

Model-Based Decoding of Reaching Movements for Prosthetic Systems

Caleb Kemere¹, Gopal Santhanam¹, Byron M. Yu¹, Stephen Ryu^{1,2}, Teresa Meng¹,
and Krishna V. Shenoy^{1,3}

¹Department of Electrical Engineering, Stanford University, Stanford, CA 94305

²Department of Neurosurgery, Stanford University, Stanford, CA 94305

³Neurosciences Program, Stanford University, Stanford, CA 94305

Abstract—Model-based decoding of neural activity for neuroprosthetic systems has been shown, in simulation, to provide significant gain over traditional linear filter approaches. We tested the model-based decoding approach with real neural and behavioral data and found a 18% reduction in trajectory reconstruction error compared with a linear filter. This corresponds to a 40% reduction in the number of neurons required for equivalent performance. The model-based approach further permits the combination of target-tuned plan activity with movement activity. The addition of plan activity reduced reconstruction error by 23% relative to the linear filter, corresponding to 55% reduction in the number of neurons required. Taken together, these results indicate that a decoding algorithm employing a prior model of reaching kinematics can substantially improve trajectory estimates, thereby improving prosthetic system performance.

Keywords—neural prosthetics, decode algorithms, pre-motor cortex, motor-cortex, brain-machine interfaces

I. INTRODUCTION

THE idea of enabling paralyzed or locked-in patients to control their limbs through direct brain interfaces is not a new one, but it may be that the next decade will see its realization. To this end, we have previously presented a technique for increasing the performance of neural prosthetic systems [1], [2] through the use of prior statistical movement models. In this work, we use a modified version of the model-based technique to validate our simulation results using experimentally gathered neural and movement data.

A. Goal-Directed Movements

An important and prevalent class of arm movements is end-point directed “reaches.” The trajectories of repeated reaches to a given target are stereotyped. For a prosthetic control system, the consequence of this observation is that expected reach trajectories fall on a limited manifold within the larger space of all potential trajectories. The probability distribution defined by this subspace constrains the model-based decoder as it reconstructs an intended trajectory from motor-related neural activity.

This work was funded in part by C2S2, the MARCO Focus Center for Circuit & System Solution, under MARCO contract 2003-CT-888 (C.K. and T.H.M.), by NDSEG fellowships (G.S. and B.M.Y.), the Christopher Reeve Paralysis Foundation (S.I.R. and K.V.S.) and the following awards to K.V.S.: NSF Center for Neuromorphic Systems Engineering at Caltech, ONR, Whitaker Foundation, Center for Integrated Systems at Stanford, Sloan Foundation, and Burroughs Wellcome Fund Career Award in the Biomedical Sciences. Please address correspondences to ckemere@stanford.edu.




Fig. 1. *Left*: Spike histograms showing plan and peri-movement activity to 7 peripheral targets shown in red (Unit G20040203.21.3). A model-based prosthetic decoder reconstructs a given trajectory from the neural activity of an ensemble of neurons using a prior model of reaching movements generated from data. The blue targets are those for the 22-target task used in this study. *Right*: The stereotyping of reaches is shown in the radial hand position for 50 reaches to the up-right target. The corresponding neural activity displays movement-related modulation on a trial-by-trial basis.

B. Plan and Peri-Movement Neural Activity

In motor-control regions of cortex (e.g., M1 and PMd), movement-related neural activity is typically divided into two categories. The first category, “plan” activity, is thought to arise from the initial stages of movement preparation. During the period between target presentation and movement cue in a delayed-reach paradigm, this type of activity varies with task-related aspects of the impending movement – target direction, distance, and speed [3], [4]. In tasks in which there is not an enforced delay period, similar patterns often arise during the period following the simultaneous target/movement cue [4]. The second category of neural motor activity, “movement” activity, is neural activity exhibited during on-going movements. Movement activity is thought to relate directly to the signals sent to muscles.

Functional magnetic resonance (fMRI) studies have shown that tetraplegic patients can generate similar patterns of activity as those observed in the motor cortex of healthy patients during movements [5]. While there is no evidence that the fine temporal structure of this neural activity remains unaffected by the loss of motor function, it seems reasonable to expect that plan and movement activity remain differentiated. Most work on prosthetic decoding, however, has not distinguished between these two types of activity. We have shown that the addition of plan activity to a movement activity-based decoding system can

result in significant performance gain [1], [2], [6]. In this report, we validate these previous simulation results using experimentally gathered neural and behavioral data.

