

4-1-2016

# Optimal Schedules in Multitask Motor Learning

Jeong Yoon Lee

*University of Southern California*

Youngmin Oh

*University of Southern California*

Sung Shin Kim

*University of Southern California*

Robert A. Scheidt

*Marquette University, robert.scheidt@marquette.edu*

Nicolas Schweighofer

*University of Montpellier, France*

## Optimal Schedules in Multitask Motor Learning

**Jeong Yoon Lee**

*ethielee@gmail.com*

*Computer Science, University of Southern California, Los Angeles,  
CA 90089, U.S.A.*

**Youngmin Oh**

*youngmio@usc.edu*

**Sung Shin Kim**

*sung.kim3@northwestern.edu*

*Neuroscience Graduate Program, University of Southern California,  
Los Angeles, CA 90089, U.S.A.*

**Robert A. Scheidt**

*robert.scheidt@marquette.edu*

*Biomedical Engineering, Marquette University, Milwaukee, WI 53233,  
and Physical Medicine and Rehabilitation, Northwestern University  
Feinberg School of Medicine, Chicago, IL 60611, U.S.A.*

**Nicolas Schweighofer**

*schweigh@usc.edu*

*Biokinesiology and Physical Therapy, University of Southern California,  
Los Angeles, CA 90089, U.S.A., and M2H Laboratory, Euromov,  
University of Montpellier I, 34090 Montpellier, France*

**Although scheduling multiple tasks in motor learning to maximize long-term retention of performance is of great practical importance in sports training and motor rehabilitation after brain injury, it is unclear how to do so. We propose here a novel theoretical approach that uses optimal control theory and computational models of motor adaptation to determine schedules that maximize long-term retention predictively. Using Pontryagin's maximum principle, we derived a control law that determines the trial-by-trial task choice that maximizes overall delayed retention for all tasks, as predicted by the state-space model. Simulations of a single session of adaptation with two tasks show that when task interference is high, there exists a threshold in relative task difficulty below which the alternating schedule is optimal. Only for large differences in task difficulties do optimal schedules assign more trials to the harder task.**

---

J.Y.L., Y.O., R.A.S., and N.S. contributed equally.

**However, over the parameter range tested, alternating schedules yield long-term retention performance that is only slightly inferior to performance given by the true optimal schedules. Our results thus predict that in a large number of learning situations wherein tasks interfere, intermixing tasks with an equal number of trials is an effective strategy in enhancing long-term retention.**

## 1 Introduction

---

The need for effective scheduling of multiple motor tasks is ubiquitous in activities such as sports, music, professional skill development, and motor rehabilitation after brain injury. However, how should the coach or the therapist schedule multiple tasks? Let us consider the case in which two tasks need to be practiced in a single session. Given the negatively accelerated shape of performance improvement as a function of practice (Liu, Mayer-Kress, & Newell, 2003), a simple possibility would be to practice one task until it reaches some performance criterion and then practice the other task. There is robust evidence, however, that such blocked schedules are detrimental to long-term retention (Schmidt & Lee, 2005). In contrast, intermixing the two tasks reduces initial learning speed but enhances long-term retention (Schmidt & Lee, 2005; Schweighofer et al., 2011). But if one task is more difficult than the other (we will propose an operationalized definition of *difficulty* below, but let us assume for the moment that difficulty is measured by the initial rate of change in performance) or if the learner has prior experience with one of the two tasks, trial-by-trial gains in the easier task will soon plateau. This “labor in vain” will possibly yield overall poorer retention because of insufficient training on the second task. The more difficult task should therefore receive a greater number of trials. However, adding trials to one task will increase the length of trial blocks for this task, and such blocked schedules may decrease long-term retention.

How, then, can we resolve the conundrum of increasing the number of trials for the more difficult task while also minimizing the deleterious effect of long blocks of same-task trials that must inevitably arise in the schedule? One possibility is to select the task at each trial based on predicted performance on the next trials (see Huang, Shadmehr, & Diedrichsen, 2008; Simon, Cullen, & Lee, 2002). Unfortunately, current performance is known to be a poor predictor of long-term retention (Joiner & Smith, 2008). Task selection must therefore be based on long-term retention. We previously showed that adaptive schedules based on performance measured on delayed-retention tests substantially improves learning compared to scheduling based on current performance (Choi, Qi, Gordon, & Schweighofer, 2008). In that previous study, however, scheduling was based on heuristics and was determined “postdictively,” that is, after performance on long-term retention test was available. To further enhance retention, it would thus be desirable to schedule the tasks predictively—early in training and without the need to wait

for long-term retention data to be available. Determining such schedules must therefore be based on predictions of long-term retention generated by computational models of motor memory.

Here, because of the availability of sound computational models, we use motor adaptation as a proxy for motor learning. *Motor adaptation* is defined as changes in motor performance that allow the motor system to regain its former capabilities in altered circumstances. Previous computational models suggest that motor adaptation occurs at multiple timescales. In the two-state model (Smith, Ghazizadeh, & Shadmehr, 2006), a fast learning process (FLP) contributes to fast initial learning but also forgets quickly. A slow learning process (SLP) contributes to long-term retention (Joiner & Smith, 2008) but learns slowly. Each process has a single state to store the accumulated adaptation. Such two-state models cannot explain dual- or multiple-task adaptation, however, because sufficient adaptation to a new task overrides adaptation of a previous task. When given contextual cues and sufficient trials, humans can simultaneously adapt to two visuomotor rotations (Choi et al., 2008; Imamizu et al., 2007; Lee & Schweighofer, 2009), two saccadic gains (Shelhamer, Aboukhalil, & Clendaniel, 2005), and in some conditions two opposite force fields (Hirashima & Nozaki, 2012; Osu, Hirai, Yoshioka, & Kawato, 2004). The MODular Selection And Identification for Control (MOSAIC) model (Wolpert & Kawato, 1998) naturally accounts for dual or multiple adaptations, via nonlinear switching among multiple parallel internal models based on “responsibility signals,” which estimate the extent to which each model should act to capture the behavior in the current situational context. The responsibility signals have the property that they lie between 0 and 1, and their sum over the models is exactly 1. In previous work (Lee & Schweighofer, 2009), we proposed a model with a fast process that contains a single state arranged in parallel with multiple slow processes switched by contextual cues. We now extend that model to include responsibility signals that control learning within multiple adaptive systems.

