Probabilistically Modeling and Decoding Neural Population Activity in Motor Cortex

Michael J. Black and **John P. Donoghue**
Departments of Computer Science and Neuroscience
Brown University, Providence RI 02912

Abstract

This paper introduces and summarizes recent work on probabilistic models of motor cortical activity and methods for inferring, or decoding, hand movements from this activity. A simple generalization of previous encoding models is presented in which neural firing rates are represented as a linear function of hand movements. A Bayesian approach is taken to exploit this generative model of firing rates for the purpose of inferring hand kinematics. In particular, we consider approximations of the encoding problem that allow efficient inference of hand movement using a Kalman filter. Decoding results are presented and the use of these methods for neural prosthetic cursor control is discussed.

Keywords: neural prosthesis, motor cortex, neural decoding, Bayesian inference, Kalman filter, brain-machine interface.

0.1 Introduction

> *"One might think of the computer in this case as a prosthetic device. Just as a man who has his arm amputated can receive a mechanical equivalent of the lost arm, so a brain-damaged man can receive a mechanical aid to overcome the effects of brain damage. It makes the computer a high-class wooden leg."*
>
> Michael Crichton, The Terminal Man [4].

Two fundamental shifts in neuroscience have recently led to a deeper understanding of the neural control of movement and are enabling the development of neural prosthesis that can assist the severely disabled by directly connecting their central nervous systems with assistive devices internal or external to the body. The first of these shifts is the result of new electrode array technology that allows the chronic implantation of hundreds of microelectrodes in the cortex that can sense and ultimately transmit outside the body the activity of populations of neurons. The second shift is part of a movement toward the study of more natural stimuli and behaviors. In contrast to previous work in neuroscience in which the activity of a single cell is correlated with a simple (e.g. one dimensional) change in behavior, today neuroscientists can observe large populations of cortical cells and how they respond during rich behavioral tasks. With richness comes the cost of complexity that makes modeling and understanding the relationship between neural activity and behavior challenging. Neural population recordings can be thought of as a high dimensional time-varying signal while motor behavior can similarly be thought of as a high dimensional time series corresponding to the biomechanical parameters of body pose and motion. We view the problem of modeling the neural code for prosthetic applications as one of learning a probabilistic model relating these high dimensional signals.

This approach is summarized in Figure 1. We focus here on neural firing rates $\mathbf{z}_t = [z_{1,t} \dots z_{n,t}]$ of a population of n cells recorded in primary motor cortex in monkeys and relate this activity to a vector of kinematics $\mathbf{x}_t$ representing the monkey's hand pose and movement at an instant in time t.[1] More generally we want to know the relationship between an entire sequence of firing rates $Z_t = [\mathbf{z}_t \dots \mathbf{z}_1]$ and hand movements $X_t = [\mathbf{x}_t \dots \mathbf{x}_1]$ from time 1 to t. In general we see the

---

[1] While here we focus on firing rates, the probabilistic modeling framework is more general and equally well applies to spike trains or other neural signals such as local field potentials. Focusing on rates however will simplify our probabilistic modeling problem. The same can be said for hand kinematics; for example, instead, we might model biomechanical properties of the arm dynamics.

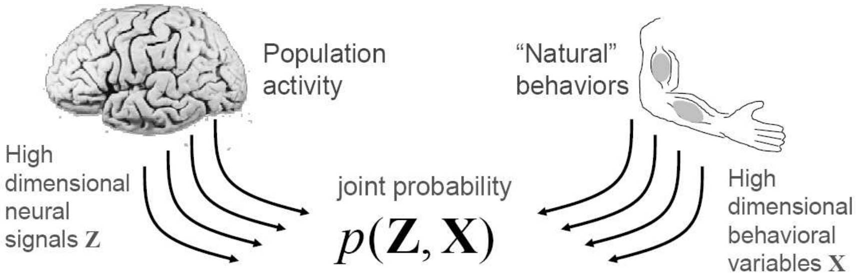

Figure 1 The problem of motor cortical modeling for prosthetic applications can be viewed as one of learning the joint probability of neural population activity and motor behavior. Neural data might correspond to spikes, firing rates, local field potentials or electrocorticograms. Motor behaviors might correspond to joint angles, muscle activation, limb pose or kinematic parameters. Here we focus on probabilistically modeling motor cortical firing rates and hand kinematics (position, velocity and acceleration).

problem as one of modeling the joint probability $p(Z_t, X_t)$ of neural activity and hand motion. From such a general model a variety of quantities can be computed and statistical properties of the model analyzed. Here we focus on the problem of decoding, or inference, of hand kinematics from firing activity. The probabilistic approach allows us to exploit a variety of well understood and powerful tools for probabilistic inference.

