

QUERYING THE USER PROPERLY FOR HIGH-PERFORMANCE BRAIN-MACHINE INTERFACES: RECURSIVE ESTIMATION, CONTROL, AND FEEDBACK INFORMATION-THEORETIC PERSPECTIVES

Cyrus Omar^{*}, Miles Johnson[†], Timothy W. Bretl[‡], Todd P. Coleman[§]

University of Illinois

ABSTRACT

We propose a complementary approach to the design of neural prosthetic interfaces that goes beyond the standard approach of estimating desired control signals from neural activity. We exploit the fact that for a user's intended application, the dynamics of the prosthetic in fact impact subsequent desired control inputs. We illustrate that changing the dynamic response of a prosthetic device can make specific tasks significantly easier to accomplish. Our approach relies upon principles from stochastic control and feedback information theory, and we illustrate its effectiveness both theoretically and experimentally - in terms of spelling words from a menu of characters using binary surface electromyography classification.

Index Terms— neural prosthetics, feedback information theory, stochastic control, interface design

1. INTRODUCTION

Neural interfaces use estimates of brain or muscle activity to generate control inputs for a prosthetic device. Many different sensors are now available to measure neural activity, both invasive (such as intracortical devices that observe ensemble spiking of individual neurons) and non-invasive (such as electroencephalography and electromyography). These sensors have been used to control a growing number of prosthetic devices that include computer cursors, cell phones, robotic manipulators, and wheelchairs. Applications include both restoring lost function (for example, with a neuroprosthetic limb) and enhancing or augmenting normal function (for example, with subvocal speech in noisy environments).

One significant challenge in the design of neural interfaces is the fundamental uncertainty about user intent. In particular, the only way that the user can communicate their intent to the prosthetic is through sensed neural activity, which is typically very noisy. As a consequence, most previous work has focused on improving the measurement and estimation of neural activity.

^{*}C. Omar (comar2@uiuc.edu) is with the CS Department, University of Illinois, Urbana, IL.

[†]M. Johnson (mjohns50@uiuc.edu) is with the Aerospace Engineering Department, University of Illinois, Urbana, IL.

[‡]T. Bretl (tbretl@uiuc.edu) is with the Aerospace Engineering Department, University of Illinois, Urbana, IL. Support provided by a Seed Grant from the Center for Healthy Minds, funded through NIH/NIA under Award No. P30-AG023101.

[§]T. Coleman (colemant@uiuc.edu) is with the ECE Department and Neuroscience Program, University of Illinois, Urbana, IL. Support provided by a Seed Grant from the Center for Healthy Minds, funded through NIH/NIA under Award No. P30-AG023101.

Here we take a complementary approach, focusing on adaptively providing sensory feedback to query the user in order to generate better control inputs. We view a neural interface as a dynamic system connecting a user with a prosthetic, described using the framework of stochastic control (Section 3). The policy we develop has roots in feedback information theory. We model user behavior in the context of specific tasks, and change the actuation policy of the interface to make it easier for the user to accomplish these tasks. The example we provide here is not particularly profound. However, our plan here is to motivate the use of this new degree of freedom by showing in some simple settings how significant performance enhancement can ensue. Specifically, we show here that system performance can be improved by a control policy that takes advantage of how user intent itself depends on the state of the prosthetic (Section 4). This dependence effectively allows the interface to have more information about intent than from measurement of neural activity alone.




Fig. 1. A screenshot of the text interface.

We demonstrate our approach in initial experiments with a text interface (Section 5), shown in Figure 1. The neural sensor is a surface electromyography device, trained to recognize sub-vocal expression of the words “left” and “right.” The prosthetic is a 1-D cursor, used to select characters from a menu of fixed length. The menu consists of the 26 letters of the English alphabet and a space character (displayed as an underscore), equally spaced along a line. The cursor is placed on this line, its position updated in response to neural activity. The task is to spell a sequence of words as quickly and accurately as possible. Results are presented for one female human subject. These results demonstrate that, even without enhancing the measurement of neural activity, the performance of a neural interface can be significantly improved. Although this observation - in terms of our example - may be trivial to some, it does not appear to have been readily exploited in the design of neural prostheses. Here we illustrate this simple observation to draw attention to the paradigm; we are currently investigating non-trivial extensions, such as incorporating this dynamic querying into the

neural control of continuous-movement prostheses. We also note that these results can clearly be improved by exploiting the statistical structure of language, similar to the “Dasher” system [1]. Such a technique can naturally be tied into our approach in a complementary fashion; indeed, our research group has begun to exploit the structure of language to enhance performance of P300 communication prostheses [2] using stochastic control.