Computational models of motor adaptation allow us to predict long-term retention performance for a task given a specific training schedule and therefore enable us to compare the effectiveness of different schedules. How then can we find schedules that maximize long-term retention? A naive approach would be to select the best schedule after comparison of all possible schedules. This approach becomes rapidly intractable, however, as the number of trials grows. For instance, for 2 tasks and 100 total trials, the number of possible schedules is  $2^{100} > 10^{30}$ . Even if we could evaluate 1 billion schedules per second, finding the optimal schedule would take longer than a thousand times the age of the universe! Thus, a brute-force search is clearly impossible for schedules longer than short schedules.

Here, we propose a novel theoretical and computationally tractable method to determine training schedules that maximizes long-term retention. Our method uses a combined approach of computational models of

motor adaptation and optimal control theory. Optimal control theory deals with the problem of finding a control law for a given system to achieve an optimality criterion. In our example of single-session adaptation training for two tasks, the optimality criterion is to maximize the predicted slow processes of both tasks at the end of training (we made this choice because the slow process, but not the fast process or the overall level of adaptation, correlates with long-term retention; Joiner & Smith, 2008). The optimal control law then determines the choice of the task to be presented at every trial. We validate our method in simulations of a single session of adaptation with two tasks, with various lengths of training, and with various relative task difficulty levels. We compared the results with those of a genetic algorithm (GA) optimization method and, for the specific case of a small schedule with 20 trials, with those of a brute-force search.

## 2 Materials and Methods

---

The purpose of this study is to combine computational models of motor adaptation and analysis techniques from optimal control theory to identify multitask training schedules that maximize long-term retention of learning. In this section, we first describe possible models of motor adaptation and a formulation of the problem to be solved, then the optimal control method to determine the schedules, and finally our simulation setup. Note that while the approach we describe is not tied to any particular computational model, the models of adaptation dynamics used here are linear with respect to trials (i.e., discretized time; see Scheidt, Dingwell, & Mussa-Ivaldi, 2001; Judkins & Scheidt, 2014).

**2.1 Modeling Multitask Motor Learning.** Whereas the conceptual MOSAIC model of Wolpert and Kawato (1998) accounts for multiple adaptations by switching among multiple parallel and independent internal models (see the related multiple parallel model in Figure 1A), experimental results from a recent study requiring dual task learning support a refined model with a single fast adaptive state arranged in parallel with multiple slow processes switched on the basis of contextual cues (Lee & Schweighofer, 2009). Here, we extend this 1FnS model to accommodate different learning and forgetting rates for the different tasks while also allowing the task-dependent modules to compete in determining behavior.

Consider the special case of an adaptation paradigm with two tasks, and therefore employ an adaptation model having one common fast state and two slow states, each specific to one task (i.e., a 1F2S model). The learning dynamics for the 1F2S model are described by equations 2.1 through 2.4. Specifically, the update equation for the shared fast state is given by

$$x_{k+1}^f = a^f \cdot x_k^f + b^f \cdot e_k, \quad (2.1)$$

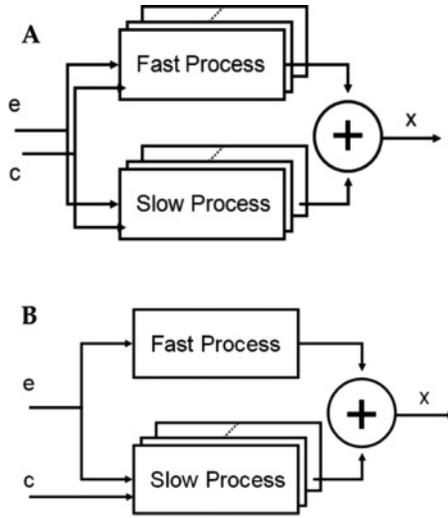


Figure 1: Two examples of multiple-task adaptation models. (A) The  $n$ -fast  $n$ -slow model (nFnS) model akin to the MOSAIC model (Wolpert & Kawato, 1998). Error  $e$  updates the selected pair of fast and slow processes corresponding to contextual cue  $c$ , which protects the unselected pairs of fast and slow processes from interference. (B) The 1-fast  $n$ -slow model (1FnS) model.  $e$  updates a common fast process and one of  $N$  parallel slow processes selected by  $c$ , which protects the unselected slow processes from interference.

where  $x_k^f$  corresponds to the fast state on trial  $k$ , constants  $a^f$  and  $b^f$  correspond to the state retention and error gain parameters, respectively, and  $e_k$  corresponds to the performance error on trial  $k$ . The update equations for the slow states are given by

$$\begin{aligned} x_{k+1}^{s1} &= a^{s1} \cdot x_k^{s1} + b^{s1} \cdot u_k^{s1} \cdot e_k, \\ x_{k+1}^{s2} &= a^{s2} \cdot x_k^{s2} + b^{s2} \cdot u_k^{s2} \cdot e_k, \end{aligned} \quad (2.2)$$

where the variables  $u_k^{s1}$  and  $u_k^{s2}$  are mutually exclusive task selection variables that determine which task influences performance on trial  $k$  and which slow state is to be updated based on the performance error. In our model,  $u_k^{s1}$  and  $u_k^{s2}$  are determined by contextual cues. As described below, the values of  $u_k^{s1}$  and  $u_k^{s2}$  reflect the result of a competition between responsibility signals  $r_k^{s1}$  and  $r_k^{s2}$  associated with the slow state components of the adaptation model. Performance on trial  $k$  is given by

$$y_k = x_k^f + x_k^{s1} \cdot u_k^{s1} + x_k^{s2} \cdot u_k^{s2}, \quad (2.3)$$

whereas the performance error is given by

$$e_k = (\hat{y}^{s1} - y_k) \cdot u_k^{s1} + (\hat{y}^{s2} - y_k) \cdot u_k^{s2}, \quad (2.4)$$

where  $\hat{y}^{s1}$  and  $\hat{y}^{s2}$  correspond to the desired motor outputs for tasks 1 and 2, respectively. For the special case of two tasks, it is possible to define a single task selection variable  $u_k$  (no superscript) as  $u_k = u_k^{s1}$  such that  $u_k^{s2} = 1 - u_k$  (i.e.,  $u_k^{s1}$  and  $u_k^{s2}$  sum to 1). To enforce the exclusivity condition such that only one model is selected on any given trial, we further constrain the task selection variables such that  $u_k^{s1} \cdot u_k^{s2} = 0$  or, equivalently,

$$C(u_k) = u_k(1 - u_k) = 0. \quad (2.5)$$

*2.1.1 Schedules That Maximize Long-Term Retention in Multitask Motor Learning.* To derive the optimal schedule, it is necessary to specify an optimality criterion or “cost function”  $J_k$ , typically defined as the sum of path costs (i.e., the cost rate  $l(\cdot)$ ) and final costs (i.e., boundary costs  $h(\cdot)$ ) (cf., Bryson & Ho, 1969). This cost function is subject to dynamic constraints described by equations 2.1 to 2.3 and a constraint on the task selection variable described by equation 2.5.

