Increasing the Performance of Cortically-Controlled Prostheses

Krishna V. Shenoy[1,2], Gopal Santhanam[1], Stephen I. Ryu[1,3], Afsheen Afshar[1,4], Byron M. Yu[1],
Vikash Gilja[5], Michael D. Linderman[1], Rachel S. Kalmar[2], John P. Cunningham[1],
Caleb T. Kemere[1], Aaron P. Batista[1,2], Mark M. Churchland[1,2], Teresa H. Meng[1]

[1]Department of Electrical Engineering, [2]Neurosciences Program, [3]Department of Neurosurgery,
[4]Medical Scientist Training Program and [5]Department of Computer Science
Stanford University, Stanford, CA 94305, USA

Abstract— **Neural prostheses have received considerable attention due to their potential to dramatically improve the quality of life of severely disabled patients. Cortically-controlled prostheses are able to translate neural activity from cerebral cortex into control signals for guiding computer cursors or prosthetic limbs. Non-invasive and invasive electrode techniques can be used to measure neural activity, with the latter promising considerably higher levels of performance and therefore functionality to patients. We review here some of our recent experimental and computational work aimed at establishing a principled design methodology to increase electrode-based cortical prosthesis performance to near theoretical limits. Studies discussed include translating unprecedentedly brief periods of "plan" activity into high information rate (6.5 bits/s) control signals, improving decode algorithms and optimizing visual target locations for further performance increases, and recording from chronically implanted arrays in freely behaving monkeys to characterize neuron stability. Taken together, these results should substantially increase the clinical viability of cortical prostheses.**

Keywords— **Brain-computer interface, brain-machine interface, neural prostheses, premotor cortex.**

I. Introduction

EACH year hundreds of thousands of people suffer from neurological injuries and disease, resulting in the permanent loss of motor function. In many cases, the disability is so severe that it is not even possible to feed oneself or communicate. Though surgical and medical interventions have made it possible to repair peripheral nerves and promote recovery in many cases, most central nervous system impairments still do not have effective treatments. Electronic medical systems that interface with the nervous system, termed neural prostheses, have started to fill some of these treatment gaps.

The ultimate goal of any prosthesis is to restore normal function. Though complete restoration is ideal, prostheses are clinically viable when the anticipated quality of life improvement outweighs the potential risks. Since neural prostheses must often measure or perturb neurons in the central nervous system, non-invasive techniques are particularly attractive and have been investigated extensively. Invasive electrode techniques have also become a major research thrust due to their high signal quality (i.e., action potentials from individual neurons). If, as a result, prosthetic performance can substantially surpass what is possible with non-invasive measurement techniques, and surgical risk can be sufficiently minimized, it is anticipated that invasive electrode-based prostheses will find more widespread use. Thus we need to understand the fundamental, neurobiologically dictated performance limits of electrode-based neural prostheses and establish a principled design methodology to achieve these theoretical limits. If successful, these advances should help accelerate the translation of laboratory proof-of-concept systems into widespread clinical use. We review below some of our recent experimental and computational work aimed at increasing the performance of cortically-controlled prostheses.

A. Motor and Communication Prostheses

Figure 1 illustrates the basic operating principle behind motor and communication prostheses. Neural activity from various brain regions is electronically processed to create control signals for enacting the desired movement. When permanently-implanted arrays of electrodes are employed, it is possible to identify individual neurons near the tip of each electrode through a process termed action potential (spike) sorting [7, 8]. After determining how each neuron responds before and during a movement, typically accomplished by correlating arm movements with neural activity, estimation (decode) algorithms can infer the desired movement from only the neural activity [9, 10]. The system can then generate control signals to continuously guide a paralyzed arm or a prosthetic arm through space (motor prosthesis), or to position a computer cursor on the desired letter on a keyboard (communication prosthesis).

Motor and communication prostheses are quite similar conceptually, but important differences critically affect their design. Motor prostheses must produce movement trajectories as accurately as possible in order to enact precisely

Fig. 1. Concept sketch of cortically-controlled prostheses. Several cortical areas in rhesus monkeys, and in homologous areas in humans, participate in the preparation and execution of natural arm movements. Areas include the medial intra-parietal area (MIP) / parietal reach region (PRR) with largely plan activity [1, 2], the dorsal aspect of pre-motor cortex (PMd) with both plan and movement activity [3], and motor cortex (M1) with largely movement activity [4–6]. Neural prostheses measure electrical neural signals (action and local field potentials) using arrays of chronically-implanted electrodes, extract action potential times for each neuron on every electrode and, finally, estimate the desired arm movement and generate control signals for guiding prosthetic devices.

