

Policies for neural prosthetic control: initial experiments with a text interface

Cyrus Omar, Miles Johnson, Timothy Bretl, and Todd Coleman

Abstract—Neural interfaces use estimates of brain or muscle activity to generate control inputs for a prosthetic device. Most previous work focuses on estimating neural activity more accurately. This paper focuses on generating better control inputs. It shows that changing the dynamic response of a prosthetic device can make specific tasks easier to accomplish. It also presents experimental results for which neural activity is measured using surface electromyography, the prosthetic is a 1-D cursor, and the task is to spell words from a menu of characters.

I. 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 noisy and low-bandwidth. As a consequence, most previous work has focused on improving the measurement and characterization of neural activity.

In this paper we take a complementary approach, focusing on generating 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 III). We model user behavior in the context of specific tasks, and change the dynamic response of the interface to make it easier for the user to accomplish these tasks. In particular, we show 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 IV). This dependence effectively allows the interface to have more

This work was partially supported by the Pilot Grant Program of the Center for Healthy Minds, funded through NIH/NIA under Award No. P30-AG023101, University of Illinois at Urbana-Champaign.

C. Omar is with the Department of Computer Science, M. Johnson and T. Bretl are with the Department of Aerospace Engineering, and T. Coleman is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL. {comar2,mjohns50,tbretl,colemant}@uiuc.edu




Fig. 1. A screenshot of the text interface.

information about intent than from measurement of neural activity alone.

We demonstrate our approach in initial experiments with a text interface (Section V), shown in Figure 1. This interface is similar to an existing and highly successful one called “Dasher” [26]. 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. These characters are not equally spaced—instead, the portion of the menu occupied by each character is determined by its conditional probability. This probability increases the longer the cursor remains on any given character. As the corresponding interval grows large, it is recursively subdivided into a second menu of characters, then a third, etc. We use a simple bigram frequency table to generate conditional probabilities given previous characters, though more sophisticated language models can be accommodated. Effectively, this interface maps each possible sequence of characters to a point on the unit interval.

The Dasher interface was designed for use with continuous 1-D or 2-D inputs. Experienced users can type up to 35 words per minute using a mouse, 25 words per minute using an eye tracker, and 16 words per minute using a “breath mouse,” which measures the movement of a user’s chest as they breathe [26]. Recent work has integrated Dasher with a neural sensor that uses electroencephalography to generate continuous 1-D signals [11]. However, the reported typing speed falls dramatically, to about 1-2 words per minute, or about 5-10 characters per minute. Currently, this level of performance is typical of other neural text interfaces.

In the experiments we describe here, the neural sensor is a surface electromyography device, trained to recognize subvocal expression of the words “left” and “right.” These words are recognized correctly only 70-80% of the time, which again is typical of neural interfaces. Rather than emulate the Dasher interface—for example, moving the cursor a fixed distance left or right in response to the sensed neural command—we chose a bisection strategy, moving the cursor to the median of the posterior distribution. Our contribution

was not to develop this strategy, which in any case is intuitive and has been used for other applications. Instead, our contribution is in using this strategy as an example to show, in both theory and experiment, that even without enhancing the measurement of neural activity, the performance of a neural interface can be significantly improved by changing the dynamic response of the prosthetic device. In particular, we show that the expected error in character selection—measured as the distance between the desired character and the cursor—decays exponentially rather than linearly.

We previously published a companion paper to this one, which describes many of the same concepts [18]. We extend our previous work here by using the Dasher-like text interface, which has improved our results significantly.

II. BACKGROUND

A. Neural interfaces

The basic communication nodes in the human nervous system are *neurons*, which transmit information about the sensory system, cognition, and motor planning. *Neural interfaces* use measurements of these neural signals to control computers or machines. These interfaces are classified by whether they use invasive or non-invasive methods to extract measurements. *Invasive* methods use intracortical sensors, often placed in the primary motor cortex, that can observe ensemble spiking of individual neurons. For example, these methods have enabled primates to move robotic arms [22], [25] and tetraplegic human patients to move a cursor and read email [12]. *Non-invasive* methods place electrodes on the surface of the skin to record signals such as the *electroencephalograph* (EEG) and the *electromyograph* (EMG), which reflect activity from large groups of neurons and muscle fibers, respectively. For example, these methods have enabled both cursor-control and text entry [9], [23], [24], [27].

B. 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. In some cases it may not be possible for a speaker to produce sounds, for example if they are paralyzed below the neck, or for a listener to hear sounds, for example if they are working in a noisy environment. However, in these cases the muscle activity in the tongue and throat is still present, and is correlated with the intended speech. Work has been done to improve signal processing and speech recognition for surface-EMG [16]. Some of this work has begun to integrate small, portable EMG sensors, like the one we use, with prosthetic devices [3], [14], [15]. Our work focuses instead on improving performance by changing the dynamic response of the prosthetic device, rather than by enhancing the recognition of words.