II. NEUROPROSTHETIC DECODER ALGORITHMS

The goal of the prosthetic decoder is to reconstruct the trajectory of an intended movement, $x(t)$, from the neural activity, $\eta_n(t)$ observed in an ensemble of N neurons, $\mathbf{N}(t) = [\eta_n(t)], n = \{1, \dots, N\}$. The movement-related modulation of the neural units in our population led us to reconstruct trajectories in the velocity domain and then integrate the results to produce the final trajectory. We collect neural activity during measured physical reaches. Algorithm performance is computed by comparing the trajectory reconstructed from the neural activity to the recorded physical one.

A. Model-less Trajectory Reconstruction

The primary reconstruction algorithm in the neuroprosthetic literature is the linear filter or the closely related “population vector” [7]–[9]. The reconstructed trajectory, $\hat{x}_f(t)$, is

$$\hat{x}_f(t) = \sum_{n=1}^N \sum_{k=0}^{N_f-1} \mathbf{H}_{kn} \eta_n(t-k) \quad (1)$$

where the inner sum represents the convolution of the neural activity of neuron n with the (N_f) -tap subfilter in filter matrix \mathbf{H} . During training, we calculate the mean firing rate for each neuron and form a zero-mean value by subtracting it from each observation. We found empirically that the linear filter algorithm performed better if trained and implemented using the mean-subtracted firing rate in place of the basic observations. Methods for estimating the minimum mean-square error (MMSE) optimal filter coefficients from training data have been presented elsewhere [8]. In this work, we used a 10 tap filter, which corresponded to a 250 ms history.

B. Model-Based Trajectory Reconstruction

In model-based decoding, a prior model for reaching movements is used to constrain the reconstruction process. Define $x(t)$ as the hand trajectory at time t , and \mathbf{M}_t and \mathbf{P}_t as the history of observed neural movement and plan activity up to time t , respectively. The MMSE model-based estimate can be written as

$$\hat{x}(t) = \sum_{r \in \mathcal{R}} \mathbb{E}(x(t) | r, \mathbf{M}_t, \mathbf{P}_t) \Pr(r | \mathbf{M}_t, \mathbf{P}_t), \quad (2)$$

where the sum is taken over all reach goals, r , from the family of possible goals, \mathcal{R} . The two elements within the sum correspond to the estimate conditioned on a particular reach goal and the likelihood of that goal given the data. In [6], we described the use of a hidden Markov model (HMM) for reaches in concert with a vector velocity, or “cosine-tuning”, model for neural activity. However, the

neural data set we use in this work did not support this type of simple instantaneous model for relating neural firing to hand kinematics. Thus, we had to modify the previous algorithm. In the Appendix we describe the data-based method used for estimating the likelihood $\Pr(r | \mathbf{M}_t, \mathbf{P}_t)$. In addition, we found the following approximation to the optimal conditional estimate performed quite well.

$$\mathbb{E}(x(t) | r, \mathbf{M}_t, \mathbf{P}_t) \approx \left(C_r^{-1}(t) + C_f^{-1} \right)^{-1} \left(C_r^{-1}(t) \mu_r(t) + C_f^{-1} \hat{x}_f(t) \right) \quad (3)$$

where C_f is the covariance of the linear filter estimate, and $\mu_r(t)$ and $C_r(t)$, are the parameters for the time dependent distribution which describes the reach model for endpoint r .

III. EXPERIMENTAL DATA

A rhesus monkey (*Macaca mulatta*) was trained to perform center-out delayed reaching tasks to visual targets projected on a fronto-parallel screen. Each trial began with the monkey touching a central target. The monkey was then shown a target between 50 and 100 mm from center. Following a random 200 to 1000 ms “plan period,” a second visual cue instructed him to reach to the target. After holding the target for 200 ms, he was given a liquid reward. This trial structure is presented in Fig. 1.