Our goal is to maximize long-term retention of performance for both tasks. Long-term retention for any given task depends on the final state of the slow process for that task, as current performance itself is not a good indicator of long-term retention (Joiner & Smith, 2008). The intermediate cost rate  $l(\cdot)$  is therefore zero, and we define the final cost as the average mean square difference between the desired performance and the slow process for each task at the end of training. The cost function  $J$  is therefore

$$J(\mathbf{x}) = h_L = h(x_L^{s1}, x_L^{s2}) = [(\varepsilon_L^{s1})^2 + (\varepsilon_L^{s2})^2]/2 \quad (2.6)$$

where we define the slow state errors  $\varepsilon_k^{s1} = (x_k^{s1} - \hat{y}^{s1})$  and  $\varepsilon_k^{s2} = (x_k^{s2} - \hat{y}^{s2})$  and  $L$  as the total number of trials in the training sequence. Thus, the scheduling problem is solved by minimizing the difference between the values of the slow state memories and their desired values at the end of the training schedule. That is, the optimal training schedule  $u_{1:L}^*$  is the one that minimizes equation 2.6 over all possible training schedules, thereby maximizing long-term retention driven by task-specific, slowly decaying motor memories:

$$u_{1:L}^* = \arg \min_{u_{1:L}}(J). \quad (2.7)$$

## 2.2 Deriving Optimal Schedules via Pontryagin's Maximum Principle.

We used Pontryagin's maximum principle (Kirk, 1992) and the 1F2S model to determine the optimal schedule for the two-task adaptation paradigm. Pontryagin's maximum principle allows a difficult optimization problem over many time steps (e.g., see equation 2.7) to be reduced to a series of simpler optimization problems over single time steps (see below). To do so, one can define an augmented Hamiltonian function  $H(\cdot)$  (Kirk, 1992) as a weighted sum of the cost function  $J(\cdot)$  and additional costs associated with the dynamic and control constraints:

$$H_k(\mathbf{x}_k, u_k, \lambda_{k+1}, \gamma_{k+1}) = J(\mathbf{x}) + f(\mathbf{x}_k, u_k)^T \cdot \lambda_{k+1} + C(u_k) \cdot \gamma_{k+1}, \quad (2.8)$$

where  $\mathbf{x}_k = \{x_k^f, \varepsilon_k^{s1}, \varepsilon_k^{s2}\}^T$  is the state vector,  $f(\mathbf{x}_k, u_k)$  represents the system dynamics corresponding to equations 2.1 and 2.2,  $\lambda$  is the Lagrange multiplier vector whose elements are the costates associated with the state constraints, and  $\gamma$  is the costate associated with the control constraint. When the Hamiltonian is minimized with respect to small changes in the state and control variables, three desirable conditions are satisfied: the final cost is minimized, the state transition dynamics are enforced, and the control constraint is enforced. We use the Hamiltonian to identify the training schedule that satisfies equation 2.7 while also satisfying the constraints imposed by the system dynamics and control constraint.

More specifically, taking the partial derivative of the Hamiltonian with respect to the  $\lambda$  costates and the  $\gamma$  costate yields the constraints (see equations 2.1–2.2 and 2.5, respectively). Taking the partial derivative of the Hamiltonian with respect to the states  $x_k$  yields the costate update equations for our minimization problem:

$$\lambda_k^1 = \frac{\partial H}{\partial x^f} = (a^f - b^f) \cdot \lambda_{k+1}^1 - b^{s1} \cdot r_k \cdot \lambda_{k+1}^2 - b^{s2} \cdot (1 - r_k) \cdot \lambda_{k+1}^3, \quad (2.9)$$

$$\lambda_k^2 = \frac{\partial H}{\partial \varepsilon^{s1}} = -b^f \cdot r_k \cdot \lambda_{k+1}^1 + (a^{s1} - b^{s1} \cdot r_k) \cdot \lambda_{k+1}^2 + \varepsilon_k^{s1}, \quad (2.10)$$

$$\lambda_k^3 = \frac{\partial H}{\partial \varepsilon^{s2}} = -b^f \cdot (1 - r_k) \cdot \lambda_{k+1}^1 + (a^{s2} - b^{s2} \cdot (1 - r_k)) \cdot \lambda_{k+1}^3 + \varepsilon_k^{s2}, \quad (2.11)$$

$$\begin{aligned} \gamma_k = \frac{\partial H}{\partial r} &= (-b^f \cdot \varepsilon_k^{s1} + b^f \cdot \varepsilon_k^{s2}) \cdot \lambda_{k+1}^1 + (-b^{s1} \cdot x_k^f - b^{s1} \cdot \varepsilon_k^{s1}) \cdot \lambda_{k+1}^2 + \dots \\ &+ (-b^{s2} \cdot x_k^f - b^{s2} \cdot \varepsilon_k^{s2}) \cdot \lambda_{k+1}^3 + (1 - 2r_k) \cdot \gamma_{k+1}. \end{aligned} \quad (2.12)$$

Here, we have defined a differentiable "responsibility signal"  $r_{k'}$  which corresponds to the discrete task selection variable  $u_k$  with its exclusivity constraint relaxed. This step is necessary so that the partial derivative of

the Hamiltonian with respect to variations in the task selection sequence be nonsingular, according to equation 2.12. Responsibility signals in multimodule adaptive systems have the properties that they lie between zero and one and sum to one over all contributing models (Wolpert & Kawato, 1998). Hence, the responsibility signal represents the extent to which each model accounts for the behavior of the system (task 1:  $r_k$ ; task 2:  $1 - r_k$ ).

