7th Annual International IEEE EMBS Conference on Neural Engineering
Montpellier, France, 22 - 24 April, 2015

System Identification of Brain-Machine Interface Control Using a Cursor Jump Perturbation

Sergey D. Stavisky*, *EMBS Student Member*, Jonathan C. Kao*, *EMBS Student Member*,
Jordan M. Sorokin, Stephen I. Ryu, *IEEE Member*,
and Krishna V. Shenoy, *IEEE Senior Member*

***Abstract*— Inspired by control theoretic approaches to studying motor control, we experimentally measured how a brain-machine interface (BMI) user responds to an unexpected perturbation. We randomly applied a step cursor position offset while a monkey controlled a BMI cursor using decoded motor cortical spiking activity. The subject was able to rapidly correct for these perturbations and (re)acquire the target regardless of when in the trial this cursor jump occurred. We observed a corrective neural response in motor cortex starting 115 ms after the cursor jump. At no time did the neural response to detecting this externally-induced error manifest itself (through the decoder) as a deleterious velocity change pushing the cursor away from the target. These results show that a user of a high-performance BMI can make rapid, accurate corrections to errors and that, insofar as the neural computations needed to counteract the error may involve motor cortex, these computations do not appear to interfere with BMI cursor control.**

I. Introduction

Brain-machine interfaces (BMIs) decode neural activity to allow a user to directly control an effector such as a computer cursor [1], [2] or robotic arm [3], [4]. Inspired by these studies in non-human primates, intracortical BMIs are now being tested in pilot clinical trials to restore communication [5], [6] and movement [7], [8] to people with paralysis. Further progress in BMI capabilities will require greater scientific understanding of how the brain controls a BMI and incorporates it into its motor schema. We pursued this question using an approach that has proven highly fruitful in basic motor neuroscience: perturb the motor system and observe how it responds to, and corrects for, errors [9]–[11]. We adapted the systems identification technique of applying an unexpected step perturbation to the controlled plant (e.g., a torque applied to the arm during reaching [12]) to the context of BMI control by introducing a sudden BMI cursor position offset ("cursor jump"), while a monkey was trying to bring the cursor to a target. Although several groups have studied how the sensorimotor system responds to a continual perturbation consisting of changing the neural-to-kinematic mapping of the BMI decoder [13]–[16], to our knowledge this is the first time that the response to a step perturbation has been studied.

We focused this initial investigation on characterizing how well a BMI user is able to detect and subsequently correct for the error introduced by the perturbation. This question will become increasingly important as BMI users perform more complex tasks in which unexpected errors can and will occur – for example, if a BMI-driven prosthetic arm is bumped as it reaches towards a cup. We were specifically interested in two questions. First, we measured how quickly a BMI user relying on visual feedback of the cursor position begins to respond to a perturbation. This control loop latency sets a fundamental limit to how controllable the BMI will be in the face of changing task goals. Second, we wanted to know whether the neural correlates of detecting the cursor jump and initiating its correction could themselves interfere with ongoing BMI control and exacerbate the error. This is a germane concern for the following reason: although error detection/correction introduces a strong change in motor cortical neural activity [17]–[19], there are many possible mechanisms (e.g., [20]) in place to prevent this activity from driving unwanted activations of the muscles. The brain's response to having an arm bumped does not cause the arm to momentarily twitch uncontrollably. During BMI use, however, the same neural population involved in correcting an error is also directly connected to the BMI output; this raises the worrisome possibility that the neural computation of a corrective response may itself introduce additional error.

II. Methods

A. Behavioral Tasks

All procedures and experiments were approved by the Stanford University Institutional Animal Care and Use Committee. A rhesus macaque (monkey J) was trained to sit head-fixed in a primate chair and perform 2D target acquisition tasks by controlling an on-screen cursor with either his hand position or via a BMI in order to obtain a liquid reward. The task was displayed with a latency of 7 ± 4 ms as in [2]. We used a "hand controlled training, hand free BMI control" paradigm [21] in which the monkey was free to move his arm during BMI control. At the start of each experiment the monkey first performed a center-out-and-back Radial 8 Target task which provided training data for a

S.D.S and J.C.K contributed equally to this work. Research supported by National Science Foundation Graduate Research Fellowship (S.D.S., J.C.K.), NSF IGERT (S.D.S), NIH Pioneer Award 8DP1HD075623, NIH T-RO1 award NS076460, and DARPA REPAIR award N66001-10-C-2010 (K.V.S.).