A commercially available 96-channel silicon electrode array (Cyberkinetics Inc., Foxborough, MA, USA) was microsurgically implanted into the right hemisphere premotor cortex with standard neurosurgical techniques. The array is connected to a data acquisition system (Cerebus, Cyberkinetics Inc.) that provides online recording and processing of neural signals. Online manual spike sorting was performed by setting a voltage threshold trigger to obtain waveforms and time-amplitude hoops to isolate sets of waveforms defined as units. Multiple units, both single-neuron and multi-neuron units, could be isolated from each channel. We also recorded 3D arm position (60 samples/s, 0.1 mm resolution) using an infrared optical tracking system (Polaris, Northern Digital, Inc., Waterloo, ON, CA). After smoothing the hand trajectories with a 20-tap low-pass FIR filter (20 Hz cutoff) and binning neural activity, we discretized our time base at 25 ms. Protocols were approved by the Stanford University Institutional Animal Care and Use Committee.

IV. RESULTS

We tested the model-based decoding paradigm on data gathered during center-out reaches to 22 endpoints arranged in a square grid (blue circles in Fig. 1). Since the position of the hand was calculated by integrating the reconstructed velocity and the monkey did not center his hand identically at the beginning of each trial, we aligned each trajectory used for both training and testing by subtracting the initial hand location from the recorded movement. To test the algorithms, we used “N-fold cross-validation,” where 40

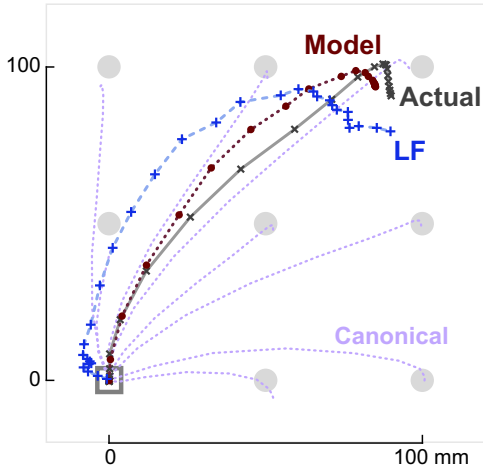


Fig. 2. Linear filter and model-based (M+P) reconstructions of an actual trajectory (Trial G20040119.697, MSE for LF and model-based are 3.5 and 2.1 cm^2). Line markers represent 25 ms steps. In the background are canonical trajectories for targets in the upper-right quadrant.

randomly selected trials were used for testing and the remainder, about 400, for training. For statistical significance, the process was repeated 30 times.

An example comparison of a model-based and linear filter reconstruction illustrating performance improvement is shown in Fig. 2. Comparing the first and second rows in Table I, the full model-based decoding algorithm reduces mean square reconstruction error by 18%, from 6.0 to 4.9 cm^2). To highlight the significance of this decrease, we randomly subsampled our neural population to yield smaller ensembles. Fig. 3 shows reconstruction error as the number of available units increases. To achieve the maximum performance of the linear filter (6 cm^2 with 111 units), the model-based decoder requires only 65 neural units (green solid line in Fig. 3).

TABLE I
TRAJECTORY RECONSTRUCTION MSE (cm^2)

Algorithm	Data Set [# units]		
	Grid [111]	7-Direction [65] [130]	
Linear Filter	6.0	11.6	9.7
Model-Based (M)	4.9	2.7	2.1
Model-Based (M+P)	4.6	2.4	2.6
Plan Only (Accuracy)	5.8 (64%)	4.6 (92%)	2.4 (97%)
Error Floor	1.6	1.5	1.5

We can establish an error floor for the portion of the model-based decoder purely derived from the reach model. The average MSE of the appropriate time-aligned canonical trajectory for a given reach is shown as the purple dashed line in Fig. 3 and the last row of Table I.

The addition of well-tuned plan activity should always assist reconstruction if the error of the movement-activity-

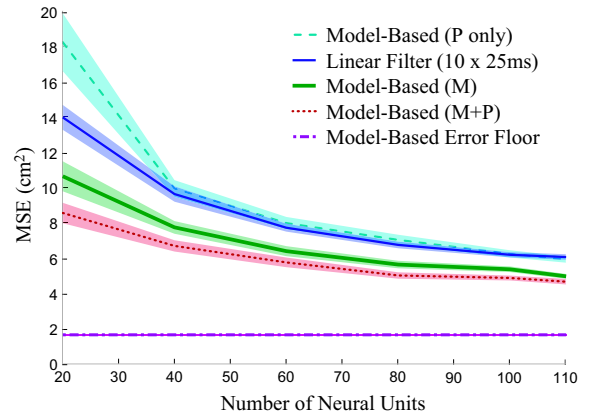


Fig. 3. Comparison of algorithms (mean \pm sem) as the quantity of neural information varies. Smaller ensembles of neural units were formed randomly, and performance was averaged (30 repetitions) over 40 randomly selected test trials and 400 training trials.

