

A State-Space Analysis for Reconstruction of Goal-Directed Movements Using Neural Signals

Lakshminarayan Srinivasan

ls2@neurostat.mgh.harvard.edu

Neuroscience Statistics Research Laboratory, Department of Anesthesia and Critical Care, Massachusetts General Hospital, Charlestown, MA 02129, and Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, and Division of Health Sciences and Technology, 77 Massachusetts Avenue, Massachusetts Institute of Technology, Cambridge, MA 02139, and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, U.S.A.

Uri T. Eden

tzvi@neurostat.mgh.harvard.edu

Neuroscience Statistics Research Laboratory, Department of Anesthesia and Critical Care, Massachusetts General Hospital, Charlestown, MA 02129, and Harvard/MIT Division of Health Sciences and Technology, Cambridge, MA 02139, U.S.A.

Alan S. Willsky

willsky@mit.edu

Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

Emery N. Brown

brown@neurostat.mgh.harvard.edu

Neuroscience Statistics Research Laboratory, Department of Anesthesia and Critical Care, Massachusetts General Hospital, Charlestown, MA 02129, and Division of Health Sciences and Technology, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, U.S.A.

The execution of reaching movements involves the coordinated activity of multiple brain regions that relate variously to the desired target and a path of arm states to achieve that target. These arm states may represent positions, velocities, torques, or other quantities. Estimation has been previously applied to neural activity in reconstructing the target separately from the path. However, the target and path are not independent.

Because arm movements are limited by finite muscle contractility, knowledge of the target constrains the path of states that leads to the target. In this letter, we derive and illustrate a state equation to capture this basic dependency between target and path. The solution is described for discrete-time linear systems and gaussian increments with known target arrival time. The resulting analysis enables the use of estimation to study how brain regions that relate variously to target and path together specify a trajectory. The corresponding reconstruction procedure may also be useful in brain-driven prosthetic devices to generate control signals for goal-directed movements.

1 Introduction

An arm reach can be described by a number of factors, including the desired hand target and the duration of the movement. We reach when moving to pick up the telephone or lift a glass of water. The duration of a reach can be specified explicitly (Todorov & Jordan, 2002) or emerge implicitly from additional constraints such as target accuracy (Harris & Wolpert, 1998). Arm kinematics and dynamics during reaching motion have been studied through their correlation with neural activity in related brain regions, including motor cortex (Moran & Schwartz, 1999), posterior parietal cortex (Andersen & Buneo, 2002), basal ganglia (Turner & Anderson, 1997), and cerebellum (Greger, Norris, & Thach, 2003). Separate studies have developed control models to describe the observed movements without regard to neural activity (Todorov, 2004a). An emerging area of interest is the fusion of these two approaches to evaluate neural activity in terms of the control of arm movement to target locations (Todorov, 2000; Kemere & Meng, 2005). While several brain areas have been implicated separately in the planning and execution of reaches, further study is necessary to elucidate how these regions coordinate their electrical activity to achieve the muscle activation required for reaching. In this letter, we develop state-space estimation to provide a unified framework to evaluate reach planning and execution-related activity.

Primate electrophysiology during reaching movements has focused on primary motor cortex (M1) and posterior parietal cortex, regions that represent elements of path and target, respectively. Lesion studies previously identified M1 with motor execution (Nudo, Wise, SiFuentes, & Milliken, 1996) and PPC with movement planning (Geschwind & Damasio, 1985). Several experiments have characterized the relationship between M1 neuronal activity, arm positions, and velocities (Georgopoulos, Kalaska, Caminiti, & Massey, 1982; Schwartz, 1992; Moran & Schwartz, 1999; Paninski, Fellows, Hatsopoulos, & Donoghue, 2004), and forces (Georgopoulos, Ashe, Smyrnis, & Taira, 1992; Taira, Boline, Smyrnis, Georgopoulos, & Ashe, 1995; Li, Padoa-Schioppa, & Bizzi, 2001). PPC is described as relating broadly to the formation of intent and specifically to the transformation of sensory cues into movement goals (Andersen &

Buneo, 2002). More recent experiments are beginning to elucidate the role of premotor cortical areas in motion planning and execution (Schwartz, Moran, & Reina, 2004), including interactions with PPC (Wise, Boussaoud, Johnson, & Caminiti, 1997). Explicit regression analyses have also been performed to relate motor cortical activity to features of both target and path (Fu, Suarez, & Ebner, 1993; Ashe & Georgopoulos, 1994).

In parallel, theoretical models for the planning and execution of reaches have developed to include different concepts in control engineering and robotics. A common starting point is the state equation, a differential equation that describes how the arm moves due to passive sources like joint tension and user-controlled forces such as muscle activation. The state equation is used to prescribe a path or a sequence of forces to complete the reach based on the minimization of some cost function that depends on variables such as energy, accuracy, or time. Many reach models specify control sequences computed prior to movement that assume a noise-free state equation and perfect observations of arm state (Hogan, 1984; Uno, Kawato, & Suzuki, 1989; Nakano et al., 1999). The execution of trajectories planned by these models can be envisioned in the face of random perturbations by equilibrium-point control, where each prescribed point in the trajectory is sequentially made steady with arm tension. Recently, reach models have been developed that explicitly account for noisy dynamics and observations (Harris & Wolpert, 1998; Todorov, 2004b). Based on stochastic optimal control theory, the most recent arm models (Todorov & Jordan, 2002; Todorov, 2004b) choose control forces based on estimates of path history and cost-to-go, the price associated with various ways of completing the reach. A general review of control-based models is provided in Todorov (2004a).

Estimation has been used to relate neural activity with aspects of free arm movements (Georgopoulos, Kettner, & Schwartz, 1988; Paninski et al., 2004). Alternate models of neural response in a specific brain region can be compared by mean squared error (MSE). Reconstruction of a measured parameter is one way to characterize neural activity in a brain region. Learning rates can be related explicitly and simultaneously to continuous and discrete behavioral responses using an estimation framework (Smith et al., 2004). Mutual information is a related alternative that has been prevalent in the characterization of neural responses to sensory stimuli (Warland, Reinagel, & Meister, 1997). Both MSE and conditional entropy (calculated in determining mutual information) are functions of the uncertainty in an estimate given neural observations, and MSE rises with conditional entropy for gaussian distributions. These two methods were recently coupled to calculate the conditional entropy associated with recursively computed estimates on neural data (Barbieri et al., 2004).

Estimation algorithms form the interface between brain and machine in the control of neural prosthetics, bearing directly on the clinical treatment of patients with motor deficits. Prototype systems have employed either estimation of free arm movement (Carmena et al., 2003; Taylor, Tillery, & Schwartz, 2002; Wu, Shaikhouni, Donoghue, & Black, 2004) or

target location (Musallam, Corneil, Greger, Scherberger, & Andersen, 2004; Santhanam, Ryu, Yu, Afshar, & Shenoy, 2005). Most recently, several estimation procedures were proposed to combine these two approaches and specifically facilitate reaching movements for brain-controlled prosthetics (Srinivasan, Eden, Willsky, & Brown, 2005; Cowan & Taylor, 2005; Yu, Santhanam, Ryu, & Shenoy, 2005; Kemere & Meng, 2005).

Two probability densities are used implicitly in estimation. The first density describes the probability of neural activity conditioned on relevant covariates like stimulus intensities or kinematic variables. This density arises through the observation equation in estimation and as an explicit function in information-theoretic measurements. The second density describes the interdependence of the relevant covariates before any neural activity is recorded. This density arises through the state equation in estimation and as a prior on stimulus values in the information-theoretic characterization of sensory neurons. In experiments that calculate mutual information between neural activity and independent stimulus parameters, this second probability density is commonly chosen to be uniform. In the study of reaching movements, the complete prior density on target and path variables cannot be uniform because the target and the path state at all times in the trajectory are dependent. A state equation naturally expresses these constraints and serves as a point of departure for analysis based on estimation.

In this letter, we develop a discrete-time state equation that relates target state and path states under weak assumptions about a reach. Specifically, the result represents the extension of the linear state-space description of free arm movement with no additional constraints. The states of the target or path refer to any vector of measurements of the arm at a particular point in time, such as joint torque, joint angle, hand velocity, and elbow position. This method supports arbitrary order, time-varying linear difference equations, which can be used to approximate more complicated state equation dynamics. The approach is based on the continuous-time results by Castanon, Levy, and Willsky (1985) in surveillance theory and draws on the discrete time derivation of a backward Markov process described by Verghese and Kailath (1979). Unlike existing theoretical models of reaching movement, we do not begin with an assumed control model or employ cost functions to constrain a motion to target. The resulting reach state equation is a probabilistic description of all paths of a particular temporal duration that start and end at states specified with uncertainty.

We first develop a form of the reach state equation that incorporates one prescient observation on the target state. We then extend this result to describe an augmented state equation that includes the target state itself. This augmented state equation supports recursive estimates of path and target that fully integrate ongoing neural observations of path and target. Sample trajectories from the reach state equation are shown. We then demonstrate the estimation of reaching movements by incorporating the reach state equation into a point process filter (Eden, Frank, Barbieri, Solo, & Brown,

2004). We conclude by discussing the applicability of our approach to the study of motion planning and execution, as well as to the control of neural prosthetics.

