

Motor Adaptation as a Greedy Optimization of Error and Effort

Jeremy L. Emken,² Raul Benitez,^{2,3} Athanasios Sideris,¹ James E. Bobrow,¹ and David J. Reinkensmeyer^{1,2}

¹Mechanical and Aerospace Engineering Department and ²Biomedical Engineering Department, University of California, Irvine, California; and ³Automatic Control Department, Universitat Politècnica de Catalunya, Barcelona, Spain

Submitted 12 October 2006; accepted in final form 24 March 2007

Emken JL, Benitez R, Sideris A, Bobrow JE, Reinkensmeyer DJ. Motor adaptation as a greedy optimization of error and effort. *J Neurophysiol* 97: 3997–4006, 2007. First published March 28, 2007; doi:10.1152/jn.01095.2006. Motor adaptation to a novel dynamic environment is primarily thought of as a process in which the nervous system learns to anticipate the environmental forces to eliminate kinematic error. Here we show that motor adaptation can more generally be modeled as a process in which the motor system greedily minimizes a cost function that is the weighted sum of kinematic error and effort. The learning dynamics predicted by this minimization process are a linear, auto-regressive equation with only one state, which has been identified previously as providing a good fit to data from force-field-type experiments. Thus we provide a new theoretical result that shows how these previously identified learning dynamics can be viewed as arising from an optimization of error and effort. We also show that the coefficients of the learning dynamics must fall within a specific range for the optimization model to be valid and verify with experimental data from walking in a force field that they indeed fall in this range. Finally, we attempted to falsify the model by performing experiments in two conditions (repeated exposure to a force field, exposure to force fields of different strengths) for which the single-state, auto-regressive equation might be expected to not fit the data well. We found however that the equation adequately captured the pattern of errors and thus conclude that motor adaptation to a force field can be approximated as an optimization of effort and error for a range of experimental conditions.

INTRODUCTION

A problem solved deftly by the human motor system each day is adapting established movement patterns to altered dynamic environments. Examples of this process include learning to swing a new tennis racquet, reaching after picking up a baby, and walking in a coordinated fashion after exchanging boots for slippers. Anticipatory control plays a key role in this process. For example, when people are unexpectedly exposed to a novel force field created by a robot or other mechanical device, their reaching paths initially become curved but gradually straighten with practice (Ghez and Sainburg 1995; Lackner and Dizio 1998; Shadmehr and Mussa-Ivaldi 1994). When the force field is unexpectedly removed, the reaching paths become distorted in a way that indicates the motor system was anticipating the force field. This anticipation has been attributed to the construction of an “internal model” that is expressed as a mapping between limb state, rather than time (Conditt and Mussa-Ivaldi 1999), and muscle activation, rather than force (Takahashi et al. 2006). The motor system is capable of generalizing the internal model to different parts of the workspace (Shadmehr and Moussavi 2000), different movement

speeds (Goodbody and Wolpert 1998), and different tasks (Conditt et al. 1997).

Although anticipatory cancellation of environmental dynamics is a key aspect of motor adaptation, it remains unclear how this process relates to the minimization of energy, a process that accounts for a wide range of experimental findings in motor control. The kinematic paths taken during reaching (Nakano et al. 1999), walking (Anderson and Pandey 2001), and even waddling by penguins (Griffin and Kram 2000) appear to minimize the energetic cost of movement. Indeed if the motor system was concerned about energy expenditure, then one would also expect to see evidence of this concern during adaptation to novel force fields. However, evidence along these lines is limited. In one experiment, the motor system reduced the force it produced when a stiff virtual channel replaced a perpendicular, viscous force field following adaptation during reaching (Scheidt et al. 2000). In another experiment, the steady-state kinematic error during reaching increased with the magnitude of the applied force field (Lai et al. 2003). These two experiments suggest that the elimination of kinematic error by predictive control processes is not the only goal during motor adaptation. However, they leave the details of how the motor system coordinates internal model formation and energy optimization unclear.

Here we show that the motor system coordinates these two processes by minimizing a cost function that includes muscle activation and kinematic error terms on a movement-to-movement basis. This minimization occurs in what can be termed a “greedy” fashion because it considers only the immediate cost of the next movement rather than the overall cost of multiple future movements. We present data that support this optimization model for the task of adapting to a viscous force field during walking. Portions of this work were published in abstract form (Reinkensmeyer et al. 2004).

Theoretical result: optimization model of motor adaptation to a dynamic environment

Recent studies have shown that the evolution of kinematic errors during reaching or walking in a viscous force field is well captured by an autoregressive equation with external input (ARX) that relates the current and previous force to kinematic performance (Donchin et al. 2003; Emken and Reinkensmeyer 2005; Scheidt et al. 2001; Thoroughman and Shadmehr 2000)

$$e_{i+1} = a_1 e_i + b_1 F_i + b_0 F_{i+1} \quad (1)$$

where e_i is a scalar measure of the trajectory error on the i th movement, F_i is a scalar measure of the external force field on

Address for reprint requests and other correspondence: D. J. Reinkensmeyer, 3225 Engineering Gateway, University of California, Irvine, Mechanical and Aerospace Engineering, Irvine, CA 92697-3975 (E-mail: dreinken@uci.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

the i th movement, and a_1, b_0, b_1 are constant parameters. These dynamics are consistent with the formation of an internal model to an externally applied force field by an error-based learning controller. For instance, they predict the presence of a kinematic after effect when the force field is unexpectedly removed as would be expected if the motor system were anticipating the field.

In this section, we address two related questions concerning these autoregressive dynamics. The first question is: “what are the overall control goals that cause the motor system to adapt to an external force field using these dynamics?” The second question is: “what is the neural learning law that implements these control goals?” In the model proposed here, the on-line learning dynamics are driven by an error-based learning rule that implements the overall goal of minimizing both error and muscle activation.

The learning rule can be viewed as being part of a sensorimotor loop with the following structure. First, given a desired trajectory, a controller with a feedback learning rule determines the next muscle force by using the performance error and muscle force in the last step (Fig. 1). The controller contains an inverse model of the limb dynamics that accounts for muscle activation under null-field conditions and a learning rule that drives adaptation when the limb or external dynamics change. The learning rule can be viewed as being composed of a next-force estimator and an optimization controller that minimizes a cost function on a trial-to-trial basis (Fig. 1, inset). The muscle force specified by the learning rule is applied to the limb dynamics, which in turn can be perturbed through the addition of an external force field. Finally, feeding an efferent copy of the muscle command and proprioceptive information of the limb position back into the learning rule closes the loop. We show that this control structure yields to a linear autoregressive equation that is the same as one identified experimentally (Eq. 1) (Scheidt et al. 2001). Further, our model predicts a range of admissible values for the linear coefficients in Eq. 1 not defined previously.

