



Published in final edited form as:

J Neural Eng. 2011 August ; 8(4): 045005. doi:10.1088/1741-2560/8/4/045005.

Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex

Cynthia A. Chestek¹, Vikash Gilja², Paul Nuyujukian^{3,4}, Justin D. Foster¹, Joline M. Fan³, Matthew T. Kaufman⁵, Mark M. Churchland¹, Zuley Rivera-Alvidrez¹, John P. Cunningham¹, Stephen I. Ryu⁶, and Krishna V. Shenoy^{1,3,5,*}

¹Department of Electrical Engineering Stanford University, Stanford, CA

²Department of Computer Science Stanford University, Stanford, CA

³Department of Bioengineering Stanford University, Stanford, CA

⁴Stanford Medical School Stanford University, Stanford, CA

⁵Neurosciences Program Stanford University, Stanford, CA

⁶Department of Neurosurgery Palo Alto Medical Foundation, Palo Alto, CA

Abstract

Cortically-controlled prosthetic systems aim to help disabled patients by translating neural signals from the brain into control signals for guiding prosthetic devices. Recent reports have demonstrated reasonably high levels of performance and control of computer cursors and prosthetic limbs, but to achieve true clinical viability the long-term operation of these systems must be better understood. In particular, the quality and stability of the electrically-recorded neural signals requires further characterization. Here, we quantify action potential changes and offline neural decoder performance over 382 days of recording from 4 intracortical arrays in 3 animals. Action potential amplitude decreased by 2.4% per month on average over the course of 9.4, 10.4, and 31.7 months in 3 animals. During most time periods, decoder performance was not well correlated with action potential amplitude ($p > 0.05$ for 3 of 4 arrays). In two arrays from one animal, action potential amplitude declined by an average of 37% over the first 2 months after implant. However, when using simple threshold crossing events rather than well isolated action potentials, no corresponding performance loss was observed during this time using an offline decoder. One of these arrays was effectively used for online prosthetic experiments over the following year. Substantial short-term variations in waveforms were quantified using a wireless system for contiguous recording in one animal, and compared within and between days for all three animals. Overall, this study suggests that action potential amplitude declines more slowly than previously supposed, and performance can be maintained over the course of multiple years when decoding from threshold crossing events rather than isolated action potentials. This suggests that neural prosthetic systems may provide high performance over multiple years in human clinical trials.

1. Introduction

Neural prostheses, also termed Brain Machine Interfaces (BMIs) or Brain Computer Interfaces (BCIs), are an emerging class of medical technology with the potential to improve the quality of life for severely disabled patients. In particular, intracortical signals from multi-electrode arrays that penetrate 1-2 mm into cortex can provide useful control signals.

*corresponding author shenoy@stanford.edu.

Performance from these systems can be high, which motivates determining how long it can be maintained. For example, discrete selections of visual targets on a screen can be made at a rate of 6.5 bps [1], which approximately corresponds to typing on a keyboard at approximately 15 words per minute. For continuous control, cortical signals have been used to guide a computer cursor to targets on a screen [2–13]. Recent work has demonstrated increased speed and reduced variability [13–15]. Further motivating clinical translation, cortical signals have also been used to control the arm and hand on an anthropomorphic robotic arm for self-feeding in primates [16]. offline, information from other degrees of freedom, such as fingers and wrist, have been extracted as well [17]. The clinical need for such control signals is high, since there is a substantial population of paralyzed patients, and state of the art prosthetic limbs currently have more degrees of freedom than one can easily control with conventional approaches [18].

However, to conduct a risk-benefit assessment for human patients, it is important to quantify the stability of cortical BMI performance. The major focus of previous work on this question has been to characterize the tissue response to electrode implantation. On the shortest time scale, after implantation, effects such as edema and hemorrhage are observed [19]. Within weeks, the brain largely recovers from the acute effects [20,21], but it is possible that the extent of this trauma determines the number of viable neurons surrounding the electrode for long term recording [22]. This early period is associated with minimal recorded neural activity [23]. On the timescale of 1-4 months, there is also an encapsulation response, in which astrocytes and microglia form a tight insulative sheath around the electrode [21,24], which can remain in place for years [25]. This has been suggested as the primary mechanism for the degradation of single unit activity over time [26]. One study reported large declines in single unit activity using microwire probes on similar timescales [23]. This effect may reach steady state as early as 12 weeks post implant [21,24,25,27]. However, many BMI studies have been completed using the “Utah” array (as well as microwires) for over 6 months. The Utah array, which is a monolithic array of 100 silicon microelectrodes, is particularly important to understand because it has been FDA approved for use in human patients [7,8,28]. This long-term use raises the question of how histological and electrophysiological measurements are related. More generally, it also raises the question of what fraction of overall signal quality degradation is caused by biological response as opposed to materials and engineering failures.

The discrepancy regarding timescale might be related to the perceived importance of single unit activity (clearly differentiable action potentials arising from individual neurons). Many histological and electrophysiological studies have assumed that single unit activity will be well correlated with BMI performance. Until recently, this was a reasonable assumption since early BMI results emphasized the use of single units [1–3]. In these systems, the activity of individual neurons is differentiated by separating action potentials with distinctive waveforms, a mathematical process termed “spike sorting”. However, with present algorithms, spike sorting leads to very little improvement in BMI performance [11, 29–31]. In fact, high performance studies recently reported have used only “threshold crossing” events, in which the firing rate of one unit per electrode is determined by counting the number of times that the voltage falls below some multiple of the RMS voltage, regardless of whether this activity results from more than one neuron [11,13–15]. Since these studies used at least one array implanted for more than one year, it is important to revisit whether decline in the amplitude of action potentials (V_{pp}) actually corresponds to performance loss.

Several previous studies have examined these questions. One previous study examined day-to-day stability in recordings from the Utah array in a feline model and saw substantial interday changes in neural waveforms which correlated with measures of inflammation [32].

