A PARALLEL POINT-PROCESS FILTER FOR ESTIMATION OF GOAL-DIRECTED MOVEMENTS FROM NEURAL SIGNALS

Maryam Modir Shanechi[†], Gregory W. Wornell[†], Ziv Williams[‡], and Emery N. Brown[‡] [*]

† EECS MIT, ‡ Massachusetts General Hospital, ∗ HST MIT

ABSTRACT

Brain machine interfaces work by mapping the relevant neural activity to the intended movement known as 'decoding'. Here, we develop a recursive Bayesian decoder for goal-directed movements from neural observations, which exploits the optimal feedback control model of the sensorimotor system to build better prior state-space models. These controlled state models depend on the movement duration that is not known a priori. We thus consider a discretization of the task duration and develop a decoder consisting of a bank of parallel point-process filters, each combining the neural observation with the controlled state model of a discretization point. The final reconstruction is made by optimally combining these filter estimates. Using very coarse discretization and hence only a few parallel branches, our decoder reduces the root mean square (RMS) error in trajectory reconstruction in reaches made by a rhesus monkey by approximately 40%.

Index Terms— Brain machine interfaces, recursive Bayesian filters, neural signal processing, optimal control

1. INTRODUCTION

Information about an intended movement is encoded in the brain in the form of ensemble neural spiking activity modeled well by point processes [1]. It has been shown that motor cortical neural signals can be used to individually decode both the kinematics of a movement and its higher level information such as intended target. Consequently, one can improve the decoding accuracy of movement kinematics by exploiting this information about its goal.

One approach to decoding a movement trajectory is based on recursive Bayesian estimation; see, e.g., [2,3]. A recursive Bayesian decoder consists of two probabilistic models: the prior model on the time sequence of kinematic states, and the observation model relating the neural signal to these states. Examples are point-process filters [4] and Kalman filters [3].

We consider the problem of decoding a goal-directed movement of *unknown* duration and develop a recursive Bayesian decoder that takes advantage of the known goal information to reduce the mean-square error (MSE) in trajectory reconstruction. There are two major components to

our decoder. The first component aims at building better prior state-space models for the kinematics by exploiting an optimal feedback control model of the sensorimotor system shown to explain many of the observed phenomena [5,6]. [1]

The optimal feedback-controlled state-space model (as well as alternative goal-directed state-space models in [2,8]) depends on the movement duration that is not known a priori to the external observer of the neural signal. Hence the second component of our algorithm addresses this uncertainty in goal-directed movements in contrast to other work that assume this timing is known [2,8]. We address this by exploiting a parallel bank of point-process filters that calculate not only causal estimates of the state at each time based on the neural observations, but also the likelihood of the arrival time based on these observations. [2] Since these filters run in parallel, the time to generate the overall estimate is on the order of the run time of a single filter. We test our algorithm on real goal-directed movements performed by a rhesus monkey by simulating the neural activity.

2. PROBLEM STATEMENT AND NOTATION

We denote the sequence of kinematic states by $\mathbf{x}_0, \cdots, \mathbf{x}_t$ and the neural point process observations of the ensemble of C neurons by $\mathbf{N}_1, \cdots, \mathbf{N}_t$ where $\mathbf{N}_t = (N_t^1, \cdots, N_t^C)$ is the binary spike events of the C neurons at time t. The point-process observation model is given by [1]

$$p(\mathbf{N}_t|\mathbf{x}_{1:t}, \mathbf{H}_t) = \prod_c (\lambda^c(t|\mathbf{x}_t, \mathbf{H}_t^c)\Delta)^{N_t^c} e^{-\lambda^c(t|\mathbf{x}_t, \mathbf{H}_t^c)\Delta} \quad (1)$$

where $\mathbf{H}_t^c = N_{1:t-1}^c$, $\mathbf{H}_t = \mathbf{N}_{1:t-1} = \mathbf{H}_t^{1:C}$, Δ is the time increment and $\lambda^c(t|\mathbf{x}_t, \mathbf{H}_t^c)$ is the modeled conditional intensity function of the cth neuron at time t (as will be discussed in Section 7). We have assumed that the observations from the C neurons are conditionally independent.

