}$ is a vector of offsets for each electrode for reach direction j and covariance matrix $V_j \in \mathbb{R}^{E \times E}$, where V_j is a diagonal matrix with the values of $v_{e,j}$ along its diagonal. We can then calculate the mean, $\mathbf{m}_{t|j}$, and covariance matrix, $\sum_{t|j}$, of

$P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t, y_{d,t}=j)$ of the left hand side of Eq. 7 as

$$\sum_{t|j} = (V_j^{-1} + \sum_{t-1}^{-1})^{-1} \quad (10)$$

$$\mathbf{m}_{t|j} = \sum_{t|j} (V_j^{-1}(\mathbf{x}_{d,t} - \mathbf{o}_j) + \sum_{t-1}^{-1} \boldsymbol{\mu}_{t-1}), \quad (11)$$

where the notation $\cdot^{-1}$ indicates matrix inversion. The scalar $c_{t|j}$ is then the integral over the numerator of Eq. 9 and can be calculated as

$$c_{t|j} = \int P(\mathbf{x}_{d,t} | \mathbf{b}_d, y_{d,t}=j) P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1}) d\mathbf{b}_d \\ = \sqrt{(2\pi)^E \det(\sum_{t|j})} P(\mathbf{x}_{d,t} | \mathbf{b}_d = \mathbf{m}_{t|j}, y_{d,t}=j) P(\mathbf{b}_d = \mathbf{m}_{t|j} | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1}), \quad (12)$$

where $P(\mathbf{x}_{d,t} | \mathbf{b}_d = \mathbf{m}_{t|j}, y_{d,t}=j)$ is obtained by evaluating Eq. 7 for the current observation at $\mathbf{b}_d = \mathbf{m}_{t|j}$ and $P(\mathbf{b}_d = \mathbf{m}_{t|j} | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ is obtained by evaluating the posterior over base values from the last trial at the same point, $\mathbf{b}_d = \mathbf{m}_{t|j}$.

Calculating $P(\mathbf{b}_d | \{\mathbf{x}_s\}_{s=1}^t)$. Finally, we compute $P(\mathbf{b}_d | \{\mathbf{x}_s\}_{s=1}^t)$, the posterior distribution over the base values for all electrodes. An analytical expression for this can be written as

$$P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t) = \sum_{j=1}^J P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t, y_{d,t}=j) P(y_{d,t}=j | \{\mathbf{x}_{d,s}\}_{s=1}^t), \quad (13)$$

where the two terms on the right hand side of Eq. 13 are obtained from Eqs. 8 and 9. The right hand side of equation Eq. 13 is a mixture of multivariate Gaussians. This will be problematic as in general there will be an exponentially increasing number of components in the representation of $P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t)$ as t increases. Therefore, we approximate the right hand side of Eq. 13 with mean $\mathbf{m}_t$ and covariance matrix Σ_t using moment matching as has been done in previous hypothesis merging algorithms (Blom and Bar-Shalom, 1988). We use the standard formulas for moment matching with two multivariate Gaussians to write

$$\mathbf{m}_t = \sum_{j=1}^J \mathbf{m}_{t|j} P(y_{d,t}=j | \{\mathbf{x}_{d,s}\}_{s=1}^t) \quad (14)$$

$$\Sigma_t = \sum_{j=1}^J [\sum_{t|j} + (\mathbf{m}_{t|j} - \mathbf{m}_t)(\mathbf{m}_{t|j} - \mathbf{m}_t)^T] P(y_{d,t}=j | \{\mathbf{x}_{d,s}\}_{s=1}^t), \quad (15)$$

where $P(y_{d,t}=j | \{\mathbf{x}_{d,s}\}_{s=1}^t)$ is calculated from Eq. 8. This completes the decoding procedure for trial t . For each new trial, this set of calculations is repeated in an iterative manner.

2.4.2. Learning the Parameters for the Self-Recalibrating Classifier—While the self-recalibrating classifier is designed to produce stable day-to-day classification without daily retraining, it must be initially trained once. To perform this initial training there must exist a small number of days for which neural activity, $\mathbf{x}_{d,t}$, has been recorded along with associated reach directions, $y_{d,t}$. We will use the notation $\mathcal{D}_{\text{Train}}$ to refer to this set of days. To avoid any possible confusion, we stress that after the classifier has been trained on this initial set of days, it requires no further supervised retraining and will be equipped to autonomously calibrate itself on new data that it encounters during normal decoding use.

The purpose of the initial training is to learn the mean, m_e , and variance, s_e , characterizing the distribution of base values for each electrode from Eq. 3 and the offset, $o_{e,j}$, and variance, $v_{e,j}$, characterizing the observation model in Eq. 6 for each electrode-class pair. Fitting the self-recalibrating classifier is more difficult than fitting the standard classifier because the daily base values for each electrode, $b_{d,e}$, cannot be directly observed. Instead, these values are randomly drawn each day from the distribution specified by Eq. 3, and after the base values are drawn on each day, they specify the mean of the Gaussian distributions in Eq. 6 from which neural spike counts are drawn. Thus, the neural data we directly observe is essentially one step removed from the distribution that is characterized by the parameters we desire to estimate. Fig. 3 illustrates the conceptual difference between the standard and self-recalibrating classifier models.

The expectation-maximization (EM) algorithm (Dempster et al., 1977) is a general procedure for finding the maximum likelihood model parameters in the presence of

unobserved variables, in this case $b_{d,e}$. The EM algorithm is an iterative procedure which is initialized with an initial guess for the model parameters. In this case, these are the m_e , s_e , $o_{e,j}$ and $v_{e,j}$ parameters. Each iteration of the algorithm produces new estimates for these parameters through a two step procedure. In the first step (referred to as the E-step) of each iteration, the latest parameter estimates and observed data are used to form a posterior distribution over the unobserved variables conditioned on the observed data. For the self-recalibrating classifier, the unobserved variables are the daily base values for each electrode, $b_{d,e}$, and the observed data are the neural observations, $x_{d,t}$, and reach directions, $y_{d,t}$, for all trials on all days in the training set. In the second step of each iteration (referred to as the M-step), parameter values are updated so that the expectation of the log likelihood of the observed and unobserved data, computed using the posterior distribution found in the first step of each iteration, is maximized. This algorithm is known to converge to a local optimum. See Supplementary Methods for specific details of the EM algorithm used in this work.

2.4.3. Robust Classification Through Outlier Detection—The classification procedure described above can be modified to improve robustness on real world data. Key to the probabilistic model underlying the self-recalibrating classifier is the specification that firing rates can change between but not within days. While minor changes in firing rate throughout a day are tolerable, large changes, which may be observed in real data, can be problematic.

In particular, the self-recalibrating classifier represents knowledge of the base values for a day obtained from classifying trials 1 through t with the distribution $P(\mathbf{b}_d | \{x_{d,s}\}_{s=1}^t)$. This distribution is refined for each trial that is classified, and over time the uncertainty for the base values for a day, represented in the covariance matrix Σ_t , is reduced. This presents no problem in normal operation. However, if the firing rate of an electrode suddenly changes after the uncertainty in base values for a day is reduced, spike counts recorded on that electrode will strongly and erroneously affect classification.

The classification procedure can be modified by adding an additional step to be performed before any of the steps already described when classifying reach direction for trial t . In this step, electrodes with large firing rate changes are flagged, and the uncertainty in the estimate for the base value for those electrodes is increased to a predetermined value. This prevents those electrodes from dominating when determining reach direction probability and maintains classification accuracy. As classification continues, uncertainty around the base value estimates for the flagged electrodes is reduced starting from the predetermined value as the self-recalibrating classifier refines the distribution $P(\mathbf{b}_d | \{x_{d,s}\}_{s=1}^t)$.