C. 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 signals [10]. 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 [19], [26]. Our work is similar, but focuses on changing how a 1-D cursor moves through a menu of fixed size and shape.

D. 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 [2]. Our goal in this paper is to choose the response of a 1-D cursor to neural activity in a way that maximizes the number of characters per minute, or equivalently, words per minute, that can be typed by the user. This problem has roots in the field of information theory. Our solution approach, *probabilistic bisection*, has been suggested as a way to do sequential coding [13], and later modified to provide guaranteed performance bounds [5]. More recently this strategy has been applied to machine learning, specifically to *active learning* [4], [6], [7], [21]. This work has also been revisited and generalized to indicate precisely how to design feedback strategies that maximize performance measures [8], [17], [20]. In particular, sensory feedback should provide information about the posterior distribution on user intent, in other words the interface’s belief about user intent, and measured neural activity should be interpreted as steering this distribution.

III. 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.

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




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

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. These six variables evolve according to four generative statistical models

$$\text{behavior } x_{k+1} \sim P(x_{k+1}|x_0, \dots, x_k, f_0, \dots, f_k, \theta)$$

$$\text{observation } y_k \sim P(y_k|x_k)$$

$$\text{actuation } p_{k+1} \sim P(p_{k+1}|p_k, u_k)$$

$$\text{feedback } f_k \sim P(f_k|p_k)$$

and one time-varying deterministic function

$$\text{policy } u_k = \mu_k(y_0, \dots, y_k, p_0, \dots, p_k, u_0, \dots, u_{k-1}, f_0, \dots, f_{k-1}).$$

Notice that both user intent and the desired control input can only be observed indirectly, through the noisy measurement y_k . We assume that this measurement depends neither on the state of the prosthetic nor on any previous intent x_0, \dots, x_{k-1} . We abbreviate by

$$I_k = (y_0, \dots, y_k, p_0, \dots, p_k, u_0, \dots, u_{k-1}, f_0, \dots, f_{k-1})$$

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. Our ultimate goal in this paper is to maximize the number of characters per minute that a user can type. As a heuristic, we will choose as a metric to minimize an expected 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. We will use this dependence to better estimate the underlying user intent θ .

IV. APPLICATION TO THE TEXT INTERFACE

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

We model this problem using our framework from Section III. 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 . We are using this metric as a heuristic to maximize the number of characters per minute that can be typed by a user. We choose this metric in particular because it allows us to give performance guarantees.

A. Noiseless case

In this case, we assume that the user never makes a mistake and that our measurement of neural activity is perfect, so both $\alpha = 0$ and $\beta = 0$. As a result, we know the desired movement direction $x_k = y_k$ exactly. So a reasonable control policy, which we call *fixed offset*, might be to move one step in this direction after every measurement:

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

This policy is easy both to understand and to implement. It is clear that over a finite horizon N , the cost of this policy is bounded by

$$J_\pi \leq \begin{cases} \max_{\theta, p_0 \in M} (|\theta - p_0| - N) & N < n - 1 \\ 0 & \text{otherwise} \end{cases}$$

$$= \begin{cases} (n - 1) - N & N < n - 1 \\ 0 & \text{otherwise.} \end{cases}$$

So, the cost of this policy decreases *linearly* with N .

However, recall that user intent depends on the state of the prosthetic. In particular, our model of behavior assumes that the desired movement direction is always toward the desired letter *from the current position of the cursor*. So every measurement gives us information not only about x_k , but also about θ . This information is disregarded by the “reasonable” policy we just described. Equivalently, we can say that this policy was derived under the assumption that user intent does not depend on the state of the prosthetic, a common simplification.

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 $r_0 = n$. We know that after each measurement, our estimate of θ will be uniformly distributed in an interval $(l_k, r_k) \subset (l_0, r_0)$ of equal or lesser width. We adopt a *bisection* policy and choose the input according to

$$\text{policy } \pi' \quad u_k = \begin{cases} \frac{l_k + r_k}{2} - p_k & \text{if } l_k + r_k \text{ is even} \\ \frac{l_k + r_k \pm 1}{2} - p_k & \text{if } l_k + r_k \text{ is odd,} \end{cases}$$

where we flip a coin to choose $+1$ or -1 in the latter case. It is easy to show that over a finite horizon N , the cost of this policy is bounded by

$$J_{\pi'} \leq \begin{cases} n2^{-N} & N < (\log n / \log 2) \\ 0 & \text{otherwise.} \end{cases}$$

So, the cost of this policy decreases *exponentially* with N .

