35th Annual International Conference of the IEEE EMBS
Osaka, Japan, 3 - 7 July, 2013

Online Adaptive Decoding of Intended Movements with a Hybrid Kinetic and Kinematic Brain Machine Interface

Aaron J. Suminski, *Member, IEEE*, Andrew H. Fagg, *Member*, *IEEE*, Francis R. Willett, Matthew Bodenhamer, and Nicholas G. Hatsopoulos

Abstract—**Traditional brain machine interfaces for control of a prosthesis have typically focused on the kinematics of movement, rather than the dynamics. BMI decoders that extract the forces and/or torques to be applied by a prosthesis have the potential for giving the patient a much richer level of control across different dynamic scenarios or even scenarios in which the dynamics of the limb/environment are changing. However, it is a challenge to train a decoder that is able to capture this richness given the small amount of calibration data that is usually feasible to collect *a priori*. In this work, we propose that kinetic decoders should be continuously calibrated based on how they are used by the subject. Both intended hand position and joint torques are decoded simultaneously as a monkey performs a random target pursuit task. The deviation between intended and actual hand position is used as an estimate of error in the recently decoded joint torques. In turn, these errors are used to drive a gradient descent algorithm for improving the torque decoder parameters. We show that this approach is able to quickly restore the functionality of a torque decoder following substantial corruption with Gaussian noise.**

I. Introduction

Many brain machine interfaces (BMIs) are used to control physical systems like a prosthesis. In these BMIs it is traditional to first decode either an 'intended' position or velocity signal and then submit this intended state to a control law (e.g., a PD controller) that attempts to match this state through the appropriate selection of joint torques [1-2]. This approach suffers from three limitations. First, it entirely reactive and thus its response is delayed: an error must exist between the desired and actual states in order for the controller to respond. Second, it does not take into account the dynamics of interaction between the arm and the environment, like those that occur when picking up a heavy object. A BMI should enable the subject to anticipate such forces proactively by providing control channels to influence these forces directly. Finally, the decoders used to predict the 'intended' movements are trained in a single context making them incapable of adapting to novel environments without subjecting the users to a recalibration session.

One approach to mitigate these concerns is a hybrid control scheme that pairs a feedforward controller capable of tracking the desired trajectory without delay with a feedback controller that corrects for any inaccuracies in the realized movement trajectory resulting from a mismatch between the feedforward model and the true dynamics. We have recently demonstrated viability of the hybrid approach for using BMIs to control robot arms [3]. Our hybrid neural decoder estimates "intended" arm positions/velocities and joint torques from recent MI activity. The decoded torque is interpreted as a feed-forward control signal, while the decoded position is interpreted as a reference signal for a proportional-derivative (PD) feedback controller.

Here, we present an online-adaptive, hybrid neural decoder that may be capable of improving task performance when placed in novel dynamic scenarios. We extend our previous work by adapting the torque decoder as it is being used to command an arm in a reaching task. Inspired by the Feedback Error Learning (FEL) model of Kawato and Gomi [4], we interpret position errors (differences between actual and decoded position) as 1) an error signal on which the PD controller will operate and 2) an estimate of recent errors made by the torque decoder. We use this latter error signal to drive a quasi-supervised learning process that alters the parameters of the torque decoder. We show that this approach is able to restore the functionality of a torque decoder after it has been corrupted with Gaussian noise.

II. Methods

A. Behavioral Task

One adult male rhesus macaque was trained to sit in a primate chair and to control a cursor in a two–dimensional workspace using the KINARM, a two-link robotic exoskeleton (BKIN Technologies, Kingston, ON). The animal's arm was abducted 90 degrees and rested in the exoskeleton such that all movements were made within the horizontal plane. Visual feedback of the monkey's movements was available via a cursor projected onto a horizontal projection screen which blocked direct vision of the arm. The position of the cursor was controlled by one of two sources: either the position of the monkey's hand or the endpoint of a simulated robot that was controlled by a BMI.

The animal performed a random target pursuit (RTP) task

Manuscript received February 4, 2013. This work was supported by NIH NINDS R01 N545853-01 to NGH and AHF.

A. J. Suminski and N. G. Hatsopoulos are with the Department of Organismal Biology and Anatomy at the University of Chicago, Chicago, IL 60637 USA (asuminski@uchicago.edu and nicho@uchicago.edu).

