

Filtering Compensation for Delays and Prediction Errors during Sensorimotor Control

F. Crevecoeur

frederic.crevecoeur@uclouvain.be

Institute of Information and Communication Technologies, Electronics and Applied Mathematics, University of Louvain, Louvain-la-Neuve 1348, Belgium, and Institute of Neuroscience, University of Louvain, Brussels 1200, Belgium

M. Gevers

michel.gevers@uclouvain.be

Institute of Information and Communication Technologies, Electronics and Applied Mathematics, University of Louvain, Louvain-la-Neuve 1348, Belgium

Compensating for sensorimotor noise and for temporal delays has been identified as a major function of the nervous system. Although these aspects have often been described separately in the frameworks of optimal cue combination or motor prediction during movement planning, control-theoretic models suggest that these two operations are performed simultaneously, and mounting evidence supports that motor commands are based on sensory predictions rather than sensory states. In this letter, we study the benefit of state estimation for predictive sensorimotor control. More precisely, we combine explicit compensation for sensorimotor delays and optimal estimation derived in the context of Kalman filtering. We show, based on simulations of human-inspired eye and arm movements, that filtering sensory predictions improves the stability margin of the system against prediction errors due to low-dimensional predictions or to errors in the delay estimate. These simulations also highlight that prediction errors qualitatively account for a broad variety of movement disorders typically associated with cerebellar dysfunctions. We suggest that adaptive filtering in cerebellum, instead of often-assumed feedforward predictions, may achieve simple compensation for sensorimotor delays and support stable closed-loop control of movements.

1 Introduction ---

The apparent ease with which we perform most daily tasks masks the difficulties of movement control. Even simple tasks such as reaching to grasp an object or performing a saccade toward a visual stimulus are supported by efficient closed-loop neural control. Several challenges arise for the brain, of which the following two have been the focus of much research effort:

the compensation for sensorimotor delays and the filtering of sensorimotor noise. Although these two functions have often been dissociated experimentally, theoretical considerations and recent results indicate that estimation and prediction are performed simultaneously in the brain.

On the one hand, sensorimotor delays are too long to explain skillful motor behavior based on delayed-feedback control only (Crevecoeur & Scott, 2014; Kawato, 1999), which suggests that the brain uses internal models and predictions (Miall & Wolpert, 1996; Wolpert & Ghahramani, 2000). One proposed mechanism for predictive control was the Smith predictor (Miall, Weir, Wolpert, & Stein, 1993). In this framework, a model of the biomechanical system is used to predict the consequences of motor commands ahead of sensory feedback. Movements are then controlled based on internal estimates, which removes the delay from the feedback loop provided that the internal model is correct (Zhong, 2010). One difficulty with this approach, and with other prediction techniques, is that they are excessively sensitive to model errors, such that an arbitrarily small mismatch in model parameters can lead to unstable closed-loop control (Michiels & Niculescu, 2007; Mondié & Michiels, 2003). Clearly, assuming that the brain has perfect knowledge of the peripheral motor system is not realistic. Thus, although the Smith predictor captures the principles of predictive neural control, it is unlikely to be implemented in neural circuits as such. Another compensatory mechanism must be at play to mitigate the impact of prediction errors.

On the other hand, with regard to sensorimotor noise, a wealth of research has shown that the brain uses near-optimal estimates of variables such as hand location or of forces applied to the limb, obtained by combining prior knowledge with available sensory information in a statistically efficient manner (Angelaki, Gu, & DeAngelis, 2009; Koording, 2007). This function has been described in the framework of Bayesian statistics in numerous studies on perceptual or motor decisions (Wolpert & Landy, 2012), but also, and importantly, during online control (Izawa & Shadmehr, 2008; Kording & Wolpert, 2004; Wolpert, Ghahramani, & Jordan, 1995). While these previous studies highlighted an efficient handling of uncertainty in the environment and in the nervous system, it is still critical to take sensorimotor delays into account to avoid unstable closed-loop control (Crevecoeur & Scott, 2013). Consistent with this hypothesis, recent results have shown that the combination of internal priors with sensory feedback depends not only on the variance of the feedback signals but also on their temporal delays (Crevecoeur, Munoz, & Scott, 2016).

Thus, predictive control requires compensation for prediction errors, and optimal estimation must include a compensation for temporal delays. In theory, an optimal (Bayesian) integration of prior and delayed feedback is achieved by estimating the present and past states, which for discrete time systems can be achieved in an augmented system (Challa, Evans, & Wang, 2003). The integration of multiple, delayed measurements can also be

optimally integrated using the same approach. However, system augmentation is suitable for digital control, and it is difficult to translate it to continuous time systems or to systems for which present and past estimates may not be jointly available, as in the sensorimotor system. In such a case, when optimal delay compensation cannot be achieved, prediction errors unavoidably occur, and the question arises as to whether closed-loop control remains feasible. Here we investigate this question in theory by analyzing the impact of approximate delay compensation in closed-loop control models of eye and arm movements subject to multiplicative noise in the command (Jones, Hamilton, & Wolpert, 2002) and affected by temporal delays compatible with each motor system.

A seminal study by Miall and colleagues (1993) addressed how a mismatch between the controller and plant delays affected the behavior of simple systems, such as low-pass filters and integrators combined with single gain control. This study also considered relatively long delays (≥ 150 ms) in comparison with visual (~ 100 ms) or somatosensory delays (~ 50 ms in the upper limb) (Scott, 2016). In addition, this formalism did not include stochastic disturbances or uncertainty in the process, which is known to affect planning and control (Harris & Wolpert, 1998; Todorov & Jordan, 2002). To address these limitations, we revisit the problem of predictive control subject to prediction errors in stochastic control models of eye and arm movements. Our approach to this problem is based on finite spectrum assignment (FSA) (Michiels & Niculescu, 2007; Niculescu & Gu, 2004), which in the context of linear and deterministic systems can be stabilized by using a low-pass filter (Mondié & Michiels, 2003). Following the approach of Mondié and Michiels (2003), we show that probabilistic filtering of the predicted state can compensate for prediction errors, including errors that result from a low-dimensional approximation of the system dynamics or from a delay mismatch between the observer and the true system.

Our analyses of the interplay between filtering and prediction errors provide insights into biologically plausible control mechanisms. Indeed, we show that delay compensation may be performed by low-pass filtering of motor commands over the delay interval (such as a moving average), which is easy to link to adaptive filtering in neural circuits. Furthermore, we illustrate how combining priors and predictions improves the stability margin of predictive control and how errors in the delay generate over- or undershoots qualitatively comparable to human movements. Thus, approximate delay compensation can still generate efficient movements in human-inspired models. In addition, this approach captures explicitly the influence of both uncertainty and time in state estimation, which represents a useful framework for modeling the impact of these variables during sensorimotor control. Finally, we found that prediction errors generated a wide spectrum of movement dysmetria that typically characterize disorders observed in disease and, in particular, in cerebellar ataxia and essential tremor. In this regard, adaptive filtering for stable closed-loop control may describe

the function of cerebellum more accurately than the often-assumed feedforward predictions.

We first present the problem definition and provide the derivation of the control algorithm. The section continues with an illustration based on a Langevin process, which shows analytically how the model captures the dependency of the optimal filter on both signal variances and temporal delays. We then illustrate the performance of the algorithm on bio-inspired models of eye and upper limb movements to investigate the robustness of this control algorithm against prediction errors.

2 Model

2.1 Definitions and Assumptions. We consider a class of linear-time-invariant stochastic processes suitable for modeling simple systems such as the control of a single joint or the control of eye movements. Let $W(t)$ represent a p -dimensional Wiener process with unit variance. The class of systems considered is as follows:

$$dx(t) = (Ax(t) + Bu(t))dt + \alpha(t)dW, \quad (2.1)$$

$$y(t) = x(t - \delta t) + \sigma(t), \quad (2.2)$$

where $x(t) \in \mathbb{R}^n$ is the state vector, $u(t) \in \mathbb{R}^m$ is the control input, and the matrices A and B are known, constant, and of appropriate dimension. The function $\alpha(t)$ is an $n \times p$ -dimensional, real-valued function multiplying the instantaneous variation of $W(t)$. We assume that $\alpha(t)$ behaves sufficiently well that the noise disturbances exist for all sample paths and have finite variance over δt . The derivation that follows is valid for any such function $\alpha(t)$. In particular, we will allow $\alpha(t)$ to depend on the control function, $u(t)$, in order to capture the property of sensorimotor systems that the intensity of motor noise increases with the motor commands (Harris & Wolpert, 1998; Jones et al., 2002) (i.e., signal-dependent noise; for more details, see section 5, equation 5.3).

