A Comparison of Neural Feature Extraction Methods for Brain-Machine Interfaces

Timothy P. Gilmour, *Student Member, IEEE*, Lavanya Krishnan, Roger P. Gaumond, *Member, IEEE* and Ryan S. Clement, *Member, IEEE*

***Abstract*— Brain-machine interfaces (BMIs) have shown promise in augmenting people's control of their surroundings, especially for those suffering from paralysis due to neurological disorders. This paper describes an experiment using the rodent model to explore information available in neural signals recorded from chronically implanted intracortical microelectrode arrays. In offline experiments, a number of neural feature extraction methods were utilized to obtain neural activity vectors (NAVs) describing the activity of the underlying neural population while rats performed a discrimination task. The methods evaluated included standard techniques such as binned spike rates and local field potential spectra as well as more novel approaches including matched-filter energy, raw signal spectra, and an autocorrelation energy measure (AEM) approach. Support vector machines (SVMs) were trained offline to classify left from right going movements by utilizing features contained in the NAVs obtained by the different methods. Each method was evaluated for accuracy and robustness. Results show that most algorithms worked well for decoding neural signals both during and prior to movement, with spectral methods providing the best stability.**

I. Introduction

Over the past few decades brain-machine interfaces (BMIs) have been designed to control many types of outputs, including computer cursors, robotic arms, and selection of letters from an onscreen keyboard [1]. The main neural signal recording approaches include non-invasive surface electrode recording of the electroencephalogram (EEG), or the more invasive intracortical microelectrode array approach. Both techniques have shown promise in various applications. Although EEGs are inexpensive and safe, they are limited to low frequency recording and low spatial resolution because of the filtering effect of the skull and skin [2]. The most invasive technique, intracortical implantation, uses electrode arrays buried inside the brain to record the action potential and local field activities of neurons near the recording tips. Because of the high speed and high resolution interface with the brain, the research described in this paper deals with intracortical implant-based BMIs [3].

The main feature-extraction method used by current BMI systems is binned spike firing rates. To estimate the spike rates, neural action potential events (spikes) are detected within background noise by imposing a voltage threshold on each channel. When the amplitude of the signal crosses the threshold, a spike is registered at that time index, and the estimate of the spike rate comes from the number of threshold crossings within fixed time bins.

From low pass filtered versions of the continuously recorded signal local field potentials (LFPs) can be obtained. LFPs are low frequency cortical oscillations, believed to represent the summed synaptic input currents to the neurons in the vicinity of the electrode. LFPs have shown good longevity and information rates comparable to rates from binned spikes [4].

Another recent and novel method for forming a neural activity vector is the autocorrelation energy measure (AEM) [5]. This approach samples local populations of neurons like the LFP method, but uses the estimated autocorrelations of local neurons and "noise" to estimate the amount of signal energy that is due to neural activity, potentially offering better longevity than binned spikes as neural signal power drops below the noise power over time.

The experiments described in this paper were designed to directly compare several literature-standard feature-extraction methods like binned spike rates and LFPs alongside the newer AEM and spectral methods. The overall project involved recording neural signals from a rat moving left or right in response to a visual stimulus and testing the offline accuracy and robustness over time of various feature-extraction methods in classifying the direction of the rats' movement. This forced-decision paradigm was ideal for testing the accuracy of the different neural activity vector (NAV)-forming methods in binary classification (cf. [6]). Results show that the AEM algorithm performs quite well compared to the other feature-extraction methods.

II. Materials and Methods

A. Subjects

Three Sprague-Dawley rats (Harlan; weighing between 400 and 550 grams) were implanted for the study. All procedures and protocols followed NIH guide for the Care and Use of Animals and were approved by Penn State's Institute for Animal Use and Care Committee (IACUC). The rats were placed on a limited water schedule, and maintained on a modified light-dark cycle to ensure that their period of highest wakefulness (early nocturnal) coincided with the experimental sessions.