The optimal control sequence can be determined using an approach based on Todorov (2007). First, Given an arbitrary initial sequence of responsibilities  $r_{1:L}$ , generate the task-exclusive sequence  $u_k$  by enforcing a winner-take-all competition on  $r_k$  (e.g., by rounding up or down to 1 or 0). Iterate the system dynamics forward in time (i.e., trial by trial) to obtain a candidate sequence of states. Second, with the resulting responsibility and state sequences defined, iterate equations 2.9 through 2.12 backward in time to obtain the costate sequences. At each time step (trial), improve the candidate responsibility sequence via gradient descent of the Hamiltonian:

$$r_k = r_k - \alpha \cdot dr_k, \quad (2.13)$$

where  $\alpha$  is a small update rate, and

$$dr_k = \frac{\partial H}{\partial r}. \quad (2.14)$$

These two steps are repeated until  $r_k$  has converged to  $r_{1:L}^*$ .

*2.2.1 Comparison with Genetic Algorithm and Brute-Force Search Methods.* Although the deterministic Pontryagin's maximum principle yields the true optimal result in theory, our simulation results are not guaranteed to always return the true optimal. This is because the result depends on the initial schedule, and iteration stops when incremental reduction in cost becomes smaller than a given threshold (e.g.,  $10^{-10}$ ), which conceivably could settle into a local, rather than global, minimum. In order to verify the validity of our theoretical methods, we applied a genetic algorithm (GA) method to determine optimal schedules and then compared the results with those from Pontryagin's maximum principle. The GA is a stochastic optimization algorithm: a pool of schedules (i.e., genes) in each "generation" of the simulation can exchange a random portion of the schedule (genetic crossover) and can randomly change bits of the schedule (genetic mutation). Only schedules with better performance—those schedules that minimize the cost of equation 2.6—survive to the next generation (i.e. eliticism).

In addition, for a small number of total trials  $K = 20$ , we performed a brute-force search to calculate costs of all possible  $2^{20}$  binary schedules (see appendix B). We then compared the optimal schedule obtained using Pontryagin's maximum principle method with the true optimal schedule from the brute-force search.

**2.3 Simulations.** We first computed optimal schedules in simulations of a single-session adaptation paradigm with two tasks in which we varied both the relative task difficulties between the two tasks and the total number of training trials (20, 40, and 80 trials). Although difference in task difficulty is commonly experienced in actual motor learning, it is unclear how to operationalize it. Is a task difficult because initial progress is slow, later progress is slow, or final performance is low—or because of all of these? Here, to simplify, we modeled task difficulty with a single difficulty parameter that simultaneously affects initial change in performance, later change in performance, and final performance. Specifically, we defined relative task difficulty of the second task compared to the first task with a difficulty parameter  $d$  that affects both fast and slow process learning rates:

$$b^{f2} = b^{f1}/d; \quad b^{s2} = b^{s1}/d. \quad (2.15)$$

Note that we used two different fast learning rates ( $b^{f1}$  and  $b^{f2}$ ) to update the common fast process motor memory. For example, if the more difficult task is twice as difficult as the easier task, the same error results in only half the increase of the fast process. While fixing the learning rates of one task (the “easy task”), we increased task difficulty of the other task (the “difficult task”) from  $d = 1.0$  to  $d = 5.0$  as steps of 0.1. To simplify, the two tasks were assumed to have opposite signs with the same magnitude, such that  $\hat{y}^{s1} = 1$  and  $\hat{y}^{s2} = -1$ . The default parameter set, estimated in a previous visuomotor rotation experiment (Lee, 2011), was taken as  $a^f = 0.965$ ,  $a^s = 0.993$  and  $b^f = 0.597$ ,  $b^s = 0.114$ . To extend the validity of our results to other types of adaptation, we performed a sensitivity analysis for these parameters (see appendix A).

We also compared the costs of the two tasks after optimal, alternating, and blocked schedules. We set the initial schedule as the alternating schedule, because this schedule maximizes long-term retention in the case of equal task difficulties in the 1FnS model (Schweighofer et al., 2011). Starting with the alternating schedule, Pontryagin’s algorithm was repeated until the cost reduction became smaller than  $10^{-10}$  from one iteration to the next with the update rate  $\alpha = 0.05$ .

After we obtained the simulated optimal schedule (as a series of 0s and 1s, with 1 coding for presentation of the easy task and 0 for presentation of the difficult task), we computed the switching index and the percentage of trials for the difficult task. The switching index is the number of task switches divided by the maximum possible number of switches. Thus, for the initial alternating schedule, the percentage of trials for the difficult task is 50% and the switching index is 1. For the blocked schedules, the percentage of trials for the difficult task is still 50%, but the switching index is low and equal to  $1/(\text{total trial number} - 1)$ . In order to discount computational boundary effects deriving from the unavoidable arbitrary assignment of the

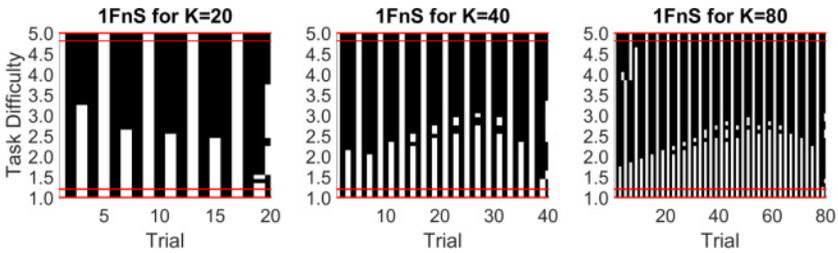


Figure 2: Simulated optimal schedules for total trial numbers 20, 40, and 80. White blocks show the trials at which the easy task is scheduled and the black blocks for the difficult task. Two optimal schedules for relative task difficulties  $d = 1.0$  and  $5.0$  are shown as examples within the red bars. Note that for  $d = 1.0$ , the optimal schedule is the alternating schedule regardless of the total number of training trials.

costate values at trial  $L + 1$  (see equations 2.9 to 2.12), the last two simulated trials were excluded from calculations of switching index and percentage of scheduled trials for the more difficult task.

For the GA simulations, we set the rate of crossover at 0.8 and the rates of mutation and eliticism at 0.03. We repeated this algorithm running through 1000 generations, starting from a population of 1000 random schedules. We finally chose the schedule in the last generation that minimizes the cost of equation 2.6. We then compared the schedules and cost obtained via GA and those obtained via the Pontryagin's maximum principle method.