The variable $\sigma(t)$ represents an n -dimensional zero-mean gaussian random variable with covariance Σ . Thus equation 2.2 expresses that the current sensory feedback and the delayed state ($y(t)$ and $x(t - \delta t)$, respectively) have joint gaussian distribution. Finally, we assume without loss of generality that Σ is positive definite.

The problem consists in deriving a feedback controller that minimizes a quadratic cost function. This problem can be solved using linear quadratic gaussian control (LQG) (Astrom, 1970; Phillis, 1985; Todorov, 2005), which gives a control law of the form $u(t) = L(t)\hat{x}(t)$, where $\hat{x}(t)$ is an estimate of the present state of the system. The challenge is to derive the optimal estimate in the presence of noise and measurement delays.

2.2 Optimal Solution Based on System Augmentation. System augmentation reduces the delayed system to the nondelayed case and allows a standard solution derived in the context of stochastic optimal control (Fridman, 2014). In practice, the system is discretized, and a novel state vector is defined as follows (Crevecœur & Scott, 2013):

$$z_t = [x_t^T, x_{t-1}^T, \dots, x_{t-h}^T]^T, \quad (2.3)$$

with the subscript t representing the time step. Each time step corresponds to a discretization interval of dt , and the parameter h represents the delay expressed in number of time steps, that is, $hdt = \delta t$. Thus, the size of the augmented state corresponds to the size of the state multiplied by the number of discrete steps in the delay interval.

This approach allows optimal compensation for temporal delays; however, it is clearly an artificial technique that is suitable for digital applications. First, it is based on the key assumption that the delay is fixed and known, which is not desirable in this context, as we are interested in how the nervous system can handle errors in the prediction or in the estimate of the delay. In addition, the delay compensation is performed by updating the time series included in z_t through the block structure of the state estimator, which requires storing present and past estimates. In this study, we investigate the impact of prediction error that may arise from approximate representation of the dynamics during the delay interval or from errors in the estimate of the delay. Thus, we derive a near-optimal closed-loop controller, that allows manipulating the delay compensation directly. It is built on the insight that system augmentation by design achieves extrapolation of delayed sensory information to the present. Our proposed approach is thus based on an explicit extrapolation of the state measurements.

2.3 Filtering over Predictions. Here we derive an estimator that extrapolates the current sensory feedback explicitly, prior to performing optimal state estimation. Our approach is based on the premise that for deterministic systems, low-pass filtering of the predicted state is known to stabilize predictive control against prediction errors (Mondié & Michiels, 2003; Niculescu & Gu, 2004). We thus derive a similar estimate in the context of stochastic processes and use the statistics of the predicted state to design an optimal filter (Kalman filter) and mitigate the impact of prediction errors.

The current state measurement must first be extrapolated, which is obtained by solving a boundary value problem with uncertainty in the initial condition. Defining $M(t) = e^{tA}$, the unknown solution of equation 2.1 over

the delay interval with initial condition $x(t - \delta t)$ is given by

$$\begin{aligned} x(t) = & M(\delta t)x(t - \delta t) + \int_{t-\delta t}^t M(t-s)Bu(s)ds \\ & + \int_{t-\delta t}^t M(t-s)\alpha(s)dW. \end{aligned} \quad (2.4)$$

The extrapolation of the delayed state measurement computed by the controller is given by the following expression, which takes into account that the third term in equation 2.4 is zero on average:

$$\hat{x}(t|y) := M(\delta t)y(t) + \int_{t-\delta t}^t M(t-s)Bu(s)ds. \quad (2.5)$$

It is easy to rewrite the extrapolation as a function of the true state, which directly gives the extrapolation error, $e(t)$:

$$\hat{x}(t|y) = x(t) + e(t), \quad (2.6)$$

$$e(t) = M(\delta t)\sigma(t) - \int_{t-\delta t}^t M(t-s)\alpha(s)dW. \quad (2.7)$$

Thus, the extrapolation error is gaussian with zero mean and variance $V(e(t))$ given by

$$V(e(t)) = M(\delta t)\Sigma M(\delta t)^T + \int_{t-\delta t}^t M(t-s)\alpha(s)\alpha(s)^T M(t-s)^T ds. \quad (2.8)$$

Note that the cross products of $\sigma(t)$ and dW are zero on average, and the expected value of dW^2 is by definition equal to dt . The variance of $e(t)$ has two terms. The first term captures the uncertainty associated with the measurement used as an initial condition. Indeed, it expresses how the measurement noise $\sigma(t)$ propagates through the systems' unforced dynamics over δt . It is worth mentioning that the uncertainty associated with $\sigma(t)$ decays exponentially, remains constant, or increases in the dimensions in which the matrix A has negative, zero, or positive eigenvalues, respectively. The second term is the variance induced by the process noise, and thus it is a function of the system dynamics captured in $M(\cdot)$, and dependent on $\alpha(t)$.

We can now derive a minimum-variance linear estimate, by using $\hat{x}(t|y)$ as measurement instead of $y(t)$ and applying the following well-known result (Anderson & Moore, 1979): let X and Y have joint gaussian distribution, defined as

$$\begin{bmatrix} X \\ Y \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}, \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix}\right).$$

Then the conditional distribution of X given Y is gaussian and is given by

$$\begin{aligned} X|Y &\sim \mathcal{N}(\mu_{X|Y}, \Sigma_{X|Y}), \\ \mu_{X|Y} &= \mu_X + \Sigma_{XY}^T \Sigma_{YY}^{-1}(Y - \mu_Y), \\ \Sigma_{X|Y} &= \Sigma_{XX} - \Sigma_{XY} \Sigma_{YY}^{-1} \Sigma_{YX}. \end{aligned}$$

The optimal filter over the predicted state is thus obtained by applying this result to calculate the posterior distribution of the next estimate given the current estimate and the extrapolated measurement (see equations 2.6 and 2.7). The whole estimation algorithm is as follows:

1. The initial state estimate is assumed to have known gaussian distribution:

$$\hat{x}(0) \sim \mathcal{N}(x_0, \Sigma_0).$$

2. The true state trajectory over a small discretization interval of dt is

$$x(t + dt) = A_d x(t) + B_d u(t) + \xi_t, \quad (2.9)$$

where $A_d := M(dt)$, $B_d := (\int_0^{dt} M(s)ds)B$, and $\xi_t = \alpha(t)dW$. With this definition, we have that $\xi_t \sim N(0, \alpha(t)\alpha(t)^T dt)$.

3. The one-step prediction of the state at the next time step is computed by simulating the system dynamics over one time interval of dt :

$$\hat{x}_p(t + dt) = A_d \hat{x}(t) + B_d u(t).$$

4. The extrapolation of sensory signals $\hat{x}(t|y)$ and its covariance matrix $V(e(t))$ are computed through equations 2.5 and 2.8.
5. Finally, we use the expression of the conditional distribution of gaussian random variables given above to combine the one-step prediction ($\hat{x}_p(t + dt)$) with the extrapolated state ($\hat{x}(t|y)$), which gives the following recurrence and completes the definition of the estimation algorithm:

$$\Sigma_x(0) = \Sigma_0, \quad (2.10)$$

$$K(t) = A_d \Sigma_x(t) [\Sigma_x(t) + V(e(t))]^{-1}, \quad (2.11)$$

$$\hat{x}(t + dt) = \hat{x}_p(t + dt) + K(t) [\hat{x}(t|y) - \hat{x}(t)], \quad (2.12)$$

$$\Sigma_x(t + dt) = A_d \Sigma_x(t) A_d^T + E(\xi_t \xi_t^T) - K(t) \Sigma_x(t) A_d^T. \quad (2.13)$$

Observe that the filter is well defined because Σ is positive definite from assumptions and $\Sigma_x(t)$ is nonnegative; hence, $V(e(t))$ is also positive definite. This can be verified by expanding equation 2.13 and using the identity that $\Sigma_x + V(e) > \Sigma_x$ in the sense of Loewner ordering. Thus

the matrix inverted in equation 2.11 is positive definite. Observe also that this filter is equivalent to a standard Kalman filter when there is no delay, as equation 2.5 computed when $\delta t = 0$ gives $\hat{x}(t|y) = y(t)$. It is thus a natural extension of the Kalman filter to systems subject to measurement delays.