To flag erratically behaving electrodes, a distribution which predicts the probability of the spike count on trial t for electrode e given all previous neural data, $P(x_{d,e,t} | \{x_{d,s}\}_{s=1}^{t-1})$, is formed for each electrode as shown in Supplementary Methods. This can be used as a gauge for determining if the observed spike count, $x_{d,e,t}$, for a trial and an electrode is anomalous by setting upper and lower thresholds on the value of $x_{d,e,t}$ so that a normally behaving electrode will only exceed these values approximately 1% of the time. Whenever an

electrode does exceed these values, it is flagged as behaving erratically. As with any anomaly detection procedure, thresholds have to be set to balance true and false positive rates, and a false positive rate of 1% should not significantly affect classification performance.

After one or more electrodes are flagged as behaving erratically for trial t , uncertainty around the estimates for their associated base values in the distribution $P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ are increased. For ease of notation, we will momentarily switch from representing the set of base values in vector notation to set notation, so that each entry in the vector $\mathbf{b}_d$ will be represented by a corresponding member in the set $\{b_{d,e}\}_{e=1}^E$. We will refer to the set of flagged electrodes for trial t as $E_{t, \text{flagged}}$ and the set of non-flagged electrodes as $E_{t, \text{non-flagged}}$. The uncertainty for the base values of erratically behaving electrodes represented in the distribution for $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ is reset in a two step procedure. In the first step, the distribution over only non-flagged electrodes,

$P(\{b_{d,e}\}_{e \in E_{t, \text{non-flagged}}} | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$, is formed. Formally, this is accomplished by forming the marginal distribution over the electrodes in $E_{t, \text{non-flagged}}$ from $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ and is easily accomplished since $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ is a Gaussian distribution. In the second step, a new distribution over the base values for all electrodes, which we refer to as $P^*(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$, is formed as

$$P^*(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1}) = P(\{b_{d,e}\}_{e \in E_{t, \text{non-flagged}}} | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1}) \prod_{e \in E_{t, \text{flagged}}} P'(b_{d,e}), \quad (16)$$

where $P'(b_{d,e})$ are Gaussian distributions. The mean of the distribution $P'(b_{d,e})$ is set to, $m_{e,t-1}$, which is the element for electrode e of the mean vector for the distribution $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$. This ensures that the mean of the final distribution, $P^*(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$, is the same as that of original distribution $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$. The variance of the distribution $P'(b_{d,e})$ is set to s_e , the variance of the distribution that base values are modeled as being drawn from at the start of each decoding day. Computationally, these two steps can be carried out through a series of simple manipulations of the covariance matrix for $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ and increase the uncertainty around base values for suspect electrodes in $P^*(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$. After the distribution for $P^*(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ is formed, classification, as described by the algorithm above, is performed where $P(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ is replaced with $P^*(\{b_{d,e}\}_{e=1}^E | \{\mathbf{x}_{d,s}\}_{s=1}^{t-1})$ as a starting point when decoding trial t .

2.5. A Simplified Self-Recalibrating Classifier

The SR classifier addresses the problem of base values, $\mathbf{b}_d$, which can drift from day to day by starting with a prior distribution, $P(\mathbf{b}_d)$, over these values and then forming refined distributions, $P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t)$, over base values based on neural activity seen throughout the

decoding day. When predicting reach direction for a trial, the *SR* classifier then considers the probability of reach directions under all possible values for $\mathbf{b}_d$ and weights its estimate according to how likely each value of $\mathbf{b}_d$ is, given by $P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t)$.

Thus, the *SR* classifier never assumes any single value of $\mathbf{b}_d$ is correct, but decodes by considering an infinite number $\mathbf{b}_d$ values, each weighted by some probability. This represents a fully probabilistic means of decoding. However, it is not the only possible means of accounting for drifting tuning parameters. In particular, we can forego modeling uncertainty in $\mathbf{b}_d$ and instead decode with a single point estimate for $\mathbf{b}_d$ formed for each trial. For the simplified classifier, $\mathbf{b}_d$ is defined so that $b_{d,e}$ is the average number of spike counts on electrode e for any trial on day d . We can express this in mathematical notation as $b_{d,e} = \mathbb{E}[x_{d,e,t}]$.

An estimate for $b_{d,e}$ on trial t , which we refer to as $\hat{b}_{d,e,t}$, can be formed by taking the empirical average of all of the observed spike counts for electrode e on day d up to and including trial t . This estimate can be recursively formed as

$$n_{d,t} = n_{d,t-1} + 1, \quad (17)$$

$$\hat{b}_{d,e,t} = \frac{n_{d,t-1} \hat{b}_{d,e,t-1} + x_{d,e,t}}{n_{d,t}}. \quad (18)$$

We must select initial values for $n_{d,0}$ and $b_{d,e,0}$, where the hat has been dropped on $b_{d,e,0}$ to indicate this is a parameter of the model specifying a value with which to initialize Eq. 18. If we select $n_{d,0} = 0$, $\hat{b}_{d,e,t}$ will be the true empirical average for trials 1 through t . However, this has one disadvantage. Because of the very fact that Eq. 18 forms estimates as the empirical average of spike counts, at the start of the day there will be no previous trials with which to average spike counts for initial trials with and hence values for $\hat{b}_{d,e,t}$ will be strongly influenced by the spike counts of each trial. This will result in large fluctuations in $\hat{b}_{d,e,t}$ for the early trials of a decoding session. An alternative is to select $b_{d,e,0}$ so that it represents the expected average spike count on electrode e for any given day and assign this some initial weight by setting $n_{d,0} > 0$. Intuitively, we can think of starting each day with an initial estimate of the average spike count for the electrode that was obtained as the empirical average of $n_{d,0}$ “virtual samples.” This initial estimate is analogous to the prior distribution for the fully probabilistic classifier specified in Eq. 3. Importantly, as shown in Supplementary Methods, as long as $n_{d,0}$ and $b_{d,e,0}$ are chosen to be finite, $\hat{b}_{d,e,t}$ approaches the true mean firing spike count for electrode e as the number of trials observed during a day grows. While this is reassuring, there are still values of these parameters that will result in optimal decoding, and below we discuss how values for $n_{d,0}$ and $b_{d,e,0}$ can be learned from the training data.

After an estimate for $\hat{b}_{d,e,t}$ is formed for each trial, an estimate for the class means, $\mu_{d,e,j}$, can be obtained with Eq. 5, and a classification of reach direction can be made using the

standard classifier described in section 2.3 with the updated class means. The SR_S classifier can be conceptualized as consisting of two algorithms acting together. The first algorithm tracks base values, and hence tuning parameters, using an accumulated average of spike counts. This simple first algorithm is independent of and provides input in the form of updated tuning parameters to the second algorithm, which serves as the decoding core and is nothing more than the standard classifier. Fig. 4 presents a schematic depiction of how the SR_S classifier decodes a single trial and depicts the separation of the self-recalibrating and classification portions of the SR_S decoding algorithm. See Supplementary Methods for a more formal description of the relationship between the SR and SR_S classifiers. We note this method of tracking baseline firing rates is similar to the work of Li et al. (Li et al., 2011), who in their work on adaptive decoders for continuous control signals experimented with tracking changes in baseline firing rates by measuring the empirical average of spike counts. However, since their algorithm runs in a batch mode and does not update parameters on a single trial basis as our does, they do not introduce the concept of selecting initial values and weights to begin the empirical average for a day with.

2.5.1. Learning Model Parameters for the Simplified Self-Recalibrating

Classifier—As with the SR classifier, the SR_S classifier must also be trained once on an initial set of training days. The goal of this training is to learn the electrode-class variances, $v_{e,j}$, and the electrode-class offsets, $o_{e,j}$, for all electrodes and reach directions. In addition, values for $b_{d,e,0}$, must be learned for each electrode, and the initial weight, n_0 , which will be assigned to these means must be determined.

