

Learning To Move Amid Uncertainty

Robert A. Scheidt^{1,2}, Jonathan B. Dingwell² and Ferdinando A. Mussa-Ivaldi²

¹Department of Biomedical Engineering
Marquette University
Milwaukee, WI 53201

²Sensory Motor Performance Program
Rehabilitation Institute of Chicago
Chicago, IL 60611

Telephone: (414) 288-6124
Fax: (414) 288-7938
email: scheidt@ieee.org

Running Title: Learning to move amid uncertainty

Keywords: motor adaptation, incremental learning, stochastic time series, human performance

Address for correspondence:
Robert A. Scheidt
Department of Biomedical Engineering
Olin Engineering Center, 303
P.O. Box 1881
Marquette University
Milwaukee, WI 53201-1881

Number of words in abstract:	268
Number of text pages:	22
Number of figures:	9
Number of tables:	1

Abstract

We studied how subjects learned to make movements against unpredictable perturbations. Twelve healthy human subjects made goal-directed reaching movements in the horizontal plane while holding the handle of a two-joint robotic manipulator. The robot generated viscous force fields that perturbed the limb perpendicular to the desired direction of movement. The amplitude (but not the direction) of the viscous field varied randomly from trial to trial. Systems identification techniques were employed to characterize how subjects adapted to these random perturbations. Subject performance was quantified primarily using the peak deviation from a straight-line hand path. Subjects adapted their arm movements to the sequence of random force field amplitudes. This adaptive response compensated for the approximate mean from the random sequence of perturbations and did not depend on the statistical distribution of that sequence. Subjects did not adapt by directly counteracting the mean field strength itself on each trial, but rather by using information about perturbations and movement errors from a limited number of previous trials to adjust motor commands on subsequent trials. This strategy permitted subjects to achieve near-optimal performance (defined as minimizing movement errors in a least-squares sense) while maintaining computational efficiency. A simple model using information about movement errors and perturbation amplitudes from a single previous trial predicted subject performance in stochastic environments with a high degree of fidelity and further predicted key performance features observed in non-stochastic environments. This suggests that the neural structures modified during motor adaptation require only short-term memory. Explicit representations regarding movements made more than a few trials in the past are not used in generating optimal motor responses on any given trial.

Introduction

A remarkable and well studied ability of the human brain is that of adapting the execution of limb movements to physical changes in operating conditions such as those that naturally occur during growth, aging, and exposure to altered mechanical environments (Bock 1990; Happee 1993; Shadmehr and Mussa-Ivaldi 1994; Lackner and Dizio 1994; Dizio and Lackner 1995; Conditt et al. 1997a; Goodbody and Wolpert 1998; Thoroughman and Shadmehr 1999; Scheidt and Rymer 2000). This process is known as motor adaptation. Motor adaptation is a form of learning that evolves over a series of movements whereby some original performance of a given task is restored in the presence of an external perturbation. This ability to adapt to environmental changes has played an important role in human survival. A species unable to compensate for prevailing winds or the refraction of light through water would be ill suited to use the basic tools (such as spears and nets) necessary for fighting off foes and obtaining food. In such instances, environmental perturbations influence the control of upper limb movement in an unpredictable way.

A number of studies have investigated the processes involved in motor adaptation by exposing subjects to specific perturbations and quantifying the changes in their responses over time. For example, some experiments have explored the changes in reaching and pointing movements of the hand induced by displacements or deformations of the visual field (Helmholtz 1925; Held and Freedman 1963; Flanagan and Rao 1995; Wolpert et al. 1995). Other experiments have perturbed the moving arm with mechanical disturbances that emulated the effects of inertial loads and/or viscoelastic media (Bock 1990; Shadmehr and Mussa-Ivaldi 1994; Lackner and Dizio 1994; Shadmehr and Brashers-Krug 1997). Each of these studies employed perturbations with fixed and repeatable structures. For example, Shadmehr and Mussa-Ivaldi (1994) used a robotic device to apply mechanical forces to the hand. These forces had a fixed linear dependence on the speed of the subject's hand.

However, the perturbations that people encounter in everyday life do not always have a repeatable and consistent structure. Consider, for example, a worker whose job might be to sort packages of varying size and weight into bins, bags or slots. Each of these packages will have different inertial properties and will impose different loads upon the arm as it moves towards the desired target position. If the worker carries out this task for a prolonged time, is it reasonable to expect some adaptation to take place? In this case, the perturbations are not fixed but vary from object to object and follow a given statistical distribution depending both on the object properties and

upon the sequence of movements in the task. Can the motor system adapt to a variable environment? And if so, how is this adaptation accomplished? Does the motor system use information it acquires on a trial-by-trial basis, or does it attempt to extract some definable statistical property about the perturbations it encounters, such as the mean or the most likely (i.e. the mode) perturbation? Can subject behavior in a stochastic environment reveal how the neural mechanisms involved in motor adaptation use information from previous trials to modify motor commands on subsequent trials? These questions were addressed in a set of experiments that employed engineering methods of systems identification and a robotic system to generate sequences of perturbing force fields having magnitudes that varied randomly from trial-to-trial.

In the present experiments, adaptation was examined in the context of goal-directed reaching movements. Twelve subjects executed reaching movements between two targets in the horizontal plane while holding the handle of a two-joint robotic manipulator. The robot applied perturbing forces to the arm during each movement. The amplitude (but not the direction) of the perturbing force field varied randomly from trial to trial. Each subject's motor response to the sequence of perturbing fields was quantified using the peak deviation from a straight-line hand path. The trial-to-trial sequences of motor errors were analyzed and the results demonstrated that subjects did adapt their motor behavior in response to the random sequences of force fields presented at the hand. Furthermore, subjects compensated for the approximate mean field of the stochastic sequence. This behavior did not depend on specific distribution properties of the sequence. Finally, subjects accomplished this adaptation by using memories of the most recent perturbations and the most recent performances only. Adaptation was not accomplished by directly counteracting the mean field strength on each individual trial. The present findings are consistent with recent experiments that suggested a prominent function of prefrontal cortex in the early stages of motor adaptation to perturbing fields (Shadmehr and Holcomb 1997). Portions of this work have appeared in abstract form (Scheidt et al. 1999).

Methods

Twelve human subjects with no known neuromotor disorders consented to participate in this study. Subjects execute half-second, 20 cm reaching movements with their dominant arm in the horizontal plane while holding the handle of a two-joint, robotic manipulator (Figure 1A). The robot was comprised of a 5-bar linkage with torque

motors controlled by a dedicated PC (Scheidt et al. 2000). Subjects were instructed to “reach from the ‘beginning target’ to the ‘ending target’ in one half second.” The computer provided qualitative feedback of movement duration after each trial (either too fast: <0.45 s, too slow: >0.55 s, or just right: 0.45 to 0.55 seconds). Subjects were instructed to relax after each movement while the manipulandum moved the hand slowly back to the beginning target. This protocol was designed for allowing subjects to experience the limb’s mechanical environment along a limited set of trajectories. Reaching movements were directed away from the subject’s body along a line (the positive Y-axis) passing through the center of rotation of the shoulder. The subjects’ arms were supported against gravity by a sling attached to the 8-ft ceiling. The support was adjusted so that the upper arm was abducted by 90 degrees. The shoulders were restrained using a Velcro torso support. “Beginning” and “end” targets corresponding to a 20 cm reach in the plane of the arm were presented on a computer monitor just above the manipulandum. The position of the hand was displayed as a small cursor on the overhead monitor. Subjects could see both their arm and the cursor representing it at all times.

Insert Figure 1 Near Here

The robotic manipulator applied perturbing force fields to the arm during each movement. A *Perpendicular Viscous Field* was designed to deflect the hand perpendicularly from its intended path with a force proportional to hand velocity along its path (Figure 1B). The forces applied to the subject’s hand during the i -th movement were defined:

$$\begin{bmatrix} F_X \\ F_Y \end{bmatrix} = B_i \begin{bmatrix} 0 & -1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (1)$$

where \dot{x} and \dot{y} were the two components of the hand velocity along the medial/lateral (x) and proximal/distal (y) directions, F_X and F_Y were the two components of the force applied by the robot along the same directions. B_i was a random real number between 0 and 30 Newton second/meter such that the amplitude (but not the direction) of the perturbing force field varied randomly from trial to trial. Movements were always made along the positive y -direction and perturbing forces were always directed to the left. It must be stressed that during each movement, subjects experienced variable forces that depended linearly on their instantaneous hand speed. However, the

magnitude of the environmental impedance, B_i , remained constant for the duration of each individual movement and changed only between trials. Subjects could experience peak hand forces up to 30 or 40 N in this field. Subjects could perform the task easily in the time allotted, however reaching accuracy was influenced by the perturbations. During each trial, instantaneous hand positions were recorded using rotational encoders on the robot's motors and hand forces were recorded using a 6 degree-of-freedom load cell mounted at the handle of the robot.

Insert Figure 2 Near Here

Two stochastic perturbation sequences were used. In experiment 1, four subjects were presented with a sequence of 200 trials in which the force field gain, B_i (Figure 2A), followed a Gaussian distribution (Figure 2B). This distribution had a non-zero mean corresponding to information about the perturbation sequence that subjects might learn. The mean perturbation amplitude was 15.2 Ns/m with a variance of 24.7 Ns/m. This sequence was designed to insure insignificant correlation between perturbation magnitudes on consecutive trials separated by more than 40 trials (Figure 2C). The significance of each correlation term was evaluated by comparing the correlation magnitude at each integer lag value to an estimate of the 95% confidence interval bounding zero correlation ($2\sigma \cong 2/\sqrt{N}$; Box et al. 1994). All four subjects were exposed to the same sequence of perturbations. In experiment 2, eight subjects were presented with a sequence of 400 trials with a bimodal probability density function (Figure 2D and 2E). This bimodal sequence was constructed by shuffling together two unimodal sequences with individual Gaussian distributions having means of 6 Ns/m (175 trials) and 25 Ns/m (225 trials), respectively. While each individual sub-population contained no significant correlations between perturbations separated by as many as 40 trials, the shuffling process used to construct the bimodal population gave rise to spurious correlations at trial lags < 40 trials (Figure 2F). All eight subjects were exposed to the same sequence of perturbations. This bimodal stochastic perturbation sequence had greatly differing mean (15.5 Ns/m) and mode (25 Ns/m) values and was constructed to distinguish whether subjects adapt more closely to the mean or the mode of a given perturbation sequence, or whether adaptation gets "trapped" in the smaller, local maximum designed into the bimodal probability distribution function).

Data Analysis

Simple measures of kinematic and kinetic behavior were used to assess subject motor performance on each

trial during this goal-directed reaching task. “Movement error” was defined as the peak deviation of the hand from a straight-line trajectory passing between the initial and final targets (Krakauer et al. 1999). Movement error was used to quantify *kinematic* performance, assuming that subjects intended to make straight-line movements of their hands. This measure of motor performance has previously been found to motivate motor adaptation during reaching (Scheidt et al. 2000). The peak hand force that was generated perpendicular to the direction of movement quantified *dynamic* performance.

An exponential function was fitted to the trial series of movement errors to characterize the rate at which subjects compensated for the random sequence of perturbation gains. This model had three free parameters: gain, A , time-constant, I , and offset, C :

$$E_i = Ae^{(-i/I)} + C \quad (2)$$

where E_i was the computed movement error on trial i . The exponential captured the overall rate of change in movement error, while the constant (C) described any steady-state bias in these time series. The free parameters of this model were fit using a simplex search algorithm (Press et al. 1988).