The filter defined above is based on a prediction of the state for simplicity (FSA), as we consider a state-feedback controller. The difference between FSA and the Smith predictor is that FSA predicts the state of the system, whereas the Smith predictor predicts the output signal, which can be a linear function of the state. With the output signal of equation 2.2, these two approaches are identical in the context of this letter. When a more general output equation of the form $y'(t) = Hx(t - \delta t) + \sigma(t)$ is considered, it is necessary to reconstruct the state vector at $t - \delta t$ with an observer prior to the extrapolation step (see equation 2.5). Note that such an observer need not be optimal to derive the filter above, provided that its output corresponds to equation 2.2 as is the case for general observers.

2.4 Optimal Control. Thus far, we have derived an optimal filter for the predicted state, assuming that the control function $u(t)$, was known. We now derive the optimal controller separately assuming that an optimal estimate of the state is available through the recurrence above (see equations 2.10 to 2.13). It is important to emphasize that this approach is based on a heuristic assumption that the separation principle applied, allowing us to solve the estimation and control problems independently. However, this assumption is not true in general. For instance, when considering multiplicative noise as in sensorimotor systems (e.g., $\alpha(t) \propto u(t)$), the optimal control and optimal filter depend on each other. In the linear case, the estimator and controller can numerically be approximated by iterating the sequence of nonadaptive Kalman gains and control gains until convergence (Phillis, 1985; Todorov, 2005). Our approach is based on first deriving the control gains as if the separation principle applied, and then optimizing the Kalman gains relative to the ongoing control function in each simulation run through the time-varying $V(e(t))$. Thus, our approach is globally suboptimal for linear systems, but we show in the next section that the loss is not substantial. The advantages are that it allows adapting the estimate online following modeled but unexpected disturbances without recalculating the control and estimation gains. It also handles general noise functions captured in $\alpha(t)$.

We consider a finite-horizon control problem that consists in minimizing a cost function of the following form (expressed here in discrete time):

$$J(x, u) = E \left[x_N^T Q_N x_N + \sum_{t=1}^{N-1} x_t^T Q_t x_t + u_t^T R_t u_t \right], \quad (2.14)$$

where N is the time horizon in number of time steps; Q_t and R_t are nonnegative and positive-definite cost matrices, respectively; and $E[\cdot]$ denotes the expected value of the argument. Assuming additive noise, an optimal controller that minimizes $J(x, u)$ can be obtained in the framework of stochastic optimal control as follows (Astrom, 1970; Todorov, 2005):

$$u(t) = -(R + B_d S(t + dt) B_d^T)^{-1} B_d^T S(t + dt) A \hat{x}(t). \quad (2.15)$$

The matrices $S(t)$ are computed offline through a backward-time recurrence starting with $S(Ndt) = Q_N$ (see Todorov, 2005, for details).

2.5 Low-Dimensional Approximation and Implementation. The adaptive filter defined in equations 2.10 to 2.13 is infinite dimensional because in the estimate $\hat{x}(t|y)$ defined by equation 2.5, $u(t)$ is integrated over δt , and the infinite-dimensional property of the integral comes from the fact that it is defined on a space of continuous functions. In order to handle a low-dimensional estimator design, the extrapolation must be replaced by a finite-dimensional approximation of the integral, or quadrature rule. One such rule can be obtained by dividing the delay interval of δt into h subintervals of dt and by summing over each time step of dt used in the computer simulation. This approach corresponds to the following quadrature rule:

$$\hat{x}(t|y) = M(\delta t)y(t) + \sum_{k=1}^h M(t - kdt)Bu(kdt)dt, \quad (2.16)$$

using $h = \delta t/dt$ (see section 5 for numerical values). Observe that in principle, this quadrature rule has the same dimension as the system augmentation technique (see equation 2.3) because we used the same discretization step, and the number of steps in the quadrature rule corresponds to the number of states used in the augmentation technique. Thus, both the number of operations in equation 2.16 and the number of operations involved in the augmentation technique scale with the dimension of the state and the number of time steps during the delay interval.

We may reduce the dimension of the approximation by considering that the motor command is constant over δt and define $B_\delta := \sum_{k=1}^h M(t - kdt)B$. This yields the following one-step approximation of $\hat{x}(t|y)$:

$$\hat{x}(t|y) \cong M(\delta t)y(t) + B_\delta \bar{u}_s. \quad (2.17)$$

The variable $\bar{u}_s$ is a constant approximating the control function $u(s)$ for $t - \delta t \leq s \leq t$. We use the average of the control function, that is, $\bar{u}_s = (\delta t)^{-1} \int_{t-\delta t}^t u(s)ds$, which is also infinite-dimensional in principle for continuous time models but easier to approximate than equation 2.5 in neural




Figure 1: Model pathway for delay compensation in sensorimotor control based on the algebraic operation associated with state estimation. The blue symbols refer to the variables and matrices as defined in the text. The explicit dependencies of M on δt and of K on t were omitted for clarity. The rectangle represents a low-pass filter, or leaky integrator, which approximates $u(s)$ for $t - \delta t \leq s \leq t$.

circuits through low-pass filtering or leaky integrators with appropriate time constants.

Thus far, our developments consisted in replacing the augmentation technique by an explicit extrapolation of the state measurement. The dimensionality of the augmentation corresponds to the number of time steps in the quadrature rule presented in equation 2.16 when the same discretization step is used. Importantly, this dimension can be reduced by approximating the sum in equation 2.16 in one step, which can be done in theory, and approximated in practice with a summary version of $u(s)$, for $t - \delta t \leq s \leq t$. These expressions will allow us to investigate the robustness of the closed-loop controller against extrapolation errors. In addition, we will also investigate the robustness against an error in the internal estimate of the delay by assuming that the controller uses $\hat{\delta t} \neq \delta t$, which has an impact on the estimate in equation 2.17 through the matrices $M(\hat{\delta t})$ and B_δ .

We will show that the state estimator defined in equations 2.10 to 2.13, combined with the low-dimensional quadrature rule given in equation 2.17, constitutes a good candidate model of neural compensation for sensorimotor delays. Indeed, this algorithm involves only multiplications of matrices and vectors, along with an operation that approximates $u(s)$ over $t - \delta t \leq s \leq t$. A schematic block representation of this estimator is shown in Figure 1. The inputs to this pathway consist of the motor command, available through the corollary discharge, the state estimate, and the sensory feedback, which are all available in the brain (Shadmehr, Smith, & Krakauer, 2010; Wolpert, Diedrichsen, & Flanagan, 2011; Wurtz, 2008). The operations are matrix-vector multiplications and additions, plus the averaging, or low-pass filtering, represented by the rectangular box. The remaining computational operations consist in adjusting K over time as a function of $u(t)$ according to equations 2.11 and 2.13. In order to validate this schematic

pathway for studying the impact of prediction errors during motor control, we must show that it generates behaviors compatible with biological movements. This is done in the next section.

3 Applications

3.1 Langevin Process. To provide intuition about how dynamics, variance, and delays affect the uncertainty about the prediction (see equation 2.5), we begin this section with a scalar example for which an analytical expression can be derived, showing in a simple formula how delay compensation affects state estimation when uncertainty about the delay interval is considered. Langevin's equation describes the velocity of a particle subject to random disturbing forces and against a frictional damping force. The scalar stochastic differential equation for the one-dimensional velocity is

$$dx = -Gxdt + \alpha dW, \quad (3.1)$$

where $G > 0$ represents the viscous constant and α is the intensity of the process noise capturing the random perturbations. The state measurement is defined as in equation 2.2. Because α and Σ are constant in this example, the variance of the extrapolation error is also constant. We call it V_e . This variance can be written, after rearranging terms, as follows:

$$V_e = \frac{\alpha^2}{2G} + (\Sigma - \frac{\alpha^2}{2G})e^{-2G\delta t}. \quad (3.2)$$

It can be observed that V_e is equal to Σ when $\delta t = 0$, and it is asymptotically equal to $\alpha^2/2G$ as $\delta t \rightarrow +\infty$. Thus, V_e increases, remains constant, or decreases with the delay, dependent on whether Σ is smaller than, equal to, or greater than $\alpha^2/2G$. The fact that the extrapolation variance decreases over time when $\Sigma > \alpha^2/2G$ is counterintuitive, as it suggests that the extrapolation becomes more reliable with longer delays. This is true in principle; however, the reason the extrapolation becomes more reliable is simply that it tends to zero due to the stable dynamics. As a consequence, the true state and the estimate converge to zero. In any case, even when $\Sigma > \alpha^2/2G$, it is always useful to integrate measurements of the state in the state estimator because $x(t|y)$ and $x(t)$ are correlated (see equation 2.6), and thus the projection of $x(t)$ onto the linear space generated by $x(t|y)$ will always decrease the estimation error.