However, they did not quantify the effects on decoder performance with this animal model. A number of papers have attempted to track single units across a few days [28,33–35]. Most relevant to the current study, one study examined the ratio of V_{pp} to the standard deviation in voltage in three macaques with Utah arrays, and found no significant decline [36]. However, if the decline is more subtle than previously supposed, the particular metric used to measure it may not detect the effect over the time periods studied, and it would not be noted subjectively during experiments if it was not also reflected in decoder performance. This study also did not attempt to systematically examine the information content on these arrays over many months, as measured by decoder performance over time.

In this report, we analyzed neural data from 382 days in 3 monkeys and 4 intracortical arrays over various periods during 3.6, 9.4, 10.4 and 31.7 months of implantation. These data were characterized in terms of the peak to peak voltage of action potentials as well as offline decoder performance. Voltage decline over long time scales was slow but significant, with several noticeable “events” in which there were rapid changes. Over time, decoder performance was rarely correlated with V_{pp} , when simple threshold crossings were used. On shorter timescales, large but reversible changes in V_{pp} were quantified. This study empirically demonstrates not only that average V_{pp} decays more slowly than previously supposed, but that decoder performance may decline even more slowly. In particular, it declines in a less direct fashion when using threshold crossings (i.e. multi-unit activity) instead of relying on well-isolated action potentials originating from single neurons.

2. Methods

2.1. Behavioral Tasks

All protocols were approved by the Stanford University Institutional Animal Care and Use Committee. We trained three rhesus macaques to perform 2D reaches to visual targets, as shown in Figure 1a. Two monkeys, I and L, were prospectively selected for this study after having completed previous studies [37–39]. A third animal, J, was retrospectively selected for this study on account of having an observed V_{pp} decline during the course of a previous study [40] as well as substantial longitudinal data from an ongoing study [13]. For these three animals, we were interested in quantitatively comparing prosthetic performance over time. Thus, we analyzed data from stereotypical tasks, during each time such a task was run on a subset of days for more than 30 days.

These data consist of 382 daily datasets from 3 animals and 4 arrays across 3.6, 9.4, 10.4 and 31.7 months of recording, as shown in Figure 2. To our knowledge this represents the largest cortical dataset analyzed for array stability. Waveform voltage data could be obtained from all 382 datasets. A subset of this data, 184 datasets, included long periods of identical tasks. These periods are shown in the dashed boxes in Figure 2. During these periods, decoder performance could be quantitatively compared over time. Two animals had multiple periods where they performed a consistent task over more than 30 days. Point-to-point reaches were performed in complete darkness except for the illuminated target. Hand position was optically tracked with a reflective bead attached to the distal joint of the index or middle finger and measured at 60 samples/sec (nominal sub-millimeter resolution) using a Polaris system (NDI, Waterloo, Ontario, Canada). For the non-BMI datasets, the task was sequenced using Tempo software (Reflective Computing, St. Louis, MO). For the pre-BMI training trials, which consisted of real center-out reaches, the task was sequenced using custom Matlab and Simulink code executed on a real-time xPC target. All monkeys were highly trained for many months prior to implantation, such that there was likely little to no learning occurring during the present study. For all animals the task was a simple center-out reach task [41]. However, animals varied in the number of targets (4,7,8) and radius of the targets (80-120 cm), in addition to other small changes. However, performance measures for

a specific animal were internally consistent, allowing quantitative comparisons over time, but not between animals. The number of trials used was identical for each animal, and included at least 25 trials per condition. All animals received a liquid reward for successful trials.

2.2. Electrophysiology Recordings

One or two commercially available silicon “Utah” arrays (Blackrock Technologies Inc., Salt Lake City, UT, USA) containing 100 electrodes each were implanted into PMd/M1 of these animals using standard neurosurgical techniques [42]. Neural data were recorded using a Cerebus system (Blackrock Microsystems, Salt Lake City, UT) while the monkey was participating in a neuroscience experiment for 1-5 hours during the day. For monkey L only, wireless data were recorded using HermesC [37,43], pictured in Figure 1b. These data were recorded on 39 particular days while the animal was in the home cage from 20 channels at a time. Wireless data were obtained 91% of the time except during the in-rig experiments, animal transfers, and rare equipment failures. While this system has the capability of recording threshold-crossing data from all 20 channels simultaneously, only the single broadband channel was used in this study. The device was programmed to switch between each of 20 channels every 23 seconds and record data at 15.7 ksp/s such that waveform shape could be recorded 24 hours a day at a duty cycle of approximately 4%. Therefore, the wireless data were obtained at half of the sampling rate as the wired recordings and were only recorded for 23 seconds every 8 minutes.

Vpp from spike sorted neurons was required for many analyses presented in this paper. Due to the large size of the dataset, an automated spike sorting algorithm was required. Neural units were isolated offline using noise whitened principal components and a mixture of Gaussian models for clustering [30,44] as illustrated in Figure 1c. This spike sorting algorithm was augmented with a bipolar shape heuristic to eliminate any non-neural artifacts that might appear as sorted units. A large proportion of the data were manually verified to ensure that no artifacts were included. When well-isolated single units were required for an analysis, unit quality was rated manually for a small number of datasets. Units were classified as putative single units, single units with some contamination by other units, and multi-units. The number of putative single units in early datasets was 34, 27, and 28/38 for animals I, L, and J/J (Monkey J had 2 arrays).

For measuring offline decoder performance, as described below, data were not spike sorted. Instead, a threshold was set at $-4.5\times$ the RMS voltage on a particular electrode, and this multi-unit’s activity was decoded similarly to a single unit. Several prior studies have reported little or no gain in performance using sorted spikes rather than threshold crossings [11,30], as discussed above. This was verified using 39 datasets that were manually spike sorted, resulting in only a 4% performance improvement. Using threshold crossings also simplified standardization across animals and time periods, as data for monkeys L and J were collected as part of ongoing online BMI experiments, in which $-4.5\times$ RMS threshold crossings were used.