The goal of the decoder is to causally calculate the state posterior density, i.e., $p(\mathbf{x}_t|\mathbf{N}_{1:t})$ based on the observations.

---

[1]This is in contrast to most work that place a random-walk prior state model on the kinematics. Alternative state-space models for goal-directed movements condition the random-walk model on being at the target at a known arrival time [2], or fit a time-invariant state model to empirical reaches to the target [7], or use a feedforward controlled state-space model [8] not taking into account the sensory feedback.

[2]Our approach can also be applied to the state models of [2,8] to generalize them and relax their assumption of known movement duration.

3. OPTIMAL FEEDBACK-CONTROLLED STATE-SPACE MODEL

To derive the prior state-space model on $\mathbf{x}_{1:t}$ in a goal-directed movement we exploit the optimal feedback control view of the sensorimotor system [5,6]. We assume that the kinematic state is generated according to the linear dynamical system

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + \mathbf{w}_t, \tag{2}$$

where $\mathbf{u}_t$ is the control input at time t, $\mathbf{w}_t$ is the zero-mean white Gaussian state noise with covariance matrix $\mathbf{W}$, and $\mathbf{A}$ and $\mathbf{B}$ are parameters of the sensorimotor system. Here we assume that the sensory feedback $\mathbf{y}_t$ is noiseless and $\mathbf{y}_t = \mathbf{x}_t$. Note that setting $\mathbf{B} = \mathbf{0}$ reduces the model to the random-walk model. One now needs to pick a cost function that will then be minimized by finding the optimal values of $\mathbf{u}_t$. The cost function in a given task should quantify its goal. For the above linear Gaussian dynamics, if we pick the cost function as a quadratic function of the state and control variables, i.e.,

$$J = \sum_{t=1}^{T} \left(\mathbf{x}_t' \mathbf{Q}_t \mathbf{x}_t + \mathbf{u}_t' \mathbf{R} \mathbf{u}_t \right), \tag{3}$$

where T is the movement duration, $\mathbf{Q}_t$ is positive semidefinite and $\mathbf{R}$ is positive definite, then the optimal control rule, $\mathbf{u}_t$, is simply a linear feedback of the state at that time [9], i.e.,

$$\mathbf{u}_t = -\mathbf{L}_t(T)\mathbf{x}_t, \tag{4}$$

where $\mathbf{L}_t$ can be found recursively and offline. This is the linear quadratic Gaussian (LQG) solution. Note that $\mathbf{Q}_t$ and $\mathbf{R}$ should be appropriately designed for an application of interest (see Section 6) and $\mathbf{L}_t(T)$ is time-varying and a function of the arrival time T. This reduces the state-space model in (2) to the optimal closed-loop controlled state-space model

$$\mathbf{x}_{t+1} = (\mathbf{A} - \mathbf{B}\mathbf{L}_t)\mathbf{x}_t + \mathbf{w}_t, \tag{5}$$

which can now be used as the prior on the kinematic states as opposed to a random-walk state model. Note that the closed-loop controlled dynamics matrix is now time-varying. We first derive our decoder for this general controlled state model and later specialize it to a reaching movement.

4. ESTIMATING THE FEEDBACK-CONTROLLED SYSTEM WITH KNOWN DURATION FROM NEURAL OBSERVATIONS

To derive the recursive Bayesian decoder, we now combine the state-space model in (5) with the point process observation model in (1) to recursively find the posterior density. For now, we assume that the arrival time and hence $\mathbf{L}_t(T)$ are known. We can write the posterior density as

$$p(\mathbf{x}_t | \mathbf{N}_{1:t}, T) = \frac{p(\mathbf{N}_t | \mathbf{x}_t, \mathbf{N}_{1:t-1}, T) \, p(\mathbf{x}_t | \mathbf{N}_{1:t-1}, T)}{p(\mathbf{N}_t | \mathbf{N}_{1:t-1}, T)} \tag{6}$$

