

Neural Recording Stability of Chronic Electrode Arrays in Freely Behaving Primates

Michael D. Linderman*, Vikash Gilja[†], Gopal Santhanam*, Afsheen Afshar*[§],
Stephen Ryu*[¶], Teresa H. Meng*, Krishna V. Shenoy*[†]

*Department of Electrical Engineering [†]Department of Computer Science

[§]School of Medicine [¶]Department of Neurosurgery [†]Neurosciences Program
Stanford University, Stanford, California, USA

Abstract—Chronically implanted electrode arrays have enabled a broad range of advances, particularly in the field of neural prosthetics. Those successes motivate development of prototype implantable prosthetic processors for long duration, continuous use in freely behaving subjects. However, traditional experimental protocols have provided limited information regarding the stability of the electrode arrays and their neural recordings. In this paper we present preliminary results derived from long duration neural recordings in a freely behaving primate which show variations in action potential shape and RMS noise across a range of time scales. These preliminary results suggest that spike sorting algorithms can no longer assume stable neural signals and will need to transition to adaptive signal processing methodologies to maximize performance.

I. INTRODUCTION

Chronically implanted electrode arrays have enabled a broad range of advances, particularly in the field of neural prosthetics. Those successes motivate development of prototype implantable prosthetic processors for long duration, continuous use in freely behaving subjects. However, the nature of current experimental protocols limit both the types and duration of recordings. As a result there is limited data available with which to characterize the stability of neural recordings over the broader range of timescales relevant to an implantable prosthetic processor.

Figure 1 summarizes the significant timescales in the life of a chronically implanted electrode array. In this paper we are concerned only with neural recording stability in the high-yield recording period during which most experiments are conducted [1]. Within this window, a continuously operating prosthetic system is potentially affected by recording instability at all three timescales (short, intermediate and long). However, current experiments, with their discrete daily recording periods, are only able to characterize variations on timescales less than a few hours and across days. In this study, we investigate variations at intermediate timescales (i.e., between daily recording sessions).

This work was supported in part by MARCO Center for Circuit & System Solutions (T.H.M.,M.D.L.), NDSEG (M.D.L.,V.G.,G.S.) and NSF (V.G.,G.S.) fellowships, Bio-X Fellowship (A.A.), Christopher Reeve Paralysis Foundation (S.I.R.,K.V.S) and the following awards to K.V.S.: NSF Center for Neuromorphic Systems Engineering at Caltech, ONR Adaptive Neural Systems, Whitaker Foundation, Center for Integrated Systems at Stanford, Sloan Foundation, and Burroughs Wellcome Fund Career Award in the Biomedical Sciences. Please address correspondence to mlinderm@stanford.edu.




Fig. 1. Summary of array lifetime and available data for recording from individual, identifiable neurons using a chronically implanted electrode array.

Characterizing the stability, or lack thereof, in neural recordings at intermediate timescale enables the principled design of spike sorting algorithms for continuous, long duration use in freely behaving subjects and provides the bridge between traditional daily recording sessions required to certify neuron identity in multi-day learning and plasticity experiments. The HermesB system, a new autonomous, long-duration neural recording system for freely behaving non-human primates, enables collection of the previously unavailable multi-day broadband datasets needed for this characterization [2].

Studies have characterized neural signal stability on both short (seconds or minutes) [3], [4] and long timescales (days) [4], [5]. Over very short time scales [3] observed variations in action potential waveform shape are a function of the interspike interval (ISI); at short ISI the waveform is typically longer (in time) and decreased in amplitude. At larger time scales, the variation in spike waveform is not as systematic, potentially arising from a number of mechanisms such as neural plasticity and physical movement in the electrode relative to the neuron [6].

We present preliminary results here quantifying the stability of neural recordings over time scales of 5 min – 54 hours. In particular we address three aspects of recording stability identified in [6]: the change in mean waveform shape over time, changes in the background noise process and changes in waveform shape due to electrode array movement. Initial

