Consideration of the functional relationship between cortex and motor periphery improves offline decoding performance

Matthew D. Best[1], Aaron J. Suminski[2], Kazutaka Takahashi[2], and Nicholas G. Hatsopoulos[1,2]

***Abstract*— Decoding neural activity to control prosthetic devices or computer interfaces is a promising avenue for rehabilitating individuals with amputation or severe spinal cord injury. In most cases, however, the local functionality of the neural tissue is not considered when designing a decoding algorithm. One way to characterize the functional specificity of a local region of motor cortex, and its output effects, is to use intracortical microstimulation. In this study, we examined how the results of an ICMS experiment relate to the performance of various offline decoders. We found evidence that units from electrodes with stimulation effects decode kinematics better than units from electrodes without stimulation effects.**

I. Introduction

Intracortical microstimulation (ICMS) is a widespread technique for characterizing the relationship between a region of motor cortex and its output behavior [1], [2], [3]. ICMS is thought to activate the neural tissue surrounding the electrode tip, and it's this activation that gives rise to ICMS evoked movements [3]. The movements evoked by ICMS in the upper limb region of primary motor cortex are often specific, and isolated to a specific joint, or set of joints [2], [3]. Parts of the upper limb are broadly localized to specific regions of the cortical surface (e.g. the hand representation is lateral to the arm representation, coarsely). Although ICMS characterizes motor outputs in the neighborhood of an electrode, information from ICMS experiments has not often directly informed studies of neural decoding.

Recently, much effort has been placed on decoding neural activity to control multiple degree of freedom brain machine interfaces [4]. One especially pertinent study demonstrates that accounting for the spatial properties of neural signals on several electrode arrays yields significant improvements in decoding performance, though this study spanned a large region of motor cortex ($\approx$ 12 mm) [7]. Mollazadeh and colleagues used ICMS to characterize the spatial distribution of limb motor representations, but were unable to explain the spatial variability in neural responses with their ICMS results [7]. We wondered, however, if the relationship between ICMS and decoding performance may be more apparent if we constrained ourselves to a smaller region of cortex, namely the 4×4 mm square spanned by our electrode array.

Accordingly, we examined the relationship between ICMS and offline decoding of joint kinematics in this study. We built decoding models that explicitly incorporated information about ICMS effects, and compared the performance of our ICMS-informed decoder to a similar model without ICMS information.

II. Methods

A. Behavioral Paradigm

One male rhesus macaque was operantly conditioned to perform a random target pursuit (RTP) task. In this task, the animal was required to move a cursor to targets appearing sequentially on a screen projected above his arm. After successfully hitting 7 targets in a row, the animal received a juice reward. Failure to hit a target in 5 seconds resulted in a failed trial and no juice reward. The cursor was controlled by a robotic exoskeleton [5]. A more complete description of the task may be found in [6].

B. Electrophysiology

We recorded unit spiking activity from one 96 channel Utah electrode array (Blackrock Microsystems, Salt Lake City, UT) implanted in the arm region of primary motor cortex (M1). Unit spiking waveforms were captured at 30 kHz (14 bit resolution) based on threshold crossings and sorted online using a hoop sorting algorithm. The average sorted waveform for each unit can be seen in figure 1.

C. Intracortical microstimulation

We used the Utah electrode array to stimulate motor cortex. Cathodal-anodal pulse trains were generated by a Blackrock current stimulator (pulse width: 0.2 ms, train-frequency: 333 Hz, train-duration 75 ms). Current intensities varied between 20 and 60 μA.

We began stimulating each electrode at 60 μA. If, after several repetitions, we did not evoke a movement, we considered that electrode inactive. However, when evoked movements were observed, we attempted to isolate the functional specificity of that electrode. If complex, or whole arm movements were evoked at that 60 μA, we would drop the current in 10 μA increments until the movement was well isolated, or no longer visible. For each electrode, the evoked movements at the lowest current that still reliably evoked movement were recorded as the behavior for that electrode. The sequence of electrodes that we stimulated through were chosen pseudorandomly.

*M.D.B, K.T., and N.G.H were supported by NIH grant R01 NS045853. A.J.S. and N.G.H. were supported by R01 NS048845.