The first term in the numerator comes from the observation model in (1) and the second term is the one-step prediction density. Note we used the fact $p(\mathbf{N}_t | \mathbf{x}_t, \mathbf{N}_{1:t-1}, T) = p(\mathbf{N}_t | \mathbf{x}_t, \mathbf{N}_{1:t-1})$. We denote the partition function by

$$g(\mathbf{N}_t | T) = p(\mathbf{N}_t | \mathbf{N}_{1:t-1}, T) \tag{7}$$

since we will exploit it later for the case of unknown arrival time. To get the recursion for the posterior density, we write the prediction density as

$$p(\mathbf{x}_t | \mathbf{N}_{1:t-1}, T) = \int p(\mathbf{x}_t | \mathbf{x}_{t-1}, T) p(\mathbf{x}_{t-1} | \mathbf{N}_{1:t-1}, T) d\mathbf{x}_{t-1} \tag{8}$$

using the conditional independence, $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{N}_{1:t-1}, T) = p(\mathbf{x}_t | \mathbf{x}_{t-1}, T)$, which comes from the optimal feedback-controlled state-space model in (5). Now the second term inside the integral is just the posterior density from the previous time step, hence substituting (8) into (6) generates the recursion. The exact expression in (6) is in general complicated. Hence here we make a Gaussian approximation to the posterior density similar to [4]. Making this approximation and since the state-space model in (5) is also Gaussian, the prediction density in (8) will be Gaussian. Let's denote the minimum MSE (MMSE) estimator, i.e., $E(\mathbf{x}_t | \mathbf{N}_{1:t}, T)$ by $\mathbf{x}_{t|t,T}$ and its covariance matrix by $\mathbf{W}_{t|t,T}$. Similarly we denote the one step prediction mean by $\mathbf{x}_{t|t-1,T} = E(\mathbf{x}_t | \mathbf{N}_{1:t-1}, T)$ and its covariance matrix by $\mathbf{W}_{t|t-1,T}$. The recursions for the MMSE estimator with this Gaussian approximation have been derived in [4]. The difference here is that our state-space model is a controlled one. The recursions of this point-process filter in our case become

$$\mathbf{x}_{t|t-1,T} = (\mathbf{A} - \mathbf{B}\mathbf{L}_t(T)) \, \mathbf{x}_{t-1|t-1,T} \tag{9}$$

$$\mathbf{W}_{t|t-1,T} = (\mathbf{A} - \mathbf{B}\mathbf{L}_t(T))\mathbf{W}_{t-1|t-1,T}(\mathbf{A} - \mathbf{B}\mathbf{L}_t(T))' + \mathbf{W} \tag{10}$$

$$\mathbf{W}_{t|t,T}^{-1} = \mathbf{W}_{t|t-1,T}^{-1} + \sum_{c=1}^{C} \left[\left(\frac{\partial \log \lambda^c}{\partial \mathbf{x}_t} \right)' \left(\frac{\partial \log \lambda^c}{\partial \mathbf{x}_t} \right) \lambda^c \Delta \right.$$
$$\left. - (N_t^c - \lambda^c \Delta) \frac{\partial^2 \log \lambda^c}{\partial \mathbf{x}_t \partial \mathbf{x}_t'} \right]_{\mathbf{x}_{t|t-1,T}} \tag{11}$$

$$\mathbf{x}_{t|t,T} = \mathbf{x}_{t|t-1,T}$$
$$+ \mathbf{W}_{t|t,T} \sum_{c=1}^{C} \left[\left(\frac{\partial \log \lambda^c}{\partial \mathbf{x}_t} \right)' (N_t^c - \lambda^c \Delta) \right]_{\mathbf{x}_{t|t-1,T}} \tag{12}$$

This gives us the feedback-controlled point-process filter (FC-PPF). Since $\mathbf{B} = \mathbf{0}$ in (2) corresponds to the random-walk model, (9)-(12) with $\mathbf{B} = \mathbf{0}$ recover the random-walk point-process filter (RW-PPF).