2 Theory

2.1 State Equation to Support Observations of Target Before Movement. The objective in this section is to construct a state equation for reaching motions that combines one observation of the target before movement with a general linear state equation for free arm movement. The resulting state equation enables estimation of the arm path that is informed by concurrent observations and one target-predictive observation, such as neural activity from brain regions related to movement execution and target planning, respectively. We begin with definitions and proceed with the derivation.

A reach of duration T time steps is defined as a sequence of vector random variables (x_0, \dots, x_T) called a trajectory. The state variable x_t represents any relevant aspects of the arm at time sample t , such as position, velocity, and joint torque. The target x_T is the final state in the trajectory. While we conventionally think of a target as a resting position for the arm, x_T more generally represents any condition on the arm at time T , such as movement drawn from a particular probability distribution of velocities.

For simplicity, we restrict our trajectory to be a Gauss-Markov process. This means that the probability density on the trajectory $p(x_0, \dots, x_T)$ is jointly gaussian and that the probability density of the state at time t conditioned on all previous states $p(x_t|x_0, \dots, x_{t-1})$ equals $p(x_t|x_{t-1})$, the state transition density. Although more general probability densities might be considered, these special restrictions are sufficient to allow for history dependency of arbitrary length. This is commonly accomplished by including the state at previous points in time in an augmented state vector (Kailath, Sayed, & Hassibi, 2000). Figure 1A is a schematic representation of the trajectory and the target observation, emphasizing that the prescient observation of target y_T is related to the trajectory states x_t only through the target state x_T .

The conditional densities of the Gauss-Markov model can alternatively be specified with observation and state equations. For a free arm movement, the state transition density $p(x_t|x_{t-1})$ can be described by a generic linear time-varying multidimensional state equation,

$$x_t = A_t x_{t-1} + w_t, \quad (2.1)$$

where the stochastic increment w_t is a zero-mean gaussian random variable with $E[w_t w_t'] = Q_t \delta_{t-\tau}$. The initial position x_0 is gaussian distributed with mean m_0 and covariance Π_0 . The prescient observation y_T of the target state x_T is corrupted by independent zero-mean gaussian noise v_T with

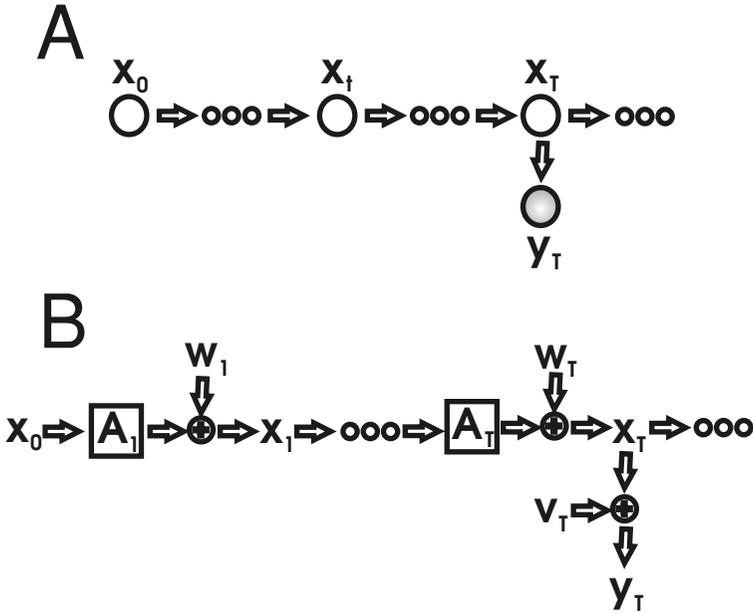


Figure 1: Alternate representations of a reaching trajectory and one observation on target. In the Markov model (A), circles represent the state of the arm at various times, and the arrangement of directed arrows indicates that the state of the arm at time t is independent of all previous states conditioned on knowledge of the state at time $t - 1$. Accordingly, the only state pointing to y_T , the prescient observation of target, is the target state x_T itself. In the system diagram (B), the specific evolution of the arm movement is described. Consistent with the state equation, the arm state x_{t-1} evolves to the next state in time x_t through the system matrix A_t , with additive noise w_t that represents additional uncertainty in aspects of the arm movement that are not explained by the system matrix. The diagram also specifies that the observation y_T of the target state x_T is corrupted by additive noise v_T .

covariance Π_T that denotes the uncertainty in target position:

$$y_T = x_T + v_T. \tag{2.2}$$

The state equation coupled with this prescient observation is described schematically in Figure 1B.

Restated, our objective is to represent the free movement state equation together with the prescient observation on target, as a Gauss-Markov model on an equivalent set of trajectory states x_t conditioned on y_T for

$t = 0, 1, \dots, T$. The consequent reach state equation is of the form

$$x_t = A_t x_{t-1} + u_t + \varepsilon_t, \tag{2.3}$$

where u_t is a drift term corresponding to the expected value of $w_t|x_{t-1}, y_T$, and the ε_t are a new set of independent, zero-mean gaussian increments whose covariances correspond to that of $w_t|x_{t-1}, y_T$. This reach state equation generates a new probability density on the trajectory of states that corresponds to the probability of the original states conditioned on the prescient observation, $p(x_0, \dots, x_T|y_T)$.

To derive this reach state equation, we calculate the state transition probability density $p(x_t|x_{t-1}, y_T)$. Because w_t is the only stochastic component of the original state equation, the new state transition density is specified by $p(w_t|x_{t-1}, y_T)$. To compute this distribution, we use the conditional density formula for jointly gaussian random variables on the joint density $p(w_t, y_T|x_{t-1})$. The resulting distribution is itself gaussian, with mean and variance given by:

$$\begin{aligned} u_t &= E[w_t|x_{t-1}, y_T] \\ &= E[w_t|x_{t-1}] + \text{cov}(w_t, y_T|x_{t-1}) \\ &\quad \times \text{cov}^{-1}(y_T, y_T|x_{t-1})(y_T - E[y_T|x_{t-1}]) \end{aligned} \tag{2.4}$$

$$\begin{aligned} \tilde{Q}_t &= \text{cov}(w_t|x_{t-1}, y_T) \\ &= \text{cov}(w_t|x_{t-1}) - \text{cov}(w_t, y_T|x_{t-1}) \\ &\quad \times \text{cov}^{-1}(y_T|x_{t-1})\text{cov}'(w_t, y_T|x_{t-1}). \end{aligned} \tag{2.5}$$

The mean u_t corresponds identically to the linear least-squares estimate of $w_t|x_{t-1}, y_T$, and the variance \tilde{Q} equals the uncertainty in this estimate.

The covariance terms in equations 2.4 and 2.5 can be computed from the following equation that relates w_t to y_T given x_{t-1} ,

$$y_T = \phi(T, t-1)x_{t-1} + \sum_{i=t}^T \phi(T, i)w_i + v_T, \tag{2.6}$$

where $\phi(t, s)$ denotes the state transition matrix that advances the state at time s to time t ,

$$\phi(t, s) = \begin{cases} \prod_{i=1+\min(t,s)}^{\max(t,s)} A_i^{\text{sign}(t-s)}, & t \neq s \\ \text{I}, & t = s \end{cases} . \tag{2.7}$$

The covariance terms are accordingly given by

$$\text{cov}(w_t | x_{t-1}) = Q_t \quad (2.8)$$

$$\text{cov}(w_t, y_T | x_{t-1}) = Q_t \phi'(T, t) \quad (2.9)$$

$$\text{cov}(y_T, y_T | x_{t-1}) = \Pi_T + \sum_{i=t}^T \phi(T, i) Q_i \phi'(T, i). \quad (2.10)$$

For notational convenience, define the following quantity:

$$\Pi(t, T) = \phi(t, T) \Pi_T \phi'(t, T) + \sum_{i=t}^T \phi(t, i) Q_i \phi'(t, i). \quad (2.11)$$

Simplifying and substituting into equations 2.4 and 2.5, we obtain the mean and covariance of the old increment given the target observation:

$$u_t = Q_t \Pi^{-1}(t, T) \phi(t, T) [y_T - \phi(T, t-1) x_{t-1}] \quad (2.12)$$

$$\tilde{Q}_t = Q_t - Q_t \Pi^{-1}(t, T) Q_t'. \quad (2.13)$$

The density on the initial state conditioned on the target observation is calculated similarly. The resulting mean and variance of the initial state are given by

$$\tilde{\Pi}_0 = (\Pi_0^{-1} + \Pi^{-1}(0, T))^{-1} \quad (2.14)$$

$$E[x_0 | y_T] = \tilde{\Pi}_0(0) (\Pi_0^{-1} m_0 + \Pi^{-1}(0, T) \phi(0, T) y_T). \quad (2.15)$$

A recursion can be obtained for equation 2.11 by writing $\Pi(t-1, T)$ in terms of $\Pi(t, T)$:

$$\Pi(t-1, T) = \phi(t-1, t) \Pi(t, T) \phi'(t-1, t) + \phi(t-1, t) Q_{t-1} \phi'(t-1, t) \quad (2.16)$$

with

$$\Pi(T, T) = \Pi_T + Q_T. \quad (2.17)$$

Complementing the new initial conditions 2.14 and 2.15, the reach state equation can be written in various equivalent forms. The following form emphasizes that the old increment w_t has been broken into the estimate u_t of w_t from y_T and remaining uncertainty ε_t ,

$$x_t = A_t x_{t-1} + u_t + \varepsilon_t \quad (2.18)$$

$$\varepsilon_t \sim N(0, \tilde{Q}_t) \quad (2.19)$$

with u_t as given in equation 2.12 and ε_t distributed as a zero-mean gaussian with covariance \hat{Q}_t .