Consider a cost function that takes into account scalar measures of current kinematic error, e_{i+1} and current muscle activation, u_{i+1} , for a given movement trajectory,

$$J = \frac{1}{2} e_{i+1}^2 + \frac{\lambda}{2} u_{i+1}^2 \tag{2}$$

where i denotes the i th step or reach in an environment and $\lambda > 0$ weights the cost of the kinematic error versus effort. For this

derivation, we assume motor command and muscle force are proportional quantities (Cheney and Fetz 1980) and represent them as a single variable u , expressed in Newtons. However, we recognize that the relationship between activation and force changes in some cases, for example, after neuromuscular fatigue (Takahashi et al. 2006) or because of the force-velocity property of muscles. The model could be extended to account for these conditions. Next, assume spring-like dynamics for the limb, with stiffness K , in response to the force-field perturbation F_i and any change in the motor command u_i

$$Ke_i = F_i + u_i \tag{3}$$

This assumption is verified experimentally in the following text (see Fig. 6A). The controller that greedily minimizes this cost function during force-field adaptation is found by determining u_{i+1} such that $dJ/du_{i+1} = 0$ in Eq. 2 and noting from Eq. 3 that e_{i+1} depends on u_{i+1} . This controller is given by

$$u_{i+1} = -\frac{F_{i+1}}{1 + \lambda K^2} \tag{4}$$

From Eq. 4, it is clear that an estimation of the next force F_{i+1} is required to apply the control law. We assume that the external force field is of the form

$$F_i = F_0 + v_i \tag{5}$$

where F_0 is a constant unknown force-field intensity and v_i is a zero-mean, independent and identically distributed (iid) noise with variance σ^2 that accounts for errors in the force perception (measurement error). We expect that the sensorimotor system will attempt to learn F_0 and compensate for it through the motor command u . Common iterative estimation schemes, such as recursive least squares and Kalman filtering (Haykin 1996) accomplish this and result in a recursive estimator in the form of a first-order, low-pass filter with time-varying coefficients (Haykin 1996, p. 285). We take here a time-invariant low-pass filter as our next force estimator

$$\hat{F}_{i+1} = \alpha \hat{F}_i + (1 - \alpha) F_i \tag{6}$$

where α weights the importance of previous estimated force \hat{F}_i and previous perceived force F_i in the recursive calculation of the next-force estimate. The estimate provided by Eq. 6 is asymptotically unbiased and therefore the mean square estima-

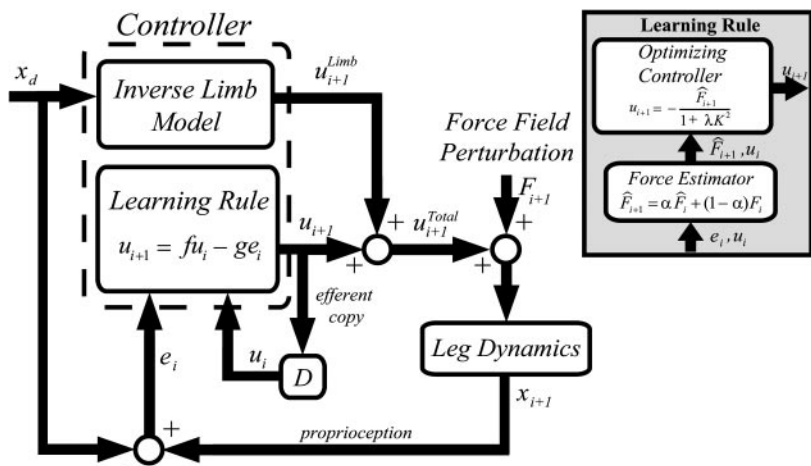


FIG. 1. Optimization framework. The learning rule can be expressed as Eq. 8, which is a combination of Eqs. 4 and 6 (inset).

tion error of F_0 is the same as the asymptotic variance, which can be shown to be

$$\hat{\sigma}^2 = \frac{1 - \alpha}{1 + \alpha} \sigma^2 \quad (7)$$

Thus the estimator reduces the variance and mean square estimation error associated with force perception, when $0 < \alpha$ and < 1 . Further, α controls the tradeoff between the accuracy of the estimate and the speed of the estimation. We note that the validity of this choice of estimator is tied to the fact that it ultimately results in the appropriate learning dynamics (i.e., Eq. 1), rather than the fact that it is related to Kalman or least squares filtering. Exploring the link with Kalman filtering is an interesting direction for future research.

Assuming that the estimate of the force field \hat{F}_{i+1} in Eq. 6 is used for F_{i+1} in the control law (Eq. 4) and using the limb dynamics equation (Eq. 3), we obtain the learning rule

$$u_{i+1} = fu_i - ge_i \quad (8)$$

where $f = (1 + \alpha\lambda K^2)/(1 + \lambda K^2)$ and $g = K(1 - \alpha)/(1 + \lambda K^2)$ are defined as the “forgetting factor” and the “learning gain,” respectively. The controller defined by Eq. 8 can be viewed as determining the motor command required to anticipate the force field, as it increments the motor command in proportion to the previous error, in the direction that reduces error. However, the controller also tries to reduce the motor command when error is small because $0 < f < 1$ (note $0 < \alpha < 1$ in Eq. 6). Thus this controller is an error-based controller with a forgetting factor. This learning process can also be viewed as driving the gradual formation of an inverse internal model of the external force field.

We can then predict the temporal evolution of trajectory errors by applying the learning rule (Eq. 8) to the spring-like dynamics of the limbs (Eq. 3) to obtain

$$e_{i+1} = \left(f - \frac{g}{K}\right)e_i - \frac{f}{K}F_i + \frac{1}{K}F_{i+1} \quad (9)$$

which recovers the autoregressive dynamics (i.e., Eq. 1) that describes error evolution during adaptation to a force field, with $a_1 = (f - g/K) = \alpha$, $b_1 = -f/K$, and $b_0 = 1/K$. The fact that $\alpha = a_1$ provides an interpretation for a_1 : namely, a_1 controls the tradeoff between speed and accuracy of the estimation of the force field.

We also note here that this optimization theory makes the following prediction regarding the ARX model: The admissible range of the autoregressive coefficients is

$$0 < a_1 < -\frac{b_1}{b_0} < 1 \quad (10)$$

The proof of this inequality is established by solving the equation $b_1 = -f/K = -(1 + \alpha\lambda K^2)/[(1 + \lambda K^2)K]$ for λ to obtain $\lambda = -(1 + Kb_1)/[K^2(Kb_1 + a_1)]$. Then because $\lambda > 0$, the only feasible solution for the coefficients is when the numerator $1 + Kb_1 > 0$ and the denominator factor $(Kb_1 + a_1) < 0$. Now, recall that $b_0 = 1/K$ and notice a further constraint: for Eq. 9 to be stable, $-1 < \alpha = a_1 < 1$. An even more restrictive constraint is $0 < \alpha = a_1 < 1$ for the estimator to reduce variance in its estimate of force. Taking this constraint, the desired relationship is obtained. Note that the requirement of Eq. 10 restricts the allowable parameter range

more than the stability requirement for Eq. 9. This additional restriction arises essentially because $\lambda > 0$ for the optimization model to make sense [if $\lambda < 0$, the optimizing controller would try to make muscle force as large as possible to make the associated cost as small (negative) as possible].