5. ESTIMATING THE FEEDBACK-CONTROLLED SYSTEM WITH UNKNOWN DURATION FROM NEURAL OBSERVATIONS

The arrival time of the controlled system is not known to the external observer only observing the neural signal. We are hence interested in the unconditional posterior density,

522

$$p(\mathbf{x}_t|\mathbf{N}_{1:t}) = \sum_{j=1}^{J} p(\mathbf{x}_t|\mathbf{N}_{1:t}, T_j)p(T_j|\mathbf{N}_{1:t}) \qquad (13)$$

where we discretize the arrival time T to J possibilities and consequently place a prior model on it given by $p_T(T_j), j = 1, \cdots, J$. This prior model (including its support) can be selected based on empirical durations observed in a given task. We need the weights

$$p(T_j|\mathbf{N}_{1:t}) = \frac{p(\mathbf{N}_{1:t}|T_j)p_T(T_j)}{p(\mathbf{N}_{1:t})} \qquad (14)$$

where $p(\mathbf{N}_{1:t})$ is independent of T_j and hence treated as a constant and $p(\mathbf{N}_{1:t}|T_j)$ represents the likelihood of the observed neural data under a state-space model with the arrival time of T_j. Hence it is the partition function for the posterior $p(\mathbf{x}_{1:t}|\mathbf{N}_{1:t}, T_j)$ and its exact computation requires an integration, which is computationally prohibitive. However, using the Gaussian approximation to the posterior, we can find this without integration as follows. Using the chain rule we have

$$p(\mathbf{N}_{1:t}|T_j) = \prod_{i=1}^{t} P(\mathbf{N}_i|\mathbf{N}_{1:i-1}, T_j) = \prod_{i=1}^{t} g(\mathbf{N}_i|T_j) \quad (15)$$

where $g(\mathbf{N}_i|T_j)$ is defined in (7) and is the ith step partition function in the recursive filter. Now exploiting the Gaussian approximation of the posterior and hence the prediction densities in (6), they are completely characterized by their means and covariances given in (9)-(12) for a given T_j. We can hence explicitly evaluate (6) at $\mathbf{x}_{i|i,T_j}$ to get

$$g(\mathbf{N}_i|T_j) = \sqrt{\frac{|\mathbf{W}_{i|i,T_j}|}{|\mathbf{W}_{i|i-1,T_j}|}} \, p(\mathbf{N}_i|\mathbf{x}_{i|i,T_j}, \mathbf{N}_{1:i-1}) \times$$

$$\exp\left[-\frac{1}{2}(\mathbf{x}_{i|i,T_j} - \mathbf{x}_{i|i-1,T_j})'\mathbf{W}_{i|i-1,T_j}^{-1}(\mathbf{x}_{i|i,T_j} - \mathbf{x}_{i|i-1,T_j})\right] \quad (16)$$

for $i = 1, \cdots, T_j$ and $j = 1, \cdots, J$ where all the quantities are known. Combining (13)-(16) gives the posterior. The MMSE estimate in the case of unknown arrival time is then given by $\mathbf{x}_{t|t} = E(\mathbf{x}_t|\mathbf{N}_{1:t}) = \sum_j p(T_j|\mathbf{N}_{1:t})\mathbf{x}_{t|t,T_j}$ where the summation is over all j for which $T_j > t$. The resulting recursive filter is shown in Fig. 1. We call this filter the feedback-controlled parallel PPF (FC-P-PPF): it consists of J parallel point-process filters, each not only calculating the MMSE estimate of $\mathbf{x}_t$ assuming a duration of T_j, but also the corresponding likelihood $p(\mathbf{N}_{1:t}|T_j)$. This approach can also be viewed as mixture modeling, a common framework in statistical inference for density estimation. This framework has also been used in [7] to combine empirically fitted state models to different targets.

6. STATE-SPACE MODEL FOR THE REACHING MOVEMENT

Having derived the decoder for a general controlled state-space model, we can specialize to different tasks by using
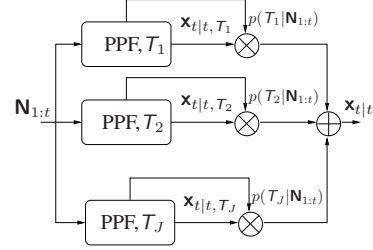

Fig. 1. Optimal feedback-controlled parallel point-process filter.