This form is suggestive of stochastic control, where u_t is the control input that examines the state at time x_{t-1} , and generates a force to place the trajectory on track to meet the observed target. Nevertheless, this form emerges purely from conditioning the free movement state equation on the target observation rather than from any specific biological modeling of motor control. Note critically that u_t is a function of x_{t-1} , so that the covariance update in a Kalman filter implementation should not ignore this term.

Alternatively, we can group the x_{t-1} terms. This form is more conducive to the standard equations for the Kalman filter prediction step:

$$x_t = B_t x_{t-1} + f_t + \varepsilon_t \tag{2.20}$$

$$B_t = [I - Q_t \Pi^{-1}(t, T)] A_t \tag{2.21}$$

$$f_t = Q_t \Pi^{-1}(t, T) \phi(t, T) y_T. \tag{2.22}$$

In both forms, the resulting reach state equation remains linear with independent gaussian errors ε_t , as detailed in the appendix. Because x_t is otherwise dependent on x_{t-1} or constants, we conclude that the reach state equation in 2.18 or 2.20 is a Markov process.

2.2 Augmented State Equation to Support Concurrent Estimation of Target. Building on the previous result, we can now construct a more versatile state equation that supports path and target estimation with concurrent observations of path and target. The previous reach state equation incorporates prescient target information into a space of current arm state x_t . We now augment the state space to include the target random variable x_T . According to this model, the state of the arm at time t is explicitly determined by the target and the state of the arm at time $t - 1$.

The reach state equation derived above suggests an approach to calculating the state transition density $p(x_t, x_T | x_{t-1}, x_T)$ that corresponds to an augmented state equation. Because x_T is trivially independent of x_t conditioned on x_T , we can equivalently calculate the transition density of $p(x_t | x_{t-1}, x_T)$. This is identical to the reach state equation derivation of $p(x_t | x_{t-1}, y_T)$ with v_T set to zero. The resulting state equation can be consolidated into vector notation to give the augmented form:

$$\begin{pmatrix} x_t \\ x_T \end{pmatrix} = \begin{pmatrix} \Psi & \Gamma \\ 0 & I \end{pmatrix} \begin{pmatrix} x_{t-1} \\ x_T \end{pmatrix} + \begin{pmatrix} \varepsilon_t \\ 0 \end{pmatrix} \tag{2.23}$$

$$\Psi = B_t \tag{2.24}$$

$$\Gamma = Q_t \Pi^{-1}(t, T) \phi(t, T) \tag{2.25}$$

$$\Pi_T = 0. \tag{2.26}$$

The initial condition on the augmented state $[x_0, x_T]'$ is the joint distribution that corresponds to our uncertainty as external observers about the true starting and target states chosen by the brain at time zero.

This augmented state equation confers additional features over the reach state equation. First, observations of the target can be incorporated throughout the duration of the reach to improve arm reconstructions. In contrast, the reach state equation incorporated one target observation before movement. Second, refined estimates of the target can be generated recursively as estimates become more informed by reach and target-related activity.

3 Results

3.1 Sample Trajectories. We proceed to illustrate the underlying structure of a reach for our goal-directed state equation, which appropriately constrains a general linear state equation to an uncertain target. We also explain how the underlying reach structure is affected by parameters of the model: reach duration, the target state observation, and target uncertainty.

The density on the set of trajectories, $p(x_t, x_{t-1}, \dots, x_0 | y_T)$, can be calculated by iteratively multiplying the transition densities $p(x_t | x_{t-1}, y_T)$ given by the state equation. This density represents our assumptions about the trajectory before receiving additional observations of neural activity during the reach. Broader probability densities on the set of trajectories imply weaker assumptions about the specific path to be executed.

We can visually examine the structure of our assumptions by plotting samples from the density on trajectories as well as the average trajectory. Sample trajectories are generated by drawing increments ε_t from the density specified in equation 2.19.

The simulated increments are accumulated at each step with $A_t x_t + u_t$, the deterministic component of the state equation 2.18. The resulting trajectory represents a sample drawn from $p(x_t, x_{t-1}, \dots, x_0 | y_T)$, the probability density on trajectories. The average trajectory is generated from the same procedure, except that the increments ε_t are set to their means, which equal zero.

We first examine sample trajectories that result from small changes in model parameters. For illustration, the states were taken to be vectors $[x, y, v_x, v_y]_t'$, representing position and velocity in each of two orthogonal directions. The original noise covariance was nonzero in the entries corresponding to velocity increment variances:

$$Q = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & q & 0 \\ 0 & 0 & 0 & q \end{pmatrix}. \quad (3.1)$$

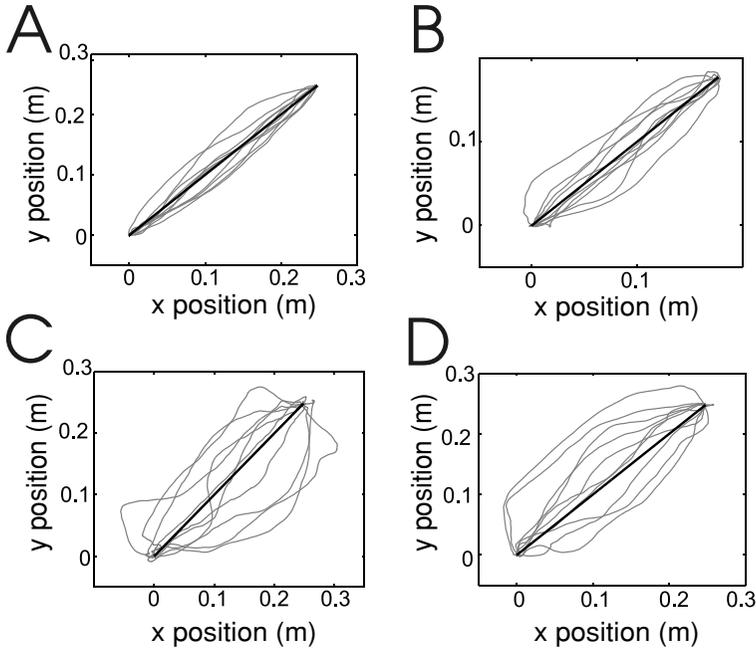


Figure 2: Sample trajectories (gray) and the true mean trajectory (black) corresponding to the reach state equation for various parameter choices. Appropriate changes in model parameters increase the observed diversity of trajectories, making the state equation a more flexible prior in reconstructing arm movements from neural signals. Parameter choices (detailed in section 3.1) were varied from (A) baseline, including (B) smaller distance to target, (C) increased time to target, and (D) increased increment uncertainty.

The uncertainty in target state Π_T was also diagonal, with

$$\Pi_T = \begin{pmatrix} r & 0 & 0 & 0 \\ 0 & r & 0 & 0 \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{pmatrix}. \tag{3.2}$$

In Figure 2, sample trajectories from the reach state equation are generated with baseline parameters (see Figure 2A) from which distance to target, reach duration, and increment uncertainty have been individually changed (see Figures 2B–2D). The baseline model parameters are given in Table 1. Parameters were individually altered from baseline as shown in Table 2.

Table 1: Sample Trajectory: Baseline Model Parameters.

Parameter	Baseline Value
Reach distance	0.35 m
Time step	0.01 sec.
Noise covariance (q)	1e-4 m ²
Reach duration	2 sec.
Target position uncertainty (r)	1e-6 m ²
Target velocity uncertainty (p)	1e-6 m ²

Table 2: Sample Position Trajectory: Altered Model Parameters.

Parameter	Altered Value	Graph
Reach distance	0.25 m	Figure 2B
Reach duration	4 sec.	Figure 2C
Noise covariance (q)	3e-4 m ²	Figure 2D

In Figure 3, sample trajectories are plotted for increasing uncertainty (r) in target position, with variances (A) 1e-4, (B) 1e-3, (C) 1e-2, and (D) 1e-1 m². This corresponds to scenarios in which observations of neural activity before movement initiation provide estimates of target position with varying certainty.

Figures 4A to 4C examine the velocity profiles in one direction generated by the reach state equation with various parameter choices. Velocity profiles from the baseline trajectory are displayed in Figure 4A, and parameters are sequentially altered from the baseline values (see Figures 4B and 4C) as shown in Table 3. Figure 4D examines the effect of target information on uncertainty in the velocity increment. The magnitude of one diagonal velocity term of the noise covariance \tilde{Q}_t is plotted over the duration of the reach for comparison against the noise covariance Q_t of the corresponding free movement state equation.