S. D. Stavisky (e-mail: sergey.stavisky@stanford.edu) and J. M. Sorokin (jorsor@stanford.edu) are with the Neurosciences Graduate Program, Stanford University, Stanford, CA 94305 USA.

J. C. Kao is with the Department of Electrical Engineering at Stanford University.

K. V. Shenoy (shenoy@stanford.edu) is with the Departments of Electrical Engineering, Bioengineering, and Neurobiology; the Neurosciences Graduate Program and the Bio-X Program; and the Stanford Neurosciences Institute, Stanford University.

S. I. Ryu (seoulman@stanford.edu) is with the Department of Neurosurgery at Palo Alto Medical Foundation, Palo Alto, CA, USA, and the Department of Electrical Engineering, Stanford University.

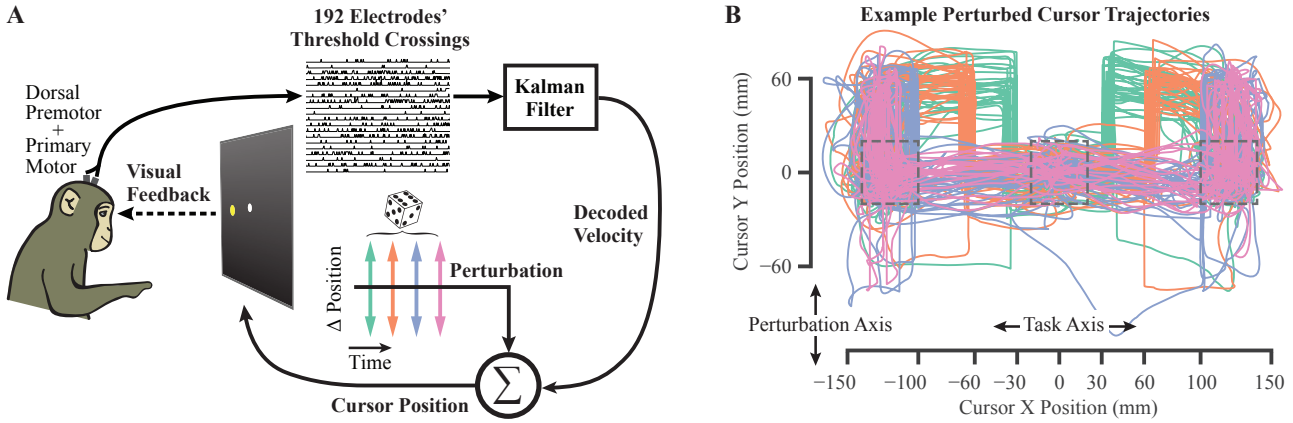

Fig. 1. (A) Overview of the BMI cursor jump experiment. A monkey controls a cursor using decoded neural activity from two 96-electrode arrays in motor cortex. He performs a two-target center-out-and-back task, and on a random 25% of outward trials the cursor is "jumped" 60 mm perpendicular to the task axis. The jump occurs after the cursor has travelled 30 mm (teal), 60 mm (orange), or 100 mm (blue) along the task axis, or during the target hold time (pink). (B) Cursor trajectories during perturbed trials from one example dataset. In this experimental session, the task axis was the horizontal dimension, and 90% (10%) of jumps were up (down). All trials originated from the center target and proceeded to the left or right targets. Target boundary boxes are shown as dashed grey squares. Dataset J-2013-10-31.

decoder. After holding the center target, one of eight peripheral targets 12 cm from the center was displayed. To acquire it, the monkey had to hold the cursor inside a 4 x 4 cm square acceptance region for a contiguous 500 ms.