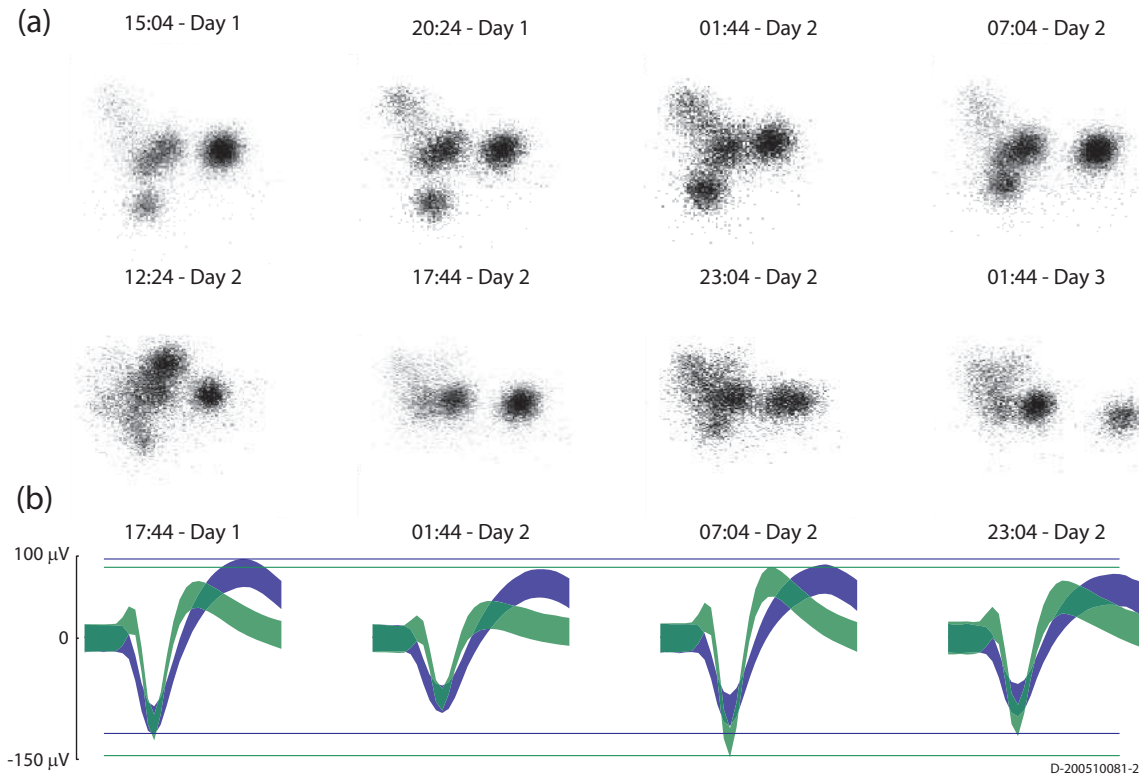


Fig. 2. Neural recordings over a period of 48 hours. (a) Histogram of spike waveform projections into a fixed NWrPCA space. PCA space determined using 20000 snippets uniformly selected across the time period. Each plot is the projection of 5 minutes of data recorded from a signal channel at the time shown (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.).

analysis shows significant variation in waveform amplitude at intermediate timescales, along with similar variation in RMS noise over the same period. In rare cases the observed variation appears to result from abrupt shifts in electrode position caused by head movement.

II. METHODS

We recorded single and multi-unit signals from a 96 channel electrode array (Cyberkinetics Neurotechnology Systems, Inc.) chronically implanted in August 2005 in the dorsal premotor cortex of an adult macaque monkey. Surgical methods are described in [7], [8]. All experiments and procedures were approved by the Stanford University Institutional Animal Care and Use Committee (IACUC). Recordings were made with an autonomous, self-contained recording setup, the HermesB system, described in [2], [9], while the monkey was freely behaving in the home cage. Each recording was 54 hours in length at a 67% duty cycle (5 minutes of recording, followed by 2.5 minutes of system sleep). Two neural channels were recorded per data set in full broadband (.5 Hz to 7.5 kHz at 12 bits at 30 kSamples/s) and a 3-axis accelerometer fixed to the monkey's head was recorded at 12 bits at 1 kSamples/s. The start date of each recording is shown in the small tag on the figure (format D-YYYYMMDD*-*).

The recorded neural signals were post-processed using the

Sahani algorithm for spike sorting [10], [11]. The spike time is identified using a threshold determined from the training data (3σ with respect to the RMS noise of filtered data). A spike waveform, or snippet, comprising a 32 sample window around the threshold event, is extracted and aligned to its center-of-mass (COM). Snippets are projected into a 4-dimensional robust, noise-whitened principle components space (NWrPCA) and clustered using a maximum *a posteriori* (MAP) clustering technique. Quantitative results are presented in the context of sort results using this algorithm.

