

Computational nature of human adaptive control during learning of reaching movements in force fields

Nikhil Bhushan, Reza Shadmehr

Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205, USA

Received: 01 October 1998 / Accepted in revised form: 26 January 1999

Abstract. Learning to make reaching movements in force fields was used as a paradigm to explore the system architecture of the biological adaptive controller. We compared the performance of a number of candidate control systems that acted on a model of the neuromuscular system of the human arm and asked how well the dynamics of the candidate system compared with the movement characteristics of 16 subjects. We found that control via a supra-spinal system that utilized an adaptive inverse model resulted in dynamics that were similar to that observed in our subjects, but lacked essential characteristics. These characteristics pointed to a different architecture where descending commands were influenced by an adaptive forward model. However, we found that control via a forward model alone also resulted in dynamics that did not match the behavior of the human arm. We considered a third control architecture where a forward model was used in conjunction with an inverse model and found that the resulting dynamics were remarkably similar to that observed in the experimental data. The essential property of this control architecture was that it predicted a complex pattern of near-discontinuities in hand trajectory in the novel force field. A nearly identical pattern was observed in our subjects, suggesting that generation of descending motor commands was likely through a control system architecture that included both adaptive forward and inverse models. We found that as subjects learned to make reaching movements, adaptation rates for the forward and inverse models could be independently estimated and the resulting changes in performance of subjects from movement to movement could be accurately accounted for. Results suggested that the adaptation of the forward model played a dominant role in the motor learning of subjects. After a period of

consolidation, the rates of adaptation in the internal models were significantly larger than those observed before the memory had consolidated. This suggested that consolidation of motor memory coincided with freeing of certain computational resources for subsequent learning.

1 Introduction

The electric fish relies on its ability to sense the weak electrical fields generated by other animals to identify prey and predator. Unfortunately, this field is not only influenced by the motion of other animals, but also by the motion of the electric fish itself. Therefore, the animal needs to be able to take into account self-generated changes to the surrounding electric field and subtract it from the measured field in order to estimate the external world. The problem might be solved if the animal could send signals from the motor centers to the sensory areas and provide a negative image of the predicted sensory changes that should be detected as a consequence of the just programmed motor output. Such a negative image would then add with the actual sensory input and result in an estimate of the external world (Sperry 1950). In fact, such signals have been detected in the electric fish (Bell 1981), shown to be adaptable (Montgomery and Bodznick 1994), and coded in a cerebellum-like structure of the animal (Bell et al. 1997).

The computation involved in predicting sensory consequences of a motor command, as exemplified by the electric fish, is termed a forward model¹ (Jordan and Rumelhart 1992). There are now a number of studies that have suggested that a forward model may also be used by the human central nervous system (CNS) to estimate sensory consequences of motor actions (Wolpert et al. 1993, 1995; Flanagan and Wing 1997). For example, during precision grasping of a small object, the grip forces change in concert with load forces that act to

Correspondence to: R. Shadmehr,
Department of Biomedical Engineering, Johns Hopkins School
of Medicine, 720 Rutland Ave/419 Traylor, Baltimore,
MD 21205-2195, USA
(e-mail: reza@bme.jhu.edu)
Tel.: +1-410-614-2458, Fax: +1-410-614-9890)

¹ In the control literature, this type of model is called an observer.

move the object (Johansson and Westling 1984; Flanagan and Wing 1997), even if the load forces are generated by one hand and the grip forces are generated by the other (Blakemore et al. 1998). Due to the delay inherent in the sensory-motor loop (Johansson and Westling 1984; Blakemore et al. 1998), a close synchrony between grip and load forces is possible only if the brain could predict the motion-dependent nature of the load forces from a forward model of the dynamics of the limb and the load.

If a forward model of the dynamics of the limb is available to the CNS, then an interesting use of such a model might be to provide a means by which the CNS could gauge the consequences of the just-programmed motor commands without having to wait for the arrival of the corresponding sensory signals (Miall et al. 1993; Darlot et al. 1996; Miall and Wolpert 1996; Ronco 1998). Gauging, in this case, means a comparison of the predicted sensory consequences of the motor commands with the desired behavior. This comparison would allow for estimation of an error signal, i.e., a measure of how far the arm is predicted to be from a desired state. Motor commands can then be modified to reduce this predicted error in advance of the corrections that would be possible if the brain had to wait for the transmission delayed afferent information from the moving arm, delays which may exceed 100 ms (Lee and Tatton 1975; Cordo et al. 1994).

In the control literature, such ideas are found, for example, in Luenberger observers and Smith predictors, which are forward models that are used for control of linear systems with time delays (Astrom and Wittenmark 1984), and in metric observers for control of nonlinear systems (Lohmiller and Slotine 1996). Whether the dynamics of the time-delayed system is linear or nonlinear, the idea remains the same: use the forward model to feed back the predicted response of the remote system immediately, as well as the error in this prediction when the real response becomes available via the transmission channel. Obviously, stability of these systems will strongly rely on the accuracy of the forward model and the ability to cancel the response from the remote system. If the human CNS uses a forward model to control the arm, then by manipulating the dynamics of the arm one can introduce a condition where the hypothetical forward model would be grossly inappropriate for the task. Here, we take this approach and ask how the resulting behavior of the human arm compares with that which would be expected if a forward model was being used for programming the descending motor commands.

While the use of a forward model may be particularly relevant for feedback control of time-delayed systems, an alternate approach is through the use of an inverse model (Atkeson 1989; Kawato 1989; Shadmehr 1990; Gomi and Kawato 1992; Katayama and Kawato 1993; Schweighofer et al. 1998). Whereas a forward model predicts the sensory consequences of motor commands, an inverse model predicts motor commands that are appropriate for a desired behavior. Inverse models are generally not considered for control of time-delayed systems because they would seem to exclude the ability of the controller to respond to errors, resulting in an open

loop controller. However, as a recent approach has illustrated (Niemeyer and Slotine 1991), if a local feedback controller stationed at the remote system is available, then reacting to the delayed error information received by the up-stream controller may be possible in a way that does not result in instability. In the case of the human arm, such a local feedback controller is thought to be present in the form of spring-like muscles and spinal reflex loops (Massaquoi and Slotine 1996). Therefore, at least in theory, adaptive control of the arm in the face of a novel mechanical load may take place through learning of a forward, an inverse, or perhaps a mixture of both types of controllers. Does the data on the behavior of the human arm when coupled to a novel load allow us to differentiate between these possibilities?

In the current report we examine the behavior of the human arm as the hand is coupled to a novel dynamical system during generation of reaching movements. We consider a reasonably realistic model of the arm's inertial and muscle dynamics, local reflexes, and delays in the various communication channels. We initially consider control of the arm via an adaptive inverse model. We find that, while the simulation results with the inverse model resemble the actual behavior of our subjects, there are also important features in hand trajectories that cannot be explained despite systematic variations in the parameters of the model. These features suggest that the arm is being controlled by factors other than those accounted for in the adaptive inverse model controller. We next consider performance of the system under the control of an adaptive forward model and find certain limitations to this approach. We demonstrate that this controller also fails to precisely account for the characteristics of the biological controller. These simulations give insights into characteristics of these two approaches and suggest a third approach, one where there is an interaction between the inverse and forward models during the process of control. We demonstrate that the behavior of the human arm when coupled with a novel mechanical system is very similar to the behavior that results when the controller relies on a combination of forward and inverse models. It appears that the biological controller relies on state estimates from a forward model which, in turn, provide an error signal that through an inverse model modulates descending motor commands.

If the biological controller is composed of a combination of forward and inverse models, then it is important to ask whether the rates of learning of these two adaptive models are the same. Can we estimate the adaptation rate of each model from the actual performance of the human subjects? An intriguing idea put forth by Jordan and Rumelhart (1992) and others (Wada and Kawato 1993; Miall and Wolpert 1996) is that if a forward model is available, it can effectively serve as a model of the controlled system. Using the forward model, the brain may simulate dynamics of the controlled system during an "off-line" period and teach itself an inverse model. The learning that may take place during the off-line period may result in a fundamentally

different control system than was apparent during the initial practice of the task. We are intrigued by this idea because it may be related to the functional (Brashers-Krug et al. 1996; Shadmehr and Brashers-Krug 1997) and neural changes (Shadmehr and Holcomb 1997, 1999) that we have observed during consolidation of motor memories. During the hours after completion of practice, the perturbation response characteristics of the human adaptive controller appears to gradually change, becoming remarkably more stable. Here, we explore this change in terms of the elements of a control system that includes adaptive forward and inverse models. We ask how the perturbation response of the arm should change as a function of adaptation rates in the forward and inverse models. We conclude with an estimate of the actual rates of adaptation for each model from data of 16 subjects that were examined during the period of consolidation.