Thus in response to the question “what are the overall control goals that cause the motor system to adapt to an external force field using the dynamics of Eq. 1”, the answer is the minimization of kinematic error and effort. This overall control goal is precisely defined as the minimization of the cost function (Eq. 2) in conjunction with the next force estimator (Eq. 6). We note that in our previous research (Reinkensmeyer et al. 2004), we used a more complicated cost function that included a penalty on the change in effort along with a simplified force estimator that assumed that the estimate of the next force would be the same as the previous one. Both approaches predict the same learning dynamics (i.e., Eq. 1). In response to the question “what is the neural learning law that implements these control goals”, the answer is an error-based learning law with a forgetting factor, as given by Eq. 8.

The novel contribution of this theoretical result is that it provides a computational connection between an overall control goal (minimization of error and effort) and the specific closed-loop learning dynamics in Eq. 1. This theoretical result is already verified experimentally, at least in part, because a main experimental prediction of this theory is Eq. 1, and Eq. 1 has been experimentally confirmed by several studies of reaching and walking in a force field (Donchin et al. 2003; Emken and Reinkensmeyer 2005; Scheidt et al. 2001; Thoroughman and Shadmehr 2000). What was missing from these previous studies that identified Eq. 1 and is now provided here is a high-level explanation of what the motor system seeks to accomplish when it adapts with the dynamics in Eq. 1. In addition, the optimization theory also predicts admissible values for the coefficients of Eq. 1, which were not predicted by previous research that identified Eq. 1. The coefficients of Eq. 1 must satisfy Eq. 10 for the optimization model to be valid. Using the experimental methods described in the next section, we verify that the coefficients identified during motor adaptation actually do satisfy Eq. 10.

We emphasize that the validity of the optimization model is essentially tied to the validity of the learning dynamics of Eq. 1—one leads to the other, with the added constraint that the coefficients of Eq. 1 must fall in some range. The optimization model could thus be falsified by showing that Eq. 1 does not fit adaptation data or by showing that Eq. 1 fits the data but its coefficients do not fall in the admissible range. We attempted to falsify the model by performing experiments in two conditions: repeated exposure to a force field and exposure to force fields of different strengths.

We were interested in the first condition because previous studies have found evidence of retention (Brashers-Krug et al. 1996; Gandolfo et al. 1996; Shadmehr and Brashers-Krug 1997), and learning at multiple time scales (Smith et al. 2006). Eq. 1, which has only one state, does not capture these time-dependent features.

We were interested in the condition of exposure to different force-field strengths because the optimization model predicts, for all but the trivial case when $\lambda = 0$, that steady-state kinematic error will increase as the forces applied by the dynamic environment increase. That is, when $\lambda < 0$, there is an

effort term in the cost function, and $f < 1$ because $f = (1 + \alpha\lambda K^2)/(1 + \lambda K^2)$ and α varies between 0 and 1. The presence of a forgetting factor not equal to 1 has the consequence that the steady-state trajectory error predicted by Eq. 9 should increase proportionally to the constant field strength $F_i = F$, with the forgetting factor f , the learning gain g , and the limb stiffness K , determining the proportionality constant

$$e_f = \frac{1-f}{K(1-f)+g} F \quad (11)$$

We tested the dependence of e_f on force strength F predicted by Eq. 11 for the task of walking in a robot-generated, perpendicularly directed, viscous force field. Specifically, we measured the forgetting factor f , stiffness K , and learning gain g , at one force-field strength and then examined whether these values predicted the dependence of error on field strength given by Eq. 11. We note that Eq. 1 by itself, without the conceptual framework of the optimization, predicts that final error depends on force-field strength, provided that $b_1 \neq -b_0$ as is suggested by the data of Scheidt et al. (2001). But the condition that $b_1 \neq -b_0$ also corresponds to the condition that $f \neq 1$, which corresponds to including an effort term in the optimization. Thus establishing the fact that final kinematic error increases predictably with force-field strength, according to Eq. 11, verifies that effort minimization takes place along with error minimization. Further, testing the applicability of the equation, with parameters identified at one force level, to predict performance at many force levels, probes its generality.

METHODS

Subjects

The University of California–Irvine, Institutional Review Board, approved all experiments. Six healthy, subjects (age 24–37 yr, 5 male) with no history of neurologic disorders participated in the experiments. Data from a previous study in which 10 healthy adults subjects walked in a viscous force field (Emken and Reinkensmeyer 2005) were also used to test the model.

Experimental apparatus and robot-generated dynamic environments

Subjects walked on a treadmill with a lightweight, 2 degrees of freedom robot secured to their left shank via a custom composite brace secured just above the ankle (Fig. 2). The brace could rotate in the sagittal plane around a passive revolute joint attached to the robot's apex, and the robot allowed a small amount of passive movement of the leg in the frontal plane. The effects of friction and gravity were cancelled in software, leaving only the effect of the robot's inertia. Details of the robotic device design and performance specifications can be found in Emken et al. 2006. Subjects wore a chest harness attached to an overhead frame as a safety precaution in case they stumbled. The harness was adjusted so that it did not provide weight support or constrain torso movement during walking. Subjects were instructed to walk as consistently as possible without looking at or thinking about the motions of their feet.

To create a novel dynamic environment, the robot was programmed to apply a vertical force proportional to the shank's (i.e., the apex of the robot) forward horizontal velocity (Eq. 12). In most experiments, the gain, B , was chosen positive such that the force field pushed the leg upward during the swing phase of gait.




FIG. 2. Picture (left) and diagram (right) of experimental setup. The robot makes use of a linear motor with 2 forcer coils and a V-shaped linkage to accommodate and drive motion of the apex in the parasagittal plane. The apex is attached through a padded brace and revolute joint to the subject's lower shank. Subjects wore a harness attached to an overhead frame to catch them in case they fell.

$$F_y = \begin{cases} -B\dot{x}, & \dot{x} < 0 \\ 0, & \dot{x} > 0 \end{cases} \quad (12)$$

The field gain was calculated for each subject and chosen to produce a peak upward force of ± 1 –10% of the subjects' body weight, depending on the experiment. Gains in this range produced obvious perturbations without causing stumbling.

Experimental protocol

Subjects first walked during a warm-up period with the robot secured to their ankle with no field applied by the robot, to become comfortable walking within the robot, and to select a comfortable walking speed. The subjects' preferred treadmill speeds ranged between 0.80 and 1.07 m/s.