The main BMI task analyzed in this study was a two-target variant of the same cursor task. The two possible targets were arranged either along the horizontal axis or vertical axis (the "task axis") and the sequence of targets was random. The monkey had 3 seconds to reach the target or else he failed the trial. The key experimental manipulation was that on a random 25% of trials towards the peripheral targets, a perturbation was applied (Fig. 1). This perturbation consisted of offsetting ("jumping") the cursor position by 60 mm perpendicular to the task axis. During perturbation trials, the jump occurred after one of four randomly selected and equally likely conditions were met. Three of the criteria were spatial: the jump happened after the cursor travelled either 30 mm, 60 mm, or 100 mm towards the target along the task axis. The fourth criterion was that the jump happened during the target hold period. Only one perturbation could happen per trial, and these perturbation conditions were interleaved throughout each session. These four criteria comprise the four different cursor jump conditions examined in this study. For technical reasons, the perturbation was actually applied on the next decode update time step following the criterion being met (i.e., within 25 ms). The time of the cursor jump command was recorded with 1 ms resolution, after which it appeared at the next monitor update.

Fourteen datasets were collected and analyzed for this study. Six of the datasets did not have the jump during hold condition; instead, they included additional conditions (not analyzed in the present study) in which the cursor jump magnitude was either 30 or 90 mm. For 9 (5) datasets, the targets were arranged along the horizontal (vertical) axis, with cursor jumps either up or down (left or right). The distribution of jump directions was 90/10% in four datasets and 50/50% in the remaining ten datasets. The hold period jump happened 125 (275) ms after the cursor entered the target in three (five) datasets; these conditions' data were similar and thus were combined into a single "jump during hold" condition.

B. Neural Recording and Decoding

Monkey J was implanted with two 96-electrode arrays (1 mm electrodes spaced 400 μm apart, Blackrock Microsystems) using standard neurosurgical techniques [22] 50-63 months prior to these experiments. One array was implanted into primary motor cortex (M1) and the other into dorsal premotor cortex (PMd) contralateral to the reaching arm. Voltage signals from each of the 192 electrodes were band-pass filtered from 250 to 7500 Hz and then processed to obtain multiunit "threshold crossings" spikes. A spike was detected whenever the voltage crossed below a threshold set at the beginning of each day to be -4.5 × rms voltage [23]. We did not spike sort to assign spikes to individual putative neurons and instead grouped together threshold crossings from a given electrode "channel"; the population activity will therefore include both single- and multiunit activity. Both arrays recorded good signals, with spikes detected on almost every electrode.

BMI control of the cursor's velocity was enabled using the Recalibrated Feedback Intention Trained Kalman Filter (ReFIT-KF) decoder and training algorithm described extensively in [2]. Briefly, at the start of each experiment we trained a position + velocity Kalman filter from a training dataset of ~500 arm-controlled Radial 8 Task trials. This first-pass decoder was used to complete an additional ~500 trials of this task under BMI control. This closed-loop data was then used to fit the final ReFIT-KF velocity decoder after several modifications of the kinematics training data to improve the estimation of the animal's true intent. Since we experimentally manipulated cursor position via the cursor jump, we did not apply the "position subtraction" ReFIT-KF operation used in [2] to avoid having the cursor jump explicitly affect decoded velocity. It is worth emphasizing that during closed-loop use this decoder rapidly converges [24] to a steady-state linear decoder of the form

$$\mathbf{x}(t) = \mathbf{M}_1\mathbf{x}(t\text{-}1) + \mathbf{M}_2\mathbf{y}(t) \qquad (1)$$

where $\mathbf{x}(t)$ is a 5 x 1 state vector of x- and y- position and velocity, and a 1 bias term. $\mathbf{y}(t)$ is a 192 x 1 vector of each channel's spike counts in the current non-overlapping 25 ms decoder time step. $\mathbf{M}_1$ is a state dynamics matrix that smooths velocity over time and lawfully integrates velocity to change position. The observation matrix $\mathbf{M}_2$ applies no neural contribution to position but has two rows of weights that linearly map each channel's firing rate to an x- and y-velocity component, after subtracting a baseline firing rate. The net effect of observed neural activity at time step t onto velocity is the summation of these velocity components.