the desired movement. Thus *continuous* prosthetic guidance is necessary, and performance measures must quantify the similarity between the prosthetic movement and the desired movement. In contrast, communication prostheses are concerned with information throughput from the subject to the world; this makes the speed and accuracy with which keys on a keyboard can be selected of primary importance. Although a continuously guided motor prosthesis could be used to convey information by moving to a key, only the key eventually struck contributes to information conveyance. Thus, simple *discrete* prosthesis positioning is sufficient for a communication prosthesis [1]. For example, if it is possible to predict which letter on a keyboard is desired, a computer cursor could be directly positioned on ("hopped to") that key as opposed to sliding out to strike the key.

B. Plan and Movement Activity

Two types of neural activity are well suited for driving prosthetic movements. Plan activity is present before arm movements begin and is believed to reflect preparatory processing required for the fast and accurate generation of movement [11]. This activity is readily observed before movement initiation in a delayed reach task. Delayed reach tasks begin by presenting a visual reach target. After a brief delay period, a "go" cue indicates that a reach may begin. Plan activity in premotor cortex appears soon after reach-target onset, persists until just after the "go" cue, and is tuned for the direction and extent of the upcoming movement [12]. Movement activity follows plan activity in a delayed

reach task, being present immediately before and during the movement. Movement activity is tuned for both the direction and speed of arm movement [13].

Until recently, both motor and communication prostheses have focused exclusively on movement activity. Movement activity can be decoded to generate instantaneous direction and speed signals, which are used to guide a prosthetic device along a trajectory [4–6]. Motor prostheses must incorporate movement activity since the goal is to recreate the desired movement path and speed.

In contrast, communication prostheses are not obliged to move the prosthetic device along a continuous path in order to strike a target such as a key on an on-screen keyboard (although several system do operate in this fashion [14–16]). Instead, if target location can be estimated directly from neural plan activity, the cursor can be positioned immediately on the desired key. Recent reports suggest that there may be considerable performance benefit to using plan activity and direct-positional prosthesis control [1, 2], and we review directly below recent confirmation [3].

II. RECENT RESULTS

A. A 6.5 bps Communication Prosthesis

To investigate how quickly and accurately a communication prosthesis can perform when driven by cortical plan activity, we conducted a series of experiments and computational simulations [3]. As shown in Fig. 2a, we trained monkeys to fixate [17] and touch central targets, and plan to reach to a visual target that could appear at one of several (2, 4, 8 or 16) different locations. Meanwhile we collected action potentials from all neurons recorded with a 96 electrode array (typically 100-200 neural units) and used the number of action potentials during an integration period (T_{int}) to predict where the monkey was planning to reach. If our prediction, made using maximum likelihood (ML) techniques and Gaussian or Poisson neural response models, matched the true target location we displayed a circle around the target, played an auditory tone and provided a liquid reward to indicate a successful trial. In this way we were able to assess how fast selections could be made, and how often the selections were correct.