A regression analysis of movement error versus perturbation amplitude was performed to determine the field strength (i.e. perturbation gain) that subjects adapted to. The strength of correlation between these two variables and the linearity of this relationship were also evaluated. The amplitude of the field strength to which subjects adapted was estimated from the zero crossing of the resulting regression line, since the perturbation gain value at which the regression line passed through zero error indicated the field strength at which subjects would exhibit error-free (straight line) trajectories. This analysis, however, provides no explanation for *how* subjects adapted to this particular field strength. Subjects could adapt by directly counteracting this ‘zero error’ perturbation magnitude itself on each and every trial; i.e. by executing a control strategy that anticipated the same constant field trial after trial. If so, movement errors would vary linearly with perturbation strength. Alternatively, subjects could employ a continuously evolving strategy of using information about perturbations and movement errors from a limited number of previous trials to adjust performance on subsequent trials. Because such a strategy could also result in a linear relationship between movement error and perturbation strength, the regression analysis described above could

not distinguish these two possibilities.

The above regression analysis was extended to evaluate the dependence of movement errors on previous perturbations and previous errors using autocorrelation and cross-correlation analyses. Specifically, the autocorrelation profile of the movement error trial sequence and the cross-correlation between the error and perturbation gain trial sequences were calculated. If subjects anticipate a constant field strength when exposed to an uncorrelated sequence of perturbations (e.g. the mean field strength), then their performance on each trial must also be uncorrelated with that of previous trials. This hypothesis was directly tested by this analysis.

These correlation analysis results were then used to guide construction of a model of motor adaptation during reaching. Specifically, movement error on each trial was modeled as a linear combination of previous movement errors as well as present and previous perturbation amplitudes. The result was a parametric model of motor adaptation that was linear in its inputs:

$$\mathbf{e}_i = \sum_{j=1}^{i-L} a_j \mathbf{e}_{i-j} + \sum_{k=0}^{i-M} b_k B_{i-k} \quad (3)$$

where a_j and b_k were coefficients weighting the relative importance of previous errors (\mathbf{e}_{i-j}) and previous perturbation magnitudes (B_{i-k}) on subsequent errors, j and k were indices of summation while L and M were limits on the number of significant terms in the model. Values for L and M were obtained directly from the correlation analysis. This model represents an autoregressive process with external input (i.e. an **ARX** model; Ljung, 1998). Terms with non-zero a_j coefficients are *autoregressive* terms because they define the dependence of the current movement error on previous movement errors. Terms with non-zero b_k coefficients are *moving average* terms of the external input because they define the dependence of the current movement error on a sliding window average of current and previous perturbation amplitudes. Because Equation 3 defines a discrete-time difference equation, the stability and steady-state behavior of this model of motor adaptation was analyzed using z-transform techniques (Oppenheim and Schaffer, 1989).

The capacity of this model to predict subjects' adaptation to the random sequence of perturbations was

evaluated. These predictions were compared to the predictions of two alternate, but viable, learning algorithms. The first alternative model accumulated an explicit representation of the mean perturbation strength by “memorizing” the perturbation sequence trial-by-trial. This explicit representation of the running-average mean perturbation was used to compensate for the perturbation on the next trial. The second alternative model was an incremental learning algorithm that utilized local weighting of the most recent movement errors to predict and compensate for the magnitude of the next perturbation. This model included the possibility of non-uniform and non-linear attention models whereby learning could either attend closely to or ignore trials where the perturbation amplitude was “surprising” or “irrelevant” (Atkeson et al., 1997). Each model was first fit from the subjects’ data from the initial 100 movements in the experiments. The performance of each model was then evaluated according to its ability to predict subject movement errors on the last 100 movements. Model performance was quantified using the variance accounted for (VAF) as a measure of goodness-of-fit:

$$VAF = 1 - \frac{\text{var}(\text{residuals})}{\text{var}(\text{data})}. \quad (4)$$

Results

Subjects compensate for the approximate mean of the random trial sequence of perturbations

An overhead view of averaged hand paths made during experiment 1 (Figure 3A; unimodal perturbation sequence) shows that subjects exhibited substantial kinematic deviations to both the left and right even though they experienced forces that pushed only to the left. To compare across trials, hand path data were aligned with respect to the onset of movement (the point in time when hand speed first exceeded 0.1 m/s; Figure 3B) and averaged into six “bins” of 5 Ns/m width each (0-5, 5-10, 10-15, etc.). Movements from trials with field strengths greater than 20 Ns/m resulted in trajectories that deviated markedly to the left (i.e. in the direction of the applied force). However, hand path deviations were consistently toward the right (i.e. opposite to the direction of the applied force) for fields with gains less than 10 Ns/m. Movements made in weaker fields had hand path errors that were approximately mirror symmetric to those made in stronger fields. Kinematic errors made in the weakest fields were nearly identical to responses observed when perturbing force fields were unexpectedly removed after adaptation (Scheidt et al, 2000). This finding is consistent with traditional measures of aftereffects of adaptation (Shadmehr and Mussa-Ivaldi, 1994). These aftereffects are a clear indication that subjects compensated for the perturbations by adopting some

automatic and predictive mechanism. Force fields roughly corresponding to both the mean (average) and mode (most likely) disturbance (10 to 15 Ns/m) resulted in movements with the least curvature. Note that these movements were only approximately straight, corresponding to the steady-state bias in movement error (constant C in equation 2).

Insert Figure 3 Near Here

Average hand speed profiles (Figure 3B) typically demonstrated two distinct peaks. The secondary peak in the hand speed profile could be the result of several mechanisms including, but not limited to, active correction under visual feedback, reflex mediated adjustments due to the mismatch between the intended and actual final joint posture (the mismatch being due to the perturbation), or the interaction of limb and manipulandum dynamics (Shadmehr and Mussa-Ivaldi, 1994). The present experiment was not designed to distinguish between these alternatives. Therefore, movements were truncated at the point of time near the end of movement where the hand speed profile reached a transient minimum (vertical line in Figure 3B, solid lines in Figure 3A returning to final target location). This limited subsequent analysis to the portion of movement that was predominantly feedforward.

An exponential function (equation 2) was fit to the movement error trial series (dark solid line in Figure 3C) confirming the presence of a steady-state bias in movement error. Figure 3C shows a rapid decrease in movement error within the first 10 to 20 trials (time constant = 2.4 trials). Time constants (λ in Equation 2) for all four subjects averaged 3.2 ± 0.74 trials (mean \pm SE mean). The residual steady-state movement error (constant C in equation 2) was observed in all four subjects (average 12.3 ± 2.7 mm), indicating that subjects compensated only approximately for the mean of the random trial sequence. These observations were consistent across all four subjects exposed to the unimodal perturbation sequence. Subjects did indeed adapt in response to the random sequence of perturbations.

Profiles of the hand forces generated perpendicular to the direction of movement (Figure 3D) provide further evidence of adaptation to the stochastic sequence of perturbations. Subject-generated forces dominated movements made in the weakest fields (the smallest force profile with biphasic shape), whereas robot-generated forces dominated movements made in the strongest fields (the largest profile with monophasic shape). The initial peak in perpendicular force generated by the subject in the weakest field (approximately 12 N) was directed opposite to the

forces imposed by the robot and was not necessary to move the hand toward the target. This excessive force caused the limb to deviate substantially from the target, producing a *kinetic aftereffect of adaptation*. Consequently, restoring forces (the negatively-directed peak in Fig 3E) were required to move the limb to the final target.

An analysis of movement error versus perturbation amplitude (Figures 3E & 3F) indicated that these two variables were well fit by a linear relationship within the range of our experiment ($r = 0.82$ in Figure 3E; $0.73 < r < 0.84$ for all four subjects). The point of zero error on these regression lines indicates the field strength that was best compensated for through the adaptive process (13.5 Nm/s in Figure 3E; 12.9 ± 1.2 Ns/m for all four subjects in Figure 3F). This adapted field strength approximated, but did not quite attain, the mean value of the distribution ($p < 0.01$; students t-test rejecting the null hypothesis $H_0: B_{adapt} = \bar{B} = 15.2$ Ns/m). Thus, subjects compensated for force fields somewhat less than (but coarsely approximating) the mean gain (i.e. $B_i \approx \bar{B}$), greatly under-compensated large force field gains (i.e. $B_i \gg \bar{B}$), and greatly over-compensated small force field gains (i.e. $B_i \ll \bar{B}$).

Insert Figure 4 Near Here

Adaptation to the approximate mean, not to the mode

An overhead view of averaged hand paths made by one subject from experiment 2 (Figure 4A; bimodal perturbation sequence) shows that again, movements were either deflected to the left or to the right. Average hand speed profiles (Figure 4B) exhibited the same biphasic pattern found in Experiment 1. Consequently, the data from experiment 2 were truncated in same way as in experiment 1. The truncated and averaged hand movements (Figure 4A) exhibited consistent deviations towards the right when movements were made in fields with strengths of ≤ 10 Ns/m and towards the left when movements were made in fields with strengths of ≥ 20 Ns/m. Force fields roughly corresponding to the mean disturbance (10 to 15 Ns/m) resulted in trajectories with the least curvature although they were not ideally straight. *These results demonstrate that adaptation to a sequence of perturbations with randomly varying magnitudes converges to the approximate mean perturbation magnitude rather than the most likely magnitude.*

Fitting an exponential function (equation 2) to the movement error trial series (dark solid line in Figure 4C)

produced results that varied widely across subjects (time constant $\lambda = 54.8 \pm 18.2$ trials (mean \pm SE mean); range = [11.2, 167] trials; N=8). Consequently, this traditional measure of learning suggests that the bimodal perturbation sequence abolished (or at least slowed) the initial rapid learning observed in experiment 1. However, as shown below, subject performance in the both experiments can be described using a single, parsimonious description of motor adaptation.

Figure 4D displays average perpendicular hand force profiles measured at the handle in the bimodal experiment. The overall shape of the profiles was similar to those observed in the unimodal experiment, with the initial peak in perpendicular force generated in the weakest fields giving rise to the undesirable deviation from the target indicative of aftereffects of adaptation. Again, this is a kinetic aftereffect of adaptation similar to that observed in the first experiment. As in the unimodal experiment, the hand force profiles were smooth and the restoring forces generated at the end of the movement did not appear to be distinct pulses.

As in the unimodal experiment, subjects exposed to the bimodal perturbation sequence exhibited a linear relationship between movement error and perturbation strength ($r = 0.85$ in Figure 4E; $0.83 < r < 0.94$ for all 8 subjects). Again, the point of zero error on these regression lines was taken as the field strength that was best compensated for through the adaptive process. These 8 subjects adapted to an average field strength of 11.33 Ns/m with a 95% confidence interval of [8.61, 14.04] Ns/m (Figure 4F). For all 8 subjects, the major and minor peaks of the bimodal probability density function (6 Ns/m and 25 Ns/m, respectively) both fell outside this 95% confidence interval. However, while the adapted field strength in the bimodal sequence was substantial, subjects did not quite compensate for the mean perturbing field ($\bar{B} = 15.5$ Ns/m), even after 300 to 400 movements. Somewhat paradoxically, subjects actually adapted to a perturbing field strength that was experienced rarely.