C. Analysis of Reponses to Cursor Jumps

The monkey almost always successfully acquired targets, despite cursor jumps, except when he ceased to engage in the task. We thus restricted our analysis to successful trials. To exclude trials where the monkey briefly ceased to engage in the task, we removed outlier trials with times to target more than three standard deviations longer the mean. Combining across the five datasets, we analyzed 16,005 unperturbed trials, 1,041 jump at 30 mm trials, 977 jump at 60 mm trials, 1,007 jump at 100 mm trials, and 582 jump during hold trials.

We quantified the effect of the perturbations on cursor control performance using two metrics. **Time to target** measured how long it took to acquire the target, not including the 500 ms required final hold time, but including incomplete earlier holds. For the jump during hold condition, we also did not include the initial hold time preceding the cursor jump, since this time was "stolen" in the sense that there was no way for the animal to minimize this penalty. **Distance to target** measured the Euclidean distance between the cursor position and the center of the target as a function of time relative to the cursor jump.

We measured how changes in neural activity evolved after a cursor jump perturbation in two ways. To see how activity evolved in the "decoder-potent" subspace, we projected the vector $\mathbf{y}(t)$, consisting of each channel's spike counts in a 25 ms bin ending at time t, through the decoder observation matrix $\mathbf{M}_2$ (1). This allowed us to measure the "neural push" [25] (i.e., how neural activity would influence decoded x- and y- velocity at each millisecond). Note that this is intended as an offline analysis of how neural activity evolves in the decoder-potent subspace, and does not exactly match the actual closed-loop BMI kinematics because the latter is only updated every 25 ms and is also smoothed by the $\mathbf{M}_1$ matrix (1). To isolate the *change* in neural push due to the monkey detecting and correcting for the cursor jump perturbation, we first subtracted from each perturbed trial's jump-aligned neural push the mean neural push on unperturbed trials towards the same target. To compute this unperturbed neural push, we aligned each unperturbed trial's neural push to when a cursor jump would have happened on said trial if the jumped trial's perturbation condition had been in effect. Thus, the time-varying neural push normally used to generate the typical trajectory to a target was subtracted away from the measured jump-induced change in neural push. To examine how neural activity changed due to perturbations without regard to the decoder, we simply measured perturbation-induced change in firing rate averaged across all channels. We subtracted each channel's trial-averaged unperturbed firing rate (aligned to when the jump would have happened as described above) from perturbed trial's jump-aligned firing rates, and then averaged this change across all channels. We note that this is a rudimentary and conservative analysis because it averages across channels and does not take the absolute value of change of firing rates. Consequently, increases and decreases of activity across channels, targets, and perturbation directions could have (but in practice did not) cancel out.

III. Results

A. Rapid Correction of a BMI Cursor Jump Perturbation

We measured how a monkey performing a two-dimensional cursor task using a high performing BMI responded to an unexpected error resulting from a cursor jump perturbation. The perturbation was applied in one of two possible directions and at one of four possible phases on 25% of trials. The monkey smoothly and rapidly corrected the trajectory error introduced by all cursor jump conditions. The example cursor trajectories in Fig. 1B show typical single-trial responses to perturbations, while Fig. 2 shows quantified average kinematic responses. For perturbations during the early (30 mm) and middle (60 mm) phase of a trial, the monkey maintained his velocity towards the target and compensated for upward (downward) cursor jumps by applying a corrective downward (upward) velocity component. Note that the immediate and continued reduction in distance to target for these conditions in Fig. 2. is due to ongoing velocity along the task axis and does not reflect a (biologically impossible) instantaneous corrective response. When perturbations occurred just as the cursor entered the target acceptance region (100 mm) or during the hold epoch, the monkey corrected for this error by bringing the cursor back towards the target with minimal movement in the direction orthogonal to the perturbation. There were no significant performance trends as a function of experimental session number, i.e. no sign of across-study learning.