III. CHARACTERIZING CHANGES IN WAVEFORM SHAPE

Preliminary results from long duration neural recordings of a freely behaving primate indicate significant variation in spike waveforms over intermediate timescales. Figure 2 shows neural recordings made over the course of 48 hours in October 2005. The upper two rows, Fig. 2a, show a time series of NWrPCA cluster plots for 5 minute data segments recorded at the times shown. Each cluster corresponds to a single neuron, and the movement (drift) of the relative distance between these clusters is readily seen by scanning across the snapshots.

The drift of the clusters in NWrPCA space is caused by changes in spike waveform shape. Figure 2b shows action potential shapes (voltage vs. time) from the same recording period. The colored region indicates the 10th–90th percentile

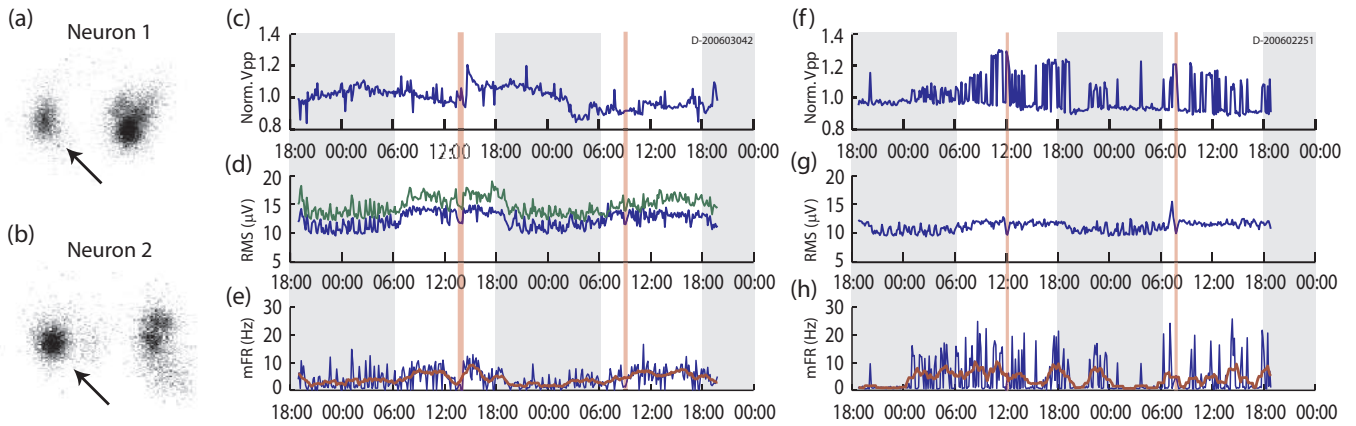


Fig. 3. Variation in V_{pp} . (a,b) Histogram of spike waveform projections into NWrPCA space. Selected neurons are indicated by arrows. (c) Normalized V_{pp} of neuron one recorded over 54 hours. (d) RMS noise of recorded channels over same period. (e) Mean firing rate of neuron #1, with 10 pt. moving average indicated by dark red line. (f-h) Same three plots for neuron #2. The wide light gray regions indicate nighttime, and the thin red regions indicate “pit stops,” when the monkey was taken from the home cage and placed in a primate chair in order to service the recording equipment.

in amplitude. The lines of constant voltage show the large changes in waveform amplitude. The waveforms show that changes in action potential shape, previously observed across once-daily recordings can occur at shorter timescales.

Figure 3c,f show the normalized waveform amplitudes for two neurons, identified in Fig. 3a,b, recorded from two different channels over different 54 hour periods. Each datapoint is the V_{pp} , or the peak-to-peak voltage, of the mean waveform for that neuron during the particular 5 minute block normalized by V_{pp} of the mean waveform for that neuron across the entire 54 hour dataset. Neurons are classified using training parameters determined separately for each block.

A single neuron is used (at a time) for this analysis to reduce the potential impact of the spike sorter. Since the relevant snippets are identified by the spike sorter, if there are very large changes it is possible to misclassify and thus obscure the variation we are looking for. However, using a well isolated unit, with a high firing rate, (histograms of the NWrPCA projections of the units used in this example are shown in Fig. 3a,b) reduces this possibility. Furthermore, the NWrPCA projections of the selected units are examined by the experimenter to ensure that snippets are not being ignored, or improperly included.