2 Plant dynamics

Here, we consider the merits of various theoretical control mechanisms that may act on a time-delayed system resembling the human arm. The psychophysical data that will be used for this comparison were gathered from subjects that made reaching movements while holding a robotic manipulandum, (Fig. 1). The movements were in the horizontal plane.

The inverse dynamics of the robot (while not being held by the subject) is described by:

$$T_r = D_r(\Phi)\ddot{\Phi} + C_r(\Phi, \dot{\Phi})\dot{\Phi}$$

$$D_r = \begin{bmatrix} k_{r1} & k_{r3} \cos(\phi_2 - \phi_1) \\ k_{r3} \cos(\phi_2 - \phi_1) & k_{r2} \end{bmatrix}$$

$$C_r = \begin{bmatrix} 0 & -k_{r3} \sin(\phi_2 - \phi_1)\dot{\phi}_2 \\ k_{r3} \sin(\phi_2 - \phi_1)\dot{\phi}_1 & 0 \end{bmatrix}$$

where T_r is joint torques on the robot due to its motion and parameters k_r are constants that depend on link lengths and mass distributions of the links that make up




Fig. 1. Schematic of the experimental set up. The figure notes the position of the robot's base and end effector (handle) with respect to a coordinate system centered on the shoulder of the subject. The positions are the same as those used for actual experiments and simulations

the robot. Similarly, the inverse dynamics of the human arm (not holding the robot), can be written as:

$$T_s = D_s(\Theta)\ddot{\Theta} + C_s(\Theta, \dot{\Theta})\dot{\Theta},$$

where matrices D_s and C_s are similar to that of the robot except that their parameters k_s depend on mass distribution and lengths of the human arm. For the robot, $k_r = [0.3189, 0.0938, 0.1262]$ kg · m², and link lengths $r_1 = 0.46$ m, $r_2 = 0.44$ m. For the human arm, $k_s = [0.265, 0.052, 0.0844]$ kg · m², and link lengths $l_1 = 0.33$ m, $l_2 = 0.32$ m (Jordan et al. 1994).

At the interaction port at the hand, where a force transducer is housed, the two systems are coupled. The interaction force acting on the hand of the subject as a result of the motion of the robot is $F_x = -(J_r^T)^{-1}T_r$. If we express the dynamics of the robot represented by Φ in terms of kinematics of the human arm Θ , and include the possibility that the robot may, in addition to its passive dynamics, impose an active force field $F(x, \dot{x})$ on the subject's hand, then the overall inverse dynamics of the system can be expressed as torques acting on the subject's arm:

$$T = A(\Theta)\ddot{\Theta} + B(\Theta, \dot{\Theta})\dot{\Theta} + J_s^T F(x, \dot{x}),$$

$$A = D_s + J_s^T (J_r^T)^{-1} D_r J_r^{-1} J_s,$$

$$B = C_s + J_s^T (J_r^T)^{-1} C_r J_r^{-1} J_s + J_s^T (J_r^T)^{-1} D_r J_r^{-1} (\dot{J}_s - \dot{J}_r J_r^{-1} J_s),$$

where $J_s = d\Theta/dx$, and $J_r = d\Phi/dx$.

We next modeled the dynamics of some of the muscles attached to the arm. The major muscles acting on the human arm during reaching movements in the current configuration include: anterior deltoid, posterior deltoid, brachialis/brachioradialis, triceps (short and long head) and biceps. To model these muscles, we considered a simplification that consisted of three muscle pairs acting around the shoulder, elbow and both joints, respectively. Since the extent of the reaching movements were small (10 cm), we assumed that the moment arms of the shoulder, elbow and double-joint muscles were constant with respect to the absolute shoulder angle, relative elbow angle, and the absolute elbow angle, respectively. The two muscles in each pair had a flexor-extensor configuration and were assumed to be identical to each other. The values for the moment arm (Murray et al. 1995) and maximum force (Karniel and Inbar 1997) of each muscle were estimated as follows: we assumed a moment arm of 5 cm for the anterior and posterior deltoids and 3 cm for all other muscles. F_{\max} was 800 N for the anterior and posterior deltoids, 700 N for brachialis and the short head of the triceps, and 1000 N for the biceps and long head of the triceps and are provided in the appendix.

We represented the force response of a single muscle fiber to an electrical impulse as the product of a normalized force response $h_i(s)$ and a constant c , where $\int_{t=0}^{\infty} h_i(t) dt = 1$. The net force $F(t)$, produced by a muscle composed of n_m fibers, each having an activation-force impulse response $c_n h_i$ and receiving electrical impulse activity A_n , is:

$$\begin{aligned}
F(t) &= \sum_{n=1}^{n_m} h_i * [c_n A_n](t) = \sum_{n=1}^{n_m} \int_0^t h_i(\rho) c_n A_n(t - \rho) d\rho \\
&= h_i * \sum_{n=1}^{n_m} [c_n A_n](t) ,
\end{aligned}$$

where $*$ denotes the convolution operation. The maximum force F_{\max} that can be generated by the muscle is equal to the maximum value for $\sum_{n=1}^{n_m} [c_n A_n]$, because $\int_{t=0}^{\infty} h_i(t) = 1$. Therefore, if we define a mean normalized electrical activity $R = \frac{\sum_{n=1}^{n_m} [c_n A_n]}{F_{\max}}$, then we can write,

$$F(t) = F_{\max} [h_i * R](t) = F_{\max} N(t) .$$

$N(t)$ is the filtered normalized electrical activity to the muscle equal to $[h_i * R]$ and having a value between 0 and 1. N directly controls the force produced by the whole muscle and, hence, will be used as the variable to represent the central motor command to the muscle. The activation/force impulse response function of the muscle was modeled as:

$$h_i(t) = \frac{t^3 e^{\left(\frac{-t}{0.01}\right)}}{6 \times 10^{-8}} . \quad (2)$$

The results above are derived for an isometric muscle and hence F is the isometric force produced by the muscle for a given neural activation. The force produced by an active muscle also depends on muscle length and velocity. The force/length relation was modeled as an active elastic element where stiffness changed proportionately with neural activation (Shadmehr and Arbib 1992). If we denote the length of the muscle as x_m and the operating length of our isometric muscle as x_{m0} , then we can represent the length-modulated force output F_a as,

$$F_a = F + FC_m(x_m - x_{m0}) .$$

The value of C_m of each muscle was derived from measured joint stiffness of the arm in intact (Gomi and Kawato 1996) and deafferented subjects (Sanes and Shadmehr 1995). We had found that in patients with large fiber sensory neuropathy, stiffness of the arm was approximately 50% of the value that we had recorded in the normal population. We therefore assumed that the intrinsic stiffness of the muscles was 50% of that measured by Gomi and Kawato (1996), with remaining stiffness contributed via a stretch reflex loop (to be described below).

We modeled the force-velocity relation of each muscle by a Hill parameterized model (Krylow et al. 1995):

$$\begin{aligned}
F_t &= \frac{bF_a + a\dot{x}_m}{b - \dot{x}_m} \quad \dot{x}_m \leq 0 \quad (\text{shortening}) \\
F_t &= \frac{b'F_a - (a' + 2F_a)\dot{x}_m}{b' - \dot{x}_m} \quad \dot{x}_m > 0 \quad (\text{lengthening}) \\
a' &= -0.4F_a \\
b' &= -b \frac{a' + F_a}{a + F_a}
\end{aligned}$$

F_t is the velocity-modulated force of the muscle given a length-modulated force F_a and muscle velocity \dot{x}_m . a and b are constants that govern viscosity of the muscle. The value of b' is derived to ensure continuity at $\dot{x}_m = 0$. For estimating a and b for each muscle, we assumed $(bF_{\max})/(ax_{m0}) = 10$ (Zajac 1989). This resulted in a muscle viscosity that was 15–20% of the muscle stiffness.

The overall input/output relation for the muscle, relating the filtered activation N with the force F_t , can be represented by a function f_M :

$$\begin{aligned}
F_t &= F + K_m(F, x_m)x_m + B_m(F, x_m, \dot{x}_m)\dot{x}_m \\
&= F_{\max}N + K_m(N, x_m)x_m + B_m(N, x_m, \dot{x}_m)\dot{x}_m \\
&= f_M(N, x_m, \dot{x}_m) , \quad (3)
\end{aligned}$$

where, K_m represents the nonlinear stiffness due to the force/length relationship and B_m the non-linear viscosity due to the force/velocity relationship for the muscle.