[1]M.D.B and N.G.H are with the Committee on Computational Neuroscience, University of Chicago, Chicago, IL 60637, USA (e-mail: {mattbest, nicho} @ uchicago.edu)

[2]K.T., A.J.S., and N.G.H. are with the Department of Organismal Biology and Anatomy, University of Chicago, IL 60637, USA (e-mail: {kazutaka, asuminski} @ uchicago.edu)

978-1-4244-7929-0/14/$26.00

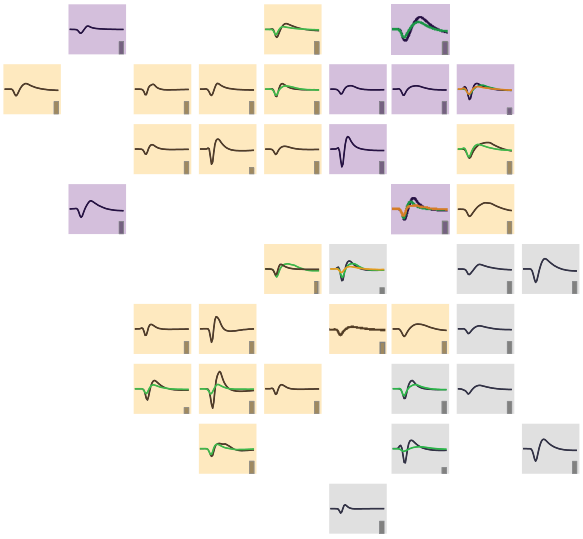

Fig. 1. Average waveforms of units on the electrode array. Background color indicates proximal (purple), distal (yellow), or no (gray) stimulation effects. The dark gray bar in the lower right corner of each panel is a scale bar indicating 125 μV. If more than one waveform was sorted from an electrode, the average waveform of each unit appears in a distinct color.

Movements evoked by microstimulation were characterized by two experimenters observing the animal. A movement was recorded only if it was visible to both experimenters and readily repeatable. We stimulated only when the animal appeared to be relaxed.

We characterized the behavior of each electrode as either proximal or distal. Shoulder and elbow movements were classified as proximal, while forearm muscle twitches and wrist movements were classified as distal. If an evoked movement contained both proximal and distal components it was classified according to the qualitative strength of each component, e.g. a movement evoking deltoid twitches, biceps twitches, and wrist flexion would be classified as proximal. In our sample, 4 out of the 27 electrodes contained both proximal and distal components (2 were classified as proximal, and 2 distal).

D. Decoding model

1) Input features: We recorded a total of 52 single units from the Utah array with a signal to noise ratio greater than 3. 39 out of those 52 units were recorded on electrodes that evoked movements when stimulated while, the remaining 13 units were recorded on 9 electrodes that did not evoke movements when stimulated. The spike counts of those 52 units, in 50 ms bins, served as input features to our model. For each moment in time, we predicted the kinematics at that moment from concurrent spike counts, as well as the spike counts in the two preceding bins (so that the spiking leads the kinematics by 50 and 100 ms, respectively). This timescale is compatible with a previous finding that the mutual information between spiking and kinematics peaks when spiking leads the kinematics by 100 ms [6]. Mathematically, we represented these input features in a matrix with size $N \times 156$ (52×3 lags) where N is the number of observations.

We normalized the spike counts in each column of our feature matrix by computing the z-score so that our input features better conformed to the assumptions of our decoding model. When cross-validating or testing, we used the mean and standard deviation of the training data to normalize the validation or test data.

2) Output features: We decoded five kinematic quantities in three coordinate frames: shoulder and elbow velocity; workspace x and y velocity; and wrist speed. We subtracted the mean of each of these quantities when fitting our model. As with the input features, we subtracted the training mean from validation and test data.

3) Behavioral data selection: We used data from all successfully completed trials (582 trials) in one experimental session. In total, this comprised 65289 observations of kinematics and spiking in 50 ms bins. Each observation was randomly and independently partitioned into one of three sets: train (comprising 70% of the data), validate (15%), and test (15%). Qualitatively, we observed that the specific allocation of data into train, validate, and test sets did not influence our findings. Therefore, the results that we report are based on one partitioning of the data.