Variability in waveform amplitude, up to 30% relative to the mean, is observed over a range of time scales. There is clear variation on the order of a single block (7.5 minutes) as well as changes on the order of several blocks, and even several hours. The high frequency “chop” in the max waveform amplitude is highly correlated to the mean firing rate (mFR), shown in Fig. 3e,h, of the neuron across the block (correlation coefficients of $-.21$ and $.69$ for neurons one and two). The variation in waveform V_{pp} includes a strong low frequency component. The normalized V_{pp} has significant power at ~ 1 cycle/day synchronized to the day/night transition. This low frequency modulation is also seen in the mFR, which is a good proxy for general physical activity of the monkey [9].

Similar characteristics have been observed for other channels (not shown), indicating the changes in waveform am-

plitude observed in Fig. 3c,f are not unique to those channels. Similar variations were not observed, however, in hour long broadband recordings described in [4]. However, those recordings were made under a more traditional experimental protocol in which a restrained monkey performed a repetitive reaching task. It is possible the more controlled and consistent environment of those recordings, in contrast to the animal freely behaving in the home cage, produces a more consistent cortical environment (e.g., less “brain bounce”) and reduced variation in waveform shape.

Changes in the cortical environment in response to subject activity, including “brain bounce,” changes in intracranial pressure (ICP) and other homeostatic factors, may actually play a significant role in recorded waveform amplitude variation. At short to intermediate timescales (i.e., longer than bursting periods), [6] suggests that array movement, or more specifically changes in the neuron-electrode distance, might play a role in waveform shape change. The large change observed at 13:00 (day 1) in Fig. 3c is coincident with a vigorous head movement, and might be the result of abrupt movement of the array. This possibility is discussed further in the following section. Fluctuations in the ICP could potentially move the cortex tissue relative to the array (or vice-versa) Confirming such a relationship is beyond the scope of this work, though may be of interest in future studies.

IV. CHARACTERIZING CHANGES IN BACKGROUND NOISE PROCESS

Figure 3d,g show the RMS voltage of filtered neural recordings from three channels recorded over five minute blocks. All spikes, identified with thresholding at 3σ of RMS noise, have been removed from the dataset prior to the RMS calculation shown. Without the spikes, the RMS value should offer a better measure of the true background noise process [12].

Even after removing the identifiable spikes, though, the RMS noise is highly correlated to neural activity (as measured by mean firing rate). The variations ($\sim 5 \mu V$) partly result from distant spike activity (i.e., neural activity sensed by the

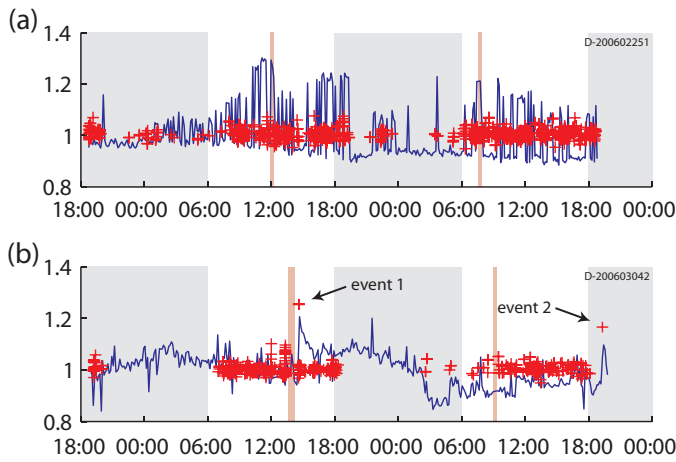


Fig. 4. Variation in waveform shape straddling high acceleration events (a,b) Local change in mean waveform amplitude ($V_{pp}^{after}/V_{pp}^{before}$) (red + symbols) for 200 snippets before and after 3g acceleration events overlaid on the normalized mean waveform amplitude (blue line) from Fig. 3c,f. The wide light gray regions indicate nighttime, and the thin red regions indicate “pit stops,” when the monkey was taken from the home cage and placed in a primate chair in order to service the recording equipment. Arrows in (b) indicate events of interest.

electrode, but the signal did not rise above the spike threshold because the spike amplitude was too small, or the neuron too far away). Depending on which data block is analyzed to set the threshold, there can be differences greater than $15 \mu V$ for a 3σ threshold.

V. CHARACTERIZING ABRUPT ELECTRODE SHIFT

An abrupt change in electrode array position in the cortex would presumably manifest itself as an abrupt change in waveform amplitude, as the neuron-electrode distance would change. If such changes do occur, we additionally presume they are correlated with high acceleration events such as vigorous head movement. Examination of recordings straddling high acceleration events show, in nearly all cases, far smaller changes in waveform amplitude than those observed over intermediate timescales suggesting that the array only rarely abruptly and substantially changes position.