3.2 Reconstructing Arm Movements During a Reach. The reach state equation can be incorporated into any estimation procedure based on probabilistic inference since it represents a recursively computed prior. Because the reach state equation minimally constrains the path to the target observation, it may be useful in the analysis of coordinated neural activity with respect to planning and execution. We illustrate the reconstruction of reaching movements from simulated neural activity using a point process filter (Eden, Frank, et al., 2004), an estimation procedure that is conducive to the description of spiking activity in particular. The extension to variants of the Kalman filter is also direct, because the reach state equation, 2.20, is written in standard Kalman filter notation.

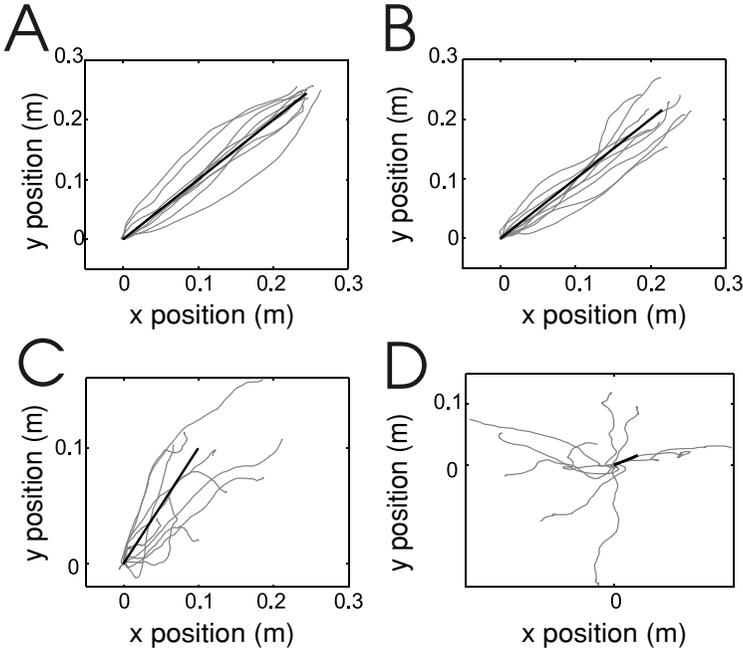


Figure 3: Sample trajectories (gray) and the true mean trajectory (black) of the reach state equation corresponding to various levels of uncertainty about target arm position. Variance in the noise v_T of the prescient observation y_T is progressively increased from (A) $1e-4$, to (B) $1e-3$, (C) $1e-2$, and (D) $1e-1$ m^2 . As target uncertainty grows, trajectories become more unrestricted, corresponding to increasing flexibility in the prior for reconstruction of arm movements.

We first simulated arm trajectories using the reach model as described in the previous section. For comparison, arm trajectories were also generated from a canonical model. This model was a family of movement profiles from which specific trajectories could be chosen that guaranteed arrival at the desired target location and time:

$$\begin{pmatrix} x \\ y \\ v_x \\ v_y \end{pmatrix}_t = \begin{pmatrix} 1 & 0 & \Delta & 0 \\ 0 & 1 & 0 & \Delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ v_x \\ v_y \end{pmatrix}_{t-1} + \left(\frac{1}{2\delta}\right) (\pi/T)^2 \cos(\pi t/T) \begin{pmatrix} 0 \\ 0 \\ x_T - x_0 \\ y_T - y_0 \end{pmatrix}. \tag{3.3}$$

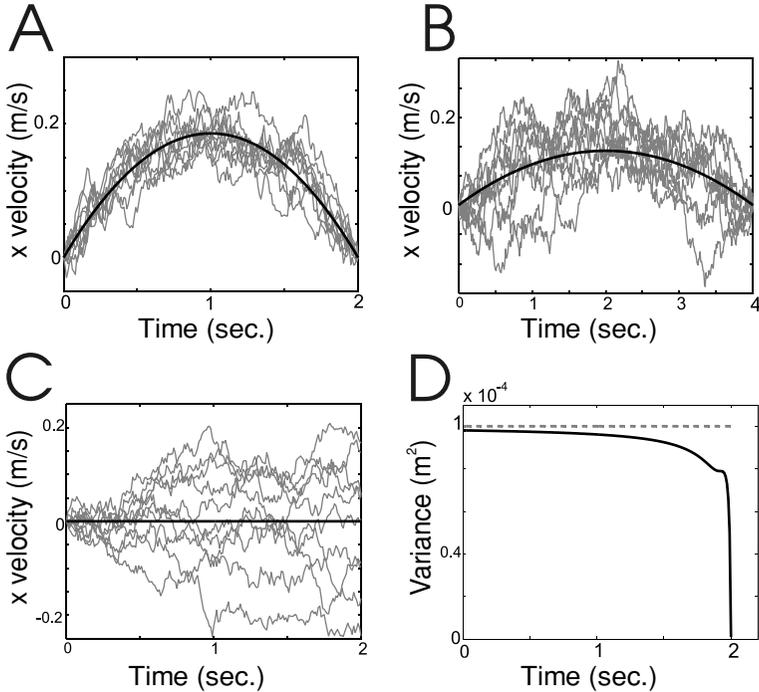


Figure 4: Sample velocity trajectories (gray) and the true mean velocity trajectory (black) generated by the reach state equation. (A) For baseline parameters (detailed in section 3.1) with reach duration of 2 seconds, the velocity profile is roughly bell shaped. (B) As reach duration increases to 4 seconds, the trajectories become more varied. (C) If uncertainty in the observed target velocity and position is large ($1e3 \text{ m}^2$ for each variance), velocity trajectories resemble samples from the free movement state equation. (D) Uncertainty in the velocity increment decreases with time due to the prescient target observation (solid line) as compared to the original velocity increment of the corresponding free movement state equation (dashed line).

This deterministic equation relates velocities $[x, y, v_x, v_y]_t$ to the time increment δ , the current time step t , and the distances in two orthogonal directions between the target and starting points, over $T + 1$ time steps.

After generating trajectories, we simulated the corresponding multiunit spiking activity from 9 neurons, a typical ensemble size for recording from a focal, single layer of cortex (Buzáki, 2004). Output from each unit in the ensemble was simulated independently as a point process with an instantaneous firing rate that was a function of the velocity. This function, referred to as the conditional intensity (Eden, Frank, et al., 2004), is equivalent to

Table 3: Sample Velocity Trajectory: Altered Model Parameters.

Parameter	Altered Value	Graph
Reach duration	4 sec.	Figure 4B
Target position and velocity uncertainty	$r = 1e3 \text{ m}^2$ $p = 1e3 \text{ m}^2$	Figure 4C

Table 4: Simulated M1 Activity: Receptive Field Parameters.

Parameter	Assignment or Interval
β_0	2.28
β_1	4.67 sec/m
θ_p	$[-\pi, \pi]$

specifying a receptive field. Our conditional intensity function is adapted from a model of primary motor cortex (Moran & Schwartz, 1999):

$$\lambda(t|v_x, v_y) = \exp(\beta_0 + \beta_1(v_x^2 + v_y^2)^{1/2} \cos(\theta - \theta_p)) \tag{3.4}$$

$$= \exp(\beta_0 + \alpha_1 v_x + \alpha_2 v_y), \tag{3.5}$$

where v_x and v_y are velocities in orthogonal directions.

The receptive field parameters were either directly assigned or drawn from uniform probability densities on specific intervals as shown in Table 4. The corresponding receptive fields had preferred directions between $-\pi$ and π , background firing rates of 10 spikes per second, and firing rates of 24.9 spikes per second at a speed of 0.2 m per second in the preferred direction.

Together with the simulated trajectory, this conditional intensity function specifies the instantaneous firing rate at each time sample based on current velocity. Spikes were then generated using the time rescaling theorem (Brown et al., 2002), where interspike intervals are first drawn from a single exponential distribution and then adjusted in proportion to the instantaneous firing rate. This method is an alternative to probabilistically thinning a homogeneous Poisson process.

The simulated spiking activity served as the input observations for the point process filter, described extensively in Eden, Frank, et al. (2004). The two defining elements of this filter are the state equation and observation equation. Our state equation is the reach model and represents the dynamics of the variables we are estimating, specified by $p(x_t|x_{t-1}, y)$. Our observation equation is the receptive field of each neuron, specified by $p(\Delta N_t|\Delta N_{1:t-1}, x_t, y_T)$. This is the probability of observing ΔN_t spikes at time t , given previous spike observations $\Delta N_{1:t-1}$, the current kinematic

state x_t , and the observation of target y_T . Because the spiking activity is described as a point process, the conditional intensity function specifies this observation density:

$$p(\Delta N_t | \Delta N_1, \dots, \Delta N_{t-1}, x_t, y_T) \approx \exp[\Delta N_t^i \log(\lambda(t|x_t)\delta) - \lambda(t|x_t)\delta], \quad (3.6)$$

where δ denotes the time increment.