Insert Figure 5 Near Here

Only recent memories contribute to adaptation:

The Gaussian-distributed random trial sequence (Figures 2A and 2B) was used to perturb subjects while adapting because this input to the motor adaptation process was both uncorrelated from trial to trial (up to a lag of 40 trials; Figure 2C) and “rich” spectrally (Marmarelis and Marmarelis 1978). Driving each subject’s motor system

with an uncorrelated trial sequence ensured that any trial-to-trial correlations observed in that subject's motor output did not originate from the perturbation sequence, but rather from information processing within the neuromotor controller. Despite this lack of correlation in the sequence of perturbing fields, significant trial-to-trial correlations were observed in subjects' motor output (Figure 5). Correlations between movement error and perturbation gain (Figure 5A) exceeded statistical significance ($> 95\%$ CI) not only on concurrent trials (i.e. lag = 0), but also on the preceding trial (lag = +1). The sign of the correlation at lag 1 was opposite that at lag 0, indicating that subjects attempted to reduce movement error on each trial by countering the previous perturbation. Significant correlations between movement error and perturbation gain extended back no more than 2 trials for all of the 4 subjects exposed to the perturbation sequence with the unimodal distribution. Movement errors on a given trial also exhibited substantial correlations with errors generated in the preceding trial (Figure 5B). By definition, the autocorrelation function is symmetric about 0 lag. Thus, correlations at any given lag are reflected at the corresponding lead with no violation of causality. Significant autocorrelation terms were found at a lag of 1 trial for 3 of the 4 subjects (the remaining subject showed no significant correlations beyond lag 0). Again, the sign of this correlation was negative indicating that subjects attempted to reduce movement errors on each trial by countering movement errors generated on the previous trial.

Similar correlation analyses were performed on the peak hand forces generated perpendicular to the intended direction of movement (Figures 5C and 5D). Correlations between peak hand force and perturbation gain (Figure 5C) exceeded statistical significance ($> 95\%$ CI) only on concurrent trials (i.e. at zero lag) and at a lag of one trial. Significant correlations between peak hand force and perturbation gain extended back no more than 2 trials for all 4 subjects in experiment 1. Correlations between movement error and peak hand force (Figure 5D) exceeded statistical significance only on concurrent trials and at a lag of one trial, a result entirely consistent with the findings of Figure 5A.

The significant correlations at zero lag (Figure 5A, 5C, and 5D) were due in part to mechanical interaction between the robot and the finite impedance of the subject's arm. Larger forces imposed by the robot on the hand resulted in both larger deviations of the hand from its intended path and in greater forces being recorded at the handle. However, significant correlations at non-zero lags cannot be explained by mechanical interactions. These lag

1 correlations indicate that subjects used explicit information regarding the strength of the perturbation from the previous trial to preprogram the motor response on each subsequent trial. Subjects also utilized information about previous *performance* to update motor behavior on subsequent trials. The lack of significant correlations beyond two previous trials indicates that explicit memory representation of more remote trials was not used during motor adaptation. Had subjects adapted to the stochastic sequence of perturbations by directly counteracting some constant field strength on each trial (e.g. the mean), their motor output would likewise be uncorrelated with the input. This hypothesis is clearly refuted by the present findings. These findings support the hypothesis that explicit information of only one or two previous trials is sufficient to allow subjects to compensate for the approximate mean field strength in a random sequence of perturbations.

Insert Table 1 Near Here

Insert Figure 6 Near Here

Predicting motor performance

The preceding analyses suggest that subject performance (quantified by movement error) exhibited on any given trial, i can be predicted solely from the field strength on that trial (B_i) and from the field strength and error exhibited on the previous trial (B_{i-1} and e_{i-1} , respectively):

$$e_i = a_1 e_{i-1} + b_0 B_i + b_1 B_{i-1}. \quad (5)$$

Regression coefficients (a_1 , b_0 , b_1) were estimated over the initial 100 trials for each subject in experiments 1 and 2 by performing a multi-linear regression of movement error on previous movement errors, previous perturbations and concurrent perturbations (Table 1). The ability of this model to predict movement error in the last 100 movements in each experimental session was evaluated to determine how well the model would generalize beyond the data set used to determine model coefficients. The ability of this model to predict performance for greatly differing perturbations sequences (i.e. unimodal vs. bimodal distribution) was also evaluated. Model performance was quantified using the variance accounted for (VAF; Equation 4). The percentage of VAF by this model from the unimodal experiment was 79% for the subject shown in Figure 6A and $71 \pm 3\%$ (mean \pm 1 SE mean) for all four subjects. The percentage of VAF by this model from the bimodal experiment was 86% for the subject shown in Figure 6B and $84 \pm 2\%$ (mean \pm 1 SE mean) for all eight subjects. Thus, a model incorporating limited explicit

memory of subject performance and perturbation magnitudes can predict movement errors with a high degree of fidelity. This strongly suggests that *motor adaptation is a continuously evolving process whereby the average field in a sequence of perturbations is compensated using properly weighted cancellation of both previous perturbation amplitudes and previous movement errors.*

Step response analysis

Although the model performed quite well in response to the stochastic input sequences from which it was originally derived, it was also important to determine how well this model could predict behavior exhibited in response to non-stochastic sequences of input perturbations. Equation 5 was used to simulate movement errors in response to a step increase in perturbation strength that included a simulated “catch trial” near the end of the input sequence (top trace in Figure 7A). This input sequence was specifically designed to mimic the constant force field gains and catch trials used in previous motor adaptation experiments (e.g. Shadmehr and Mussa-Ivaldi, 1994). Average coefficient values from the unimodal experiment (Table 1) were used to define the model parameters. When presented with a step increase in perturbation strength, the simulated movement errors rapidly approached their asymptotic value (within 3 to 4 trials) and exhibited a small steady-state error at large trial numbers, as did subjects in both experiments. The model output also exhibited the classic behavior of an “after-effect” (Shadmehr and Mussa-Ivaldi, 1994) when the catch trial was introduced at trial number 75. Furthermore, this model was also able to account for the observations of Thoroughman and Shadmehr (2000) that a single catch trial can transiently degrade the adapted state generated in response to a consistent perturbing field. Consequently, this very simple model of motor adaptation succinctly captures the fundamental behavioral characteristics exhibited in both the present experiment and in more traditional experiments of motor adaptation.

Insert Figure 7 Near Here

Interpretation of Model Coefficients

To investigate how each term of the model related to observed behaviors, the output of the model in response to the same step input (Figure 7A) was analyzed for various combinations of model parameters (Figure 7B). A model that does not rely on prior experience, and thus has no memory (i.e. $a_1 = b_1 = 0$; trace 1), can only respond to the current perturbation and fails to adapt. This is the response one would expect if subjects were only co-

contracting their limb in response to the perturbations. Increased co-contraction might decrease the magnitude of the b_0 term, but unless the limb stiffness became exceedingly large, a substantial residual offset would remain. The amount of residual steady-state error in the adaptive response is determined by relative magnitudes of b_0 and b_1 . When $b_0 = -b_1$ (trace 2), then the residual steady-state error is eliminated. On the other hand, when the autoregressive term in the model is removed ($a_1 = 0$; trace 3), the dynamics associated with initial exposure to the perturbation are eliminated while the steady-state error is reduced relative to the full model (Figure 7A). If this autoregressive term is instead doubled (trace 4), the initial transients are extended and the steady-state error is increased. Changing the sign of a_1 (traces 5 and 6) does not alter the time course of adaptation, but causes the model's response to oscillate within the envelope defined in Figures 7A (bottom) and 7B (trace 4) respectively. Note however that changing the sign of a_1 does in fact reduce (but not eliminate) the steady-state error. Finally, not all choices of parameters yield stable learning. Setting $a_1 > 1.0$ yields an unstable algorithm that never adapts (trace 7).

Discussion

The present experiments investigated the ability of unimpaired humans to adapt to a viscous, perpendicular, force field environment having force field gains that were unpredictable (and uncorrelated) from trial-to-trial. Experiments were designed to determine if subjects adapted to the mean force field gain, the most likely field, or whether adaptation would depend on other features of the perturbation sequence's probability density function. Correlation analyses were performed to determine how much motor performance on any given trial was correlated with performance on previous trials. It was found that: a) subjects adapted their motor behavior in response to the random sequence of force field gains, b) subjects compensated for the approximate mean field of the stochastic sequence, not the most likely field, and c) subjects compensated using memories of only the most recent perturbations and the most recent performances.

What did subjects adapt to?

Subjects experienced perturbing forces that were always directed toward their left. If no adaptation had occurred, then all movement errors would likewise have been directed toward the left. However, for field strengths of less than 10 Ns/m, hand path deviations were made consistently toward the right in both experiments (Figures 3A and 4A). The presence of these oppositely-directed errors (i.e. to the right) indicates that subjects were directly

opposing forces they anticipated encountering and precludes the possibility that they were merely stiffening the arm around some reference trajectory (Flash 1987; Shadmehr and Mussa-Ivaldi 1994; Conditt et al. 1997a). Furthermore, movements made in stronger-than-average force fields were under-compensated while movements made in weaker-than-average force fields were over-compensated, suggesting that subjects were compensating approximately for the mean perturbing force field in both experiments. This finding was confirmed by linear regression analysis (Figures 3E and 4E).

Learning rates in the present study (Figures 3C & 4C) were *slower* than rates reported for compensation of inertial loads (within 1 trial; Bock, 1993), but were substantially *faster* than learning rates reported for consistent (but geometrically complex) viscous environments when subjects were required to reach in several different directions (more than 100 trials; Bhushan and Shadmehr 1999). Remarkably, the learning rates observed in the experiment 1 were almost identical to the rates at which subjects regained adaptation to a predictable perturbing environment after a single “catch trial” in which the perturbing environment was unexpectedly removed (about 3 trials; Thoroughman and Shadmehr 2000). Clearly, subjects adapted to these stochastic environments.

Adaptation modeled as an autoregressive process with external input (ARX process)

Although the linear regression results demonstrated that subjects compensated for the approximate mean perturbation strength in both force field environments, they did not suggest *how* the central nervous system accomplished this adaptation. Mathematically, the mean perturbation magnitude is defined as the sum the individual magnitudes divided by the number of perturbations. Since subjects had no way of knowing all of the perturbation magnitudes until the experiment was completed, it was not possible for them to directly compute the mean field strength. Subjects could have evaluated a “running average” of all trials they experienced so far. However, this strategy would require subjects to retain either explicit working memory of *all* previously encountered perturbations or explicit memory of the average of all previous perturbations along with a running total of the number of previous perturbations. In either case, the relative importance of the most recent perturbation would decrease linearly as a function of the number of perturbations.

A less demanding alternative would be for subjects to rely only on explicit information regarding only recent

experiences. Motor performance (and consequently motor adaptation) may rely on information about past movement performance and/or past perturbations derived from a variety of sensory sources (eg. muscle spindles, Golgi tendon organs, slowly-adapting hand mechanoreceptors, vision, etc.). A general form of this model, one that depends only on information regarding movement errors and perturbation amplitudes (Equation 3) was examined in the present experiment. One important aspect of this model is that information about experiences in the distant past is retained implicitly in the autoregressive terms (i.e. if Equation 3 contains at least one non-zero a_j term).

