OPTIMAL ESTIMATION OF FEED-FORWARD-CONTROLLED LINEAR SYSTEMS

Caleb Kemere and Teresa Meng

Stanford University
Dept. of Electrical Engineering
Stanford, CA

ABSTRACT

The neuroprosthetic interface must infer an intended movement from the neural activity that would accompany it in healthy individuals. We show that an optimal estimator for a controlled system such as that responsible for human movements jointly estimates the goal and the trajectory of point-to-point movements. We demonstrate that this paradigm can achieve orders of magnitude of increased accuracy in regimes in which the interface has low SNR. With high SNR, our technique proves reliably more accurate than a typical approach which ignores the controlled nature of the system under observation. Furthermore, we show that even when the system violates the model assumptions of feed-forward linear control with additive noise, system performance remains appreciably better than the alternative.

1. INTRODUCTION

In the past decade there has been a dramatic rise in the interest given to the problem of *neuroprosthetics*. For the patient with a neurodegenerative disease, spinal cord injury, or perhaps even a missing limb, such a device would bypass the missing or damaged neural circuitry that normally transmits control signals from the cortex to the limbs. Ideally, the neuroprosthetic would decipher cortical neural activity so seamlessly that patients would be able to generate natural patterns of activity and have them correspond to natural-seeming movements. Unfortunately, though micromachining techniques have led to arrays of more than 100 electrodes, one significant roadblock is the limited quantity of neural information, and its degradation over time. Thus, we would like to maximize our ability to estimate the patient's intended movements given limited neural information.

For the class of human movement we consider of primary importance – point-to-point reaches – an optimal decoder will not only estimate the trajectory of the movement, but rather jointly estimate the trajectory and the target. We have previously shown that this approach leads to significant performance increases [1], here, we will demonstrate the optimality of the technique when certain assumptions are met.

Recent analyses of human movement suggest that the human motor control system may behave as an optimal controller [2]. Thus, one can approach the neuroprosthetic problem as an attempt to estimate the trajectory of a controlled dynamic system. Because this general problem appears novel, and may be of interest to a wider community, we will make a generic presentation of the approach, and then follow it with results from a simulation particular to the neuroprosthetic problem.

2. OPTIMAL ESTIMATION

We seek to estimate the trajectory of a system with controlled dynamics. This requires two models, a model of the system being tracked, and a model of the observation process. In general, these can be written as

$$\left.\begin{aligned} \mathbf{x}_{k+1} &= f\left(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k\right), \\ \mathbf{y}_k &= g\left(\mathbf{x}_k, \mathbf{v}_k\right), \\ \mathbf{u}_k &= \mathcal{L}_k\left(\hat{\mathbf{x}}_k\right) = \mathcal{L}_k\left(\mathbf{y}_k, \hat{\mathbf{x}}_{k-1}\right), \end{aligned}\right\} \quad \textit{Controlled System}$$

$$\mathbf{z}_k = h\left(\mathbf{x}_k, \mathbf{u}_k, \boldsymbol{\omega}_k\right), \qquad \textit{Observation}$$

where $\mathbf{x}_k$ is the observed system's state, $\mathbf{y}_k$ represents the system's internal observations of itself, $\mathbf{z}_k$ is the signal received by the external observer, $\mathbf{w}_k$, $\mathbf{v}_k$ and $\boldsymbol{\omega}_k$ are shorthand for stochastic processes, and $\mathbf{u}_k$ represents the control the system exerts on itself in order to minimize some cost function. Our problem, then, is to estimate the density

$$\Pr\left(\mathbf{x}_k \mid \mathbf{z}_0, \ldots, \mathbf{z}_k\right). \tag{1}$$

We now proceed to demonstrate that in the case of linear dynamics with Gaussian noise and feed-forward control, if the cost function minimized by the controller is quadratic in the target location, the mean of this density has a simple solution which involves joint estimation not only of the state $\mathbf{x}_k$, but also the target state.

This work was funded in part by C2S2, the MARCO Focus Center for Circuit & System Solution, under MARCO contract 2003-CT-888. Please address correspondences to ckemere@stanford.edu.