While it is possible to predict the desired target location using neural activity from a variety of different T_{int} windows [1, 2], we optimized T_{int} in order to maximize performance. We first conducted a series of delayed reach experiments to determine how rapidly movement plans form in PMd following the onset of a peripheral reach target. This time (T_{skip}) includes the visual latency for PMd to "see" the target, the time required to select the desired target from among other potential targets (not shown in Fig. 2), and the time for a plan to stabilize. Using ML based methods we found T_{skip} to be approximately 150 ms across a variety of conditions, which is in good agreement with "normalized variance" based methods [11] and dynamical systems methods [18] applied to similar data. The T_{int} window should
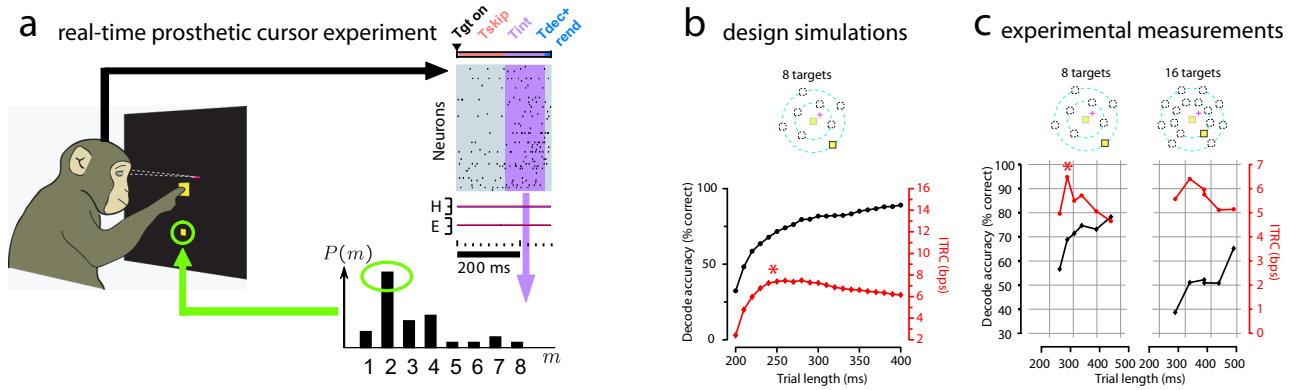
Fig. 2. Overview of recent high-performance communication prosthesis experiments, simulations and results. (**a**) Real-time prosthetic cursor placement task begins by fixating and touching central targets, followed by the appearance of a peripheral target to which the monkey plans (but does not execute) a reach. A period of neural data following this target onset (Tgt on) is set aside (T_{skip}), a period of neural data (T_{int}) is then analyzed in order to estimate the desired target (target $\hat{m}$ with highest probability, $P(\hat{m})$; $\hat{m}$=2 in this example) which could have appeared in one of 8 locations in this task and, following a brief computational decode and display rendering "overhead" period ($T_{dec+rend}$) the predicted target is encircled. (**b**) Design simulations performed with neural data to help identify best T_{int} time (data and simulations to optimize T_{skip} not shown). Single-trial decode accuracy (black) and information transfer rate capacity (ITRC, red) were calculated offline from experiment H20041118. An 8 target configuration was used in the experiment. Trial length was taken to be (T_{skip} + T_{int} + $T_{dec+rend}$) ms with T_{skip} fixed at 150 ms and $T_{dec+rend}$ fixed at 40 ms. T_{int} was varied and performance extrapolated. Red asterisk indicates the best predicted ITRC (7.7 bps) and associated T_{int} (70 ms; T_{int} = trial length - T_{skip} - $T_{dec+rend}$). (**c**) Experimentally measured single-trial accuracy (black) and ITRC (red). Performance is plotted for an 8 target configuration (to be compared with panel b) and a 16 target configuration across varying total trial lengths. All results are from monkey H and each data symbol represents online performance calculated from one experiment (many hundreds of trials). Note that within each target configuration, accuracy increases as a function of T_{int} and ITRC shows a peak (red asterisk) close that predicted by simulation (compare with asterisk in panel b). Across target configurations (2 and 4 target results not shown), average accuracy drops as a function of the number of targets while ITRC increases until saturating around 6.5 bps with both 8 and 16 target configurations. Adapted from [3].

begin immediately after T_{skip}: any sooner would introduce "plan noise" into the decoder, any later would slow down trial pace.

Next, as shown in Fig. 2b, we ran simulations to determine the optimal T_{int} duration. While the accuracy with which the reach target can be predicted continues to increase as T_{int} expands, overall performance (quantified in terms of information transmission rate capacity, ITRC, in bits/s) has a peak [3]. This is because beyond some T_{int} point of diminishing returns (i.e., the optimal T_{int}), accuracy fails to increase rapidly enough to overcome the slowdown in trial pace accompanying long T_{int} durations.

Finally, we conducted experiments to verify the existence and location of the optimal ITRC operating point. As shown in Fig. 2c, the ITRC peak occurred near the anticipated T_{int} value and 6.5 bps performance was achieved. This 6.5 bps result required optimization of several parameters including T_{skip} and T_{int} and compares favorably with previous state-of-the-art reports (e.g., 1.6 bps [19]). Intriguingly, ~6.5 bps was also achieved with a 16 target configuration which increased the task difficulty (4 bits/trial maximum, as opposed to 3) but resulted in lower single-trial accuracy. It appears that performance may be saturating; pushing beyond 6.5 bps will require other optimizations, as discussed below.