### 3 Results

We first simulated optimal schedules for the 1F2S model with increasing values of the task difficulty parameter  $d$  and for 20, 40, and 80 total trials (see Figure 2). We chose these trial numbers because they typically span the number of trials needed for asymptotic performance in visuomotor adaptation experiments. When both tasks were of similar difficulty, the alternating schedule was the optimal schedule, with half the total trials assigned to each task. As the difficulty of the second task increased, the general trend was that more trials were assigned to the difficult task (black boxes in Figure 2). The resulting small trial blocks had the tendency to be distributed evenly throughout the training sequence.

Figure 3A shows the switching index (upper row) and the percentage of trials for the difficult task (lower row) as a function of relative task difficulty. There were relatively large thresholds of task difficulty below which the alternating schedule was optimal. These thresholds were 2.5, 2.1, and 1.8 for the total trial number 20, 40, and 80, respectively. As task difficulty increased further, the switching index decreased with several

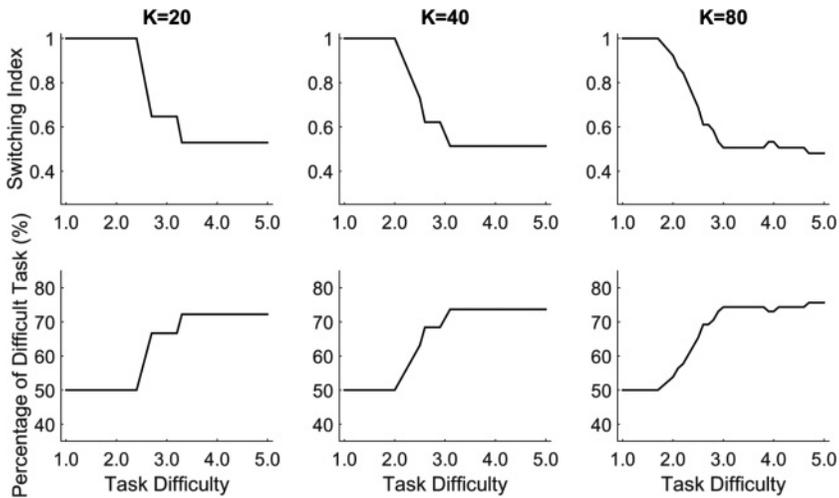


Figure 3: Nonlinear characteristics of the optimal schedules as a function of relative task difficulty for the 1F2S model. (Top row) Switching index, calculated as the number of switches of one task to the other divided by the maximum possible number of switches. (Bottom row) Percentage of number of trials for the difficult task. For the perfect alternating schedule, the switching index is 1, and the percentage of difficult task is 50%. Note that for task difficulty less than around 2.2, alternating schedules are optimal.

plateaus. For the simulated range of task difficulties, the final plateau of the switching index was around 0.53, 0.51, and 0.51 for training sequence lengths of 20, 40, and 80 trials, respectively. A similar trend can be seen in the increase in the percentage of trials for the difficult task (see Figure 3B), with a high (negative) correlation between the two quantities. This high correlation arose because the optimal schedules were small blocks of trials of the difficult task evenly separated by a single trial of the easy task, as illustrated in Figure 2. The final plateau for the percentage of the difficult task was around 72%, 74%, and 76% for the total trial number 20, 40, and 80, respectively. Overall, this indicates that three times more trials were assigned to the difficult task than the easy task when the tasks were different in difficulty by a factor of 5.

Updates of the slow and fast processes during training for optimal and alternating schedules are shown in Figure 4 for relative task difficulty  $d = 5.0$  and 40 training trials. The combined final value of the slow process states (i.e., the quadratic mean of the two slow processes) following the alternating schedule was 94% that of the optimal schedule (100% and 85% for the easy and difficult task, respectively). Therefore, the optimal schedule achieved not only better overall final retention, but also better balance between the

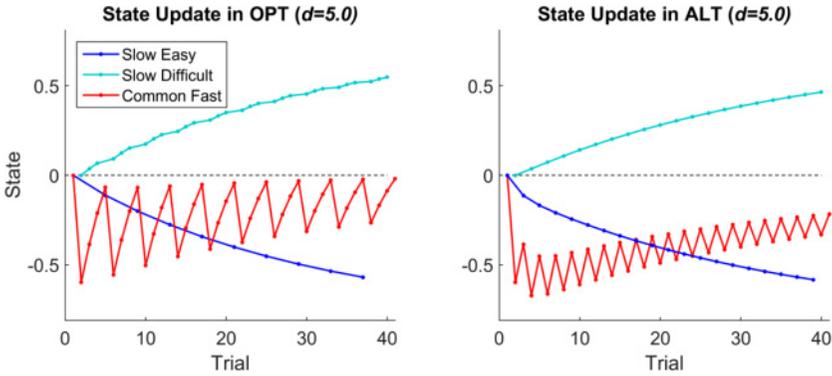


Figure 4: Examples of updating within fast and slow process motor memories in the 1F2S model for task difficulty  $d = 5.0$  and 40 total trials. (Left) Optimal schedule, OPT. (Right) Alternating schedule, ALT. Markers (dots) represent the trials assigned to the two tasks. Note how in OPT, scheduling of the easy task once every four trials (in most cases) keeps the fast process near 0 most of the time. This increases overall performance errors, which results in an increase in the update of the slow process for the difficult task (and thus, overall retention performance at the end of training). In contrast, in the ALT schedule, the fast process has a relatively high (absolute) level of activity throughout.

two tasks compared to the alternating schedule. However, these differences are relatively small, even for a large difference in task difficulty as in this example.

Figure 4 (left panel) illustrates why the optimal schedule generates (in most cases) small blocks of trials for the difficult task separated by one trial for the easy task. Separations between small blocks of the difficult task implement a compromise between assigning more trials to the difficult task and minimizing the block lengths. As a result, there is a minimal update of the fast process throughout the optimal schedule (see the red line in the left panel). This results in increased performance errors and allows greater update in the slow process of the difficult task, while not being too detrimental for the easy task, and thus optimizes final retention for both tasks.