Given a collection of days, $\mathcal{D}_{\text{Train}}$, we estimate values for $b_{d,e,0}$, $o_{e,j}$ and $v_{e,j}$ as follows. First, for each day in $\mathcal{D}_{\text{Train}}$ we estimate the average spike counts on each electrode over all reach directions, $\hat{\mu}_{d,e}$, and the average spike counts on each electrode when the subject made reaches in direction j , $\hat{\mu}_{d,e,j}$. We calculate

$$\hat{\mu}_{d,e} = \frac{\sum_{t=1}^{T_d} x_{d,e,t}}{T_d}, \quad (19)$$

and

$$\hat{\mu}_{d,e,j} = \frac{\sum_{t: y_{d,t}=j} x_{d,e,t}}{T_{d,j}}, \quad (20)$$

where T_d is the number of trials on day d , and $T_{d,j}$ is the total number of trials with reaches in direction j on day d . Having calculated these intermediate values, an estimate of $b_{d,e,0}$ is then calculated as the average of the daily average spike counts as

$$\hat{b}_{d,e,0} = \frac{\sum_{d \in \mathcal{D}_{\text{Train}}} \hat{\mu}_{d,e}}{|\mathcal{D}_{\text{Train}}|}, \quad (21)$$

where $|\mathcal{D}_{\text{Train}}|$ is the number of days in $\mathcal{D}_{\text{Train}}$. Estimates for the class-electrode offsets, $o_{e,j}$, are calculated as the average of the difference between daily class-electrode mean spike counts and the daily average spike counts. We write this as

$$\hat{o}_{e,j} = \frac{\sum_{d \in \mathcal{D}_{\text{Train}}} \hat{\mu}_{d,e,j}}{|\mathcal{D}_{\text{Train}}|} - \hat{\mu}_{d,e}. \quad (22)$$

Finally, the electrode-class variances, $v_{e,j}$, are estimated as

$$\hat{v}_{e,j} = \frac{\sum_{d \in \mathcal{D}_{\text{Train}}} \sum_{t: y_{d,t}=j} (x_{d,e,t} - \hat{\mu}_{d,e,j})^2}{(\sum_{d \in \mathcal{D}_{\text{Train}}} T_{d,j}) - 1}, \quad (23)$$

which can be understood as a generalization of the standard formula for the variance of a sample generalized to account for values of $\mu_{d,e,j}$ which may change between days.

Learning the value of n_0 The equations above are used to estimate values for the $b_{d,e,0}$, $o_{e,j}$ and $v_{e,j}$ parameters of the model. We turn to cross-validation to learn a value for n_0 , which is the amount of weight or number of virtual samples that should be assigned to the initial estimate of marginal spike counts for a day, $b_{d,e,0}$.

In the cross-validation procedure, we seek the value of n_0 which will produce the highest classification accuracy on novel data. To do this, multiple rounds of training and validation are performed on the data that is present in $\mathcal{D}_{\text{Train}}$. In each round, one day in $\mathcal{D}_{\text{Train}}$ is set aside as a validation day. On the remaining days, the $b_{d,e,0}$, $o_{e,j}$ and $v_{e,j}$ parameters of the model are learned. Importantly, these values are learned on all days in $\mathcal{D}_{\text{Train}}$ except the validation day. This is critical because cross-validation uses the validation day to simulate classification on novel data, and it is therefore important that none of the trials in the validation day are used for learning any part of the model for a round of cross-validation. After learning these parameters, a value of n_0 is selected and the classifier is used to predict reach direction for each trial in the validation day. This is repeated for a range of n_0 values. For a single round of cross-validation, one or more values of n_0 will produce the highest classification accuracy on the validation day. To select the final value of n_0 for use in the trained classifier, a round of cross validation is performed for each day in $\mathcal{D}_{\text{Train}}$, and the value of n_0 which produces the highest average accuracy across all rounds is chosen.

3. RESULTS

3.1. Tracking Tuning Parameter Drift Across Time

Figs. 5 and 6 show the results of the analysis of drifting tuning parameters for monkeys L and I, respectively. Throughout this work threshold crossings were collected in a single window per trial. In this analysis and in all of the work that follows we count spikes falling in a single 250 millisecond long window for each trial starting 150 milliseconds after the go cue. As described in section 2.2, spike counts from each electrode were then convolved with a Gaussian kernel to produce smoothed estimates of tuning parameters across trials. The

results presented here were obtained with a kernel width of 100 trials. Results obtained with kernel widths of 25, 50 and 200 trials were qualitatively similar (see Supplementary Results).

Panel A in Fig. 5 shows the traces of two tuning parameters for one electrode for monkey L. Traces for monkey I were similar but are not shown. Three important properties of tuning parameter drift are apparent in this plot. First, tuning parameters change much more between than within days. Second, tuning parameters for the same electrode tend to move up and down together across days. Finally, tuning parameters do not tend to drift progressively further from their values on day 1 over time.

Fig. 5B and Fig. 6A provide a quantitative measurement of the observation that tuning parameters appear to drift more between than within days. The drift of a tuning parameter within a single day is quantified as the standard deviation of its estimated values over all of the trials within a day, $\sigma_{d,e,j}$. This can be measured for all days, electrodes and reach directions. The distribution of all $\sigma_{d,e,j}$ values is plotted as the dashed line in each figure. Similarly, drift between days is quantified as the standard deviation of mean tuning parameter values for each day, $\sigma_{e,j}$. The distribution of $\sigma_{e,j}$ parameters is plotted as the solid line in both figures. The mean of the between-day standard deviation values is significantly larger than that of the within-day standard deviation values (permutation test, $p < .01$). This finding motivates the decision to simplify the self-recalibrating classifiers so that they ignore within-day drift but compensate for between-day drift in tuning parameters.

Fig. 5D and Fig. 6B provide evidence that tuning parameters on the same electrode drift in a correlated manner. Correlation coefficients for all pairs of daily mean tuning parameters (the black dots for the example pair of tuning parameters in Fig. 5C) on the same electrode were calculated for each electrode. We also examined correlation across electrodes by measuring the correlation of daily mean values for pairs of tuning parameters for the same reach direction across all pairs of electrodes. The distribution of these two sets of correlation values is shown in Fig. 5D and Fig. 6B. For both subjects, correlation values for tuning parameters on the same electrode concentrate near 1. This insight has important practical implications for the design of self-recalibrating classifiers which will be detailed in *Discussion*.

While tuning parameters drift in a correlated manner on the same electrode, it is important to point out that the distributions of correlation coefficients in Fig. 5D and Fig. 6B leave room for tuning parameters to move in a limited manner relative to one another while still moving up and down together. Indeed, this is visible in Fig. 5A where traces of the two tuning parameters move up and down together between days but are closer to one another on days 1–4 than they are for days 7–10, for example. We will show below that our decoders, which use a simplified design that models parameters on the same electrode as moving in perfect lockstep, achieve good performance, but future decoders could be developed to model additional aspects of tuning parameter drift. Interested readers may refer to Supplementary Results, where we quantify changes in modulation depth across time.

Fig. 5E and Fig. 6C measure the tendency of tuning parameters to progressively drift from their value on day 1. The percent change of the mean tuning parameter values from day 1, $d_{e,j}$, for days 2 and on for all tuning parameters was calculated. The average of the absolute value of the percent change in tuning parameter values across all electrodes and reach directions for each day was then plotted in these figures. When performing this analysis, for numerical reasons we removed any electrodes which had a value of less than .001 spikes/window for any reach direction. The results are shown in Fig. 5E and Fig. 6C. If tuning parameters drift in an unconstrained manner, the average of the unsigned $d_{e,j}$ values should monotonically increase with time. However, average values do not appear to systematically increase with time, supporting the conclusion that tuning parameters can drift from day to day but do so in a constrained manner.