2.1. Simplified Model

We make the assumption that system has linear dynamics:

$$\begin{aligned}\mathbf{x}_{k+1} &= \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{C}\mathbf{w}_k, \text{ and} \\ \mathbf{y}_k &= \mathbf{G}\mathbf{x}_k + \mathbf{D}\mathbf{v}_k,\end{aligned}$$

where $\mathbf{w}_k$ and $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

We are interested in point-to-point movements, in which the system seeks to achieve a final state which is non-zero only in the position term. Thus, let us expand the state vector to include this desired position.

$$\mathbb{x}_k = \begin{bmatrix}\mathbf{x}_k \\ \mathbf{p}\end{bmatrix}$$

where $\mathbf{p}$ is the the target location. We can then express the cost function which is minimized by the controller in terms of this expanded state:

$$C = \mathbb{x}_f^T \mathbb{Q}_f \mathbb{x}_f + \sum_{k=0}^{f} \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k$$

where $\mathbb{Q} = \left[\begin{smallmatrix}\mathbf{1} & \mathbf{0} & -\mathbf{1} \\ \mathbf{0} & \mathbf{W} & \mathbf{0} \\ -\mathbf{1} & \mathbf{0} & \mathbf{1}\end{smallmatrix}\right]$ and $\mathbf{W}$ is a diagonal matrix representing the relative importance of zero final values for the non-positional terms of the state vector. The system dynamics in terms of the expanded state are

$$\mathbb{x}_{k+1} = \mathbb{A}\mathbb{x}_k + \mathbb{B}\mathbf{u}_k + \mathbb{C}\mathbf{w}_k \tag{2}$$

$$\mathbf{y}_k = \mathbb{G}\mathbb{x}_k + \mathbf{D}\mathbf{v}_k \tag{3}$$

where $\mathbb{A} = \left[\begin{smallmatrix}\mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}\end{smallmatrix}\right], \mathbb{B} = \left[\begin{smallmatrix}\mathbf{B} \\ \mathbf{0}\end{smallmatrix}\right], \mathbb{C} = \left[\begin{smallmatrix}\mathbf{C} \\ \mathbf{0}\end{smallmatrix}\right]$, and $\mathbb{G} = \left[\begin{smallmatrix}\mathbf{G} \\ \mathbf{0}\end{smallmatrix}\right]$.

We recognize this as the common linear, quadratic, Gaussian (LQG) control problem. Its solution has the characteristic that at any time the optimal control can be written as a linear feedback mechanism:

$$\mathbf{u}_k = \mathbf{L}_k \hat{\mathbb{x}}_k, \tag{4}$$

where $\hat{\mathbb{x}}_k$ is the system's current estimate of its state, given observations $\mathbf{y}_{0:k}$, and $\mathbf{L}_k$ is a matrix sequence that can be pre-calculated. In general the values of $\mathbf{L}_k$ will depend on $\mathbf{p}$; we make the assumption that for small variations of the target location, the same control sequence is used.

2.2. Feed-forward Control

If there is no feedback mechanism within the system, i.e., $\mathbb{G} = \mathbf{0}$, the optimal controller may still be written as above, with the following simplification of the state estimation

$$\hat{\mathbb{x}}_{k+1} = \mathbb{A}\hat{\mathbb{x}}_k + \mathbb{B}\mathbf{u}_k. \tag{5}$$

Using (4),

$$\hat{\mathbb{x}}_{k+1} = \prod_{m=0}^{k} \left(\mathbb{A} + \mathbb{B}\mathbf{L}_m\right) \mathbb{x}_0 \tag{6}$$

$$= \boldsymbol{\Upsilon}_{k+1} \mathbb{x}_0. \tag{7}$$

Let us assume that the initial state of the system is zero. Then, the controller can be rewritten as a linear function of the current state.

$$\begin{aligned}\mathbf{u}_k &= \mathbf{L}_k \boldsymbol{\Upsilon}_k \begin{bmatrix}\mathbf{0} \\ \mathbf{p}\end{bmatrix} \\ &= \mathbf{L}_k \boldsymbol{\Upsilon}_k^0 \mathbb{x}_k\end{aligned}$$

where $\boldsymbol{\Upsilon}_k^0 = \left[\cdots \mathbf{0} \cdots \mathbf{1}\right] \circ \boldsymbol{\Upsilon}_k$ and "$\circ$" denotes the Hadamard (or element-wise) product.