4) Decoding model: We used a penalized linear regression model called ridge regression to relate neural spiking to movement kinematics [8]. Mathematically, we can express this model as an optimization problem in which we seek values of β to minimize the following objective:

$$\min_{\beta} \frac{1}{2}||K - X\beta||_2^2 + ||\Gamma\beta||_2^2 \tag{1}$$

where K is an $N \times 1$ vector of a given kinematic variable, X is the input feature matrix (with size $N \times p$), Γ is a $p \times p$ diagonal matrix with the value γ on its diagonal, and $||\cdot||_2$ denotes the Euclidean norm. A closed form solution to (1) exists and is given by:

$$\hat{\beta} = \left(X'X + \Gamma'\Gamma\right)^{-1}\left(X'K\right). \tag{2}$$

We used the validation set to find the value of γ that minimized error on the validation set. We then took the optimal $\hat{\beta}$ learned from the validation set and tested it on the test data. Performance metrics are reported from the test data. We quantified decoding performance using the coefficient of determination, R^2, given by the following formula

$$R^2 := 1 - \frac{\sum_i (K_i - X_i\hat{\beta})^2}{\sum_i (K_i - \bar{K})^2} \tag{3}$$

where $X_i\hat{\beta}$ indicates the model prediction at observation i and $\bar{K}$ is the arithmetic mean of K.

III. Results

A. Stimulation effects predict the presence of units

Out of the 96 electrodes on the Utah array, we observed stimulation effects on 48 electrodes, and recorded well isolated units on 36. 27 electrodes contained both a unit, and exhibited stimulation effects (figure 1). We wondered if the presence or absence of stimulation effects would be informative about the existence of units on that electrode.

That is, are stimulation effects statistically independent of unit recordings? We tested this by performing Fisher's exact test on the data in table I. We rejected the hypothesis that stimulation effects and and the presence of a unit on an electrode are independent events ($p < 0.001$). This suggests that if an electrode evokes movements when stimulated, it's more likely to have a unit on it.

TABLE I

PROPERTIES OF THE UTAH ARRAY. PARENTHESES INDICATE EXPECTED COUNTS UNDER STATISTICAL INDEPENDENCE

	stimulation effects	no stimulation effects	total
with units	27 (18)	9 (18)	36
without units	21 (30)	39 (30)	60
total	48	48	$n = 96$

B. Decoding from units with and without stimulation effects

We examined whether there may be some qualitative difference in decoding performance based on the presence or absence of stimulation effects at a given electrode. More simply, are units on electrodes with stimulation effects generally better at decoding than units on electrodes without stimulation effects? To answer this question we built decoding models for all of the kinematic variables based on spiking in the 13 units on electrodes without stimulation effects. We allowed the parameter, γ, to be learned from the validation set. A summary of decoding performance for each kinematic variable is presented in table II.

We performed a bootstrap analysis to assess whether or not decoding performance depends on the presence or absence of stimulation effects. Because there are substantially more units with stimulation effects (39 units) than without (13 units), we needed to sample that population. So, for each iteration of the bootstrap (1000 total), we drew a sample of 13 units, uniformly at random without replacement, from the population of units on electrodes with stimulation effects. We then quantified decoding performance for each of these bootstrapped samples. The distribution of goodness of fit values obtained by this bootstrap analysis is shown in figure 2, and summary statistics from this analysis are available in table II.

We found that each kinematic quantity had a higher median R^2 value from electrodes with stimulation effects, (table 2, column with stim effects) as compared to neurons without stimulation effects (table 2, column no stim effects). Moreover, decoding performance from the units with stimulation effects almost always exceeds that of the units without stimulation effects (table 2, column proportion).

TABLE II

DECODING PERFORMANCE SUMMARY

kinematics	no stim effects	with stim effects	proportion
shoulder velocity	$R^2 = 0.33$	$R^2 = 0.39$	0.82
elbow velocity	0.26	0.37	0.95
wrist speed	0.10	0.25	0.99
x velocity	0.26	0.44	0.98
y velocity	0.26	0.39	0.97