We modeled the spinal component of the stretch reflex loop as a linear approximation of the non-linear model proposed by Gielen and Houk (1987). The activation to the muscle through the spinal reflex pathway, N_r , was based on reciprocal inhibition of a muscle pair and was the solution of the following simultaneous equations:

$$\begin{aligned}
N_{r1} - N_{r2} &= K_s(x_m - x_{ms}) + B_s(\dot{x}_m - \dot{x}_{ms}) \\
N_{rc} &= \sqrt{N_{r1}N_{r2}}
\end{aligned}$$

where, x_{ms} , \dot{x}_{ms} were the set-point muscle length and velocity, K_s was the reflex stiffness, B_s was the reflex viscosity, and N_{rc} was the reciprocal inhibition constant. The ratio of B_s to K_s has been suggested to be approximately 0.1 (Gielen and Houk, 1987). The reflex pathway was modeled with a time delay of 0.03 s (Rothwell, 1990). The stiffness of this pathway, K_s , was set at 50% of the muscle stiffness derived from measures of Gomi and Kawato (1996). The muscle-load-spinal feedback system is summarized in Fig. 2.

The three pairs of muscles acting on the simulated arm are redundant for the purpose of joint torque generation. In order to assign a neural activation so that a desired torque could be generated, we made the further assumption that elbow torque was distributed equally between the single-joint and double-joint muscles and that, similarly, the shoulder torque was produced by an equal contribution from the single-joint and double-joint muscles.

To summarize, the essential components of the system are as follows:

1. The inertial dynamics of the arm are described via a two-joint planar model that interacts with a two-joint robotic manipulandum.
2. Muscles produce passive force via a zero-delay stiffness and viscosity mechanical response, and active force in response to neural activation via a transformation that is characterized by the impulse response of Eq. (2).
3. The spinal stretch reflex provides added stiffness and viscosity but at a time delay of 30 ms.

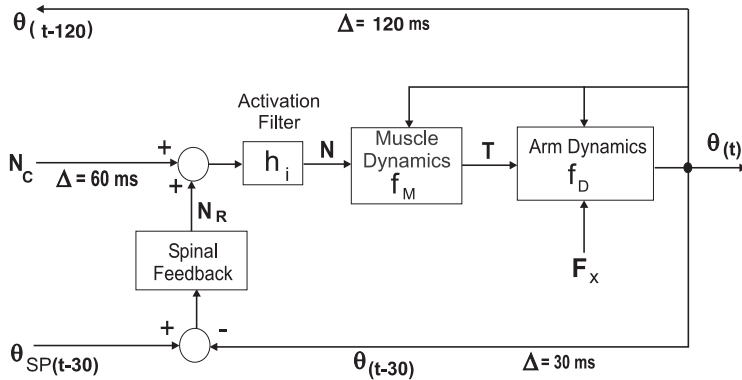


Fig. 2. Block diagram of the muscle, load, spinal feedback system. N_c is the descending neural command, N_r is the contribution of the spinal stretch reflex, modeled as a linear feedback controller, and θ_{sp} is the set point for the reflex loop in joint coordinates. Δ indicates delays in the communication channels. h_i is the filter transforming neural activation into activation dynamics of the muscle (Eq. 2)

4. State of the arm, i.e., its position and velocity, is made available to the brain after a delay of 120 ms.
5. The motor commands have a delay of 60 ms from the brain to the generation of a measurable force in the muscle.

3 Forward and inverse models for control

Here, we consider a number of controllers and demonstrate their performance on the system described above. To simulate performance of the system under various controllers, we assumed that the desired trajectory for the arm is minimum jerk (Flash and Hogan 1985) with a movement time of 0.5 s to targets placed at 10 cm in eight equally spaced directions about a starting point. We further considered two conditions. First, movements in a null field, i.e., the subject's arm is unloaded. Second, movements in a curl force field $F = B\dot{x}$, with the field described by $B_1 = \{\{0, 13\}, \{-13, 0\}\} \text{N} \cdot \text{s}^{-1} \cdot \text{m}^{-1}$ or $B_2 = -B_1$.

3.1 Feedforward control via inverse models

We initially consider a simple approach where assignments of neural activations are arrived at via an inverse muscle model so that at a desired position and velocity, the net torque acting on each joint is zero. In effect, we assign neural activations so that the equilibrium position of the system moves along the desired trajectory. The desired trajectory in hand coordinates was transformed to muscle space and the set point to the spinal feedback system was assigned to this desired trajectory. The schematic for this system is shown in Fig. 3, as is the performance of the system in the null field and force field B_1 . While the system is stable and reaches the desired target, its performance only approximates the desired behavior (a minimum jerk trajectory to the target) because the controller does not attempt to compensate for the dynamics of the arm or the field. The fairly good performance of the system when it is not coupled to a force field (Fig. 3A) demonstrates that the CNS can do quite well in controlling the arm in unloaded situations without explicit compensation for the inertial dynamics of the limb. This is in agreement with a previous

simulation result (Gribble et al. 1998). It is clear, however, that precise tracking of the desired trajectory requires compensation for the inertial dynamics of the arm and any loads that may be attached to it.

The performance of the feedforward approach can be improved if, in addition to an inverse muscle model, a model was available to compensate for the inertial dynamics of the limb. In Fig. 4, the schematics of such a controller is provided. In this approach, the inverse model of the inertial dynamics of the limb transforms the desired trajectory (x_{md}, \dot{x}_{md}) into a desired joint torque T_d , which is then converted to the desired forces in the individual muscles F_{ld} , then using the inverse muscle model, provides the motor commands N ,

$$N = \hat{f}_M^{-1}(F_{ld}, x_{md}, \dot{x}_{md}) .$$

This type of control system has been applied by a number of investigators to the problem of generating reaching movements (Atkeson 1989; Shadmehr 1990; Katayama and Kawato 1993; Stroeve 1997). The performance of this system in the null field and in field B_1 is shown in Fig. 4. By compensating for inertial dynamics of the arm, the hand almost precisely follows the desired trajectory in the null field. The reason for the small error is that the dynamics associated with the neural activation filter h_i (for each muscle) cannot be exactly inverted and is approximated here. The corrective response of the system to unmodeled dynamics of the force field, as shown in the bottom right chart in Fig. 4, is due to the spring-like mechanical properties of the muscles and the feedback dynamics of the reflex loop.

3.2 Comparison with control of the human arm

How does the performance of the controller (Fig. 4) compare with the human arm? We will show that while there are significant similarities between the two, there are crucial differences which suggest that it is unlikely that the human arm is controlled only via this feedforward system. We trained 16 subjects in the null field (800 targets), then in field B_1 (576 targets). Subjects were then presented with field B_2 (384 targets). We had previously observed that, after about 200 targets, performance reached a plateau and movements converged to the