3.2. Evaluation of Classifier Stability with Prerecorded Data

We compare the performance of the *SR*, *SR_S*, standard retrained and standard non-retrained classifiers on the two prerecorded datasets described in section 2. The large number of days in each makes these ideal datasets for evaluating classifier stability.

The standard retrained classifier is trained fresh each day using the first 400 trials of the day. After training, the classifier is used to predict reach direction on each of the remaining trials of the day. Importantly, while true reach direction is known for each trial of the datasets and is used for training, this information is not made available to the classifier when predicting reach direction (testing) and is only used after a prediction is made to evaluate classifier accuracy.

The *SR*, *SR_S* and standard non-retrained classifiers must be initially trained on a small number of days where neural activity is recorded and the desired reach direction for trials is known. After this they require no further training when classifying reach direction on days when only neural activity is available. For this analysis, all trials on the first 10 days of each dataset are used for training these classifiers. After this training, the classifiers are used to classify reach direction for trials on the remaining days of each dataset. To allow a fair head-to-head comparison with the standard retrained classifier, classification on each test day for these classifiers begins at trial 401 of each day. In this way, the set of trials that each classifier predicts reach direction for is the same on any given day. The breakup of data for training and testing each of the classifiers is shown in Fig. 7B.

When training all classifiers, any electrode with an average of less than two spike counts per trial on the training data (the first 400 trials of each day for the standard retrained classifier and all trials on days 1–10 for the standard non-retrained, *SR* and *SR_S* classifiers) is not used when subsequently decoding reach direction on test data. This step prevents these electrodes from causing problems if their average spike counts were to rise at a later point in time when performing classification. Critically, this can be done in clinical practice as the criterion for exclusion is based only on training data.

3.3. Performance of the Standard Classifier without Daily Retraining

We now examine the implications of tuning parameter drift on the standard non-retrained classifier. The performance of this classifier indicates what might happen if we simply used the standard classifier without retraining and did nothing to account for drifting tuning parameters. Of course, this is an offline analysis and subjects may be able to adapt their neural signals through cognitive strategies in a closed loop setting to improve BCI performance (Chase et al., 2009), but the results presented here may still be considered as a type of lower bound on performance without retraining.

Examination of Fig. 8, which displays the average daily classification accuracy of the standard non-retrained classifier for both subjects, reveals two important findings. First, performance of the standard non-retrained classifier is substantially lower than that of the daily retrained classifier on most days. This motivates the development of classifiers which account for drifting tuning parameters. However, classification accuracy of the standard non-retrained classifier does not significantly decline with time, which constitutes the second finding. To quantify this claim, the slope of a linear fit to daily classification accuracy is not significantly different than zero for both subjects. Specifically, the 95% confidence interval for the slope, measured in units of percent accuracy per day, is $(-.006, .006)$ for monkey L and $(-.012, .001)$ for monkey I, which contains zero for both monkeys. While this finding might at first seem surprising, it is consistent with the finding above that tuning parameters do not show a systematic tendency to progressively wander from their values on day 1 but seem to vary within a relatively limited range.

Performance of the Self-Recalibrating Classifiers—Figs. 9 and 10 show the daily classification accuracies of the *SR* and *SR_S* classifiers for each subject. We emphasize that after these classifiers were trained on days 1–10 for a dataset, they were not provided with any information about the true reach direction of trials with which to retrain themselves when predicting reach direction on the following days. Instead, using the techniques described in section 2, the *SR* and *SR_S* classifiers simultaneously predicted reach direction and the daily value of tuning parameters from neural data alone. Therefore, the classification accuracies reported here are indicative of what might be achieved in the clinic without any further daily supervised retraining or recalibration after the system is initially trained.

Table 1 compares the overall accuracies, calculated as the mean of the average daily accuracies, for all four classifiers considered in this work. The overall accuracies of the *SR* and *SR_S* classifiers were both 14% higher than that of the standard non-retrained classifier for monkey L and 13% and 16%, respectively, higher for monkey I. These performance improvements were statistically significant for both self-recalibrating classifiers and both subjects (One Sided Wilcoxon Signed-Rank Test, $p < 10^{-6}$ for both classifiers for monkey L, $p < 10^{-5}$ for both classifiers for monkey I). Further, both self-recalibrating classifiers nearly recovered the accuracy achieved by the standard retrained classifier. The overall accuracy of the *SR* and the *SR_S* classifiers were both 1% higher than that of the retrained classifier for monkey L and only 5% and 3% lower for monkey I. These differences were not statistically significant for either self-recalibrating classifier and monkey L (One Sided Wilcoxon Signed-Rank Test, $p > 0.1$ for *SR* classifier and $p > .05$ for the *SR_S* classifier),

while the small decrease in accuracies observed on the data for monkey I were (One Sided Wilcoxon Signed-Rank Test, $p < 10^{-4}$ for SR classifier and $p < 10^{-2}$ for the SR_S classifier). To be clear, the self-recalibrating classifiers decoded the exact same trials, trials 401 and on for each test day, as the retrained classifier and achieved this performance without the benefit that the retrained classifier had of being retrained on the fully labeled data of the first 400 trials of each day.

An important consideration when evaluating the performance of a self-recalibrating classifier is the number of trials required to learn tuning parameter values. A self-recalibrating classifier is of little use if its performance starts near chance each day and then takes hundreds of trials to reach acceptable prediction accuracy. Fig. 11 shows the average classification accuracy achieved by each of the self-recalibrating classifiers throughout the course of a decoding day for both subjects. These accuracies were calculated by dividing all of the trials decoded during a test day into sequential bins of twenty trials each and then calculating the average prediction accuracy in each of these bins over all days of data for a subject. The number of trials in each day varied, and the number of bins that accuracies could be calculated over was limited by the day with the fewest number of trials for each subject. Monkey L had at least 386 trials for classification on each test day and monkey I had at least 752. Examination of the plots in Fig. 11 reveal that both classifiers start a decoding day well above chance levels of performance. Further, the SR classifier shows little evidence of a “burn-in” period, where accuracy might be lower as tuning parameter values are being learned. It is likely analysis, a burn-in period of less than 20 trials might become apparent, but in any case the SR classifier rapidly achieves steady state performance. The SR_S classifier shows evidence of slightly decreased performance in the first bin of 20 trials for both subjects, and performance appears to be at asymptotic levels by trials 21–40.

Fig. 12 displays tuning parameter estimates on two electrodes produced by both self-recalibrating classifiers on the first five days of test data for monkey L. Results for monkey I were similar. For this analysis, estimates for true tuning parameters were obtained by convolved spikes with a Gaussian kernel with a standard deviation of 100 trials as described in *Methods* with one exception: days were processed separately and not concatenated together. The SR classifier produces a distribution over tuning parameter values on each trial (shown as the distributions over $\mu_{d,e,j=1}$ and $\mu_{d,e,j=2}$ in Fig. 2), and the mean of these

distributions, $\mathbb{E}[\mu_{d,e,j} | \{\mathbf{x}_{d,s}\}_{s=1}^t]$, was used to plot the traces in Fig. 12. Even though there are seven tuning parameters per electrode, traces for only two tuning parameters per electrode are shown for visual clarity. Characteristics of the way both classifiers predict tuning parameters become apparent with inspection of the plots in Fig. 12. First, on all days estimates for tuning parameters on the same electrode move up and down together. This immediately follows from the modeling assumption that tuning parameters on a given electrode are yoked to a single base value, $b_{d,e}$, in Eq. 5. Second, tuning parameter estimates vary more strongly at the beginning than at the end of a decoding session. This can be observed in the steep change in tuning parameter values at the beginning of certain days, such as day 14 for electrode 17 in Fig. 12. Intuitively, the reason for the reduced variability in tuning parameter estimates for both decoders is that confidence in estimates of base values increases as more neural activity is observed. Mathematically, this corresponds to

less uncertainty in the distribution of base values, $P(\mathbf{b}_d | \{\mathbf{x}_{d,s}\}_{s=1}^t)$, for the *SR* classifier and an increasing value of $n_{d,t}$ for the *SR_S* classifier.