The probabilistic modeling problem however is made challenging by the dimensionality of the neural population and the hand kinematics. Consequently we will make a number of explicit approximations below that will make modeling the probabilistic relationships tractable. In particular we will exploit lower dimensional parametric models and assumptions of conditional independence. These will lead us to an efficient decoding algorithm that takes as input a sequence of neural firing rates and returns a sequence of probability distributions representing possible hand motions. This decoding algorithm is used in a neural motor prosthesis that directly connects the motor cortex of a monkey to a computer cursor and enables the monkey to move the cursor under brain control. Such a device provides the foundation for a new class of cortical brain machine interfaces (BMIs) for the severely disabled and, in the near future, may be used to control other external devices such as robot arms or even the patient's own limbs through functional electrical stimulation [14].

This paper introduces and summarizes recent work on probabilistically decoding motor cortical population activity. It briefly summarizes the major issues in the field: sensing neural activity, models of cortical coding, probabilistic decoding algorithms, and applications to neural prostheses. In particular we start with the standard models of motor cortical tuning (e.g. directional tuning) and then show that these are narrow instantiations of a more general linear model relating hand motion and neural firing rates. From this generalization, we show that a well motivated decoding algorithm emerges based on Bayesian probability that provides a principled approach to decoding hand motions. One advantage of this Bayesian approach is that the assumptions made along the way are explicit in a way they are often not in competing approaches. Each of these assumptions provide an opportunity to improve the model and there have already been many such improvements which are beyond the scope of this introduction.

0.2 Sensing Neural Activity

> *"Now listen to me closely, young gentlemen. That brain is thinking. Maybe it's thinking about music. Maybe it has a great symphony all thought out or a mathematical formula that would change the world or a book that would make people kinder or the germ of an idea that would save a hundred million people from cancer. This is a very interesting problem, young gentlemen, because if this brain does hold such secrets, how in the world are we ever going to find out?"*
>
> Dalton Trumbo, Johnny Got His Gun [27].

A variety of sensing technologies allow the recording of neural activity with varying levels of temporal and spatial resolution. To record the action potentials of individual cells we use the Cyberkinetics/Utah microelectrode array shown in

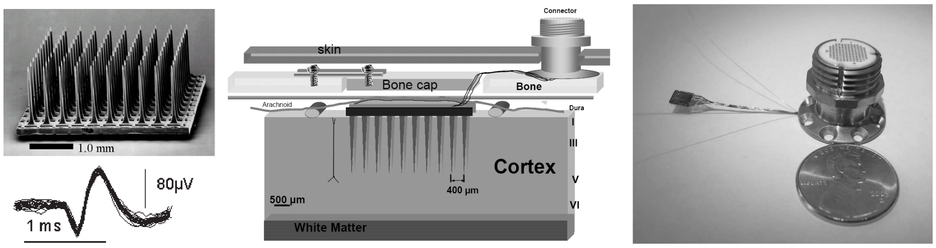

a *b* *c*

Figure 2 Implantable electrode array and connector. *(a)* Cyberkinetics/Bionic/Utah electrode array and example waveforms recorded for one cell. *(b)* Sketch of the implanted array and connector (source: J. P. Donoghue). *(c)* Size of array along with percutaneous connector in reference to a US penny.

Figure 2*a* which consists of a 10×10 grid of electrodes [15]. The array is implanted in arm area of primary motor cortex (MI) in macaque monkeys as illustrated in Figure 2*b* and data is transferred out of the brain through a percutaneous connector shown in Figure 2*c*.

The implant area satisfies a number of constraints. First, our goal is to restore movement to people who have lost the ability to control their bodies directly. It has long been known that the activity of cells in this area of the brain is modulated by arm and hand movements [8, 9]. While it may be possible to train people to use other brain regions to control movement, our working hypothesis is that it will be more “natural” and hence easier to learn to control the movement of cursors or other devices using a region of the brain already related to movement control. Second, this region is surgically accessible and on the surface of cortex facilitating implantation.

Each electrode may record the activity of zero or more neurons. The activity on each channel (electrode) is filtered and thresholded to detect action potentials. If the activity of multiple cells (units) is detected on a single channel the action potentials may be sorted based on their waveform shape and other properties using manual or automatic spike sorting techniques. A representative example of waveforms detected for an individual unit using the device is shown in Figure 2*a*. It is common to recored

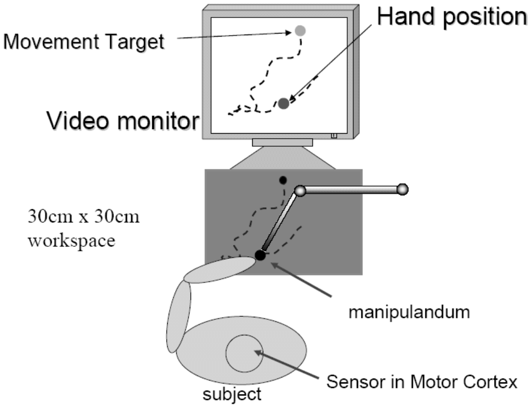

Figure 3 Experimental paradigm. Neural signals are recorded while hand motion controls a computer cursor to hit targets presented at successive random locations on a computer monitor. Source: J. P. Donoghue.

from 40-50 distinct cells from a single array. We have found however that, for neural prosthetic applications, careful spike sorting may not be necessary and it may be sufficient to use the multi-unit activity of all cells recorded on a given channel [30].