The correlation analyses (Figure 5) demonstrated that movement error on a given trial i was well predicted from the field strength on that trial (B_i) and from the field strength (B_{i-1}) and movement error (e_{i-1}) exhibited on the previous trial (Equation 5). Why do subjects compensate for previous movement errors when the step response analysis suggests that learning would be more rapid and effective if those errors were disregarded altogether (i.e. set $a_1 = 0$ in Equation 5) and the most recent perturbation was canceled exactly (i.e. set $b_1 = -b_0$ in Equation 5; Figure 7B, trace 2)? Are there unavoidable history dependencies in the proprioceptive and/or visual sensory pathways that constrain the motor learning mechanisms in their “choice” of compensatory strategies? It has been suggested that cancellation of prior movement errors is important to motor adaptation (eg. Flanagan and Rao, 1995; Wolpert et al., 1995; Scheidt et al., 2000). Equation 5 suggests that compensating for the most recent movement error *exactly* (i.e. $a_1 = -1$ in Equation 5) would be a counter-productive strategy since Equation 5 becomes unstable when $|a_1| \geq 1$. This can be seen by examining the stability of Equation 5 in the complex z -domain (Oppenheim and Schaffer, 1989). The z -transform of equation 5 is:

$$E(z) = a_1 z^{-1} E(z) + b_0 B(z) + b_1 z^{-1} B(z). \quad (6)$$

Consequently, the model's transfer function $H(z)$ is:

$$H(z) = \frac{E(z)}{B(z)} = \frac{(b_0 + b_1 z^{-1})}{(1 - a_1 z^{-1})}. \quad (7)$$

$H(z)$ has a zero at $z = -b_1/b_0$ (see Figure 7B, trace 2) and a single, real pole situated at $z = a_1$. The location of the zero in the unimodal and bimodal experiments was not significantly different at the $p < 0.05$ level ($p = 0.39$; two-sample

t-test) and the location of the system pole in the two experiments was also not significantly different ($p = 0.13$). Therefore, to the extent that this linear model captures the mechanisms of adaptation, we conclude that the process of adaptation is not sensitive to the details of the distribution of the perturbing forces (eg. mode, skewness, etc.), but only to its mean. The location of the zero at $z = -b_1/b_0$ is significant because in the steady state (i.e. when $z = 1$) the transfer function is minimized when $b_1 = -b_0$. Furthermore, for any linear system to be stable, all the *poles* of its transfer function must lie within the unit circle defined in the complex z -plane (i.e. $|z| < 1$; Oppenheim and Schaffer, 1989). Therefore $|a_1| < 1$ must be satisfied for Equation 5 to be stable. Perfect cancellation of the most recent movement error (i.e. $a_1 = 1$) would cause the motor adaptation process to become unstable (e.g. Figure 7B, trace 7).

Explicit representation of the internal model

Since the relationship between movement error and perturbation gain was reasonably well fit by a straight line (Figures 3E and 4E), the ARX model of subject performance (Equation 5) was rearranged to yield an expression for the internal model of the perturbing environment. Specifically, movement error generated on trial i was regarded as a function of the mismatch between the actual perturbation experienced on that trial and the expected (or adapted) perturbation magnitude: $e_i = f(B_i - B_{adapted_i})$. Figures 3E and 4E demonstrate that this relationship was reasonably described by a linear function:

$$e_i = k(B_i - B_{adapted_i}) \quad (8)$$

Re-arranging Equation 5 into the form of Equation 8 yields:

$$e_i = b_0 \left(B_i + \frac{a_1}{b_0} e_{i-1} + \frac{b_1}{b_0} B_{i-1} \right) \quad (9A)$$

$$= b_0 (B_i - B_{adapted_i}) \quad (9B)$$

where:
$$B_{adapted_i} = -(a_1/b_0)e_{i-1} - (b_1/b_0)B_{i-1}. \quad (10)$$

Equation 10 provides a very simple representation of the subject's prediction of the perturbation magnitude on trial i based solely on explicit information about the error and perturbation magnitude on the most recent trial. However,

even this simple representation faithfully reproduced subjects' behavior in the present experiment (Figure 6). It is worth noting that parameters estimated from the first 100 trials adequately accounted for the time series of errors up to 300 trials following parameter estimation. Thus, the adaptive behaviors observed in both the unimodal and bimodal experiments (Figures 3 and 4) were a consequence of the dynamics of a quasi-stationary process with very limited memory. Motor performance would be optimized by tuning the coefficients a_1 , b_0 , and b_1 . Both the rate of adaptation and the steady state error can be altered by appropriate modification of these coefficients (Figure 7).

How do the motor adaptation mechanisms estimate the most recent movement error and perturbation strength?

Constructing the internal representation of the perturbing field strength via equation 5 requires accurate estimation of e_{i-1} and B_{i-1} . Movement error is likely to be sensed both visually (e.g. Wolpert et al., 1995) and proprioceptively (Shadmehr and Mussa-Ivaldi, 1994; Dizio and Lackner, 2000). The current experiments were not designed to evaluate the relative contributions of different feedback modalities to motor adaptation, but rather to explore how the neural mechanisms involved in motor adaptation use information from previous movements (however that information is sensed) to modify motor commands on subsequent movements. Both visual and proprioceptive feedback appear to be important (Conditt et al., 1997b), although it is not yet clear how this feedback information is combined in driving motor adaptation.

There are at least two ways the central nervous system could estimate the most recent perturbation strength in keeping with the spirit of equation 10. The first strategy would be to estimate the field strength directly using sensory organs sensitive to the kinetic demands of the task (e.g. Golgi tendon organs, hand mechanoreceptors, and indirectly, muscle spindle receptors since they are coupled to the perturbation through limb tissues with finite impedance). In this case, B_{i-1} would be "measured" directly and equation 10 would be implemented as written. The second way of estimating the most recent perturbation strength would be to do so indirectly and recursively, using only the most recent movement error to update the previous estimation of the perturbing field. In this case, equation 10 can be reformulated:

$$B_{adapted\ i} = - (a_1 / b_0) e_{i-1} - (b_1 / b_0) B_{adapted\ i-1} \quad (11)$$

Here, movement errors drive the formation of the internal model of the perturbations. Substituting the z-

transform of equation 11 into the z -transform of equation 9B for $B_{adapted}(z)$ yields a transfer function that is identical in structure to equation 7 except that the algorithm's single pole location is shifted to $z = a_1 - (b_1/b_0)$. The coefficient values that ensure stability of the system are therefore: $|a_1 - (b_1/b_0)| < 1$. Using equation 11 and the modified equation 9B to fit subject U2's movement error data yields the coefficients: $a_1 = -0.34$, $b_0 = -4.2$, $b_1 = 4.0$. The effective pole location for this system was $z = 0.62$ (compare to Figure 7B, trace 4). Since recursive estimation of B_{i-1} via equation 11 does not alter the transfer function structure of equation 7, the correlation analyses of Figure 5 cannot distinguish between recursive estimation of $B_{adapted}$ (equation 11) and estimation of $B_{adapted}$ via proprioception of B_{i-1} (equation 10). Consequently, while force feedback gain from Golgi tendon organs is likely to be quite low (cf. Houk and Rymer, 1981), such information is not necessary to construct an internal model of the perturbing force field environments.

Insert Figure 8 Near Here

Comparison of the simple autoregressive model with alternative learning strategies

If movement error is linearly **related** to the perturbing field amplitude, $e_i = k(B_i - B_{adapted})$, then the optimal internal model in terms of least square error (i.e. the $B_{adapted}$ that minimizes the sum of e_i^2) is given by the mean field:

$$B_{optimal\ i} = \frac{1}{i-1} \sum_{j=1}^{i-1} B_j. \quad (12)$$

In the present study, movement errors were linearly related to perturbation amplitude (Figures 3E and 4E). This finding supports the empirical claim that subjects adapted their reaching movements so as to minimize deviations from a rectilinear path (Shadmehr and Mussa-Ivaldi 1994). One plausible alternative to the learning strategy described by Equation 5 would be for subjects to "explicitly" learn the mean. In this case, equation 12 is substituted into equation 5 for B_{i-1} while the dependence on previous movement errors (the e_{i-1} term) is dropped:

$$e_i = b_0 B_i + b_1 B_{optimal}. \quad (13)$$

Equation 13 was fit to subject U2's initial perturbation and movement error data set (first 100 movements). The model's performance was then evaluated in the final 100 movements of the experiment (Figure 8A). This algorithm performed respectably when perfect memory of all perturbations encountered was available (VAF = 70%). This algorithm also quickly converged to the (ideally) linear relationship between perturbation gain and movement error (Figure 8B). However, this model failed to exhibit the significant lag 1 correlations between movement error and perturbation amplitude observed experimentally (compare Figure 8C to Figure 5A). The absence of these lag 1 correlations arises from the fact that when all previous movements are considered explicitly, the contribution of any one individual perturbation to the internal model of the mean field becomes vanishingly small after only a few trials. Therefore the simple adaptive mechanism of equation 10 permits subjects to achieve near optimal performance (i.e. approximately minimizing squared errors) while maintaining computational efficiency since the autoregressive equation 10 approximates the ideal average of equation 12 with a very limited number of memory elements. (The autoregressive term in equation 10 retains indirect and exponentially weighted memory of all past perturbations due to the nested dependence of ϵ_{i-1} on previous errors). Note also that equation 12 defines optimal performance only when the underlying distribution of B_i is stationary. In the non-stationary case (such as the step input of Figure 7A) the model of equation 12 would respond very slowly, since the incremental contribution from each subsequent trial decreases progressively as the number of trials increases. Consequently, adaptation in the sense of Equation 5 strikes a desirable compromise between decreasing movement errors in a stationary but potentially unpredictable environment while allowing the motor system to respond rapidly and appropriately to long-term changes in the perturbing environment (Figure 7A).

Insert Figure 9 Near Here

A second alternative learning strategy describes what could be called "careless learning" and was motivated in part by the observation that subjects never made ideally straight movements and almost always had peak hand deviations exceeding ~1 cm (Figures 3E and 4E). Perhaps subjects considered movements with such small errors "good enough" for the specified task? This form of learning is careless in the sense that the learner does not attend to small movement errors. The internal representation of the perturbation in a careless learning model would be updated only when the learner is "surprised" (i.e. when the movement errors experienced on a given trial exceed a minimum threshold value, $\epsilon_{threshold}$). An attention model that describes how well movement errors are attended to

is:

$$U(\mathbf{e}_i) = u(\mathbf{e}_i - \mathbf{e}_{threshold}) \quad (14)$$

where $u(\cdot)$ is the unit step function. The update rule for the internal model then becomes:

$$\hat{\mathbf{B}}_{adapted\ i} = \hat{\mathbf{B}}_{adapted\ i-1} + U(\mathbf{e}_i) \left[\left(\frac{-b_1}{b_0} - 1 \right) \mathbf{B}_{adapted\ i-1} - \frac{a_1}{b_0} \mathbf{e}_{i-1} \right] \quad (15)$$

Note that if $\epsilon_{threshold} = 0$, then the update rule of equation 15 simply reverts to that of equation 11. The selection of $\epsilon_{threshold}$ specifies how “attentive” the adaptation process is to small movement errors. Setting $\epsilon_{threshold}$ large implies that the internal representation of the perturbation will adapt only on “exceptional” trials where large movement errors indicates that the model’s prediction of the most recent perturbation was grossly inaccurate.

The attention model of equation 14 with $\epsilon_{threshold} = 5$ mm is shown in Figure 9A. The careless learning algorithm (equation 15) was fit to subject U2’s initial movement error data (first 100 movements). The model’s performance was then evaluated in the final 100 movements (Figure 9B). Even though the algorithm neglected the smallest movement errors, overall performance was respectable when driven by the unimodal perturbation sequence (VAF = 64%). However, when the same model was driven by the bi-modal sequence, the algorithm’s performance suffered (Figure 9C). Movement errors were negatively biased, indicating an inability of the model to compensate for the sequence of perturbations as well as the subject did. With $\epsilon_{threshold} = 5$ mm, the careless learning algorithm compensated only for the approximate mean of the minor peak in the bimodal distribution (4.9 Ns/m, Figure 9D). Consequently, a learning algorithm that performs well in the unimodal sequence may perform poorly in the bimodal sequence unless movement errors are attended to carefully. The residual curvature observed while reaching in both stochastic perturbation sequences was likely due to biomechanical constraints and/or information processing within the motor control systems and not due to inattention to very small movement errors.