The formulation of a recursive estimation procedure from these two probability densities is the topic of Eden, Frank, et al. (2004). As with the Kalman filter, the resulting point process filter comprises a prediction step to compute $p(x_t | \Delta N_{1:t-1}, y_T)$ and an update step to compute $p(x_t | \Delta N_{1:t}, y_T)$. The reach state equation determines the mean and variance prediction steps of the point process filter, as given by

$$\hat{x}_{t|t-1} = B_t \hat{x}_{t-1|t-1} + f_t \quad (3.7)$$

$$\Lambda_{t|t-1} = B_t \Lambda_{t-1|t-1} B_t' + \tilde{Q}_t. \quad (3.8)$$

The update step remains unchanged:

$$\begin{aligned} (\Lambda_{t|t})^{-1} = & \Lambda_{t|t-1}^{-1} + \left[\left(\frac{\partial \log \lambda}{\partial x_t} \right) [\lambda \Delta \delta_k] \left(\frac{\partial \log \lambda}{\partial x_t} \right) \right. \\ & \left. - (\Delta N_t - \lambda \Delta \delta_k) \frac{\partial^2 \log \lambda}{\partial x_t^2} \right]_{x_{t|t-1}} \end{aligned} \quad (3.9)$$

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + \Lambda_{t|t} \left[\left(\frac{\partial \log \lambda}{\partial x_t} \right)' (\Delta N_t - \lambda \Delta \delta_k) \right]_{x_{t|t-1}}. \quad (3.10)$$

We compared the quality of reconstruction using the reach state equation versus the standard free arm movement state equation. The same covariance Q_t from equation 3.1 was incorporated into the free arm movement state equation 2.1 and the reach state equation 2.20. Figure 5 compares position and velocity decoding results for one simulated trial on a trajectory generated from the reach state equation. In this trial, the filter employing a reach state equation is provided the target location with relative certainty by setting both the r and p parameters of Π_T to $1e-5 \text{ m}^2$ in equation 3.2. The point process filter appears to track the actual trajectory more closely with the reach state equation than with the free movement state equation.

Next, we examined the performance of the reach model point process filter in estimating trajectories that were generated from the canonical equation, 3.3, rather than the reach state equation to determine whether the reconstruction would still perform under model violation. Decoding

performance for one trial with the canonical trajectory is illustrated in Figure 6, using the free movement state equation and the reach state equation with r and p in Π_T set to $1e-5 \text{ m}^2$, as with Figure 5. Again, the point process filter tracks the actual trajectory more closely when using the reach state equation than when using the free movement state equation.

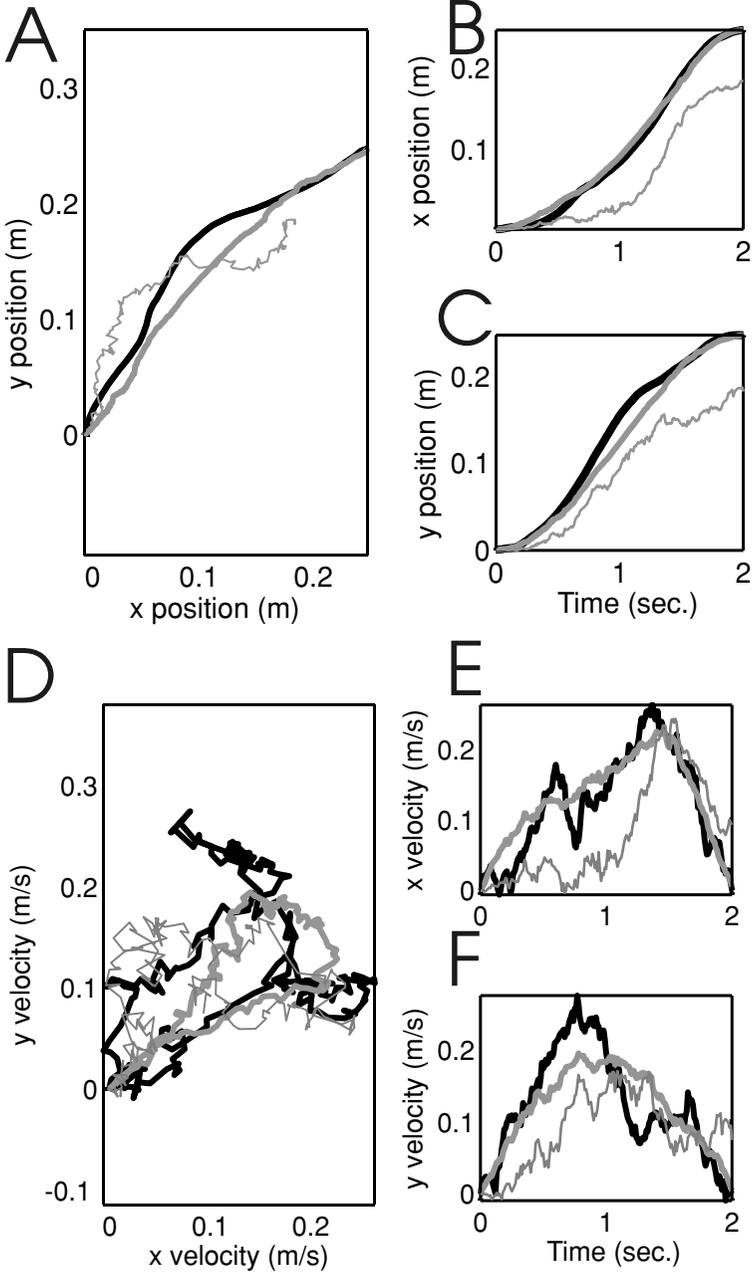
We then assessed whether incorrect and uncertain target planning information could be refined with neural activity that was informative about the path. We implemented the target-augmented state equation and examined the mean and variance of estimates of the target position as the reach progressed. Although the true target coordinates were (0.25 m, 0.25 m) on the x - y plane, the initial estimate of the target location was assigned to (1m, 1m) with a variance of 1 m^2 , large relative to the distance between the initial target estimate and correct target location. Decoding performance for one trial is illustrated in Figure 7. In Figures 7A and 7B, the estimate of the target location is shown to settle close to the true target location relative to the initial target estimate within 1.5 seconds of a 2 second reach. In Figure 7C, the variances in the position (solid) and velocity (dotted) estimates for target (black) approach the variances in estimates for the path (gray) as the reach proceeds.

Finally, we confirmed in simulation that the MSE of reconstruction using the reach state equation approaches that of the free movement state equation as the uncertainty in target position grows. One common simulated set of neural data was used to make a performance comparison between the two methods. Mean squared errors were averaged over 30 trials for the point process filter using the free and reach state equations separately. The results are plotted in Figure 8 for values of r and p in Π_T set equal and over a range from $1e-7 \text{ m}^2$ to 10 m^2 , evenly spaced by 0.2 units on a $\log_{10}(\text{m}^2)$ scale. The MSE line for the reach state equation approaches that of the free movement state equation as Π_T grows large and also flattens as Π_T approaches zero.

4 Discussion

We have developed a method for describing reaching arm movements with a general linear state equation that is constrained by its target. We first derived a reach state equation, which incorporates information about the target that is received prior to movement. This derivation was then adapted to explicitly include the target as an additional state space variable. The resulting augmented state equation supports the incorporation of target information throughout the reach as well as during the planning period.

As described in the derivation, the reach state equation is Markov. This property is guaranteed in part by the independence of noise increments that is demonstrated in the appendix. Consequently, the reach state equation is amenable to recursive estimation procedures. With no further alterations, the estimate of x_t can be obtained exclusively from the neural observation at time t and the estimate of x_{t-1} given data through time $t - 1$.