The “spikes” in the parameter estimates for an electrode produced by *SR* classifier visible in both figures occur when that electrode is flagged as behaving erratically by the mechanism for detecting anomalous spike counts, which may occur after the threshold for an electrode is reset, as described in section 2.4.3. When this occurs, the uncertainty around the base value for that electrode is increased. This allows spike counts collected immediately after an electrode is flagged to strongly influence parameter estimates, and just as at the beginning of a day, multiple trials have to pass before the uncertainty for the base value of a flagged electrode is reduced and parameter estimates vary less. The reader may notice that the plotted electrodes are flagged as behaving erratically multiple times within the same day. As explained in section 2.4.3, the false positive rate of the flagging mechanism was set to 1%; therefore, the frequency of parameter resets seen in Fig. 12 should not be understood to indicate the true frequency of anomalous events. While we do not show results for this, in practice the *SR* classifier appears to be reasonably robust without this mechanism to detect and compensate for electrodes with anomalous spike counts. In fact, classification accuracy of the *SR* classifier was markedly improved on only three days for monkey L and on no days for monkey I by its inclusion.

4. Discussion

The development of self-recalibrating classifiers for intracortical BCI has received little previous attention. In this work, we first characterized properties of tuning parameter drift which can be leveraged for the design of self-recalibrating classifiers. We provided evidence that between-day tuning parameter drift is substantially larger than within-day drift and that tuning parameters on the same electrode drift in a highly correlated manner. Further, tuning parameter drift appears constrained across time. This is consistent with the performance of a non-retrained classifier, which is unstable from day-to-day but does not tend to decline with time. We developed two novel, self-recalibrating classifiers which leverage the insights gained into tuning parameter drift to maintain stable performance over periods of 31 and 26 days of prerecorded data. This results in an approximate 15% improvement in overall accuracy compared to the standard classifier without daily retraining. This improvement was large enough that both self-recalibrating classifiers nearly recovered the accuracy achieved by the standard classifier when it was retrained daily.

The self-recalibrating classifiers presented in this work train themselves daily without knowledge of true reach directions. This is the first report of such classifiers for intracortical BCI to the authors’ knowledge. While these results must be verified in a closed-loop setting, they may represent a significant step towards the clinical viability of intracortical BCI systems. Indeed, future systems patterned off of the ideas in this paper would be ready for use almost immediately after being powered on each day as threshold levels can be set via an automated method and the classifiers update their parameters autonomously in the background during normal BCI use.

The classifiers presented here require foreknowledge of the window in time during which a user is making a decision. Previous work has demonstrated the utility of BCI systems subject to this constraint (Vaughan et al., 2006). Further, self-paced BCI systems can be developed by combining the self-recalibrating classifiers presented here with previously published algorithms to determine at each point in time whether the user is intending to make a decision or not (Achtman et al., 2007; Kemere et al., 2008).

The two self-recalibrating classifiers performed very similarly. Is one to be preferred over the other for clinical application? The SR_S classifier adds little computational cost to the standard BCI classifier. This has important implications for resource constrained systems, such as implantable medical devices. Computational hardware which supports the standard classifier will also likely support an upgrade to the simplified self-recalibrating classifier. However, the SR classifier may still be preferred in two scenarios. First, as shown in Fig. 11, the SR classifier requires fewer trials to reach stable levels of performance at the beginning of a decoding day. Second, the SR_S classifier implicitly assumes that a user makes reaches in all directions with equal frequency. If this assumption is broken, which may occur if a user reaches predominantly in one direction, tuning parameter estimation will suffer. The SR classifier bypasses this shortcoming by calculating the probability of reaching in each direction while estimating tuning parameters for each trial. In addition to these considerations, the SR classifier makes modeling assumptions explicit. Changes to these assumptions can be carried out in a straight forward manner with the SR classifier, and it not obvious how this might be done with the SR_S classifier.

Tuning parameter estimates produced by both self-recalibrating classifiers could differ significantly from those estimated as ground truth. However, this did not prevent either classifier from nearly reproducing the decoding accuracy of the daily retrained standard classifier. While this may seem surprising, it is likely classification accuracy is robust to small deviations in estimated tuning parameters as long as their relative ordering is maintained. Our models ensure this by referencing all tuning parameters on an electrode to a single base value. Preliminary analyses with more complicated models which did not force tuning parameters on the same electrode to move in perfect lockstep and also modeled drift in tuning parameters within days did not reveal markedly improved decoding performance (data not shown). Therefore, as the ultimate goal of this work is stable classification and not tuning parameter estimation, we believe the simpler models are to be preferred.

Our decoders also do not carry information about tuning parameters from one decoding day to the next. However, this does not significantly hinder decoding performance. This may seem surprising at first as Fig. 5A indicates that on some pairs of days the previous day's tuning parameter values may be predictive of the next day's values. However, we note that this is not the case for all days. In Fig. 5A, for example, there is a large jump in the value of parameters between days 4 and 5. Thus, on some days using information about the decoding parameters from the previous day would likely help the decoder learn the correct tuning parameters for the present day more quickly, and on some others it could hinder learning. Additionally, as shown in Figs. 11B and 11D, both decoders learn tuning parameter values quickly enough to allow average decoding performance to reach asymptotic levels in tens of trials. Thus, the improvement in decoding accuracy that could be realized by learning tuning

parameters more quickly is fairly minimal. Finally, formally modeling a time series of base values between days would make our EM algorithm for the SR classifier more complicated as the distribution of base values conditioned on observed data would no longer factor as a product of the distribution of base values on individual days (see Eq. S1 of the Supplementary Methods), and additional parameters for a model of transitions in base values between days would need to be learned. Though these reasons motivate the simplification to not carry tuning parameter information between days, we acknowledge this may be interesting future work.

An important aim of this work was to identify properties of the drift in intracortical BCI signals that could be exploited for the development of simplified classifiers. The finding that tuning parameters on the same electrode drift in a highly correlated manner reduces the dimensionality of the space that tuning parameters must be tracked in as algorithms can track a single base value instead all individual parameters on an electrode. High correlations in the drift of tuning parameters on the same electrode may be consistent with changes in baseline firing rates. Changes in baseline could be driven by changes in recorded action potential amplitude (Chestek et al., 2011), micromovements of the recording array relative to cortical tissue (Santhanam et al., 2007) and daily fluctuations in the value of the threshold that was used for spike detection. Whatever the mechanisms which underlie this behavior, robust classification was achieved by the relatively simple SR_5 classifier. The success of this simple approach is tantalizing because it suggests that any BCI algorithm which can be parameterized by variables which are referenced to the average spike counts of electrodes could be made self-recalibrating by following a recipe similar to that presented in section 2.5 for the SR_5 classifier. While we caution below results that pertain to adaptive classifiers should not be trivially generalized to that of regressors, it could be productive to apply this same modification to standard algorithms for regression, such as the Kalman filter (Wu et al., 2006; Wu and Hatsopoulos, 2008) or population vector algorithm (Georgopoulos et al., 1986, 1988) to render them self-recalibrating.