Relation to studies requiring adaptation to consistent perturbation sequences

A recent study of reaching movements by Thoroughman and Shadmehr (2000) examined movement errors generated by subjects exposed to predictable perturbing environments with periodic “catch trials” where the predictable perturbation was unexpectedly removed. Movement errors generated in the constant-gain curl-field just after exposure to a catch trial were substantially larger than errors generated just prior to the catch trial. This increase in error was attributed to an “unlearning” of the internal model of the environment. This increase in error decayed on subsequent movements to the same target and was undetectable by about the third trial following a catch trial. This decay rate was comparable to the rate of adaptation observed in experiment 1 from the present study, even though the perturbation sequence used in experiment 1 was random. The decay rates obtained experimentally were comparable to the rate predicted by Equation 5 (Figure 7A). Thoroughman and Shadmehr fit a system of equations to their movement error data that captured this experimentally-observed unlearning behavior. Following a rearrangement of terms and substitution of indices, it can be shown that their system of equations can be represented in the form of equation 5. The similarity in experimental observations and the successes in equivalent modeling techniques between the present study and that of Thoroughman and Shadmehr (2000) suggest that the processes involved in adapting to consistent perturbing environments are the same as those involved in adapting to stochastic perturbing environments.

In conclusion, a sequence of perpendicular viscous force fields with stochastically varying gains triggered an adaptive process that compensated for the *approximate* mean field gain from that sequence. Furthermore, the force field gain which subjects adapted to was not the most frequently experienced gain, nor was adaptation dependent on the particular distribution of perturbations. Although adaptation to the mean field gain would be optimal in the sense that squared movement error would be minimized in the steady state, this strategy is computationally costly and does not allow sufficient flexibility to accommodate efficient learning of nonstationary environments. A simple model of motor performance that depended only on movement error and perturbation gain from the previous trial (equation 5) achieved substantial reduction in movement error, while allowing a rapid and appropriate response to long-term changes in the distribution of perturbations (Figure 7). This simple model predicted subject performance with approximately 84% variance accounted for (VAF). These findings support the hypothesis that the neural structures modified as a result of motor adaptation do not explicitly retain memories of performances or perturbations beyond one or two trials in the past.

Acknowledgements

Special thanks are extended to Dr. Chris Raasch for creating Figure 1-A. This work was supported by grants: NINDS NS35673 and NIMH P50MH48185.

References

- Atkeson, C. G., Moore, A. W. and Schaal, S. Locally weighted learning. *Art. Intel. Rev.*, 11: 11-73, 1997.
- Bhushan, N. and Shadmehr, R. Evidence for learning of a forward dynamic model in human adaptive control. *Adv. Neural Inform. Proc. Systems*, 11: 3-9, 1999.
- Bock, O. Load compensation in human goal-directed arm movements. *Behav. Brain Res.*, 41: 167-177, 1990.
- Bock, O. Early stages of load compensation in human aimed arm movements. *Behav. Brain Res.*, 55: 61-68, 1993.
- Box, G. E. P., Jenkins, G. M. and Reinsel, G. C. *Time Series Analysis; Forecasting and Control*. New Jersey: Prentice Hall 1994.
- Conditt, M. A., Gandolfo, F. and Mussa-Ivaldi, F. A. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J. Neurophysiol.*, 78: 554-560, 1997a.
- Conditt, M. A., Scheidt, R. A. and Mussa-Ivaldi, F. A. Visual influence on learning arm dynamics. (Abstract), *Soc. Neurosci. Abstr.* 23, 1997b.
- Dizio, P. and Lackner, J. R. Motor adaptation to Coriolis force perturbations of reaching movements: endpoint but not trajectory adaptation transfers to the nonexposed arm. *J. Neurophysiol.*, 74: 1787-1792, 1995.
- Dizio, P. and Lackner, J. R. Congenitally blind individuals rapidly adapt to Coriolis force perturbations of their reaching movements. *J. Neurophysiol.*, 84: 2175-2180, 2000.
- Flanagan, J. R. and Rao, A. K. Trajectory adaptation to a nonlinear visuomotor transformation: Evidence of motion planning in visually perceived space. *J. Neurophysiol.*, 74: 2174-2178, 1995.

- Flash, T. The control of hand equilibrium trajectories in multi-joint arm movements. *Biol. Cyber.*, 57: 257-274, 1987.
- Goodbody, S. J. and Wolpert, D. M. Temporal and amplitude generalization in motor learning. *J. Neurophysiol.*, 79: 1825-1838, 1998.
- Happee, R. Goal-directed arm movements: III: Feedback and adaptation in response to inertial perturbations. *J. Electromyography and Kinesiology*, 3: 112-122, 1993.
- Held, R. and Freedman, S. J. Plasticity in human sensorimotor control. *Science*, 142: 455-462, 1963.
- Helmholtz, H. V. v. Treatise on physiological optics, Vol. 3. *Optical. Soc. Am.* 1925.
- Houk, J. C. and Rymer, W. Z. Neural control of muscle length and tension. In: *Handbook of Physiology. The Nervous System. Motor Control*. Baltimore, MD: Am Physiol. Soc., 1981, sect. 1, vol. II, p. 257-323.
- Krakauer, J. W., Ghilardi, M.-F. and Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neurosci.*, 2: 1026-1031, 1999.
- Lackner, J. R. and Dizio, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.*, 72: 299-313, 1994.
- Ljung, L. *System Identification. Theory for the user*. New Jersey: Prentice Hall, 1999.
- Marmarelis, P. Z. and Marmarelis, V. Z. *Analysis of Physiological Systems. The White-Noise Approach*. New York: Plenum Press 1978.

- Oppenheim, A. V. and Schaffer, R. W. *Discrete-time signal processing*. New Jersey: Prentice Hall 1989.
- Press, W. H., Flannery, B. P., Teukolsky, S. A. and Vetterling, W. T. *Numerical recipes in C; the art of scientific computing*. New York: Cambridge University Press 1988.
- Scheidt, R. A. and Mussa-Ivaldi, F. A. Time series analysis of motor adaptation. (Abstract), *Soc. Neurosci. Abstr.* 25, 1999.
- Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z. and Mussa-Ivaldi, F. A. Persistence of motor adaptation during constrained, multi-joint, arm movements. *J. Neurophysiol.*, 84: 853-862, 2000.
- Scheidt, R. A. and Rymer, W. Z. Control strategies for the Transition from Multi-Joint to Single-Joint Arm Movements Studied Using a Simple Mechanical Constraint. *J. Neurophysiol.*, 83: 1-12, 2000.
- Shadmehr, R. and Holcomb, H.H. Neural correlates of motor memory consolidation. *Science*, 277: 821-825, 1997.
- Shadmehr, R. and Brashers-Krug, T. Functional stages in the formation of human long-term motor memory. *J. Neurosci.*, 17: 409-419, 1997.
- Shadmehr, R. and Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.*, 14: 3208-3224, 1994.
- Thoroughman, K. A. and Shadmehr, R. EMG correlates of learning internal models of reaching movements. *J. Neurosci.*, 19: 8573-8588, 1999.
- Thoroughman, K. A. and Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature*, 2000.

Wolpert, D. M., Ghahramani, Z. and Jordan, M. I. Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Exp. Brain Res.*, 103: 460-470, 1995.

Tables

Table 1: Equation 5 regression coefficients and goodness-of-fit measures for both the unimodal and bimodal experiments.

subject	a_1	b_0	b_1	VAF	r^2
U1	0.29	-2.43	1.94	0.62	0.79
U2	0.43	-3.99	3.73	0.73	0.85
U3	0.45	-5.37	4.66	0.77	0.88
U4	0.09	-3.84	2.85	0.70	0.84
B1	0.53	-4.29	2.23	0.87	0.93
B2	0.52	-2.55	2.62	0.80	0.89
B3	0.39	-4.32	3.47	0.86	0.93
B4	0.49	-4.40	3.35	0.86	0.93
B5	0.29	-3.27	1.94	0.81	0.90
B6	0.52	-3.25	2.68	0.77	0.88
B7	0.53	-3.92	3.27	0.86	0.93
B8	0.62	-5.59	4.57	0.92	0.96
U: mean (n=4)	0.31	-3.91	3.29	0.71	0.84
SE mean	0.11	0.60	0.58	0.03	0.02
B: mean (n=8)	0.49	-3.95	3.02	0.84	0.92
SE mean	0.04	0.33	0.29	0.02	0.01

Figure Legends

Figure 1: *A:* Schematic representation of the 2 degree-of-freedom robotic manipulandum used in the present experiments. *B:* Graphical representation of the perpendicular field presented to the subjects. Perturbing forces were directed perpendicular to the direction of intended motion with amplitudes proportional to hand velocity along the intended movement direction. Force field gains (but not direction) varied randomly from trial to trial.

Figure 2: *A:* The trial-to-trial sequence of force field gains (B_i) used in experiment 1 (4 subjects). *B:* Unimodal Gaussian probability density function used to generate the random sequence in panel A. *C:* Autocorrelation of the Gaussian sequence shown in panel A. The two dashed lines correspond to the 95% confidence interval bounds (i.e. the 2σ limits) on the correlation magnitudes. Note that no correlation term within the first 40 trials (except the unit autocorrelation at 0 lag) exceeded the 95% bounds to attain significance at the $p < 0.05$ level. *D:* The trial-to-trial sequence of force field gains (B_i) used in experiment 2 (8 subjects). *E:* Bimodal probability density function used to generate the random sequence in panel D. *F:* Autocorrelation of the bimodal sequence shown in panel D. Dashed lines correspond to the 95% confidence interval bounds on the correlation magnitudes. While each unimodal subpopulation contained no significant correlations within < 40 trials, the shuffling process used to combine the two individual populations gave rise to spurious correlations at trial lags < 40 trials.

Figure 3: Results from experiment 1 (Unimodal perturbation sequence). *A:* Overhead view of averaged hand paths from one subject across the entire experiment. Trials were averaged into six “bins” of 5 Ns/m width (0-5, 5-10, 10-15, etc.). Trials with field strengths ≥ 25 Ns/m were under-compensated (left-most profile) while trials with field strengths ≤ 5 Ns/m were over-compensated (right-most profile). Movements were truncated at the point of time near the end of movement where the hand speed profile reached a transient minimum. Average trajectories after the time of truncation are shown with triangular symbols. *B:* Average hand speed profiles for the same subject. The vertical dashed line indicates the approximate time at which hand speed reached a transient minimum, separating the hand speed profile into two peaks. *C:* Movement errors for this subject plotted against trial number. The dark, solid line represents the exponential best-fit estimate of movement error (equation 2). *D:* Perpendicular hand force profiles obtained by averaging the data in the same manner as in panels A and B. *E:* Scatter plot of movement error versus

perturbation strength, exhibiting a nearly linear relationship ($r = 0.82$). *F*: Best-fit linear regressions from the scatter plots of all subjects from experiment 1 ($0.73 < r < 0.84$).