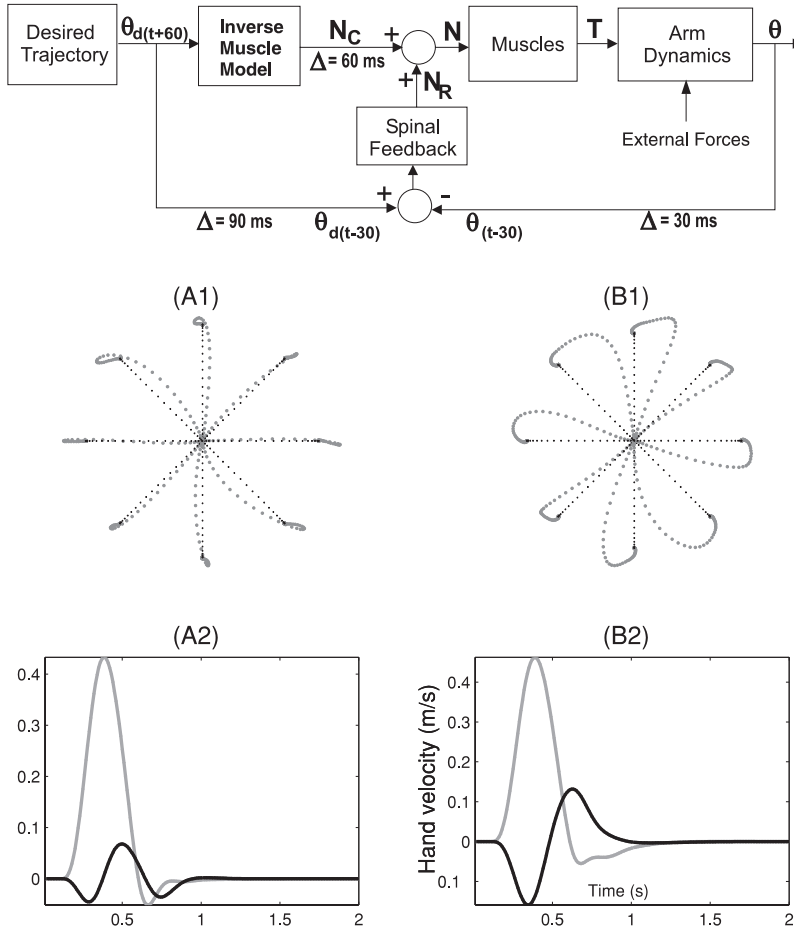


Fig. 3A,B. Block diagram illustrating a feedforward controller that utilizes an inverse muscle model for assignment of neural activations to the muscles, without accounting for dynamics of the limb. **A1** and **B1** show simulation results of hand trajectories in a null field and in force field B_1 , respectively. All movements are center-out. The *dotted straight line* is an ideal minimum-jerk motion while the other paths are the outputs of the controller. **A2** and **B2** show velocity of the hand for a movement to the bottom-most target (at -90°). The *gray line* is velocity in the direction parallel to the direction of target (i.e., along the y -axis of Fig. 1), and the *black line* is the velocity in a direction perpendicular to that of target (i.e., along the x -axis of Fig. 1)

trajectories that were observed before introduction of the force field (Shadmehr and Mussa-Ivaldi 1994). Movements were highly correlated to a minimum jerk trajectory (0.975 ± 0.004 , mean correlation coefficient \pm SD). We were interested in quantifying the performance of the subjects after adaptation to field B_1 in field B_2 , i.e., the response to large changes in system dynamics.

In Fig. 5 we plotted the performance of a typical subject after she had extensively trained in field B_1 , but was suddenly presented with field B_2 . Initially, let us consider a single movement downward. This movement appears segmented, i.e., there are points where there are sudden changes in both the derivative of hand speed and the direction of hand velocity. We identified the segmentation points, S_i , by finding places where a local minimum in the hand speed profile coincided with a maximum in the derivative of direction of velocity signal.

To compare the subject's data with that of our controller (Fig. 4), we assume that, after extensive training in field B_1 , performance of the subject has improved due to adaptation of the inverse model. When the inverse model is of field B_1 and the system is coupled to field B_2 , the resulting simulated hand trajectory is as shown in Fig. 5A. While both the subject (Fig. 5B) and the controller show stability and arrive at the target, the behavior of the two cases is very different about the segmentation points. Some of the parameters that

identify this behavior are labeled in Fig. 5B. These parameters are: λ_i (angles of the hand's trajectory about a segmentation point), d_i (distance between segmentation point $i-1$ and i), t_i (time between segmentation point $i-1$ and i), $|v|_i$ (hand speed at segmentation point i), and N_s (number of segmentation points in the entire trajectory).

In Fig. 6, we have quantified these parameters in all 16 of our subjects for their first downward movement in field B_2 . We chose this direction of movement for illustration of results because, due to the shape of the arm's stiffness (Mussa-Ivaldi et al. 1985), the effect of perturbing forces is strongest for this direction of movement, resulting in greatest position errors. This figure also shows the performance of the controller using a sensitivity-based approach where model parameters were varied. We considered a $\pm 15\%$ change in muscle viscosity, a $\pm 15\%$ change in inertia of the arm and link lengths, as well as a $\pm 50\%$ change in feedback gains and stiffness of the arm. We further assumed a $\pm 10\%$ variation in desired movement time, based on peak speed measurements for subjects. We found that, for parameters that described the behavior of the arm up until the first segmentation point, λ_1 , d_1 , and t_1 , performance of the system of Fig. 4 almost exactly matched that of the subjects. However, beyond the first segmentation point the controller of Fig. 4 could not account for the experimental data.

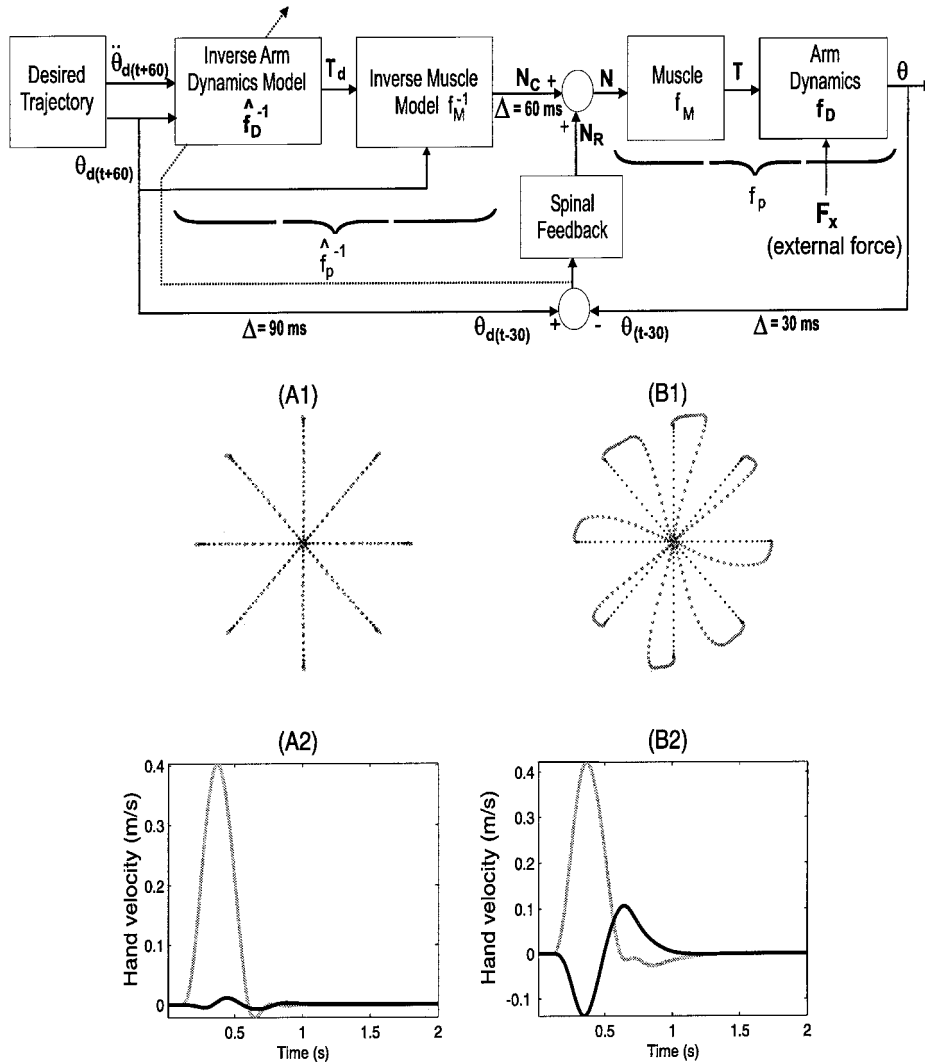


Fig. 4A,B. Block diagram illustrating a feedforward controller that utilizes an inverse dynamics model of the inertial properties of the limb, as well as an inverse muscle model, for assignment of neural activations to the muscles. **A1** and **B1** show hand trajectory of the system in a null field and in field B_1 . **A2** and **B2** show velocity of the hand for a movement to the bottom-most target. The gray line is velocity along the y-axis, and the black line is the velocity along the x axis

Based on this result, it appears that while the oscillations observed in the real arm after the first segmentation point are partly due to the intrinsic visco-elastic properties of muscles and the spinal loops, the behavior of the arm cannot be explained solely by this feedback system. In particular, note that at the first segmentation point, the arm does not move toward the target, but, on average, at $\lambda_2 = 17^\circ$ away from the target. Further note that the real arm has higher frequencies in its response to unmodeled dynamics than our controller (compare second row of Fig. 5). We will show that this is precisely the behavior that one would expect if the controller used a forward model to predict the state of the arm from time-delayed sensory feedback, and then acted to reduce that estimated error via an inverse model based controller.

