Published in final edited form as:

Ann Biomed Eng. 2013 August ; 41(8): 1613–1624. doi:10.1007/s10439-012-0669-2.

A long-memory model of motor learning in the saccadic system: A regime-switching approach

Aaron L. Wong¹ and Mark Shelhamer^{2,3}

¹Department of Neurology, The Johns Hopkins University School of Medicine, Baltimore, Maryland 21287

²Department of Otolaryngology – Head and Neck Surgery, The Johns Hopkins University School of Medicine, Baltimore, Maryland 21287

³Department of Biomedical Engineering, The Johns Hopkins University School of Medicine, Baltimore, Maryland 21287

Abstract

Maintenance of movement accuracy relies on motor learning, by which prior errors guide future behavior. One aspect of this learning process involves the accurate generation of predictions of movement outcome. These predictions can, for example, drive anticipatory movements during a predictive-saccade task. Predictive saccades are rapid eye movements made to anticipated future targets based on error information from prior movements. This predictive process exhibits long-memory (fractal) behavior, as suggested by inter-trial fluctuations. Here, we model this learning process using a regime-switching approach, which avoids the computational complexities associated with true long-memory processes. The resulting model demonstrates two fundamental characteristics. First, long-memory behavior can be mimicked by a system possessing no true long-term memory, producing model outputs consistent with human-subjects performance. In contrast, the popular two-state model, which is frequently used in motor learning, cannot replicate these findings. Second, our model suggests that apparent long-term memory arises from the trade-off between correcting for the most recent movement error and maintaining consistent long-term behavior. Thus, the model surprisingly predicts that stronger long-memory behavior correlates to faster learning during adaptation (in which systematic errors drive large behavioral changes); greater apparent long-term memory indicates more effective incorporation of error from the cumulative history across trials.

Key terms

long-term memory; predictive saccade; adaptation; fractal fluctuations

Introduction

Movements enable us to interact with the world. To that end, movements must be accurate; this requires active and continuous updating of future motor plans according to previously observed movement errors. Such corrections produce inter-trial fluctuations with memory, as each trial depends upon prior behavior. Previous research suggests that this inter-trial dependence can be quite strong. That is, for certain motor processes, there is clear evidence for the presence of both trial-by-trial error corrections^{22,26} and long-term memory^{19,26}.

One such process in which long-term memory has been previously identified is the generation of predictive saccades. Predictive saccades are rapid eye movements evoked by asking subjects to look between perfectly predictable targets at a rapid pace (0.6–1 Hz). Subjects readily anticipate the next movement, and initiate a saccade before information about the target can be processed by the visual system (latency <70msec)¹⁵. Since these movements are anticipatory and not guided by the current visual stimulus, accurate performance is maintained by accumulating error information across trials. In fact, subjects learn by assessing the accuracy of their past predictions, and using those errors to inform future behavior – these movements employ motor learning. Such learning is related to motor adaptation, which similarly relies on a prediction-based error signal^{4,28}. Thus, predictive saccades are a special case of adaptation in which the task goal is to maintain consistent movement accuracy in response to variable errors, rather than induce a gain change to decrease systematic post-saccadic errors. Indeed, the control of predictive saccades involves multiple timescales of learning, akin to that found for motor adaptation. In particular, predictive-saccade amplitudes exhibit both rapid trial-to-trial error corrections and long-term memory (fractal fluctuations, or power law decay of the power spectrum)²⁶. Furthermore, stimulus variability modulates the long-memory behavior by decreasing inter-trial correlations (see Figure 1), as the brain recognizes that prior experience no longer predicts future behavior. Similar results have been observed for the control of predictive-saccade timing¹³.

The finding of long-term memory (fractal fluctuations) implies that the motor system retains information about performance far in the past, using it for complex computations to generate behavior that exhibits long-range dependence. This seems physiologically implausible, since the motor system must deal with constant short-term changes like environmental perturbations. Instead, a relatively simple trial-by-trial adjustment process seems sufficient to accomplish the required task goals^{2,21}. There are also computational difficulties: long-term memory theoretically means movement errors made *infinitely* far in the past still strongly affect future behavior. Thus, typical long-memory models require vast amounts of information storage, which increase as a function of the process duration. This contrasts with data from human subjects: performance typically becomes more automatic and *less* computationally involved or resource-intensive as subjects perform a task for a longer duration; neural activity shifts from cortex to subcortical brain regions with repetition⁸.

On the other hand, long-term memory is appealing because it implies rich information content, high flexibility, and purposeful variability^{18,20}. It indicates that the underlying system is robust and quickly responds to errors in a manner that is well informed by prior knowledge. In this formulation, variability is not simply random noise, but instead reflects a careful balance about the ideal adaptable state to respond to change, being neither too inflexible nor unstable²³. From this standpoint, it makes sense that motor-learning exhibits long-term memory behavior.

To determine if a physiologically reasonable model can exhibit long-memory behavior, we developed a motor learning model that did not rely on fractal mathematics. This model employs the method of regime switching, a computationally simple approach that previously has been shown capable of mimicking long-term memory^{7,24}. Regime switching is based on the observation that long-range dependence almost looks statistically nonstationary: fluctuations arise from random, sparse shifts of the mean behavioral response. Introducing these shifts at random intervals yields a process that appears fractal, as it exhibits variations on all time scales. For example, one way to produce this type of switching behavior is to couple two processes together, wherein one slowly and intermittently modulates the parameters of the second.