2.3. Neural Data Analysis

For estimating prosthetic decoder performance offline, two different decoders were used: a discrete and a continuous decoder. For the discrete decoder, a naive Bayes classifier was used, similar to one we have used previously [1]. The average firing rate of the threshold crossings on a particular electrode on a particular day was modeled as a Poisson distribution with a characteristic mean firing rate λ for each angle, measured with 500 ms of neural data after the target appears. For each block of single trials in each direction, all other reaches were used to train a model which was then tested on those trials. Decoders were trained and

tested on data within the same day. To decode the target angle, the likelihood of each angle was calculated as in Equation 1:

$$P(\Theta=\theta|y_t) = \sum_{n=1}^N \frac{\lambda_{n,\theta}^{y_t} e^{-\lambda_{n,\theta}}}{y_t!} \quad (1)$$

where θ is target angle, N is the number of neurons, and y_t is the number of spikes that occurred during the integration window. The angle with the highest probability was then selected. Using this type of decoder, performance can be evaluated with percent correct.

A continuous decoder was also constructed offline using a linear filter [3, 6, 7, 33]. The firing rate and position was averaged over 100 ms bins during each trial. The linear model used the firing rates of the units at ten sequential 100 ms time lags. Position was modeled as a function of firing rate using a linear Wiener filter. In Equations 2 and 3, the firing rate matrix X has ten columns for each unit, and horizontal and vertical position is found in matrix Y .

$$Y=XB \quad (2)$$

$$\begin{bmatrix} pos(1) \\ \dots \\ pos(t) \end{bmatrix} = \begin{bmatrix} neur_1lag_1(1) & neur_1lag_2(1) & \dots & neur_Nlag_{10}(1) \\ \dots & \dots & \dots & \dots \\ neur_1lag_1(t) & \dots & \dots & neur_Nlag_{10}(t) \end{bmatrix} \begin{bmatrix} B_{neur_1lag_1} \\ \dots \\ B_{neur_Nlag_{10}} \end{bmatrix} \quad (3)$$

Both X and Y have rows corresponding to the total number of 100 ms bins in the experiment. The resulting linear decoder, matrix B , is computed through linear regression. The model was tested using cross validation with 20 folds. I.e., for each 5% of the data, neural decoders were generated using the other 95% of the data, and tested on that 5%. The final performance value was averaged across all such 5% pieces. Continuous performance was quantified using a correlation coefficient with actual hand position, as well as the average distance to the target during the decoded reach.

3. Results

3.1. Long Timescales

The simplest way to characterize electrode stability across many months is to examine the peak to peak amplitudes (Vpp) of the largest action potentials on each channel. Figure 3 shows this Vpp averaged across electrodes (with high noise electrodes removed) for 4 arrays in 3 animals. Three out of four arrays showed a statistically significant negative trend during this time ($p < 0.001$). The size of the effect was, +0.3% (I, not significant), -6.1% (J-1), -2.8% (J-2), and -0.6% (L) per month. The arrays with fewer months of recording show larger effects, which is consistent with a trend that slows down over time.

This data is characterized by long periods with a small negative change and shorter periods with large changes, which may be relevant to determining the mechanism of these changes. For example, in Monkey J, both arrays showed decline during the first 2 months after implantation. This can be seen more easily by viewing only the channels with observed isolated action potentials on day one, shown in Figure 4a, rather than all channels as in Figure 3. The overall Vpp declined by 47.0% and 28.2% ($p < 0.001$). Subjectively, during experiments well-isolated action potentials were observed to decline substantially in number. However, in one of these arrays that continued to be used for online BMI experiments over the following year it is clear that this trend did not continue at the same

rate. For the 8.5 months following, voltage declined by only 2.0% per month. Another period of decline was observed at the end of 27 months (2.2 years) of recording from Monkey L. During the initial 27 months, this array's V_{pp} declined by 17.2% (0.6% per month). However, during the following 2.5 months, as shown in Figure 4b, it declined by 35% (and then recovered by 15%). This event was concurrent with an infection at the implant site, but there is no obvious mechanism to explain the changes in voltage. In all animals, V_{pp} declined slower than what would be expected if the glial scar response sealed off the electrode in the first 12 weeks.

To determine how the voltage changes might affect decoder performance, offline performance was evaluated in two ways. First, a discrete prediction of which target the animal was reaching to provided a percent correct value. Second, an offline implementation of a continuous linear decoder mapped neural firing rates to hand position. Quality of the continuous decode was measured by correlation coefficient between predicted and actual hand position. However, correlation coefficient can remain stable in the face of large absolute changes in decoder performance. Therefore, a second metric, mean distance to target, was also calculated during the decoded reach as in [45,46].

For three out of four animals, decoder performance was not significantly correlated with V_{pp} . Figure 5 shows all three performance metrics below the voltage amplitude traces in Monkeys I and J. Performance variation was not significantly correlated with voltage variation in these arrays ($p > 0.05$ except one $p = 0.02$, which is not significant when corrected for multiple comparisons). In the remaining data from the second array in monkey J, it continued to be uncorrelated with voltage over months 4-11, during which a different task with more targets was run and decoded offline (23 datasets). Performance by itself during these periods showed no significant negative trends, despite the negative trends in voltage during the first two months for both arrays in Monkey J. These data show that action potential amplitude cannot be used as a proxy for measuring BMI performance.

Similar data for monkey L is shown in Figure 6. This array is older than the others, having been implanted for over 2.5 years. There were two periods with identical tasks across periods longer than 30 days, which are graphed on either side of the ellipses near day 350 in Figure 6, such that absolute performance cannot be compared across this boundary, but trends can be examined within each period. During the first period, variation in performance was correlated with variation in voltage for all three measures ($\rho = 0.45$, $p < 0.01$). However, there were no significant trends in performance during this time. During the second period, performance using all three metrics was highly correlated with voltage ($\rho = 0.67$, $p < 0.001$) and decreased along with the voltage ($p < 0.001$). During the final 5 days of recording, performance had nearly recovered, with an average discrete percent correct of 93% predicting which of 8 targets the animal was reaching towards and an average correlation coefficient with hand position of 0.85. This is despite the fact that V_{pp} had only recovered from approximately half of the loss.