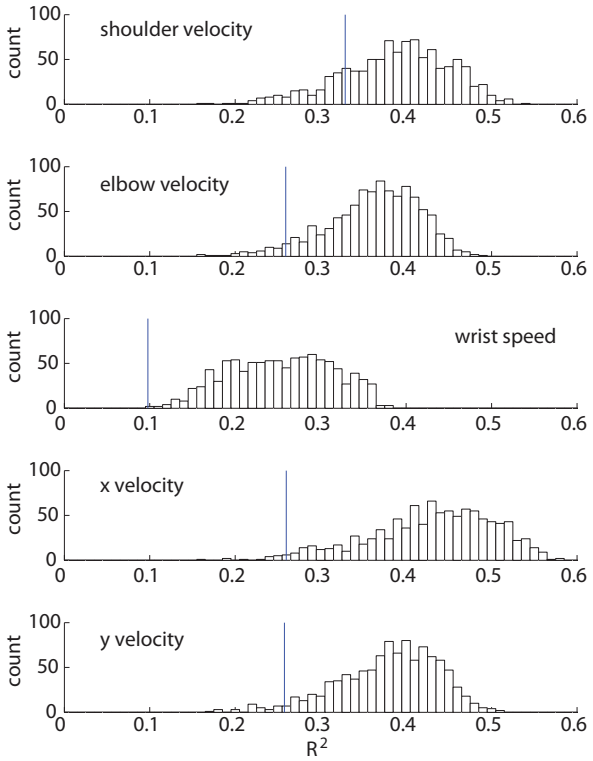

Fig. 2. Bootstrap analysis of decoding performance. We compared the decoding performance of units with and without stimulation effects. The blue vertical line indicates the R^2 of a decoder using units on electrodes without stimulation effects. The histogram depicts the R^2 values of decoders built using using a random sample of units on electrodes with stimulation effects. A decoder using units on electrodes with stimulation effects almost always outperforms a decoder using units one electrodes without stimulation effects.

Taken in sum, these results suggest that decoding performance is significantly better using a population of neurons from electrodes with stimulation effects. The discrepancy in decoding performance does not appear to be related to the quality of the neural signals in each population because the signal to noise ratios of the two populations are not significantly different (rank sum test, $p > 0.05$). Moreover, the firing rates of the two populations are also statistically indistinguishable ($p > 0.05$).

C. Does the functional specificity of ICMS benefit decoding?

In the previous section, we provided evidence that units on electrodes with ICMS effects do a better job of decoding kinematics; however, it is not clear why this is the case. One possibility is that the improvement in decoding performance can be explained by the functional specificity revealed by ICMS. That is, we hypothesize that units on electrodes that evoked shoulder movements predict shoulder velocity better than elbow velocity, and vice versa for units on electrodes that evoked elbow movements.

To this end, we performed another bootstrap analysis. From the 13 cells that evoked proximal movements, we drew a uniform random sample of 5. We then built decoders to predict shoulder velocity and elbow velocity using the neurons in that sample. We then quantified decoding performance as

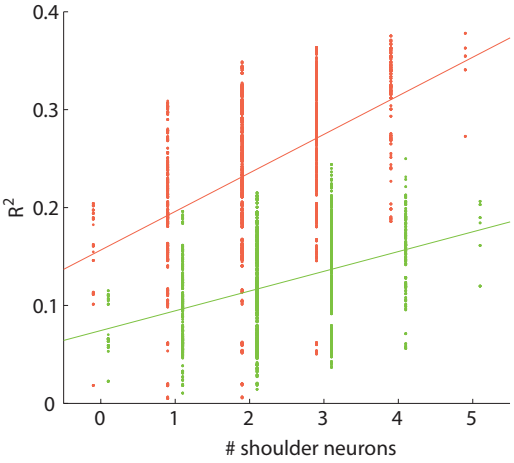

Fig. 3. We used a bootstrap procedure to select a random sample of 5 neurons on electrodes that evoked either shoulder or elbow movements, and then decoded shoulder (red) and elbow (green) velocities with those neurons. Here we present decoding performance as a function of the number of shoulder cells in the sample where a shoulder cell is defined to be a unit on an electrode that evoked shoulder movements when stimulated using ICMS.

a function of the number of cells on electrodes that evoked shoulder movements. Under our stated hypothesis, we would expect shoulder velocity decoding to improve as the number of shoulder neurons increases, while elbow velocity decoding performance would worsen. Interestingly, we find that for both shoulder and elbow velocity decoding, performance increases as the number of shoulder neurons increases (figure 3). However, the slopes of the shoulder and elbow decoders are significantly different (t test, $p < 0.001$). This suggest that as additional shoulder neurons are added, the shoulder decoder improves more than the elbow decoder, implying that the functional specificity of stimulation effects conveys some information about decoding performance in the cells.