B. Behavioral Paradigm

The rats were trained in an operant conditioning chamber which had a retractable water tube and three nose-pokeholes. Each pokehole was equipped with infrared sensors and a cue

T.P. Gilmour is a graduate student in the Department of Electrical Engineering, and L. Krishnan and R.S. Clement are with the Department of Bioengineering, The Pennsylvania State University, University Park, PA 16802 USA (e-mail: rclement@psu.edu).

This work was supported by The Pennsylvania State University and the Whitaker Foundation.

light. The rats were trained to wait in the center pokehole while either the left or right target cue light was briefly flashed, after which the rat would have to enter the correct target pokehole to receive the reward (Fig. 1).

This waiting period in the center pokehole was designed to force the rat to make its decision before the actual movement and to potentially enable the recording of a pre-movement 'intent' signal. The research described in this paper makes use of both 'pre-movement' neural data from the 300ms window just prior to leaving the center pokehole, and 'during-movement' neural data from the 300ms window just prior to entering the target pokehole (Fig. 2).

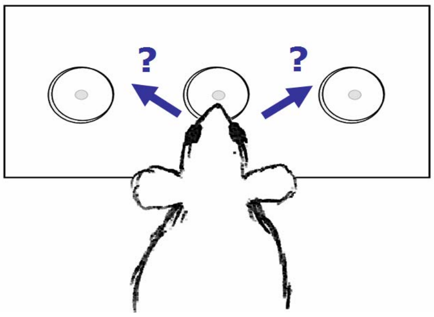

Fig. 1. Diagram of nose-poke hole set-up.

C. *Electrode Implantation and Neural Recordings*

The rats were implanted when they were acquiring water rewards in approximately 75% of the initiated trials. Subjects R16 and R17 were implanted unilaterally in primary motor cortex and subject R19 was implanted with two arrays (bilaterally) in premotor cortex. The microwire electrodes were hand fabricated and sterilized by gamma radiation. Each electrode had four rows of four polyimide-insulated 50µm tungsten microwires, spaced at approximately 500µm. The surgical implantation of the intracortical electrodes was performed under sterile conditions using standard protocols [6].

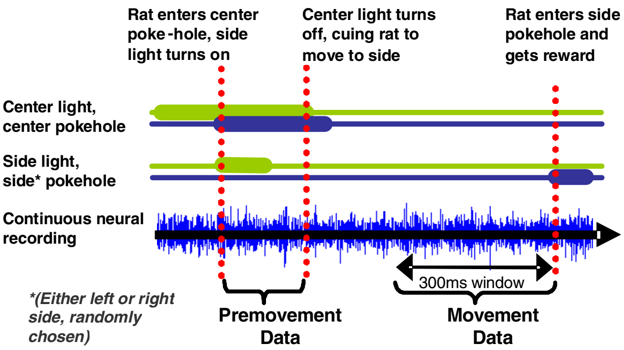

Fig. 2. Timing diagram describing behavioral paradigm

A commercial multi-channel acquisition system (Tucker-Davis Technologies) was used to collect simultaneous neural recordings from the electrodes of chronically implanted animals. After a head-mounted unity-gain buffering stage, a bio-amplifier (Medusa, TDT) digitized the signals at 25 kS/s and 16-bits. Continuous signals were digitally filtered in a DSP unit (Pentusa, TDT), down-sampled at 12 kS/s and streamed to a PC for offline analysis. The electrode impedances were also measured at 1 kHz with an electrode impedance meter (Bak Electronics).

D. *Feature Extraction*

The main analysis involved training a support vector machine (SVM) classifier on the processed data from the pokehole transition intervals to measure accuracy in predicting left and right movements.

Three processed versions of the original digitized neural signal were extracted and stored to disk. One version contained threshold-detected spikes (timestamps and 32-point waveform snippets; 25kS/s sampling), with the threshold adjusted manually in the neighborhood of two standard deviations from the signal mean for optimum detection. The second version was bandpass-filtered from 10 to 50 Hz and stored at 381 S/s, comprising the local field potential signals (LFP). The third version was bandpass-filtered from 300 to 5000 Hz and stored at 12.207 kS/s, picking up the higher spectral content of the neural action potentials in undecimated ("raw") waveform format. Also stored were the times of all pokehole entrances and exits, the on and off times of all cue lights, and the extension and retraction times of the reward.