3.3 Feedback control via a forward model

A forward model refers to a hypothetical computational network that can predict the change in the state of the arm from inputs that provide a copy of the descending

motor command and an estimate of the current state. The forward model may be in a Smith predictor in the case that the distal system is linear (Astrom and Wittenmark 1984; Miall et al. 1993) to predict the state of the arm at the current time, given delayed feedback about the actual state and the motor commands up to the current time. Alternatively, if the distal system is non-linear, as is the case here, a non-linear observer must be formulated to exactly model the forward dynamics of the system. Let us express the inertial dynamics of the arm and the dynamics of the muscles by a function f_p (Fig. 4) as follows:

$$\Sigma : \quad \ddot{x}(t) = f_p[N(t), x(t), \dot{x}(t)] \\ y(t) = [x, \dot{x}](t - t_0)$$

where N is the motor command input to the system, x is the position of the system, and y is the measured output of the system, i.e., delayed position and velocity of the limb. We have to design an observer that can estimate the current state $x(t)$ from the delayed state $x(t - t_0)$, given a particular history of descending motor commands $N_C(t)$. The observer design is as follows:

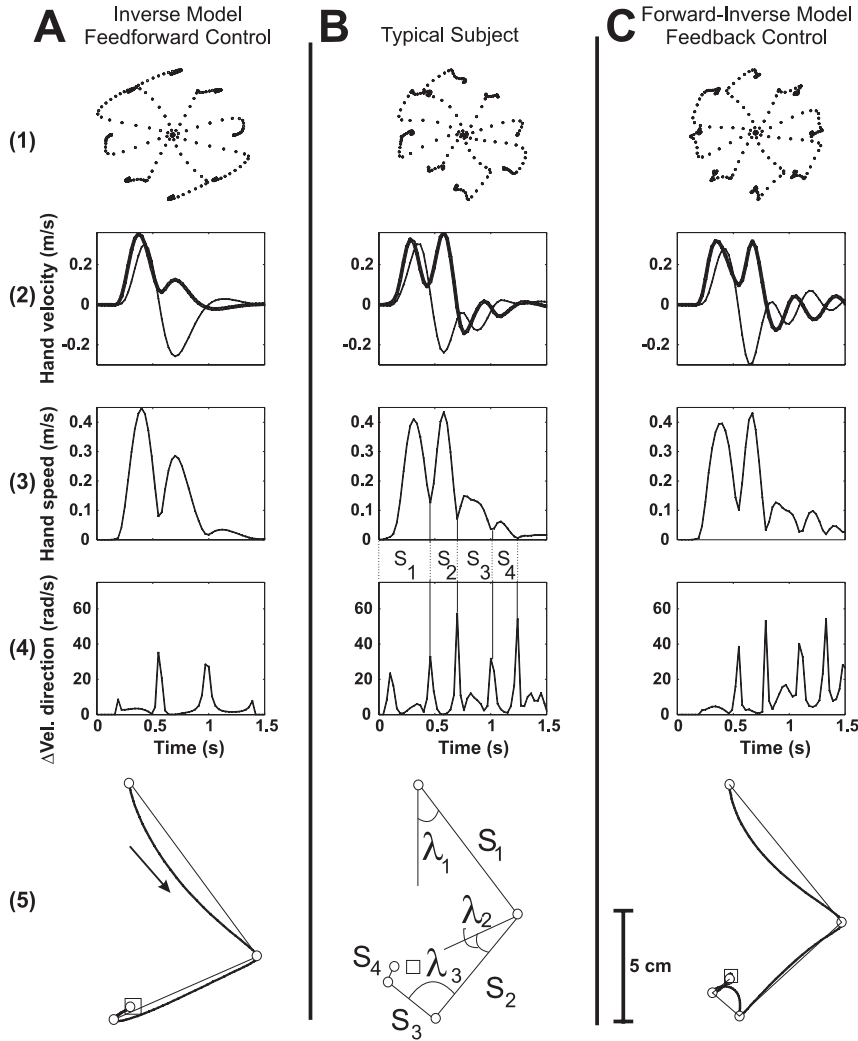


Fig. 5A–C. Trajectories are in field B_2 after subject and controllers had adapted to field B_1 . **A** Hand trajectories for the controller of Fig. 4, where only an inverse model is used. **B** Trajectories for a typical subject. **C** Trajectories for a controller corresponding to Fig. 11 (switch set to 1) which used a forward model in conjunction with an inverse model. *First row**, hand paths for eight movement directions. *Second row*, velocity along the y-axis (gray line, parallel to the direction of target) and x-axis (black line, perpendicular to the direction of target) for a movement toward a target at -90° . *Third row*, hand speed and segmentation points S_i for a movement toward -90° . *Fourth row*, derivative of velocity direction and corresponding segmentation points for a movement toward -90° . *Fifth row*, segmentation of the hand's trajectory

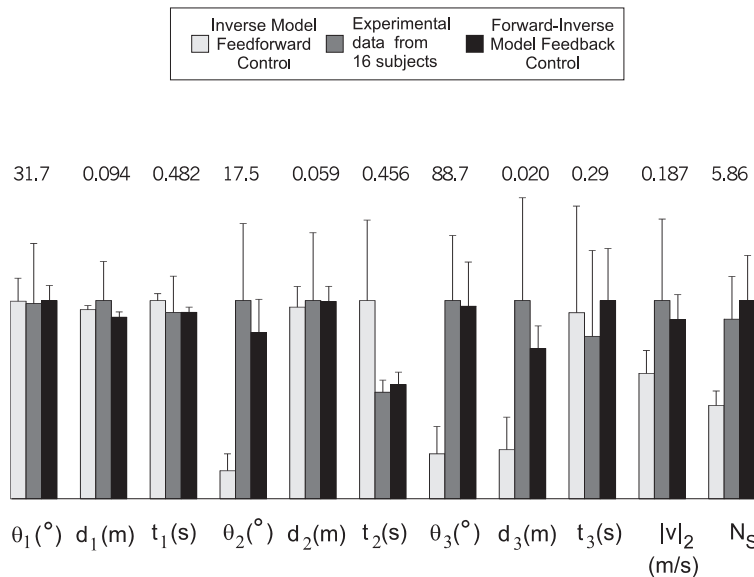


Fig. 6. Trajectory characteristics during a reaching movement toward the bottom most target (-90°) for 16 subjects in force field B_2 after adaptation to field B_1 (middle bar, dark gray). We have also plotted the results of 29 simulations of inverse model controller (light gray, corresponding to the controller in Fig. 4) and 35 simulations of the forward-inverse model feedback controller (black, corresponding to the controller in Fig. 11, switch set to 1) for the same movement. The trajectory parameters refer to the segmentation shown in Fig. 5. λ_i is angle about a segmentation point, t_i is the time to reach the i th segmentation point, d_i is the distance to the i th segmentation point, $|v|_i$ is the hand speed at the segmentation point, and N_S is the number of segmentation points in the trajectory. The value printed at the top of each bar triplet is the value at the mean for the highest bar in the triplet. Note that for λ_1 , d_1 , and t_1 , i.e., the initial part of the movement, performance of both controllers closely matched that of the experimental data. However, in later stages of the movement only the forward model based controller of Fig. 11 continued to accurately predict the experimental data




Fig. 7. Block diagram showing how a forward model \hat{f}_p of a non-linear system f_p can be used to construct an observer for a time-delayed nonlinear system where the state at time t is estimated from the measured state at time $t - t_0$ and the orderly cascade of descending commands N_c from time $t - t_0$ to t

$$\begin{aligned} \Sigma_0 : \quad & \hat{\dot{x}}(t) = \hat{f}_p[N_C(t), \check{x}(t), \check{\dot{x}}(t)] \\ & [\check{x}, \check{\dot{x}}](t - t_0) = y(t) \\ \check{\dot{x}}(t - t_0 + i\Delta) &= \check{\dot{x}}(t - t_0 + (i - 1)\Delta) + \int_{t-t_0+(i-1)\Delta}^{t-t_0+i\Delta} \hat{\dot{x}}(T) dT \\ \check{x}(t - t_0 + i\Delta) &= \check{x}[t - t_0 + (i - 1)\Delta] + \int_{t-t_0+(i-1)\Delta}^{t-t_0+i\Delta} \check{\dot{x}}(T) dT \\ \hat{x}(t) &= \check{x}(t) \\ \hat{\dot{x}}(t) &= \check{\dot{x}}(t) \\ i &= 1 \cdots \frac{t_0}{\Delta} \end{aligned}$$

where \hat{x} and $\hat{\dot{x}}$ are the outputs of the forward model, \check{x} and $\check{\dot{x}}$ are intermediate variables used by the forward model. The above equations represent the iterative solution of a non-linear differential equation \hat{f}_p at time t , given the initial state of the system y and the input N_C during the time interval $t - t_0$ to t . Δ is the discretized iteration time interval which should ideally be infinitely small. The value of Δ can be determined by the frequency response of the system and for simulations in the current study was chosen to be 0.004 s.