3.2 Simulation of Human Movements. We now present simulations of saccadic eye movements and single joint reaching of the upper limb with mechanical perturbations. These examples were chosen because the biomechanical plant can be approximated in a linear model and thus correspond




Figure 2: (a) Simulated eye trajectories with the extended LQG based on system augmentation (black) or with the proposed adaptive filter with high- (blue) or low-dimensional (red) quadrature rules (see equations 2.16 and 2.17, respectively). The cost function consisted of two fixation intervals separated by a 50 ms window for movement without any penalty on the state (see section 5 for details about the simulations). (b) Positional variance as a function of time for the three control models with the same color code as in panel a. All traces were normalized to the peak positional variance of the blue trace. (c) Expected cost from the extended LQG control model, and distribution of cost calculated for 500 simulation runs. Observe that the low-dimensional quadrature rule generates inaccurate movements with higher cost but lower positional variance.

to the class of stochastic processes defined in equation 2.1. The extrapolation was based on the two rules given in equations 2.16 and 2.17. We then varied the delay parameters used in the estimation (δt), without varying the true delay, in order to investigate numerically the sensitivity of these control models to this parameter.

Eye movements are of particular interest as the visual system is affected by relatively long delays and the eyeball has low inertia. Thus, close-loop control of the oculomotor plant is challenging and likely more sensitive to model errors than other motor systems. The simulations shown below use $\delta t = 100$ ms for the oculomotor system and $\delta t = 50$ ms for the upper limb. For the visual system, this delay is based on the latency of reflexive saccades (Munoz & Everling, 2004). For the upper limb, the delay of 50 ms corresponds to long-latency feedback supported by a pathway through cortex (Scott, 2016). Details about the models and parameters can be found in the section 5.

We first evaluate the performance of the algorithm derived above by comparing the results with those obtained in the extended LQG framework with system augmentation (Todorov, 2005) in order to quantify how it departs from the augmentation technique (see equation 2.3). When the prediction is performed with the high-dimensional quadrature rule, which corresponds to the dimension of an estimator obtained with state augmentation (see equation 2.16), the average behaviors of the two control algorithms are almost indistinguishable (see Figure 2a, black and blue traces). The difference can be observed in the peak positional variance, where an

increase of about 10% can be noticed. The positional variance then recovered similar values as in the case of extended LQG near the end of the simulation (see Figure 2b). We also measured the cost for each trajectory and found an average increase of less than 10% in comparison with the measured and predicted expected costs from extended LQG. This result is expected because our controller uses slightly greater control signals, which in turn induces more variability across simulations due to signal-dependent noise.

Comparing the two techniques, the augmentation reduces the system to the nondelayed case, but in this context, the presence of signal-dependent noise requires approximate solutions. In contrast to the explicit delay, compensation provides an exact filtering solution to a modified problem and for a given control law. In particular, the analyses in Figure 2 show that for behaviors considered in this study, the simulations obtained for both controllers are very close, and thus the filter developed in equations 2.10 to 2.13 is valid and suitable for studying approximate delay compensation in biological control.

The impact of the low-dimensional quadrature rule is also shown in Figure 2. The main effect is a reduction in the control gain, followed by larger target overshoot for the eye trajectory (see Figure 2a, red). The consequence is that the traces become less variable across trials due to reduced signal-dependent noise, whereas the average cost displays approximately a three-fold increase. We verified that the results were similar for the reaching task (simulations not shown). Note that our goal is not to establish a novel control model for which solutions are available (Bar-Shalom, Huimin, & Mahendra, 2004; Challa et al., 2003; Choi, Choi, & Chung, 2012; Todorov, 2005) but to investigate the impact of extrapolation errors, as well as of errors in the estimate of the delay. To this aim, below we use the low-dimensional quadrature rule (see equation 2.17) and modify the delay parameter used by the controller to compute the extrapolation.

Figure 3a illustrates the impact of an error in the internal estimate of the delay. The different traces were averaged across 50 simulation runs with an estimate of the delay ($\hat{\delta}t$) that was -30% or $+50\%$ of the true delay chosen for illustration (left and right, respectively). Simulations are shown with the closed-loop controller using sensory predictions only (gray, $\hat{x}_t = x(t|y)$), and sensory predictions with filtering (black, $\hat{x}_t$ from equation 2.12) to highlight how filtering reduces the impact of prediction errors. It is noticeable that varying this single parameter generates different behaviors dependent on whether the delay is over- or underestimated. In the case of underestimation (left), the true trajectory is ahead of the estimated trajectory, leading to overshoots and oscillations in the eye angle. In contrast, when the delay is overestimated (right), the controller generates trajectories that fall short of the target and gradually catch up over time. Qualitatively, the simulation of a pursuit movement displays similar properties, although the impact of errors is not as important as during saccades because the pursuit task uses lower control gains than saccadic movements.




Figure 3: (a) Simulation of saccadic eye movements with delay under- (left) or overestimation (right). The curves represent the eye position and the black/gray display refers to the estimation algorithm that was used. The estimation was either the predicted state ($\hat{x}(t|y)$, gray) or the filtered prediction (black). Observe that oscillation amplitude and frequency increase without filtering. (b). Quantification of oscillation amplitudes with or without filtering based on the simulated eye velocity.

To quantify the oscillatory behavior in the simulations, we extracted the ratio between the first and second peak velocity based on the fact that the eye plant was a second-order model (see Figure 3b). It can be observed that the filtering limits the amplitude of the oscillations, which reproduces the results previously established by Mondié and Michiels (2003) in the context of stochastic control. These authors showed that filtering mitigates the impact of prediction errors induced by a quadrature. Here we directly apply this result to stochastic processes and show that prediction errors induced by an approximated extrapolation can also be stabilized. In principle, it is clear that filtering can only improve the controller; thus, the important and novel observation is that closed-loop control can be stabilized for a relatively broad range of delay errors in human-inspired models of eye and arm movements.

We also observed that the main frequency of the oscillations was lower when the filter was applied, and the difference was consistent across over- or underestimation of the delay (average reduction oscillation frequency of 25%). It is also noticeable that it is possible to minimize the amplitude of the oscillations by using an approximation of the delay (see the black arrow, Figure 3b); thus, the delay estimate and the low-dimensional approximation can be tuned together to achieve efficient closed-loop control. Together, these results highlight that applying the filter on approximate predictions improves the stability margin of the closed-loop controller.

In theory, equation 3.2 captures analytically (in the simplest case) the interplay between the noise properties of the process and the delay for estimating its state. In the simulation presented, we concentrated on prediction errors induced by the delay while assuming that the noise parameters were fixed. Interestingly the noise properties have different effects on the




Figure 4: (a) Same as Figure 3 for simulated reaching control of upper limb. A perturbation load (force pulse: -2 Nm applied for 100 ms) was applied to the limb at the moment indicated by the black arrow to highlight closed-loop control properties. (b) Quantification of the oscillation amplitude based on the joint torque. Delay overestimation generated slower movements with shallower velocity peaks, and oscillatory behaviors were observed again for large estimation errors ($\hat{\delta t} \sim 2.5 \delta t$, simulations not shown).

simulations. To observe this, we scaled the motor and sensory noise matrices by a factor of 0.5 and 2 for the saccade task, with a delay error of $\hat{\delta t}/\delta t = 0.6$. In this range, changes in motor noise affected the variability of trajectories and estimates without affecting the closed-loop dynamics substantially. In contrast, increasing the sensory noise reduced the gain of the prediction error in the Kalman filter, which tended to reduce the oscillations induced by the prediction error (simulations not shown). A theoretical account of the impact of errors in the noise covariance matrices on state estimation can be found elsewhere (Heffes, 1966; Nishimura, 1966).