From the three signal versions recorded, six different feature vectors (also called neural activity vectors, or NAVs) were generated for each pokehole-transition event, as depicted in Fig. 3. When analysis began, 28 recording blocks containing all three signal versions were hand-picked for investigation. This manual selection ensured reasonable minimum block length and reasonably good signal quality on valid channels. Each selected block was typically at least two minutes long and contained more than 30 total pokehole transitions. Four blocks were chosen from R16 performing the task over a 12 day period, fifteen blocks from R17 spanning 23 days, six blocks from R19U (left side) spanning 11 days, and three blocks from R19L (right side) spanning 9 days.

The Binned-spikes method (SPK) estimated spike firing rates by counting the number of (unsorted) spikes over 30ms time bins. The FFT on LFP method (fLFP) used smoothed FFT values from the 10-32 Hz frequency band as the SVM vector. The FFT on Raw-waveform (fRaw) method was the

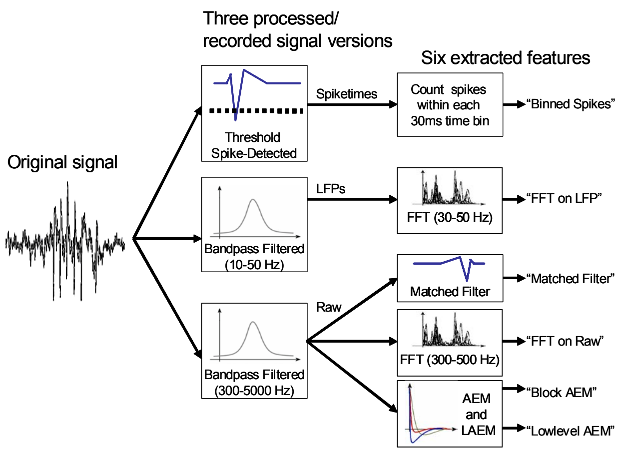

Fig. 3. Six Feature-Extraction Methods

same as the fLFP method except that it used the 100-300 Hz frequency band. These bands were selected for best discrimination by inspection of the average FFT difference between leftgoing and rightgoing trials on a typical recording block. The block-autocorrelation energy measure (bAEM) and Low-level autocorrelation energy measure (LAEM) used spectral information to estimate the fraction of signal energy that was due to neural action potentials [5]. LAEM was identical to bAEM except that it first removed spikes using a threshold. The Matched-filter energy method (MF) first created a filter template by averaging several thousand above-threshold spikes (8ms duration, aligned-peaks), for each channel. The template was then applied as a matched-filter on the incoming 30ms chunks of data. The RMS energies of each chunk of filtered data became the SVM vector.

In all the methods, the 10-point vectors generated from each of the 16 electrodes were joined to form the full 160-point vector for each pokehole-transition event.

E. Classification and Validation

The SVM classifier used a nonlinear radial basis function as a kernel to implicitly transform the 160-dimensional input vectors into a higher dimensional space in which the decision boundary was linear. Kernel-width and margin parameters were simultaneously optimized by an exhaustive search through logarithmic scales, computing 100 classifications per parameter combination and taking the mean prediction-accuracy value across all recording blocks as the performance indicator.

After computing the six types of NAVs for each pokehole transition interval in the 28 recording blocks, each NAV set was shuffled and subjected to twofold validation: half of the feature vectors were used to train the classifier, and the other half were used for testing. The random shuffling was repeated 10 times and the mean test accuracy saved for comparison.

The cross-validations were computed in two different ways: within-block and across-block. The within-block classifications trained the SVM on a subset of the data from a particular recording block and tested the SVM on the remainder of that data. The across-block classifications trained the SVM on all of the data from a particular recording block and tested the SVM on all of the data from a different recording block (several days distant), providing a measure of robustness-over-time of the different NAV algorithms.