A network to implement the forward model based observer is presented in Fig. 7. It requires multiple copies of the forward model because of the iterative nature of the method that must be used to arrive at the solution to the nonlinear differential equation describing the dynamics of the muscles/limb/load. As it is evident from this figure, formulating an accurate forward model is a surprisingly intensive process. This is because dynamics of a non-linear system are state dependent and arrival of delayed sensory feedback triggers a re-estimation of the current state². Here we arbitrarily represented the computation time as an 8 ms delay in the control system. The effect of this delay is to impose a limit on the gain of the feedback loop that will act on the output of the forward model.

For simulations of the forward model (observer) presented here, t_0 in the description of Σ_0 has a value of 210 ms. This is the look-ahead period for which the

forward model is integrated. This value represents the sum of the 150-ms delay in receiving sensory feedback from the limb and the 60-ms delay in the transmission of descending motor commands to the muscles, including the delay in activation force response of the muscle. The output of the observer represents the best knowledge of the current state of the hand. Three different modalities of control, based on three coordinate systems in which the observer might estimate state of the arm are considered here. First, we consider control via an observer that estimates muscle position and velocities. Second, an observer in joint coordinates. Third, an observer in Cartesian coordinates of the hand.

Consider the control system of Fig. 8 with the switch set to position 2. The forward model provides an estimate of the current muscle lengths and velocities based on a copy of descending commands and time-delayed sensory feedback. This estimate is compared with the desired trajectory in muscle space, resulting in an error estimate. A command N_C is computed as:

$$N_C = K_p(\hat{x}_m - x_{m_d}) + K_v(\hat{\dot{x}}_m - \dot{x}_{m_d}),$$

which acts as a linear error feedback controller. When the forward model perfectly describes the dynamics of the muscle/arm system, and when the gains K_p and K_v are sufficiently large, the effect of the forward model coupled with the linear feedback controller is to approximate an inverse model of the plant. To see this, consider the very simple linear system of Fig. 9. Here, we wish to control dynamics of a system, specified by G , via a forward model \hat{G} . Note that $y/x = 1/(1 + G)$. Therefore, $y/x \approx \hat{G}^{-1}$ when $G \gg 1$, i.e., an inverse model of the plant.

In the case of Fig. 8, where the system is non-linear, the main question is with regard to the gains K_p and K_v in the linear feedback loop. The gains need to be high if we wish to closely follow the desired trajectory, but as the gains increase, the system edges closer to instability. This is for two reasons. First, there is a small delay in the computations of the forward model and, therefore, in the feedback loop. Second, with high gains, the system becomes very sensitive to unmodeled dynamics, i.e., it will easily become unstable when the arm is coupled to an unknown load. In contrast, if the gains are low, the system barely follows the desired trajectory but is robust to unmodeled dynamics. In examining this control system, we found that even with a perfect forward model, no combination of gains on the linear feedback system could be found so that the system closely followed the desired trajectory and remained stable in field B_1 .

² To reduce computational complexity, one approach is to re-estimate the current state $\hat{x}(t)$ only when there is a larger than threshold difference between $\hat{x}(t - t_0)$ and $x(t - t_0)$. Effectively, this changes the feedback system through the forward model to an intermittent feedback controller (Ronco 1998). Application of a similar idea to tracking movements has recently been demonstrated (Hanneton et al. 1997).

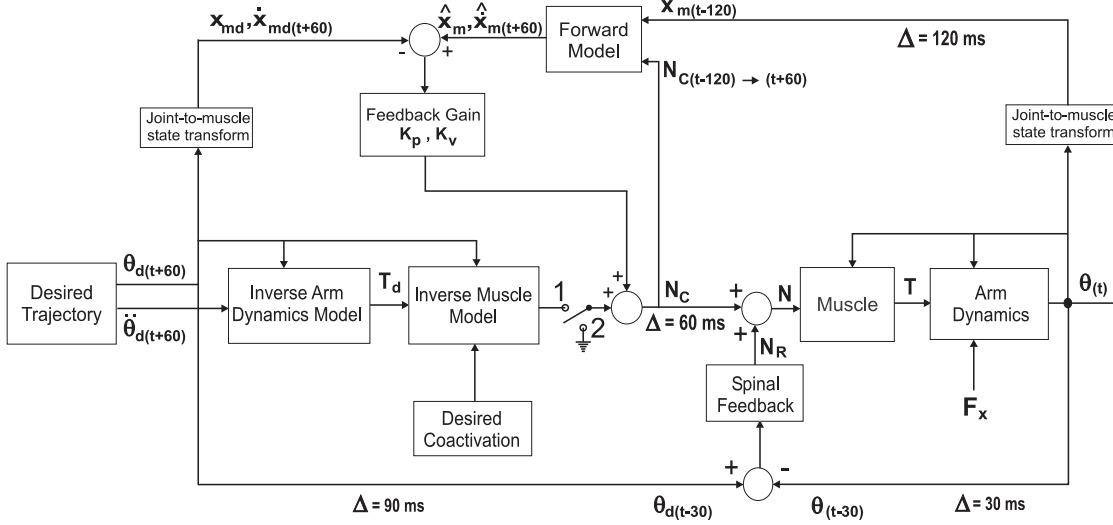


Fig. 8. A control method that uses a forward model in muscle coordinates. The switch is in position 2 for feedback control using only the forward model, and in position 1 when a forward model is used in conjunction with an inverse model




Fig. 9. A simple linear system that uses a forward model \hat{G} of system dynamics G . Note that $y/x = 1/(1 + \hat{G}) \approx \hat{G}^{-1}$, i.e., the feedback loop effectively approximates an inverse of the plant dynamics G

The main problem with the control scheme of Fig. 8 (switch set to 2) is that even with a perfect forward model, the transformation from an error in position to neural commands is a non-linear map that depends on the inertial dynamics of the arm and force-activation dynamics of the muscles. The performance of this system could be improved if an inverse muscle model was available to transform the error commands into neural activations. This control system is shown in Fig. 10, where the forward model estimates the state of the limb in joint coordinates. We arrived at the gains $K_p = 30$ N/rad and $K_v = 3$ N·s⁻¹·rad⁻¹ in the linear error feedback loop of Fig. 10 by initially setting the forward model to approximate the dynamics of the arm/muscle in the null field, then finding the gains that made the system marginally stable in field B_1 . K_p and K_v were set at 50% of this value. Performance of the resulting control system is shown in Fig. 10 for three conditions: first, when the forward model is expecting a null field and the arm moves in the null field; second, when the forward model is expecting a null field but the arm moves in field B_1 ; third, when the forward model is expecting B_1 and arm moves in field B_1 . The results demonstrate that even with a perfect forward model, because of the trade-off between gain and susceptibility to unmodeled dynamics, the arm trajectories in B_1 are far from the desired trajectory.

Similar results were found if the estimate of error in position for the forward model was transformed to neural commands through the use of both an inverse

limb model and an inverse muscle model. The control system is shown in Fig. 11 (switch set to 2), where now the forward model estimates the state of the limb in hand coordinates. We again found that with the gain on the feedback loop set to a high level, the system closely followed the desired trajectory, but was very sensitive to unmodeled dynamics. As a compromise, we set the gain at a level that was half as high as would make the system marginally stable. This resulted in a system that was stable in field B_1 but did not closely follow the desired trajectory when the forward model correctly estimated state of the limb.

In summary, if our control system is driven by only an error signal from the forward model, then several factors limit the gain of the feedback loop, which in turn prevent the system from closely following the desired trajectory even when the forward model is accurate. These factors include the non-zero computational time of the forward model (here assumed to be 8 ms), and the desire to keep the system stable when the arm is in contact with an unknown load.

3.4 Control using feedforward and feedback pathways

We next considered the performance of a system where both a feedforward pathway (consisting of the inverse model) and a feedback pathway (via the forward model) were used to generate descending motor commands. This corresponds to the case where the switch is set to position 1 in Figs. 8, 10, and 11. In this approach, the role of the forward model is not to provide the driving input to the system, but to respond only if there are dynamics in the distal plant that are not compensated for in the actions of the inverse models. We tested each controller with the gains K_p and K_v unchanged from above. Obviously, when the inverse model is perfect, the forward models pathway makes no contribution to the system and the arm moves along the desired trajectory.