2. BACKGROUND

2.1. Neural interfaces

Neural interfaces use measurements of neural signals to control computers or machines. In general, these interfaces are classified by whether they use invasive (intracortical sensors) methods [3, 4] or non-invasive (*electroencephalograph*, EEG, and the *electromyograph*, EMG) methods [5, 6, 7] to extract measurements.

2.2. Communication using surface electromyography

The interface we focus on in this paper uses surface-EMG to capture *subvocalization* or “silent speech.” Subvocal speech recognition uses analysis of muscle activity in the tongue and throat to determine what someone is trying to say, even if it is not possible for the speaker to produce sounds. A lot of work has been done to improve signal processing and speech recognition for surface-EMG [7]. Some of this work has begun to integrate small, portable EMG sensors with prosthetic devices [8]. Our work focuses instead on improving performance by changing the dynamic response of the prosthetic device, rather than by enhancing the recognition of words.

2.3. Text interfaces

In this paper we enable neural control of a text interface, a common application of surface-EMG sensors. However, unlike most other work with surface-EMG, we do not try to identify a large vocabulary of spoken words or letters. Instead, we only allow the user to give a binary input, by saying either “left” or “right.” This binary signal moves a 1-D cursor through a menu of characters. We take this approach because we want to show what is possible even with a limited neural signal. In this respect our text interface is more strongly related to those that use EEG sensors. For example, the P300 spelling paradigm measures *event-related potentials*, characteristic responses to stimuli that can be interpreted as binary (on/off) signals [9]. It randomly illuminates on-screen characters and selects the one that elicits the strongest response. The performance of this type of interface (in this case measured by the number of words per minute) is strongly affected by the size, shape, and arrangement of the on-screen menu, which can be designed based on the statistical structure of language [10, 1]. Our work is similar, but focuses on changing how a 1-D cursor moves through a menu of fixed size and shape.

2.4. Solution approach

We model a neural interface as a discrete-time dynamic system and focus on optimizing the dynamic response of this system using a stochastic control framework [11]. Our goal in this paper is to choose the response of a 1-D cursor to neural activity in a way that minimizes the number of binary signals needed to select a single character from a fixed menu. In

Section 4 we will see that this boils down to estimating which character the user wants to select based on their behavior over time.

The technique we used to address our problem, *probabilistic bisection*, has roots in the field of information theory. It was originally suggested [12, 13] as a way to perform communication across a binary symmetric channel with noiseless feedback. These works have subsequently been generalized [14] to indicate how feedback coding strategies should be developed for arbitrary noisy channels: *sensory feedback should provide information about the posterior distribution on user intent (that is, the interface’s belief about user intent), and measured neural activity should be interpreted as steering this distribution*. Communication across noisy channels with noiseless feedback has become a fundamental component to reasoning about “control over noisy channels”, an active area of research that unifies information theory and control theory. More recently this strategy has been applied to machine learning, specifically to the technique of *active learning* (also known as adaptive sampling) [15].

3. FRAMEWORK

In this paper we model a *neural interface* as a discrete-time dynamic system connecting a *user* with a *prosthetic* (Fig. 2). At each time step k , six variables describe the state of this system:

θ	user intent
x_k	desired control input to the prosthetic
y_k	measurement of neural activity
u_k	actual control input to the prosthetic
p_k	state of the prosthetic
f_k	sensory feedback.




Fig. 2. A neural interface is an interconnection between two dynamic systems.