Figure 4a,b show the normalized mean waveform amplitude (same neurons as shown in Fig. 3c,f), overlaid by local changes (red + symbols) in the mean waveform amplitude ($V_{pp}^{after}/V_{pp}^{before}$). Over a period of 50 hours, there were ~ 1700 and 800 high acceleration events for panels a and b, respectively. Acceleration events are identified with a 3g threshold. When computing the local change metric, the V_{pp} values are constructed from 200 snippets before and 200 snippets after the acceleration event. Those events with too few snippets available were dropped.

For nearly all events shown in Fig. 4a,b there is less than a 5% change in mean waveform amplitude straddling an acceleration event. There are however, two events in Fig. 4b, which show much larger changes. The NWrPCA projections of the before (blue) and after (green) snippets for the events indicated by the arrows are shown in Fig. 5a,c. The significant change in waveform shape is clearly seen in the NWrPCA




Fig. 5. Variation in waveform shape for events identified in Fig. 5b. (a) NWrPCA projection of 200 before (blue) and 200 after (green) snippets straddling acceleration event overlaid on NWrPCA histogram for all snippets in 5 minute block. (b) Local change in mean waveform amplitude ($V_{pp}^{after}/V_{pp}^{before}$) for 50 before and after snippets straddling each spike recorded from the neuron during the block in which the event occurred. The red + symbols mark the acceleration event. (c,d) Similar plots for event 2. The width of the peaks in (b,d) are a direct function of the number of spikes (50) used to calculate the mean.

projection. A second unit on this channel (the other cluster in the NWrPCA projection) shows a smaller change in amplitude (1.1 vs. 1.25 for event 1) straddling the same acceleration event suggesting that the observed variation does not result from changes in the signal conditioning pathway.

Figure 5b,d show the local change on a per spike basis (i.e., straddling each snippet recorded from the neuron) during the blocks in which the events were recorded. The close alignment between the acceleration events (indicated by the red + symbols) and the peaks in the local change metric suggest that the relationship between the change in waveform amplitude and the high acceleration event is not coincidental. Furthermore, the profile of the local change metric appears consistent with an abrupt change in array position. Well before and after the shift, the array is in a stable state, and one would expect a relatively stable waveform amplitude. Conversely, coincident with the shift, a sharp peak in the local change metric, indicating a step change in waveform amplitude, would be expected.

These results suggest that abrupt and substantial changes in electrode position can be associated with vigorous head movement, but that these abrupt and substantial changes are rare. These results are only preliminary, however, and will require more datasets and more animals for confirmation. Also, this analysis does not address all electrode-movement caused variation (e.g., variability seen in the previous section).

VI. DISCUSSION

Traditional experimental protocols that utilize discrete, daily recording periods have provided limited information regarding

neural recording stability. The daily sampling limits the potential characterization of variations to timescales of either minutes or days. The development of HermesB, a long duration, broadband neural recorder, enables nearly continuous sampling, and thus characterization at timescales of minutes, hours and days. We have shown examples from preliminary datasets of significant waveform shape and RMS noise variation at all three timescales. Both types of variation can have adverse effects on spike sorting performance, either through the use of an inappropriate threshold or outright misclassification. The improved statistical characterization of the stability of neural recordings enabled by these new long duration datasets will allow the principled design and evaluation of sorting algorithms.

Tolerance to some instabilities in neural recordings has already been incorporated into sorting algorithms. The short timescale variations in spike shape can be addressed by incorporating firing statistics into the spike sorting algorithm [13] and changes in RMS voltage (from which the threshold is typically derived) can be addressed through adaptive thresholding [12]. Long term variation, however, may require periodic retraining of the spike sorting parameters. With such readjustments, experimenters report the ability to track single neurons across months or even years (although experimenters cannot be sure the same neurons are being observed without constant tracking, a capability now available with HermesB). There does not appear to be a consensus on exactly what retraining period is required. Experiments that use discrete daily recording periods typically only update once per day. Future prosthetic systems which operate continuously will likely need to retrain more regularly.