Thus, the system dynamics can be rewritten,

$$\begin{aligned}\mathbb{x}_{k+1} &= \left(\mathbb{A} + \mathbb{B}\mathbf{L}_k \boldsymbol{\Upsilon}_k\right) \mathbb{x}_k + \mathbb{C}\mathbf{w}_k \\ &= \bar{\mathbb{A}}_k \mathbb{x}_k + \mathbb{C}\mathbf{w}_k.\end{aligned} \tag{8}$$

2.3. Optimized Tracking

Now, we will further assume that our observation process is linear in the state of the system of (8), that is

$$\mathbf{z}_k = \mathbf{H}\mathbb{x}_k + \mathbf{E}\boldsymbol{\omega}_k, \tag{9}$$

where $\boldsymbol{\omega}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. We will assume that the last columns of $\mathbf{H}$ are zero, implying that the target of the movement is not directly observed. Furthermore, let us make the simplifying assumption that the control is optimized over a family of potential movements with varying targets centered on some known average point.

Given (8) and (9), the estimator distribution, (1), is Gaussian distributed, and its mean, the MMSE optimal estimator for $\mathbb{x}_k$ given the observations $\mathbf{z}_{0:k}$ is a Kalman filter:

$$\bar{\mathbb{x}}_{k+1} = \bar{\mathbb{A}}_k \bar{\mathbb{x}}_k + \mathbf{K}_k \left(\mathbf{z}_k - \mathbf{H}\bar{\mathbb{x}}_k\right), \tag{10}$$

where one solution for the filter coefficients, $\mathbf{K}_k$ is the solution to the recursion

$$\begin{aligned}\mathbf{K}_k &= \left(\bar{\mathbb{A}}_k \mathbf{P}_k \mathbf{H}^T + \mathbb{C}\mathbf{E}^T\right)\left(\mathbf{H}\mathbf{P}_k\mathbf{H}^T + \mathbf{E}\mathbf{E}^T\right)^{-1} \\ \mathbf{P}_{k+1} &= \bar{\mathbb{A}}_k \mathbf{P}_k \bar{\mathbb{A}}_k^T + \mathbb{C}\mathbb{C}^T - \mathbf{K}_k\left(\bar{\mathbb{A}}_k \mathbf{P}_k \mathbf{H}^T + \mathbb{C}\mathbf{E}^T\right)^T.\end{aligned}$$

The inversion in the standard form of the recursion given above renders it vulnerable to numerical instability. As a result, for the simulations which follow, we used a more stable "square-root" variant [3].

It is important to keep in mind that the filter we present is optimal only because it is jointly estimating the target and the trajectory of the movement. Thus, what is remarkable about this formulation is that our estimate of the target of

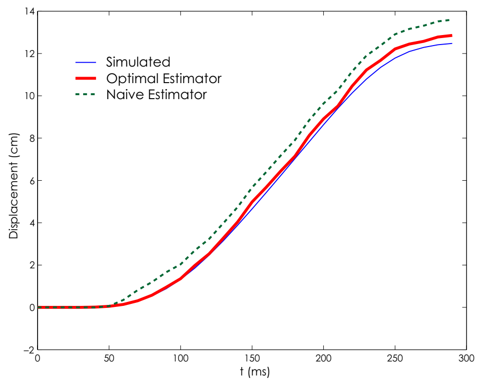

Fig. 1: A simulated arm trajectory reconstructed using the optimal joint estimator compared with that achieved by a naive approach. Model system was identical to the additive noise model of [4]. Observation SNR was $\sim$ 10dB.

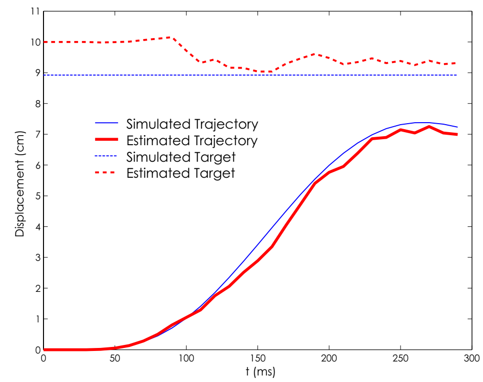

Fig. 2: A simulated arm trajectory reconstructed using the optimal estimator showing the time evolution of the target estimate. System and observation models as in Fig.1, but with feed-forward control only.

a given movement improves in accuracy over the course of the movement, which in turn increases the accuracy of the tracking of the instantaneous trajectory. To demonstrate what we will call the "optimal joint estimator," we used a model of human movements to simulate observations of a controlled linear system, and in the next section present the results of these simulations under different conditions.