We then systematically studied the difference among optimal, alternating, and blocked schedules. Although the alternating schedule was optimal only up to a certain threshold (as shown in Figure 3), Figure 5 shows that the alternating schedule achieved almost as much final retention as the optimal schedules for a wide range of task difficulties: costs of the alternating schedule are almost same as those of the optimal schedules up to task difficulty  $d = 3.0$ , and 110%, 120%, and 130% of those of the optimal schedules for 20, 40, and 80 total trials, respectively, at task difficulty  $d = 5.0$ . Thus, for a wide range of task difficulties tested for the 1F2S model,

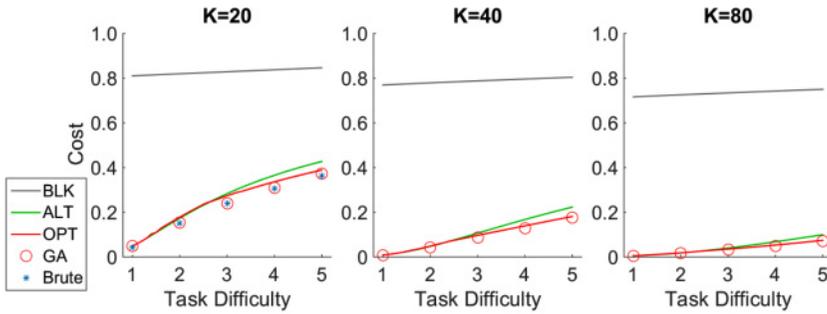


Figure 5: Costs at the end of training for  $K = 20, 40,$  and  $80$  total trials as a function of task difficulty for three different schedules: the optimal schedule (OPT), the alternating schedule (ALT), and the blocked schedule (BLK). Note how the alternating schedules yield similar costs as the optimal schedules in most cases. The five circles in each panel show costs generated by the genetic algorithm. The five stars in the  $K = 20$  panel represent costs calculated from the true global optimum using brute-force search.

the alternating schedule is practically as effective as the optimal schedule in increasing the accumulated learning within the state variables of the slow processes.

To verify the validity of our results overall, we used two approaches. First, we adopted a computationally expensive genetic algorithm (GA) approach to determine optimal schedules “experimentally” and compared the results with these of Pontryagin’s maximum principle method for all schedule lengths. Figure 5 shows that the two methods provide almost identical performance results for a range of task difficulties and different total number of trials (less than 1% difference in performance), although the schedules found by the two methods could differ slightly as parameter  $d$  increased. Second, for the small schedule with 20 total trials, we performed a brute-force search of all possible schedules for the relative task difficulty parameter  $d = 4$  (see appendix B). Such a search shows that the true optimal is very near the optima found by the maximum principles and the GA method for a range of difficulties (see Figure 5, left). In addition, the brute search reveals how close the alternating schedule is to the true optimal, even with large relative difficulty between tasks (see Figure 7). Finally, comparing the true optimal schedule from the brute-force search and the schedule from the maximum principles shows small differences that barely affect long-term retention, as both schedules have very similar costs (see Figure 7).

#### 4 Discussion

Our study made three novel contributions. The first contribution is a theoretical method to optimize multitask motor learning. To determine the

optimal schedules, we have used Pontryagin's maximum principle with constraints on the system states and the command. We have validated the results of this deterministic method in the two-task adaptation paradigm with the results of a stochastic method based on genetic algorithms. Although we determined the schedules in motor adaptation tasks, these optimal schedules can be applied to any types of learning (e.g., motor learning in healthy subjects, motor retraining after stroke, associative learning, declarative learning) for which the state and control variables can be represented in differentiable form (e.g., see equations 2.9 to 2.12). Thus, our method can also be applied to motor rehabilitation to determine the schedule of multiple tasks training, as state-space models of recovery and rehabilitation have been proposed and validated (Casadio & Sanguineti, 2012; Hidaka, Han, Wolf, Winstein, & Schweighofer, 2012; Scheidt & Stoeckmann, 2007). Similarly, at least in theory, this method could also be used to schedule multiple tasks in association experiments, and even in certain cognitive experiments, as long as state-space models are applicable (Smith & Brown, 2003; Kording, Tenenbaum, & Shadmehr, 2007).

The second contribution is that we showed that under conditions of task interference in the fast process, there exists a threshold in relative task difficulty below which the alternating schedule is the true optimal schedule. The third contribution is that for a large range of task difficulties, we found that there is little difference in long-term retention following optimal and alternating schedules. In addition, our results shed light on the well-established contextual interference (CI) effect (Schmidt & Lee, 2005; Shea & Morgan, 1979), in which intermixing tasks during training led to enhanced retention compared to learning tasks sequentially. Our results suggest that the CI can be observed even for tasks of different difficulties. When interference is high, the alternating schedules are clearly superior to the blocked schedules.

What is the mechanism leading to the task difficulty thresholds below which the alternating schedule is the true optimal schedule? In our simulation of the 1F2S model with two opposing tasks, presenting the other task reduces activity in the common fast process. As a result, overall performance gains are reduced, resulting in greater error in the next trial; this in turn results in greater update of the error-driven slow processes. Therefore, when task difficulty differs but stays below threshold, the gain from high switching probability in the alternating schedule is greater than the loss of update in the difficult task resulting from assigning equal numbers of trials to both tasks, hence creating the threshold.

Comparison of the results from Pontryagin's maximum principle and the GA method show very similar final costs for a broad range of relative task difficulties (see Figure 5), which for the special case of 20 total trials are only slightly greater than the true optimal cost found by the brute-force search. As long as tasks alternate with small blocks of the difficult task intercalated between single trials of the easy task, retention is very high and difference

in cost with the optimal schedule minimal. Note that besides brute-force search for a smaller schedule, we have used two optimization techniques: the deterministic Pontryagin's maximum principle and a stochastic GA method. A third possible method, dynamic programming, could also be used to determine optimal schedules. We leave for future work the exploration of dynamic programming to determine optimal schedules in motor adaptation.

Our study has a number of limitations that could also be addressed in future work. First, because this is a simulation study, our results depend on our choice of models. To increase the validity of our results, we have performed several sensitivity analyses whereby we have varied the relative task difficulty, the number of trials, and learning and forgetting rates. Overall, the results show the existence of a difficulty ratio threshold below which the alternating schedule is nearly as effective in increasing long-term retention as the optimal schedule.

A second limitation is that we have studied optimal schedules for only two tasks in a single session, in which adaptation occurs at least at two different timescales. Studies of memory consolidation over multiple days show that additional processes with much longer timescales may play important roles during long-term motor learning (Criscimagna-Hemminger & Shadmehr, 2008). We leave scheduling of multiple tasks and scheduling over multiple sessions for future work. A third limitation of our study is that we have studied optimal scheduling only for multiple motor adaptation tasks with no generalization between tasks. Generalization effects can be implemented by adding parameters to the slow process (Tanaka, Krakauer, & Sejnowski, 2012), and optimal schedules could be determined with this new model.