the suitable musculoskeletal state model and appropriate cost functions. For a reaching movement, the cost function should enforce end-point positional accuracy, stopping condition, and energetic efficiency. Denoting the desired final position by x^* and taking the state to be $\mathbf{x}_t = \begin{bmatrix} x_t & v_t & a_t \end{bmatrix}'$ where the components represent position, velocity and force respectively, similar to [6] we take this cost function to be

$$J = w_x(x_T - x^*)^2 + w_v v_T^2 + w_a a_T^2 + w_r \sum_{t=1}^{T} u_t^2 \quad (17)$$

The first 3 weights in the cost function are chosen to equally penalize these terms on average and the last weight to fit the biomechanical data. We adapt the first order lowpass muscle-like system in [6] for the dynamical system in (2),

$$\begin{bmatrix} x_{t+1} \\ v_{t+1} \\ a_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & \Delta & 0 \\ 0 & 1 - \frac{b\Delta}{m} & \frac{\Delta}{m} \\ 0 & 0 & 1 - \frac{\Delta}{\tau} \end{bmatrix} \begin{bmatrix} x_t \\ v_t \\ a_t \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{\Delta}{\tau} \end{bmatrix} u_t + \begin{bmatrix} 0 \\ 0 \\ w_t \end{bmatrix} \quad (18)$$

where the parameters b, τ, and m come from biomechanics given in [6] and the state noise is only non-zero in the force dimension. . The feedback matrices $\mathbf{L}_t(T_j)$ can now be easily found for all j from the recursive solution of LQG by either augmenting the state to include x^* or using a non-zero set point version of LQG.

7. RESULTS

Here we demonstrate the performance of our algorithm for both known and unknown arrival times and compare it with RW-PPF. Our data consisted of reaching trajectories performed by a rhesus monkey. The only parameter of (18) fitted to data was the state noise covariance $\mathbf{W}$, which represents the variability in reaching to the same goal over different trials. [3] We simulated the point-process neural signal of $C = 20$ neurons for a given trajectory using the time-rescaling theorem and taking the cosine tuning model of the conditional intensity function for a two-dimensional movement [1],

---

[3]We fitted two noise variances to the two-dimensional empirical trajectories for both random-walk and feedback-controlled state models: one parallel to the target direction and one perpendicular to it. Doing so in the case of the random-walk model gives it a substantial advantage as in general a random-walk model does not exploit the target information and hence should use the same noise variance in both directions. Hence the improvements shown here are only lower bounds on the true improvement one would get by using the target information and hence FC-P-PPF and FC-PPF over RW-PPF with no target information.

Table 1. RMS error (cm) in decoded trajectory of 55 real reaching movements.

Known T			Unknown T		
RW-PPF	FC-PPF	ratio	RW-PPF	FC-P-PPF	ratio
1.3225	0.8849	0.67	1.7106	1.0374	0.60

$$\lambda_t^c = e^{\alpha_0 + \alpha_1 |\mathbf{v}_t| \cos(\theta_t - \theta_p^c)} = e^{\alpha_0 + \alpha_y^c v_y(t) + \alpha_x^c v_x(t)}$$

where θ_p^c is the preferred angle of the cth neuron sampled randomly from $[-\pi, \pi]$, θ_t is the movement angle at t, and $v_y(t)$ and $v_x(t)$ are the velocities in the y and x directions. Here, α_0 and α_1 are chosen to have a background firing rate of 5 Hz and maximum firing rate of 30 Hz for each neuron.

We used 55 trajectories of different durations in the range of 150-400 ms from the real reaching movements and simulated 100 realizations (trials) of the point-process neural signal for each of them. Hence the window of uncertainty for the arrival time is 150-400 ms.

Fig. 2 shows one of these 55 trajectories and 15 sample decoded trajectories for both known and unknown durations. We can see that even for a known duration, the FC-PPF performs better than RW-PPF since it places a better prior on the state. With an unknown duration, the FC-P-PPF with only 4 parallel components, does better than RW-PPF for two reasons: 1. It places a better controlled prior on the states. 2. It combines the estimates of the parallel filters, each designed for a different arrival time, with time-varying optimal weights that are updated purely based on the neural observation. Fig. 3 shows these optimal weights in FC-P-PPF and the decoded $v_y(t)$. Here the true arrival time is at 300 ms and as we can see the corresponding calculated weight dominates up to 300 ms after which of course the only component left in the estimate is the 400 ms branch of the filter. Also, comparing the decoded v_y for FC-P-PPF and RW-PPF we can see that FC-P-PPF correctly brings the estimated velocity close to zero at the end of movement as opposed to RW-PPF. Note that the arrival time is discretized very coarsely and only 4 parallel branches are used.