3.2. Short Timescales

In addition to long term trends, short-term variation in neural waveforms was observed over the course of hours and between days. Figure 7 shows data from the wireless incage recordings. The example electrode in Figure 7A had single unit action potentials that were sometimes visible and sometimes undifferentiable from the multi-unit hash. These short term changes occur despite the fairly stable average unit voltage shown in Figures 3 through 6. Voltage data for 19 channels of wireless recording are shown in Figure 7B. For each electrode, the 95% envelope of voltage waveform snippets was found and normalized to the average voltage on that electrode across all days so that percentage deviations could be

examined. Peak to peak voltages ranged from 66% to 236% of these mean values. The speed of changes is apparent from the abrupt changes between adjacent 1 hour bins.

During experiments in which the animal sat in a primate chair with a fixed head position, smaller changes in neural waveforms were also observed. Figure 8 shows normalized voltage data from hand verified single units in one dataset per animal. Over the course of the experiment, the Vpp's progressively deviate from their initial value, as shown by the standard deviation traces. The final standard deviation from the initial values was 9.0%, 6.4%, and 14.8% over 1, 2, and 5 hours.

For all experimental days and animals, the data were spike sorted in an automated fashion, and the Vpp of the largest sorted unit on each channel was calculated. Changes in these waveform amplitudes were measured within single experimental sessions, as well as across many days. A histogram of changes that occurred (per hour) over 1-5 hrs while the animals were seated in the rig is shown in Figure 9a. With 21310 unit-day pairs, this distribution shows an average absolute change of 3.2% per hour, though 6% of the waveforms changed by over 10%. Across days, changes were more substantial with an average absolute change from the mean voltage of 23.7%, as shown in Figure 9b. Again, there were outlying changes that could be very large with 14.6% of the waveforms changing by over 50% between two given days. These results are consistent with previous studies in which large changes were observed between days [32, 36]. The wireless data above provides the fast timecourse of these changes.

4. Discussion

Action potential amplitude (Vpp) declines reported in this study were slow, with an average decline of 2.4% per month. Older implants had slower rates of decline, consistent with a slowing trend. These data are consistent with a prior study [36], which described no trend presumably because these changes are rare, slow, and subtle enough to require many months of data. Also, no significant trend is visible using all sorted units rather than only the largest units on each electrode. For monkey J, these data may also be consistent with tissue response studies in which there is a period of rapid decline in the amplitudes of single units during the first 12 weeks [21, 23, 27]. To our knowledge this is the first study to correlate decoder performance with array voltage over multi-month timescales, though preliminary work has appeared previously [47, 48]. Vpp was not well correlated with performance when using threshold crossings, which suggests that the presence and amplitude of single units cannot be used as a measure of BMI performance. While the performance discussed here was assessed with offline decodes, Monkeys J and L were simultaneously participating in online BMI experiments. During this time, no noticeable decline in online decoder performance was noted subjectively, except during the period shown in Figure 6, where offline performance declined as well. While offline decodes may not be a good predictor of absolute online performance [45], they may still be useful for estimating the information content over time. Also, the arrays presented here were representative of the arrays observed in our laboratory.

In this study, the performance was analyzed while retraining the decoder every day. This improves performance because of the short term waveform changes shown in Figures 7 through 9, which lead to a changing subset of neurons comprising the multi-unit. It has been suggested that this may be an obstacle for clinical systems. During online monkey experiments, this requires approximately 15 min of training data and is fully automated in software. No manual supervision of spike sorting is required when using threshold crossings. High performance has also been demonstrated with fewer training trials [10]. Other signals sources such as ECoG have demonstrated stable performance using the same

model parameters [49,50] for up to 8 months [51], but this may reflect a tradeoff between performance and stability. It might be possible that after some number of years, ECoG signals would outperform intracortical array signals on the same task, but that remains unconfirmed. However, with micro-ECoG implants, it might be possible to increase the amount of cortical coverage beyond what is possible using intracortical arrays for applications such as finger decoding [52–54].

To improve intracortical arrays so that they can last for decades, it is important to characterize the type and speed of signal changes. On short timescales, reversible changes in action potential amplitude can occur very quickly. Motion is likely responsible for some of these changes [55]. However, their speed would suggest movements of many tens of μm in a homogenous medium [56]. This does not include non-linear effects, such as touching a neuron. If there are substantial movements of the array, this underscores the importance of the glial sheath, which enables an electrode to move with less damage to surrounding neural tissue [21]. Changes in electrode or tissue impedance could also result in waveform changes [24]. No matter what the cause, these changes make it very difficult to maintain single units over time.

Fortunately, spike sorting does not substantially improve decoder performance [11]. This could be because only a subset of channels have multiple units, and only a subset of those have opposed directional tuning [30]. On average, combining units may tend to decrease the Poisson noise. Also, human spike sorters have been shown to ignore units containing useful information on account of poor isolation [29]. In this study, threshold crossing events were detected whenever voltage went below $-4.5 \times$ the RMS voltage on that particular channel. In addition to eliminating the need for supervised spike sorting, this approach enabled online BMI experiments to be conducted with older arrays than previous work [13].

Neural data from all arrays in this study were recorded for more than 12 weeks, at which time the glial sheath has probably already formed. The slow long-term decline in action potential voltage across many months may reflect consolidation of this process, neuron migration away from the site, or accumulating adverse events such as head accelerations or infections. Importantly, V_{pp} declines on the timescales of months and years may have as much or more to do with materials and engineering failures, and future research is needed to assess this contribution. In this study, large changes in voltage were observed without concurrent changes in performance when using threshold crossings. A number of effects, for example gliosis or loss of electrode insulation, could cause action potential amplitude to decrease. However, if there was no cell migration away from the electrode, one would still expect the same population of neurons to be recorded at a lower voltage. Therefore, one would expect performance using threshold crossings to remain constant, as described in the Results. While single units might no longer be clearly differentiable, the remaining activity on these channels would still be primarily neural, and probably above the inherent noise of the electrode and amplifiers. If the voltage declined into this noise floor or an alternative process became dominant, the performance would eventually correlate with voltage, similar to the effect observed in Monkey L at 27 months post implant. Of course, while this description is consistent with data presented here, only a small number of arrays were examined. Future work examining multi-year old arrays may be warranted. This study also predicts that one could simulate older arrays by using threshold crossings even while single units are present. Wireless systems may be helpful in running array lifetime studies concurrently with conventional experiments [38]. Additionally, fully implantable devices may become crucial for long term device stability [43,57,58]. In this and previous studies, array stability at multiple years is likely confounded by the ability to maintain a transcutaneous link.