only decoder is above this level. The pink dotted line in Fig. 3 depicts the performance of the system when information from the plan period is added. While the addition of plan activity does reduce the average MSE an additional 5%, it does not aid the decoding process as much as expected. Furthermore, the cyan dashed line in Fig. 3 represents the reconstruction error when the decoder chooses the optimally aligned canonical trajectory from only plan activity. Notice that it lies well above the error floor. In examining the plan activity for the 22-target grid task, we found that the plan activity alone only predicted the proper endpoint with about 64% accuracy. Thus, to further investigate the utility of integrated plan activity, we tested the algorithm on a data set in which the plan activity was able to predict the movement endpoint with probability greater than 90%. This is shown as the last column in Table I.

The experimental paradigm for this second data set was center-out reaches to 7 targets arranged radially 10 cm from center (red circles in Fig. 1). In this data set, both plan and movement activity were able to accurately predict the correct endpoint. Thus, we see that the performance of the decoder with only plan or only movement activity was quite close to the error floor (2.1 and 2.4 cm^2 vs. 1.5 cm^2 error floor). However, because the movement activity was already very accurate, combining it with plan activity which was less accurate yielded poorer performance. The middle column presents the results of the same data set when only 65 randomly chosen units were made available to the decoder. In this case, the movement activity was less accurate, and thus the benefit of the addition of plan activity is apparent, as it was for the 22-target task. In a motor area in which plan activity dominated movement activity in decoding accuracy, the benefits of combined decoding could be more dramatic.

V. CONCLUSION

The performance of neuroprosthetic interfaces can be improved significantly through the use of a prior model

of reaching movements. There are two qualitative reasons for the increase in performance. First, the model-based decoder is aware of the distinction between reaching and not-reaching states. Thus, it is effective at reproducing the initial stationary portions of a reach. In contrast, the linear filter decodes the constant background of neural activity as small background movements. The second advantage of the model-based decoding paradigm is that, despite using the linear filter output, the prior reach model makes it more robust to the non-linearity of neural tuning. In the case of very linear and Gaussian neural activity, we would expect the performance gain over linear filters to be smaller. While these advantages do come at the cost of increased computation, for prosthetic interfaces this cost is relatively small when compared to the cost and risk of interfacing double the number of neural units. Thus, for neuroprosthetic interfaces which primarily decode reaching movements, the model-based approach, along with combining plan and movement activity, is an attractive option.

VI. APPENDIX: MODEL-BASED DERIVATION

We begin by assuming that the trajectory of the intended at time t , $x(t)$, arises from a family of goal-directed reaching movements, \mathcal{R} , and thus the *a posteriori* density can be written as

$$\Pr(x(t) | \mathbf{N}_t) = \sum_{r \in \mathcal{R}} \Pr(x(t) | r, \mathbf{M}_t, \mathbf{P}_t) \Pr(r | \mathbf{M}_t, \mathbf{P}_t). \quad (4)$$

where $\mathbf{N}(t) = \{\mathbf{M}_t, \mathbf{P}_t\} = [\{\mathbf{M}(\tau), \mathbf{P}(\tau)\}]$, $\tau = \{0, \dots, t\}$ are the movement and plan activity observed up to time t . For many reaches, the goal, r , is simply the endpoint of the movement. In general, however, it could include planned variations in path or speed.