To model the relationship between neural firing rates and behavior we used neural spiking activity recorded while a monkey performed a 2D cursor control task [22]. The monkey's hand motion and neural activity were recorded simultaneously and were used to learn a probabilistic model as described below. The task involved moving a manipulandum on a 2D plane to control the motion of a feedback cursor displayed on a computer monitor (Figure 3). In contrast to previous studies that focused on center-out reaching tasks [3, 25] this data was from a sequential random tracking task in which a target appeared on the screen and the monkey was free to move the feedback cursor as they liked to "hit" the target. When a target was acquired it disappeared and a new target appeared in a new random location. Target locations were drawn independently and identically from a uniform distribution over the 2D range of the $30cm \times 30cm$ workspace. See [22] for more information on the sequential random tracking task.

0.3 Encoding

> *If spikes are the language of the brain, we would like to provide a dictionary ... perhaps even ... the analog of a thesaurus.*
>
> Rieke, *et al.*, Spikes [19]

To model what aspects of movement are represented (encoded) by the brain we adopt a probabilistic approach and learn a *generative model* of neural activity. In particular, we seek a function $f(\cdot)$ of the hand kinematics, $\mathbf{x}_t$ at time t, that "explains" the observed neural firing rates

$$\mathbf{z}_t = f(\mathbf{x}_t) + \mathbf{q}_t \tag{1}$$

where we expect the firing activity $\mathbf{z}_t$ to be noisy observations of a stochastic process and where $\mathbf{q}_t$ is a noise vector drawn from some distribution. Note that this *generative model* is descriptive rather than mechanistic – it does not say how the spatio-temporal dynamics of neural networks encode movement.

With generative approach, the problem of modeling the neural code has four components:

1. What neural data should be modeled (e.g. spikes, rates, local field potentials, etc.)?

2. What behavioral variables are important (e.g. joint angles, torques, muscle activation, hand direction, etc.)?

3. What functional relationship between behavior and neural activity is appropriate (e.g. linear or any number of non-linear functions)?

4. What model the "noise" should be used (noise may arise from the stochastic nature of the neurons as well as electrical noise, failures in spike detection/sorting and more amorphous inadequacies of the functional model)?

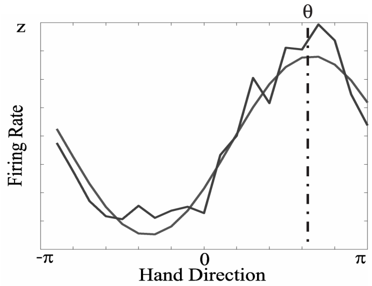

Figure 4 Cosine tuning. The firing rate of a cell (jagged curve) as a function of hand direction θ_t. This data is well fit by a so called “cosine tuning function” (smooth curve). The direction of maximal firing, θ, is referred to as the preferred direction.

In addressing the first question, here we focus on firing rates computed from spike counts in non-overlapping $70ms$ time bins. Firing rates of cells in MI have been long known to be modulated by hand motions and provide a reasonable input signal for neural decoding. While we could work with spike trains, this complicates the probabilistic modeling problem [32].

The next choice pertains to the behavioral variables $\mathbf{x}_t$ we wish to model. Candidates here might include limb joint angles, torques, or muscle activity. While each of these has been shown to be correlated with neural firing rates, there is a simpler representation for the control of computer cursors: hand position, velocity and acceleration. These kinematic parameters have also been shown to be related to modulation of firing rates. The choice here however is not completely independent of the next problem which is the choice of the function f.

While f could be an arbitrary function (e.g. as embodied in an artificial neural network (ANN) [29]) we can impose some constraints on its choice. Low dimensional parametric models, particularly linear ones, are desirable because they are easy to fit to relatively small amounts of data without overfitting. A second design criterion might be “interpretability” which ANN’s lack.