This study suggests that the average action potential amplitude on electrode arrays may decay more slowly than previously supposed and that decoder performance can decay at an even slower rate. This is enabled by the use of simple threshold crossings events rather than well isolated action potentials from single neurons. This further enables and motivates the translation of cortical BMI work into quality of life improvements for disabled patients.

Acknowledgments

We thank M. Risch and S. Kang for expert surgical assistance and veterinary care, D. Haven for technical consultation, and S. Eisensee for administrative support.

This work was supported by National Science Foundation graduate research fellowships (C.A.C., M.T.K., J.M.F., V.G.), NDESG Fellowship (V.G.), Burroughs Wellcome Fund Career Awards in the Biomedical Sciences (M.M.C., K.V.S.), the Christopher Reeve Paralysis Foundation (S.I.R., K.V.S.), a Stanford University Graduate Fellowship (J.P.C., Z.A.R., C.A.C., J.M.F., J.D.F.), a CONACYT and DARE fellowship (Z.A.R.), a Stanford NIH Medical Scientist Training Program grant and Soros Fellowship (P.N.), and the following awards to K. V. S.: Stanford CIS, Sloan Foundation, NIH CRCNS R01-NS054283, McKnight Endowment Fund for Neuroscience, NIH Director's Pioneer Award 1DP1OD006409, DARPA Revolutionizing Prosthetics program contract N66001-06-C-8005, and DARPA REPAIR contract N66001-10-C-2010.

References

- [1]. Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. A high-performance brain-computer interface. *Nature*. 2006; 442:195–198. [PubMed: 16838020]
- [2]. Taylor DM, Tillery S.I. Helms, Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. *Science*. 2002; 296:1829–1832. [PubMed: 12052948]
- [3]. Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP. Instant neural control of a movement signal. *Nature*. 2002; 416:141–142. [PubMed: 11894084]
- [4]. Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, Patil PG, Henriquez CS, Nicolelis MAL. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*. 2003; 1:193–208.
- [5]. Lebedev MA, Carmena JM, O'Doherty JE, Zacksenhouse M, Henriquez CS, Principe JC, Nicolelis MAL. Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *J Neurosci*. 2005; 25(19):4681–4693. [PubMed: 15888644]
- [6]. Carmena JM, Lebedev MA, Henriquez CS, Nicolelis MA. Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci*. 2005; 25(46):10712–10716. [PubMed: 16291944]
- [7]. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*. 2006; 442:164–171. [PubMed: 16838014]
- [8]. Kim S, Simeral JD, Hochberg LR, Donoghue JP, Black MJ. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J Neural Engineering*. 2008; 5:455–476.
- [9]. Mulliken GH, Musallam S, Andersen RA. Decoding trajectories from posterior parietal cortex ensembles. *J Neurosci*. 2008; 28:12913–12926. [PubMed: 19036985]
- [10]. Chase SM, Schwartz AB, Kass RE. Bias, optimal linear estimation, and the differences between open-loop simulation and closed loop performance of spiking based brain computer interface algorithms. *Neural Networks*. 2009; 22(9):1203–1213. [PubMed: 19502004]
- [11]. Fraser GW, Chase SM, Whitford A, Schwartz AB. Control of a brain-computer interface without spike sorting. *J Neural Engineering*. 2009; 6:1–8.
- [12]. Ganguly K, Carmena JM. Emergence of a stable cortical map for neuroprosthetics. *PLoS Biology*. 2009; 7(7):1–13.
- [13]. Gilja V, Nuyujukian P, Chestek C, Cunningham J, Yu B, Ryu S, Shenoy K. High-performance continuous neural cursor control enabled by a feedback control perspective. *Conference Abstract: Computational and Systems Neuroscience*. 2010