III. Results

The accuracies for within-block classifications based on “movement” data are shown in Fig. 4, and the accuracies for within-block classifications based on “pre-movement” data are shown in Fig. 5. A direct comparison of the classification accuracies based on data from movement or pre-movement periods is shown in Fig. 6.

The across-block mean classification accuracies for movement data are shown in Fig. 7, separated by rat (Rat 17 had the most data, Rat 16 had the least data; Rat 19 had left and right bilateral implants, labeled R19U and R19L, respectively). All NAV methods are shown except bAEM, which was almost identical to LAEM. The across-block accuracies for pre-movement data are not shown, but these accuracies generally tapered off to 50% (chance level) after time intervals of only one day, in all rats.

IV. Discussion and Conclusion

As shown in Figs 4 and 5, the classification performance of the SVM on each NAV method was quite good,

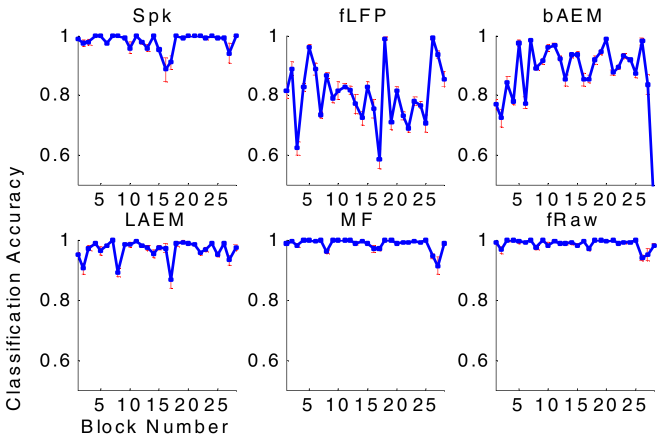

Fig. 4. "Movement" within-block mean classification accuracies for twofold cross-validation for each NAV method and for each recording block, with standard error (SE) bars.

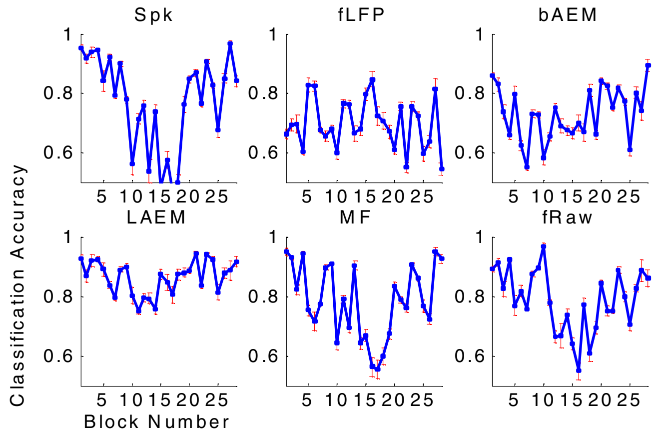

Fig. 5. "Pre-movement" within-block mean classification accuracies for twofold cross-validation for each NAV method and for each recording block, with standard error (SE) bars.

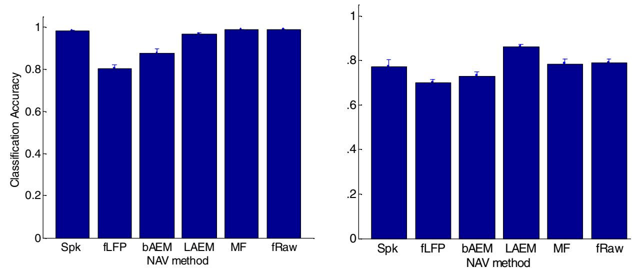

Fig. 6. Within-block mean classification accuracies for movement data (left) and pre-movement data (right), with SE errorbars.