Finally, in a practical application of our study, determination of the optimal schedule would largely depend on accurate parameter estimation, including learning rates, forgetting rates, and degree of interference between tasks. In particular, we expect that accurate parameter estimation would be crucial when determining schedules for tasks of vastly different difficulties. Extrapolating our finding suggests that for this case, the optimal schedule could be truly superior to the alternating schedule. However, in most practical applications with tasks of similar difficulty, our simulation results suggest that the alternating schedule may be a near-optimal choice for enhancing long-term retention of motor learning.

Our study makes three counterintuitive yet practical predictions for a large range of tasks. First, therapists, coaches, and teachers should design the training schedule to include interfering tasks. Second, tasks should be scheduled alternatively or pseudorandomly; the details of the schedules do not matter to a great extent as long as switching occurs frequently and more or less evenly. Finally, if only a single training session is available, trainers can ignore task difficulty (unless extremely different) and assign a similar number of trials for all tasks according to an alternating schedule.

## Appendix A: Model Parameter Sensitivity Analysis

---

One of the main findings on the optimal schedule with the 1F2S model was the existence of a threshold of task difficulty below which the alternating schedule is optimal (see Figure 3). This result was obtained with a specific set of parameters in a visuomotor experiment. However, depending on the specifics of the experimental protocol of the task, or depending on the adaptation tasks (e.g., visuomotor rotation, saccadic adaptation), parameters will vary. Here we show that the existence of the threshold is a general phenomenon across a wide range of learning and retention parameters. In the main text, we ran simulations based on the estimated set of parameters from a visuomotor rotation task:

$$\text{Easy task: } a^f = 0.965, a^s = 0.993, b^f = 0.597, b^s = 0.114$$

$$\text{Difficult task: } a^f = 0.965, a^s = 0.993, b^f = 0.597/d, b^s = 0.114/d,$$

where  $d > 1$  defines a relative task difficulty of the more difficult task. The threshold for this set of  $a$  and  $b$  values was  $d = 2.3$ . In order to simplify the sensitivity analysis, we introduced two variables that define relative values of learning and retention parameters between fast and slow processes:  $b^f/b^s$  (the ratio of fast learning gain to slow learning gain) and  $\log_{10}(\tau^s/\tau^f)$  (the logarithmic ratio of a slow time constant compared to a fast time constant), where the time constants are defined from a retention parameter,  $\tau = \frac{1}{1-a}$ . In simulations, we fixed  $\tau^f$  at 28.57 and  $b^s$  at 0.114.

Figure 6 shows that the threshold is greater than 1 for a large range model of parameters. The threshold decreases as the slow time constant increases and increases as  $b^f/b^s$  increases. Summarizing, the result shows that the alternating schedule is optimal unless one of the tasks is a lot more difficult than the other (up to thresholds), and this holds true for a wide range of parameters, with the thresholds depending on the ratio of fast and slow learning gains.

## Appendix B: Brute-Force Search

---

Our main optimization algorithm produced schedules that could have become trapped in a local minimum. Here, our goal is to examine how close the final cost obtained from our optimization algorithm is to the true global minimum cost. We performed a brute-force search of the optimal schedule for  $K = 20$  total trials (for 40 and 80 trials, the brute-force search becomes computationally prohibitive). We generated all possible  $2^{20} \simeq 10^6$  schedules of  $K = 20$  and calculated the final cost at the end of training for each of these schedules. We then sorted these  $2^{20}$  costs from the smallest to the largest. We defined the rank of each schedule as the order in this sorted list (rank 1 as the smallest).

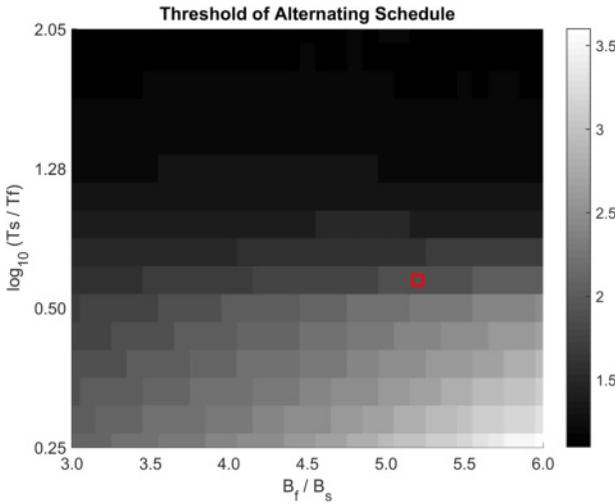


Figure 6: Threshold of task difficulty below which the alternating schedule is optimal. The gray scale shows the threshold. The horizontal axis represents the ratio of fast and slow gains, and the vertical axis represents the logarithmic ratio of slow and fast time constants. For all simulations,  $\tau^f$  was fixed at 28.57 and  $b^s$  was fixed at 0.114, the parameters of the original simulation in the main text. The small red square indicates the parameters closest to the original parameter set used in the main text.

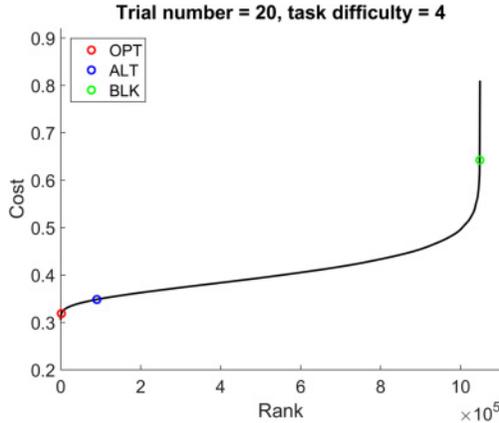


Figure 7: Cost from all possible schedules with 20 total trials for the 1FnS model and relative difficulty between tasks  $d = 4$ , sorted in ascending order (rank). Note that the optimal schedule OPT has a cost that is very close to the true minimum, the alternating schedule ALT has a cost and rank that is close to the true optimal, and the blocked schedule BLK has a very high rank and cost.