Regime switching has some precedent in the motor learning literature. Zarahn et al.²⁹ analyzed motor adaptation during a reaching task and concluded that the best model fit was one in which the parameters modulated over time – that is, learning parameters changed throughout the session. These modulations of the learning parameters could arise from a higher-level process regulating learning and forgetting rates. Using this approach, it may be possible to generate a computationally reasonable model of the motor-prediction process that exhibits the type of apparent long-term memory observed in human subjects. Such a model is instructive in demonstrating how information from movement errors that occurred far in the past can still exert a large influence on future behavior, and suggests new ways of thinking about both movement variability and the presence of long-term inter-trial correlations. We therefore chose to base our modeling efforts on this approach – which has some precedent in simulating psychological response sequences – and investigated the applicability of such a model to the field of motor control. Thus, we not only examined the regime-switching model's ability to replicate predictive-saccade data, but also investigated its ability to reproduce saccade-adaptation results. Where the model succeeds and where it fails is instructive for understanding its limitations, and in guiding its application in other modeling efforts to explore motor learning.

Materials and Methods

Computational models were constructed in Matlab (MathWorks, Natick, MA). Human-subjects data are adopted from a previous study by the authors²⁶, and were used to fit model parameters. No specific experimental data were used for modeling motor adaptation.

Human-subjects experimentation and data analysis

Ten subjects participated in one or more experimental sessions consisting of at least one block of predictive saccades (16 experiments were performed). All subjects gave their written informed consent according to the local institutional review board. Subjects performing multiple experiment sessions were tested several months apart to minimize learning between sessions. Saccades were recorded using a scleral search coil (Skalar Medical BV, Delft, The Netherlands). Subjects sat in a dark room on a stationary chair, and head movements were minimized using a bite-bar. Paradigms were controlled using experimental software developed in-house. Targets generally consisted of a red laser dot 0.2 mm in diameter that was back-projected onto a screen 100 cm in front of the subject. Eye-tracking data were analyzed offline via an interactive computer program that selected saccades according to a velocity threshold of 15°/s. Saccade amplitudes were measured, then arranged in chronological order to form time series for analysis.

During predictive-saccade blocks, a target was alternately presented to either side of the vertical midline (typically, $\pm 5^\circ$) for 300, 500, or 1000 trials. Stimuli were generally paced at 0.9 Hz. In one session consisting of three 300-target blocks, stimulus variability was introduced in the form of catch trials in which the target was displaced 2° from its original location in the same direction as the current saccade (see Fig. 1). Catch trials were randomly interspersed throughout the block. Subjects experienced three conditions in which the proportion of catch trials differed (0%, 10%, or 20% catch trials). In all other sessions, the stimuli were completely predictable.

Long-term memory was assessed by examining saccade amplitude fluctuations across all time scales using the power spectrum of the data. Since it is the Fourier transform of the autocorrelation function, the power spectrum conveys how information is retained across trials. Long-term memory was quantified by measuring the slope of the power spectrum when plotted on log-log axes (i.e., assessing the power-law decay of the power spectrum). This yields α , the frequency scaling exponent¹⁷. Hurst exponents, H , were computed

according to the rescaled range technique¹ to independently verify α values. H is mathematically related to α ^{11–12}; consistency of these values according to the relationship $H = (1+\alpha)/2$ confirms the presence of long-term memory. Simulations were further evaluated using a modified bootstrap technique to estimate the length of the apparent long-term memory^{25–26}. In this technique, the data are divided into sections of length ΔN , with the start of the first segment randomly selected. The slope of the power spectrum is computed for each section and then averaged to produce an estimated scaling exponent value, α_{est} . As ΔN is varied, the smallest ΔN for which α_{est} matches the measured α value for the full data set equals the longest continuous section of data required to produce the observed long-term memory behavior.

Short-term performance was evaluated by examining trial-to-trial error corrections. The gain change (gain of the next saccade divided by gain of the current saccade) was plotted against the error on the current saccade; 95% confidence ellipses (CE_{95}) were then fit to these graphs. The tilt of the CE_{95} major axis (i.e., the major-axis angle) describes the extent of the correction produced in response to a given error. Thus, this angle estimates the quality of trial-by-trial error corrections in the data. A CE_{95} major-axis angle of -3.01 radians corresponds to a process in which errors are completely corrected on each trial; this is estimated by noting, for example, that a 1° hypometric error (undershoot) for a 10° saccade should produce a gain correction of 1.11.

Model simulation and analysis

The regime-switching model^{7,24} served as the foundation of our predictor model. This model is simply meant to guide our thoughts about motor learning; we do not intend that its component processes readily map to specific brain areas.

A general overview of the model is provided here; additional equations and details are provided in the Supplemental Methods. The model consists of three component processes (Fig. 2; summarized in Table 1). First, the “mean” process controls regime switching by estimating the average desired output. Second, the “error-correction” process adjusts the final output in response to the error on the most recent trial. Finally, the “error-threshold” process regulates long-term performance, as assessed by cumulative errors. The first two model components are essential to simulate prediction behavior, and are described first. This pair of learning processes sum to produce the final output, $\hat{y}_i$, for each trial in response to the input stimulus y_i . That is,

$$\hat{y}_{i+1} = \lambda_1 \hat{\mu}_i + \lambda_2 \hat{x}_i + \varepsilon_{\hat{y}} \quad [1]$$

where $\hat{\mu}_i$ and $\hat{x}_i$ are the mean and trial-by-trial component processes respectively, and λ_1 and λ_2 are mixing coefficients. The $\varepsilon_{\hat{y}}$ term accounts for estimation errors, including errors associated with evaluating post-saccadic accuracy.

