

Learning the Contributions of the Motor, Premotor, and Posterior Parietal Cortices for Hand Trajectory Reconstruction in a Brain Machine Interface

Justin C. Sanchez¹, Deniz Erdogmus², Yadunandana Rao²,
Jose C. Principe², Miguel Nicolelis³, Johan Wessberg³

Departments of ¹Biomedical and ²Electrical and Computer Engineering
University of Florida, Gainesville, FL 32611
[justin, deniz, yadu, principe]@cnel.ufl.edu

³Department of Neurobiology, Duke University, Durham, NC 27710
[nicoleli, wessberg]@neuro.duke.edu

Abstract The ability to record, in real-time, the activity of hundreds of cortical neurons gives the ability to selectively study the function of clusters of cortical neurons in Brain Machine Interface (BMI) experiments. We have demonstrated using a recursive multilayer perceptron (RMLP) that using the appropriate signal processing theory in a well-chosen parsimonious model, we can develop constructs that agree with basic physiological modeling of neural control. By looking through the trained model, we have found interesting relationships between the neuronal firing and the movement. The RMLP allows us to continuously study the relationship between neural activity and behavior without the active interference of the experimenter. The findings presented in this study offer an opportunity for the neuroscience community to compare the cortical interactions as constructed by the RMLP to what is known about motor neurophysiology.

Keywords - Brain Machine Interface, BMI, neural control, population coding, recurrent neural network, neural ensemble, motor control

I. INTRODUCTION

With the development of the ability to record, in real-time, the activity of hundreds of cortical neurons, we have the freedom to selectively pick from the ensemble the neurons that we study in Brain Machine Interface (BMI) experiments. These experiments provide multi-unit recordings of behaviorally active animals, where not only neural activity, but also the corresponding reaching movements in 3D space are simultaneously recorded. The contributions of the activity of subsets of this large ensemble can now be directly studied relative to the continuous behavior of the animal over long periods of time. This new field of BMIs holds the promise of a new framework to help understand the neural code due to the availability of time synchronized, high-resolution quantitative data between multi-unit recordings and behavior.

Using input/output models we can relate a subset of the spatio-temporal spike trains that code the intent and the act of moving the hand (input signals) with fine timing relations of the hand movements (desired response). The model parameters (weights) are adjusted to minimize the difference between the model output and hand movements using a statistical criterion such as mean-square error. We repeat the

presentation of input training samples until a stopping criterion is met. At the end of training, one ends up with a model that is able to predict the desired hand positions from the spike trains. Many other BMI groups have demonstrated neural control of devices with chronically implanted electrodes. Most notably Nicolelis *et al.* have trained linear FIR filters and Time-Delay Neural Networks (TDNN) to predict the hand position of a primate in 3-D space [1]. Neural cursor control using linear filters trained with least squares has also been explored by Donoghue *et al.* [2]. More recently Schwartz, *et al.* has also shown neural cursor control using a modified version of the Population Vector Algorithm (PVA) [3]. Chapin and colleagues utilized a Recursive Multilayer Perceptron (RMLP) to predict lever pressing from ensembles of rat cortical neurons [4]. We have also demonstrated that the hand position of a primate can be predicted with high accuracy using a RMLP, Multiple Linear Models, and a Kalman filter [5] [6]. In our data, the RMLP model has produced the best reconstruction of hand movement from neuronal firing patterns when measured in terms of correlation coefficient, target acquisition, and signal to error ratio. This body of work indicates that a reasonably accurate functional mapping between neural activity and behavior can be constructed, and we hypothesize here that it can be further used to interpret relationships between the spike train and hand movements.

We will demonstrate below using an RMLP, that using the appropriate signal processing theory in a well-chosen parsimonious model, we can develop constructs that agree with basic physiological modeling of neural control. However, we would like to emphasize that the relevance of this analysis to neurophysiology cannot be guaranteed and it either requires further scrutiny by comparing the model predictions with established neurophysiological principles. At the very least, these predictions shall be used as a mere guiding hypothesis for further research work in neurobiology.

II. METHODOLOGY

In this experiment we are interested in the cortical regions that contribute to the triad of movements defined in Fig. 1. Each reaching movement can be segmented as a reach

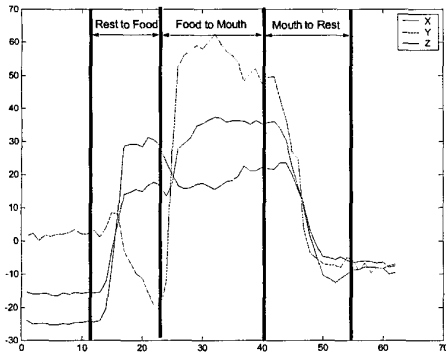


Figure 1. One movement segmented into rest/food, food/mouth, and mouth/rest motions.