The subsequent experimental protocols conformed to the following general paradigm. Experiments began with an initial null phase in which the subjects walked in a null field (N) in which the robot applied no forces to the leg. After this initial stage, subjects were presented with a force stage during which the force field (F) was turned and held on for a number of steps. Subjects then again stepped in the absence of the force field in a second null field (N). In some cases, catch trials, in which the field strength was altered to an arbitrary value or set to zero for a single step, were randomly introduced into both the null and force-field stages.

Two specific experiments were performed by altering the number of N-F-N sequences, the length of the individual stages, the strength and direction of the F stages, or the strength, number, and spacing of catch trials (Table 1). *Experiment A* was designed to examine the effect of repeated exposure on direct effect height and speed of adaptation. For this experiment, we reanalyzed previously published data (Emken and Reinkensmeyer 2005) with a new error measure described in the following text. *Experiment B* was designed to test the prediction of field strength on direct effect size and final steady-state error.

Data analysis

The model of motor adaptation proposed above requires that muscle force and kinematic error be summarized with scalar values for each step in the force field. Previous studies have used the peak force, and the maximum perpendicular error (Scheidt et al. 2001), the distance after 300 ms from movement initiation (Shadmehr and Brashers-Krug 1997) and the area circumscribed by the hand path and the straight line between the targets (Takahashi et al. 2001) as scalar measures of force and kinematic error. We systematically studied

TABLE 1. *Experimental protocols*

Experiment	Subjects	Protocol
A: repeated exposure	10	100N/13CT-(25N-25F) ₁₀ —where the subscript 10 indicates that the protocol in parentheses was repeated 10 times. This experiment was reported previously (Emken and Reinkensmeyer 2005) and reanalyzed here with a new error measure.
B: error dependence on field strength	6	(120N-(100F-100N) ₄) ₂ where the subscripts 4 & 2 indicate that the protocol in the inner parentheses was repeated twice, with 8 different Fs corresponding to ±2, 5, 8, and 11% of the subjects' body weight. In each of the two runs, there were two positive and two negative force stages.

#N, number of steps in null field; #F, number of steps in force field; #CT, number of randomly interspersed catch trials.

whether there was an optimal time during the movement to take scalar measures of force and kinematic error by examining how the ability of the model (Eq. 1) to fit the data depended on when the scalar measurement was made. Specifically, we varied the timing of the error and force measurements from 75 to 350 ms after the zero crossing of the x velocity. The error-force pair with the highest mean R^2 was found to be the force applied by the robot 100 ms following the start of forward ankle movement and the resulting step height error 200 ms later. These measures of force and error significantly outperformed the use of peak force and peak step height measures that we used previously (Emken and Reinkensmeyer 2005) by increasing the R^2 for the fit of (Eq. 1) to the data from *experiment A* by 12% ($P < 0.02$, 1-way ANOVA). Residual analysis (Montgomery et al. 2001) indicated that the measures were linear between the 0.1 and 0.9 cumulative probability points and flattened at both ends, suggesting that the residual distribution obtained with these measures did not deviate substantially from normality, but was slightly light-tailed. The 200-ms lag between the optimal times at which force and position are best measured is interesting because it corresponds approximately to the leg's voluntary reaction time. This timing for measuring position error thus allows the force perturbation to create a positional deviation for as long as possible before the potential confound of voluntary reaction can influence limb kinematics.

Given these measures of force and error, multiple linear regression was used to identify the coefficients a_1 , b_0 , and b_1 of Eq. 1. Then, f , g , and K were calculated as follows

$$K = \frac{1}{b_0} \quad f = -\frac{b_1}{b_0} \quad g = K(f - a_1) \quad (13)$$

In turn, the parameters of the cost function were calculated as

$$\lambda = \frac{K(1 - \alpha) - g}{gK^2} \quad \alpha = a_1 \quad (14)$$

Model performance as a function of model complexity was compared by multiple linear regression (MLR) and K -fold cross validation using data from Emken and Reinkensmeyer (2005). For both analyses, 13 different models were assumed to represent the data. For each model, MLR was performed to extract the R^2 correlation coefficient and regression coefficients. In all cases, only models whose regression was significant were chosen for comparison ($P < 0.05$). All assumptions to perform classical linear regression were true with the exception that the residuals were often auto-correlated. This stems from the fact that the model is autoregressive in nature and that there are often drifts in the baseline of both the null and force-field step heights that

are likely a function of changes in walking posture. Due to this possible misinterpretation, we chose to analyze the models with a cross validation approach as well. For each candidate model structure, we computed fitted models based on 9 of the 10 subjects and computed the performance of each of the fitted models on the tenth subject. Actual subject performance as quantified by mean square error (MSE) is defined as the mean of the square differences between the actual and estimated subject performance. By averaging the 10 MSEs obtained for each model structure, we obtain an estimate for the expected MSE error for that structure.

RESULTS

The main theoretical contributions of this paper are 1) a theory that demonstrates that the autoregressive dynamics that have been previously identified for force-field adaptation (Eq. 1) are mathematically identical to a minimization of a cost function containing kinematic error and muscle effort terms and 2) the derivation of an admissible range for the autoregressive coefficients of Eq. 1, for the minimization framework to hold (i.e., Eq. 10). We also performed two experiments designed to test the validity of Eqs. 1 and 10 and thus of the optimization model also.

Do the learning coefficients fall in the admissible range predicted by the optimization model?

We analyzed data from two experiments (A: repeated exposure to a force field, B: exposure to force fields of different strengths) to identify coefficients for Eq. 1 and test whether they fell within the admissible range. Equation 1 fit the data adequately well, capturing $83 \pm 0.052\%$ of the variance of the step height error for *experiment A* in which 10 subjects were exposed to 10 consecutive constant strength force fields, each followed by a period of washout (Table 2). The identified forgetting factor f for this experiment was significantly less than one ($P < 0.0001$, t -test, mean = 0.76 ± 0.21), consistent with the inclusion of an effort term the cost function. Average weighting parameters from the cost function and the model parameters from (Eqs. 1, 13, and 14) are given in Table 2. When we plotted the parameters for both *experiments A* and B,

TABLE 2. *Model parameters and fit*

Parameter	a_1	b_0	b_1	R^2	K	f	g	λ	α
Units	—	cm/N	cm/N	—	N/cm	—	N/cm	cm ² /N ²	—
<i>Experiment A</i> *	0.49 ± 0.094	0.34 ± 0.071	-0.26 ± 0.040	0.83 ± 0.052	3.0 ± 0.62	0.76 ± 0.21	0.80 ± 0.80	0.31 ± 0.46	0.66 ± 0.19
<i>Experiment B</i> †	0.41 ± 0.16	0.20 ± 0.071	-0.14 ± 0.068	0.77 ± 0.094	5.9 ± 3.0	0.64 ± 0.17	1.2 ± 1.0	0.30 ± 0.29	0.53 ± 0.15

*Data reanalyzed from baseline experiment of (Emken and Reinkensmeyer 2005) using a new error measure. †Values are from the first of the two runs for all subjects and are representative of both runs ± one standard deviation.

the autoregressive coefficients ($a_1, -b_1/b_0$) fell in the admissible range defined by Eq. 10 for every subject (Fig. 3).