In terms of interpretability, linear models have a distinct advantage in that they are a generalization of well known models of motor cortical coding. One of the

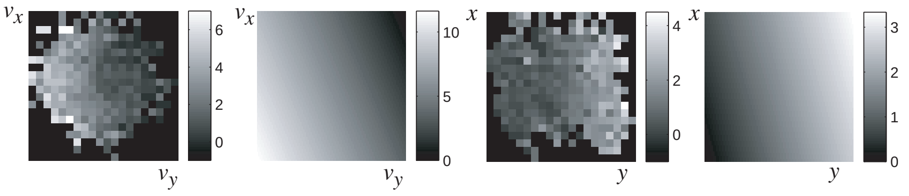

a *b*

Figure 5 Linear tuning functions. *(a)* Firing rate as a function of hand *velocity* for one cell. Light colors correspond to higher firing rates than dark colors. Note that black corresponds to regions of velocity space that were never observed. On the left of *(a)* is a normalized histogram of the firing rates while on the right is the linear fit to this data. *(b)* A different cells shows approximately linear tuning with respect to hand *position* on a 2D plane.

hallmarks of cells in arm area of MI is that they are “directionally tuned” [8, 20]. This theory of motor cortical coding suggests that cells have a “preferred direction” and when the hand moves in this direction a cell’s firing rate is maximal. This is illustrated in Figure 4 for a representative cell from our data. Mathematically the firing rate, z_t, of a cell at time t can be expressed as the following function of hand direction θ_t

$$z_t = h_0 + h\,\cos(\theta_t - \theta) = h_0 + h_x \cos(\theta_t) + h_y \sin(\theta_t) \tag{2}$$

where the h_i are scalar values that can be fit to the data for a particular cell. Note that this equation is in the same form as our generative model above but that there is no explicit model of the noise.

The story does not end with directional tuning however. Moran and Schwartz [16] for example noted that firing rates of MI cells increase with the *speed* at which a hand movement is performed; that is

$$z_t = s_t(h_0 + h\,\cos(\theta_t - \theta)) = h_0^* + h_x^* v_{t,x} + h_y^* v_{t_y} \tag{3}$$

where the h_i^* are, once again, scalar values and $v_{t,x}$ and $v_{t,y}$ represent the *velocity* of the hand in the x and y direction respectively. Figure 5a illustrates this roughly linear velocity tuning for one motor cortical neuron.

The above equation then suggests that the firing rate of these cells is simply a linear function of hand velocity. Again this is not the whole story. Firing rates of these cells may also be linearly tuned to hand position [12] and hand acceleration [5] and possibly even higher order derivatives of the hand motion [34]. Figure 5*b* shows the firing activity of a cell that is roughly linearly tuned to position. For a thorough treatment see [18].

Taken together these findings suggest that firing rates may be approximated as a linear combination of simple hand kinematics (position, velocity and acceleration); that is

$$\mathbf{z}_t = H\mathbf{x}_t \tag{4}$$

where, if $\mathbf{z}_t$ is a vector of n cells' firing rates and $\mathbf{x}_t = [x_t, y_t, v_{t,x}, v_{t,y}, a_{t,x}, a_{t,y}]^T$ contains the hand kinematics at time t, H is a $n \times 6$ matrix that relates hand pose/motion to firing rates. The inclusion of the all these kinematic terms (position, velocity and acceleration) in the model turns out to be important. It has been noted that not all cells in primary motor cortex are equally tuned to each of these variables; some cells are modulated more by one variable or another [18].

It is important to note that this model is a strict generalization of the traditional model of directional tuning. Previous decoding models such as the population vector method rely on tuning for direction or speed and direction [20, 21]. These parameters are included in the linear model along with position and acceleration.

We now come to the final design choice in the generative framework namely: what noise model should we use? Note first that firing rates are strictly positive and, over relatively small time windows, exhibit a roughly Poisson distribution. As a mathematical convenience however we would prefer to model the noise as Gaussian which will admit efficient inference algorithms as described in the following section.

To facilitate such a model we first center the firing rates by subtracting the vector of mean firing rates from all the data; the firing rates are no longer strictly positive. We do the same for the hand kinematics. We then approximate the noise as Gaussian; that is $\mathbf{q}_t \sim N(0, Q)$.

Unlike previous approaches, this generative model explicitly (if only approximately) models the noise in the observations. In particular we take Q to be a full error covariance matrix which models correlations in the noise among the cells. This is critical for accurate modeling since any model is going to be an approximation to the truth and there may be other, hidden, causes of firing rate modulation that may cause correlated errors in the observed firing rates.

0.4 Decoding

If I could find ... a code which translates the relation between the reading of the encephalograph and the mental image ... the brain could communicate with me.

Curt Siodmak, Donovan's Brain [24].