Another important finding was that tuning parameters drift much more between than within days. This suggests that self-recalibrating classifiers can be simplified to capture only between-day changes, as was done in this work. Recently Perge et al. investigated intracortical BCI signal stability within the context of decoding continuous control signals (i.e., regression) in closed loop human experiments (Perge et al., 2013). While they do not explicitly compare within-day changes to between-day changes in BCI signals, they do report reduced decoding accuracy due to within-day changes in neural spike rates. The differing conclusions may be explained by the relative sensitivity of regressors and classifiers to changes in BCI signals. Even small signal changes in BCI signals will likely be reflected in the decoded output of a regressor, whereas the performance of a classifier will only be affected if the relative ordering of the probabilities assigned to each class (reach direction in this work) are altered. This should lend caution when applying our findings on tuning parameter drift to the continuous decoding scenario. While tuning parameters do move up and down together between days on the same electrode, other aspects of tuning, such as modulation depth, as shown in Supplementary Results, can also change. For the purposes of developing self-recalibrating classifiers for our datasets, we were able to

simplify our classifiers and ignore these other aspects of drift. However, it may be that these simplifications would be more detrimental in the context of regression. This is not to degrade the importance of the present study. As discussed in the introduction, BCI classifiers offer unique and important clinical promise, and insights that allow for the unique simplification of BCI classifiers may facilitate their use in non-laboratory settings.

Recently Flint et al. examined BCI performance of a non-retrained continuous decoder of arm movements in a closed loop setting (Flint et al., 2013) when either local field potentials or threshold crossings were used for control signals. We focus on their results for threshold crossings. Two findings of their work are of direct relevance to the present study. First, in an offline analysis of hand-control data, they report highly variable day-to-day performance but little change in average performance with time as decoders age after the initial day they were trained on. This is consistent with our similar finding for a discrete decoder. Second, when they measure correlation coefficients between a variety of measures of closed loop decoder performance and time, they show non-degrading performance over time for one subject and a mild degradation of performance for a second. Beyond these overall trends, day-to-day variability in performance is still apparent. While this may be partly due to variability in subject motivation, it is also quite possible that the use of self-recalibrating classifiers could have provided reduced day-to-day variability in performance in this closed loop setting.

There are at least two possible reasons tuning parameter drift is constrained across time, as we observed in this work. First, it may be the case that tuning parameters are fixed across days but our estimates varied from day to day due to fluctuations in threshold settings based on a finite sample (two seconds) of data. However, this likely does not entirely explain the changes in tuning parameters that we observe because the relationship between multiunit activity and arm kinematics can change even when threshold settings are held fixed across days (Flint et al., 2013). A second tantalizing possibility is that micromovements of the recording array might underlie such tuning parameter changes so that electrodes are in a random position within a relatively stable neural environment. If this is the case, the development of recording arrays that maintain a more fixed position relative to brain tissue may hold promise for increasing the fundamental stability of signals for intracortical BCI use.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by National Science Foundation graduate research fellowships (CAC, VG), NDSEG fellowships (VG, WB), Burroughs Wellcome Fund Career Awards in the Biomedical Sciences (KVS), the Christopher Reeve Paralysis Foundation (KVS), a Stanford University Graduate Fellowship (CAC), a Texas Instruments Stanford Graduate Fellowship (JDF), a Stanford NIH Medical Scientist Training Program grant and Soros Fellowship (PN), NIH grant T90 DA022762 (WB), NIH grant R90 DA023426-06 (WB) and the following awards to KVS: Stanford CIS, Sloan Foundation, NIH CRCNS R01-NS054283, McKnight Endowment Fund for Neuroscience, NIH Directors Pioneer Award 1DP1OD006409, DARPA Revolutionizing Prosthetics program contract N66001-06-C-8005, and DARPA REPAIR contract N66001-10-C-2010.

References

- Achtman N, Afshar A, Santhanam G, Byron MY, Ryu SI, Shenoy KV. Free-paced high-performance brain–computer interfaces. *Journal of Neural Engineering*. 2007; 4:336–347. [PubMed: 17873435]
- Aggarwal V, Acharya S, Tenore F, Shin HC, Etienne-Cummings R, Schieber MH, Thakor NV. Asynchronous decoding of dexterous finger movements using m1 neurons. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2008; 16(1):3–14. [PubMed: 18303800]
- Blom H, Bar-Shalom Y. The interacting multiple model algorithm for systems with markovian switching coefficients. *Automatic Control, IEEE Transactions on*. 1988; 33(8):780–783.
- Blumberg, J.; Rickert, J.; Waldert, S.; Schulze-Bonhage, A.; Aertsen, A.; Mehring, C. Adaptive classification for brain computer interfaces. *Proc. of the 29th Annual International Conf. of the IEEE EMBS*; Lyon, France. 2007. p. 2536-2539.
- Brunner P, Ritaccio A, Emrich J, Bischof H, Schalk G. Rapid communication with a “p300” matrix speller using electrocorticographic signals (ecog). *Frontiers in Neuroscience*. 2011; 5:1–9. [PubMed: 21390287]
- Chase SM, Kass RE, Schwartz AB. Behavioral and neural correlates of visuomotor adaptation observed through a brain–computer interface in primary motor cortex. *Journal of Neurophysiology*. 2012; 108(2):624–644. [PubMed: 22496532]
- Chase SM, Schwartz AB, Kass RE. Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain–computer interface algorithms. *Neural Networks*. 2009; 22(9):1203–1213. [PubMed: 19502004]
- Chestek C, Batista A, Santhanam G, Yu B, Afshar A, Cunningham J, Gilja V, Ryu S, Churchland M, Shenoy K. Single-neuron stability during repeated reaching in macaque premotor cortex. *Journal of Neuroscience*. 2007; 27(40):10742–10750. [PubMed: 17913908]
- Chestek C, Gilja V, Nuyujukian P, Foster J, Fan J, Kaufman M, Churchland M, Rivera-Alvidrez Z, Cunningham J, Ryu S, et al. Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *Journal of Neural Engineering*. 2011; 8(4):045005. [PubMed: 21775782]
- Chestek, CA.; Cunningham, JP.; Gilja, V.; Nuyujukian, P.; Ryu, SI.; Shenoy, KV. Neural prosthetic systems: current problems and future directions. *Proc. of the 31st Annual International Conf. of the IEEE EMBS*; Minneapolis, MN. 2009. p. 3369-3375.
- Chestek CA, Gilja V, Blabe CH, Foster BL, Shenoy KV, Parvizi J, Henderson JM. Hand posture classification using electrocorticography signals in the gamma band over human sensorimotor brain areas. *Journal of Neural Engineering*. 2013; 10(2):026002. [PubMed: 23369953]
- Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ, Velliste M, Boninger ML, Schwartz AB. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*. 2012; 381:557–564.
- Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B*. 1977; 39(1):1–38.
- Eden, U.; Truccolo, W.; Fellows, M.; Donoghue, J.; Brown, E. Reconstruction of hand movement trajectories from a dynamic ensemble of spiking motor cortical neurons. *Proc. of the 26th Annual International Conf. of the IEEE EMBS*; San Francisco, California. 2004a. p. 4017-4020.
- Eden UT, Frank LM, Barbieri R, Solo V, Brown EN. Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Computation*. 2004b; 16(5):971–998. [PubMed: 15070506]
- Flint RD, Wright ZA, Scheid MR, Slutzky MW. Long term, stable brain machine interface performance using local field potentials and multiunit spikes. *Journal of neural engineering*. 2013; 10(5):056005. [PubMed: 23918061]
- Fraser G, Chase S, Whitford A, Schwartz A. Control of a brain–computer interface without spike sorting. *Journal of Neural Engineering*. 2009; 6:055004. [PubMed: 19721186]
- Galán F, Nutton M, Lew E, Ferrez PW, Vanacker G, Philips J, Millán JdR. A brain-actuated wheelchair: asynchronous and non-invasive brain–computer interfaces for continuous control of robots. *Clinical Neurophysiology*. 2008; 119(9):2159–2169. [PubMed: 18621580]
- Ganguly K, Carmena JM. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*. 2009; 7(7):e1000153. [PubMed: 19621062]