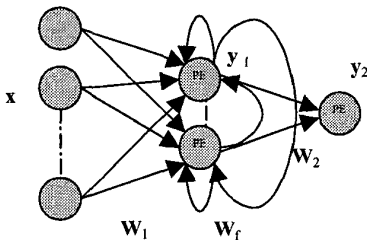


Figure 2. Fully connected, state recurrent neural network

to food (rest/food), a reach from food to mouth (food/mouth), and a reach from mouth to rest (mouth/rest). By training the RMLP model using the neurons associated with different combinations of sampled cortical areas, and observing how the network output changes during testing we can build a set of constructs and compare with established neurophysiological principles.

The RMLP modeling approach assumes that there exists an unknown system that maps spike trains into hand positions, and we can adapt a nonlinear model that can approximate the desired relationship. We can parameterize this unknown system by training the RMLP with multichannel neuronal firing times from up to 104 cells that were collected synchronously at Duke University using owl monkeys (*Aotus trivirgatus*). Microwire electrodes were implanted in the posterior parietal (PP) (Area 1), primary motor (M1) (Area 2, Area 4), and premotor dorsal (PMd) (Area 3). During the neural recording process, sixty-four electrodes were implanted with each cortical region receiving sixteen electrodes [7]. From each electrode, one to four neurons can be discriminated. The firing times of single neurons were recorded while the monkey performed the 3-D reaching task.

Neuronal firings, binned (added) in non-overlapping windows of 100ms, were directly used as inputs to the RMLP. The monkey's hand position, used as the network desired signal, was also recorded (with a time shared clock) and digitized with 200Hz sampling rate. In order to take the

monkey's reaction time into account, the spike trains were delayed by 0.23 seconds with respect to the hand position. This delay was chosen based on loose neurophysiologic reasoning, and should be subject to optimization in future studies. The architecture which uses this data can accept an input layer with any combination of up to 104 channels, a hidden layer of nonlinear processing elements (PEs), (in this case tanh), and an output layer of three linear PEs.

$$y_1(t) = f(W_1 x(t) + W_f y_1(t-1)) \quad (1)$$

$$y_2(t) = W_2 y_1(t) + b_2 \quad (2)$$

Fig. 2 depicts the topology of the recurrent network that is used in our studies. Each hidden layer PE is connected to every other hidden PE using a unit time delay. We can see in (1) that the state produced at the output of the first hidden layer is a nonlinear function of a weighted combination (including a bias) of the current input and the previous state. The feedback of the state allows for continuous representations on multiple timescales. The output layer is a simple linear combination, shown in (2), of the hidden layer states. Memory is created by feeding back the states of the hidden PEs among themselves. Each of the hidden PEs outputs can be thought of as a nonlinear adaptive basis of the output space utilized to project the high dimensional data. These projections are then linearly combined to form the outputs of the RMLP that will predict the desired hand movements.

The RMLP first proposed in [8] differs from a feedforward MLP since it contains feedback connections in its hidden layer. The MLP from which this topology is derived has been shown to be a universal mapper in \mathfrak{R}^n [9]. The time delay neural network (TDNN) has been also shown to be a universal mapper in myopic functional spaces [10]. Although no theoretical work to prove the universal approximation of the recurrent MLP is known, we expect it to display the same universality because it can be unfolded in a TDNN [11]. Hence this network when properly dimensioned and trained has the power to find an arbitrary non-linear mapping.

The RMLP was trained with backpropagation through time (BPTT) with a trajectory of 30 samples and learning rates of 0.01, 0.01, and 0.001 for the input, feedback, and output layers respectively. Momentum learning was also implemented with a rate of 0.7. One hundred Monte Carlo simulations with different initial conditions were conducted with 20,010 consecutive bins (2,001 secs) of neuronal data to improve the chances of obtaining the global optimum. During training all Monte Carlo simulations achieved a similar mean square error level, and the one with the smallest error was chosen. In testing, the network parameters were fixed and 3,000 consecutive bins (300 secs) of novel neuronal data were fed in the network to predict new hand trajectories. Fig. 3a shows the output of the trained network with 3-D hand position decomposed into X, Y, and Z