B. Noisy case

In the noiseless case, we showed that a change in the control policy could improve performance, even without enhancing our measurement of neural activity. Here, we show an identical result when both α and β are non-zero. Even in this case, when both behavior and measurements are noisy, we can do no better on average than moving one step in the direction of y_k after every measurement, as in

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

without taking advantage of how intent, and in particular the user’s desired control input, depends on the state of the prosthetic. At each time step k , the cursor will move in the wrong direction, away from θ , with probability

$$\gamma = (1 - \alpha)\beta + \alpha(1 - \beta).$$

So the expected distance moved toward the target after k steps is at least

$$k(1(1 - \gamma) - 1\gamma) = k(1 - 2\gamma),$$

with a slightly larger value being obtained as p_k nears 1 or n . So we have

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

for all $N < n - 1$. Note that although this bound captures the trend, we could obtain a slightly tighter one, as well as a bound for $N \geq n - 1$. So, the cost of this policy again decreases *linearly* with N .

Just as before, the fact that intent depends on the state of the prosthetic can be used to improve performance. In particular, 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

$$P(\theta|y_0, \dots, y_k) = \eta \cdot \begin{cases} (1 - \gamma) \cdot P(\theta|y_0, \dots, y_{k-1}) & \text{for all } \theta \in (p_{k+1}, \dots, n) \\ \gamma \cdot P(\theta|y_0, \dots, y_{k-1}) & \text{for all } \theta \in (1, \dots, p_k) \end{cases}$$

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 distribution over θ , so u_k satisfies both

$$\sum_{i=1}^{p_k + u_k - 1} p(\theta = i|y_0, \dots, y_k) \leq \frac{1}{2}$$

and

$$\sum_{i=p_k + u_k}^n p(\theta = i|y_0, \dots, y_k) > \frac{1}{2}.$$

We know that this policy has cost bounded by

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

over a finite horizon N , for some constant $c > 0$ [5], [7], [20]. 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 desired inputs x_0, \dots, x_{k-1} and feedback signals y_1, \dots, y_{k-1} . This interpretation is fundamental to developing good control schemes [8], [20].

V. RESULTS

A. Experimental setup

In the previous section we presented two control policies, which we called “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 two healthy volunteers, 1 male and 1 female, aged 20, with normal vision.

TABLE I
ERROR RATES

	Subject 104	Subject 105
User error (α)	6%	4%
Device error (β)	5%	5%

TABLE II
AVERAGE PERFORMANCE

	Subject 104		Subject 105	
	bisection	fixed offset	bisection	fixed offset
chars. / minute	9.9	8.57	10.1	8.45
bits / minute	47.5	40.7	48.2	40.2

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 [1]. 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.

2) *Parameterization:* To implement the probabilistic bisection control policy, we need to know the values of α and β . These values measure the likelihood of mistakes made by the user and the neural sensor, respectively. They were computed in advance from experimental data.

To find α , subjects used the left and right arrow keys on a standard keyboard as input, rather than the surface-EMG sensor. They entered 200 characters from the short story “The Door in the Wall” by H.G. Wells. If the subject provided no input for 1 second, the letter highlighted at that time was accepted. Bernoulli noise with parameter $p = 0.85$ was added to the subject’s input in order to account for potential effects of noise on user accuracy. For this first trial, we set $\alpha = 0.00$ and $\beta = 0.15$. The estimate for α was then taken to be the error of the subject’s actual input, neglecting the additional noise, by comparing the cursor position to the position of the target letter at each time step.

To find β , subjects used the surface-EMG sensor. First, each subject trained this sensor to recognize subvocal expression of the two words “left” and “right.” Then, they were presented with a series of 50 prompts, each asking the subject, with equal probability, to subvocalize one of these two words. The resulting error rate was used for β .

Table I summarizes the results of these experiments.

3) *Online tests:* Each subject was asked to write three phrases: “HEAD HURTS”, “WHAT CHANCE”, and “MINE MOVED”. 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 1.5 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.




Fig. 3. Number of characters selected per minute for each phrase using bisection algorithm, computed as an average over all of the runs for a particular phrase. (Subject 104 is in blue and 105 in red)




Fig. 4. Number of characters selected per minute for each phrase using simple fixed offset input, computed as an average over all of the runs for a particular phrase. (Subject 104 is in blue and 105 in red)

B. Data and analysis

We measured subject performance as the number of characters selected per minute. Figure 4 shows a running average over all trials with a 50-letter horizon. Average characters per minute across all phrases and the corresponding average effective bits per minute are summarized in Table II.

Despite widely varying subject characteristics, including user and device error rate, clear trends emerge regarding the performance of our two policies. Figure 4 shows that probabilistic bisection improved performance for both subjects, allowing significantly higher average throughput than fixed offset for subjects 104 and 105.