Is Eq. 1 the simplest linear equation that predicts the learning dynamics?

Because the validity of the optimization model rests on the ability of Eq. 1 to predict errors, we thought it prudent to revisit the question of whether linear difference equations that were simpler or more complex than Eq. 1 better explained the error dynamics for this type of experiment. Scheidt et al. (2001) addressed this question by identifying the error and force variables that are correlated for different trial-to-trial lags. Here we tested a broad range of linear models with error and force terms at different lags to see which fit the data best with the fewest terms. Models with fewer terms than model 7 explained $\geq 10\%$ less of the variance (Fig. 4, model 7). More complex models that looked back in time more than one step did not explain significantly more of the variance. MSE dropped and then leveled off at model 7, indicating that more complex models did fit significantly better. An increase in MSE was not seen with more complex models suggesting that over fitting was not a concern. Thus model 7 (Eq. 1) was the simplest model that captured the data well.

Effect of repeated exposure

The autoregressive equation (Eq. 1) predicts that the motor system will respond in a stereotypical fashion each time it is exposed to a force field. Specifically, because the equation's coefficients are assumed constant, it predicts that the motor system will respond to changes in field strength in a rote fashion blind to all previous steps in the force field except the most recent one. Analysis of the data from experiment A in which 10 subjects were exposed to the force field 10 consecutive times with null fields interspersed between each exposure, demonstrated that subjects indeed responded in a rote fashion to the field (Fig. 5). They did not have smaller initial

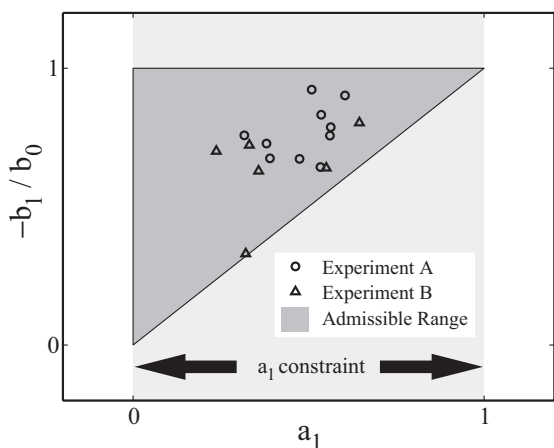


FIG. 3. Admissible range for model parameters defined by Eq. 1. The parameter a_1 in Eq. 1 must fall in the region labeled “ a_1 constraint” for the estimator of Eq. 6 to reduce both the variance and mean square error in its estimation process. This is more restrictive than the basic stability constraint for Eq. 1, which is that $-1 < a_1 < 1$. For the optimization to be valid, the ratio $-b_1/b_0$ must satisfy Eq. 10, which together with the a_1 constraint is shown as the gray shaded triangular region. \circ and Δ , actual parameters identified for individual subjects in experiments A and B, respectively.

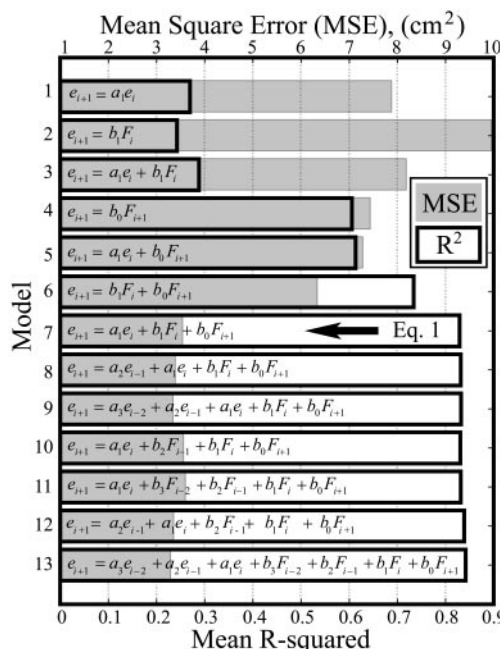


FIG. 4. Model performance. Mean square error and mean R^2 values across subjects when models of different order were used to fit the data from experiment A. Grey bars show the average mean square error (MSE) calculated via cross validation of the 10 subjects. Contributions of individual terms to the goodness of fit can be inferred by comparing the goodness of fit for models that differ by only the term of interest.

step height errors ($P = 0.82$, linear regression) nor did they show a trend to learn more quickly with repeated exposure to the field ($P = 0.60$, linear regression). The equation slightly overestimated the rate of adaptation and underestimated the initial rate of de-adaptation, but these misestimates were not statistically different from the model (Fig. 5).

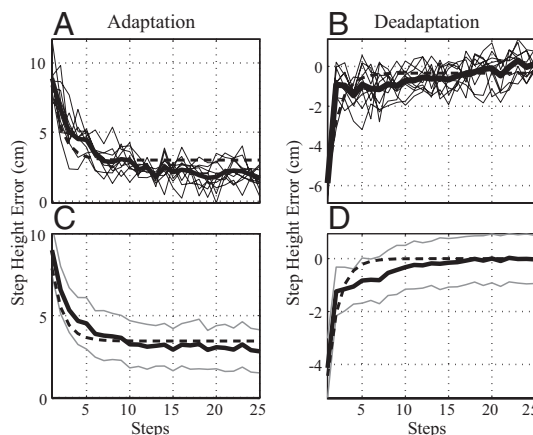


FIG. 5. Adaptation and deadadaptation on repeated exposure to a constant strength force field (experiment A). A and B: adaptation and deadadaptation of a representative subject on 10 consecutive field exposures. Thick line represents the mean of 10 error evolutions and thin lines represent individual error evolutions. C and D: mean adaptation and deadadaptation (thick solid line) across 10 subjects, each with 10 consecutive adaptations with thin gray lines showing 95% confidence interval. The stiffness, learning gain, and forgetting factor were identified by regressing Eq. 1 with data from the baseline experiment for a representative subject (A and B) and 10 subjects as a group and averaged (C and D), and then were used to predict the adaptation and deadadaptation for a single field exposure. This prediction is shown as a thick dashed line in all subfigures.