Simulations of upper limb control are shown in Figure 4 and emphasize similar results. Because the delay is shorter and the limb has higher inertia, the impact of using a low-dimensional quadrature rule was hardly visible (see equation 2.16). Figure 4 illustrates the effect of filtering by comparing trajectories with or without filtering for 20 degree reaching movements with a perturbation pulse applied during movement (-1 Nm applied for 100 ms; the onset time is represented with the black arrow). The simulations were generated with varying levels of delay error. When the delay is overestimated ($\hat{\delta t} > \delta t$), the extrapolation is larger than the true state, and the overestimation in turn reduces the control gain. This effect results in shallower movements with slow corrections that do not oscillate until large errors are considered (simulations not shown). Thus, the analysis in Figure 4 concentrated on delay underestimation, which has more impact on control.

In this situation, when the true delay is underestimated ($\hat{\delta t} < \delta t$), the system starts oscillating because the extrapolated state is behind the true state




Figure 5: (a) Simulations of movement oscillations in three distinct loading conditions following pulses of 0.5 Nm applied at the beginning of the simulated task. The cost function penalized deviation from the origin angle throughout the time horizon. The limb inertia was $I = 0.08$ (black), $I = 0.1$ (blue), and $I = 0.12$ (red) [Kgm²]. The simulations were generated with a delay error of 20%: $\delta t = 0.8 \delta t$. (b) Power spectral density of the acceleration signals in each loading condition with similar color code as in panel a. The traces peaked on average at 4.3 Hz, and the gray rectangle highlights that there is less than 1 Hz change across conditions for the peak frequency. The vertical arrow is aligned with the average main frequency. (c) Power spectral density of the acceleration traces for a fixed loading condition ($I = 0.1$ [Kgm²]) across distinct values of delay error: $\delta t / \delta t$ took values 0.6 (solid black), 0.7 (solid gray), 0.8 (dashed black), and 0.9 (dashed gray).

and the controller erroneously increases the control gain. An example is given in Figure 4a for a value of the delay estimate chosen such that the control based on the predicted state only is unstable, whereas the filtered prediction with the same delay error is still stable due to a larger stability margin. The amplitude of the oscillations with or without filtering for delay underestimation is shown in Figure 4b.

To investigate the model properties in more detail, we sought to characterize the oscillations across a range of loading conditions. This simulation was clearly motivated by the variety of tremor types across clinical conditions (e.g., physiologic, essential), and we were interested to see if the model could provide insight into the origin of any kind of tremor. We simulated a posture task in which the controller had to maintain the limb at the origin angle ($\theta = 0$) and applied a small pulse at the beginning of the simulation (see Figure 5a). We focused again on delay underestimation as it evoked clearer oscillations and had more impact on control, but in biological control, it is clear that both over- and underestimation may occur. These simulations were computed with a relative delay error of 20%, in three distinct conditions of limb inertia ($I = 0.1$ Kgm², and $\pm 20\%$). For a limb of 2.5 Kg such as the wrist and hand, a change of 20% corresponds to a loading of 500 g, as used in previous reports on tremor (Elble, 1986). Interestingly,

we found that the peak in the power spectral density was relatively invariant (less than 1 Hz difference) across the three loading conditions (see Figure 5b). The peaks occurred also at the same frequency when we varied the delay error (see Figure 5c). This is due to the fact that the system remains a delayed system with a term corresponding to the true delay for any (including arbitrarily small) delay error (Michiels & Niculescu, 2007). Qualitative similarities with essential tremor are discussed below.

Importantly, we found that for all simulated tasks, the compensation for sensorimotor delays was necessary. Indeed, using $\delta t = 0$ corresponds to a situation in which the sensory feedback (see Figure 1; $y(t)$) is directly compared to the one step prediction of the next state ($\hat{x}_p(t + dt)$) without any delay compensation (see the red box in Figure 1). In this case, all simulations displayed unstable oscillations that increased in amplitude over time and eventually diverged toward infinity. This was observed even for simulations of reaching movements, despite the fact that the upper limb is less prone to instability due to shorter delays and higher inertia. These results suggest that it is necessary to compensate for delays when combining priors and feedback. Alternative control models are thus constrained by the fact that combining priors and feedback without delay compensation may lead to unstable closed-loop control. Previous work shows that it is possible to use delayed state-feedback control to model human motor corrections (Crevecoeur & Scott, 2014). However, since there is evidence of continuous state feedback control and of the existence of internal priors supporting motor responses (Crevecoeur & Scott, 2013; Crevecoeur & Kurtzer, 2018), these control models are not suitable to describe the neural mechanisms underlying long-latency responses to mechanical perturbations.

Finally, we observed that the approach presented could provide a heuristic for combining asynchronous and delayed measurements of the process, such as when vision and proprioception must be combined (Crevecoeur et al., 2016). This problem can in principle be solved based on system augmentation (Challa et al., 2003), which allows updating the posterior distribution of the present and previous states when a novel but delayed measurement becomes available. Here we sought to test whether independent extrapolations of two measurements could be combined with the filter already defined. Such independent extrapolation is suboptimal in theory because it does not correct simultaneously the present and previous estimates. However, it is relevant in the context of sensorimotor control because it remains unclear when or how sensory information from distinct modalities is combined in the neural feedback controller (Oostwoud Wijdenes & Medendorp, 2017). Thus it is interesting to ask whether suboptimal delay compensation may approach the optimal solution.

More precisely, we consider that the measurement vector is

$$y(t) = \begin{bmatrix} x(t - \delta t_1) \\ x(t - \delta t_2) \end{bmatrix} + \begin{bmatrix} \sigma_1(t) \\ \sigma_2(t) \end{bmatrix}, \quad (3.3)$$

which, after extrapolation of each measurement and considering equations 2.7 and 2.8, gives the following measurement signal (I_n is the identity matrix of size n):

$$\hat{x}(t|y_1, y_2) = \begin{bmatrix} I_n \\ I_n \end{bmatrix} x(t) + \begin{bmatrix} e_1(t) \\ e_2(t) \end{bmatrix}, \quad (3.4)$$

with e_1 and e_2 corresponding to each measurement, and the filtering can be performed following standard procedures. It should be emphasized that the noise terms are not independent, because assuming that $\delta t_2 > \delta t_1$, the two extrapolations are performed over a common interval from $t - \delta t_1$ to t , during which the same process noises affect the dynamics. Using equation 2.8, we have

$$\begin{aligned} E[[e_1, e_2]^T [e_1, e_2]] \\ = \begin{bmatrix} M_1 \Sigma_1 M_1^T + I(t - \delta t_1, t) & I(t - \delta t_1, t) \\ I(t - \delta t_1, t) & M_2 \Sigma_2 M_2^T + I(t - \delta t_2, t) \end{bmatrix}, \end{aligned} \quad (3.5)$$

where $M_i := M(\delta t_i)$ and

$$I(u, v) := \int_u^v M(t - s) \alpha(s) \alpha(s)^T M(t - s)^T ds. \quad (3.6)$$

Using the measurement signal from equation 3.4 captures the fact that the latency and gain of feedback responses to perturbations are similar with or without the most delayed measurement, whereas the positional variance across simulations is reduced in the case of combined measurements. Such behavior was observed when participants were instructed to track visually the motion of their hand while their arm was mechanically perturbed (Crevecoeur et al., 2016). This latter study reported similar response latency but reduced estimation variance, which cannot be captured if the different sensor signals are linearly combined without taking their respective delays into account. However, we observed that this independent extrapolation process was numerically fragile and, with a low quadrature rule and the parameters considered, occasionally yielded unstable trajectories. The conditions under which approximate predictions in the form of equation 2.17 can stabilize control with multiple measurements delays should provide informative constraints on the accuracy of neural representations associated with eye-hand coordination.

To summarize, we have shown that combining prediction with optimal filtering can capture closed-loop control of human-inspired models of eye and upper limb movements. In this context, we have highlighted that sensory prediction errors arising when the process integral is approximated over the delay interval, or when the delay is not known exactly, generated

movement over- or undershoots that can be partially mitigated by the application of a filter that considers the noise properties of the extrapolation and the process dynamics. Thus, combining priors with predictions about movement outcome improves the stability margin of the closed-loop controller.