The performance penalty introduced by the cursor jump, in terms of additional time needed to make a successful target acquisition, increased as a function of how late into the trial

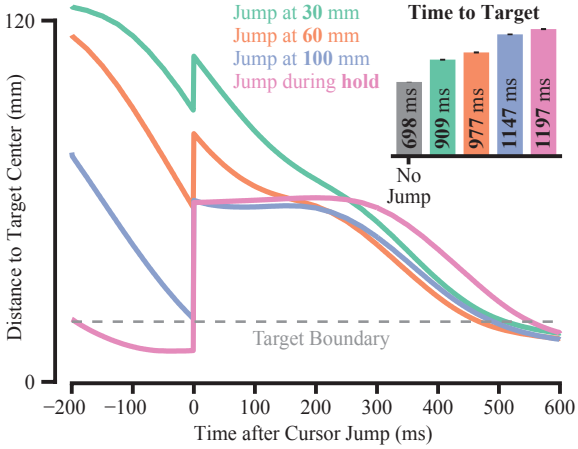

Fig. 2. Kinematics of responses to cursor jumps. Mean distance to target is plotted as a function of time relative to the cursor jump. The monkey was able to rapidly correct the cursor position error and acquire the target. Data are combined across all fourteen experimental sessions. (Inset) Mean ± SEM time to target for unperturbed and cursor jump trials. All conditions' means were significantly different from one another's ($p < 0.001$).

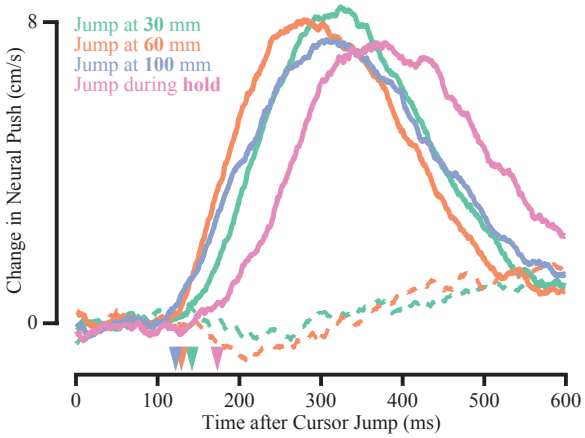

Fig. 3. Neural responses to cursor jump perturbations. Solid traces show how the population neural response to a cursor jump affects decoded velocity ("neural push") in the direction opposing the perturbation. Downward arrowheads mark the response latency, defined as when a significant change ($p < 0.001$) in neural push was first detected. These latencies were 142, 129, 122, and 173 ms for the 30, 60, 100 mm, and hold jump conditions, respectively. Dashed traces show (minimal) change in neural push along the task axis (i.e., orthogonal to the perturbation) for the two conditions where continued velocity towards the target was required. All traces are means across trials from all experimental sessions.

the perturbation occurred (Fig. 2 inset). This penalty was 210 ± 9 ms (difference of means ± std of this difference) for jumps at 30 mm towards the target, 279 ± 10 ms for jumps at 60 mm, 449 ± 8 ms for jumps at 100 mm, and 499 ± 10 ms if the jump happened during the target hold.

B. Neural Response to Cursor Jump Perturbation

We next examined the neural response to a cursor jump perturbation. We specifically were interested in how the neural response evolved as viewed through the decoder, which maps firing rates to BMI cursor velocity. Although all of the recorded channels contributed differently to the decoder, we can summarize the net effect of cursor jump-evoked population activity change as a change in "neural push" contribution to velocity (Fig. 3). This response began as soon as ~115 ms (after accounting for an additional 7 ms monitor refresh latency) following the cursor jump and pushed the cursor in the direction opposing the perturbation. For the jump during hold condition, the corrective response occurred considerably slower (~165 ms after perturbation).