For example, if the prosthetic is a robotic manipulator, then we might define θ as the intended trajectory, x_k as a set of desired joint angles, y_k as neural spiking in a region of motor cortex, u_k as a set of joint torques, p_k as the actual angle and angular velocity of each joint, and f_k as the position of the end-effector. To be concise, we in general denote $z^k \triangleq (z_0, \dots, z_k)$. These six variables evolve according to four generative statistical models

behavior	$x_{k+1} \sim P(x_{k+1} x^k, f^k, \theta)$
observation	$y_k \sim P(y_k x_k)$
actuation	$p_{k+1} \sim P(p_{k+1} p_k, u_k)$
feedback	$f_k \sim P(f_k p_k)$

and one time-varying deterministic function

$$\text{policy } u_k = \mu_k(y^k, p^k, u^{k-1}, f^{k-1}).$$

Notice that both user intent and the desired control input can only be observed indirectly, through the noisy measure-

ment y_k . We assume that this measurement depends neither on the state of the prosthetic nor on any previous intent x^{k-1} . We abbreviate by

$$I_k = (y^k, p^k, u^{k-1}, f^{k-1})$$

the vector of all information available when selecting a control input u_k , where $I_0 = (y_0, p_0)$. In this paper, we will assume that models of behavior, observation, and actuation are given. For simplicity, we will also assume that feedback is full-state and error-free, so $f_k = p_k$. Then, we will be interested in choosing a policy $\pi = \{\mu_0(I_0), \dots, \mu_{N-1}(I_{N-1})\}$ that optimizes some performance metric, for example a *minimax* cost of the form

$$J_\pi = \max_{\theta, x_0, p_0} E \{g(x_N, p_N)\}.$$

We will show that performance can be improved by a policy that takes advantage of how the user's desired control input depends on the state of the prosthetic. In particular, we will use this dependence to better estimate the underlying user intent θ .

4. APPLICATION TO THE TEXT INTERFACE

Consider the problem of selecting a single letter from a menu, using the text interface described in Section 1. The prosthetic is a 1-D cursor that highlights letters in the menu. The neural sensor is a surface-EMG device, trained to recognize subvocal expression of the words *left* and *right*. Whenever the surface-EMG device detects a word, the cursor moves in response; otherwise, it remains motionless. If the cursor highlights the same letter for a significant length of time, that letter is selected, and the process repeats.

We model this problem using our framework from Section 3. Assume the menu of letters has length n . We refer to each letter by its position and describe the menu by the ordered set $M = (1, \dots, n)$. We refer to the words *left* and *right* by the integers -1 and $+1$, respectively. The state of the prosthetic is the current position $p_k \in M$ of the cursor. The control input is the amount $u_k \in \mathbb{Z}$ that the cursor moves at each time step k , where we require that $p_k + u_k \in M$. The measurement is a word $y_k \in \{-1, +1\}$. The feedback is the entire state $f_k = p_k$. The user's desired control signal is the movement direction $x_k \in \{-1, +1\}$. Finally, the user's underlying intent is a letter $\theta \in M$, which we assume remains fixed until it is selected.

We assume that these variables evolve according to the following four models:

$$\text{behavior } x_{k+1} = \begin{cases} \text{sign}(\theta - f_k) & \text{w/probability } 1 - \alpha \\ -\text{sign}(\theta - f_k) & \text{w/probability } \alpha \end{cases}$$

$$\text{observation } y_k = \begin{cases} x_k & \text{w/probability } 1 - \beta \\ -x_k & \text{w/probability } \beta \end{cases}$$

$$\text{actuation } p_{k+1} = p_k + u_k$$

$$\text{feedback } f_k = p_k$$

The model of behavior says that the user always wants to move the cursor toward the desired letter, but has some chance $0 \leq \alpha < 1/2$ of making a mistake. The model of observation says that our measurement of the desired movement direction also has some chance $0 \leq \beta < 1/2$ of being wrong. The model of actuation says that the cursor moves exactly the distance specified by the interface. The model of feedback says that the user perfectly observes the current position of the cursor.

Our goal is to select a policy

$$\pi = \{\mu_0(I_0), \dots, \mu_{N-1}(I_{N-1})\}$$

for computing $u_k = \mu_k(I_k)$ at each time step k that minimizes the cost

$$J_\pi = \max_{\theta, p_0 \in M} E \{|\theta - p_N|\}.$$

over a finite horizon N . In other words, our goal is to bound the worst-case expected error in cursor position, regardless of the desired letter. In particular, we will be interested in the rate at which this bound decreases with N , for a given policy π .