4 Discussion

We have studied the interplay between sensory prediction and state estimation in human-inspired models of the eye and upper limb movements. The presence of sensorimotor delays motivated the hypothesis that the nervous system uses predictions (Miall et al., 1993), often expressed as feedforward control (Kawato, 1999). However, even with predicted feedback in the case of the Smith predictor or with predicted state in the case of FSA, the system remains a closed-loop system, and delays remain problematic when the prediction is not perfect. The delay compensation involves two filtering operations. The first operation is the weighting of one-step predictions with the extrapolated state measurement, which takes the form of a standard Kalman filter (see equation 2.12). The second filtering operation is the computation of the extrapolated state measurement $\hat{x}(t|y)$, which in theory requires integrating the system dynamics over the delay interval (see equation 2.5), but can be approximated with a moving average of the control function as in equation 2.17, or other classes of filters. We have illustrated based on simulations that an approximate extrapolation can still generate movements compatible with human movements, and we have highlighted how the Kalman filter reduces the impact of prediction errors by improving the stability margin. Thus, in addition to reducing uncertainty for perception, optimal estimation also critically improves the stability margin of predictive neural control.

Our approach was mainly motivated by evidence that the nervous system extrapolates sensory signals to adjust motor commands to the expected instantaneous state of the body or environment. Evidence for internal predictions during visual tracking is widespread. Indeed, previous studies showed that both sensory and extraretinal information about the eye and target trajectories are used during online tracking (Bennett, de Xivry, Barnes, & Lefevre, 2007; Blohm, Missal, & Lefevre, 2005; de Brouwer, Yuksel, Blohm, Missal, & Lefevre, 2002), or when anticipating future events such as trajectories following collisions (Badler, Lefevre, & Missal, 2010). This hypothesis also accounts for the perceptual loss that occurs around the time of saccades, suggesting that sensory extrapolation is an important component of visual processing (Crevecoeur & Kording, 2017). Regarding upper limb motor control, the same hypothesis accounted for the use of priors in motor responses to perturbations of distinct profiles (Crevecoeur & Scott, 2013). In this reference, long-latency responses (approximately 50 ms in the upper limb) depended on the expected profile, which is compatible

with an extrapolation of the sensed perturbation-related motion. Thus, we included the extrapolation of sensory information explicitly in the model to account for this function of the nervous system (see equation 2.5).

A clear limitation of our developments is that we focused on linear systems. Of course, the neuromusculoskeletal system is nonlinear. Mathematical challenges aside, the algorithm proposed above is in principle also applicable to nonlinear models of the sensorimotor system. An interesting question for future studies is to investigate whether or how much filtering tolerates prediction errors that arise when one ignores the nonlinearities of the biomechanical plant.

In spite of its simplicity, a compelling aspect of the model is that it captured a broad spectrum of motor disorders by varying a single parameter. Indeed, assuming that the delay compensation mechanisms are similar for eye and limb movements, then movement dysmetria, including hyper- and hypometria, nystagmus, oscillations, and limb tremor, can be explained by assuming that the estimated delay is not equal to the true delay. In practice, such error simply comes down to an inaccurate calibration of the averaging or filtering operation that extrapolates the sensory feedback (see Figure 1, red box).

Regarding the neural basis of the delay compensation, cerebellum has often been associated with motor predictions and state estimation (Bastian, 2011; Kawato, 1999; Miall, Christensen, Cain, & Stanley, 2007; Miall et al., 1993; Wolpert, Miall, & Kawato, 1998). What reminds one of cerebellum again in the context of this letter is that, qualitatively, the dysmetria obtained by varying the internal estimate of the delay has been associated with cerebellar ataxia in clinical populations or following inactivation through cooling experiments performed with nonhuman primates (NHPs) (Bodranghien et al., 2016; Diedrichsen & Bastian, 2014; Mackay & Murphy, 1979). Regarding the control of eye movements, saccadic dysmetria as well as nystagmus were reported in cerebellar ataxia (Gomez et al., 1997; Leech, Gresty, Hess, & Rudge, 1977; Zee, Yee, Cogan, Robinson, & Engel, 1976). This is consistent with the idea that a wrong estimate of the delay in cerebellum alters saccade amplitudes through its involvement in the online control of saccades (Kheradmand & Zee, 2011), which in turn generates movement errors and oscillations.

Turning to the control of the upper limb movements, oscillations were observed in cerebellar patients following reaching movements (Diener, Hore, Ivry, & Dichgans, 1993), similar to those reproduced in Figure 4. Qualitatively, similar oscillations were documented in NHPs following inactivation of cerebellum (Flament, Vilis, & Hore, 1984; Hore & Flament, 1988; Hore & Vilis, 1984). As well, a range of deficits including hyper- and hypometria characterized goal-directed reaching in cerebellar patients (Bhanpuri, Okamura, & Bastian, 2014), a deficit that could be reproduced here by varying the estimated delay.

Of course, other parameters may generate undershoots or oscillations in movements such as errors in the parameters of the internal models (e.g., eye or limb inertia, viscous forces), and it is difficult to derive predictions that are specific to the compensation of sensorimotor delays, and not to other sources of errors. We focused here on approximate delay compensation and thus highlighted that errors in this parameter can also potentially generate movement disorders.

To investigate further the nature of oscillations induced in this model, we simulated a posture task and extracted the main frequency following small perturbation pulses (see Figure 5). Interestingly, we found a main frequency of approximately 4 Hz, which was invariant across loading conditions and errors in the delay. This result bears a striking similarity to previously documented cases of essential tremor, associated with main frequencies in the range 4–8 Hz, and independent of loading (Elble, 2003, 2013). These observations pointed to a neurogenic source of essential tremor that involves a cortico-cerebellar loop (Muthuraman et al., 2018). Thus, poor delay compensation captures oscillations of the closed-loop system at a frequency determined by the fixed and true delay and relatively insensitive to loading or internal errors. In this view, it is not clear that any brain region could be unambiguously identified as the source of oscillations in essential tremor, because the error can be small, but the details of processing can make the closed-loop system oscillate.

This framework is also consistent with the role of cerebellum in coordinating multiple effectors (Bastian, Martin, Keating, & Thach, 1996; Bo, Block, Clark, & Bastian, 2008; Diedrichsen, Criscimagna-Hemminger, & Shadmehr, 2007), as uncoordinated movements can also arise from errors in the estimated delay. Assuming that cerebellum hosts a pathway similar to that of Figure 1, then all of these motor dysfunctions can be explained by inaccurate tuning of the sensory extrapolation, corresponding in principle to an error in the estimate of the delay or, equivalently, an estimate that is not adjusted to the true difference between the sensory feedback and the current state of the body.

Finally, a strength of the model is that the low-pass or averaging operation that is required to compensate for the delay may be directly linked to the adaptive filtering performed in cerebellar microcircuits (Dean, Porrill, Ekerot, & Jorntell, 2010). Thus adaptive filtering of sensory extrapolations may achieve delay compensation and support predictive motor control. The model also suggests that the time constants of the adaptive filters in cerebellum should be related to the temporal delay and to the mechanical properties of the associated motor systems (see equation 2.17). To conclude, we suggest that errors in the estimated delays, or a poor tuning of the filtering operation required in the extrapolation of sensory feedback, may constitute the underlying cause of the wide variety of movement disorders generally associated with cerebellar dysfunctions.

5 Methods

5.1 Continuous-Time Models. The continuous-time state-space representation of the eye dynamics was a second-order model defined as follows:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1/(\tau_1 \tau_2) & -(\tau_1 + \tau_2)/(\tau_1 \tau_2) \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1/(\tau_1 \tau_2) \end{bmatrix} u. \quad (5.1)$$

The state variables x_1 and x_2 represent the eye angle and velocity, respectively; u is the motor command; and the dot operator represents the time derivative. The time constants used in the model were based on previous modeling work (Robinson, Gordon, & Gordon, 1986): we used $\tau_1 = 224$ ms and $\tau_2 = 13$ ms. The explicit dependency of x and u on time was omitted for clarity.

For the upper limb, the continuous-time differential equation was

$$\begin{bmatrix} \dot{\theta} \\ \ddot{\theta} \\ \dot{T} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & -G/I & 1/I \\ 0 & 0 & -1/\tau \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \\ T \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1/\tau \end{bmatrix} u, \quad (5.2)$$

with θ being the joint angle and T the muscle force. The parameters were taken from previous experimental testing (Brown, Cheng, & Loeb, 1999; Crevecoeur & Scott, 2014): $G = 0.14$ Nms (velocity-dependent term capturing viscous forces in muscles); $I = 0.1$ Kgm² (limb inertia); and $\tau = 60$ ms (muscle time constant).