- [14]. Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Ryu SI, Shenoy KV. A high-performance continuous cortically controlled prosthesis enabled by feedback control design. *Soc. for Neurosci. Abstr.* 2010; 20.6
- [15]. Nuyujukian P, Gilja V, Chestek CA, Cunningham JP, Fan JM, Yu BM, Ryu SI, Shenoy KV. Generalization and robustness of a continuous cortically-controlled prosthesis enabled by feedback control design. *Soc. for Neurosci. Abstr.* 2010; 20.7
- [16]. Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. *Nature.* 2008; 453:1098–1101. [PubMed: 18509337]
- [17]. Vargas-Irwin CE, Shakhnarovich G, Yadollahpour P, Mislow JMK, Black MJ, Donoghue JP. Decoding complete reach and grasp actions from local primary motor cortex populations. *J Neurosci.* 2010; 30(29):9659–9669. [PubMed: 20660249]
- [18]. Kuiken TA, Li G, Lock BA, Lipschutz RD, Miller LA, Stubblefield KA, Englehart KB. Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms. *J of American Medical Assoc.* 2009; 301(6):619–628.
- [19]. Schmidt S, Horch K, Normann R. Biocompatibility of silicon-based electrode arrays implanted in feline cortical tissue. *J of Biomedical Materials Research.* 1993; 27(11):1393–1399.
- [20]. Stensaas SS, Stensaas LJ. Histopathological evaluation of materials implanted in the cerebral cortex. *Acta Neuropathologica.* 1978; 41:145–155. [PubMed: 636844]
- [21]. Turner JN, Shain W, Szarowski DH, Andersen M, Martins S, Isaacson M, Craighead H. Cerebral astrocyte reponse to micromachined silicon implants. *Experimental Neurology.* 1999; 155(1):33–49. [PubMed: 10192775]
- [22]. Edell DJ, Toi VV, McNeal VM, Clark LD. Factors influencing the biocompatibility of insertable silicon microshafts in cerebral cortex. *IEEE Transactions on Biomedical Engineering.* 1992; 39(6):635–643. [PubMed: 1601445]
- [23]. Williams JC, Rennaker RL, Kipke DR. Long-term neural recording characteristics of wire microelectrode arrays implanted in cerebral cortex. *Brain Research: Brain Research Protocols.* 1999; 4:303–313. [PubMed: 10592339]
- [24]. Winslow BD, Tresco PA. Quantitative analysis of the tissue response to chronically implanted microwire electrodes in rat cortex. *Biomaterials.* 2010; 31:1558–1567. [PubMed: 19963267]
- [25]. Griffith RW, Humphrey DR. Long-term gliosis around chronically implanted platinum electrodes in the rhesus macaque motor cortex. *Neuroscience Letters.* 2006; 406:81–86. [PubMed: 16905255]
- [26]. Polikov VS, Tresco PA, Reichert WM. Response of brain tissue to chronically implanted neural electrodes. *J of Neurosci Methods.* 2005; 148:1–18. [PubMed: 16198003]
- [27]. Szarowski DH, Andersen MD, Spence AJ, Isaacson M, Craighead HG, Turner JN, Shain W. Brain responses to micro-machined silicon devices. *Brain Research.* 2003; 983(1-2):23–35. [PubMed: 12914963]
- [28]. Truccolo W, Friehs GM, Donoghue JP, Hochberg LR. Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *J Neurosci.* 2008; 28(5):1163–1178. [PubMed: 18234894]
- [29]. Wood F, Fellows M, Donoghue J, Black MJ. Automatic spike sorting for neural decoding. *IEEE EMBS.* 2004:4009–4012.
- [30]. Santhanam, G.; Sahani, M.; Ryu, SI.; Shenoy, KV. An extensible infrastructure for fully automated spike sorting during online experiments. *Proceedings of the 26th Annual International Conference of the IEEE EMBS; San Francisco, CA.* 2004. p. 4380-4384.
- [31]. Stark E, Abeles M. Predicting movement from multiunit activity. *J Neurosci.* 2007; 27(31):8387–8394. [PubMed: 17670985]
- [32]. Parker-Ure RA, Christensen MB, House P, Tresco PA, Greger B. Micro-electrode functionality inversely correlated with markers of inflammation and not with markers of neuronal cell viability. *Soc. for Neurosci. Abstr.* Nov 18.2008
- [33]. Chestek* CA, Batista* AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV. Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci.* 2007; 27:10742–10750. [PubMed: 17913908]

- [34]. Dickey AS, Suminski A, Amit Y, Hatsopoulos HG. Single-unit stability using chronically implanted multielectrode arrays. *J Neurophysiol.* 2009; 102:1331–1339. [PubMed: 19535480]
- [35]. Serruya M, Hatsopoulos NG, Fellows MR, Paninski L, Donoghue JP. Robustness of neuroprosthetic decoding algorithms. *Biological Cybernetics.* 2003; 88:219–228. [PubMed: 12647229]
- [36]. Suner S, Fellows MR, Vargas-Irwin C, Nakata GK, Donoghue JP. Reliability of signals from a chronically implanted silicon-based electrode array in non-human primate primary motor cortex. *IEEE Trans. Neural Sys. Rehabil. Eng.* 2005:524–541.
- [37]. Chestek* CA, Gilja* V, Nuyujukian P, Kier R, Solzbacher F, Ryu SI, Harrison RR, Shenoy KV. HermesC: Low-power wireless neural recording system for freely moving primates. *IEEE Trans. in Neural. Sys. and Rehabil. Eng.* 2009; 17:330–338. (In Press).
- [38]. Gilja* V, Chestek* C, Nuyujukian P, Foster J, Shenoy KV. Autonomous head-mounted electrophysiology systems for freely behaving primates. *Curr Opin Neurobiol.* 2010 (In press).
- [39]. Rivera-Alvidrez Z, Kalmar R, Ryu SI, Shenoy KV. Low-dimensional neural features predict muscle emg signals. *IEEE EMBS.* 2010 (In press).
- [40]. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron.* 2010 (In press).
- [41]. Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci.* 1982; 2(11):1527–1537. [PubMed: 7143039]
- [42]. Hatsopoulos N, Joshi J, O’Leary JG. Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *J Neurophysiol.* 2004; 92:1165–1174. [PubMed: 15277601]
- [43]. Harrison RR, Kier RJ, Chestek CA, Gilja V, Nuyujukian P, Ryu S, Greger B, Solzbacher F, Shenoy KV. Wireless neural recording with single low-power integrated circuit. *IEEE Trans. Neural Sys. Rehabil. Eng.* 2009; 17:322–329.
- [44]. Sahani, M. Ph.D. Thesis, Computational and Neural Systems. California Institute of Technology; 1999. Latent variable models for neural data analysis.
- [45]. Cunningham JP, Nuyujukian P, Gilja V, Chestek CA, Ryu SI, Shenoy KV. A closed-loop human simulator for investigating the role of feedback-control in brain-machine interfaces. *J Neurophysiol.* (In Press).
- [46]. Cunningham JP, Nuyujukian P, Gilja V, Chestek CA, Ryu SI, Shenoy KV. A closed-loop human simulator for understanding the role of feedback-control and its relevance for brain-machine interfaces. *Soc. for Neurosci. Abstr.* 2010; 20.5
- [47]. Chestek* CA, Cunningham* JP, Gilja V, Nuyujukian P, Ryu SI, Shenoy KV. Neural prosthetic systems: Current problems and future directions. *Proc IEEE EMBS.* 2009:3369–3375.
- [48]. Chestek CA, Gilja V, Nuyujukian P, Foster JD, Kaufman MT, Ryu SI, Shenoy KV. Impact of spike waveform instability on performance measured across seven weeks using a wireless system. *Soc. for Neurosci. Abstr.* 2010; 85.1
- [49]. Blakely T, Miller KJ, Zanos SP, Rao R, Ojemann JG. Robust long-term control of an electrocorticographic brain-computer interface with fixed parameters. *Neurosurgical Focus.* 2009; 27(1):1–5.
- [50]. Schalk G, Miller KJ, Anderson NR, Wilson JA, Smyth MD, Ojemann JG, Moran DW, Wolpaw JR, Leuthardt EC. Two-dimensional movement control using electrocorticograph signals in humans. *J Neural Engineering.* 2008; 5(1):75–84.
- [51]. Chao ZC, Nagasaka Y, Fujii N. Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. *Frontiers in Neuroengineering.* 2010; 3(3):1–10. [PubMed: 20162033]
- [52]. Kubanek J, Miller KJ, Ojemann JG, Wolpaw JR, Schalk G. Decoding flexion of individual fingers using electrocorticographic signals in humans. *J Neural Engineering.* 2009; 6:1–14.
- [53]. Miller KJ, Zanos S, Fetz EE, Nijs M, Ojemann JG. Decoupling the cortical power spectrum reveals real-time representation of individual finger movements in humans. *J Neurosci.* 2009; 29(10):3132–3137. [PubMed: 19279250]