The goal of motor-cortical decoding is to recover the intended movement, for example hand kinematics $\mathbf{x}_t$, given a sequence of observed firing rates $Z_t = [\mathbf{z}_t \ldots \mathbf{z}_1]$. Probabilistically we would like to represent the *a posteriori* probability of the hand motion $p(\mathbf{x}_t | Z_t)$. To represent this probability we first make a few simplifying assumptions that prove quite reasonable in practice. For example, we assume that the hand kinematics at time t are independent of those at time $t-2$ and earlier conditioned on $\mathbf{x}_{t-1}$. This gives a simple form for the *a priori* probability of hand kinematics

$$p(\mathbf{x}_t | X_{t-1}) = p(\mathbf{x}_t | \mathbf{x}_{t-1}, \ldots, \mathbf{x}_1) = p(\mathbf{x}_t | \mathbf{x}_{t-1}). \tag{5}$$

We also assume that, given the kinematics $\mathbf{x}_t$ at time t, the firing rates at time t are conditionally independent of the hand kinematics at earlier times. This gives a simple form for the *likelihood* of firing rates conditioned on hand kinematics

$$p(\mathbf{z}_t|X_t) = p(\mathbf{z}_t|\mathbf{x}_t). \tag{6}$$

With these assumptions, Bayes rule can be used to derive an expression for the posterior probability in terms of the likelihood the prior

$$p(\mathbf{x}_t|Z_t) \propto p(\mathbf{z}_t|\mathbf{x}_t) \int p(\mathbf{x}_t|\mathbf{x}_{t-1}) p(\mathbf{x}_{t-1}|Z_{t-1}) d\mathbf{x}_{t-1}. \tag{7}$$

A “decoded” value for $\mathbf{x}_t$ can then be obtained by either computing the expected value or the maximum a posteriori value of $p(\mathbf{x}_t|Z_t)$.

This Bayesian formulation is very general and the likelihood and prior can be arbitrary. In the general case, the integral in (7) is problematic and must be computed using Monte Carlo sampling methods. For the recursive estimation of $p(\mathbf{x}_t|Z_t)$ this inference takes the form of a “particle filter” which has been applied to neural decoding [2, 6, 7]. These methods however are computationally intensive and not yet appropriate for real-time decoding.

By making a few more simplifying assumptions however, inference with this Bayesian formulation becomes straightforward. In particular we observe that the prior probability of hand motions in our task is well approximated by a linear Gaussian model; that is

$$\mathbf{x}_t = A\mathbf{x}_{t-1} + \mathbf{w}_t \tag{8}$$

where A is known as a state matrix that models the change in kinematics from one time to the next and the noise, $\mathbf{w}_t \sim N(0, W)$, is normally distributed with mean zero and covariance W.

Assuming that the kinematics $\mathbf{x}_0$ is normally distributed at time 0, then $\mathbf{x}_t$ is normally distributed. This is convenient since it implies that firing rates $\mathbf{z}_t = H\mathbf{x}_t + \mathbf{q}_t$ conditioned on $\mathbf{x}_t$ are also normally distributed. While this assumption of Gaussian-distributed firing rates is only approximation, performing a square root transformation of the firing rates improves the approximation; for more details the reader is referred to [7, 34].

With these assumptions, the likelihood term in (7) becomes

$$p(\mathbf{z}_t|\mathbf{x}_t) \propto \exp\left(-\frac{1}{2}(\mathbf{z}_t - H\mathbf{x}_t)^T Q^{-1}(\mathbf{z}_t - H\mathbf{x}_t)\right). \tag{9}$$

The assumptions tell us how firing rates are generated from intended hand movements. Bayes rule tells us how to take such a generative model of firing rates and "turn it around" for the purpose of decoding hand kinematics from observed firing rates.

The linear and Gaussian assumptions mean that fitting the parameters H, Q, A and W is straightforward via least squares regression on training data [34]. Also, given linear Gaussian expressions for the likelihood and prior, the resulting posterior is also Gaussian. Estimating this Gaussian posterior can be done very easily and efficiently using the Kalman filter [11, 28] since the update of the posterior at each time instant can be performed in closed form. For details of the algorithm and its implementation for neural decoding, the reader is referred to [34].

A few example reconstructions of hand trajectories are shown in Figure 6 in which we display the expected hand kinematics, $\mathbf{x}_t$, at each time instant computed from test data not used to train the model. Reconstructed hand trajectories qualitatively match the true trajectories and quantitatively compare favorably to the state of the art (see [34]). The Kalman filter provides a computationally efficient and accurate method for neural decoding that is directly derived from our models of the neural

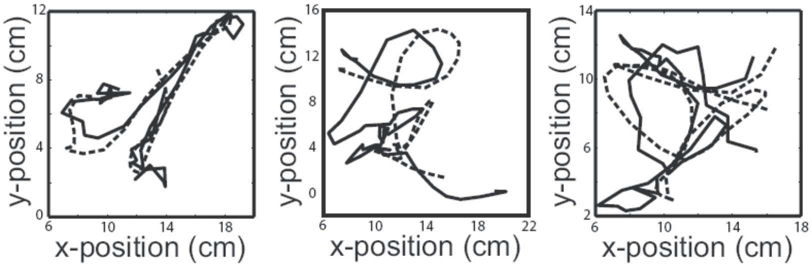

Figure 6 Reconstructed trajectories (portions of $1min$ test data – each plot shows 50 time instants ($3.5s$)): true target trajectory (dashed) and reconstruction using the Kalman filter (solid); from [34].

code. Experiments in monkeys show that the method provides effective on-line cursor control [35]. In particular, Wu *et al.* [35] showed a 50% improvement in the number of targets a monkey could hit in a given period of time using the Kalman filter as compared with a more traditional, non-generative, linear regression method [3, 22].