We found no evidence of a transient, maladaptive change in neural push that could reflect (unwanted) influence on decoded velocity by neural correlates of processing the perturbation-induced error. This effect would have manifest itself in our analysis as a negative neural push along the perturbation axis (transiently magnifying the perturbation) or as an early occurring negative change in neural push along the task axis during the 30 mm or 60 mm conditions (interfering with the monkey's ability to continue generating neural push to move the cursor towards the target). We did observe slight changes in neural push along the task axis later, likely reflecting the monkey re-aiming the cursor towards the center of the target from its newly offset position.

Finally, we looked directly for an early perturbation-induced neural activity transient that did not affect decoded velocity. Fig. 4 shows that there was indeed an initial perturbation-evoked change of firing rate which preceded the earliest change in neural push. This response was particularly pronounced for the jump at 100 mm and jump during hold conditions, which may reflect the response being easier to detect in these conditions because of their lower overall firing rates at the time of the perturbation. This initial transient diminished before firing rates again increased, now as a decoder-potent error correcting response. Note that in this analysis (unlike in Fig. 3) we are plotting change in neural push *magnitude* (i.e. changes in both x- and y- velocity) so as to explicitly compare overall firing rate change to how neural activity changes in the decoder-potent dimensions, without restricting the latter to only the perturbation axis.

IV. Discussion

A monkey controlling a BMI was able to correct for relatively large perturbations quickly and accurately. This effective error correction was enabled by two factors. Firstly, there was a rapidly occuring and strong corrective neural response beginning 115 ms after the cursor jump. This measurement is consistent with feedback latency reported in [15] and is only slightly slower than the latency of visually-driven responses in motor cortex (80-100 ms) during arm control [17], [18]. This suggests that the additional computational overhead of mapping error to the appropriate BMI corrective response (at least in this high performing biomimetic BMI) is modest. This latency is nonetheless considerably slower than the ~50 ms needed for prioriception-mediated corrective M1 responses [12]; BMIs with prioprioceptive write-in may potentially enable even faster error corrections [27]. Secondly, the neural response pushed the cursor in the direction opposing the perturbation as soon as the response began to influence decoded velocity. We observed no maladaptive velocity transient due to, for example, "leakage" of the initial neural computations into the "output-potent" subspace (borrowing concepts from [20]). Rather, the earliest neural response was confined to the "decoder-null" subspace.

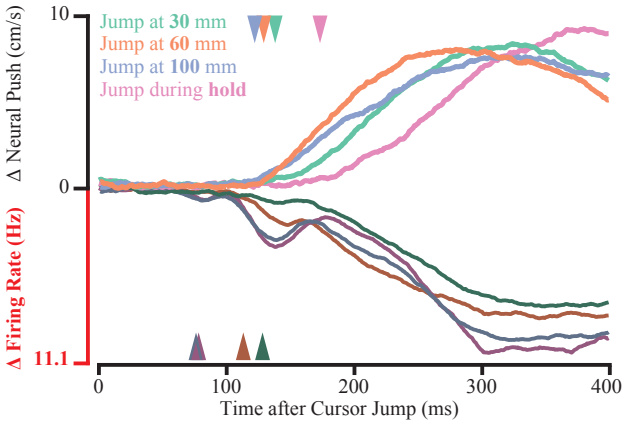

Fig. 4. Comparing overall perturbation-induced neural response to the projection of this response onto the decoder. Brighter upper traces show change in decoder-potent neural push magnitude. Darker lower traces show change of firing rate averaged across the recorded population, i.e., average firing rate on cursor jump trials minus average firing rate on unperturbed trials. Downward arrowheads mark the response latency ($p < 0.001$) for changes in neural push (earliest of change along task axis or perturbation axis). Upward arrowheads mark the response latency for changes in firing rate (128, 113, 76, and 78 ms for the 30, 60, 100 mm, and hold jump conditions, respectively). Traces are means across all recorded trials.