F. R. Willett is with the Department of Biomedical Engineering at the Case Western Reserve University, Cleveland, OH 44106 USA (frw9@case.edu).

A. H. Fagg and M. Bodenhamer are with the School of Computer Science at the University of Oklahoma, Norman, OK 73019 USA (fagg@cs.ou.edu and mbodenhamer@cs.ou.edu).

978-1-4577-0216-7/13/$26.00 ©2013 IEEE

requiring repetitive movements of the visual cursor (6 mm diameter) to square targets (2.25cm^2). Targets appeared at a random location within the workspace (14 by 13.2 cm), and each time the cursor acquired one, a new target appeared in a random location. In order to complete a successful trial and receive a juice reward the monkey was required to sequentially acquire between two and seven targets. A trial was aborted if any movement took longer than 5s.

B. *Hybrid BMI*

Here, we utilized a BMI that decoded both kinematic and kinetic signals. The specific details of this implementation have been described previously [3]. In brief, the Hybrid BMI (Fig. 1) converts neural activity into elbow and shoulder joint torques in order to drive the movement of a two-link virtual arm that simulates the dynamics of the monkey's arm and exoskeleton. The hand position of the simulated arm (X_C) is updated based on the current joint angles, velocities, and accelerations of the virtual arm and the joint torques generated by the BMI. The torques, τ, are a weighted combination of the torques generated by the individual components of the BMI:

$$\tau = k_t * \tau_t + \tau_{pd}\,, \tag{1}$$

where τ_t is the prediction of the monkey's intended shoulder and elbow joint torques made by the torque decoder and τ_{pd} is the torque generated by the position-derivative controller, and k_t is the gain coefficient to control the mixture of the component torques.

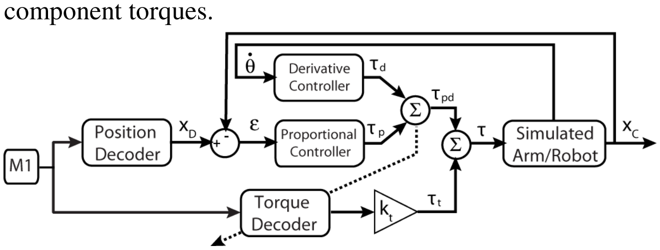

Figure 1: The Hybrid BMI uses neural activity recorded from the primary motor cortex to drive the simulated arm in Cartesian space which the monkey then uses to hit targets. X_D and X_C are two-element column vectors containing the X and Y components of the endpoint of the simulated arm, ε is a two-element column vector containing the X and Y components of an error signal, τ_t, τ_{pd} and τ are two-element column vectors containing joint torque terms, and k_t, is a scalar gain term. The torque output of the PD controller, τ_{pd}, is used as an error signal to modify the torque decoder (dashed arrow).

The proportional (position) and derivative (velocity) controllers function together to move the simulated arm towards the position decoder's prediction of the monkey's intended hand position, X_D. The position controller uses the error signal, $\varepsilon = X_D - X_C$, to generate τ_p. Similarly, the velocity controller generates joint torques proportional to the current angular velocities of the simulated arm. The individual gain parameters of the PD controller were tuned in simulation prior to the experiments.

The position and torque decoders, implemented as Wiener Filters, predict the monkey's intended hand position, X_D, and the monkey's intended joint torque, τ_t. In our approach, an estimate of "intended" hand position/joint torque is reconstructed from a linear combination of binned spike counts from the available neurons. We employ a history of B = 20 bins of Δt = 50ms each for every neuron, giving the Wiener Filters access to a total of one second of neural spiking history. Specifically, signal k (X or Y hand position, or elbow or shoulder torque) at discrete time bin t, is reconstructed as follows:

$$S_k(t) = W_k\, F(t)\,, \tag{2}$$

where W_k is the set of parameters for filter k, and $F(t)$ is a column vector of size BxC that contains the spike counts for C neurons and B time bins (covering time [t-(B-1)Δt,…,t]). As in our previous work, the coefficients for the decoder are solved for analytically using ridge regression that trades prediction accuracy on the training set for a smoother prediction surface [3].