The regime-switching component, $\hat{\mu}_i$, regulates average performance and changes infrequently. It keeps performance stable by accumulating error information from the time, τ , that the process was last updated. Clearing the stored memory after each update creates an inherent performance-based “window” that limits the extent of explicitly retained information. This process is defined as

$$\hat{\mu}_{i+1} = \begin{cases} \beta_1 \hat{\mu}_i + \beta_2 \sum_{j=i-\tau+1}^i \varphi_j (y_j - \hat{y}_j) + \varepsilon_\mu & \text{cumerr} \geq \text{ErrThresh} \\ \hat{\mu}_i & \text{otherwise} \end{cases} \quad [2]$$

Eqn. 2 contains a second pair of mixing coefficients, β_1 and β_2 . A discounting term, φ_j , decays exponentially across trials such that recent errors are weighted more heavily. Decay terms are normalized such that the sum of all φ_j is 1.0. Normally-distributed observation noise, ε_μ , reflects uncertainty in estimating average desired performance.

The mean process monitors behavior and makes adjustments when task performance declines. That is, $\hat{\mu}_i$ updates when a weighted cumulative error summation, *cumerr*, reaches a threshold, *ErrThresh*. This is the condition in eqn. 2, *cumerr* $\geq$ *ErrThresh*. Thus, the “switching” portion of the model does not depend on a switching probability, as in other regime-switching models. Instead, errors accumulate until a threshold is reached, at which time the model resets the value of $\hat{\mu}_i$ according to the stored error information. The integrator includes a parameter to make it leaky, so that past errors have a weaker influence on future behavior. Furthermore, errors are weighted according to their magnitude. In this manner, a “cost” is assigned to errors, reflecting the idea that errors are not equally important. The employed cost function assigns more weight to small errors than large errors, consistent with the idea that adaptation is more robust when observed errors remain small^{5,27}. This novel feature distinguishes this model from other common error-based learning models (e.g., static state-space models^{14,21}), which assume that larger errors promote greater learning. Such a cost function that weights errors in direct proportion to their magnitude was also examined (data not shown), but it could not be made to fit the human-subjects data.

The second component process, $\hat{x}_i$, has no long-memory component. It is merely a trial-by-trial update that uses the most recent error to improve the accuracy of the mean process, enabling the system to respond to transient errors such as noise. This process is defined as:

$$\hat{x}_{i+1} = \xi_1 \hat{x}_i + \xi_2 w_i (y_i - \hat{\mu}_i). \quad [3]$$

In eqn. 3, ξ_1 and ξ_2 are a third pair of mixing coefficients, and w_i is the same cost-function weight mentioned above. The trial-by-trial process, combined with the mean process, together are responsible for driving the model output according to prior error information.

The third component is a simple error-tolerance process that describes how strongly the model responds to errors (i.e., it controls the *ErrThresh* value). The error threshold increases (up to a maximum value) as long as errors remain small, representing decreased error sensitivity with task repetition. *ErrThresh* is reset when errors become too large; a large error “surprises” the system and causes it to resume close monitoring of performance. Note, this component process is inactive during predictive-saccade generation, causing the error threshold to remain at its maximal value for the entire simulation. The error threshold only changes during adaptation simulations. Thus, this component can be represented by a constant threshold value for the predictive-saccade task. In short, the regime-switching process paired with the trial-by-trial process are sufficient to produce fluctuations across all time scales, giving the predictor model its long-term memory characteristics.

Model simulations were analyzed using the same techniques as those described for the human-subjects data. One hundred model simulations were produced for each analysis. Both short-term and long-term model behavior were examined. Long-term memory was assessed

by computing the α value of the power spectrum¹⁷, and verified using the Hurst exponent, H^{11-12} . The bootstrap technique was applied to estimate the length of the apparent long-term memory, ΔN^{25-26} . Trial-to-trial error corrections were examined by measuring the angle of the 95% confidence-ellipse major axis (CE_{95}).

The predictor model contains seven parameters that were optimized during model-fitting: three pairs of mixing coefficients, and the error-accumulation leaky-integrator decay rate. Initial conditions were randomly selected such that the initial model output had a mean approximately equal to the initial stimulus input. Parameter fits were judged on four criteria that together characterized the human-subjects data collected previously²⁶. First, simulation data had to have the appropriate long-term dynamics, assessed by α values. Second, the model also had to display appropriate short-term dynamics, with trial-to-trial error corrections producing appropriate CE_{95} major-axis angles. Third, the model output had to be similar to that of the incoming stimulus ($\sim 10^\circ$ amplitude). Finally, the long-memory process had to modulate appropriately such that α decreased as stimulus variability increased.

Once the model parameters were set, the model was then used to explore a secondary learning process: motor adaptation. The model was tested with a conventional adaptation stimulus that consisted of a sudden introduction of exaggerated errors of fixed magnitude (adaptation phase) followed by the sudden removal of those errors (washout phase). This enabled us to examine how well a regime-switching model could respond to a nonstationary, adapting stimulus. Additionally, we explored how inter-trial information storage was modulated in the stationary predictive-saccade task and during the acquisition phase of adaptation, when persistent long-memory may be counter-productive to learning. Adaptation rate was assessed by measuring the time constant of an exponential regression fit of the simulation data. These simulations provided insight into the ways that long-term memory processes might relate more broadly to motor learning.