- [54]. Acharya S, Fifer MS, Benz HL, Crone NE, Thakor NV. Electrographic amplitude predicts finger positions during slow grasping motions of the hand. *J Neural Engineering*. 2010; 7:1–13.
- [55]. Santhanam G, Linderman MD, Gilja V, Afshar A, Ryu SI, Meng TH, Shenoy KV. Hermesb: A continuous neural recording system for freely behaving primates. *IEEE Trans. in Biomed. Eng.* 2007; 54:2037–2050.
- [56]. Moffitt MA, McIntyre CC. Model-based analysis of cortical recording with silicon microelectrodes. *Clinical Neurophysiology*. 2005; 116:2240–2250. [PubMed: 16055377]
- [57]. Kim S, Bhandari R, Klein M, Negi S, Rieth L, Tathireddy P, Toepper M, Oppermann H, Solzbacher F. Integrated wireless neural interface based on the Utah electrode array. *Biomedical Microdevices*. 2009; 11(2):453–466. [PubMed: 19067174]
- [58]. Borton DA, Song YK, Patterson WR, Bull CW, Park S, Laiwalla F, Donoghue JP, Nurmikko AV. Wireless, high-bandwidth recordings from non-human primate motor cortex using a scalable 16-ch implantable microsystem. *IEEE EMBS*. 2009:5531–5534.

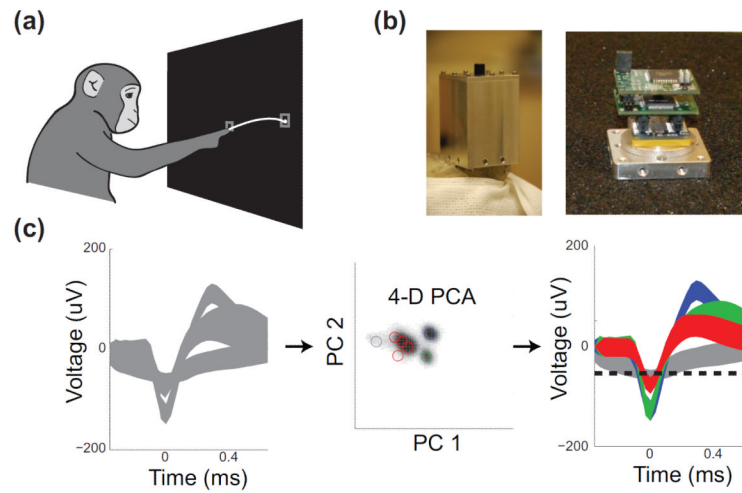


Figure 1.

(a) Illustration of experimental setup. Animals made 2D reaches to radial targets. (b) Wireless recording device used for Monkey L [37] (c) Neural data was processed through various filters, and through a principal component based automatic spike sorter [30]. Data were first filtered, peak-aligned and noise whitened in the leftmost panel. Then, waveforms were projected into 4D principle component space, where units were clustered using an expectation-maximization based mixture of gaussian models, producing the classification shown on the right.




Figure 2. Chart showing when datasets were collected. Tick marks represent 1 day of recording. Gray line represents time each animal was implanted. Dashed boxes indicate times where a consistent center out reach task was performed over more than 30 days such that offline BMI performance could be analyzed. Color and number denotes different but self-consistent tasks. Period with wireless data denoted by black arrow.

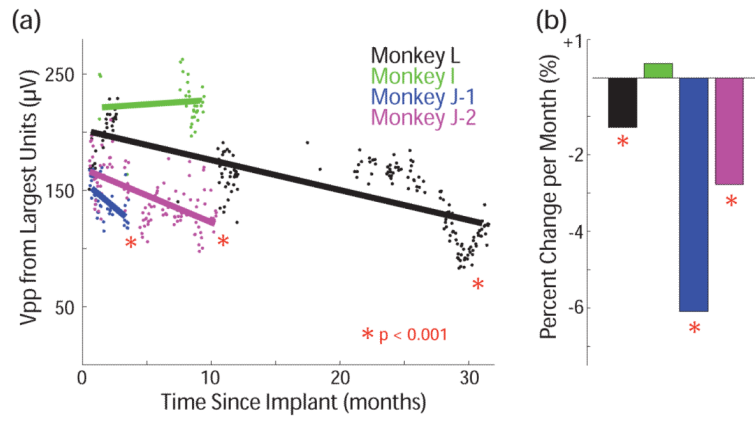


Figure 3. (a) Voltage from largest unit on each channel averaged across channels for 4 arrays in 3 animals, where color denotes a specific array. (b) Average change per month in each array.

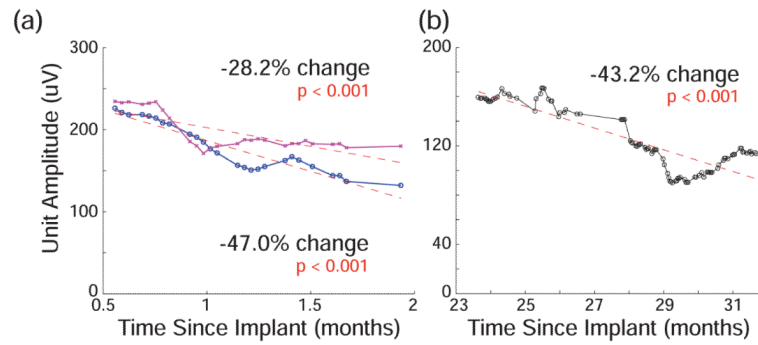


Figure 4.

