

Inferring Attentional State and Kinematics from Motor Cortical Firing Rates

Frank Wood[†] Prabhat[‡] John P. Donoghue^{§*} Michael J. Black[†]

[†]Department of Computer Science, Brown University, Providence, RI, USA

[‡]Center for Computation and Visualization, Brown University, Providence, RI, USA

[§]Department of Neuroscience, Brown University, Providence, RI, USA

Abstract—Recent methods for motor cortical decoding have demonstrated relatively accurate reconstructions of hand trajectory from small populations of neurons in primary motor cortex. Decoding results are often reported only for periods when the subject is attending to the task. In a neural prosthetic interface, however, the subject must be able to switch between controlling a device or performing other mental functions. In this work we demonstrate a method for detecting whether or not a subject is attending to a motor control task. Using the firing activity of the same neural population used for decoding hand kinematics we demonstrate that a Fisher linear discriminant performs well in classifying the attentional state of a monkey. We use the output of this classifier to augment a hidden state in a first order Markov model and use particle filtering to recursively infer hand kinematics and attentional state conditioned on neural firing rates. We demonstrate high accuracy on test data where a monkey switches between attending to a task and not. By decoding a discrete “state” in addition to hand kinematics our proposed classification and estimation scheme may enable real-world neuroprosthetic functions such as “hold”, “click”, and “turn off/on”.

Index Terms—Particle filtering, pattern classification, neural decoding, neuroprosthetics.

I. INTRODUCTION

Neuroprosthesis research and development has matured to the point that neuroprosthetic devices are being tested in humans [1], however many ways to improve neural decoding and prostheses remain. For instance, algorithmic improvements continue to be made over population vector [2] and linear filtering [3], [4] decoding approaches. Kalman filters [5], particle filters [6], artificial neural networks, and support vector machines [7] have all been shown to be more accurate, though often at the cost of increased computational complexity. Looking forward, improvements to neural recording and signal processing technology, neural representational and computational models, and engineering solutions to real world complexities are increasingly important to the development of effective neural prosthetic devices.

The issue we address in this paper is one of the “real world” engineering issues critical to neuroprosthesis development: How does one “turn off” a neural prosthesis? In other words, if a user of a prosthesis either wishes no longer to use the device, or turns their attention away from use of the device we might wish for it to temporarily stop decoding

or “turn off”. This is akin to writing a letter and then turning one’s attention to other tasks such as typing or grasping a cup of coffee. Motor cortical neurons continue to fire but the pattern of activity will be different than during the writing task. Analogously, a neural motor prosthesis should be able to recognize when a subject is no longer interested in a particular control task by analyzing their neural activity. In current systems, the direct connection between cortical firing rates and decoded output means that this is not the case.

One could imagine various biological methods for switching a prosthesis on/off. Our focus, however, is on finding an automatic way of detecting user attentional states using neural firing rates only. We posit that this could lead to a more natural interface. We found that a simple linear function of a short history of firing rates is sufficient to detect and classify whether or not a monkey is performing a neural control task. When the subject attends to the task we automatically detect this and decode the subject’s hand motion in a standard way but, when the subject stops performing the task, the detector recognizes this and stops decoding. This suggests that a simple function of recent firing activity may be used as an “on/off” switch for neuroprosthetic devices. We develop a probabilistic algorithm to simultaneously decode both discrete attentional states and continuous control signals from motor cortical activity.

II. METHODS

A. Recording

The task and recording setup used to make these hand-labeled recordings was similar to that used in [8] for the on-line neural control of 2D cursor motion. We used two independent recordings. Each was from a trained monkey moving a two-joint manipulandum on a 2D plane to control the motion of a feedback cursor on a computer screen. Hand kinematics and neural activity were simultaneously recorded while the animal performed a sequential random tracking task [8]. Firing rate was estimated by binning spikes into 70 msec time bins and hand position, velocity, and acceleration were computed for each bin.