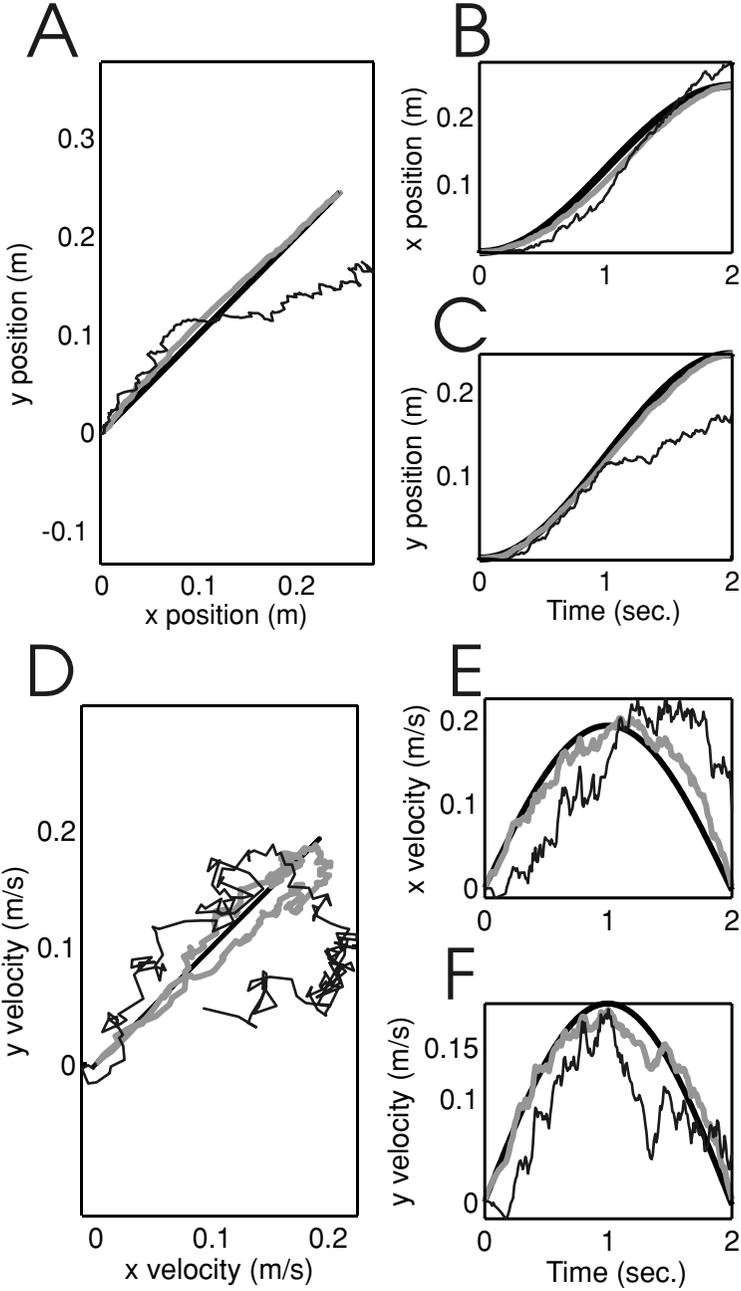


The form of the reach state equation presented in 2.18 is particularly suggestive of stochastic control. In fact, the u_t component in equation 2.18 is the solution to the standard linear quadratic control problem. This represents a duality between estimation and control (Kailath et al., 2000). In this interpretation, the reach state equation is a model for the way in which the subject dynamically plans their path from a current position to the target. The stochastic increment ε_t represents our uncertainty as external observers, about the precise control strategy being employed. The variable u_t takes the role of a control input that represents the adjustments that the subject is expected to make to return the trajectory to a path that puts it on track to the target. In the reach state equation, u_t is a function of the state x_{t-1} and target observation y_T . In the augmented state equation, u_t is a function of x_{t-1} and the actual target x_T rather than the target observation y_T .

Various parameters work together to determine our uncertainty in the control strategy, including the increment variance in the original free movement state equation, distance to target, time to target, and target uncertainty. Together, these parameters determine whether the state equation at any given time forces the trajectory toward a particular target or whether the trajectory tends to proceed in a relatively unconstrained fashion. Figures 2 and 3 describe the variation in trajectories that can be generated by modulating these parameters, from very directed movements to paths with nearly unconstrained directions.

The reach state equation in its simplest form is sufficient to generate, on average, bell-shaped velocity profiles that are similar to those observed in natural arm reaching (Morasso, 1981; Soechting & Lacquaniti, 1981). Models of reaching movement that are based on optimization of specific cost functions, examples of which include Hogan (1984) Uno et al. (1989), Hoff and Arbib (1993), and Harris and Wolpert (1998), also generate these bell-shaped velocity profiles. It has been previously noted in a literature review (Todorov, 2004a) that these various methods implicitly or explicitly optimize a smoothness constraint. In our reach state equation, the

Figure 5: Reconstruction of reaching arm movements from simulated spiking activity. The reach state equation was used to generate trajectories, from which spiking activity was simulated with a receptive field model of primary motor cortex. Point process filter reconstructions using a free movement state equation (thin gray) and a reach state equation (thick gray) were compared against true movement values (black). Trajectories of x and y arm positions were plotted against each other (A) and as a function of time (B, C). Additionally, trajectories of x and y arm velocities were plotted against each other (D) and as a function of time (E, F). In these examples, target location is known almost perfectly to the reconstruction that uses the reach state equation, with position and velocity variances of $1e-5 \text{ m}^2$.



bell-shaped velocity profile emerges implicitly from the zero-mean gaussian increment of the original free movement state equation. This probability density sets a probabilistic smoothness constraint, where it is more likely that the state at consecutive time steps will be similar.

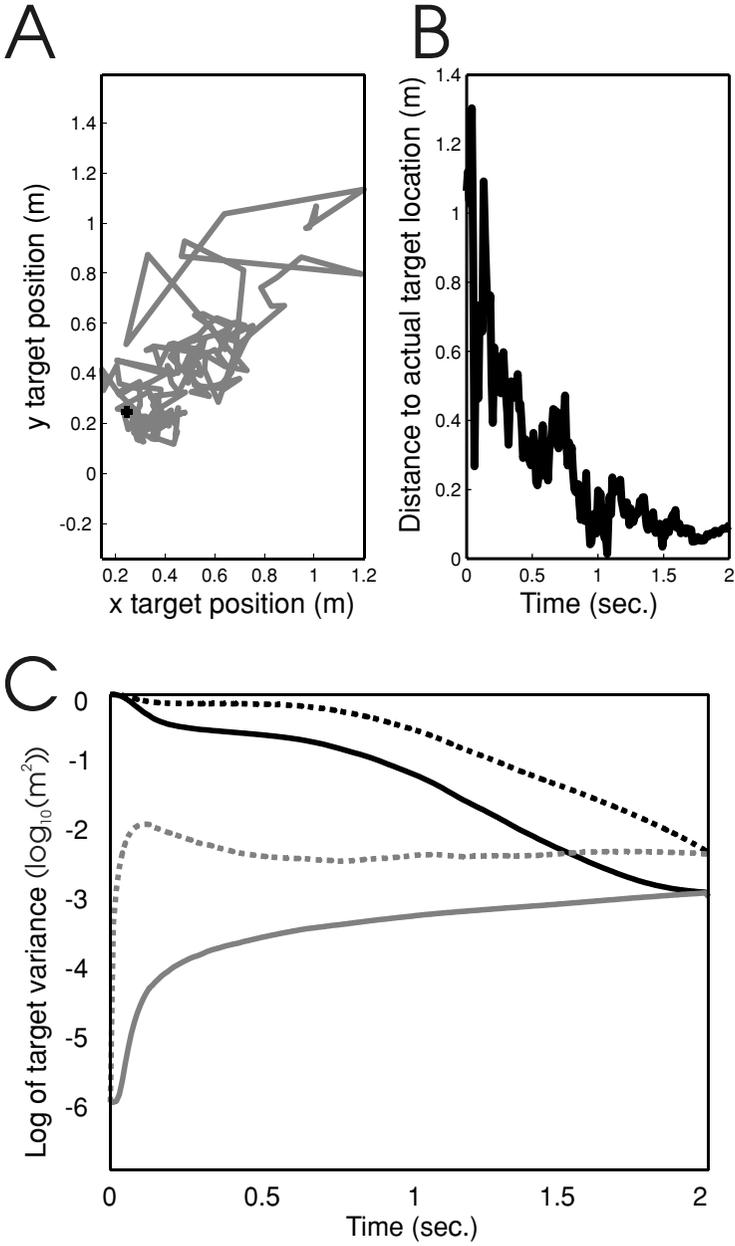
Additionally, symmetry in the profile emerges from the choice of a constant, invertible matrix A_t in equation 2.18 and equal mean starting and ending velocities, as with trajectories in Figure 4A. Optimal control models have previously reproduced the skewed velocity profiles (Hoff, 1992) that occur in experiments (Milner & Ijaz, 1990) where the target must be acquired with increased precision. With the reach state equation, skewed profiles may require the appropriate choice of time-varying components such as A_t and w_t . When the arrival time grows longer (see Figure 4B) or the end point becomes less constrained (see Figure 4C) in the reach state equation, the trajectory tends to resemble a sample path of the free movement state equation, as intended by construction.

As the reaching motion approaches the target arrival time, our sense of the subject's control strategy becomes clearer, because we know the intended target with some uncertainty. We also know that the path must converge to this target soon. Furthermore, we can calculate the control signal that would achieve this goal based on the system dynamics represented by the A_t matrices in equation 2.18. Figure 4D illustrates that the uncertainty in the control strategy, represented by the variance in the stochastic increment ε_t , decreases over the duration of the reach based on y_T , the prescient observation of the target. In contrast, the free movement state equation maintains constant uncertainty in the control strategy as the reach progresses because it is not informed about the target location.

Because the reach state equation incorporates target information, it is able to perform better than the equivalent free movement state equation that is uninformed about target. This is illustrated in Figure 5, where closer tracking is achieved over the entire reach when the state equation is informed about target than otherwise.

This reach model and its augmented form are minimally constrained linear state equations. In a probabilistic sense, this means that the estimation prior at each step is only as narrow (or broad) as implied by the original free

Figure 6: Reconstruction in the face of model violation. Trajectories are generated with an appropriately scaled cosine velocity profile. Again, results are compared for point process filtering using free (thin black) and reach (thick gray) movement state equations against true values (thick black). As with Figure 5, trajectories of x and y arm positions were plotted against each other (A) and as a function of time (B, C). Similarly, trajectories of x and y arm velocities were plotted against each other (D) and as a function of time (E, F). Position and velocity variances of the target observation are $1e-5 \text{ m}^2$.



movement state equation and observations of path and target. In contrast, most reach models based on specific control strategies (Todorov, 2004b), cost functions (Todorov, 2004a), or canonical models (Kemere, Santhanam, Yu, Shenoy, & Meng, 2002; Cowan & Taylor, 2005) place additional constraints on the path that make the estimation prior more exclusive of alternate paths to target. An exception is Kemere and Meng (2005), which uses the linear quadratic control solution that provides identical average trajectories to the reach state equation, based on the estimation control duality (Kailath et al., 2000) although the resulting increment variances are different. As depicted in Figure 6, estimation with a reach state equation is able to perform under model violation, where arm movements are generated by a different model while still taking advantage of the target information.