Results

Experiment data

Subjects produced predictive saccades with appropriate latencies across all blocks (mean $\pm$ s.d., -27.61 ± 83.21). In assessing the long-term memory behavior of these data, subjects exhibited α values of 0.37 ± 0.14 (significantly different from 0.0; t-test, $p < 0.01$) and H values of 0.73 ± 0.04 (significantly different from 0.5; t-test, $p < 0.01$). According to the bootstrap analysis, significant inter-trial correlations extended about 73 ± 7 trials in the past. In agreement with this, trial-to-trial error corrections are incomplete: the CE_{95} major-axis was oriented at -3.05 ± 0.02 radians (significantly different from -3.01 radians representing complete correction for errors; t-test, $p = 0.01$). This suggests that information from many trials influences the production of the next predictive-saccade amplitude. That is, subjects do not simply correct for the error on the most recent trial, but also consider performance from trials that occurred far in the past. Hence, these data strongly indicate the presence of (persistent) long-term memory in the fluctuations of predictive-saccade amplitudes.

Model data

With regard to the model, a small range of parameter values yielded reasonable fits to the human-subjects data. Simulation data are presented for a single representative parameter set in Fig. 3. In general, the model could generate data exhibiting an average α value of 0.38 ± 0.09 (see Fig. 3B), and a mean H value of 0.77 ± 0.04 . These values are similar to those measured from the human-subjects: while the H value is marginally different (Wilcoxon rank-sum test, $p = 0.05$), α is not (Wilcoxon rank-sum test, $p = 0.58$). Thus, the model does a reasonable job of simulating the long-memory properties of the predictive-saccade data.

Although α values approaching zero imply greater trial-to-trial fluctuations (the process becomes more akin to white noise), introducing too much random noise into the model actually makes the simulated data appear to have a strong long-memory component. This seems counterintuitive, but arises because random perturbations accumulate across trials (in a random-walk sense), causing the mean of the process to slowly vary. This creates the *appearance* of long-term fluctuations, and is further exaggerated by increasing trial-to-trial learning.

The bootstrap analysis quantified exactly how “long” the long-memory process extended. For the simulation data, ΔN was 55 ± 54.2 trials, which was not different from the human-subjects data (Wilcoxon rank-sum test, $p = 0.36$). This large ΔN value is quite impressive considering that the size of the window across which information is explicitly retained is much smaller (Fig. 3A, third panel from the top). The average model window size at the time when the switches occurred was 9.62 ± 1.84 trials (largest window size across all simulations is 18 trials), which is significantly smaller than ΔN (Wilcoxon rank-sum, $p < 0.01$). This shows that the model can reproduce apparent long-memory behavior without explicitly retaining a large amount of prior information. In this manner, the motor system could similarly generate movement sequences that appear to exhibit long-range dependence without allocating vast neural resources.

Second, the model mimics trial-to-trial error corrections. Simulations exhibited a CE_{95} major-axis angle of -3.06 ± 0.01 radians (Fig. 3C). This is similar to that of the human-subjects data (Wilcoxon rank-sum test, $p = 0.08$), and different from the ideal correction angle of -3.01 radians (Wilcoxon signed-rank test, $p < 0.01$). The observed major-axis angle indicates that single-trial errors were under-corrected. This is consistent with human-subjects data, and can be attributed to the influence of the long-memory process.

Third, the model output matches the stimulus input. Model simulations were $\sim 7\%$ hypometric to the stimulus input (average $\hat{y}$ value was $9.28^\circ \pm 0.02^\circ$ for a 10° stimulus; see Fig. 3A). Note, this hypometria was not purposely added to the model as an intentional bias, but arose because the model has some tolerance for errors.

Finally, the model exhibits a decreasing α value (a reduction in the long-range dependence) with increasing stimulus variability. Although small, α changes appropriately in response to stimulus variability, (α changes from 0.38 ± 0.09 to 0.36 ± 0.13 ; t-test, $p = 0.04$). This modulation of α is difficult to achieve using other model designs, such as when choosing a cost function that assigns more weight to large errors than to small errors (data not shown).

A brief analysis was conducted to explore the sensitivity of the model to parameter changes. It was found that some model parameters were not particularly critical, while other parameters generated quite large changes in the model output. For example, a 1% change of the most sensitive parameter value, β_1 , caused almost a 45% change of the resulting average α value. On the other hand, the average window size at the time of the regime switch is significantly less sensitive to parameter changes – a 1% change of β_1 caused only a 30% change of the window size (see Supplemental Methods for more details). Thus, the length of the apparent long-memory may be much more sensitive to parameter variations than is the average window size (which relates to the amount of explicitly stored information). This reflects the fact that the dynamic regime-switching process, and to a lesser extent the modulation of the window size, imbues the system with apparent long memory.

Thus, a predictor model was constructed using a regime-switching approach, which could exhibit long-term memory without requiring a heavy computational burden or vast error-information storage. Instead, long-term memory was mimicked by simple component processes that introduced small nonstationarities into the model output according to

measured performance statistics within a narrow temporal window. This model was capable of producing time series data that satisfactorily mimicked human-subjects data during a standard predictive-saccade task. Therefore, the proposed model formulation comprises a reasonable means of generating the complexities in inter-trial movement fluctuations observed during human behavior.