Dependence of initial and steady-state stepping error on force-field strength

A key consequence of including an effort term in the optimization model is that the steady-state stepping error after adaptation will increase with the strength of the force field (*Eq. 11*) because the model proposed here suggests that the motor system cares not only about kinematic error but also muscle effort. This relationship was experimentally verified using data from *experiment B* in which the force-field strength was systematically varied during eight exposures to the force field (*Fig. 6, C and D*, $R^2 = 0.99$, $P < 0.0001$, linear regression). The forgetting factor identified for one force amplitude allowed detailed prediction of the relationship between force and kinematic error to all other force amplitudes studied. Thus the optimization model provides a precise high-level explanation for the form of the dependence of steady-state kinematic error on field strength, demonstrating its validity across a range of field strengths.

Validity of the linear spring model of the leg/robot dynamics

A key assumption in the derivation of the autoregressive model (*Eq. 9*) from a minimization of the cost function (*Eq. 2*) is that the leg behaves like a linear spring (*Eq. 3*) in response to the force field for the stepping task. *Figure 6A* shows that the initial step height error (i.e., the “direct effect”) experienced when the force field was turned on unexpectedly was a linear function of the force-field strength ($R^2 = 0.97$, $P < 0.0001$, linear regression), verifying this assumption. The linear spring model well predicted the initial error for a range of force-field strengths (*Fig. 6B*).

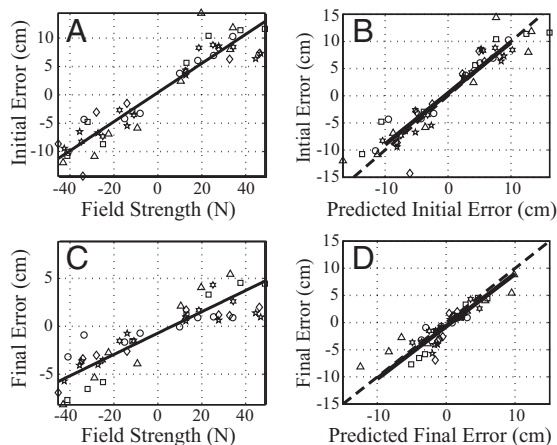


FIG. 6. Initial and steady-state step height errors when adapting to force fields with different strengths (*experiment B*). *A*: initial step height error increased with peak force-field strength. *B*: measured initial step height error vs. the initial step height error predicted by the model. *C*: final step height error, which was quantified as the mean of the last half stepping errors, increased linearly with peak force-field strength. *D*: measured final step height error vs. the final error predicted by the model according to equation (*Eq. 11*). The stiffness, learning gain, and forgetting factor were identified by regressing (*Eq. 1*) with data from a 7% of peak body weight field and then were used to predict the direct effect and final error for the ± 2 , 5, 8, and 11% peak body weight fields. For both *B* and *D*, the solid line is the best-fit line across subjects and the dashed line's slope is 1. Symbols represent data from 6 different subjects.

DISCUSSION

We have shown that motor adaptation to an externally applied force field is mathematically equivalent to a process that greedily minimizes a weighted sum of kinematic error and muscular effort. Optimization of kinematics or effort alone fails to capture the observed learning dynamics. The optimization can be implemented in one simple computation by an error-based learning law with a forgetting factor. This learning algorithm yields a linear autoregressive error equation (i.e., closed-looped learning dynamics) identical to one previously identified for reaching and for walking (Donchin et al. 2003; Emken and Reinkensmeyer 2005; Scheidt et al. 2001; Thoroughman and Shadmehr 2000). We show here that the coefficients of this equation must fall within an admissible range for the optimization model to be valid; we found that the experimentally identified coefficients for two experimental protocols involving 16 subjects did indeed fall in this range. Thus if one accepts that the autoregressive *Eq. 1* is a good fit to the data, as demonstrated before and re-checked here, the theoretical result given here shows that an appropriate high-level interpretation of these learning dynamics is that they minimize error and effort.

Further, we examined an experimental condition in which the rote learning dynamics of *Eq. 1* might not be expected to apply. We found that adaptation and de-adaptation occur in a stereotypical fashion when subjects are exposed repeatedly to a force field (i.e., no evidence of retention after a washout period and re-exposure). Finally, we found that steady-state kinematic error increases with increasing force-field strength in a way precisely predicted by the cost function. It is somewhat remarkable to us that an equation as simple as *Eq. 1*, a first-order linear equation, was competent to predict learning dynamics in these experimental conditions. This means, equivalently, that the optimization model is also applicable to these experimental conditions, given that the coefficients identified for *Eq. 1* for these experiments fell in the admissible range.

Relationship to some previous work in motor adaptation, optimization, and supervised learning

The proposed learning algorithm only uses local information of error and effort during the optimization. Thus in the framework proposed here, motor adaptation is driven by local gradients of the cost function rather than by the search of globally optimal solutions. Such an approach is only “one step ahead” (or “greedy”) in the sense that the learning dynamics are driven by a recursive update rule that does not necessarily yield to the best possible control action. It is unclear how well such a greedy optimization would generalize to more complex tasks and dynamic environments and whether the motor system can resort to more sophisticated optimization techniques. However, it is striking that such a simple process can capture the dynamics of motor adaptation.

Optimization principles have been widely proposed in motor control research to explain motor phenomena (for an extensive review, see Todorov 2004). For example, the spatiotemporal properties of reaching trajectories can be explained by a controller that minimizes jerk (Hogan 1984), torque change (Uno et al. 1989), or reaching variance in the presence of signal dependent noise (Harris and Wolpert 1998). Minimization of energy can explain locomotor trajectories (Anderson and

Pandy 2001). Optimal feedback control strategies can account for the way that the motor system manages redundant degrees of freedom in reaching and grasping tasks (Todorov and Jordan 2002). The present work is different from this previous work in that it does not consider either the details of the movement trajectory or the use of redundant degrees of freedom—the trajectory is abstracted by a single scalar measure—“the trajectory error.” Instead the focus of the optimization model presented here is in explaining how this trajectory error evolves in the presence of an external, perturbing force field. The fact that different cost functions, including the one found here, can explain different aspects of motor behaviors may indicate that the motor system changes the weighting of different costs across different tasks and types of movements (Bays and Wolpert 2007).

The model proposed also corresponds to a supervised learning algorithm because it uses information about errors to adjust its output. Different supervised learning schemes have been proposed for the acquisition of inverse models in feedforward motor control (Jordan 1996; Kawato and Gomi 1992; Wolpert and Kawato 1998). The most relevant is the feedback error learning theory proposed by Kawato et al. (1987) that suggests that the inverse model is acquired by using a feedback controller proportional to the kinematic error. The optimization model presented here extends such an error-based learning law by including a forgetting factor in the error-based learning law. This forgetting factor is directly linked to the inclusion of an effort term in the cost function that is minimized.