The target-augmented state equation also allows neural activity related to the path to inform estimates of the target. This is illustrated in Figure 7, where the initial estimate of target position was assigned to be incorrect and with large uncertainty (variance). Consequently, the estimate of the target location relied in large part on neural activity that related to the path. The augmented state equation projects current path information forward in time to refine target estimates. As a result, the estimated target location in Figure 7B settled close to the actual target location 0.5 second before completion of the 2 second reach. The remaining distance between the target location estimate and the actual target location is limited by the extent to which path-related neurons provide good path estimates. For example, path-related neural activity that is relatively uninformative about the path will result in poor final estimates of the target when combined only with poor initial target information. Because the target in the augmented state equation is simply the final point of the path, the variance in the target estimate plotted in Figure 7C approaches that of the path estimate as the reach proceeds to the arrival time T .

Figure 7: Target estimation with the augmented state equation for one trial. The initial estimate of the target is intentionally set to be incorrect at (1 m, 1 m) and with variance of 1 m^2 that is large relative to the distance to the true target location at (0.25 m, 0.25 m). Subsequent target estimates are produced using simulated neural spiking activity that relates directly to the path rather than the target. (A) Estimates of the target position are plotted in gray on the x - y axis, with the actual target marked as a black cross. (B) Distances from target estimates to the actual target location are plotted in meters against time. (C) Variances in estimates of target (black) and path (gray) are plotted on a logarithmic scale over the duration of one reach for position (solid) and velocity (dashed). These target estimate variances reduce with observations consisting of only simulated primary motor cortical activity relating to path.

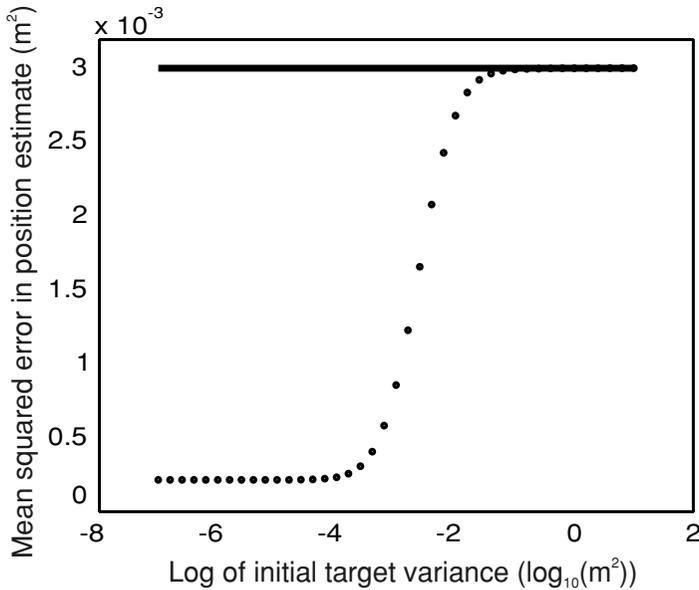


Figure 8: Performance comparison between two approaches to estimation on the same simulated set of neural data. MSE of position reconstruction is plotted versus \log_{10} of uncertainty (variance) in the prescient observation of target. For each of 30 trials, receptive field parameters, trajectory, and spiking activity were simulated anew. For each target variance, MSE is averaged over reconstructions from the 10 trials. In the case of large initial target uncertainty, the MSE for reconstruction with the reach state equation (dotted) asymptotes to that of the free movement state equation (solid). The MSE for reconstruction with the reach state equation also asymptotes as initial target uncertainty diminishes.

The reach state equation in 2.18 or 2.20 reduces to the original free movement state equation in the limit that the prescient target observation is completely uncertain. This explains the trend in Figure 8, where MSE in trajectory estimates with the reach state equation approaches that of the free movement state equation. Estimates were produced from a common simulated set of neural data to allow performance comparison between these two approaches.

Filtering with the reach and augmented state equations, 2.18 and 2.23, respectively, bears resemblance to fixed interval smoothing. Fixed interval smoothing refers to a class of estimation procedures that produce maximum a posteriori estimates of trajectory values over an interval with observations of the trajectory over the entire interval (Kailath et al., 2000). In filtering with the reach state equation, estimates at a given time t are based on data received through time t and the single prescient observation y_T on the target

state x_T . In filtering with the augmented state equation, estimates of x_t are based on data received through time t and potentially multiple prescient observations on x_T . While these three filter types employ observations of future states in maximum a posteriori estimates, there are important distinctions in terms of which observations are used and allowance for multiple sequential observations of a single state, such as with x_T in the augmented state equation.

Although parallels exist to stochastic control, there is a sharp distinction between the results of this article and a control-based state equation (Todorov, 2004b; Kemere & Meng, 2005). First, the reach state equation was derived as the natural extension of a free movement state equation, with no further assumptions. In contrast, control-based state equations are derived by assuming a specific form for the brain's controller and choosing the parameters that optimize some cost function. Second, the increment in the reach state equation approaches zero for perfectly known targets. The increment of control-based state equations persists and represents system properties rather than our uncertainty about the control signal. Third, the reach state equation describes the target state in the most general sense, including the possibility of nonzero velocities. While this can be accommodated in the control framework, the classical notion of a reaching motion has been to a target with zero velocity.

Distinctions between the reach state equation and control-based state equations are especially important in considering the study of reaching motions. Recursive estimation coupled with a state equation that relates target to path provides a convenient tool for the analysis of neural data recorded during planning and execution of goal-directed movements. The state-space estimation framework can assess the extent to which neural data and an observation equation improve the reconstruction beyond information about the movement built into the state equation.

Classically, control-based state equations have been developed to explain as many features about reaching movements as possible without any neural data. In contrast, the reach state equation was developed to extend the free movement state equation with no further assumptions. Both approaches represent different levels of detail in a spectrum of models for the dynamics that drive the observed neural activity in brain regions that coordinate movement. These models can be used to clarify the roles of various brain regions or the validity of alternate neural spiking relationships.

The reach and augmented state equations may also provide improved control in brain machine interfaces (Srinivasan et al., 2005) by allowing the user to specify a target explicitly with neural signals or implicitly through the probability density of potential targets in a workspace. This and other recent approaches (Cowan & Taylor, 2005; Yu et al., 2005; Kemere & Meng, 2005) are hybrids between target-based control prosthetics (Musallam et al., 2004; Santhanam et al., 2005) and path-based control prosthetics (Carmena et al., 2003; Taylor et al., 2002; Wu, Shaikhouni, et al., 2004), perhaps most

relevant when neither approach alone is sufficient for the desired level of control using available recording hardware to complete a task. Additionally, the method could support more robust receptive field estimates in the face of disappearing units due to neuronal death or tissue retraction (Eden, Truccolo, Fellows, Donoghue, & Brown, 2004). The flexibility of the reach and augmented state equations over more specific reach models might also allow the user to employ the same reaching algorithm to navigate obstacles in acquiring the target.

In developing the method further for scientific and clinical application, it is important to consider limitations of the equations presented in this article. Importantly, both the augmented and reach state equation are written for the prescient observation of a target with known arrival time T . We are currently developing a principled approach to accommodate uncertain arrival time, although uncertainty in the target velocity estimate might be a convenient surrogate. Also, the calculations were simplified greatly by assuming a linear free-arm-movement state equation with gaussian increments. This may not be possible if linear approximation is insufficient to describe the nonlinear dynamics of a movement. Finally, additional experimental work will be needed to elucidate the appropriate observation equations, recording sites, and patient rehabilitation regimen that would enhance the clinical application of this hybrid approach to control prosthetics.

Appendix: Proof of Independent Increments in the Reach State Equation

The new increments are defined as $\varepsilon_t = w_t - E[w_t|y_T, x_{t-1}]$. Substituting equation 2.6 into an equation that is equivalent to equation 2.12, we can rewrite the new increments as

$$\varepsilon_t = w_t - Q_t \phi(T, t)' S_t^{-1} \left(\sum_{i=t}^T \phi(T, i) w_i + v_T \right), \quad (\text{A.1})$$

where $S_t = R_T + \sum_{i=t}^T \phi(T, i) Q_i \phi(T, i)'$ and R_T is the covariance of the observation random variable y_T , with $R_T = \phi(T, t-1) V_{t-1} \phi'(T, t-1) + \sum_{i=t}^T \phi(T, i) Q_i \phi'(T, i) + \Pi_T$. Therefore, ε_t can be written entirely in terms of the future increments $\{w_i\}_{i=t}^T$ and v_T . For $s < t$,

$$E[\varepsilon_t \varepsilon_s'] = E \left[\left(w_t - Q_t \phi(T, t)' S_t^{-1} \left(\sum_{i=t}^T \phi(T, i) w_i + v_T \right) \right) \times \left(w_s - Q_s \phi(T, s)' S_s^{-1} \left(\sum_{i=s}^T \phi(T, i) w_i + v_T \right) \right)' \right]$$

$$\begin{aligned}
&= -Q_t \phi(T, t)' S_s^{-1} \phi(T, s) Q_s + Q_t \phi(T, t)' S_t^{-1} \\
&\quad \times \left(\sum_{i=t}^T \phi(T, i) Q_i \phi(T, i)' + R_T \right) S_s^{-1} \phi(T, s) Q_s \\
&= -Q_t \phi(T, t)' S_s^{-1} \phi(T, s) Q_s + Q_t \phi(T, t)' S_s^{-1} \phi(T, s) Q_s = 0.
\end{aligned}
\tag{A.2}$$

Acknowledgments

L.S. thanks Ali Shoeb, Benjie Limketkai, Ashish Khisti, Gopal Santhanam, and Rengaswamy Srinivasan for helpful discussions. We thank Julie Scott and Riccardo Barbieri for help in preparing the manuscript. This research is supported in part by the NIH Medical Scientist Training Program Fellowship to L.S. and NIH grant R01 DA015644 to E.N.B.