3. SIMULATION RESULTS

For simulation purposes, we used the model of one-dimensional human reaching movements presented in [4] (except that the control cost term was decreased from 1e-5 to 1e-10). In contrast to Section 2.1, this model further encompasses feedback-control with signal-dependent noise – we simulated both with and without these assumption-violating enhancements. We assumed that the target was chosen from a Gaussian distribution centered around 10 cm displacement, with standard deviation 1 cm. A common model for the activity of neurons in the arm-areas of the motor cortex has their rate of activation vary linearly with hand velocity. Thus, the observations are taken as the noise-corrupted velocity (i.e., in (9), $\mathbf{H} = [0\ 1\ 0\ 0\ 0]$).

For comparative purposes, let us consider an alternative approach. A naive observer, ignoring the fact that the system under observation was being intelligently controlled, might also attempt to use a Kalman filter, treating the control signals as noise, and increasing the noise variance system parameter to account for the consequent excess variation in the system (see, e.g., [5]). Fig. 1 compares a sample trajectory generated by the model system with reconstructions from the optimal joint estimator and this naive approach.

Fig.2 depicts another trajectory reconstruction. The dotted and dashed lines reflect the novel aspect of our technique. The dotted line is the target of the movement while the dashed line is the estimate of this target. Notice how as the movement progresses, this estimate transitions from its initial value, the mean of the target distribution, towards the correct value.

The relative utility of optimal tracking depends on the relationship between the quality of the observation process and the parameters of the system being tracked. Fig.3 depicts the relationship between the average observation signal-to-noise ratio (SNR) and the average accuracy of trajectory reconstructions (mean-square error of position). Results using the naive estimator are included for comparison. As would be expected, as the SNR of the observed signal increases, the performance of the naive and optimal strategies begin to converge. With medium to low SNR, the optimal joint estimator begins to perform significantly better than the naive alternative. As the SNR decreases, the average error of the optimal approach tends toward the bound corresponding to simply choosing the mean trajectory. Thus, the actual value of the error bound is determined by the distribution of targets and the internal system noise. The thickening lines show that as the internal system noise decreases, the error bound also decreases. The lowest line shows the error bound in the case of zero target variation and very low internal noise.

In Fig. 4 we depict the effect of violations of the model assumptions. In addition to the pure feed-forward system, we simulated the actual feedback controlled models from [4], with either additive or multiplicative internal noise. Interestingly, even though the optimal joint estimator assumes purely feed-forward control, the performance benefits over the naive method are still quite large under feedback-control with additive system noise. Even in the case of the full human-like model with feedback control and multiplicative noise, if the system model of the optimal joint estimator is modified by choosing an appropriate value for the noise term ($\mathbb{C}$ in (8)), the performance increase remains signifi-

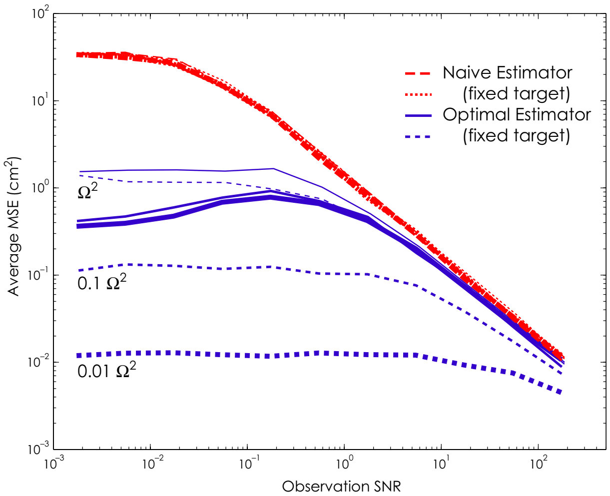

Fig. 3: Reconstruction accuracy and SNR. Increasing line widths correspond to decreasing internal system variation, where the initial value, Ω^2, is that specifi ed in [4]. To isolate the effect of internal system noise, the dashed lines correspond to performance in the case where the movement target is fi xed at 10 cm.