Some previous studies of motor adaptation proposed error-based learning laws that implicitly (Donchin et al. 2003; Scheidt et al. 2001; Thoroughman and Shadmehr 2000) or explicitly (Emken and Reinkensmeyer 2005; Liu and Reinkensmeyer 2004; Smith et al. 2006; Emken, Benitez, and Reinkensmeyer, 2007) included forgetting factors [also called “slacking” (Reinkensmeyer et al. 2004) or “retention” factors (Smith et al. 2006)] in error-based learning laws. The present paper extends this research by making a connection between such forgetting factors and an effort term in a cost function that is minimized.

A prediction of the presence of an effort term in the proposed optimization is the existence of a linear relationship between final kinematic error and external field strength. The measurements of step height adaptation for walking in different strength force fields confirmed the accuracy of this prediction. This finding is consistent with previous work for reaching that showed that steady-state kinematic error increased with field strength (Lai et al. 2003). However, the present work is novel in that it specifies the functional form and an interesting high level interpretation for the dependence.

We specified here that the muscular force that is minimized is an absolute muscular force, but it could also be expressed as a relative force, such as percent maximum voluntary force. Minimizing an absolute force carries the implication of metabolic cost minimization, in which case the error-effort tradeoff might be more observable for movements that involve large muscles, such as walking, versus movements that use small muscles, such as eye movement. Minimizing a relative force is similar to the concept of minimizing sense of effort or perhaps fatigue. The optimization model is amenable to both formulations, since the weighting coefficient in the cost function could be adjusted to account for a relative versus absolute formulation.

Falsifying the error/effort optimization model of motor adaptation

It is interesting to speculate about how the optimization model proposed here might be shown to be invalid. One interesting direction for testing the model concerns an experimental finding from Scheidt et al. (2000), who measured how subjects behaved when they moved in a stiff, virtual channel after adapting to a perpendicular viscous force field. The forces that the subjects had been generating to cancel the force field were now reacted by the channel so that subjects could not perceive the switch in environment. However, the lateral forces against the channel were unnecessary to move accurately to the target. Scheidt et al. found that subjects slowly decreased the contact force against the channel with reaching practice. This decrease is consistent with the presence of a forgetting factor that scales down previous muscle activations when kinematic error is small as proposed here.

However, if we compare the current results with those of (Scheidt et al. 2000) more closely, we find a discrepancy. Specifically, when kinematic error is zero, then the motor command u will decay according to a simplified version of Eq. 8 $u_{i+1} = fu_i$, which has the solution

$$u_i = e^{-i/\tau} u_0 \quad \tau = \frac{-1}{\ln(f)} \Rightarrow f = e^{-1/\tau} \quad (15)$$

Using the mean value for f found here for stepping (0.76), $\tau = 3.6$ trials. However, (Scheidt et al. 2000) found the time constant of decay of force in the virtual channel (i.e., “0-error”) condition to be 138 trials. The difference might be explained by differences in protocol (walking in a relatively strong force field in gravity vs. reaching in a relatively weak force field with the arm supported). Alternately, the optimization model proposed here may not adequately explain the behavior of the motor system when kinematic errors are small; more research is clearly needed.

Another possible path for falsifying the model is to identify history-dependent changes in adaptation that depend on experiences more than one movement back in time. There is indeed evidence that the motor system can retain an internal model learned in one session, reducing error in a subsequent session (Brashers-Krug et al. 1996; Gandolfo et al. 1996; Shadmehr and Brashers-Krug 1997). Similar to Caithness et al. (2004), however, we found no evidence of retention after repeated exposure to the force field after deadadaptation in the null field. The explanation offered by Caithness, which we concur with, is that the experimental protocol used here hindered contextual association of the force field with the robot by applying a null field immediately following any force-field stage. However, it has been shown that retention in robotic reaching tasks is possible as long as the robot applies consistent forces that are linked with appropriate contextual clues (Wada et al. 2003). The optimization model proposed here cannot account for such retention. It also cannot account for changes in retention observed when subjects simply rest their hands between force-field exposures (Takahashi et al. 2006).

Another possible way the optimization model can be shown to have limited applicability is by showing that the dynamics of Eq. 1 do not capture the learning dynamics as well as other models. Indeed the autoregressive dynamics in Eq. 1 appear to consistently overestimate the initial learning rate and then

underestimate the subsequent learning rate (Fig. 5), although this mismatch was not statistically significant with the small number of subjects studied here. A recent model that includes two interacting adaptive processes, each with a forgetting (or "retention" factor), may better capture these dynamics (Smith et al. 2006). Determining how such a two time constant model links to higher level optimization goals is an interesting area for future research. We note that one of the linear equations that we tested (model 12 in Fig. 4.) is equivalent to the two-state, gain-independent model in (Smith et al. 2006). We did not find this equation to be better at explaining data as a whole. This may be because the added value of the higher-order terms in this equation is primarily in fitting those periods of time when error changes rapidly, which were periods of short durations in the present experiments. In other words, our experimental paradigms may have been biased toward a model that fits well steady-state behavior at different force levels.

Surprisingly to us, we found that the coefficients of Eq. 1 must lie within an admissible range, defined by Eq. 10, for the optimization model to have validity. The fact that Eq. 1 could well fit experimental data, but its coefficients be inconsistent with the optimization model, means that the optimization model is falsifiable in a straightforward way. The coefficients identified for the stepping tasks studied here did indeed fall in the admissible range. Finding tasks in which they do not fall in this range may give insight into when the motor system does not use an optimization approach for motor adaptation.

Finally, other shortcomings of the model include the fact that we ignored the known dependence of limb stiffness on force (Gomi and Osu 1998; Perreault et al. 2001). The model likely still worked fairly well because the forces we applied to cause adaptation were small relative to the forces already generated by the legs to control walking. Larger force perturbations, and a nonconstant stiffness, might be accounted for by a cost-function weighting factor λ that changes with force, or other forms of costs besides quadratic. The model also cannot account for explicit impedance control, which has been observed in unstable and noisy environments (Burdet et al. 2001; Takahashi et al. 2001), and incorporating impedance considerations in the cost function is a direction for future research. We also chose to study a repetitive movement paradigm (stepping on a treadmill at a constant speed) for which the desired leg trajectory is approximately the same each movement. However, it is known that when people reach along different trajectories, experiences along one trajectory can generalize to others (Conditt et al. 1997; Goodbody and Wolpert 1998; Shadmehr and Mussa-Ivaldi 1994). An interesting question is how the optimization model can be extended to account for such generalization.

In summary, it is likely that experimental paradigms can be generated for which Eq. 1 (or the admissibility condition for its coefficients) does not fit the data well, and thus for which the link to the error/effort optimization model of adaptation is broken. Indeed several such paradigms already appear to exist. What will likely remain true, however, is that the error/effort optimization model provides a good explanation for steady-state behavior after adaptation, and a reasonable approximation to transient behavior, during first exposure to a range of novel dynamic environments.

ACKNOWLEDGMENTS

Present address of J. L. Emken: California Institute of Technology, MC 216-76, Pasadena, CA, 91125 (E-mail: emken@caltech.edu.).