C. *Feedback Error Learning*

In hybrid control, any failing of the feedforward controller, the torque decoder in our implementation, to estimate the correct feedforward torques manifests itself as a position (and velocity) error that accumulates slowly with time. Hence, a positional error that is observed at time t, can be due to a feedforward torque error that was made prior to t. The Feedback Error Learning assumption addresses this credit assignment problem by assigning some responsibility for positional/velocity errors at time t over the recent window of torque estimates. Specifically, we can describe a squared torque error for a single joint and for time t as:

$$E_k(t) = \frac{1}{2}\sum_{n=0}^{T-1} \rho(T-1-n)[W_k\, F(t-n\Delta t) - \tau_k(t-n\Delta t)]^2\,, \tag{3}$$

where $k \in \{shoulder, elbow\}$, T is the number of bins that are considered as affecting position at time t, $\rho(n)$ is a weighting function defined over the bins, in which $\sum_{n=0}^{T-1}\rho(n) = 1$, $\tau_k(t)$ is the correct (but unknown) torque at t, and $W_k\, F(t) - \tau_k(t)$ is the error between decoded and this correct torque. Given this definition of cost, one can formulate the search for an appropriate set of decoder parameters as an incremental gradient descent problem, specifically:

$$\frac{\partial E_k(t)}{\partial W_k} = \sum_{n=0}^{T-1} \rho(T-1-n)[W_k\, F(t-n\Delta t) - \tau_k(t-n\Delta t)]F(t-n\Delta t). \tag{4}$$

By the FEL assumption:

$$W_k\, F(t-n\Delta t) - \tau_k(t-n\Delta t) \approx \tau_{pd,k}(t)\,. \tag{5}$$

As a result, the update to the estimate of filter parameters based on a single observation is:

$$W_k \leftarrow W_k + \alpha\, \tau_{pd,k}(t)\sum_{n=0}^{T} \rho(T-1-n)\, F(t-n\Delta t)\,, \tag{6}$$

where α is a learning rate parameter.

The weighting function $\rho(n) \propto \exp(-\beta n)$ grants most of the responsibility for the current positional error to the most recently decoded torque, and exponentially less responsibility to decoded torques made in the past. This formulation also has the advantage that the above sum can be computed incrementally.

D. Experimental Procedure

Prior to the experiments, the monkey performed a visual observation task where he held his arm still while observing movements of the cursor hitting targets. These movements were recorded previously while the monkey performed the RTP task with his own arm. During the observation period, we recorded four minutes and twenty seconds (5200, 50ms bins) of spiking activity and movement data to train the decoders. The torques used to train the decoder were estimated from the observed arm trajectories and an inverse dynamics model that included both the KINARM and the monkey arm. During the experiments, the monkey used a BMI to move the cursor (via the simulated arm) in the same task based on the activity of an ensemble of recorded motor cortical neurons.

We performed a single experiment (a total of 3 sessions) to investigate the ability of the Adaptive Hybrid BMI to improve BMI performance after an unexpected external perturbation to the decoder. The experiment consisted of two conditions arranged in an ABA paradigm. In the first condition the monkey moved the cursor using an unperturbed version of the Hybrid BMI in order to establish a performance baseline. Next, we corrupted the torque decoder by adding zero-mean, unit variance, noise drawn from a Gaussian distribution to its coefficients. During the second condition, we enabled the adaptation algorithm and evaluated its ability restore the functionality lost after the decoder corruption. Following an extended learning period, the adaptation algorithm was disabled and the original torque decoder was restored.

E. Electrophysiology

The monkey was chronically implanted with a 100-electrode microelectrode array (Blackrock Microsystems, Inc., Salt Lake City, UT) in MI contralateral to the arm used for the task. The electrodes on the array were 1.0 mm in length and were coated with iridium oxide. The neural activity used to train the decoders and operate the BMI was comprised of single and multiunit spiking events. Only waveforms that crossed a user defined threshold and sorted online using a time-amplitude technique were used for real-time decoding. All of the surgical and behavioral procedures were approved by the University of Chicago Institutional Animal Care and Use Committee and conform to the principles outlined in the *Guide for the Care and Use of Laboratory Animals*.

F. Kinematic Analyses

We used two kinematic measures to evaluate the effect of the decoder adaptation algorithm on BMI performance. The normalized time-to-target metric is defined as the time difference between consecutive target hits divided by the Euclidean distance between the targets. The normalized path length metric is defined as the length of the path the cursor traverses between consecutive targets divided by the Euclidean distance between those targets.