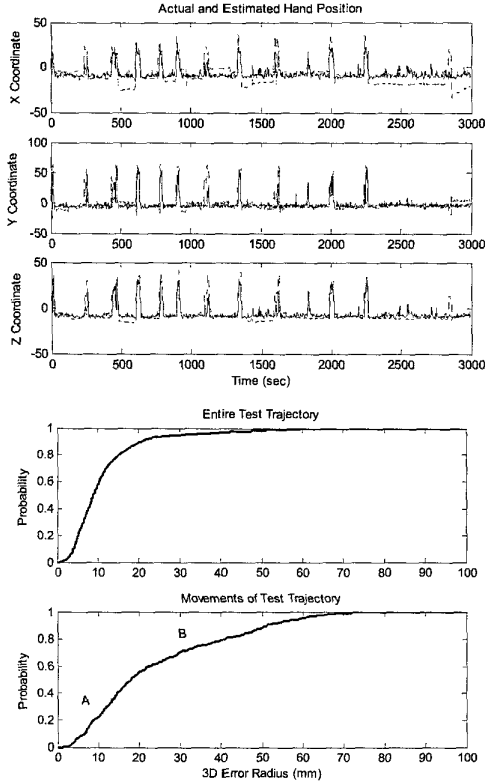


Figure 3. a) Actual (dotted) and estimated (solid) hand coordinates for the testing period using RMLP. b) Cumulative error probability (Probability that the error radius is smaller than the value on the horizontal axis.)

coordinates. We can see from the plots that the RMLP repeatedly produced accurate estimations of the peak values of the movements in the test set. In [9] we quantify the accuracy of the mapping in terms of correlation coefficient and signal to error ratio. The part of the trajectory that is most difficult to model is the X coordinate when the hand is at rest close to the animal's body. We also conclude that for the short period of observation (5 min) in this study, there is no noticeable progressive degradation in time of the model.

To evaluate the performance of the RMLP we propose a more specific figure of merit that emphasizes the accuracy of the reach movement. We first train the RMLP with the entire ensemble of neurons, and plot the probability of finding a network output (using a test dataset of novel neuronal data) within a 3-D radius around the desired data point. The first subplot of Fig. 3b shows the error probability for the entire test trajectory as a function of 3-D error radius as computed by (3) and (4). The dynamic range of the movement is approximately 100mm. Since we are training with the entire ensemble of neurons, we are sending to the model the maximal amount of available information; as a result, this evaluation represents the best performance possible.

$$\mathbf{e} = \mathbf{d} - \mathbf{y} \quad (3)$$

$$\text{Performance} = P\{\|\mathbf{e}\|_2 \leq r\} \quad (4)$$

Since target acquisition is critical for real BMI implementation, a plot of the error probability for only the movements is shown in the second subplot of Fig. 2b. The movement error probability curve contains two linear regions (this means the underlying PDF is roughly a mixture of two uniform distributions: region A (0-20 mm) can be attributed to the variance of the predictions at the peak values while region B can be attributed to peaks not captured by the model). We feel that the results of the RMLP movement estimations may be further improved for BMI applications. However, at this point in our study the RMLP model produces the best estimations when compared to other modeling techniques (FIR, local linear models, Kalman filters, and TDNN) [5] [6].

Now that we have an estimate of the best performance we can achieve with the full ensemble of neurons, we trained fifteen recurrent neural networks using all combinations of the base set of four cortical regions and neuronal firing counts. The same portions of training and testing data as well as network parameters were utilized. After training, the RMLP weights were frozen and five minutes of novel neuronal data was presented and output trajectories were produced.

III. RESULTS

The X, Y, and Z network outputs (bold) as well as the desired X, Y, and Z coordinates for one movement and each of the fifteen networks are plotted in Fig. 4. Cortical area 1 captured the reach rest/food but could not achieve the large displacement for food/mouth. Area 1 shows a poor fit (smaller network output) in the food/mouth region. Area 2 does not display any correlation to the desired trajectory even though neuronal firing in this region is nonzero. Sharp changes in the network output appear in movement transitions for the network trained with area 3. Area 4 accurately captures all three peaks in the food/mouth and mouth/rest regions, but misses the beginning of movement. Networks trained with combinations of the base cortical areas display the following trends in the hand trajectory reconstruction:

- All combinations trained without area 1 do not capture the reach rest/food.
- All combinations trained without area 4 do not capture the reach food/mouth
- All combinations trained with area 2 perform as well as without area 2. Only small variations in the rest/food trajectories are evident.
- The network trained with areas 1, 3, and 4 has smaller ripples during movement than the network trained with only areas 1 and 4.