B. Attentional States

During a typical recording session the monkey spends only a portion of the time actually performing the task. Periodically they may attend away from the task, often letting go of the manipulandum to do various other things with their arm. When they are not gripping the manipulandum,

This work was supported by NIH-NINDS R01 NS 50967-01 as part of the NSF/NIH Collaborative Research in Computational Neuroscience Program.

* John Donoghue is a director, minority shareholder (<3%), and consultant to Cyberkinetics Neurotechnology Systems. The company develops human neurotechnology and sells neurophysiology research equipment, including the recording arrays used in this study.

Correspondence F. Wood (fwood@cs.brown.edu)

its motion is zero and these periods of inactivity are easy to determine by inspection.

We hand classified the the recordings into two states at each time point: if the monkey was performing the experimental task we labeled that point as belonging to class -1 , otherwise, if it was doing something else we labeled it as belonging to class 1 . The determination of whether or not the monkey was performing the task was made by detecting significant regions of constant position and zero velocity. We assume that the monkey was not attending to the task during such segments. Unfortunately we have no way of knowing what the monkey was doing during these time periods; the only question we can ask is whether or not the monkey's neural activity is detectably different during these periods. If so then detecting attentional state shifts could be used to turn on and off the decoding algorithm.

C. Estimation and Inference

Let $X = \{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N\}$ be a time ordered sequence of hand positions and let $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_N\}$ be an indication of the corresponding attentional state of the monkey at each time step. Let the correspondence between indicator and attentional state be $\gamma_i = \{-1, 1\}$ where $\gamma_i = -1$ if the monkey was performing to the task and $\gamma_i = 1$ if the monkey was doing anything else. Let the set of states be $S = \{\vec{s}_1, \vec{s}_2, \dots, \vec{s}_N\}$ where $\vec{s}_i = [\vec{x}_i \gamma_i]$. Let the observations $Z = \{\vec{z}_1, \vec{z}_2, \dots, \vec{z}_N\}$ be a corresponding sequence of firing rates where each $\vec{z}_i = [\vec{z}_i^1, \vec{z}_i^2, \dots, \vec{z}_i^K]$ is a vector of the firing rates of K independent units or multi-units.

1) *Detection*: To detect the attentional state of the monkey we first hand label a training observation sequence of length N by assigning values to γ_i . Using this hand-labeled sequence we learn a Fisher linear discriminant [9] of the form $\langle \vec{w}, \vec{z}_i \rangle + b$ between the “go-labeled” states \mathcal{D}_{-1} and the “stop-labeled” states \mathcal{D}_1 . The Fisher linear discriminant attempts to learn a linear projection direction \vec{w} and offset b that simultaneously maximizes the inter-class distance and minimizes the intra-class scatter. We define the intra-class scatter matrix to be

$$\Sigma_c = \sum_{\hat{z}_i \in \mathcal{D}_c, h < i < N} (\hat{z}_i - m_c)(\hat{z}_i - m_c)^T$$

where $c = \{-1, 1\}$ is a class label, $\hat{z}_i = [\vec{z}_{i-h}, \dots, \vec{z}_i]$ is a history of the h most recent observations, and $m_c = \frac{1}{n_c} \sum_{\hat{z}_i \in \mathcal{D}_c} \hat{z}_i$ is the mean observation history of class c .

It can be shown that the optimal projection is $w = \Sigma_W^{-1}(m_{-1} - m_1)$ where $\Sigma_W = \Sigma_{-1} + \Sigma_1$. The optimal threshold b is chosen to be the point at which the posteriors of two one dimensional Gaussians fitted as a mixture model to each projected class are equal. We do not need to explicitly calculate b for our purposes.

Once trained, such a classifier can noisily distinguish stop and go states using only a short history of firing rates.