References

- Andersen, R. A., & Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, 25, 189–220.
- Ashe, J., & Georgopoulos, A. P. (1994). Movement parameters and neural activity in motor cortex and area 5. *Cerebral Cortex*, 4, 590–600.
- Barbieri, R., Frank, L. M., Nguyen, D. P., Quirk, M. C., Solo, V., Wilson, M. A., & Brown, E. N. (2004). Dynamic analyses of information encoding by neural ensembles. *Neural Computation*, 16(2), 277–308.
- Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E., & Frank, L. M. (2002). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14(2), 325–346.
- Buzáki, G. (2004). Large-scale recording on neuronal ensembles. *Nature Neuroscience*, 7, 446–451.
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., & Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *Public Library of Science Biology*, 1(2), 193–208.
- Castanon, D. A., Levy, B. C., & Willsky, A. S. (1985). Algorithms for the incorporation of predictive information in surveillance theory. *International Journal of Systems Science*, 16(3), 367–382.
- Cowan, T. M., & Taylor, D. M. (2005). Predicting reach goal in a continuous workspace for command of a brain-controlled upper-limb neuroprosthesis. *Proceedings of the 2nd International IEEE EMBS Conference on Neural Engineering*. Piscataway, NJ: IEEE.
- Eden, U. T., Frank, L. M., Barbieri, R., Solo, V., & Brown, E. N. (2004). Dynamic analyses of neural encoding by point process adaptive filtering. *Neural Computation*, 16(5), 971–998.

- Eden, U. T., Truccolo, W., Fellows, M. R., Donoghue, J. P., & Brown, E. N. (2004). Reconstruction of hand movement trajectories from a dynamic ensemble of spiking motor cortical neurons. *Proceedings of the 26th Annual International Conference of the IEEE EMBS*. Piscataway, NJ: IEEE.
- Fu, Q.-G., Suarez, J. I., & Ebner, T. J. (1993). Neuronal specification of direction and distance during reaching movements in the premotor area and primary motor cortex of monkeys. *Journal of Neurophysiology*, *70*(5), 2097–2116.
- Georgopoulos, A. P., Ashe, J., Smyrnis, N., & Taira, M. (1992). Motor cortex and the coding of force. *Science*, *256*, 1692–1695.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, *2*, 1527–1537.
- Georgopoulos, A., Kettner, R., & Schwartz, A. (1988). Primary motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, *8*(8), 2928–2939.
- Geschwind, N., & Damasio, A. R. (1985). Apraxia. In P. J. Vinken, G. W. Bruyn, & H. L. Klawans (Eds.), *Handbook of clinical neurology* (pp. 423–432). Amsterdam: Elsevier.
- Greger, B., Norris, S. A., & Thach, W. T. (2003). Spike firing in the lateral cerebellar cortex correlated with movement and motor parameters irrespective of the effector limb. *Journal of Neurophysiology*, *91*, 576–582.
- Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*, 780–784.
- Hoff, B. (1992). *A computational description of the organization of human reaching and prehension*. Unpublished doctoral dissertation, University of Southern California.
- Hoff, B., & Arbib, M. A. (1993). Models of trajectory formation and temporal interaction of reach and grasp. *Journal of Motor Behavior*, *25*, 175–192.
- Hogan, N. (1984). An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, *4*, 2745–2754.
- Kailath, T., Sayed, A. H., & Hassibi, B. (2000). *Linear estimation*. Upper Saddle River, NJ: Prentice Hall.
- Kemere, C., & Meng, T. H. (2005). Optimal estimation of feed-forward-controlled linear systems. *IEEE International Conference on Acoustics, Speech and Signal Processing*. Piscataway, NJ: IEEE.
- Kemere, C., Santhanam, G., Yu, B. M., Shenoy, K. V., & Meng, T. H. (2002). Decoding of plan and peri-movement neural signals in prosthetic systems. *IEEE Workshop on Signal Processing Systems*. Piscataway, NJ: IEEE.
- Li, C. R., Padoa-Schioppa, C., & Bizzi, E. (2001). Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron*, *30*, 593–607.
- Milner, T. E., & Ijaz, M. M. (1990). The effect of accuracy constraints on three-dimensional movement kinematics. *Neuroscience*, *35*, 365–374.
- Moran, D. W., & Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, *82*(5), 2676–2692.
- Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, *42*, 223–227.

- Mussallam, S., Corneil, B. D., Greger, B., Scherberger, H., & Andersen, R. A. (2004). Cognitive control signals for neural prosthetics. *Science*, *305*, 258–262.
- Nakano, E., Imamizu, H., Osu, R., Uno, Y., Gomi, H., Yoshioka, T., & Kawato, M. (1999). Quantitative examinations of internal representations for arm trajectory planning: Minimum commanded torque change model. *Journal of Neurophysiology*, *81*, 2140–2155.
- Nudo, R. J., Wise, B. M., SiFuentes, F., & Milliken, G. W. (1996). Neural substrates for the effects of rehabilitation training on motor recovery after ischemic infarct. *Science*, *272*, 1791–1794.
- Paninski, L., Fellows, M. R., Hatsopoulos, N. G., & Donoghue, J. P. (2004). Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology*, *91*, 515–532.
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2005). A high performance neurally-controlled cursor positioning system. *IEEE Engineering in Medicine and Biology (EMBS) 2nd International Conference on Neural Engineering*. Piscataway, NJ: IEEE.
- Schwartz, A. B. (1992). Motor cortical activity during drawing movements: Single-unit activity during sinusoid tracing. *Journal of Neurophysiology*, *68*, 528–541.
- Schwartz, A. B., Moran, D. W., & Reina, G. A. (2004). Differential representation of perception and action in the frontal cortex. *Science*, *303*, 380–383.
- Smith, A. C., Frank, L. M., Wirth, S., Yanike, M., Hu, D., Kubota, Y., Graybiel, A. M., Suzuki, W. A., & Brown, E. N. (2004). Dynamic analyses of learning in behavioral experiments. *Journal of Neuroscience*, *24*(2), 447–461.
- Soechting, J. F., & Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. *Journal of Neuroscience*, *1*, 710–720.
- Srinivasan, L., Eden, U. T., Willsky, A. S., & Brown, E. N. (2005). Goal-directed state equation for tracking reaching movements using neural signals. *Proceedings of the 2nd International IEEE EMBS Conference on Neural Engineering*. Piscataway, NJ: IEEE.
- Taira, M., Boline, J., Smyrnis, N., Georgopoulos, A. P., & Ashe, J. (1995). On the relations between single cell activity in the motor cortex and the direction and magnitude of three-dimensional static isometric force. *Experimental Brain Research*, *109*, 367–376.
- Taylor, D. M., Tillery, S. I. H., & Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science*, *296*, 1829–1832.
- Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: A model. *Nature Neuroscience*, *3*(4), 391–398.
- Todorov, E. (2004a). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907–915.
- Todorov, E. (2004b). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, *17*, 1084–1108.
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*, 1226–1235.
- Turner, R. S., & Anderson, M. E. (1997). Pallidal discharge related to the kinematics of reaching movements in two dimensions. *Journal of Neurophysiology*, *77*, 1051–1074.

- Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement: Minimum torque-change model. *Biol. Cybern*, *61*, 89–101.
- Vergheze, G., & Kailath, T. (1979). A further note on backwards Markovian models. *IEEE Transactions on Information Theory*, *IT-25*(1), 121–124.
- Warland, D. K., Reinagel, P., & Meister, M. (1997). Decoding visual information from a population of retinal ganglion cells. *Journal of Neurophysiology*, *78*, 2336–2350.
- Wise, S. P., Boussaoud, D., Johnson, P. B., & Caminiti, R. (1997). Premotor and parietal cortex: Corticocortical connectivity and combinatorial computations. *Annual Review of Neuroscience*, *20*, 25–42.
- Wu, W., Shaikhouni, A., Donoghue, J. P., & Black, M. J. (2004). Closed-loop neural control of cursor motion using a Kalman filter. *The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. Piscataway, NJ: IEEE.
- Yu, B. M., Santhanam, G., Ryu, S. I., & Shenoy, K. V. (2005). Feedback-directed state transition for recursive Bayesian estimation of goal-directed trajectories. *Computational and Systems Neuroscience (COSYNE) meeting*. Salt Lake City, UT. Available online at <http://www.cosyne.org/program05/291.html>.