Figure 4: Results from experiment 2 (Bimodal perturbation sequence). *A*: Overhead view of averaged hand paths from one subject. Trials were averaged into six “bins” of 5 Ns/m width (0-5, 5-10, 10-15, etc.). As in experiment 1, trials with field strengths ≥ 25 Ns/m were under-compensated (left-most profile) while trials with field strengths ≤ 5 Ns/m were over-compensated (right-most profile). Movements were truncated at the point of time near the end of movement where the hand speed profile reached a transient minimum. Average trajectories after the time of truncation are shown with triangular symbols. *B*: Average hand speed profiles for the same subject. The vertical dashed line indicates the approximate time at which hand speed reached a transient minimum, separating the hand speed profile into two peaks. *C*: Movement errors for this subject plotted against trial number. The dark, solid line represents the exponential best-fit estimate of movement error (equation 2). Note that movement error was not well fit by a falling exponential function in the bimodal experiment. *D*: Perpendicular hand force profiles obtained by averaging the data in the same manner as in panels A and B. *E*: Scatter plot of movement error versus perturbation strength, exhibiting a nearly linear relationship ($r = 0.85$). *F*: Best-fit linear regressions from the scatter plots of all subjects from experiment 2 ($0.83 < r < 0.94$).

Figure 5: Correlation analysis of motor performance during adaptation for a typical subject from experiment 1. *A*: Cross-correlation magnitude between movement error and perturbation gain. Dashed lines correspond to 95% confidence interval bounds (i.e. the 2σ limits) on the correlation magnitudes. Movement error on a given trial correlated with perturbation gain on that same trial and with perturbation gain on the previous trial, as subjects compensated for the most recently experienced perturbation. *B*: Autocorrelogram of movement error. Movement error on a given trial correlated with movement error on the previous trial. *C*: Cross-correlation magnitude between peak perpendicular hand force and the perturbation gain, exhibiting a lag-1 correlation with perturbation amplitude. *D*: Cross-correlation magnitude between peak perpendicular hand force and movement error, which also exhibited a significant correlation with previous movement error (lag = +1), consistent with panel A.

Figure 6: Within-subject comparison of predicted (ARX model; Equation 5) movement errors with actual

movement errors during the final 100 movements of the experiment. Model parameters were evaluated from the first block of 100 trials. **A:** Prediction of performance for a typical subject in the unimodal perturbation sequence. The dark line represents actual subject performance, while the thin line represents the model prediction. The model accounted for 79% of the data variance. **B:** Prediction of subject performance in the bimodal perturbation sequence (same line types as in A). The model accounted for 89% of the data variance.

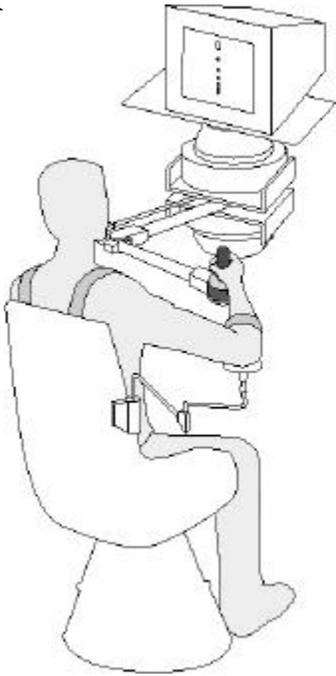
Figure 7: Step response analysis of a family of autoregressive models based on equation 5. The coefficients for each model are described above the corresponding response curves. **A:** (top) Input sequence corresponding to a step increase in the strength of force field perturbation, similar to the perturbations traditionally used to study motor adaptation. Note that a “catch trial” (impulse) was included near the end of the sequence to allow a direct comparison of the model’s response to catch trial behavior described in the text. (bottom) Response of a model derived from the average coefficients of the four subjects participating in the unimodal experiment. Note that the model exhibits a steady-state error as trial numbers become large as did subjects in both experiments. **B:** Step response of a model with no memory; the response on any given trial is only dependent on the current perturbation magnitude (1). Response of a model that precisely compensates for the most recent perturbation without either attenuation or amplification; $b_1 = b_0$, 2). Response of a model with no autoregressive term; $a_1 = 0$, (3). Response of a model with its autoregressive term doubled (4). Model responses when the sign of the autoregressive term has been inverted (5). Model responses when the autoregressive term has been both inverted and doubled (6). Responses of an unstable model $a_1 > 1.0$ (7).

Figure 8: Simulation results of a learning algorithm that accumulates an explicit representation of the mean perturbing field (Equation 12). **A:** Model performance in the unimodal perturbation sequence. The thick line represents subject performance while the thin line represents the model prediction. The model accounted for 70% of the data variance. **B:** Scatter plot of algorithm predictions of movement errors (triangles) and subject performance (filled dots) vs. perturbation amplitude in the unimodal sequence of perturbations. Thick lines represent the linear regressions fitting the dependence of algorithm-predicted movement errors and subject-generated movement errors on perturbation amplitude (zero-crossings of 14.4 and 14.6 Ns/m, respectively) **C:** Cross-correlation magnitude between simulated movement error and perturbation gain for the learning algorithm that accumulates an explicit

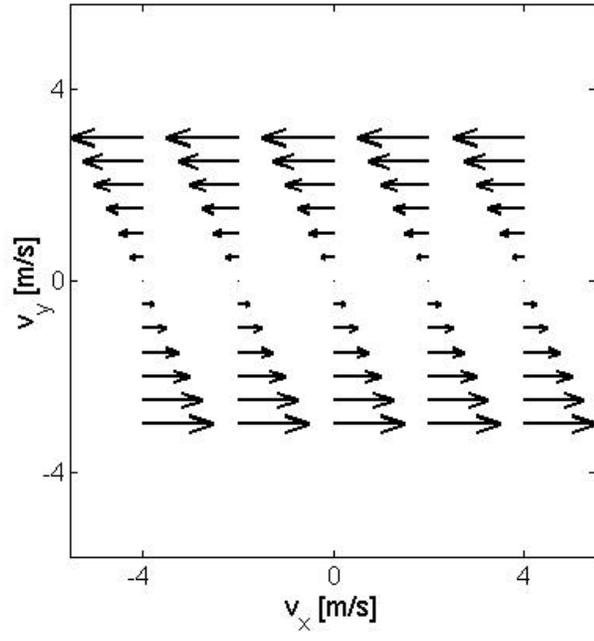
representation of the mean perturbing field. The two dashed lines correspond to the 95% confidence interval bounds (i.e. the 2σ limits) on the correlation magnitudes. Compensation for the explicit average perturbation did not predict the significant lag-1 correlation between movement errors and perturbation gain seen in the subjects' data (e.g. Figure 5A).

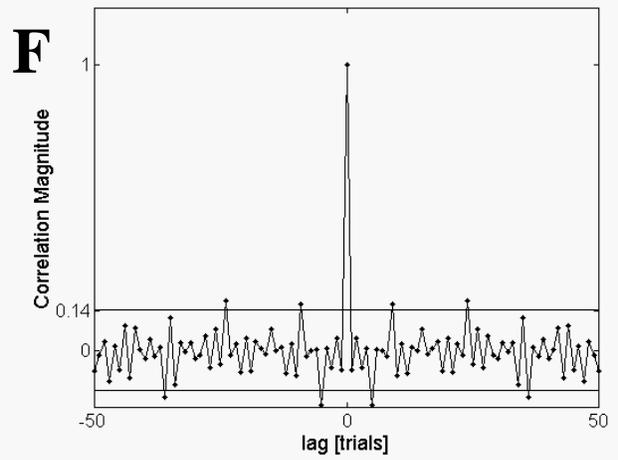
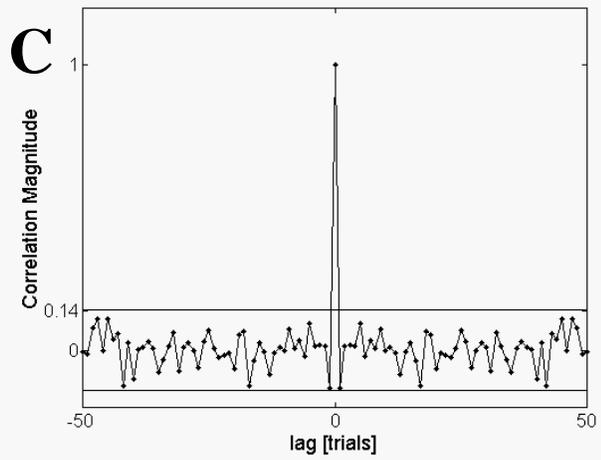
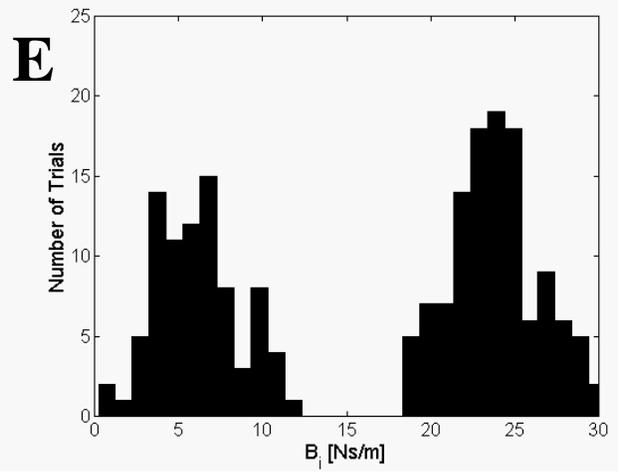
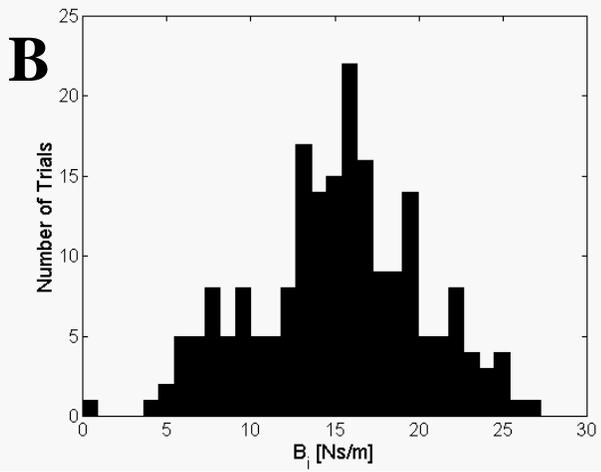
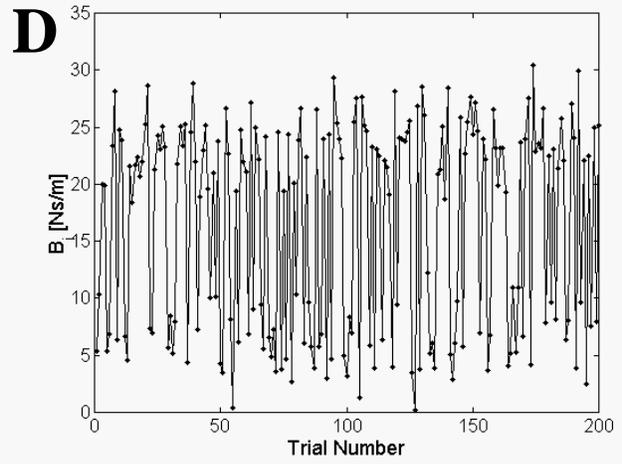
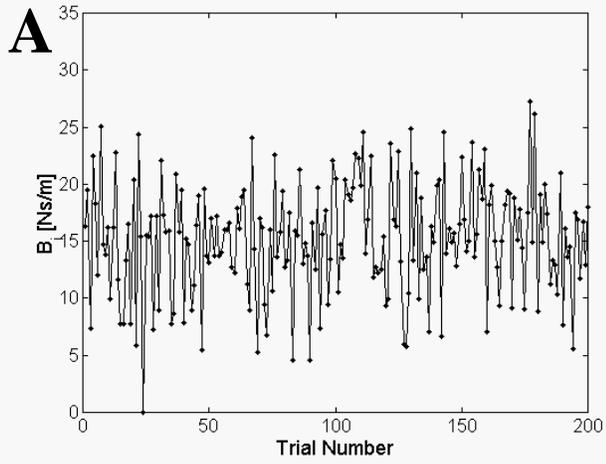
Figure 9: Simulation results of a learning algorithm that updates its internal model when the field strength experienced on a given trial deviates from the predicted field by more than some minimum value (equation 15). **A:** The attention model $U(\mathbf{e}_i)$ for a learning algorithm that attends only to movement errors that exceeded $\epsilon_{\text{threshold}} = 5$ mm. **B:** Comparison of algorithm and subject performance in the unimodal sequence of perturbations. The thick, line represents actual subject performance while the thin line represents the model's prediction of subject performance. The model accounted for 64% of the data variance. **C:** Comparison of algorithm and subject performance in the bimodal perturbation sequence. Line types are as in panel B. **D:** Scatter plot of algorithm predictions of movement errors (triangles) and subject performance (filled dots) vs. perturbation amplitude in the bimodal sequence. The thick lines represent the linear regressions of algorithm-predicted movement errors on perturbation amplitude (left-most diagonal line) and the of subject-generated movement errors on perturbation amplitude (right-most diagonal line).