The standard understanding of motor cortical activity is that it is the culmination of a serially ordered control system [4]. Thus, we have the following dependency model.

$$r \begin{cases} \rightarrow \mathbf{P}_t \\ \rightarrow x(t) \end{cases} \rightarrow \mathbf{M}_t \quad (5)$$

There are two important consequences of this structure. First, the actual movement is independent of the plan activity given the movement goal.

$$\Pr(x(t) | r, \mathbf{M}_t, \mathbf{P}) = \Pr(x(t) | r, \mathbf{M}_t) \quad (6)$$

Secondly, the implication of the lower path (5) is that movement activity is conditionally independent of the goal, given the actual movement.

$$\Pr(\mathbf{M}_t | x(t), r) = \Pr(\mathbf{M}_t | x(t)) \quad (7)$$

Therefore, using Bayes' rule, we can write

$$\Pr(x(t) | r, \mathbf{M}_t) = \frac{\Pr(\mathbf{M}_t | x(t)) \Pr(x(t) | r)}{\int_x \Pr(x(t), \mathbf{M}_t | r)}. \quad (8)$$

To proceed, the model-likelihood density, $\Pr(r | \mathbf{M}_t, \mathbf{P}_t)$, and the densities relating neural activity to

movement (the ‘‘neural tuning’’ density) and movement to goal (the ‘‘reach model’’ density) are required.

Reach Model Density: We estimated the reach model density, $\Pr(x(t) | r)$, directly from our data. For each endpoint, a canonical reach was generated. Specifically, the mean and covariance of the velocity at each point along the movement were calculated, beginning 150 ms before the movement onset. To account for a variable reaction time, we assume that the arm can remain at time ‘‘0’’ for an arbitrarily long period. Thus, for all t ,

$$\Pr(x(t - \tau_0) | r, \tau_0) \sim \mathcal{N}(\mu_r(t - \tau_0), C_r(t - \tau_0)), \quad (9)$$

where $\mu_r(s)$ and $C_r(s)$ are the canonical mean and covariance at time s in the canonical movement to endpoint r , and τ_0 represents the reaction time.

Neural Tuning Density: In the case of the ‘‘cosine-tuning’’ model used in our previous work,

$$\tilde{\mathbf{N}}(t) \approx \mathbf{G}x(t) + w(t), \quad (10)$$

where \mathbf{G} represents the preferred movement directions of the neuron ensemble, and $\tilde{\mathbf{N}}(t)$ is the mean subtracted firing of the ensemble. If $w(t)$ is normally distributed with mean zero and covariance Ω , and $x(t)$ has a prior normal distribution with mean μ and covariance C_x , the MMSE estimate is

$$\begin{aligned} \check{x}(t) &= \mu + C_x \mathbf{G}^T (\mathbf{G} \Omega \mathbf{G}^T + C_x)^{-1} (\mathbf{N}(t) - \mathbf{G}\mu) \\ &= (C_x^{-1} + C^{-1})^{-1} (C_x^{-1} \mu + C^{-1} \hat{x}(t)), \end{aligned} \quad (11)$$

where $\hat{x}(t)$ and C are the mean and covariance of the minimum variance unbiased (MVU) estimator for $x(t)$ in the absence of a prior.

Our data did not strongly support a simple instantaneous linear model. Unfortunately, as the reach model is time dependent, the use of tuning models with more complex time relationships (i.e., which relate firing to more than just instantaneous kinematic variables) would require computationally intensive variational methods to numerically solve (8). However, the second form of (11) does not depend explicitly on the observed neural activity. The activity is first linearly transformed into the reconstruction space, then combined with the prior. Replacing the MVU estimator in (11) with the output of the linear filter, \hat{x}_f , and its covariance, C_f , yields (3),

$$\begin{aligned} \mathbb{E}(x(t) | r, \mathbf{M}_t, \mathbf{P}_t) &\approx \\ & \left(C_r^{-1}(t) + C_f^{-1} \right)^{-1} \left(C_r^{-1}(t) \mu_r(t) + C_f^{-1} \hat{x}_f(t) \right), \end{aligned}$$

which results in a weighted combination of the linear filter output and the model prediction.

Model-Likelihood Density: The model-likelihood density, $\Pr(r | \mathbf{M}_t, \mathbf{P}_t)$, relates the observed neural activity to a particular canonical reach, and thus is the foundation of the model-based decoder. The likelihood of an observation given an endpoint, $\Pr(r | \mathbf{P}_t, \mathbf{M}_t)$, is simply the normalized product of the probabilities of the observed plan activity

and the observed movement activity calculated using the statistics for endpoint r .

$$\Pr(r | \mathbf{P}_t, \mathbf{M}_t) = \frac{\Pr(\mathbf{P}_t | r) \Pr(\mathbf{M}_t | r)}{\sum_{r \in \mathcal{R}} \Pr(\mathbf{P}_t | r) \Pr(\mathbf{M}_t | r)} \quad (12)$$