- Georgopoulos AP, Kettner RE, Schwartz AB. Primate motor cortex and free arm movements to visual targets in three-dimensional space. ii. coding of the direction of movement by a neuronal population. *The Journal of Neuroscience*. 1988; 8(8):2928–2937. [PubMed: 3411362]
- Georgopoulos AP, Schwartz AB, Kettner RE, et al. Neuronal population coding of movement direction. *Science*. 1986; 233(4771):1416–1419. [PubMed: 3749885]
- Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy KV. A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*. 2012; 15:1752–1757.
- Hochberg L, Bacher D, Jarosiewicz B, Masse N, Simeral J, Vogel J, Haddadin S, Liu J, Cash S, van der Smagt P, Donoghue J. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*. 2012; 485(7398):372–375. [PubMed: 22596161]
- Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB. Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences*. 2008; 105(49):19486–19491.
- Kemere C, Santhanam G, Yu BM, Afshar A, Ryu S, Meng T, Shenoy K. Detecting neural-state transitions using hidden markov models for motor cortical prostheses. *Journal of Neurophysiology*. 2008; 100(4):2441–2452. [PubMed: 18614757]
- Leeb R, Friedman D, Müller-Putz G, Scherer R, Slater M, Pfurtscheller G. Self-paced (asynchronous) bci control of a wheelchair in virtual environments: a case study with a tetraplegic. *Computational Intelligence and Neuroscience*. 2007; 2007:79642. [PubMed: 18368142]
- Li Z, O'Doherty J, Lebedev M, Nicolelis M. Adaptive decoding for brain-machine interfaces through bayesian parameter updates. *Neural computation*. 2011; 23(11):3162–3204. [PubMed: 21919788]
- Li Z, O'Doherty JE, Hanson TL, Lebedev MA, Henriquez CS, Nicolelis MA. Unscented kalman filter for brain-machine interfaces. *PLoS One*. 2009; 4(7):e6243. [PubMed: 19603074]
- Maynard E, Hatsopoulos N, Ojakangas C, Acuna B, Sanes J, Normann R, Donoghue J. Neuronal interactions improve cortical population coding of movement direction. *The Journal of Neuroscience*. 1999; 19(18):8083–8093. [PubMed: 10479708]
- Musallam S, Corneil B, Greger B, Scherberger H, Andersen R. Cognitive control signals for neural prosthetics. *Science*. 2004; 305(5681):258–262. [PubMed: 15247483]
- Ohnishi K, Weir R, Kuiken T. Neural machine interfaces for controlling multifunctional powered upper-limb prostheses. *Expert Review of Medical Devices*. 2007; 4(1):43–53. [PubMed: 17187470]
- Pascual J, Vidaurre C, Kawanabe M. Investigating eeg non-stationarities with robust pca and its application to improve bci performance. *International Journal of Bioelectromagnetism*. 2011; 13(1):50–51.
- Perge JA, Homer ML, Malik WQ, Cash S, Eskandar E, Friebs G, Donoghue JP, Hochberg LR. Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *Journal of Neural Engineering*. 2013; 10(3):036004. [PubMed: 23574741]
- Pistohl T, Schulze-Bonhage A, Aertsen A, Mehring C, Ball T. Decoding natural grasp types from human ecog. *Neuroimage*. 2012; 59(1):248–260. [PubMed: 21763434]
- Polikov V, Tresco P, Reichert W. Response of brain tissue to chronically implanted neural electrodes. *Journal of Neuroscience Methods*. 2005; 148(1):1–18. [PubMed: 16198003]
- Santhanam G, Linderman M, Gilja V, Afshar A, Ryu S, Meng T, Shenoy K. HermesB: a continuous neural recording system for freely behaving primates. *IEEE Transactions on Biomedical Engineering*. 2007; 54(11):2037–2050. [PubMed: 18018699]
- Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. A high-performance brain-computer interface. *Nature*. 2006; 442(7099):195–198. [PubMed: 16838020]
- Santhanam G, Yu B, Gilja V, Ryu S, Afshar A, Sahani M, Shenoy K. Factor-analysis methods for higher-performance neural prostheses. *Journal of Neurophysiology*. 2009; 102(2):1315–1330. [PubMed: 19297518]
- Scherer R, Muller G, Neuper C, Graimann B, Pfurtscheller G. An asynchronously controlled eeg-based virtual keyboard: improvement of the spelling rate. *IEEE Transactions on Biomedical Engineering*. 2004; 51(6):979–984. [PubMed: 15188868]

- Shenoy KV, Meeker D, Cao S, Kureshi SA, Pesaran B, Buneo CA, Batista AP, Mitra PP, Burdick JW, Andersen RA. Neural prosthetic control signals from plan activity. *Neuroreport*. 2003; 14(4):591–596. [PubMed: 12657892]
- Shenoy P, Krauledat M, Blankertz B, Rao R, Müller K. Towards adaptive classification for bci. *Journal of Neural Engineering*. 2006; 3:R13–R23. [PubMed: 16510936]
- Srinivasan L, Eden U, Mitter S, Brown E. General-purpose filter design for neural prosthetic devices. *Journal of Neurophysiology*. 2007; 98(4):2456–2475. [PubMed: 17522167]
- Sykacek P.; Roberts, S. Adaptive classification by variational kalman filtering. In: Thrun, S.; Becker, S.; Obermayer, K., editors. *Advances in Neural Information Processing Systems (NIPS)*. Vol. 15. Cambridge, MA: MIT Press; 2002. p. 737–744.
- Sykacek P, Roberts S, Stokes M. Adaptive bci based on variational bayesian kalman filtering: an empirical evaluation. *IEEE Transactions on Biomedical Engineering*. 2004; 51(5):719–727. [PubMed: 15132497]
- Taylor DM, Tillery SIH, Schwartz AB. Direct cortical control of 3d neuroprosthetic devices. *Science*. 2002; 296(5574):1829–1832. [PubMed: 12052948]
- Vaughan TM, McFarland DJ, Schalk G, Sarnacki WA, Krusienski DJ, Sellers EW, Wolpaw JR. The wadsworth bci research and development program: at home with bci. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2006; 14(2):229–233. [PubMed: 16792301]
- Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. *Nature*. 2008; 453(7198):1098–1101. [PubMed: 18509337]
- Vidaurre C, Blankertz B. Towards a cure for bci illiteracy. *Brain Topography*. 2010; 23(2):194–198. [PubMed: 19946737]
- Vidaurre, C.; Schlögl, A.; Blankertz, B.; Kawanabe, M.; Müller, K. Unsupervised adaptation of the lda classifier for brain-computer interfaces. *Proceedings of the 4th International Brain-Computer Interface Workshop and Training Course*; Graz, Austria. 2008. p. 122–127.
- Wang, W.; Degenhart, AD.; Collinger, JL.; Vinjamuri, R.; Sudre, GP.; Adelson, P.; Holder, DL.; Leuthardt, E.; Moran, D.; Boninger, ML.; Schwartz, A.; Crammond, D.; Tyler-Kabara, EC.; Weber, D. Human motor cortical activity recorded with micro-ecog electrodes, during individual finger movements. *Proc. of the 31st Annual International Conf. of the IEEE EMBS*; Minneapolis, MN. 2009. p. 586–589.
- Wang Z, Gunduz A, Brunner P, Ritaccio A, Ji Q, Schalk G. Decoding onset and direction of movements using electrocorticographic (ecog) signals in humans. *Frontiers in Neuroengineering*. 2012; 5:15. [PubMed: 22891058]
- Wu W, Gao Y, Bienenstock E, Donoghue JP, Black MJ. Bayesian population decoding of motor cortical activity using a kalman filter. *Neural Computation*. 2006; 18(1):80–118. [PubMed: 16354382]
- Wu W, Hatsopoulos N. Real-time decoding of nonstationary neural activity in motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2008; 16(3):213–222. [PubMed: 18586600]

Figure 1.