The quality of the trained spike sorting parameters is paramount. Poor classification parameters, and thus poor classification performance, will affect all aspects of neural prosthetic system performance. This does not imply that systems should retrain arbitrarily often. Frequent retraining can have significant costs. For advanced spike sorting algorithms [10], the training algorithm is computationally expensive. Although a power feasibility study has shown that the power consumption of the algorithm in [10] is small relative to real-time classification, it was assumed that retraining would be required only every 12 hours [11]. If a much shorter training period is required, the power consumption of training could quickly become significant.

Sorting algorithms with an adaptive training approach that continuously integrates over an extended period, as opposed to discrete retraining, might be the best approach in light of the instability of neural recordings. A suitable adaptive algorithm would have an effective training interval short enough to track variations in waveform shape and background process, without the cost of traditional discrete retraining. The apparent rarity of abrupt changes in waveform shape due to rapid array movement may reduce the occurrence of a potential problem scenario in which abrupt retraining is required. Nonetheless, the rare presence of abrupt changes in waveform shape does suggest that to maximize spike classification accuracy, any

algorithm might benefit from the ability to initiate discrete retraining when step changes in the waveform shape are observed.

VII. CONCLUSION

Given the stark changes in NWrPCA projection shown in Fig. 2 it seems likely that current spike sorting algorithms will be affected by waveform variation. New adaptive approaches may be necessary. However, the actual end-to-end effects of these waveform changes are unknown. It is possible that the variations characterized in this paper do not actually result in a degradation of spike sorting performance. The preliminary characterization presented in this paper is the first step of a detailed statistical analysis of recording stability that will enable construction of high fidelity synthetic data sets. These datasets will enable the rigorous evaluation of end-to-end performance of current spike sorting methodologies and the principled design of future algorithms.

ACKNOWLEDGMENT

The authors would like to thank Caleb Kemere for many helpful discussions and Mackenzie Risch for expert veterinary care.

REFERENCES

- [1] A. B. Schwartz, "Cortical neural prosthetics," *Ann. Rev. of Neuroscience*, vol. 27, pp. 487–507, 2004.
- [2] M. D. Linderman, *et al.*, "An autonomous, broadband, multi-channel neural recording system for freely behaving primates," in *Proc. of Conf. of IEEE EMBS*, 2006, submitted.
- [3] M. S. Fee, P. P. Mitra, and D. Kleinfeld, "Variability of extracellular spike waveforms of cortical neurons," *J. Neurophysiol.*, vol. 76, pp. 3823–3833, 1996.
- [4] S. Suner, *et al.*, "Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex," *IEEE Trans. Neural Syst. Rehab. Eng.*, vol. 13, pp. 524–541, 2005.
- [5] J. C. Williams, R. L. Rennaker, and D. R. Kipke, "Stability of chronic multichannel neural recordings: Implications for a long term neural interface," *Neurocomputing*, vol. 26-27, pp. 1069–1076, 1999.
- [6] M. S. Lewicki, "A review of methods for spike sorting: the detection and classification of neural action potentials," *Network: Comput. Neural Syst.*, vol. 9, pp. R53–R78, 1998.
- [7] N. Hatsopoulos, J. Joshi, and J. G. O'Leary, "Decoding continuous and discrete behaviors using motor and premotor cortical ensembles," *J. Neurophysiol.*, vol. 92, pp. 1165–1174, 2004.
- [8] M. Churchland, *et al.*, "Neural variability in premotor cortex provides a signature of motor preparation," *J. Neuroscience*, vol. 26, pp. 3697–3712, 2006.
- [9] V. Gilja, *et al.*, "Multiday electrophysiological recordings from freely behaving primates," in *Proc. of Conf. of IEEE EMBS*, 2006, submitted.
- [10] M. Sahani, "Latent variable models for neural data analysis," Ph.D. dissertation, California Institute of Technology, 1999.
- [11] Z. S. Zumsteg, *et al.*, "Power feasibility of implantable digital spike sorting circuits for neural prosthetic systems," *IEEE Trans. Neural Syst. Rehab. Eng.*, vol. 13, pp. 272–279, 2005.
- [12] P. T. Watkins, G. Santhanam, K. V. Shenoy, and R. R. Harrison, "Validation of adaptive threshold spike detector for neural recording," in *Proc. of Conf. of IEEE EMBS*, 2004, pp. 4079–4082.
- [13] C. Pouzat, M. Delescluse, P. Viot, and J. Diebolt, "Improved spike-sorting by modeling firing statistics and burst-dependent spike amplitude attenuation: A markov chain monte carlo approach," *J. Neurophysiol.*, vol. 91, pp. 2910–2928, 2004.