B. Beyond 6.5 bps — Countering Response Non-Stationarity

While 6.5 bps represents a substantial increase in performance, it fell 1.2 bps short of our expectations based on simulations (Fig. 2b). Further analysis revealed that this was due in part to changes in the response of individual neurons: firing rates often became more or less intense

depending on when, in a high-speed sequence of prosthetic target trials, the response was measured [20]. Fig. 3 shows one neuron (left panel) that responds more to the 230° reach target when it appears first in a "chain" than when it appears third. Similarly, the neuron in the right panel is less responsive to the 310° target when it appears first than when it appears third. Fortunately, by recognizing this and other (e.g., motivational state) response non-stationarities, adaptive "auto-normalizing" algorithms can partly counter these effects [21].

C. Beyond 6.5 bps — Optimizing Target Placement

The performance results above were obtained using "canonical" reach target layouts, which are typically symmetric target configurations which reflect the experimenter's best guess as to which targets will be maximally discriminible. Canonical layouts are independent of the tuning properties of the neurons available and can limit performance. We recently sought to increase decode accuracy by judiciously selecting the locations of the reach targets based on the characteristics of the neural population at hand. We developed an optimal target placement (OTP) algorithm that approximately maximizes decode accuracy with respect to target locations [22]. Using maximum likelihood decoding, the optimal target placement algorithm yielded up to 11 and 12% accuracy improvement for two and sixteen targets, respectively. The OTP solution for a sixteen target experiment is shown in Fig. 4. For four and eight targets, gains were more modest (5 and 3%, respectively) as the target layouts found by the algorithm closely resembled the canonical layouts. Thus,
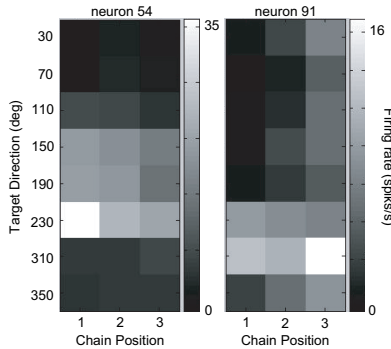
Fig. 3. Directional tuning curves as a function of chain position in a high-speed prosthetic cursor experiment. Average firing rate from two neurons (one neuron in left panel; one in right panel) from an eight-target experiment as illustrated in Fig. 2a are plotted versus the order presented ("chain position"). Figure from [20].

the algorithm can serve not only to find target layouts that outperform canonical layouts, but it can also confirm or help select among multiple canonical layouts. These results indicate that the OTP algorithm is a valuable tool for further increasing prosthesis performance.

D. Beyond 6.5 bps — Internally Selected Targets

The performance results above were also obtained by visually cueing a sequence of targets. By increasing the speed with which these targets were presented, the maximum rate of *externally cued* reach planning was assessed. The true rate of planning reaches to spatial targets may be considerably higher if these targets are selected internally. To conduct experiments where targets are displayed and monkeys are able to plan reaches to these targets at will, we must first develop algorithms capable of automatically detecting when neurons are planning reaches. We have recently reported preliminary results demonstrating this detection [23], using simple ML estimation techniques [1]. Considerably better "state estimation" is possible using more sophisticated Hidden Markov Models [24]. With these algorithms we hope to measure the maximum rate of internal target selection and planning.

E. Neural Stability & Adaptive Algorithms

Finally we must consider the stability of chronically-implanted electrode arrays and how best to design adaptive algorithms to contend with any instability or non-stationarity. While it is well recognized that electrode arrays are able to record from the same cortical area for many months or even a few years without substantial signal degradation [28], it is also well recognized that immunological responses, and likely other still unidentified mechanisms, cause complete or near complete signal loss on the time scale of months to years [29]. Several efforts are underway to design novel materials and coatings to dramatically improve electrode stability. However, the stability of neural activity, mechanisms of instabilities, and the design of adaptive spike sorting and decoding algorithms remain poorly understood.
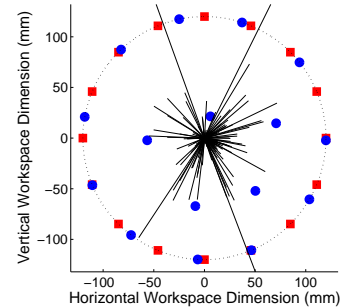

Fig. 4. Sixteen target optimal target placement (OTP) example. Blue circles: OTP solution; red squares: uniformly spaced targets around a circle, as an experimenter might use without OTP methods. Note, two concentric rings with eight uniformly spaced targets each is another likely target pattern lacking OTP methods. Black lines show the preferred direction (line angle) and tuning strength ($\propto$ line length) for each of 189 neural units recorded from a chronic electrode array. Workspace bound shown as a dotted line (120 mm radius). From [22].