A representative trial of the center out reaching task (Monkey L, trial ID: R,2008-12-26,mm, 109). The top row is the behavioral task. The middle row contains spike rasters across 96 electrodes and the bottom row presents the timecourse of the hand position. Full range of scale for the hand position is 15 cm.

Figure 2.

An illustration of the probabilistic distribution over tuning parameters for the self-recalibrating classifier at three different points in time during a decoding session on day d for an example where spike counts recorded on one electrode are used to predict reaches in one of two directions. Panel A shows the distribution over the base value for the day, $b_{d,e}$, for the electrode at the start of the decoding day before any neural activity is available. This distribution is a Gaussian distribution with mean m_e and variance s_e . In all panels, distributions over $b_{d,e}$ are shown in red to indicate that this is the distribution the self-recalibrating classifier explicitly represents. Panel A also shows the distributions over $\mu_{d,e,j=1}$ and $\mu_{d,e,j=2}$, which are the mean number of spike counts conditioned on reaching in directions 1 and 2. The classifier does not explicitly represent these distributions, but they can be recovered from the distribution over $b_{d,e}$ with the formula $\mu_{d,e,j} = b_{d,e} + o_{e,j}$. Panel B shows the distributions over $b_{d,e}$, $\mu_{d,e,j=1}$ and $\mu_{d,e,j=2}$ after spike counts for one trial have been observed. Importantly, note that the true reach direction associated with this spike

count is not made available to the classifier. Panel C shows distributions over the same quantities after 10 trials have been observed. The solid dots in panels B and C represent spike counts observed on individual trials.

Figure 3.

Graphical models for the generation of a single spike count on electrode e for trial t on day d for the standard (A) and self-recalibrating (B) classifiers. Circles denote random variables and dots indicate model parameters. In the standard classifier, the electrode-class means, $\mu_{e,j}$, and variances, $v_{e,j}$, for all reach directions directly characterize the distribution over observed spike counts, $x_{d,e,t}$. For the self-recalibrating classifier, the base value for electrode e , $b_{d,e}$, is randomly generated each day from a Gaussian distribution with mean m_e and variance s_e . For a given base value, spike counts are produced according to a distribution with electrode-class means $\mu_{e,j} = b_{d,e} + o_{e,j}$ and variances $v_{e,j}$. In this figure, the notation $\{v_{e,j}\}_{j=1}^J$ and $\{\mu_{e,j}\}_{j=1}^J$ indicates the set of variances and means, respectively, for all reach directions for electrode e .

Figure 4.

An illustration of how the SR_5 classifier predicts reach direction for a single trial. Starting from the top, spike counts, $\mathbf{x}_{d,t}$, for trial t are recorded and the SR_5 decoder updates its estimate for the day's base values, $\hat{b}_{d,e,t}$, using a simple accumulated average. Updated estimates for tuning parameters, $\hat{\mu}_{d,e,j}$, are then calculated from the updated estimates of the base values. Finally, the standard classifier is used with the updated tuning parameters to predict reach direction.

Figure 5.

Statistical structure of changes in tuning parameters across days for monkey L. Panel A: Example traces of $\mu_{d,e,j,t}$ for two representative tuning parameters on the same electrode across the first 25 days of data. Background shading represents day boundaries. Panel B: Distribution of the between ($s_{e,j}$ for all e and j) and within-day ($s_{d,e,j}$ for all d , e , and j) standard deviations of tuning parameters. A small number (.08%) of within-day standard deviations are outside the range of this plot. Panel C: Daily means (black dots) and associated ellipses which indicate regions in which tuning parameters for each day fall with approximately 95% probability for the same $\mu_{d,e,j,t}$ parameters shown in panel A on each of the 41 days of recorded data. Panel D: Distribution of correlation coefficients computed across and within electrode pairs. To relate panel C to D, the correlation of the cloud of 41 black points in (C) is one of the $96 \times \binom{7}{2}$ pairs of correlations that was used to construct the “same electrode” distribution in panel D. Panel E: Average values with 95% confidence

intervals of unsigned percent change of mean tuning parameter values on days 2–41 from their corresponding value on day 1. Three electrodes were withheld from this analysis due to the small value of the tuning parameter for one or more reach directions on day 1. See text for details.

Monkey I

Figure 6.

Summary of the statistical structure of changes in tuning parameters across days for monkey I. Panels A, B and C correspond to panels B, D and E in Fig. 5. No electrodes were withheld when calculating the average shown in panel C. For space reasons, panels showing representative parameters traces and the spread of values for a pair of tuning parameters are not shown for monkey I.

Figure 7.

Partitioning of data when testing classifiers with (A) and without (B) daily retraining as described in Methods. Each column of rectangles represents an experimental day of collected data, and each rectangle represents a trial.

Figure 8.

Daily accuracy values for the standard non-retrained and retrained classifiers for both subjects with 95% confidence intervals. The retrained classifier was retrained each day while the non-retrained classifier was trained on days preceding those decoded here and then held fixed. See Fig. 7 for a depiction of the partitioning of data when training and testing each classifier. The dashed line in each panel shows a linear fit of daily accuracy of the non-retrained classifier with test day. The slope of this fit is not significantly different from zero for either subject. Arrows on the right show overall accuracies of each classifier.

Figure 9.

Daily performance of the self-recalibrating (top panel) and simplified self-recalibrating (bottom panel) classifiers with 95% confidence intervals for data from monkey L. Arrows on the right show overall accuracies of each classifier. Results for the standard classifier with and without daily retraining are reproduced from Fig. 8 for reference.

Figure 10.

Daily performance of the self-recalibrating (top panel) and simplified self-recalibrating (bottom panel) classifiers with 95% confidence intervals for data from monkey I. Arrows on the right show overall accuracies of each classifier. Results for the standard classifier with and without daily retraining are reproduced from Fig. 8 for reference.

Figure 11.

Average accuracy throughout a decoding session for the SR and SR_S classifiers for both subjects with 95% confidence intervals. Accuracy values are calculated as the average classification accuracy in sequential bins of 20 trials over all test days for a subject. There were 386 or more trials for classification on each test day for monkey L and 752 or more on each test day for monkey I. Panels A and C show results for monkey L, while panels B and D show results for monkey I.

Electrode 17

Electrode 60

Figure 12.

Estimates of tuning parameters on two representative electrodes (left and right columns) produced by the *SR* (top) and *SR_S* (bottom) classifiers on the first five days of test data for monkey L compared to smoothed estimates obtained by convolving spikes with a Gaussian kernel. The “spikes” in the tuning parameter estimate of the *SR* classifier occur when the mechanism for detecting anomalous spike counts is triggered and the uncertainty around the parameter estimates for an electrode is reset. The false positive rate for detecting anomalous spikes was set to approximately 1%, and the rate of spikes seen in this plot should not be understood as the true rate of anomalous events.

Table 1

Average daily accuracy with 95% confidence intervals for the standard non-retrained, retrained and both self-recalibrating classifiers for both subjects.

Classifier	Monkey L	Monkey I
Non-Retrained	60 $\pm$ 5%	62 $\pm$ 5%
Retrained	73 $\pm$ 2%	80 $\pm$ 2%
Self-Recalibrating	74 $\pm$ 2%	75 $\pm$ 1%
Simplified Self-Recalibrating	74 $\pm$ 2%	77 $\pm$ 1%