GRANTS

This work was supported by National Institute of Standards and Technology Advanced Technology Program Grant 00-00-4906, National Center for Research Resources Grant M01RR00827, National Institute on Disability and Rehabilitation Research Grant H133E020732, an Achievements Rewards for College Scientists Foundation Scholarship for J. L. Emken, and a postdoctoral Balsells fellowship from the California-Catalonia Engineering Program for R. Benitez.

REFERENCES

- Anderson FC, Pandy MG. Dynamic optimization of human walking. *J Biomech Eng* 123: 381–390, 2001.
- Bays PM, Wolpert DM. Computational principles of sensorimotor control that minimize uncertainty and variability. *J Physiol* 387–396, 2007.
- Brashers-Krug T, Shadmehr R, Bizzi E. Consolidation in human motor memory. *Nature* 382: 252–255, 1996.
- Burdet E, Osu R, Franklin DW, Milner TE, Kawato M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414: 446–449, 2001.
- Caithness G, Osu R, Bays P, Chase H, Klassen J, Kawato M, Wolpert DM, Flanagan JR. Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *J Neurosci* 24: 8662–8671, 2004.
- Cheney PD, Fetz EE. Functional classes of primate corticomotoneuronal cells and their relation to active force. *J Neurophysiol* 44: 773–791, 1980.
- Conditt MA, Gandolfo F, Mussa-Ivaldi FA. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78: 554–560, 1997.
- Conditt MA, Mussa-Ivaldi FA. Central representation of time during motor learning. *Proc Natl Acad Sci USA* 96: 11625–11630, 1999.
- Donchin O, Francis JT, Shadmehr R. Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J Neurosci* 23: 9032–9045, 2003.
- Emken JL, Benitez R, Reinkensmeyer DJ. Human-robot cooperative movement training: Learning a novel sensory motor transformation during walking with robotic assistance-as-needed. *J NeuroEngin and Rehab* 4: 8, 2007.
- Emken JL, Reinkensmeyer DJ. Robot-enhanced motor learning: accelerating internal model formation during locomotion by transient dynamic amplification. *IEEE Trans Neural Systems and Rehab Engng* 13: 33–39, 2005.
- Emken JL, Wynne JH, Harkema SJ, Reinkensmeyer DJ. A robotic device for manipulating human stepping. *IEEE Trans Robotics* 22: 185–189, 2006.
- Gandolfo F, Mussa-Ivaldi FA, Bizzi E. Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843–3846, 1996.
- Ghez C, Sainburg R. Proprioceptive control of interjoint coordination. *Can J Physiol Pharmacol* 73: 273–284, 1995.
- Gomi H, Osu R. Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *J Neurosci* 18: 8965–8978, 1998.
- Goodbody SJ, Wolpert DM. Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825–1838, 1998.
- Griffin TM, Kram R. Penguin waddling is not wasteful. *Nature* 408:929, 2000.
- Harris CM, Wolpert DM. Singal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Haykin SS. *Adaptive Filter Theory*. Upper Saddle River, NJ: Prentice Hall, 1996.
- Hogan N. An organizing principle for a class of voluntary movements. *J Neurosci* 4: 2745–2754, 1984.
- Jordan MI. Computational aspects of motor control and motor learning. In: *Handbook of Perception and Action: Motor Skills*, edited by Heuer H and Keele S. New York: Academic, 1996.
- Kawato M, Furukawa K, Suzuki R. A hierarchical neural-network model for control and learning of voluntary movement. *Biol Cybern* 57: 169–185, 1987.
- Kawato M, Gomi H. The cerebellum and VOR/OKR learning models. *Trends Neurosci* 15: 445–453, 1992.
- Lackner JR, Dizio P. Gravitoinertial force background level affects adaptation to Coriolis force perturbations of reaching movements. *J Neurophysiol* 80: 546–553, 1998.

- Lai EJ, Hodgson AJ, Milner TE.** Influence of interaction force levels on degree of motor adaptation in a stable dynamic force field. *Exp Brain Res* 153: 76–83, 2003.
- Liu J, Reinkensmeyer DJ.** Motor Adaptation as an optimal combination of computational strategies. *26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, San Francisco, CA, 2004*, p. 4025–4028.
- Montgomery DC, Peck EA, Vining GG.** *Introduction to Linear Regression Analysis*. New York: Wiley, 2001.
- Nakano E, Imamizu H, Osu R, Uno Y, Gomi H, Yoshioka T, Kawato M.** Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. *J Neurophysiol* 81: 2140–2155, 1999.
- Perreault EJ, Kirsch RF, Crago PE.** Effects of voluntary force generation on the elastic components of endpoint stiffness. *Exp Brain Res* 141: 312–323, 2001.
- Reinkensmeyer DJ, Emken JL, Liu J, Bobrow JE.** The nervous system appears to minimize a weighted sum of kinematic error, force, and change in force when adapting to viscous environments during reaching and stepping. *Adv Computat Mot Control* 3 [http://www.bme.jhu.edu/acmc/].
- Scheidt RA, Dingwell JB, Mussa-Ivaldi FA.** Learning to move amid uncertainty. *J Neurophysiol* 86: 971–985, 2001.
- Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA.** Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84: 853–862, 2000.
- Shadmehr R, Brashers-Krug T.** Functional stages in the formation of human long-term motor memory. *J Neurosci* 17: 409–419, 1997.
- Shadmehr R, Moussavi ZMK.** Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20: 7807–7815, 2000.
- Shadmehr R, Mussa-Ivaldi FA.** Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208–3224, 1994.
- Smith MA, Ghazizadeh A, Shadmehr R.** Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4: 1035–1043, 2006.
- Takahashi CD, Nemet D, Rose-Gotttron CM, Larson JK, Cooper DM, Reinkensmeyer DJ.** Effect of muscle fatigue on internal model formation and retention during reaching with the arm. *J Appl Physiol* 100: 695–706, 2006.
- Takahashi CD, Scheidt RA, Reinkensmeyer DJ.** Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J Neurophysiol* 86: 1047–1051, 2001.
- Thoroughman KA, Shadmehr R.** Learning of action through adaptive combination of motor primitives. *Nature* 407: 742–747, 2000.
- Todorov E.** Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.
- Todorov E, Jordan M.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
- Uno Y, Kawato M, Suzuki R.** Formation and control of optimal trajectory in human multijoint arm movement: minimum torque-change model. *Biol Cybern* 61: 89–101, 1989.
- Wada Y, Kawabata Y, Kotosaka S, Yamamoto K, Kitazawa S, Kawato M.** Acquisition and contextual switching of multiple internal models for different viscous force fields. *Neurosci Res* 46: 319–331, 2003.
- Wolpert D, Kawato M.** Multiple paired forward and inverse models for motor control. *Neural Netw* 11: 1317–1329, 1998.