There is one additional detail that is relevant for accurate decoding. Changes in the firing rates of the cells tend to precede the observed activity. Consequently it is appropriate to train the model with a built in lag j such that

$$\mathbf{z}_{t-j} = H\mathbf{x}_t + \mathbf{q}_t. \tag{10}$$

A fixed lag of approximately $140ms$ improves decoding accuracy. The lag for each cell however may differ and fitting individual lags improves decoding further but complicates learning the model parameters [34]. Wu *et al.* [34] found that the Kalman filter with a $140ms$ lag reconstructed hand trajectories for this data with with a mean squared error (MSE) in hand position of $5.87cm^2$ while a non-uniform lag, optimized for each cell, reduced the MSE to $4.76cm^2$.

They also observed the value of representing a full error covariance matrix in the generative model. Using only a diagonal covariance matrix, which assumes conditional independence of the firing rates of different cells, resulted in an increase in the MSE from $5.87cm^2$ to $6.91cm^2$.

0.5 Interfaces

The big machine.... Operated by remote control.... Operated by the electromagnetic impulses of individual Krell brains.

The Forbidden Planet [1].

There have now been numerous demonstrations of neural control of devices using different recording technologies and different decoding algorithms [3, 10, 21, 22, 25, 35]. In the case of cortical implants these methods can be classified according to two kinds of interfaces: discrete or continuous.

In the discrete task a monkey has one of a fixed number of targets which they must select by either direct arm motion or neural signals [17, 23]. Neural decoding in this case reduces to a discrete classification task. Furthermore, in the case that all the targets are equally likely (i.e. the prior is uninformative), Bayesian classification reduces to maximum-likelihood classification. Given a population of neurons in primary motor cortex or pre-motor areas this classification task can be performed extremely accurately. In fact, under brain control monkeys can respond more rapidly than by making actual arm motions and they quickly learn to perform target selection without moving their arm [17, 23].

A variety of interfaces have been developed for disabled people using discrete selection such as this (though using EEG and not neural implants). Interfaces based on selection of a small number of states (e.g. binary) can be cumbersome to use. It is not yet known however how many discrete states can be recognized from a neural population of a given size.

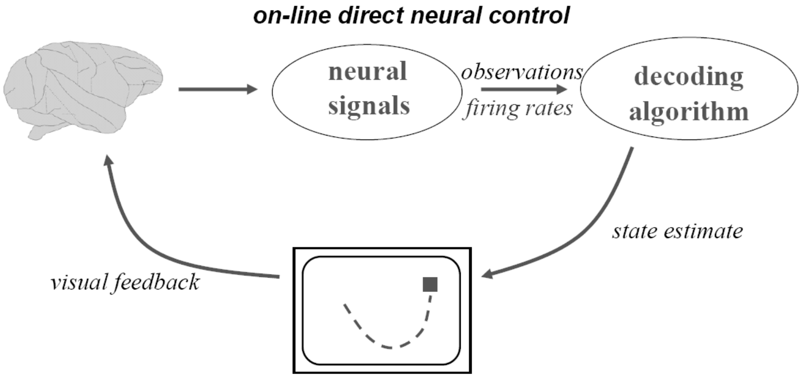

Figure 7 Closed-loop neural cursor control. Neural signals directly control cursor motion while a subject receives feedback about the cursor position through their visual system. In our case the neural signals are population firing rates and the decoding algorithm is the Kalman filter.

The alternative which we have pursued here is to recover a continuous control signal. The closed-loop control task is illustrated in Figure 7 where the brain controls a 2D cursor position on a computer screen and a monkey (or human) receives visual feedback by viewing the cursor on a monitor. We suspect that for robot control tasks (e.g. moving a wheelchair or robot arm) that continuous control will be preferable because it is inherently more flexible. It is also, however, more noise prone so there is a tradeoff that gives higher spatial resolution with less accuracy. The tradeoffs between discrete and continuous methods and their relevance for rehabilitation applications deserve further study.