III. Results

We sought to demonstrate that our Adaptive Hybrid BMI was capable of restoring BMI performance following an unexpected external perturbation to the coefficients of the torque decoder. The monkey was readily able to acquire targets using the Hybrid BMI prior to corruption of the torque decoders (Fig. 2, negative time points). Corruption of the torque decoder parameters (Fig. 2, vertical dashed line at t =0) resulted in functional paralysis of the BMI as the monkey was unable to modulate the cursor position and acquire targets. Upon enabling of the adaptive algorithm (Fig. 2, 2^{nd} dashed vertical line), normal cursor movement was restored within 10 seconds. The monkey was able to acquire targets at a similar rate to the pre corruption period, approximately 20 seconds after the adaptive algorithm was enabled.

We were interested in characterizing how the adaptive BMI was able to compensate for a perturbation that eliminated the monkey's ability to effectively control the BMI cursor. Fig. 3 shows time-resolved and average performance data for a representative experiment. During the pre-learning period (Fig. 3, blue), BMI performance reached a steady state that was disrupted by the decoder corruption (Fig. 3, Learning On). Early in the learning period (Fig. 3, gray) performance degraded and then gradually improved over the course of many trials until adaptation was turned off (Fig. 3, Learning Off). The time for the cursor to reach target and path length were significantly reduced during the learning period (two sample t-test, $p < 0.01$ and $p < 0.05$, respectively). There was no difference between the performance in the late learning phase and the pre/post learning phases.

IV. Discussion

Brain-machine interfaces that provide a patient with direct access to the kinetics of prosthetic limb control could increase the range of tasks that are feasible to learn and perform. This increase in the richness of control also comes with an increase in the complexity of the decoder above that of kinematic decoders. The Feedback Error Learning approach to updating the torque decoder on-the-fly enables the decoder to integrate experience from the dynamic scenarios that are actually encountered by the patient. Hence, all dynamic scenarios do not have to be anticipated *a priori* during a pre-use calibration phase.

Our results demonstrate that the FEL approach allowed the monkey to regain control of the BMI after a catastrophic failure due to external perturbation of the torque decoders.

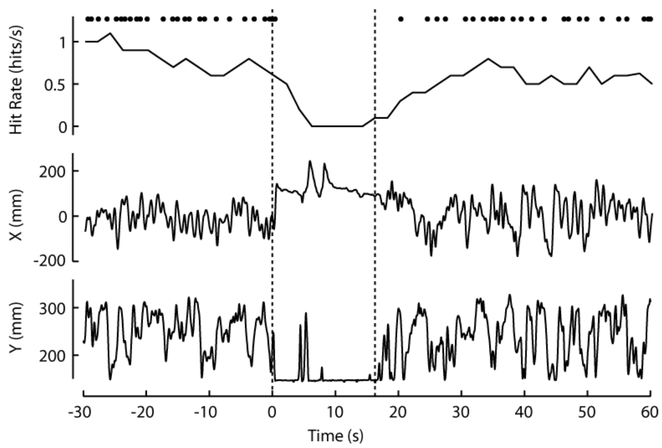

Figure 2: Hit rate (top), X position (middle) and Y position (bottom) generated by the online, adaptive hybrid BMI prior to and after corruption of the torque decoder. Corruption of the torque decoder (vertical dashed line, t = 0) rendered that monkey unable to use the BMI. Performance was rapidly restored when the adaptive algorithm was enabled (vertical dashed line, t ~ 17s). Similar trends were observed in each of the 3 experiments.

This is especially true in first 30 seconds after the adaptation algorithm was enabled. What is less clear is the contribution of the adaptive decoder to the performance gains observed over longer time courses like that shown in Fig. 3. Consistent with previous reports [5], we expect that both the adaptive algorithm and the monkey would contribute to performance improvement on longer timescales. Future work will attempt to tease apart the individual contributions of the monkey and the algorithm in order to optimize BMI performance.

Online adaptive decoders have previously been proposed [5-6]. Here, kinematic decoders are updated on-the-fly based on differences between decoded velocity and an expectation of what that velocity should be given knowledge of the goal. However, as we move toward *in situ* adaptive decoders, the intended goal of movement may not be know explicitly. Instead, our learning algorithms must infer this intent in other ways. The FEL approach does just this. Errors in the torque decoder are estimated, in part, from the decoded position. One limitation, however, is that we assume that the position decoder itself is of high quality and does not require continued adaptation. This is one subject of future work.