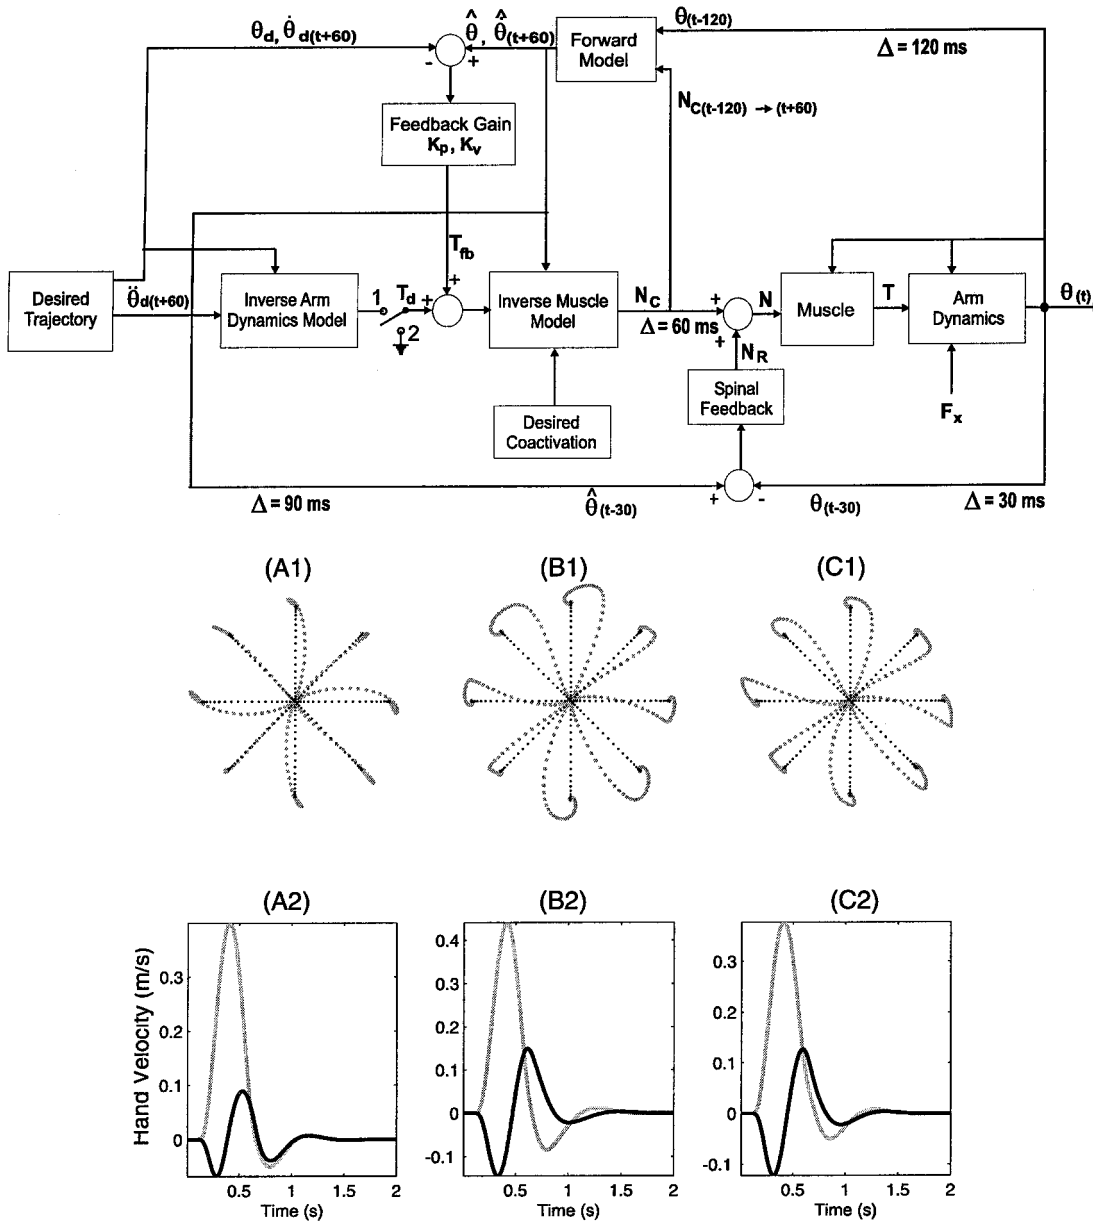


Fig. 10A–C. A control scheme that uses a forward model in joint coordinates. The switch is in position 2 for control via only the forward model (FM), and in position 1 for control via both the forward and inverse models (IM). Trajectories are for switch in position 2. In **A1** and **A2**, FM expects null field, arm moves in the null field. In **B1** and **B2**, FM expects null field, arm moves in force field B_1 . In **C1** and **C2**, FM expects B_1 , arm moves in B_1 . Hand paths are represented as dots at 20-ms intervals (the straight path is the desired trajectory). Hand velocities are for a downward movement. Gains on the linear feedback error controller, $K_p = 30 \text{ N/rad}$ and $K_v = 3 \text{ N} \cdot \text{S}^{-1} \cdot \text{rad}^{-1}$, were set at 50% of the value for which the system was marginally stable. Even with a perfect forward model, the system is not able to follow the desired trajectory

However, the main question is, how does the system behave when there are unmodeled dynamics? To explore this, we re-examined the data from our subjects during the condition where they had trained in field B_1 , but were suddenly presented with B_2 . This presents the largest error in expected dynamics because $B_2 = -B_1$, providing us with the best opportunity to ask whether the behavior of the biological controller could be explained by the influence of the forward model.

The data from a typical subject was shown in Fig. 5 and we had previously concluded that performance of the controller in Fig. 4 could not account for this be-

havior. The resulting trajectories for the controller of Fig. 11 (switch set to 1) are shown in Fig. 5C. Here, we assumed that after training in B_1 , both the forward and inverse models accurately represented the dynamics of B_1 . Without any modification to the parameters of the system, we found a remarkable similarity between the actual and simulated trajectories when the field was changed to B_2 . In particular, note the higher frequencies in the response of the system (2nd row of Fig. 5C) and the behavior about the segmentation points. We performed a sensitivity analysis by varying the parameters of the model: we considered a $\pm 15\%$ change in muscle