Comparison to the standard two-state model

To confirm that the moderate complexity of the predictor model is necessary, simulations were compared to those of a two-state, state-space model. The two-state model is frequently employed to describe motor learning^{14,21}. Motor plans arise from the sum of two underlying processes with fixed time courses: a slow process that learns gradually and retains that learning well, and a fast process that learns rapidly but quickly forgets what it learned. These processes imbue the model with dynamics across more than one time scale, which is necessary when modeling motor adaptation²¹.

On the surface, the two-state model and the predictor model are similar. Both consist of two underlying processes that describe how past errors regulate future behavior. The greatest difference between them lies in the nature of the “slow” process: the mean process of the predictor model has no single characteristic time course. However, the predictor model also has more parameters and greater model complexity (i.e., explicit storage of several previous trials). On the other hand, the two-state model seems insufficient to produce fractal behavior because it only possesses two time courses of learning; a true fractal process exhibits dynamics across all time scales. If the two-state model can generate long-memory behavior, the additional complexity of the predictor model would be unnecessary. To resolve this question, simulations of the two-state model were compared to those of the predictor model. Parameters of the two-state model were taken from Smith and Shadmehr²¹.

The two-state model exhibits power spectra that decay with a non-zero α value, and trial-by-trial error corrections appear under-corrected. However, power spectra do not exhibit power-law decay (i.e., they are not linear on a log-log plot, but rather exhibit two scaling regions separated by an inflection point). This can be quantified by examining regression fits to power spectra, comparing a single linear regression to a piecewise-linear regression with a single inflection point. No restrictions were placed on the nature of the piecewise fit, such as continuity of the two halves of the regression, to allow for the best possible fit. The Bayesian Information Criterion (BIC), which accounts for the goodness of fit and the number of fitting parameters, was used to compare model responses (Fig. 4).

Of 100 simulations, the piecewise-linear regression (with an inflection point) fit the two-state model power spectra best 95% of the time according to the BIC. The two scaling regions identified in the power spectrum may reflect the two fixed time courses of the two-state model, one operating at low frequencies (the slow process), and one at high frequencies (the fast process). In contrast, the piecewise-linear model only fits the predictor-model power spectra better 16% of the time. This confirms that the mean process in the predictor model influences performance dynamics across numerous time scales, and imbues the model output with a power-law decay of fluctuations across frequency. Thus, the two-state model cannot fully capture the subtle but complex long-memory dynamics of motor learning, which is the strength of regime-switching models like the predictor model.

Model response to an adaptation stimulus

Although not intended to simulate nonstationary processes, the predictor model is a motor-learning model. It appears to describe human-subjects behavior during a stationary predictive-saccade task well, particularly with regard to the presence of long-memory

behavior. Thus, we wanted to examine whether a model based on the regime-switching approach could also describe the nonstationary motor-adaptation process. Adaptation simulations enabled an investigation of the potential relationship between prediction and adaptation. In particular, we explored the manner in which the trade-off between prior information and current errors supports two seemingly different behavioral goals: consistent steady-state performance over long periods of time, and rapid adaptive adjustments in behavior in response to novel errors.

The model responds reasonably to an adaptation stimulus (see Fig. 5). Although there is a short, initial delay in the model response, the output changes in a smooth, exponential fashion in response to the adaptation stimulus. The rate of adaptation is of a similar order of magnitude to that observed for human subjects¹⁶. However, the model learns ~100% of the gain change, which is unrealistic. Furthermore, the model does not behave appropriately during the washout phase: the model response does not return to baseline, and the rate of washout is not faster than the rate of learning as might be expected⁶. This likely results from the symmetric nature of the cost function and the lack of an explicit bias promoting undershoot of the stimulus: positive and negative errors are given equal weight, so washout stops when the *absolute value* of the error is the same as it was prior to adaptation. Thus, while this model reasonably simulates the acquisition phase of adaptation, it cannot account for unlearning. Nevertheless, this affords the opportunity to explore how information might be retained during learning.

Varying the model parameters affects both the predictive-saccade α value and the adaptation learning rate. That is, the model implies a relationship between these seemingly distinct measures of learning dynamics. To explore this relationship, the model was simulated using a range of parameters. The long-memory parameter α was measured during a given prediction simulation, and the learning rate was assessed during a corresponding adaptation simulation using the same parameter set. A significant correlation was observed between the long-memory α value and the rate of adaptation ($p < 0.01$; Fig. 6). Unexpectedly, the model suggests that a subject exhibiting more low-frequency fluctuations during a prediction task (larger α value) will also adapt faster. This is counterintuitive because α is usually thought to describe persistence, with larger α values reflecting a more stable system (less high-frequency fluctuations). Recall, however, that the model can generate nonzero α values by *increasing* low-frequency fluctuations, which arise when greater dependence on incoming information produces larger nonstationarities. Under those conditions, this correlation makes sense: larger α values and faster adaptation rates both reflect the weighting of current error information over prior performance information. In contrast, smaller α values (fewer low-frequency fluctuations) correspond to slower adaptation, implying that future performance depends primarily upon past errors – a behavior-stabilizing approach. In short, the model suggests that if prediction and adaptation share a common underlying process, that learning process depends on the regulation of current and prior error information. Furthermore, these shared dynamics should be evident in measured α values during a predictive-saccade task as well as learning rates during an adaptation task. Finally, it suggests a way to classify individuals according to their learning style: some subjects adapt more rapidly because they rely predominantly upon incoming information, while others adapt more slowly because they discount current movement errors in favor of maintaining stable long-term behavioral consistency.