The two models were augmented with the target location. The state vector for the upper limb was further augmented with an external torque (T_E) representing the perturbation applied to the limb. This external torque is assumed to follow step functions, that is, $\dot{T}_E = 0$. This assumption is necessary to reproduce responses to perturbations following a step function, as it must be estimated to generate active compensation for the sustained load applied to the limb. It is also compatible with the behavioral observations that internal priors about the perturbation profile influence long-latency responses (Crevecoeur & Scott, 2013).

5.2 Delays and Noise Parameters. Each state-space representation was transformed into a discrete time control system following equations 2.1, 2.2, and 2.9, with discretization step of $dt = 5$ ms. Thus, for the model of the eye plant, the delay of 100 ms corresponded to $h = 20$ time steps, whereas the delay of 50 ms for the upper limb corresponded to $h = 10$ time steps. We then considered the presence of signal-dependent noise as observed in the sensorimotor system (Harris & Wolpert, 1998), such that the function $\alpha(t)$ was proportional to the control input, u_t . We used the following noise expression,

$$\xi_t = \alpha_1 B_d u(t) \varepsilon_{1,t} + \alpha_2 B_d \varepsilon_{2,t}, \quad (5.3)$$

with $\alpha_1 = 0.08$ or 0.006 for the oculomotor system or upper limb system, respectively; $\alpha_2 = 0.03$ or 0.003 ; and $\varepsilon_{i,t}$ are independent gaussian random variables with zero mean and unit variance. For the sensory noise, we used $\sigma(t) \sim \mathcal{N}(0, \Sigma)$ at each time step with $\Sigma = 0.003 \times I_{n1}$ or $0.001 \times I_{n2}$ for the oculomotor or upper limb system, respectively (I_{ni} represents the identity matrix of appropriate dimensions).

For the upper limb control system, we added a prediction noise that affects the one-step prediction $\hat{x}_p(t + dt)$, without affecting the process dynamics. This prediction noise is required by the presence of an external torque (T_E) that is not controllable through the change in control u_t . Since there is no process noise affecting this variable and because it is not controllable, it is necessary to consider that there is uncertainty about this variable through the internal estimate to allow the estimator to track possible changes in this variable due to external disturbances. This prediction noise was assumed to follow a gaussian distribution with zero mean and covariance matrix set to 0.001 times the identity matrix of the same dimension as the state vector.

These parameters were adjusted manually so that the simulations were numerically well conditioned. Indeed, numerical errors can propagate through the recursions (see equations 2.10–2.13), leading to bad conditioning of the matrix being inverted even if it is mathematically well defined. These noise parameters have an impact on the variability of the trajectories. In principle, the stability margin may depend on the noise parameters, but the improvement obtained with filtering should not depend on these values. A thorough characterization of the relationship between the noise parameters and the stability radius relative to the estimated delay is an interesting question for prospective work.

5.3 Cost Functions. The cost function for all simulations corresponded to equation 2.14. We used $R = 10^{-2}$ and $R = 10^{-4}$ for the oculomotor and upper limb systems, respectively. For the simulations of the oculomotor system, the cost function always consisted in penalizing deviation from the target,

$$x(t)^T Q x(t) = c(x_1(t) - x_1^*)^2, \quad (5.4)$$

where x_1^* is the target angle and c is a parameter equal to 100 during fixation and 0 otherwise. Saccades were simulated by considering two fixation periods separated by an interval of 50 ms without any penalty on the state ($c = 0$). The pursuit task was simulated as a fixation task with a sudden jump in the target position and velocity occurring during the simulation run.

For the reaching task (see Figure 4), we considered a cost function similar to equation 5.4 with $c = 1$ for 300 ms following the movement time, during which c was set to 0. The movement times for the reaching movements represented in Figure 4 was set to 500 ms. The posture task (see Figure 5) was simulated with a target angle of zero for 800 ms, and no movement time. The external perturbation was generated during each simulation run as a force pulse. The pulse amplitude was -2 Nm for the reaching task and -0.5 N for the posture task. In each case, the force pulse was applied for 100 ms.

Acknowledgments

We thank Andrew Pruszynski for comments on an earlier version of the manuscript. F.C. is supported by the F.R.S.-FNRS (Belgium).

References

- Anderson, B. D. O., & Moore, J. D. (1979). *Optimal filtering*. Engelwood Cliffs, NJ: Prentice Hall.
- Angelaki, D. E., Gu, Y., & DeAngelis, G. C. (2009). Multisensory integration: Psychophysics, neurophysiology, and computation. *Current Opinion in Neurobiology*, 19(4), 252–258. doi:10.1016/j.conb.2009.06.008
- Astrom, K. J. (1970). *Introduction to stochastic control theory*. New York: Academic Press.
- Badler, J., Lefevre, P., & Missal, M. (2010). Causality attribution biases oculomotor responses. *J. Neurosci.*, 30(31), 10517–10525. doi:10.1523/JNEUROSCI.1733-10.2010
- Bar-Shalom, Y., Huimin, C., & Mahendra, M. (2004). One-step solution for the multistep out-of-sequence measurement problem in tracking. *IEEE Transactions on Aerospace and Electronic System*, 40(1), 27–37.
- Bastian, A. J. (2011). Moving, sensing and learning with cerebellar damage. *Curr. Opin. Neurobiol.*, 21(4), 596–601. doi:10.1016/j.conb.2011.06.007
- Bastian, A. J., Martin, T. A., Keating, J. G., & Thach, W. T. (1996). Cerebellar ataxia: Abnormal control of interaction torques across multiple joints. *J. Neurophysiol.*, 76(1), 492–509.
- Bennett, S. J., de Xivry, J.-J. O., Barnes, G. R., & Lefevre, P. (2007). Target acceleration can be extracted and represented within the predictive drive to ocular pursuit. *Journal of Neurophysiology*, 98(3), 1405–1414. doi:10.1152/jn.00132.2007
- Bhanpuri, N. H., Okamura, A. M., & Bastian, A. J. (2014). Predicting and correcting ataxia using a model of cerebellar function. *Brain*, 137(Pt. 7), 1931–1944. doi:10.1093/brain/awu115
- Blohm, G., Missal, M., & Lefevre, P. (2005). Processing of retinal and extraretinal signals for memory-guided saccades during smooth pursuit. *J. Neurophysiol.*, 93(3), 1510–1522. doi:10.1152/jn.00543.2004
- Bo, J., Block, H. J., Clark, J. E., & Bastian, A. J. (2008). A cerebellar deficit in sensorimotor prediction explains movement timing variability. *J. Neurophysiol.*, 100(5), 2825–2832. doi:10.1152/jn.90221.2008