Figure 7 shows the costs of all possible schedules in ascending order (with relative task difficulty  $d = 4$ ). Circles represent the corresponding ranks and costs of our algorithm's optimal schedule, alternating schedule, and blocked schedules. Ranks of these schedules were 0.13%, 8.62%, and 99.97% of the  $2^{20}$  schedules, respectively. Corresponding costs were 0.3186, 0.3485, and 0.6422, respectively, while the true minimum cost was 0.3073. Thus, both ranks and cost show that the schedule determined by the optimization algorithm is very close to the true optimal. In addition, the alternating schedule has both very low cost and rank.

## Acknowledgments

---

This work was funded by grants NSF BCS 1031899, MC-IIF 299687, and NIH HD053727.

## References

---

- Bryson, A. E., & Ho, Y-C. (1969). *Applied optimal control: Optimization, estimation and control*. Waltham, MA: Blaisdell Publishing.
- Casadio, M., & Sanguineti, V. (2012). Learning, retention, and slacking: A model of the dynamics of recovery in robot therapy. *IEEE Trans. Neural Syst. Rehabil. Eng.*, *20*(3), 286–296. doi:10.1109/TNSRE.2012.2190827
- Choi, Y., Qi, F., Gordon, J., & Schweighofer, N. (2008). Performance-based adaptive schedules enhance motor learning. *J. Mot. Behav.*, *40*(4), 273–280. doi:10.3200/JMBR.40.4.273-280
- Criscimagna-Hemminger, S. E., & Shadmehr, R. (2008). Consolidation patterns of human motor memory. *J. Neurosci.*, *28*(39), 9610–9618. doi:28/39/9610
- Hidaka, Y., Han, C. E., Wolf, S. L., Winstein, C. J., & Schweighofer, N. (2012). Use it and improve it or lose it: Interactions between arm function and use in humans post-stroke. *PLoS Comput. Biol.*, *8*(2), e1002343. doi:10.1371/journal.pcbi.1002343
- Hirashima, M., & Nozaki, D. (2012). Distinct motor plans form and retrieve distinct motor memories for physically identical movements. *Curr. Biol.*, *22*(5), 432–436. doi:10.1016/j.cub.2012.01.042
- Huang, V. S., Shadmehr, R., & Diedrichsen, J. (2008). Active learning: Learning a motor skill without a coach. *Journal of Neurophysiology*, *100*, 879–887.
- Imamizu, H., Sugimoto, N., Osu, R., Tsutsui, K., Sugiyama, K., Wada, Y., & Kawato, M. (2007). Explicit contextual information selectively contributes to predictive switching of internal models. *Exp. Brain. Res.*, *181*(3), 395–408. doi:10.1007/s00221-007-0940-1
- Joiner, W. M., & Smith, M. A. (2008). Long-term retention explained by a model of short-term learning in the adaptive control of reaching. *J. Neurophysiol.*, *100*(5), 2948–2955. doi:90706.2008
- Judkins, T., & Scheidt, R. A. (2014). Visuo-proprioceptive interactions during adaptation of the human reach. *J. Neurophysiol.*, *111*(4), 868–887. doi:10.1152/jn.00314.2012
- Kirk, D. E. (1992). *Optimal control theory: An introduction*. Englewood Cliffs, NJ: Prentice Hall.

- Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.*, *10*(6), 779–786. doi:nn1901
- Lee, J. Y. (2011). *Modeling motor memory to enhance multiple task learning*. Doctoral dissertation, University of Southern California.
- Lee, J. Y., & Schweighofer, N. (2009). Dual adaptation supports a parallel architecture of motor memory. *J. Neurosci.*, *29*(33), 10396–10404.
- Liu, Y. T., Mayer-Kress, G., & Newell, K. M. (2003). Beyond curve fitting: A dynamical systems account of exponential learning in a discrete timing task. *J. Mot. Behav.*, *35*(2), 197–207.
- Osu, R., Hirai, S., Yoshioka, T., & Kawato, M. (2004). Random presentation enables subjects to adapt to two opposing forces on the hand. *Nat. Neurosci.*, *7*(2), 111–112.
- Scheidt, R. A., Dingwell, J. B., & Mussa-Ivaldi, F. A. (2001). Learning to move amid uncertainty. *J. Neurophysiol.*, *86*(2), 971–985.
- Scheidt, R. A., & Stoekmann, T. (2007). Reach adaptation and final position control amid environmental uncertainty after stroke. *J. Neurophysiol.*, *97*(4), 2824–2836.
- Schmidt, R. A., & Lee, T. D. (2005). *Motor control and learning: A behavioral emphasis* (4th ed. ). Champaign, IL: Human Kinetics.
- Schweighofer, N., Lee, J. Y., Goh, H. T., Choi, Y., Kim, S. S., Stewart, J. C., . . . Winstein, C. J. (2011). Mechanisms of the contextual interference effect in individuals poststroke. *J. Neurophysiol.*, *106*(5), 2632–2641. doi:10.1152/jn.00399.2011
- Shea, J. B., & Morgan, R. L. (1979). Contextual interference effects on the acquisition, retention, and transfer of motor skills. *J. Exp. Psychol. (Hum. Learn.)*, *3*, 179–187.
- Shelhamer, M., Aboukhalil, A., & Clendaniel, R. (2005). Context-specific adaptation of saccade gain is enhanced with rest intervals between changes in context state. *Ann. N.Y. Acad. Sci.*, *1039*, 166–175. doi:1039/1/166
- Simon, D. A., Cullen, J. D., & Lee, T. D. (2002). Win-shift/lose-stay: Contingent switching as an alternative to random practice? *Journal of Sport and Exercise Psychology*, *24*, S116.
- Smith, A. C., & Brown, E. N. (2003). Estimating a state-space model from point process observations. *Neural Comput.*, *15*(5), 965–991. doi:10.1162/089976603765202622
- Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.*, *4*(6), e179.
- Tanaka, H., Krakauer, J. W., & Sejnowski, T. J. (2012). Generalization and multi-rate models of motor adaptation. *Neural Comput.*, *24*(4), 939–966. doi:10.1162/NECO\_a\_00262
- Todorov, E. (2007). Optimal control theory. In K. Doya (Ed.), *Bayesian brain: Probabilistic approaches to neural coding* (pp. 269–298). Cambridge, MA: MIT Press.
- Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Netw.*, *11*(7–8), 1317–1329.

Copyright of Neural Computation is the property of MIT Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.