(a) Voltage from largest unit averaged across the 28 and 38 channels on two arrays from Monkey J that had observed single unit activity in initial dataset during first 2 months of recording (b) Similar data from Monkey L after 2 years of recording. All data is smoothed across a 10 day averaging window to reveal trends.

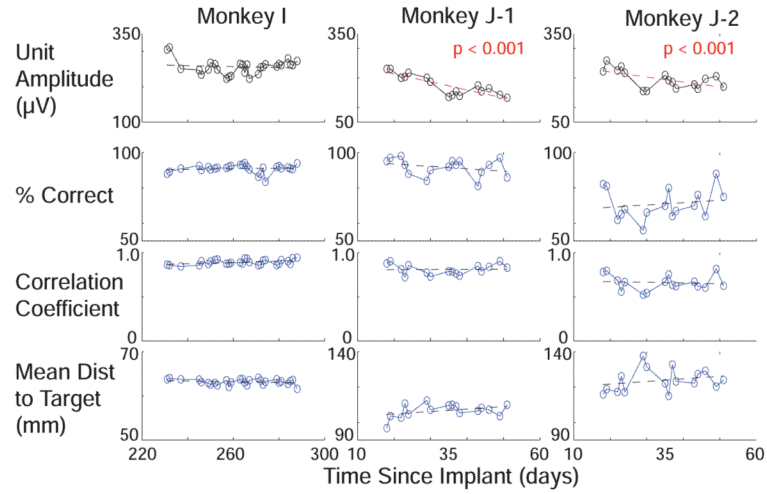


Figure 5.

Top row shows average voltage amplitude across electrodes with observed single unit activity on day one on 3 arrays in 2 animals during periods where the same task was performed for more than 20 days. The number of single unit channels was 34, 28, and 38 electrodes for I, J-A, and J-B. Lower rows show offline decoder performance. For discrete target decoders, performance is measured by percent correct. For a continuous linear decode, performance is measured with correlation coefficient to actual hand position and mean distance to target over the course of the trial. Red lines denote significant trends. No smoothing filters were used.

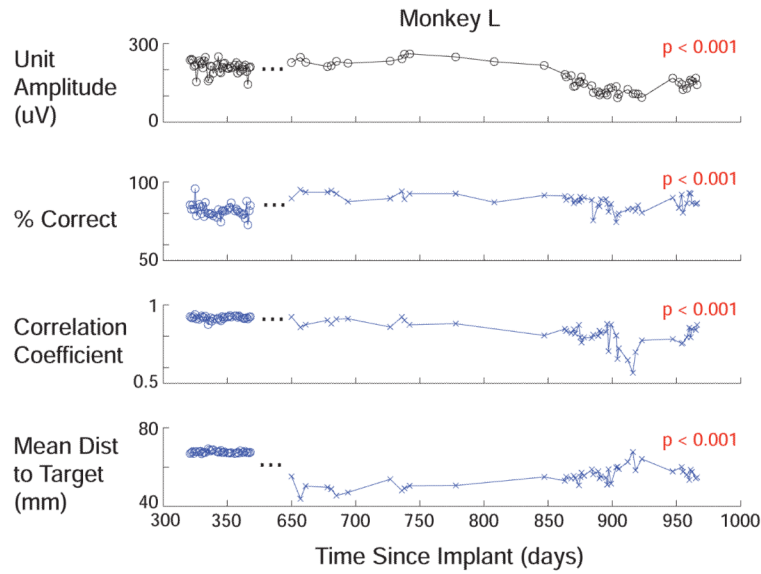


Figure 6.

Top row shows voltage amplitude from largest unit on each electrode averaged across electrodes for Monkey L. Ellipses denotes discontinuity on x axis near month 12. The performance data to the left and right of these ellipses come from different tasks, denoted by symbols, with different average performance, but are self-consistent on either side. Data is more sparse at the beginning of the second period because various training paradigms were attempted, and only a subset of the days included an identical 8 cm center out task with more than 200 trials, though this became standard later on. The same performance metrics as Figure 5 are shown below. On the first task, no significant trends are present. On the second task, voltage declines and recovers significantly ($p < 0.001$), and is significantly correlated with changes in performance using all 3 metrics ($p < 0.001$). No smoothing filters were used.




Figure 7.

(a) Example waveforms across days from wireless dataset. Unit is regularly visible and occasionally disappears. (b) Average voltage in 1 hour bins calculated on 19 of 20 active channels across 13 days of wireless recording starting 328 days after implantation. Voltage is normalized to mean voltage on that channel. White space denotes time without wireless recording, usually during daily experimental session. Arrows denote the channel which was plotted in (a).




Figure 8.

Changes in single unit amplitude during normal experiments while seated with a fixed head position. Voltage normalized to the size of the single unit at the beginning of the experiment to show percentage change in three animals. Red lines denote the standard deviation from zero.

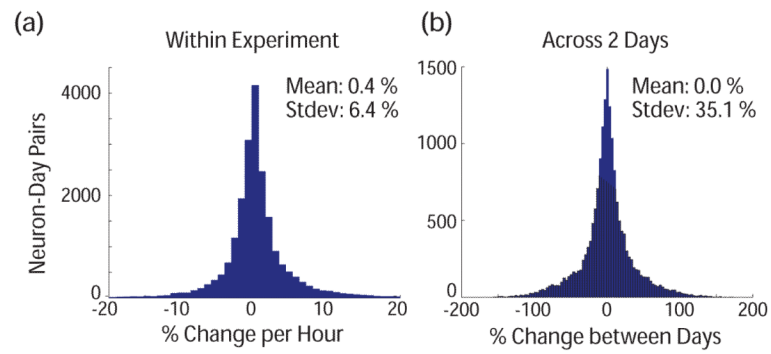


Figure 9. Distribution of percent changes in waveform voltage amplitude from largest unit on an electrode (a) Per hour during 1 - 5 hour experiments (b) Between 2 experimental days on the same electrode.