VI. 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.

VII. ACKNOWLEDGMENTS

The authors thank Ed Maclin and Michael Callahan for their helpful discussions and Katie Snell, David Rockwood, and Daniel Chang for their assistance with the experiments.

REFERENCES

- [1] The Audeo. Ambient Corp., <http://www.theaudio.com/tech.html>.
- [2] D. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific, third edition, 2005.
- [3] D. R. Brown, III, R. Ludwig, A. Pelteku, G. Bogdanov, and K. Keenaghan. A novel non-acoustic voiced speech sensor. *Measurement Science and Technology*, 15:1291–1302, 2004.
- [4] B. Bryan, J. Schneider, R. C. Nichol, C. J. Miller, C. R. Genovese, and L. Wasserman. Active learning for identifying function threshold boundaries. In *NIPS*, 2005.
- [5] M. Burnashev and K. Zigangirov. An interval estimation problem for controlled observations. *Problems in Information Transmission*, 10:223–231, 1974. Translated from *Problemy Peredachi Informatsii*, 10(3):51–61, July–September, 1974.
- [6] R. Castro, R. Willett, and R. Nowak. Faster rates in regression via active learning. In *NIPS*, 2005.
- [7] R. M. Castro and R. D. Nowak. Active learning and sampling. In *Foundations and Applications of Sensor Management*. Springer-Verlag, 2007.
- [8] N. Elia. When Bode meets Shannon: control-oriented feedback communication schemes. *IEEE Trans. Autom. Control*, 49(9):1477–1488, Sept. 2004.
- [9] G. E. Fabiani, D. J. McFarland, J. R. Wolpaw, and G. Pfurtscheller. Conversion of eeg activity into cursor movement by a brain-computer interface (bci). *IEEE Trans. Neural Syst. Rehabil. Eng.*, 12(3):331–338, 2004.
- [10] L. Farwell and E. Donchin. Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalogr. Clin. Neurophys.*, 70:510–523, 1988.
- [11] E. Felton, N. L. Lewis, S. A. Wills, R. G. Radwin, and J. C. Williams. Neural signal based control of the dasher writing system. In *IEEE EMBS Conf. Neural Engr.*, 2007.
- [12] 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*, 442:164–171, July 2006.
- [13] M. Horstein. Sequential transmission using noiseless feedback. *IEEE Trans. Inf. Theory*, 9:136–143, July 1963.
- [14] C. C. Jorgensen and K. Binsted. Web browser control using EMG based sub vocal speech recognition. In *Int. Conf. System Sciences*, 2005.
- [15] C. C. Jorgensen, D. D. Lee, and S. Agabon. Sub auditory speech recognition based on emg signals. In *Int. Joint Conf. on Neural Networks*, pages 3128–3133, 2003.
- [16] L. Maier-Hein. *Speech Recognition Using Surface Electromyography*. PhD thesis, Universität Karlsruhe, Karlsruhe, Deutschland, July 2005.
- [17] N. C. Martins and M. A. Dahleh. Feedback control in the presence of noisy channels: “Bode-like” fundamental limitations of performance. *IEEE Trans. Autom. Control*, May 2008. To Appear.
- [18] C. Omar, M. Johnson, T. Bretl, and T. Coleman. Querying the user properly for high-performance brain-machine interfaces: recursive estimation, control, and feedback information-theoretic perspectives. In *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Las Vegas, NV, 2008.
- [19] H. Serby, E. Yom-Tov, and G. F. Inbar. An improved P300-based brain-computer interface. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 13(1):89–98, Mar. 2005.
- [20] O. Shayevitz and M. Feder. Communication with feedback via posterior matching. In *IEEE Int. Symp. Info. Th.*, Nice, France, June 2007.
- [21] A. Singh, R. Nowak, and P. Ramanathan. Active learning for adaptive mobile sensing networks. In *Int. Conf. Information Processing in Sensor Networks*, pages 60–68, Apr. 2006.
- [22] D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz. Direct cortical control of 3d neuroprosthetic devices. *Science*, 296:1829–1832, 2002.
- [23] M. Thulasidas, C. Guan, and J. Wu. Robust classification of EEG signal for brain-computer interface. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 14(1):24–29, Mar. 2006.
- [24] T. M. Vaughan, D. J. McFarland, G. Schalk, W. A. Sarnacki, D. J. Krusienski, E. W. Sellers, and J. R. Wolpaw. The wadsworth BCI research and development program: At home with BCI. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 14(2):229–233, 2006.
- [25] J. Wessberg, C. Stambaugh, J. Kralik, P. Beck, M. Laubach, 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.
- [26] S. A. Wills and D. J. C. Mackay. DASHER—an efficient writing system for brain-computer interfaces. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 14(2):244–246, June 2006.
- [27] J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan. Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6):767–791, 2002.