2) *Decoding*: We view the decoding problem as a statistical inference problem in which we would like a Bayesian estimate of the posterior $p(\vec{s}_i | \vec{s}_{i-1:1}, \vec{z}_{i:1})$ at every timestep. Making certain independence and first order Markov assumptions leads to a recursive estimate of the posterior:

$$p(\vec{s}_i | \vec{s}_{i-1:1}, \vec{z}_{i:1}) = \kappa p(\vec{z}_i | \vec{s}_i) \int p(\vec{s}_i | \vec{s}_{i-1}) p(\vec{s}_{i-1} | \vec{z}_{i-1}) \delta \vec{s}_{i-1}$$

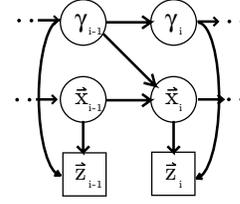


Fig. 1. Graphical model illustrating the conditional dependencies between the discrete “attentional” state, the continuous kinematic state, and the firing rate specified by our model. At timestep i the hand position is \vec{x}_i , the attentional state is γ_i , and the firing rate is \vec{z}_i . Arcs in the graph illustrate conditional dependencies, circles indicate random variables or vectors, and squares indicate observations.

where κ is a normalizing constant.

Two statistical models must be specified to make use of this recursion: an observation model, or likelihood, $p(\vec{z}_i | \vec{s}_i)$ and a state model, or temporal prior, $p(\vec{s}_i | \vec{s}_{i-1})$. For inference we adopt a particle filtering framework in which the posterior is represented as a weighted set of samples and is updated using Monte Carlo integration [6], [10].

To deal with our mixed discrete/continuous state vector, we implement a mixed-state, or switching, particle filter [11] by choosing state and observation models that factor according to the graphical model in Fig. 1. The state and observation models along with their factorizations are

$$\begin{aligned} p(\vec{s}_i | \vec{s}_{i-1}) &= p(\vec{x}_i, \gamma_i | \vec{x}_{i-1}, \gamma_{i-1}) \\ &= p(\vec{x}_i | \vec{x}_{i-1}, \gamma_{i-1}) p(\gamma_i | \gamma_{i-1}) \end{aligned}$$

$$\begin{aligned} p(\vec{z}_i | \vec{s}_i) &= p(\vec{z}_i | \vec{x}_i, \gamma_i) \\ &= p(\vec{z}_i | \vec{x}_i) p(\vec{z}_i | \gamma_i) \end{aligned}$$

where $p(\gamma_i | \gamma_{i-1})$ is the conditional probability of the monkey transitioning between performing or not performing the task, $p(\vec{z}_i | \gamma_i)$ is the generative firing model induced by the Fisher linear discriminant, and

$$p(\vec{x}_i | \vec{x}_{i-1}, \gamma_{i-1}) = \begin{cases} p(\vec{x}_i | \vec{x}_{i-1}) & \text{if } \gamma_i = -1 \\ \delta(\vec{x}_i - 0) & \text{if } \gamma_i = 1 \end{cases}$$

is the hand position at time i conditioned on both it's prior position and the prior attentional state of the monkey where δ is the Kronecker delta function.

The Fisher linear discriminant specifies a two component Gaussian mixture model over the linear projection with means μ_{-1}, μ_1 and variances σ_{-1}, σ_1 . If $G(a; b, c)$ is the probability of a under a Gaussian distribution with mean b and variance c then.

$$p(\gamma_i | \vec{z}_i) = \frac{p(\gamma_i) G(w^T \vec{z}_i; \mu_i, \sigma_i)}{\sum_j p(\gamma_j) G(w^T \vec{z}_j; \mu_j, \sigma_j)}$$

This yields the generative firing rate model induced by the Fisher linear discriminant

$$p(\vec{z}_i | \gamma_i) = \frac{p(\gamma_i | \vec{z}_i) p(\vec{z}_i)}{p(\gamma_i)}$$

where $p(\vec{z}_i)$ can be ignored in the particle filtering framework.