Discussion

The model presented here arises from an effort to construct a physiologically reasonable model of the predictive motor-learning process. This process, as described previously²⁶, is characterized primarily by rapid trial-by-trial error corrections and persistent long-range

dependence. The latter of these features is particularly troubling from a biological standpoint, since long-memory processes typically require a heavy computational burden. Here, we demonstrate that a regime-switching approach can mimic long-memory behavior via subtle data nonstationarities. To that end, we developed a simple two-component model based on that approach, which includes both a rapidly changing process to keep the model accurate in the face of noise and a slowly changing process that dictates the average model output based on prior information. We then evaluated the applicability of this regime-switching approach to general questions of motor control. We determined that although it appropriately reproduces our major results on saccade prediction, it falls short on replicating certain aspects of saccade adaptation. We suspect that this may be because an additional process must be invoked when errors fall outside a normal expected range, as when artificially introduced by an adaptation stimulus.

The model presented here reasonably replicates some of the characteristics that describe the control of predictive saccades. In particular, it reproduces the short-term and long-range dynamics observed in human-subjects data. This is consistent with previous work indicating that a regime-switching model can produce apparent long-memory behavior^{7,24}, and was a major motivation for selecting this basic structure for the model. Our model also behaves appropriately in the presence of stimulus variability, in that long-term correlations decrease when variability is introduced. This arises because the model continuously monitors performance, but the response to each error is tempered by knowledge about prior behavior. In other words, the model has some tolerance for errors. High error-sensitivity would increase trial-to-trial error corrections, which can accumulate to yield greater nonstationarities (and therefore greater low-frequency fluctuations). This can cause α to increase instead of decrease with additional stimulus variability. Thus, the ability to respond appropriately to stimulus variability – in a manner consistent with human data – turns out to be a non-trivial feature to reproduce.

Aside from producing data that appear to exhibit long-term memory while containing no true long-memory process, this model teaches two additional lessons about motor learning in general. First, the popular two-state adaptation model²¹ cannot produce similar long-memory behavior. This suggests that the two-state model is not a complete description of the dynamics of motor learning, when modeling the behavior of any motor system. In particular, one important shortcoming is that these models can only exhibit two distinct time courses of learning, which is insufficient to replicate long-term memory behavior. Of course, additional time courses can be added to mimic long-memory behavior, as in a so-called aggregation model^{7,9,24}. Aggregation models are based on the idea that motor learning engages multiple brain regions simultaneously, all of which operate at different time scales. Hence, these models are more biologically realistic than regime-switching models. However, aggregation models are challenging to construct because the dynamics of each brain region (including their time courses) must be known. To produce long-term memory, at least one time course of learning must be longer than the longest continuous experimental session. In contrast, regime-switching models can generate long-memory behavior for any number of trials with only two component processes. Thus, they have a distinct advantage over expanded two-state (aggregation) models in representing long-memory behavior, in simulations that explore the implications and benefits of motor behavior that exhibits such complex learning dynamics.

In a broader sense, our model suggests that some aspects of the dynamics of sensorimotor behavior are not captured by more conventional models. Perhaps the experimental protocols used in most previous sensorimotor research do not employ stimuli of sufficient range or variability to invoke the higher-order dynamics that we include in our model – dynamics that involve nonstationarity and apparent long memory. Indeed, if the regime-switching

portion of our model were never invoked, the resulting simulation would look like the typical trial-by-trial learning process that is frequently used to model sensorimotor behavior. In contrast, our experimental protocols (i.e., the predictive-saccade task) require accumulated performance statistics across large numbers of trials, invoking these higher-order dynamics. Adaptation data may simply represent one class of motor-learning behaviors, which in general exhibit complex patterns of inter-trial fluctuations. Capturing subtle aspects of this process, such as long-term memory, has led us to add this additional component to the model. In this sense, our model can be seen as a more general model of sensorimotor learning, and may provide additional insights when applied to other motor systems.

Second, the model suggests that there is a fundamental relationship between prediction and the acquisition phase of adaptation. Since we did not construct this model or fit parameters with the intent to simulate adaptation, it is not surprising that the regime-switching model cannot reproduce all characteristic adaptation features – in particular, the washout phase exhibits inappropriate dynamics. On the other hand, the model reasonably replicates the acquisition phase of learning: simulated adaptation rates are similar to those typically observed for humans¹⁶. This is consistent with the idea that learning depends upon modulating the use of prior and current error information to guide performance. Furthermore, the model predicts that the rate of learning is tied to the system's long-memory behavior. We posit that this relationship may be preserved across all motor systems, and is not unique to saccades. That is, all motor-learning processes must involve information storage across trials. Our model suggests that this retention is achieved as a trade-off between updating movements in response to the most recent error, and maintaining stable behavior by accumulating performance statistics about a reasonable number of prior trials. Thus, more rapid adaptation may correspond to *greater* apparent long-term memory (reflected in larger α values measured during the predictive-saccade task). This counterintuitive correlation reflects the weighting of current errors over prior information in driving future behavior. It remains to be seen if this model prediction can be experimentally confirmed. If so, it would support the hypothesis that motor prediction and adaptation processes are related by learning mechanisms that balance the demands of single-trial movement accuracy and long-term behavioral consistency. The apparent long-term memory observed in behavioral fluctuations may be a characteristic signature of this trade-off. Such a relationship would facilitate the evaluation of learning behavior prior to any exposure to an adaptation stimulus, which is useful when designing individualized rehabilitation or motor-learning paradigms.