To answer these questions, we recently built a light-weight, low-power, autonomous recording system capable of recording broadband neural signals and 3D head accelerations [25–27]. The entire system fits within a small head-mounted enclosure, as shown in Fig. 5a. Preliminary data suggest that rapid accelerations can cause abrupt shifts in action potential waveforms as well as longer timescale changes, as shown in Fig. 5b. By tracking neural waveforms continuously it should be possible to certify neurons as being the same across days. This is critical for multi-day plasticity experiments and to maintain prosthetic decode performance, which would be compromised by uncertain neuron identity.

III. CONCLUSION

In an attempt to increase the clinical viability of cortically-controlled communication prostheses, we have designed and demonstrated a system capable of 6.5 bps performance. Further performance gains are achieved with a better understanding of cortical dynamics and response non-stationarities, employing optimal target placement algorithms, exploring internally selected target tasks, and by characterizing electrode-neuron stability so as to optimally design adaptive neuron-identification algorithms.

REFERENCES

[1] K. V. Shenoy, D. Meeker, S. Cao, S. A. Kureshi, B. Pesaran, P. Mitra, C. A. Buneo, A. P. Batista, J. W. Burdick, and R. A. Andersen, "Neural prosthetic control signals from plan activity," *NeuroReport*, vol. 14, pp. 591–596, 2003.

[2] S. Musallam, B. D. Corneil, B. Greger, H. Scherberger, and R. A. Andersen, "Cognitive control signals for neural prosthetics," *Science*, vol. 305, pp. 258–262, 2004.

[3] G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy, "A high-performance brain computer interface," *Nature, in press*, 2006.

Fig. 5. Neural recordings from an unrestrained monkey over a period of 48 hours. (**a**) Illustration of HermesB enclosure mounted on monkey's head. (**b**) Spike waveforms of two neurons for selected 5 minute blocks. Colored region indicates 10th–90th percentile in amplitude. Horizontal lines indicate maximum and minimum voltage for each unit. Waveforms shown are recorded from a single channel using the same signal conditioning path. Note that between 17:44 (day 1) and 07:04 (day 2) V_{pp}, the peak-to-peak voltage, of the green waveform increases, while V_{pp} of the blue waveform decreases, showing that waveform changes cannot be attributed to fluctuations in signal conditioning pathway (connectorization, amplifiers, ADC, etc.). Adapted from [25–27].

[4] D. M. Taylor, S. I. Helms Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, pp. 1829–1832, 2002.

[5] M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J.P. Donoghue, "Instant neural control of a movement signal," *Nature*, vol. 416, pp. 141–142, 2002.

[6] J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis, "Learning to control a brain-machine interface for reaching and grasping by primates," *PLoS Biology*, vol. 1, pp. 193–208, 2003.

[7] G. Santhanam, M. Sahani, S. I. Ryu, and K. V. Shenoy, "An extensible infrastructure for fully automated spike sorting during online experiments," *Proc. 26th Ann. International Conf. IEEE EMBS, San Francisco, CA*, pp. 4380–4384, 2004.

[8] Z. S. Zumsteg, C. Kemere, S. O'Driscoll, G. Santhanam, R. E. Ahmed, K.V. Shenoy, and T. H. Meng, "Power feasibility of implantable digital spike sorting circuits for neural prosthetic systems," *IEEE Trans. in Neural Systems and Rehab. Eng.*, 2005.

[9] C. T. Kemere, K. V. Shenoy, and T. H. Meng, "Model-based neural decoding of reaching movements: a maximum likelihood approach," *IEEE Trans. on Biomed. Eng.*, vol. 51, pp. 925–932, 2004.

[10] B. M. Yu, C. Kemere, G. Santhanam, A. Afshar, S. I. Ryu, T. H. Meng, M. Sahani, and K. V. Shenoy, "Mixture of trajectory models for neural decoding of goal-directed movements," *Program No. 520.18. 2005 Abstract Viewer/Itinerary Planner. Washington, DC: Soc. for Neurosci.*, Nov. 2005.

[11] M. M. Churchland, B. M. Yu, S. I. Ryu, G. Santhanam, and K. V. Shenoy, "Neural variability in premotor cortex provides a signature of motor preparation," *J. Neurosci.*, vol. 26, pp. 3697–3712, 2006.

[12] J. Messier and J. F. Kalaska, "Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task," *J Neurophysiol*, vol. 84, pp. 152–165, 2000.