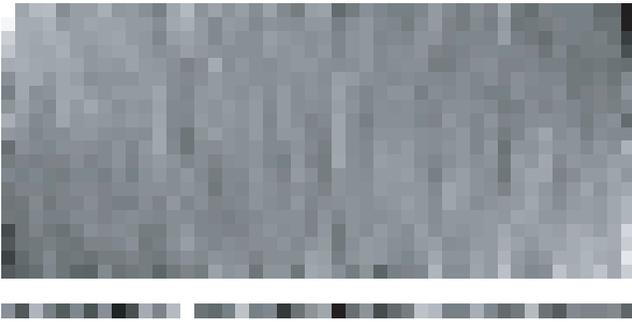


Fig. 3. The optimal linear projection of neural firing history as determined by the Fisher linear discriminant objective. To generate these figures we decomposed the linear projection \vec{w} in $\langle \vec{w}, \vec{z}_i \rangle + b$ into two parts: $\vec{\mu}$ is the mean weight given to each neuron and \vec{r} is the residual weight after $\vec{\mu}$ has been removed, i.e. $\langle \vec{r} + \vec{\mu}, \vec{z}_i \rangle + b$. The top figure shows \vec{r} (y -axis, 20 time bin history, recent rates at the top) for 46 neurons (x -axis). The bottom figure shows $\vec{\mu}$ which is constant for each neuron. In both figures light colors correspond to high values. We interpret \vec{r} as the “tuning” of the cell w.r.t. attentional state and $\vec{\mu}$ as the constant importance of each cell to the classification. The columns of \vec{r} and $\vec{\mu}$ have been sorted to more clearly illustrate the “tuning” of individual cells w.r.t. classification. The cells on the left are weighted in the classification so that modulation in their most recent firing rates counts most. The opposite is true for cells on the right. Cells in the middle may modulate their overall firing rate or not participate in the classification.

We learn the the marginal probabilities $p(\gamma_i)$ and the transition probabilities $p(\gamma_i|\gamma_{i-1})$ between the discrete states from training data and assume linear Gaussian models for $p(\vec{z}_i|\vec{x}_i)$ and $p(\vec{x}_i|\vec{x}_{i-1})$.

$$\begin{aligned} \vec{z}_i - H\vec{x}_i &\sim N(0, Q) \\ \vec{x}_i - A\vec{x}_{i-1} &\sim N(0, W) \end{aligned}$$

The matrices A, W, H , and Q are learned by linear regression; see [5] for details.

D. Experiments

A 42-cell 10 minute recording was used to determine the optimal history to use for the Fisher discriminant. Cross-validated classifier performance on held out data was used to determine the optimal firing history window size which was found to be 20. A separate 46-cell, 21 minute recording was sorted and binned into 18,000 70 msec bins. Of that, 15,000 bins were used to train the classifier and the state and observation models for the particle filter. A separate 3,000 bin segment from the same recording was used for testing.

The particle filter implementation we chose is the Sequential Importance Sampling particle filter with resampling [12]. For every reported result we used 1000 particles.

III. RESULTS

A. Classification

The linear discriminant can be used to classify each time instant using the preceding firing rate history over 20 time bins. This simple classification scheme, with no temporal model of the task state, achieves quite accurate results on the test data: 0.6% false positives (i.e. classifying the subject as performing the task when they are not), 8.0% false negatives (i.e. classifying them as not performing when they are), and 96.5% overall correct classifications. Decoding performance with a temporal model is described below. Figure 3 shows the

	x cc	y cc	x mse	y mse
Particle Filter	0.34	0.32	29.39	36.62
Switching PF	0.76	0.68	8.39	8.68
Kalman Filter	0.45	0.45	34.74	17.64

TABLE I
DECODING PERFORMANCE

optimal linear separating \vec{w} as determined from the training data.

B. Decoding

Figure 2 illustrates typical decoding performance using a particle filter. Displayed here is the expected value of the system state at each time instant in a test sequence. We show results for a traditional particle filter decoder in which the state variable contains only the hand kinematics (bottom) and the results with our augmented state space (top). By default, we let the hand X^{pos}, Y^{pos} be zero for states, \vec{s}_i , in which $\gamma_i = 1$ (i.e. the monkey is not attending to the task). The top figures show the decoding X and Y hand positions as well as the decoded attentional state (performing the task or not). Additionally the figure shows sign output of the Fisher linear discriminant classifier compared with the same estimated by our augmented particle filtering method.