Here we have chosen a *minimax* cost primarily because it allows us to give performance guarantees. Other costs (such as min-time, sum-squared error, or error probability) lead to similar results in practice. If we do not take into account how intent depends on the state of the prosthetic, we can do no better on average than moving one step in the direction of y_k after every measurement, as in

$$\text{policy } \pi \quad u_k = y_k.$$

So at each time step k , the cursor will move in the wrong direction (away from θ) with probability $\gamma = (1 - \alpha)\beta + \alpha(1 - \beta)$ and the expected distance moved toward the target after k steps is at least $k(1 - 2\gamma)$. So we have

$$J_\pi \leq (n - 1) - N(1 - 2\gamma)$$

for all $N < n - 1$, and the cost of this policy decreases *linearly* with N .

If we instead take advantage of this dependence, we can achieve better performance. In particular, assume that our estimate of θ begins uniformly distributed in the ordered set $M = (l_0, \dots, r_0)$ where we define $l_0 = 1$ and $r_0 = n$. Assume that at time $k - 1$, our estimate of θ has some distribution $P(\theta|y_0, \dots, y_{k-1})$ over the finite set M . Then after measuring $y_k = +1$, we can update this distribution with Bayes' rule to find

$$P(\theta|y^k) = \eta \cdot \begin{cases} (1 - \gamma) \cdot P(\theta|y^{k-1}), & \forall \theta \in (p_{k+1}, \dots, n) \\ \gamma \cdot P(\theta|y^{k-1}) & \forall \theta \in (1, \dots, p_k) \end{cases}$$

The update for $y_k = -1$ is analogous. We adopt a *probabilistic bisection* policy, which we denote by π' , and choose the input u_k that places p_{k+1} at the median of the posterior distribution over θ , so u_k satisfies both:

$$\sum_{i=1}^{p_k + u_k - 1} p(\theta = i|y^k) \leq \frac{1}{2} \quad \text{and} \quad \sum_{i=p_k + u_k}^n p(\theta = i|y^k) > \frac{1}{2}.$$

Then we know from [13, 15, 14] that this policy has cost bounded by

$$J_{\pi'} \leq n2^{-cN}$$

over a finite horizon N , for some constant $c > 0$. In other words, the cost of this policy once again decreases *exponentially* with N .

Note that one way to interpret this policy is from a feedback information-theory perspective: the user's desired control input x_k at each time k is independent of previous feedback signals y_1, \dots, y_{k-1} , so the query of the user is such that his response is *the most informative*. This is not a coincidence, and is fundamental to developing good control schemes in the presence of noisy communication channels [16],[14, SecII]. This scheme can be extended for virtually




Fig. 3. Number of commands required per character, computed as a running average with a 50-character horizon. The dotted line is the average over all 200 characters.

any neural interface pertaining to choices [14] over a finite alphabet. It can be used even more generally for systems with continuous dynamics - although it is not always guaranteed to be the optimal solution to a stochastic control problem (which in general suffers from the curse of dimensionality). Nonetheless, its operation is intimately tied to recursive estimation - in that the optimal current query position can be expressed as a function of the previous query position and the most recent observation [14, Lemma 1] - and thus has computational ease. We also espouse using this approach more generally because of its information-theoretic performance guarantees (i.e. error exponents).

5. RESULTS

5.1. Experimental setup

In the previous section we presented two control policies (“fixed offset” and “probabilistic bisection”) for a 1-D menu-based text interface. The first policy moves one step to the left or right in response to neural activity; the second policy moves a variable distance based on an estimate of the desired character. These policies were evaluated in experiments with one healthy 20-year old female volunteers with normal vision.

5.1.1. Signal acquisition

In order to extract neural signals, we use a commercially available, self-contained, single-channel surface-EMG sensor designed to be worn around the neck (“The Audeo”). This sensor amplifies, filters, digitizes and transmits the EMG signal to a computer via a standard USB interface. Software available with the device allows a user to train it to recognize subvocal expression of the words “left” and “right.” This sensor was originally designed to enable a paralyzed patient to drive a motorized wheelchair. Additional information regarding the device can be found at <http://www.theaudeo.com/tech.html>.