We were surprised to find that the monkey did not correct errors occuring later in the trial with more vigor, as might be expected under many optimal feedback control cost functions [10]. Rather, the time to target penalty increased with later cursor jumps. Due to the geometry of the task, the jump offset increased straight-line distance to the target more if it happened closer to the target. Since the neural response did not increase faster or reach a greater maximum for later perturbations, it follows that the time to correct the error was longer for the later perturbations. This could be due to a lack of motivation on the part of the monkey to make stronger corrections, or alternatively because of ceiling effects in terms of how strong a neural corrective response could be generated. Future work will be needed to differentiate between these possibilities. Furthermore, we found that the corrective response to perturbations during the hold time was delayed compared to corrective responses during cursor movement, and that the response to jumps midway through the trajectory increased fastest. This raises the possibility that the motor system is better primed to make strong corrective BMI responses during movement compared to when it is attempting to hold the cursor in place.

Incorporating this knowledge about how the brain responds to perturbations into future decoder designs may lead to more robust BMIs that can explicitly facilitate a user's error corrections.

ACKNOWLEDGMENT

We thank M. Mazariegos, M. Wechsler, C. Sherman, & L. Yates for expert surgical assistance & veterinary care, B. Oskotsky for IT support, and B. Davis, & E. Casteneda for administrative assistance.

REFERENCES

[1] K. Ganguly and J. M. Carmena, "Emergence of a stable cortical map for neuroprosthetic control," *PLoS Biol.*, vol. 7, no. 7, p. e1000153, Jul. 2009.

[2] V. Gilja, P. Nuyujukian, C. A. Chestek, J. P. Cunningham, B. M. Yu, J. M. Fan, M. M. Churchland, M. T. Kaufman, J. C. Kao, S. I. Ryu, and K. V Shenoy, "A high-performance neural prosthesis enabled by control algorithm design," *Nat. Neurosci.*, vol. 15, no. 12, pp. 1752–7, Dec. 2012.

[3] D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, no. 5574, pp. 1829–32, Jun. 2002.

[4] 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 Biol.*, vol. 1, no. 2, pp. 193–208, Nov. 2003.

[5] D. Bacher, B. Jarosiewicz, N. Y. Masse, S. D. Stavisky, J. D. Simeral, K. Newell, E. M. Oakley, S. S. Cash, G. Friehs, and L. R. Hochberg, "Neural Point-and-Click Communication by a Person With Incomplete Locked-In Syndrome," *Neurorehabil. Neural Repair*, Nov. 2014.

[6] C. Pandarinath, P. Nuyujukian, V. Gilja, C. H. Blabe, J. A. Perge, B. Jarosiewicz, L. R. Hochberg, K. V Shenoy, and J. M. Henderson, "Application of a high performance intracortical brain computer interface for communication in a person with amyotrophic lateral sclerosis," in *Program No. 252.08 Neuroscience 2014 Abstracts*, 2014, pp. 50–52.

[7] L. R. Hochberg, D. Bacher, B. Jarosiewicz, N. Y. Masse, J. D. Simeral, J. Vogel, S. Haddadin, J. Liu, S. S. Cash, P. van der Smagt, and J. P. Donoghue, "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm," *Nature*, vol. 485, no. 7398, pp. 372–5, May 2012.

[8] J. L. Collinger, B. Wodlinger, J. E. Downey, W. Wang, E. C. Tyler-Kabara, D. J. Weber, A. J. C. McMorland, M. Velliste, M. L. Boninger, and A. B. Schwartz, "High-performance neuroprosthetic control by an individual with tetraplegia.," *Lancet*, vol. 381, no. 9866, pp. 557–64, Feb. 2013.

[9] D. M. Wolpert and R. C. Miall, "Forward Models for Physiological Motor Control," *Neural networks*, vol. 9, no. 8, pp. 1265–1279, Nov. 1996.

[10] S. H. Scott, "Optimal feedback control and the neural basis of volitional motor control," *Nat. Rev. Neurosci.*, vol. 5, no. 7, pp. 532–46, Jul. 2004.