By producing apparent long-memory behavior without the high computational burden and memory-storage costs, we have demonstrated the feasibility of producing behavioral fluctuations that exhibit long-term memory. This computational simplicity could enable the brain to balance single-trial errors against long-term stability, imbuing it with the flexibility to respond appropriately to sudden but enduring changes while being tolerant to the presence of transient outliers. Such advantages may explain why long-term memory can be observed in the fluctuations of many other human behaviors, such as gait stride intervals¹⁰ or postural control³. Our model suggests that these behaviors are updated according to a very simple, performance-based scheme in which error statistics are tracked only within a finite time window. This implies that there may be an optimal schedule for the presentation of errors during motor learning paradigms – measureable by evaluating long-term memory behavior and fitting this type of regime-switching model to the data – that can be chosen to produce the greatest impact on learning. Thus, it is useful to understand the nuances of how the brain produces movement fluctuations that exhibit long-term memory – especially without incurring the high costs typically associated with true long-range dependence – to support robust motor learning.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by NSF grant BCS-0615106, NIH grant R21-EY019713, and NIH grant T32 DC000023.

References

1. Basingthwaite, JB.; Liebovitch, LS.; West, BJ. Fractal physiology. New York: Published for the American Physiological Society by Oxford University Press; 1994.
2. Cheng S, Sabes PN. Modeling sensorimotor learning with linear dynamical systems. *Neural Comput.* 2006; 18:760–793. [PubMed: 16494690]
3. Collins JJ, De Luca CJ. Random walking during quiet standing. *Phys Rev Lett.* 1994; 73:764–767. [PubMed: 10057531]
4. Collins T, Wallman J. The relative importance of retinal error and prediction in saccadic adaptation. *J Neurophysiol.* 2012
5. Criscimagna-Hemminger SE, Bastian AJ, Shadmehr R. Size of error affects cerebellar contributions to motor learning. *J Neurophysiol.* 2010; 103:2275–2284. [PubMed: 20164398]
6. Davidson PR, Wolpert DM. Scaling down motor memories: de-adaptation after motor learning. *Neurosci Lett.* 2004; 370:102–107. [PubMed: 15488303]
7. Diebold FX, Inoue A. Long memory and regime switching. *J Econometrics.* 2001; 105:131–159.
8. Floyer-Lea A, Matthews PM. Changing brain networks for visuomotor control with increased movement automaticity. *J Neurophysiol.* 2004; 92:2405–2412. [PubMed: 15381748]
9. Granger CWJ. Long memory relationships and the aggregation of dynamic models. *J Econometrics.* 1980; 14:227–238.
10. Hausdorff JM, Peng CK, Ladin Z, Wei JY, Goldberger AL. Is walking a random walk? Evidence for long-range correlations in stride interval of human gait. *J Appl Physiol.* 1995; 78:349–358. [PubMed: 7713836]
11. Hurst HE. Long-term storage capacity of reservoirs. *Trans Amer Soc Civil Eng.* 1951; 116:770–808.
12. Hurst, HE.; Black, RP.; Samayka, YM. London: Constable; 1965. Long-term storage: an experimental case study.
13. Joiner WM, Shelhamer M. Responses to noisy periodic stimuli reveal properties of a neural predictor. *J Neurophysiol.* 2006; 96:2121–2126. [PubMed: 16790594]
14. Kording, K.; Tenenbaum, JB.; Shadmehr, R. Multiple timescales and uncertainty in motor adaptation. In: Scholkopf, B.; Platt, J.; Hoffman, T., editors. *Advances in neural information processing systems*. Vol. 19. Cambridge, MA: MIT Press; 2007. p. 745-752.
15. Leigh, RJ.; Zee, DS. The neurology of eye movements. New York: Oxford University Press; 2006.
16. Miller JM, Anstis T, Templeton WB. Saccadic plasticity: parametric adaptive control by retinal feedback. *J Exp Psychol Hum Percept Perform.* 1981; 7:356–366. [PubMed: 6453929]
17. Rangarajan G, Ding M. Integrated approach to the assessment of long range correlation in time series data. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics.* 2000; 61:4991–5001. [PubMed: 11031542]
18. Riley MA, Turvey MT. Variability of determinism in motor behavior. *J Mot Behav.* 2002; 34:99–125. [PubMed: 12057885]
19. Shelhamer M, Joiner WM. Saccades exhibit abrupt transition between reactive and predictive; predictive saccade sequences have long-term correlations. *J Neurophysiol.* 2003; 90:2763–2769. [PubMed: 14534279]
20. Slifkin AB, Newell KM. Noise, information transmission, and force variability. *J Exp Psychol Hum Percept Perform.* 1999; 25:837–851. [PubMed: 10385989]