Table 1 shows the RMS error in the decoded trajectory over the 55 real reaching movements. For known arrival time, FC-PPF reduces the RMS error by almost 33% and for unknown arrival time, FC-P-PPF reduces it by almost 40%.

8. REFERENCES

[1] W. Truccolo, U. Eden, M. Fellows, J. Donoghue, and E. N. Brown, "A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects," *J. Neurophysiol.*, pp. 1074–1089, 2005.

[2] L. Srinivasan, U. Eden, A. Willsky, and E. N. Brown, "A state-space analysis for reconstruction of goal-directed movements using neural signals," *Neural Comp.*, pp. 2465–2494, 2006.

[3] W. Wu, Y. Gao, E. Bienenstock, J. Donoghue, and M. Black, "Bayesian population decoding of motor cortical activity using a Kalman filter," *Neural Comp.*, pp. 80–118, 2005.
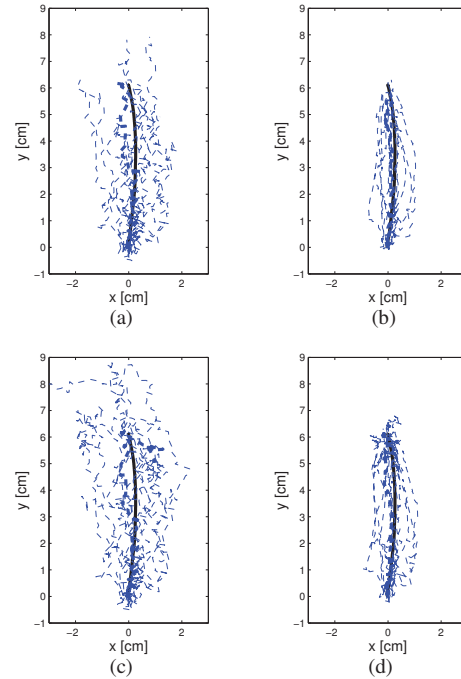
Fig. 2. True trajectory (solid) and 15 sample decoded ones (dashed) with known duration in (a) and (b) and unknown duration in (c) and (d). (a),(c) RW-PPF, (b) FC-PPF, (d) FC-P-PPF.
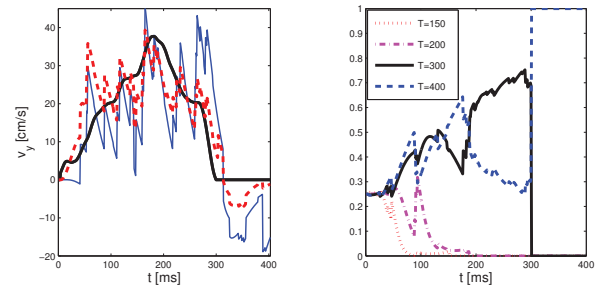

Fig. 3. Left: True v_y (solid) with unknown duration and sample decoded velocities with FC-P-PPF (thick dashed) and RW-PPF (thin solid). (b) Combining weights vs. time. True duration is 300 ms.

[4] U. Eden, L. Frank, R. Barbieri, V. Solo, and E. N. Brown, "Dynamic analysis of neural encoding by point process adaptive filtering," *Neural Comp.*, 2004.

[5] E. Todorov, "Optimality principles in sensorimotor control," *Nature Neuroscience*, pp. 907–915, Sep. 2004.

[6] D. Liu and E. Todorov, "Evidence for the flexible sensorimotor strategies predicted by optimal feedback control," *J. Neurosci.*, pp. 9354–9368, Aug. 2007.

[7] B. M. Yu *et al.*, "Mixture of trajectory models for neural decoding of goal-directed movements," *J. Neurophysiol.*, pp. 3763–3780, 2007.

[8] C. Kemere and T. Meng, "Optimal estimation of feed-forward-controlled linear systems," in *Proc. ICASSP*, Mar. 2005, pp. 353–356.

[9] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.

524