[11] R. Shadmehr, M. A. Smith, and J. W. Krakauer, "Error correction, sensory prediction, and adaptation in motor control," *Annu. Rev. Neurosci.*, vol. 33, pp. 89–108, Jan. 2010.

[12] J. A. Pruszynski, I. Kurtzer, J. Y. Nashed, M. Omrani, B. Brouwer, and S. H. Scott, "Primary motor cortex underlies multi-joint integration for fast feedback control," *Nature*, vol. 478, no. 7369, pp. 387–90, Oct. 2011.

[13] B. Jarosiewicz, S. M. Chase, G. W. Fraser, M. Velliste, R. E. Kass, and A. B. Schwartz, "Functional network reorganization during learning in a brain-computer interface paradigm," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 105, no. 49, pp. 19486–91, Dec. 2008.

[14] S. M. Chase, R. E. Kass, and A. B. Schwartz, "Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex," *J. Neurophysiol.*, vol. 108, no. 2, pp. 624–44, Jul. 2012.

[15] M. D. Golub, B. M. Yu, and S. M. Chase, "Internal models engaged by brain-computer interface control," in *34th Annual International Conference of the IEEE EMBS*, 2012, vol. 2012, pp. 1327–30.

[16] P. T. Sadtler, K. M. Quick, M. D. Golub, S. M. Chase, S. I. Ryu, E. C. Tyler-Kabara, B. M. Yu, and A. P. Batista, "Neural constraints on learning," *Nature*, vol. 512, no. 7515, pp. 423–426, Aug. 2014.

[17] A. Georgopoulos, J. Kalaska, R. Caminiti, and J. Massey, "Interruption of Motor Cortical Discharges Subserving Aimed Arm Movements," *Exp. Brain Res.*, vol. 49, pp. 327–340, 1983.

[18] P. S. Archambault, S. Ferrari-Toniolo, and A. Battaglia-Mayer, "Online control of hand trajectory and evolution of motor intention in the parietofrontal system," *J. Neurosci.*, vol. 31, no. 2, pp. 742–52, Jan. 2011.

[19] J. Pruszynski, M. Omrani, and S. Scott, "Goal-Dependent Modulation of Fast Feedback Responses in Primary Motor Cortex," *J. Neurosci.*, vol. 34, no. 13, pp. 4608–4617, Mar. 2014.

[20] M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V Shenoy, "Cortical activity in the null space: permitting preparation without movement," *Nat. Neurosci.*, vol. 17, no. 3, pp. 440–8, Mar. 2014.

[21] K. V. Shenoy and J. M. Carmena, "Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces," *Neuron*, vol. 84, no. 4, pp. 665–680, Nov. 2014.

[22] J. M. Fan, P. Nuyujukian, J. C. Kao, C. A. Chestek, S. I. Ryu, and K. V Shenoy, "Intention estimation in brain–machine interfaces," *J. Neural Eng.*, vol. 11, no. 1, p. 016004, Feb. 2014.

[23] C. A. Chestek, V. Gilja, P. Nuyujukian, J. D. Foster, J. M. Fan, M. T. Kaufman, M. M. Churchland, Z. Rivera-Alvidrez, J. P. Cunningham, S. I. Ryu, and K. V Shenoy, "Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex," *J. Neural Eng.*, vol. 8, no. 4, p. 045005, Aug. 2011.

[24] W. Q. Malik, W. Truccolo, E. N. Brown, and L. R. Hochberg, "Efficient decoding with steady-state Kalman filter in neural interface systems," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 1, pp. 25–34, Feb. 2011.

[25] S. D. Stavisky, J. C. Kao, P. Nuyujukian, S. I. Ryu, and K. V Shenoy, "Hybrid Decoding of Both Spikes and Low-Frequency Local Field Potentials for Brain-Machine Interfaces," in *36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2014, pp. 3041–3044.

[26] M. C. Dadarlat, J. E. O'Doherty, and P. N. Sabes, "A learning-based approach to artificial sensory feedback leads to optimal integration," *Nat. Neurosci.*, Nov. 2014.