We estimated canonical movement activity vectors for each neuron using a method similar to that used for the canonical trajectories. For each neuron and each endpoint, we calculated the mean smoothed time-dependent firing rate for the movement period. We assumed that the movement activity follows a Poisson process with this time-dependent rate vector. As above, we account for the reaction time by allowing the rate corresponding to the initial window to be repeated. For each possible endpoint, the reaction time offset which maximizes the likelihood of the reach is chosen.

$$\Pr(\mathbf{M}_t | r) = \max_{\tau_0} \Pr_{\text{Pois}} \left(\mathbf{M}_t | \lambda_t = \overline{\mathbf{M}}_{t+\tau_0} \right), \quad (13)$$

where the probability function on the right-hand side is in vectorized form: λ_t is mean of the Poisson distribution, and $\overline{\mathbf{M}}_t$ is the average binned firing rates for movement activity up to time t in a reach to endpoint r .

We modeled the rate of plan activity (spikes/s) over the entire interval from the target onset until the presentation of the go cue – excluding the initial 150 ms to account for a potential visual transient – with a normal distribution.

$$\Pr(\mathbf{P}_t | r) \sim \mathcal{N}(\mu_r, C_r), \quad (14)$$

where μ_r and C_r are the mean and covariance of firing rate of plan activity before movements to endpoint r .

The model-based estimate in (2) can thus be found using (12) and the approximation in (11).

ACKNOWLEDGEMENT

The authors would like to thank Missy Howard for expert surgical assistance and veterinary care.

REFERENCES

- [1] C. T. Kemere, G. Santhanam, B. M. Yu, K. V. Shenoy, and T. H. Meng, ‘Decoding of plan and peri-movement neural signals in prosthetic systems,’ in *Proc. IEEE Workshop on Signal Processing Systems (SIPS ’02)*, San Diego, CA, Oct. 2002, pp. 276–283.
- [2] C. T. Kemere, K. V. Shenoy, and T. H. Meng, ‘Model-based neural decoding of reaching movements: A maximum likelihood approach,’ *IEEE Trans. Biomed. Eng.*, in press.
- [3] K. V. Shenoy, M. M. Churchland, G. Santhanam, B. M. Yu, and S. I. Ryu, ‘Influence of movement speed on plan activity in monkey pre-motor cortex and implications for high-performance neural prosthetic system design,’ in *Proc. IEEE Engineering in Medicine and Biology Society 25th Annual Conference (EMBS ’03)*, Cancun, Mexico, Sept. 2003, pp. 2079–2082.
- [4] D. J. Crammond and J. F. Kalaska, ‘Prior information in motor and premotor cortex: Activity during the delay period and effect on pre-movement activity,’ *J. Neurophysiology*, vol. 84, no. 2, pp. 986–1005, Aug. 2000.
- [5] S. Shoham, E. Halgren, E. M. Maynard, and R. A. Normann, ‘Motor-cortical activity in tetraplegics,’ *Nature*, vol. 413, p. 793, 25 Oct. 2001.
- [6] C. T. Kemere, M. Sahani, and T. H. Meng, ‘Robust neural decoding of reaching movements for prosthetic systems,’ in *Proc. IEEE Engineering in Medicine and Biology Society 25th Annual Conference (EMBS ’03)*, Cancun, Mexico, Sept. 2003, pp. 2079–2082.
- [7] D. Taylor, S.I.Helms-Tillery, and A. Schwartz, ‘Direct cortical control of 3D neuroprosthetic devices,’ *Science*, vol. 296, no. 3, pp. 1829–1832, June 2002.
- [8] M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. P. Donoghue, ‘Instant neural control of a movement signal,’ *Nature*, vol. 416, no. 6877, pp. 141–142, Mar 2002.
- [9] J. Carmena, M. Lebedev, R. E. Crist, J. E. O’Doherty, D. M. Santucci, D. Dimitrov, P. Patil, C. S. Henriquez, and M. A. Nicolelis, ‘Learning to control a brain-machine interface for reaching and grasping by primates.’ *PLoS Biol.*, vol. 1, no. 2, pp. 193–208, Nov. 2003.