21. Smith MA, Ghazizadeh A, Shadmehr R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* 2006; 4:e179. [PubMed: 16700627]
22. Srimal R, Diedrichsen J, Ryklin EB, Curtis CE. Obligatory adaptation of saccade gains. *J Neurophysiol.* 2008; 99:1554–1558. [PubMed: 18234985]
23. Stergiou N, Decker LM. Human movement variability, nonlinear dynamics, and pathology: Is there a connection? *Hum Mov Sci.* 2011
24. Wagenmakers EJ, Farrell S, Ratcliff R. Estimation and interpretation of $1/\alpha$ noise in human cognition. *Psychon Bull Rev.* 2004; 11:579–615. [PubMed: 15581115]
25. Wang GD, Antar G, Devynck P. The Hurst exponent and long-time correlation. *Phys Plasmas.* 2000; 7:1181–1183.
26. Wong AL, Shelhamer M. Exploring the fundamental dynamics of error-based motor learning using a stationary predictive-saccade task. *PLoS ONE.* 2011; 6:e25225. [PubMed: 21966462]
27. Wong AL, Shelhamer M. Saccade adaptation improves in response to a gradually introduced stimulus perturbation. *Neurosci Lett.* 2011; 500:207–211. [PubMed: 21741440]
28. Wong AL, Shelhamer M. Sensorimotor adaptation error signals are derived from realistic predictions of movement outcomes. *J Neurophysiol.* 2011; 105:1130–1140. [PubMed: 21123665]
29. Zarah E, Weston GD, Liang J, Mazzoni P, Krakauer JW. Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. *J Neurophysiol.* 2008; 100:2537–2548. [PubMed: 18596178]

Figure 1.

Modulation of long-range dependence with stimulus reliability. As stimulus variability (and hence stimulus uncertainty) increases, long-range dependence decreases. Stimulus variability was modulated by selecting a different proportion of catch trials to be presented to subjects during each predictive-saccade block. A: Sample eye (gray solid line) and target (dashed black line) position traces during a catch trial experiment. Catch trials are denoted by arrows. B: Autocorrelation functions from one example subject for three conditions of increasing stimulus variability show that as variability increases, the width of the autocorrelation function decreases (i.e., inter-trial correlations fall off faster, reflecting less retention of information from trial to trial). C: Across 5 subjects, the α value decreases (successive predictive-saccade amplitudes become less correlated, or behavioral fluctuations appear more like a zero-memory, white-noise process) as stimulus variability increases.

Figure 2.

A wiring diagram of the predictor model. The three component processes are denoted, although we find that highest-level (error-threshold) process is only engaged during adaptation simulations. The primary model components supporting predictive-saccade behavior are the “mean” process and the trial-by-trial process. A full description of this model is provided in the text, and summarized in Table 1.

Figure 3.

Sample predictive-saccade simulation of a “good” parameter set. A: The simulation is plotted in its entirety, although for analysis purposes only the last 500 data points of the simulation are considered to remove the influence of the model’s arbitrary initial conditions. Plotted from top to bottom: the model output $\hat{y}$ and the mean component process $\hat{\mu}$ along with the stimulus input; error on each trial; the weight of the error according to the cost function; the window size (number of trials over which error information is accrued to reset the $\hat{\mu}$ estimate; and the error threshold value as controlled by the third component (error-tolerance) process (which remains constant for these simulations). B: Power spectrum for these data, showing a measured α value of 0.38 with a confidence interval of (0.26, 0.50). C: Trial-by-trial error-correction analysis with the CE_{95} plotted; the major-axis angle for this simulation is -3.06 with a confidence interval of $(-3.065, -3.069)$. Also displayed for comparison is the ideal CE_{95} ellipse (gray) with a major-axis angle of ~ 3.01 radians; there is a significant difference between the simulation and the ideal ellipse.

Figure 4.

Comparison between the predictor model and the two-state model. Although the two-state model produces power spectra that decay with nonzero slopes, the power spectra do not exhibit strict power-law decays. (A) Regression analysis reveals that the two-state model exhibits a marked inflection point such that power spectra are best described with piecewise-linear regression fits (PW). (B) In contrast, single linear fits (LIN) best describe the power spectra computed for simulations of the predictor model. Fits are evaluated according to the Bayesian Information Criterion (BIC).

Figure 5.

Simulation of the predictor model in response to an adaptation stimulus. A: Model output in response to a traditional adaptation paradigm in which the requested saccade gain increases by 25%, followed by a washout phase where the gain returns to its initial condition. Qualitatively, the model responds reasonably well; the simulation amplitude follows an exponential time course during the adaptation phase. B: Error-related parameters during the simulation; from top to bottom: error on each trial; and the cumulative error threshold as controlled by the third (error-tolerance) component process. Note that the error threshold modulates slowly, unlike during the prediction simulation when the error threshold is maintained at a constant value.

Figure 6.

Correlation between the prediction long-memory α value and the adaptation rate. Simulation data are plotted for a range of parameter values. There is a significant correlation between the α -value assessed during a simulation of a standard prediction task, and the learning rate measured with those same parameters during an adaptation simulation. This implies that if prediction and adaptation share a common underlying learning process, the dynamics of that process can be similarly captured in these two learning parameters (the α value and the adaptation rate).

Table 1

Summary of model component processes

Component process	Time scale	Responsibility
“error-tolerance” process	slow	Adjust the cumulative-error threshold used to determine when regime switching occurs, based on prior performance
“mean” process	moderate	When error-accumulation reaches threshold, re- estimate the “mean” of the process by averaging over prior output values (only including values since the last regime switch; at most, this is ~18 trials)
“error-correction” process	fast	Rapidly adjust model output to address single-trial variability (errors and noise)