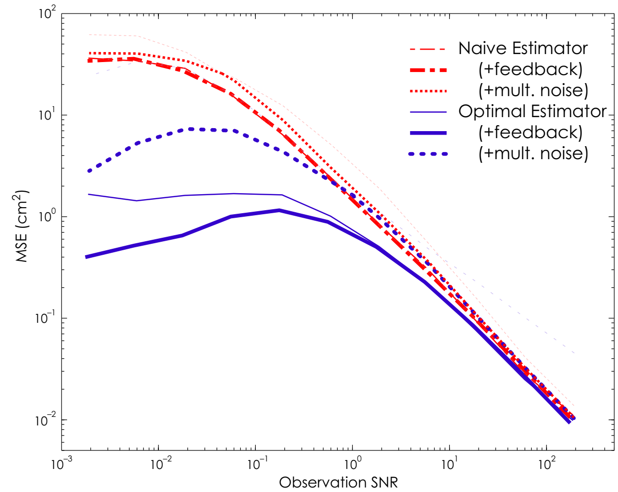

Fig. 4: Reconstruction accuracy and SNR in simulated systems where model assumptions are violated. In the case of multiplicative noise, the best parameter choices were highly dependent on SNR for both the optimal and naive approaches – the lower bounds are depicted.

cant except at the highest levels of SNR.

4. DISCUSSION

We have presented an approach for tracking controlled dynamic systems which is optimal under the conditions presented. For a model of the human motor system – of particular interest to us – we have shown in simulation that even if the system under observation is actually employing feedback control our estimator can outperform the paradigm which ignores the controlled aspect of the movements. Furthermore, even when the noise process within the system is radically different, multiplicative rather than additive, we found that performance remains excellent. While not presented here, we have used a similar tracking strategy elsewhere [1] to show that if information is present from some other source regarding the target of the movement – corresponding to the $\mathbf{H}$ matrix in (9) having nonzero values in the final columns – we can use it to further improve the estimator accuracy.

We would like to extend this result in two key dimensions. First, to enhance the generality of the joint estimation approach, we are pursuing an analytical solution for observing a feedback controlled system. Furthermore, we assumed a single control for a family of target locations; it would be better to allow for multiple task- or target- dependent controllers. Though this appears to result in the loss of analytical tractability, there is significant potential in various Monte-Carlo techniques. Secondly, to increase the utility of this type of optimal estimation in the particular application of neuroprosthetics, as neural signals are both non-linear and non-Gaussian distributed, techniques for joint system model and observation process identification are needed.

In regimes of lower SNR, the optimal joint estimation strategy becomes quite useful – the trajectories it produces are orders of magnitude more accurate than those of a more naive approach. Thus, while the particular application of neuroprosthetics is a quite interesting and important one, we are excited that this paradigm may find uses in a wide range of areas.

5. REFERENCES

[1] C. Kemere, , G. Santhanam, B. M. Yu, S. I. Ryu, T. H. Meng, and K. V. Shenoy, 'Model-based decoding of reaching movements for prosthetic systems," in *Proc. of the 26th IEEE EMBS Ann. Conf.*, San Francisco, CA, Sept. 2004, pp. 4524–4528.

[2] S. H. Scott, 'Optimal feedback control and the neural basis of volitional motor control," *Nature Reviews Neuroscience*, vol. 5, no. 7, pp. 532–546, July 2004.

[3] M. Morf and T. Kailath, 'Square-root algorithms for least-squares estimation," *IEEE Trans. Automat. Contr.*, vol. 20, no. 4, pp. 487–497, Aug. 1975.

[4] E. Todorov, 'Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system," *Neural Computation*, submitted for publication.

[5] W. Wu, M. Black, Y. Gao, E. Bienenstock, M. Serruya, A. Shaikhouni, and J. Donoghue, 'Neural decoding of cursor motion using a Kalman fi lter," in *in Neural Info. Proc. Sys., 15, MIT Press*, 2003.