The decoder corruption paradigm that we have used in this work is a stand-in for changes that might occur over a time-course of days or weeks in the behavior of individual neurons or a drift in the set of neurons that might be available to the decoder. Our experimental results suggest that the FEL approach could aid a dynamic decoder to compensate for these more realistic changes. Furthermore, we have shown in simulation that FEL can compensate for changes in the dynamics of arm behavior, for example, through the introduction of a curl field [7]. In future experiments, we are planning a range of dynamic scenarios, including those involving the grasping and placement of heavy objects.

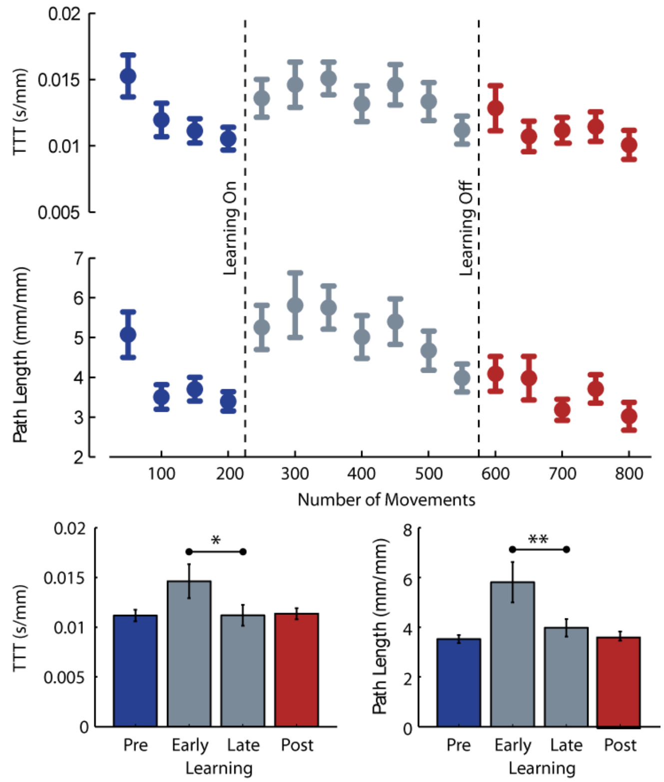

Figure 3: (Left) Time resolved BMI performance prior to decoder adaptation (Pre, blue) during decoder adaptation (Learning, gray) and following decoder adaptation (Post, red). Dashed vertical lines indicate the start and end of decoder adaptation (Learning On and Learning Off, respectively). (Right) Average BMI performance prior to decoder adaptation, during early and late learning and following adaptation. Error bars denote ± 1SE about the mean performance. * and ** indicated significant differences at the alpha = 0.1 and 0.05 levels, respectively.

REFERENCES

[1] J. M. Carmena, M. A. Lebedev, R. E. Christ, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez and M. A. L. Nicholelis, "Learning to control a brain-machine interface for reaching and grasping by primates," *PLOS Biol*, vol. 1, pp. e42, 2003

[2] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford and A. B. Schwartz, "Cortical control of a prosthetic arm for self-feeding," *Nature*, vol. 453, pp. 1098-101, 2008

[3] A. J. Suminski, F. R. Willett, A. H. Fagg, M. Bodenhamer and N. G. Hatsopoulos, "Continuous decoding of intended movements with a hybrid kinetic and kinematic brain machine interface," *Conf Proc IEEE Eng Med Biol Soc. 2011*, pp. 5802-06, 2011.

[4] M. Kawato and H. Gomi, "A computational Model of 4 regions of the cerebellum base on feedback-error learning," *Biol Cyvern*, vol. 68, pp. 95-103, 1992.

[5] A. L. Orsborn, S. Dangi, H. G. Moorman and J. M. Carmena, " Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions," *IEEE Trans Neural Syst Rehabil Eng*, vol. 20, pp. 468–77, 2012.

[6] V. Gilja, P. Nuyujukian, C. A. Chestek, J. P. Cunningham, B. M. Yu, J. M. Fan, M. M. Chruchland, M. T. Kaufman, J. C. Rao, S. I. Ryu and K. V. Shenoy, "A high-performance neural prosthesis enabled by control algorithm design," *Nat Neurosci*, vol. 15, pp. 1752-7, 2012..

[7] M. Bodenhamer, F. R. Willett, A. J. Suminski, N. G. Hatsopoulos and A. H. Fagg, "A feedback error learning approach to online-adaptive decoding for dynamic prosthetic arm dontrol," *Cosyne Abstracts 2012*, Salt Lake City, USA, 2012