There are several factors that may explain why both shoulder and elbow decoding performance improves with the number of shoulder neurons. First, the intersegmental dynamics of the limb imply that movements of the elbow will also create torques at the shoulder, and similarly, shoulder movements will create elbow torques if the elbow angle is maintained [5]. Thus, shoulder neurons may be highly sensitive to elbow movements because they create torques at the shoulder. Another possibility is that cells on elbow electrodes are simply less predictive by chance alone. There are a total of 7 units classified as elbow cells, so we are almost certainly under sampling the true distribution of units on electrodes that evoke elbow movements.

IV. Discussion

In this paper, we have shown that decoding performance from units on electrodes with stimulation effects is better than decoding performance from units on electrodes without stimulation effects. We suggest that at least some of this improvement may arise because the evoked movements from ICMS on a given electrode reveal information about the tuning properties of the units we record through that electrode.

One limitation of the present study, however, is that these results are based on data obtained from one monkey over a relatively short span of time. In future work, we plan to recapitulate the present findings in a second animal, while also considering how the relationship between ICMS and decoding performance changes over time. Lastly, we would like to more systematically consider the relationship between decoding performance and stimulation threshold. That is, do electrodes with the lowest stimulation thresholds contain units with the best decoding performance?

A. Relevance for a clinical neuroprosthesis

Performing our ICMS experiment, however, would be difficult in a person with severe motor disability receiving a neural prosthesis as the patient likely has either spinal cord damage, or an amputation. In this scenario, other techniques could be employed to determine the functional specificity of each electrode. Chiefly, it has been shown that motor cortex is active during imagined movements [9]. Thus, the response properties of units may be revealed by instructing the patient to imagine making various types of arm movements.

Acknowledgment

The authors would like to thank J. Coles for assistance in the care and training of laboratory animals and A. Dickey for useful comments and discussion while preparing the manuscript.

References

[1] S. D. Stoney, W. D. Thompson, and H. Asanuma. Excitation of Pyramidal Tract Cells by Intracortical Microstimulation: Effective Extent of Stimulating Current. Journal of Neurophysiology 31, no. 5 (September 1, 1968): 659669.

[2] H. C. Kwan, W. A. MacKay, J. T. Murphy, and Y. C. Wong. Spatial Organization of Precentral Cortex in Awake Primates. II. Motor Outputs. Journal of Neurophysiology 41, no. 5 (September 1, 1978): 1120–1131.

[3] M. C. Park, A. Belhaj-Saif, M. Gordon, and P. D. Cheney. Consistent Features in the Forelimb Representation of Primary Motor Cortex in Rhesus Macaques. The Journal of Neuroscience 21, no. 8 (April 15, 2001): 2784–2792.

[4] K. Balasubramanian, J. Southerland, M. Vaidya, K. Qian, A. Eleryan, A.H. Fagg, M. Sluzky, K. Oweiss, N. G. Hatsopoulos. "Operant conditioning of a multiple degree-of-freedom brain-machine interface in a primate model of amputation," Engineering in Medicine and Biology Society (EMBC), 2013 35th Annual International Conference of the IEEE , vol., no., pp.303,306, 3-7 July 2013

[5] S. H. Scott. Apparatus for Measuring and Perturbing Shoulder and Elbow Joint Positions and Torques during Reaching. Journal of Neuroscience Methods 89, no. 2 (July 15, 1999): 119127.

[6] W. Wu, and N. G. Hatsopoulos. Evidence against a Single Coordinate System Representation in the Motor Cortex. Experimental Brain Research 175, no. 2 (November 1, 2006): 197210.

[7] M. Mollazadeh, V. Aggarwal, A. G. Davidson, A. J. Law, N. V. Thakor, and M. H. Schieber. Spatiotemporal Variation of Multiple Neurophysiological Signals in the Primary Motor Cortex during Dexterous Reach-to-Grasp Movements. The Journal of Neuroscience 31, no. 43 (October 26, 2011): 1553115543.

[8] A. E. Hoerl, and R. W. Kennard. Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics 12, no. 1 (1970): 5567.

[9] B. Blankertz, G. Dornhege, M. Krauledat, K. Mller, and G. Curio. The Non-Invasive Berlin Brain Computer Interface: Fast Acquisition of Effective Performance in Untrained Subjects. NeuroImage 37, no. 2 (August 15, 2007): 539550.