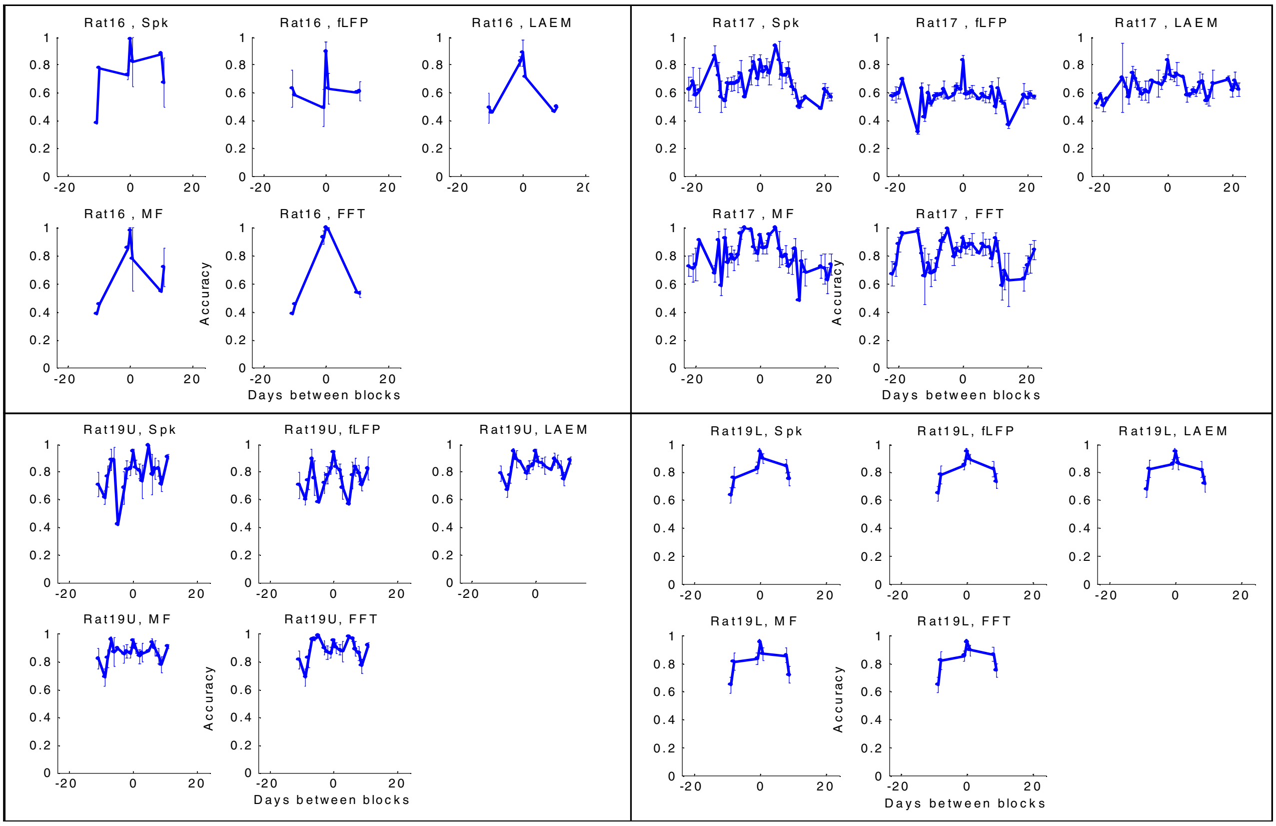

Fig. 7. Across-block mean classification accuracies for movement data (SE errorbars).

especially for the neural data collected while the animal was moving. The within-block classifications were quite accurate overall. Even for the premovement case, the mean classification accuracies were all significantly above chance (independent-sample Student's T test on the worst-performing fLFP compared against 50% gave T=12.5, $p<0.0001$).

For the across-block classifications (Fig 7), most NAV methods showed good performance at short time lags between the training and testing sets, but tapered off in accuracy rapidly as the time between training and testing sets grew larger than one or two days. Furthermore, the classification accuracy for data recorded during movement were higher on average than for premovement data. This may simply reflect differences in the strength of the neural modulation, or there may be other higher-level causes for the drop in accuracy after time such as movement of the electrodes relative to certain trained neurons or even neuroplastic retraining of the nearby neurons.