5.1.2. Online tests

The subject was asked to write the same passage using surface-EMG that they had written using keyboard input. No additional noise was added to the neural signal. They used the fixed-step policy first, then repeated the experiment with the bisection policy. The length of time before letter acceptance was increased to 4 seconds for both policies, in order to capture the additional amount of time it takes to generate a signal with surface-EMG as compared to keyboard input.

5.2. Data and analysis

We measured subject performance as the number of characters selected per minute. Figure 3 shows a running average over all trials with a 50-letter horizon, along with the number of neural commands given, again as a running average. It shows that probabilistic bisection improved performance, allowing significantly higher average throughput than fixed offset (by an average of 112%). Figure 3 also shows that there was also a significant difference between policies in the effort the subject had to exert, measured as the number of commands required to generate a character (by an average of 63%).

6. CONCLUSION

In this paper we considered the design of neural interfaces, which use estimates of brain or muscle activity to generate control inputs for a prosthetic device. We showed that changing the dynamic response of the prosthetic can make specific tasks easier to accomplish, even without improving the measurement of neural activity. We verified our results in experiments with a text interface, where neural activity was measured using surface electromyography and the prosthetic was a 1-D cursor used to spell words from a menu of characters. We are currently extending our approach to a wider array of neural sensors and prosthetic devices, in particular devices with more complex dynamics.

Acknowledgements

The authors would like to thank Ed Maclin and Michael Callahan for their helpful discussions and Katie Snell and David Rockwood for their assistance with the experiments.

References

- [1] S. A. Wills and D. J. C. Mackay, “DASHER—an efficient writing system for brain-computer interfaces,” vol. 14, no. 2, pp. 244–246, Jun. 2006.
- [2] J. A. Jarzebowski, L. Srinivasan, and T. P. Coleman, “A stochastic control approach to variable length menus in p300 neural communication prostheses,” in *Society for Neuroscience Annual Meeting*, San Diego, CA, Nov. 2007, selected for Oral Presentation. In preparation for journal submission.
- [3] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue, “Neuronal ensemble control of prosthetic devices by a human with tetraplegia,” *Nature*, vol. 442, pp. 164–171, July 2006.
- [4] S. G. R. SI, Y. BM, A. A, and S. KV, “A high-performance brain-computer interface,” *Nature*, vol. 442, pp. 195–198, 2006.
- [5] J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan, “Brain-computer interfaces for communication and control,” *Clinical Neurophysiology*, vol. 113, no. 6, pp. 767–791, 2002.
- [6] M. Thulasidas, C. Guan, and J. Wu, “Robust classification of EEG signal for brain-computer interface,” vol. 14, no. 1, pp. 24–29, Mar. 2006.
- [7] L. Maier-Hein, “Speech recognition using surface electromyography,” Ph.D. dissertation, Universität Karlsruhe, Karlsruhe, Deutschland, Jul. 2005.
- [8] C. C. Jorgensen and K. Binsted, “Web browser control using EMG based sub vocal speech recognition,” in *Int. Conf. System Sciences*, 2005.
- [9] L. Farwell and E. Donchin, “Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials,” *Electroencephalogr. Clin. Neurophys.*, vol. 70, pp. 510–523, 1988.
- [10] H. Serby, E. Yom-Tov, and G. F. Inbar, “An improved P300-based brain-computer interface,” vol. 13, no. 1, pp. 89–98, Mar. 2005.
- [11] D. Bertsekas, *Dynamic programming and optimal control*, 3rd ed. Athena Scientific, 2005, vol. 1.
- [12] M. Horstein, “Sequential transmission using noiseless feedback,” *IEEE Transactions on Information Theor.*, vol. 9, pp. 136–143, Jul. 1963.
- [13] M. Burnashev and K. Zigangirov, “An interval estimation problem for controlled observations,” *Problems in Information Transmission*, vol. 10, pp. 223–231.
- [14] O. Shayevitz and M. Feder, “Communication with feedback via posterior matching,” in *IEEE Int. Symp. Info. Th.*, Nice, France, Jun. 2007.
- [15] R. M. Castro and R. D. Nowak, “Active learning and sampling,” in *Foundations and Applications of Sensor Management*. Springer-Verlag, 2007.
- [16] N. Elia, “When Bode meets Shannon: control-oriented feedback communication schemes,” *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1477–1488, Sep. 2004.