One promising direction combines discrete and continuous control in a single interface [31]. The Bayesian decoding framework can easily accommodate a mixed state space with both continuous (2D) and discrete (task-orientated) parameters. The generative model then involves first selecting the task (continuous or discrete) and then generating the observations conditioned on the task. Decoding is slightly more complicated but can be achieved using a switching Kalman filter [33] or particle filter [2, 6]. Recently, Wood *et al.* [31] used such an approach to decode whether or not a monkey was performing a 2D control task and, if so, to decode the hand state with a linear Gaussian model. Such an approach holds promise for flexible brain-machine interfaces in which the user can switch between a variety of functions or control modes.

0.6 Discussion and Conclusions

The probabilistic modeling of the neural code presents many challenges. Beyond the simple linear Gaussian models explored here there is likely an advantage in modeling the non-Gaussian and non-linear nature of neural activity [7, 13, 33]. Beyond firing rates we may wish to formulate probabilistic models of spike trains [26]. Efficient learning and decoding methods however do not currently exist for non-Gaussian, non-linear models of point processes. There is an opportunity here to develop new machine learning methods for capturing the high dimensional relationship between motor behavior and neural firing.

Moreover, here we only consider information from primary motor cortex. Additional information may be obtained from pre-motor and parietal areas. The Bayesian framework we have proposed provides a solid foundation on which to integrate sources of information from various brain areas in a principled way.

The approach does not, however, necessarily provide any new insight into how the brain controls movement. Like the approaches it generalizes (e.g. the population vector method), the relationships between firing rates and kinematics are purely descriptive. One cannot infer, for example, that the brain is somehow implementing a Kalman filter. Rather, all these methods describe attributes of the neural computation and not the computation itself.

This paper only hints at the prosthetic applications of these methods. While Bayesian methods have been used for closed-loop neural control of cursors by monkeys [34], the use of this or any decoding method in paralyzed humans remains to be explored. Particularly important in the case of paralyzed humans will be the issue of training and adaptation. Training data for the encoding model here, for example, will have to rely on imagined movement. Whether human users will be able to adapt their neural signals to improve control with a given decoder remains to be seen and may prove critical for practical motor-cortical control of devices.

While current methods provide a proof of concept that cortical implants can provide reliable control signals over extended periods of time there is still much work to be done. Current continuous decoding results still provide a somewhat "jerky" reconstruction – new decoding/control algorithms for damping the cursor reconstruction may enable a wider range of applications. The great challenge however is to move beyond simple 2D or 3D cursor control to ultimately give patients high-dimensional control of devices such as dexterous robot hands.

Acknowledgments

The research summarized here is the product of interdisciplinary research with a large number of collaborators who we wish to thank. They include Elie Bienenstock, Matthew Fellows, Mijail Serruya, Yun Gao, Wei Wu, Frank Wood, Jessica Fisher, Shy Shoham, Carlos Vargas-Irwin, Ammar Shaikhouni, David Mumford, Arto Nurmikko, Beth Travers, Gerhard Friehs, and Liam Paninski.

This work was funded by the DARPA BioInfoMicro Program, NIH NINDS Neural Prosthesis Program and Grant NS25074, NIH-NINDS N01-NS-2-2345, NSF ITR award 0113679, NIH-NINDS R01 NS 50967-01 as part of the NSF/NIH Collaborative Research in Computational Neuroscience Program, the Veteran's Administration grant #A3772C, and the Office of Naval Research award N0014-04-1-082.

References

1. *The Forbidden Planet.* MGM, 1956.
2. A. E. Brockwell, A. L. Rojas, and R. E. Kass. Recursive bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology*, 91:1899–1907, 2004.
3. J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS, Biology*, 1(2):001–016, 2003.
4. M. Crichton. *The Terminal Man.* Knopf, 1972.
5. D. Flament and J. Hore. Relations of motor cortex neural discharge to kinematics of passive and active elbow movements in the monkey. *Journal of Neurophysiology*, 60(4):1268–1284, 1988.
6. Y. Gao, M. J. Black, E. Bienenstock, S. Shoham, and J. P. Donoghue. Probabilistic inference of hand motion from neural activity in motor cortex. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 213–220, Cambridge, MA, 2002. MIT Press.
7. Y. Gao, M. J. Black, E. Bienenstock, W. Wu, and J. P. Donoghue. A quantitative comparison of linear and non-linear models of motor cortical activity for the encoding and decoding of arm motions. In *1st International IEEE/EMBS Conference on Neural Engineering*, pages 189–192, Capri, Italy, March 2003.
8. A. Georgopoulos, J. Kalaska, R. Caminiti, and J. Massey. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 8(11):1527–1537, 1982.
9. A. Georgopoulos, A. Schwartz, and R. Kettner. Neural population coding of movement direction. *Science*, 233:1416–1419, 1986.
10. S. Helms Tillery, D. Taylor, R. Isaacs, and A. Schwartz. Online control of a prosthetic arm from motor cortical signals. In *Society for Neuroscience Abst.*, volume 26, 2000.
11. R. E. Kalman. A new approach to linear filtering and prediction problems. *Trans. ASME, Journal of Basic Engineering*, 82:35–45, 1960.
12. R. Kettner, A. Schwartz, and A. Georgopoulos. Primary motor cortex and free arm movements to visual targets in three-dimensional space. iii. positional gradients and population coding of movement direction from various movement origins. *Journal of Neuroscience*, 8(8):2938–2947, 1988.
13. S-.P. Kim, J. C. Sanchez, Y. N. Rao, D. Ergodmus, J. C. Principe, J. M. Carmena, M. A. Lebedev, and M. A. L. Nicolelis. A comparison of optimal MIMO linear and nonlinear models for brain-machine interfaces. *Journal of Neural Engineering*, submitted, 2005.
14. R. Lauer, P. Peckham, K. Kilgore, and W. Heetderks. Applications of cortical signals to neuroprosthetic control: A critical review. *IEEE Trans. Rehab. Eng.*, 8(2):205–208, June 2000.
15. E. Maynard, C. Nordhausen, and R. Normann. The Utah intracortical electrode array: A recording structure for potential brain-computer interfaces. *Electroencephalography and Clinical Neurophysiology*, 102:228–239, 1997.
16. D. Moran and A. Schwartz. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82(5):2676–2692, 1999.