- Bodranghien, F., Bastian, A., Casali, C., Hallett, M., Louis, E. D., Manto, M., . . . van Dun, K. (2016). Consensus paper: Revisiting the symptoms and signs of cerebellar syndrome. *Cerebellum*, 15(3), 369–391. doi:10.1007/s12311-015-0687-3
- Brown, I. E., Cheng, E. J., & Loeb, G. E. (1999). Measured and modeled properties of mammalian skeletal muscle. II. The effects of stimulus frequency on force-length and force-velocity relationships. *Journal of Muscle Research and Cell Motility*, 20(7), 627–643. doi:10.1023/a:1005585030764
- Challa, S., Evans, R. J., & Wang, X. (2003). A Bayesian solution and its approximation to out-of-sequence measurement problems. *Information Fusion*, 4, 185–189.
- Choi, M., Choi, J., & Chung, W. K. (2012). State estimation with delayed measurements incorporating time-delay uncertainty. *IET Control Theory and Applications*, 6(15), 2351–2361.
- Crevecoeur, F., & Kording, K. P. (2017). Saccadic suppression as a perceptual consequence of efficient sensorimotor estimation. *Elife*, 6. doi:10.7554/eLife.25073
- Crevecoeur, F., & Kurtzer, I. L. (2018). Long-latency reflexes for inter-effector coordination reflect a continuous state-feedback controller. *Journal of Neurophysiology*, 120, 2466–2483. <https://www.physiology.org/doi/full/10.1152/jn.00205.2018>
- Crevecoeur, F., Munoz, D. P., & Scott, S. H. (2016). Dynamic multisensory integration: Somatosensory speed trumps visual accuracy during feedback control. *J. Neurosci.*, 36(33), 8598–8611. doi:10.1523/JNEUROSCI.0184-16.2016
- Crevecoeur, F., & Scott, S. H. (2013). Priors engaged in long-latency responses to mechanical perturbations suggest a rapid update in state estimation. *PLoS Computational Biology*, 9(8). doi:10.1371/journal.pcbi.1003177
- Crevecoeur, F., & Scott, S. H. (2014). Beyond muscles stiffness: Importance of state estimation to account for very fast motor corrections. *PLoS Computational Biology*, 10(10), e1003869.
- Crevecoeur, F., Sepulchre, R. J., Thonnard, J. L., & Lefevre, P. (2011). Improving the state estimation for optimal control of stochastic processes subject to multiplicative noise. *Automatica*, 47(3), 595–596. doi:10.1016/j.automatica.2011.01.026
- de Brouwer, S., Yuksel, D., Blohm, G., Missal, M., & Lefevre, P. (2002). What triggers catch-up saccades during visual tracking? *J. Neurophysiol.*, 87(3), 1646–1650.
- Dean, P., Porrill, J., Ekerot, C. F., & Jorntell, H. (2010). The cerebellar microcircuit as an adaptive filter: Experimental and computational evidence. *Nat. Rev. Neurosci.*, 11(1), 30–43. doi:10.1038/nrn2756
- Diedrichsen, J., & Bastian, A. J. (2014). Cerebellar function. In M. S. Gazzaniga & G. R. Mangun (Eds.), *The cognitive neurosciences* (pp. 451–460). Cambridge, MA: MIT Press.
- Diedrichsen, J., Criscimagna-Hemminger, S. E., & Shadmehr, R. (2007). Dissociating timing and coordination as functions of the cerebellum. *Journal of Neuroscience*, 27(23), 6291–6301. doi:10.1523/jneurosci.0061-07.2007
- Diener, H. C., Hore, J., Ivry, R., & Dichgans, J. (1993). Cerebellar dysfunction of movement and perception. *Can. J. Neurol. Sci.*, 20 (Suppl. 3), S62–S69.
- Elble, R. J. (1986). Physiologic and essential tremor. *Neurology*, 36(2), 225–231.
- Elble, R. J. (2003). Characteristics of physiologic tremor in young and elderly adults. *Clin. Neurophysiol.*, 114(4), 624–635.
- Elble, R. J. (2013). What is essential tremor? *Curr. Neurol. Neurosci. Rep.*, 13(6), 353. doi:10.1007/s11910-013-0353-4

- Flament, D., Vilis, T., & Hore, J. (1984). Dependence of cerebellar tremor on proprioceptive but not visual feedback. *Experimental Neurology*, 84(2), 314–325. doi:10.1016/0014-4886(84)90228-0
- Fridman, E. (2014). *Introduction to time-delay systems: Analysis and control*. New York: Springer.
- Gomez, C. M., Thompson, R. M., Gammack, J. T., Perlman, S. L., Dobyns, W. B., Truwit, C. L., . . . Anderson, J. H. (1997). Spinocerebellar ataxia type 6: Gaze-evoked and vertical nystagmus, Purkinje cell degeneration, and variable age of onset. *Ann. Neurol.*, 42(6), 933–950. doi:10.1002/ana.410420616
- Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394(6695), 780–784.
- Heffes, H. (1966). The Effect of erroneous models on the Kalman filter response. *IEEE Transactions on Automatic Control*, 11(3), 541–543.
- Hore, J., & Flament, D. (1988). Changes in motor cortex neural discharge associated with the development of cerebellar limb ataxia. *J. Neurophysiol.*, 60(4), 1285–1302.
- Hore, J., & Vilis, T. (1984). Loss of set in muscle responses to limb perturbations during cerebellar dysfunction. *Journal of Neurophysiology*, 51(6), 1137–1148.
- Izawa, J., & Shadmehr, R. (2008). On-line processing of uncertain information in visuomotor control. *Journal of Neuroscience*, 28(44), 11360–11368. doi:10.1523/jneurosci.3063-08.2008
- Jones, K. E., Hamilton, A. F. D., & Wolpert, D. M. (2002). Sources of signal-dependent noise during isometric force production. *Journal of Neurophysiology*, 88(3), 1533–1544. doi:10.1152/jn.00985.2001
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6), 718–727.
- Kheradmand, A., & Zee, D. S. (2011). Cerebellum and ocular motor control. *Front. Neurol.*, 2, 53. doi:10.3389/fneur.2011.00053
- Koerding, K. (2007). Decision theory: What “should” the nervous system do? *Science*, 318(5850), 606–610. doi:10.1126/science.1142998
- Kording, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247. doi:10.1038/nature02169
- Leech, J., Gresty, M., Hess, K., & Rudge, P. (1977). Gaze failure, drifting eye movements, and centripetal nystagmus in cerebellar disease. *Br. J. Ophthalmol.*, 61(12), 774–781.
- Mackay, W. A., & Murphy, J. T. (1979). Cerebellar modulation of reflex gain. *Progress in Neurobiology*, 13(4), 6291–6301. doi:10.1016/0301-0082(79)90004-2
- Miall, R. C., Christensen, L. O. D., Cain, O., & Stanley, J. (2007). Disruption of state estimation in the human lateral cerebellum. *PloS Biology*, 5, 2733–2744. doi:10.1371/journal.pbio.0050316|ISSN 1544-9173
- Miall, R. C., Weir, D. J., Wolpert, D. M., & Stein, J. F. (1993). Is the cerebellum a Smith predictor? *Journal of Motor Behavior*, 25(3), 203–216.
- Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9(8), 1265–1279.
- Michiels, W., & Niculescu, S.-L. (2007). *Stability and stabilization of time-delay systems*. Philadelphia: Society for Industrial and Applied Mathematics.
- Mondié, S., & Michiels, W. (2003). Finite spectrum assignment of unstable time-delay systems with a safe implementation. *IEEE Transactions on Automatic Control*, 48(12), 2207–2212.

- Munoz, D. P., & Everling, S. (2004). Look away: The anti-saccade task and the voluntary control of eye movement. *Nature Reviews Neuroscience*, 5(3), 218–228. doi:10.1038/nrn1345
- Muthuraman, M., Raethjen, J., Koirala, N., Anwar, A. R., Mideksa, K. G., Elble, R., . . . Deuschl, G. (2018). Cerebello-cortical network fingerprints differ between essential, Parkinson's and mimicked tremors. *Brain*, 141(6), 1770–1781. doi:10.1093/brain/awy098
- Niculescu, S.-I., & Gu, K. (2004). *Advances in time-delay systems*. Berlin: Springer-Verlag.
- Nishimura, T. (1966). On the a priori information in sequential estimation problems. *IEEE Transactions on Automatic Control*, 11(2), 197–204.
- Oostwoud Wijdenes, L., & Medendorp, W. P. (2017). State estimation for early feed-back responses in reaching: Intramodal or multimodal? *Front. Integr. Neurosci.*, 11, 38. doi:10.3389/fnint.2017.00038
- Phillis, Y. A. (1985). Controller design of systems with multiplicative noise. *IEEE Transactions on Automatic Control*, AC-30(10), 1017–1019.
- Robinson, D. A., Gordon, J. L., & Gordon, S. E. (1986). A model of the smooth pursuit eye movement system. *Biol. Cybern.*, 55(1), 43–57.
- Scott, S. H. (2016). A functional taxonomy of bottom-up sensory feedback processing for motor actions. *Trends Neurosci.*, 39(8), 512–526. doi:10.1016/j.tins.2016.06.001
- Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33, 89–108. doi:10.1146/annurev-neuro-060909-153135
- Todorov, E. (2005). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, 17(5), 1084–1108.
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.*, 5(11), 1226–1235. doi:10.1038/nrn963
- Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12(12), 739–751. doi:10.1038/nrn3112
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, 3, 1212–1217.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269(5232), 1880–1882.
- Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, 22(6), 996–1003. doi:10.1016/j.conb.2012.05.003
- Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, 2(9), 338–347.
- Wurtz, R. H. (2008). Neuronal mechanisms of visual stability. *Vision Res.*, 48(20), 2070–2089. doi:10.1016/j.visres.2008.03.021
- Zee, D. S., Yee, R. D., Cogan, D. G., Robinson, D. A., & Engel, W. K. (1976). Ocular motor abnormalities in hereditary cerebellar ataxia. *Brain*, 99(2), 207–234.
- Zhong, Q. C. (2010). *Robust control of time-delay systems*. London: Springer-Verlag.