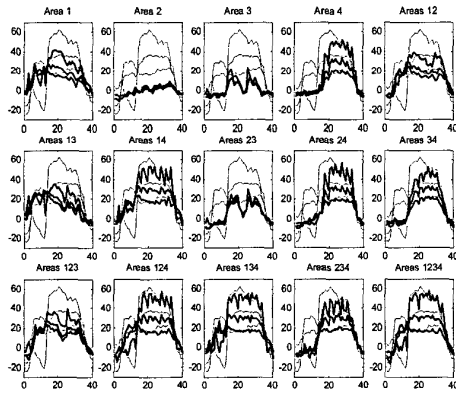


Figure 4. Testing output X, Y, and Z trajectories (bold) for one desired movement (light) from fifteen RMLPs trained with neuronal firing counts from all combinations of four cortical areas.

IV. DISCUSSION

In our development of BMIs we have adopted the RMLP as a tool not only to discover the mappings from spike trains to hand movements, but also for gaining insight into the neural coding strategies. The RMLP allows us to continuously study the relationship between neural activity and behavior without the active interference of the experimenter. We simply choose neural data to input to the model and if a mapping with the behavior exists, the universal approximating capabilities of the network will construct the relationship. In this paradigm we are passively observing how the network responds to inputs from different areas of the brain.

Even though we may not know the exact details of the neural code, our understanding of the inner workings of the model allow us to draw conclusions about some of the relationships between neuronal firing and behavior. In Fig. 4, we have observed that RMLPs trained without the neurons in the posterior parietal cortex fail to capture the reach to the food. However, when these neurons are included in the model, the network is able to produce the mapping and we conclude that some aspect of this group's firing is correlated with the movement. The same observation holds true with the RMLPs trained without neurons of the primary motor cortex. The absence of neural activity correlated with the behavior does not allow the RMLP to capture the reach from food to mouth to rest.

The results in this study offer the opportunity to speculate about the role selected cortical areas in the RMLP and how the conclusions connect to what is known about motor neurophysiology. However, we would like to emphasize that this analysis requires further scrutiny by the neuroscience community. Despite these issues by looking through the trained model, we have found interesting relationships between the neuronal firing and the movement. However, we must be cautious when drawing conclusions

from a model trained for a specific reaching task and a specific set of neural recordings. It is by far not all encompassing and representative of all movements and all neuronal firing patterns. In light of the overwhelming challenges associated with understanding neural coding, we still claim that the model based approach developed here to understand the strategies taken by motor cortex neurons is a small step towards the goal.

ACKNOWLEDGMENT

This work was supported by DARPA sponsored grant # ONR-450595112.

REFERENCES

1. Wessberg, J., et al., *Real-time prediction of hand trajectory by ensembles of cortical neurons in primates*. Nature, 2000. **408**(6810): p. 361-365.
2. Serruya, M.D., et al., *Brain-machine interface: Instant neural control of a movement signal*. Nature, 2002. **416**: p. 141-142.
3. Taylor, D.M., S.I.H. Tillery, and A.B. Schwartz, *Direct cortical control of 3D neuroprosthetic devices*. Science, 2002. **296**(5574): p. 1829-1832.
4. Chapin, J.K., et al., *Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex*. Nature Neuroscience, 1999. **2**(7): p. 664-670.
5. Sanchez, J.C., et al. *Input-Output Mapping Performance of Linear and Nonlinear Models for Estimating Hand Trajectories from Cortical Neuronal Firing Patterns*. in *NNSP*. 2002. Martigny, Switzerland: IEEE.
6. Sanchez, J.C., et al. *A Comparison Between Nonlinear Mappings and Linear State Estimation to Model the Relation from Motor Cortical Neuronal Firing to Hand Movements*. in *SAB Workshop on Motor Control in Humans and Robots: on the Interplay of Real Brains and Artificial Devices*. 2002. University of Edinburgh, Scotland.
7. Nicolelis, M.A., et al., *Reconstructing the engram: simultaneous, multisite, many single neuron recordings*. Neuron, 1997. **18**(4): p. 529-537.
8. Puskorius, G.V., et al., *Dynamic neural network methods applied to on-vehicle idle speed control*. Proceedings of the IEEE, 1996. **84**(10): p. 1407-1420.
9. Cybenko, G., *Approximation by superpositions of a sigmoidal function*. Mathematics of Control, Signals, and Systems, 1989. **2**: p. 303-314.
10. Sandberg, I.W. and L. Xu, *Uniform approximation of multidimensional myopic maps*. IEEE Transactions on Circuits and Systems, 1997. **44**: p. 477-485.
11. Principe, J.C., N.R. Euliano, and W.C. Lefebvre, *Neural and adaptive systems: fundamentals through simulations*. 2000, New York: Wiley.