Table I show modified correlation coefficients (cc) and mean square error (mse) for an unmodified particle filter and our switching particle filter. The modified correlation coefficient was calculated as follows: in regions where the true class (as hand labeled) was 1 (or “stop”) we defined the correlation coefficient to be 1 if our estimated state was “stop” and zero otherwise. In regions where the true class was -1 it was calculated as usual. For mean square error we made a similar modification: if the true class was 1 and our estimate agreed we assigned 0 error. If our estimate did not agree (i.e. we were actively decoding) mean square error was calculated as usual. These modifications extend these measures to appropriately account for the mixed continuous and discrete state space. Using these measures, the naive particle filter (without state switching) and the Kalman filter (results also shown for comparison in Fig 2) will be penalized relative to the switching model in regions where the switching particle filter correctly infers the attentional state of the monkey.

IV. DISCUSSION AND FUTURE WORK

We found that a short history of motor cortical firing rates could be used to classify whether or not a monkey was performing a tracking task. In particular, we found that a simple linear classifier proved sufficient for reliable classification. Moreover we found the same population of cells could be used for inferring hand motion. Exploiting these observations we formulated an “augmented state” or “state-switching” particle filter [11] to simultaneously infer the hand kinematics and whether or not the monkey was performing the task. Integrating discrete attentional state information with continuous hand kinematics led to substantially improved decoding performance. The particle filter is able to probabilistically turn on and off the decoder by detecting when the monkey is attending to the hand

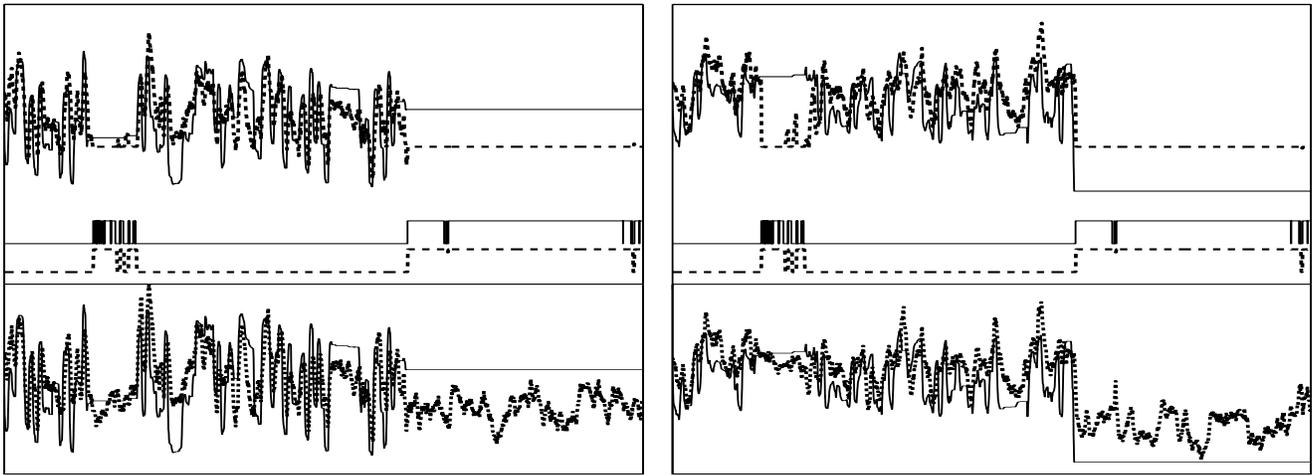


Fig. 2. Decoding results from an unmodified particle filter (lower figures) and from a particle filter with attentional-state inference (upper figures) on a typical 3,000 time bin recording segment. The left figures are for X^{pos} and the right figures are for Y^{pos} . In all figures the solid line is the actual position and the dotted line is the decoded position. In the upper figures the offset solid line is the task state determined by the simple classifier $-1, +1$ and the offset dotted line is the task state as estimated by the particle filter which incorporates a temporal prior. In this segment the monkey performs the task, stops performing briefly, returns to the task, then again stops performing for the remainder. The unmodified particle filter continues to decode over time periods when the monkey is not performing the task, whereas the modified particle filter detects when the monkey stops performing the task and stops decoding.

movement task. Future work could explore related decoding methods such as the switching Kalman filter [13] which may be appropriate for our proposed augmented likelihood and priors models.