Comparing the means of the different NAV method accuracies for across-block movement data showed that both the mean MF accuracy and the mean fRaw accuracy were significantly greater than the mean SPK accuracy classification (T=6.32, $p<0.001$, and T=6.68, $p<0.001$, respectively). This suggests that the MF and fRaw spectral NAV methods are more robust neural features across days than binned spikes.

The novel Low-level AEM algorithm generally performed well in extracting usable neural response information from recordings, even after the spikes had been cropped out, suggesting the potential for continued usefulness in long-term BCI applications after encapsulation occurs. In fact, for premovement data, the Low-level AEM algorithm performed even better than the other methods (Figure 6), possibly because the prior cropping of large amplitude spikes removed other non-neural voltage artifacts that were detrimental to left-right classification.

Traditional methods such as binned spikes and FFT on LFP generally performed well in the experiment, concurring with previous research that these methods are useful in BMI work [6, 7]. Although FFT on LFP tended to provide the lowest classification accuracies (Figure 6), this performance could possibly be improved by further research into the best parameters (FFT length, frequency band, vector length, SVM parameters, etc.) for optimal performance.

Another conclusion is that spectral methods provide just as much discriminatory information as the thesholded-spike method during movements. This is significant for several reasons. First, spectral methods such as matched filtering could be computationally easier than spike thresholding if the sorting of spikes is required. Second, spectral methods show promise to be more robust than spike thresholding as spike amplitudes decrease with electrode encapsulation. The excellent performance of the "FFT on Raw" method may be partially attributable to the supervision available in picking the best left/right discriminatory frequency band, since the other feature extraction methods did not utilize task-specific information in this manner.

A final conclusion is that the neural data from both movement periods and premovement periods could be successfully classified offline by these methods, suggesting that the classifiers were extracting intentionality information and not just muscle activation information. This is tentative because there is always the possibility that the classification was based on tiny invisible muscle movements or sensory processing rather than true neural intent. However, the classification prior to task-related movement is a very positive result, and an important step toward human BCIs, where muscle activation information will not always be available.

In summary, experimental results show that while there was no clear best feature-extraction method for all circumstances, spectral methods such as matched-filter energy and FFT on the continuous signal provided the most robust accuracies across multiple days. The SVM-classification system based on features extracted from motor cortex delivers the capability to predict a rat's behavior in a forced binary choice paradigm with excellent accuracy.

ACKNOWLEDGMENT

Authors would like to acknowledge the help of Dr. Jeff Dodds and the animal care staff at Penn State as well as the assistance provided Jeremy White, Jonathan Lawrence, Sudharshana Seshadri, Natasha Tirko, Kirstin Tawse, and Michelle Matzko.

REFERENCES

[1] D. M. Taylor, S. I. Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, pp. 1829-32, 2002.

[2] G. M. Friehs, V. A. Zerris, C. L. Ojakangas, M. R. Fellows, and J. P. Donoghue, "Brain-machine and brain-computer interfaces," *Stroke*, vol. 35, pp. 2702-5, 2004.

[3] T. Gilmour, "A Comparison of Neural Feature Extraction Methods for Brain Machine Interfaces," in *Electrical Engineering*. University Park, PA: The Pennsylvania State University, 2005.

[4] C. Mehring, J. Rickert, E. Vaadia, S. Cardosa de Oliveira, A. Aertsen, and S. Rotter, "Inference of hand movements from local field potentials in monkey motor cortex," *Nat Neurosci*, vol. 6, pp. 1253-4, 2003.

[5] R. P. Gaumond, R. Clement, R. Silva, and D. Sander, "Estimation of neural energy in microelectrode signals," *J Neural Eng*, vol. 1, pp. 127-34, 2004.

[6] B. P. Olson, J. Si, J. Hu, and J. He, "Closed-loop cortical control of direction using support vector machines," *Neural Systems and Rehabilitation Engineering, IEEE Transactions on [see also IEEE Trans. on Rehabilitation Engineering]*, vol. 13, pp. 72-80, 2005.

[7] H. Scherberger, M. R. Jarvis, and R. A. Andersen, "Cortical local field potential encodes movement intentions in the posterior parietal cortex," *Neuron*, vol. 46, pp. 347-54, 2005.