17. S. Musallam, B. D. Corneil, B. Greger, H. Scherberger, and R. A. Andersen. Cognitive control signals for neural prosthetics. *Science*, 305(568), 2004.
18. L. Paninski, M. Fellows, N. Hatsopoulos, and J. P. Donoghue. Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology*, 91:515–532, 2004.
19. F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code*. MIT Press, 1999.
20. A. Schwartz, R. Kettner, and A. Georgopoulos. Primary motor cortex and free arm movements to visual targets in three-dimensional space. i. relations between single cell discharge and direction of movement. *Journal of Neuroscience*, 8(8):2913–2927, 1988.
21. A. Schwartz, D. Taylor, and S. Helms Tillery. Extraction algorithms for cortical control of arm prosthetics. *Current Opinion in Neurobiology*, 11:701–707, 2001.
22. M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. P. Donoghue. Brain-machine interface: Instant neural control of a movement signal. *Nature*, 416:141–142, 2002.
23. K. V. Shenoy, D. Meeker, S. Cao, S. A. Kureshi, B. Pesaran, C. A. Buneo, A. P. Batista, P. P. Mitra, J. W. Burdick, and R. A. Andersen. Neural prosthetic control signals from plan activity. *NeuroReport*, 14(4):591–597, March 2003.
24. C. Siodmak. *Donovan's Brain*. Knopf, 1942.
25. D. Taylor, S. Helms Tillery, and A. Schwartz. Direct cortical control of 3D neuroprosthetic devices. *Science*, 296(5574):1829–1832, 2002.
26. W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown. A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *Journal of Neurophysiology*, 93:1074–1089, 2005.
27. D. Trumbo. *Johnny Got His Gun*. Bantam Windstone, 1982.
28. G. Welch and G. Bishop. An introduction to the Kalman filter. Technical Report 95–041, University of North Carolina at Chapel Hill, 2001.
29. J. Wessberg, C. Stambaugh, J. Kralik, Laubach M. Beck, P., J. Chapin, J. Kim, S. Biggs, M. Srinivasan, and M. Nicolelis. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408:361–365, 2000.
30. F. Wood, M. J. Black, C. Vargas-Irwin, M. Fellows, and J. P. Donoghue. On the variability of manual spike sorting. *IEEE Transactions on Biomedical Engineering*, 51(6):912–918, 2004.
31. F. Wood, Prabhat, J. P. Donoghue, and M. J. Black. Inferring attentional state and kinematics from motor cortical firing rates. In *27th Int. Conf. of the IEEE Engr. in Med. and Bio. Soc.*, pages 149–152, September 2005.
32. F. Wood, S. Roth, and M. J. Black. Modeling neural population spiking activity with gibbs distributions. In *Advances in Neural Information Processing Systems 15*. MIT Press, to appear, 2005.
33. W. Wu, M. J. Black, D. Mumford, Y. Gao, E. Bienenstock, and J. P. Donoghue. Modeling and decoding motor cortical activity using a switching Kalman filter. *IEEE Transactions on Biomedical Engineering*, 51(6):933–942, 2004.

34. W. Wu, Y. Gao, E. Bienenstock, J. P. Donoghue, and M. J. Black. Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18:80–118, 2005.

35. W. Wu, A. Shaikhouni, J. P. Donoghue, and M. J. Black. Closed-loop neural control of cursor motion using a Kalman filter. In *Proc. IEEE Engineering in Medicine and Biology Society*, pages 4126–4129, Sep 2004.