It remains to be understood what features in the most recent firing rate are most relevant to the classification of task state. Figure 3 suggests that a number of cell types and activity patterns may be relevant in considering discrete state decoding versus continuous decoding. The ability to decode this type of state information may provide a way of implementing richer neural control systems that include discrete operations such as “hold” or “click” in addition to continuous cursor motion. Additionally if the classification is sufficiently robust, this signal might be used as a way of turning such a prosthesis on and off which could be important in a clinical setting. For example a neural prosthetic that uses the proposed decoding method might enable its user to both move a cursor to a desired location and “click” in that location by detecting different neural firing rate states. Exploring how many states can be detected and how accurately and efficiently this detection can be achieved remains future work.

ACKNOWLEDGMENT

The authors thank M. Serruya and M. Fellows for providing the data for our research.

REFERENCES

- [1] (2005) Cyberkinetics Inc. Foxboro, MA. [Online]. Available: <http://www.cyberkineticsinc.com/>

- [2] A. Georgopoulos, A. Schwartz, and R. Kettner, “Neuronal population coding of movement direction,” *Science*, vol. 233, pp. 1416–1419, 1986.
- [3] M. Serruya, N. Hatsopoulos, M. Fellows, L. Paninski, and J. Donoghue, “Robustness of neuroprosthetic decoding algorithms,” *Biological Cybernetics*, 88(3):201–209, 2003.
- [4] Y. Gao, M. Black, E. Bienenstock, W. Wu, and J. Donoghue, “A quantitative comparison of linear and non-linear models of motor cortical activity for the encoding and decoding of arm motions,” *1st Int. IEEE/EMBS Conf. on Neural Eng.*, Mar. 2003, pp. 189–192.
- [5] W. Wu, M. Black, Y. Gao, E. Bienenstock, M. Serruya, A. Shaikhouni, and J. P. Donoghue, “Neural decoding of cursor motion using a Kalman filter,” *NIPS 15*, 2003, pp. 133–140.
- [6] A. Brockwell, A. Rojas, and R. Kass, “Recursive Bayesian decoding of motor cortical signals by particle filtering,” *J. Neurophysiology*, 91(4):1899–1907, 2004.
- [7] L. Shpigelman, Y. Singer, R. Paz, and E. Vaadia, “Spikernels: Predicting arm movements by embedding population spike rate patterns in inner-product spaces,” *Neural Comp.*, 17(3):671–690, 2004.
- [8] M. Serruya, N. Hatsopoulos, L. Paninski, M. Fellows, and J. Donoghue, “Brain-machine interface: Instant neural control of a movement signal,” *Nature*, vol. 416, pp. 141–142, 2002.
- [9] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. John Wiley and Sons, 2001.
- [10] Y. Gao, M. Black, E. Bienenstock, S. Shoham, and J. Donoghue, “Probabilistic inference of arm motion from neural activity in motor cortex,” *NIPS 14*, pp. 221–228, 2002.
- [11] M. Isard and A. Blake, “A mixed-state Condensation tracker with automatic model-switching,” *IEEE Int. Conf. on Computer Vision*, 1998, pp. 107–112.
- [12] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Springer, 2001.
- [13] W. Wu, M. Black, D. Mumford, Y. Gao, E. Bienenstock, and J. Donoghue, “Modeling and decoding motor cortical activity using a switching Kalman filter,” *IEEE Trans. Biomedical Eng.*, 51(6):933–942, 2004.