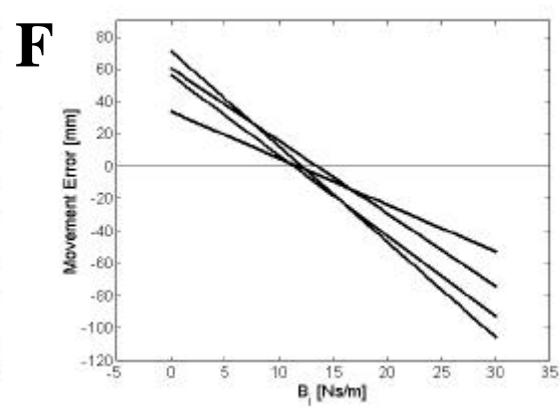
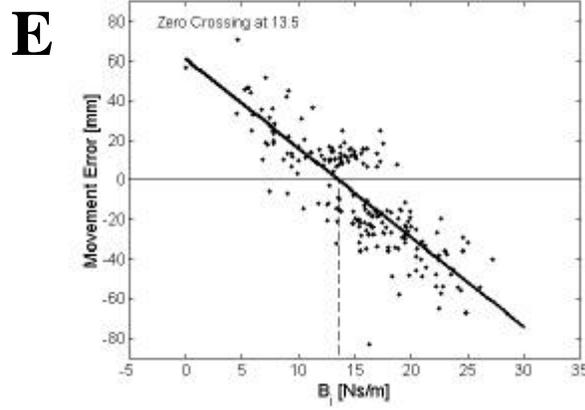
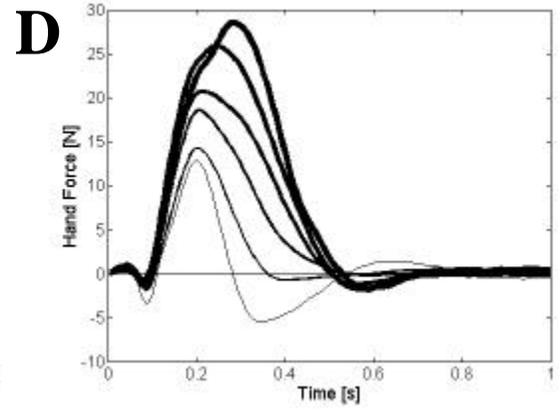
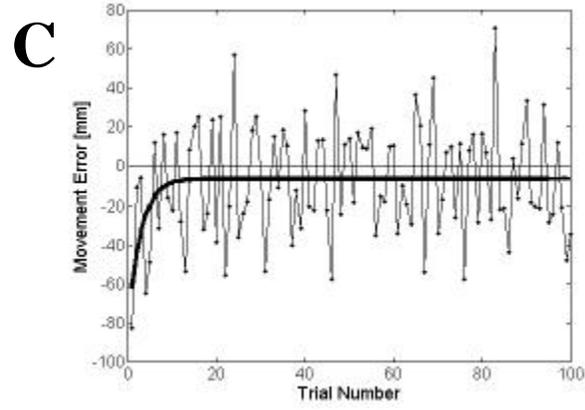
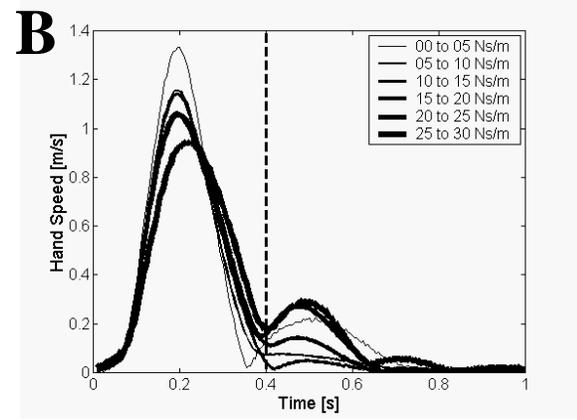
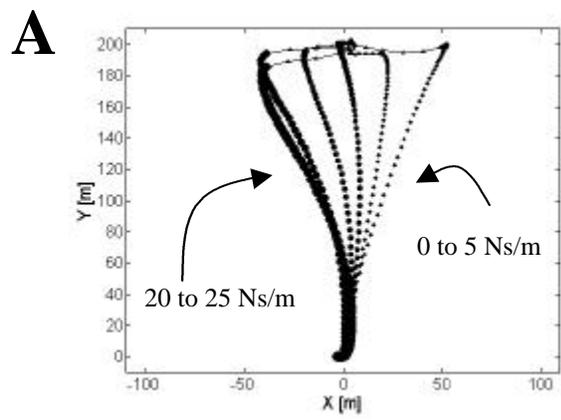
A

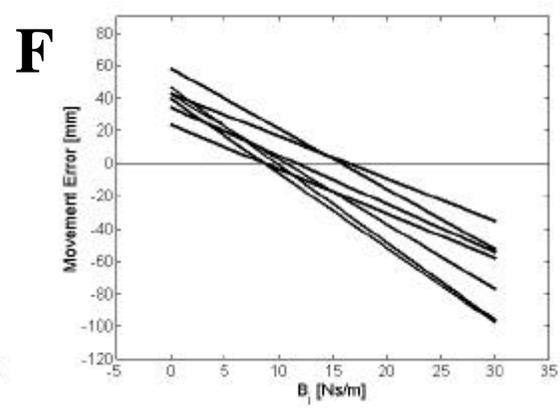
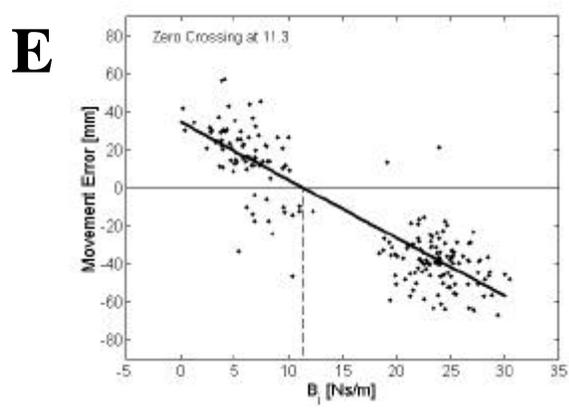
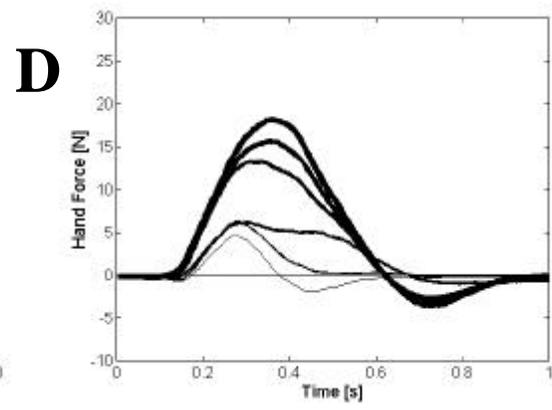
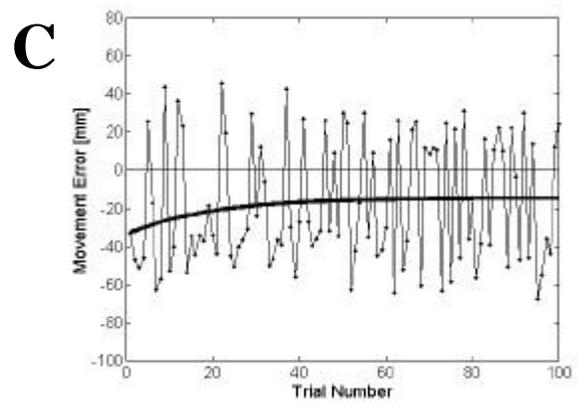
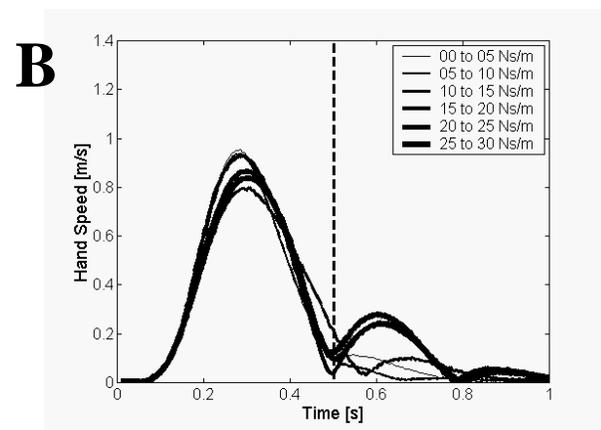
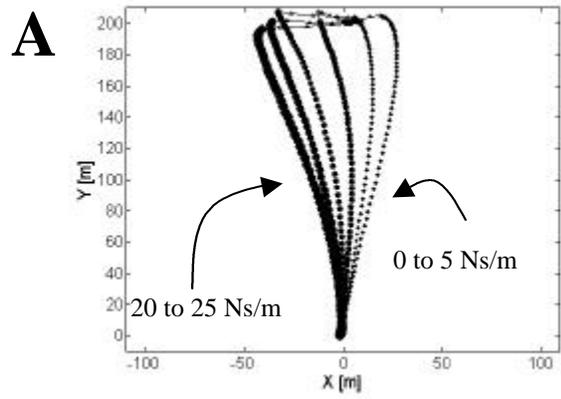