[13] D. W. Moran and A.B. Schwartz, "Motor cortical activity during drawing movements: population representation during sprial tracing," *J. Neurophysiol.*, vol. 82, no. 5, pp. 2693–2704, 1999.

[14] P. R. Kennedy and R. A. E. Bakay, "Restoration of neural output from a paralyzed patient by a direct brain connection," *NeuroReport*, vol. 9, pp. 1707–1711, 1998.

[15] E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Morann, "A brain-computer interface using electrocorticographic signals in humans," *J. Neural Eng.*, vol. 1, pp. 63–71, 2004.

[16] J.R. Wolpaw and D.J. McFarland, "Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans," *PNAS*, vol. 101, no. 51, pp. 17849–17854, 2004.

[17] A. P. Batista, B. M. Yu, G. Santhanam, S. I. Ryu, A. Afshar, and K. V. Shenoy, "Influence of eye position on endpoint decoding accuracy in dorsal premotor cortex," *2006 Abstract Viewer/Itinerary Planner, Atlanta, GA: Soc. for Neurosci.*, Oct. 2006.

[18] B. M. Yu, A. Afshar, G. Santhanam, S.I. Ryu, K. V. Shenoy, and M. Sahani, "Extracting dynamical structure embedded in neural activity," *Proceedings of Neural Information Processing Society (NIPS), in press*, 2006.

[19] D. M. Taylor, S. I. Helms Tillery, and A. B. Schwartz, "Information conveyed through brain-control: cursor vs. robot," *IEEE Trans. on Neural Sys. and Rehab. Engr.*, vol. 11, pp. 195–199, 2003.

[20] R. S. Kalmar, V. Gilja, G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy, "PMd delay activity during rapid sequential movement plans," *Program No. 519.17. 2005 Abstract Viewer/Itinerary Planner. Washington, DC: Soc. for Neurosci.*, Nov. 2005.

[21] V. Gilja, R. S. Kalmar, G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy, "Trial-by-trial mean normalization improves plan period reach target decoding," *Program No. 519.18. 2005 Abstract Viewer/Itinerary Planner. Washington, DC: Soc. for Neurosci.*, Nov. 2005.

[22] J. P. Cunningham, B. M. Yu, and K. V. Shenoy, "Optimal target placement for neural communication prosthses," in *Proc. of the 28th Annual International Conf. of the IEEE EMBS, New York, NY*, 2006, submitted.

[23] A. Afshar, N. Achtman, G. Santhanam, S. I. Ryu, B. M. Yu, and K. V. Shenoy, "Free-paced target estimation in a delayed reach task," *Program No. 401.13. 2005 Abstract Viewer/Itinerary Planner. Washington, DC: Soc. for Neurosci.*, Nov. 2005.

[24] C. T. Kemere, B. M. Yu, G. Santhanam, S. I. Ryu, A. Afshar, T. H. Meng, and K. V. Shenoy, "Hidden markov models for spatial and temporal estimation for prosthetic control," *2006 Abstract Viewer/Itinerary Planner, Atlanta, GA: Soc. for Neurosci.*, Oct. 2006.

[25] M. D. Linderman, V. Gilja, G. Santhanam, A. Afshar, S. Ryu, T. H. Meng, and K. V. Shenoy, "An autonomous, broadband, multi-channel neural recording system for freely behaving primates," in *Proc. of the 28th Annual International Conf. of the IEEE EMBS, New York, NY*, 2006, submitted.

[26] V. Gilja, M. D. Linderman, G. Santhanam, A. Afshar, S. Ryu, T. H. Meng, and K. V. Shenoy, "Multiday electrophysiological recordings from freely behaving primates," in *Proceedings of the 28th Annual International Conference of the IEEE EMBS, New York, NY*, 2006, submitted.

[27] M. D. Linderman, V. Gilja, G. Santhanam, A. Afshar, S. Ryu, T. H. Meng, and K. V. Shenoy, "Neural recording stability of chronic electrode arrays in freely behaving primates," in *Proc. of the 28th Annual International Conf. of the IEEE EMBS, New York, NY*, 2006, submitted.

[28] S. Suner, M. R. Fellows, C. Vargas-Irwin, G. K. Nakata, and J. P. Donoghue, "Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex.," *IEEE Trans. in Neural Syss and Rehab. Eng.*, vol. 13, pp. 524–541, 2005.

[29] V. S. Polikov, P. A. Tresco, and W. M. Reichert, "Response of brain tissue to chronically implanted neural electrodes," *J. Neurosci. Methods*, vol. 148, pp. 1–18, 2005.