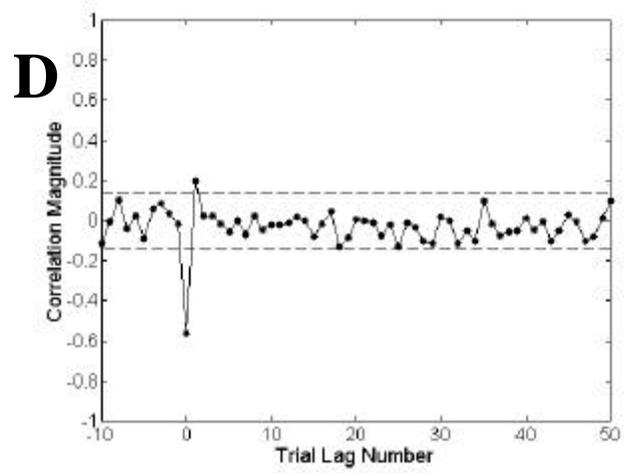
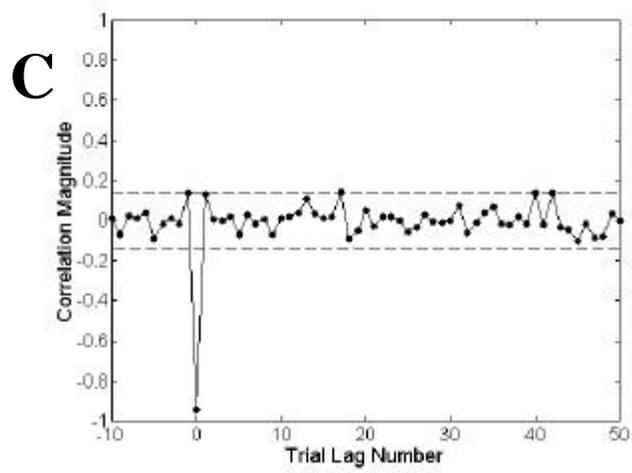
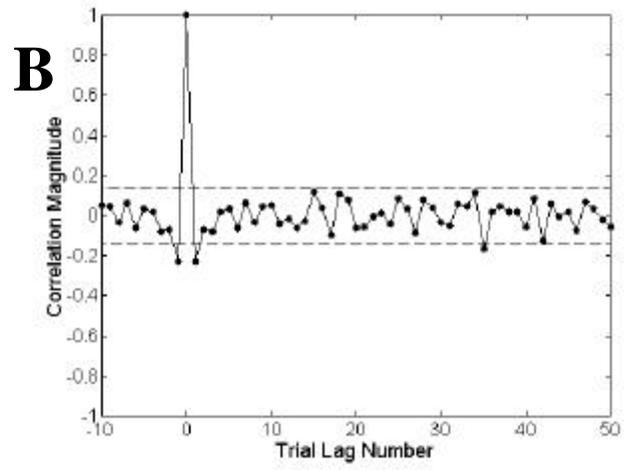
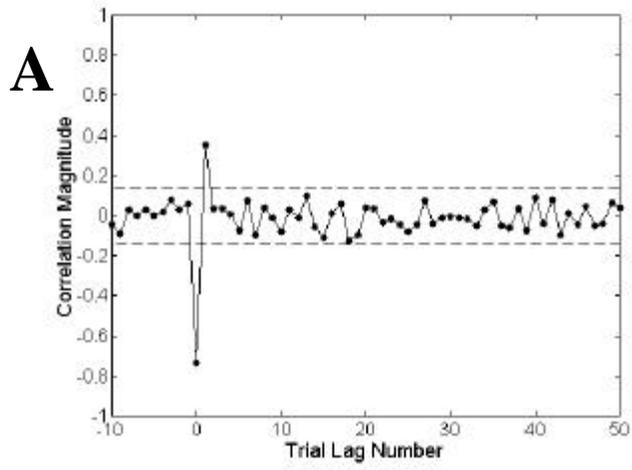
B



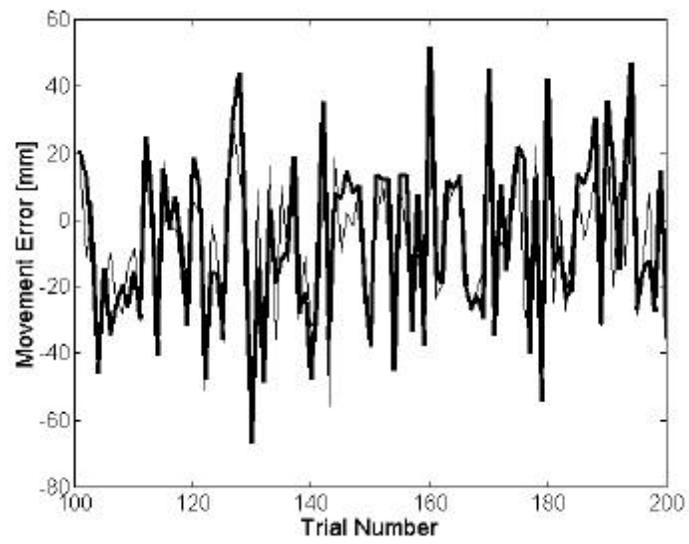




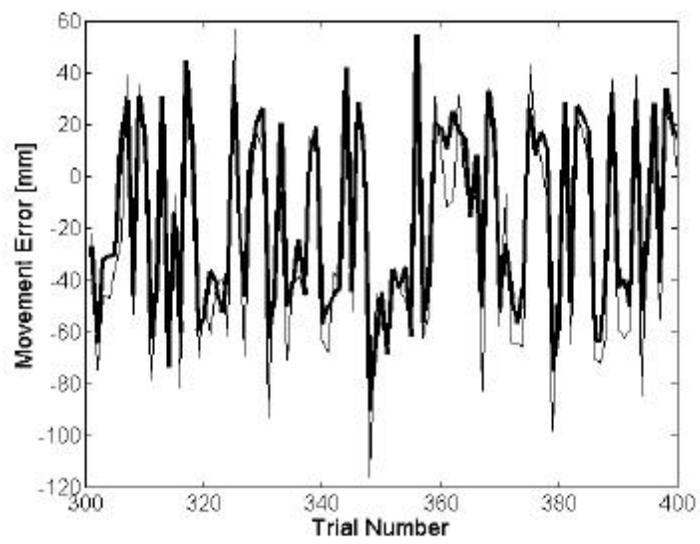


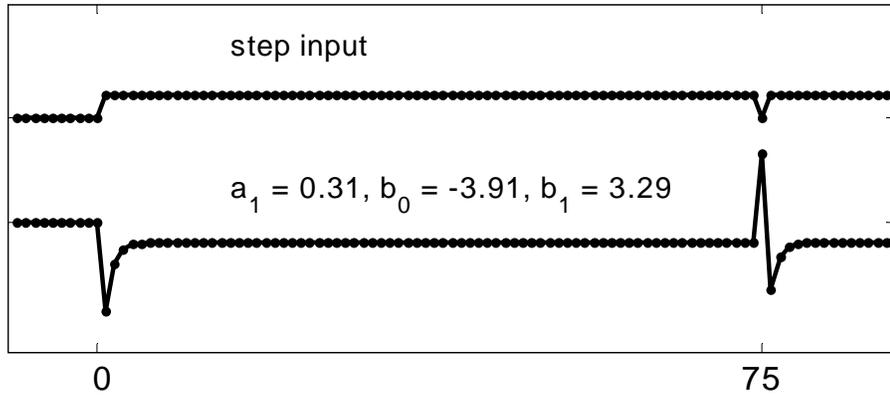
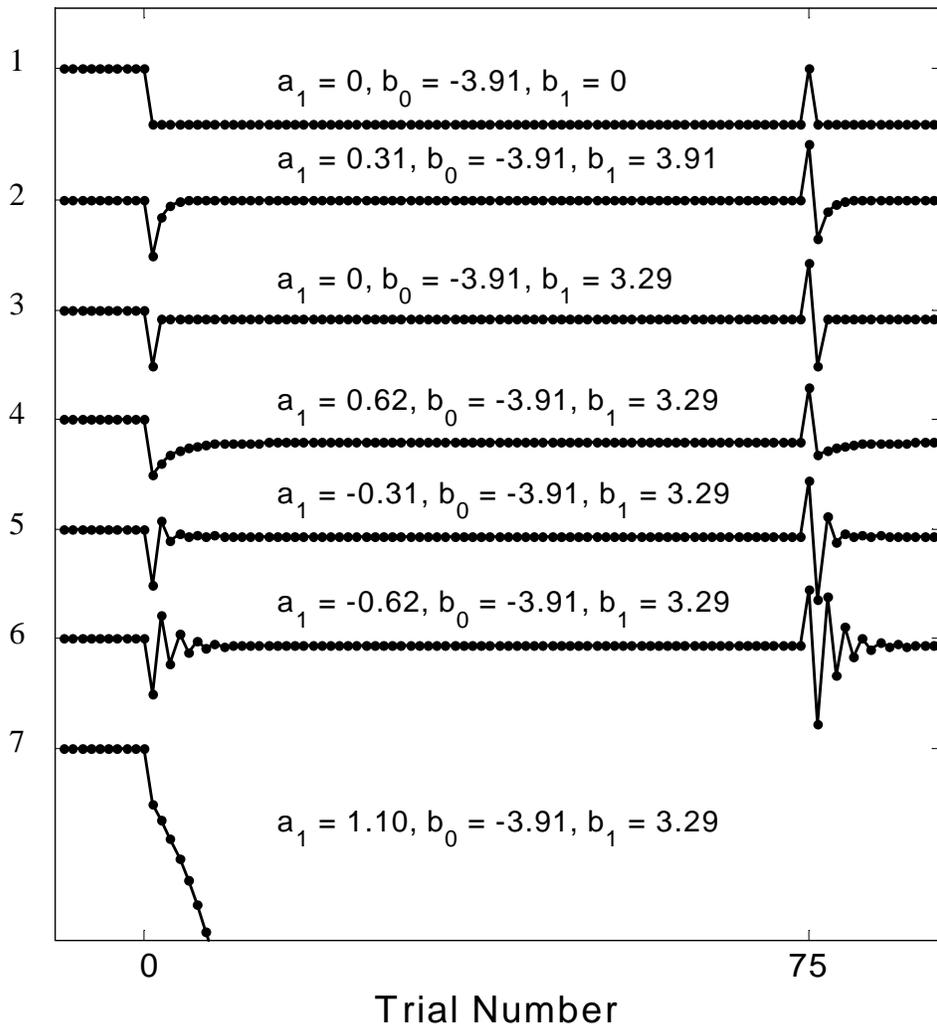


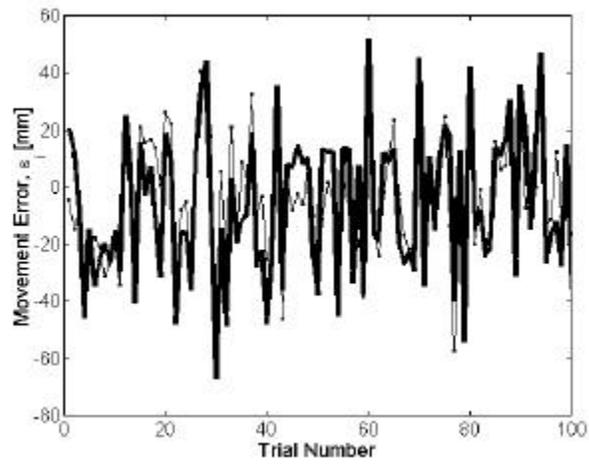
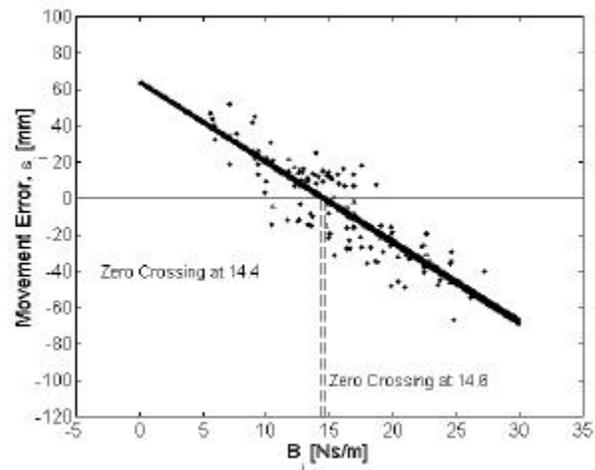
A



B



A**B**

A**B****C**