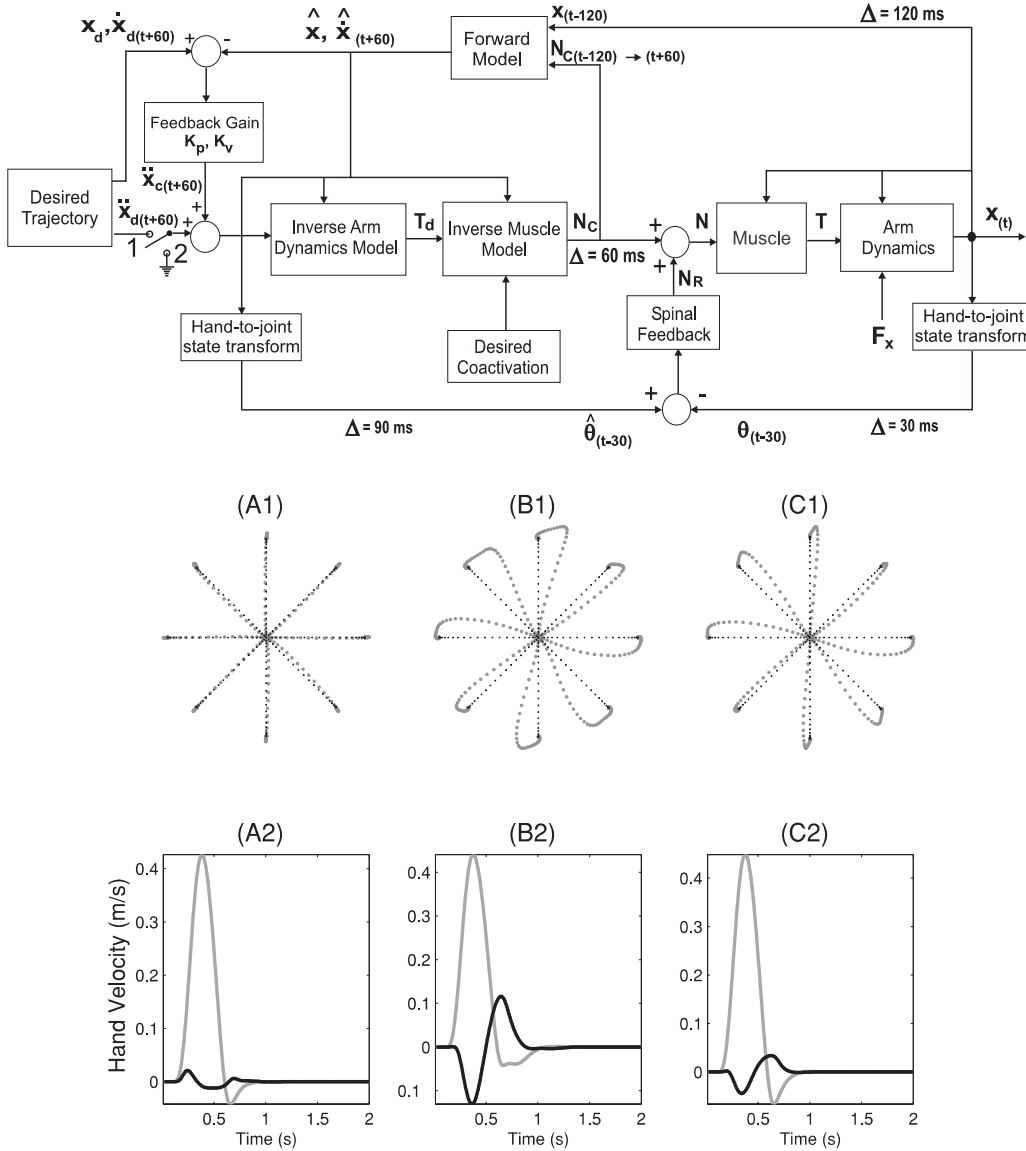


Fig. 11A–C. A control scheme that uses a forward model in hand coordinates. The switch is in position 2 for feedback control via the forward model, and in position 1 for control via both the forward and inverse models. Simulated trajectories for switch in position 2. In **A1** and **A2**, FM expects null field, arm moving in the null field. In **B1** and **B2**, FM expects null field, arm moves in force field B_1 . In **C1** and **C2**, FM expects B_1 , arm moves in B_1 . Gains on the linear feedback error controller, $K_p = 500 \text{ s}^{-2}$ and $K_v = 50 \text{ s}^{-1}$, were set at 50% of the value for which the system was marginally stable. Even with a perfect forward model, the system is not able to follow the desired trajectory

viscosity, a $\pm 15\%$ change in inertia of the arm and link lengths, as well as $\pm 50\%$ change in feedback gains (both the spinal loop and the linear controller attached to the forward model) and the stiffness of the arm. We further assumed a $\pm 10\%$ variation in desired movement time based on peak speed measurements for subjects. The results of this approach were quantified via the effect of model parameter variations on movement parameters, and are summarized in Fig. 6. This figure allows for a comparison of the data from all our subjects with the simulation results. Every parameter appears to be accurately predicted by the behavior of the control system in Fig. 11 (switch at position 1). Similar results were found when the control architectures in Figs. 8 or 10 were used. For this reason, for the remainder of this

report we will concentrate on the system of Fig. 11 as a prototype.

Why does the arm behave as it does about the segmentation points? When only an inverse model is available (Fig. 4), the muscles and reflex pathways are programmed based on the desired trajectory of the arm. This implies that the equilibrium position for the muscles is the desired target at $t = 500 \text{ ms}$ and the corrective action for $t > 500 \text{ ms}$ is like a visco-elastic system pulling the arm directly toward the target. However, when a forward model is used in conjunction with the inverse model (Fig. 11), the descending neural commands rely on the estimated trajectory instead of the desired trajectory for the system. Furthermore, the corrective actions taken by the system rely on both the spinal loop/muscle

visco-elastic properties, and the predicted error signal from the forward model. In this situation, there are three reasons for the behavior about the segmentation points in Fig. 5C: first, the external force field B_2 pushing the hand in an anti-clockwise direction; second, the forward model incorrectly anticipates a clockwise field B_1 and generates position estimates accordingly; third, the inverse model incorrectly generates additional torques in the anti-clockwise direction to counteract the clockwise field B_1 . Therefore, both the wrong inverse and forward models are contributing to the behavior about the segmentation points. Which is more important?

To assess the relative role of forward and inverse models in the controller of Fig. 11, simulations in force field B_2 were carried out for three conditions: first, we assumed that after training in B_1 , only the forward model might have adapted to field B_1 , i.e., inverse model continued to expect a null field; second, we assumed that only the inverse model might have adapted to B_1 , i.e., forward model expected a null field; third, both models had adapted to B_1 . The results of simulations in field B_2 are shown in Fig. 12. We note that if, after training, only the inverse model has adapted, we fail to see a significant segmentation pattern. The pattern is observed only when the forward model has adapted, regardless of the state of the inverse model. This establishes that the segmentation behavior is mainly a result of adaptation of the forward model to force field B_1 and is not significantly affected by the state of the inverse model.

It is now a question of how the forward model contributes to the behavior about the segmentation points. The state estimates are used as input to detect errors in position and velocity, and provide state input to the inverse model. Which is the main cause of the segmentation? We simulated the behavior of the system in field B_2 with the forward model expecting field B_1 and the inverse model correctly modeling field B_2 . The results of

the simulation are plotted in Fig. 13. The estimated trajectories are shown along with the desired and actual trajectories in Fig. 13 A–C. It is difficult to visualize the control process through only these estimates; therefore, in Fig. 13 D–F the desired, corrective and actual acceleration signals are plotted as vectors. It is immediately apparent from Fig. 13E that the cause of the segmentation behavior is inappropriate feedback from the forward model that tries to accelerate the hand in an anti-clockwise direction away from the target. This establishes that through error feedback, the forward model generates incorrect state estimates which are the main reason behind the behavior about the segmentation points.

3.5 Robustness to measurement noise

While the primary purpose of our report is to account for the behavior of the human arm, it is worth noting some of the other properties of the controller of Fig. 11 from a purely practical perspective. In the design of a control system that relies on a forward model (or observer), two concerns are paramount: robustness to unmodeled dynamics, and robustness to measurement noise. In the above discussion we presented the behavior of the system when the distal dynamics were substantially unmodeled and found the system to be stable. Moreover, the proposed control system appeared to closely account for the behavior of the biological controller. What happens if measurements of state are noisy? A forward model is expected to be particularly susceptible to measurement noise because it is attempting to predict the future state of a non-linear process from some initial conditions. If these initial conditions are measured by noisy sensors, how will the performance of the control system be affected?

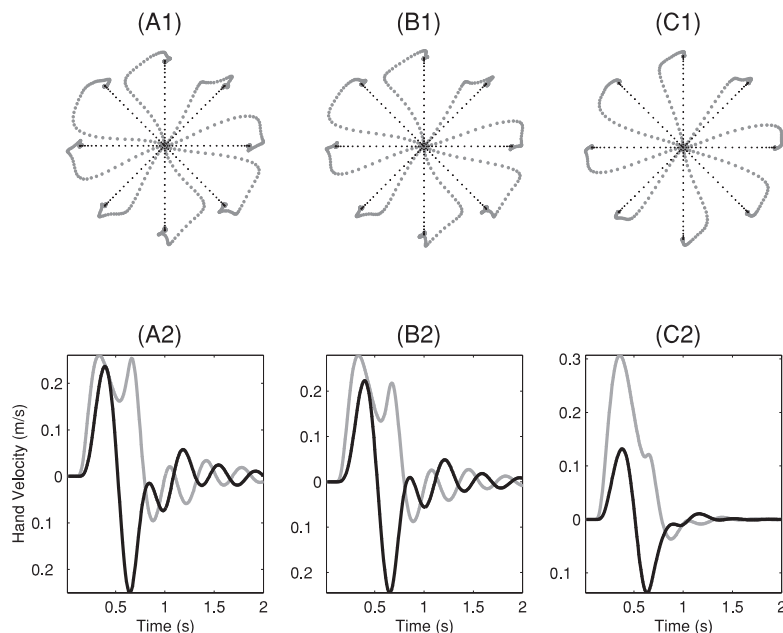


Fig. 12A–C. Trajectories for the controller of Fig. 11 (switch set to 1). All movements are in field B_2 . We considered three different states of adaptation of the inverse model (IM) and the forward model (FM). In **A1** and **A2**, IM = B_1 , FM = B_1 . In **B1** and **B2**, IM = null field, FM = B_1 . In **C1** and **C2**, IM = B_1 , FM = null field. The term null implies that the model compensates for only the inertial dynamics of the limb. Note that the segmentation behavior and the high frequency in the response of the system are present regardless of the state of adaptation of the inverse model. However, the segmentation is present only if the forward model has adapted

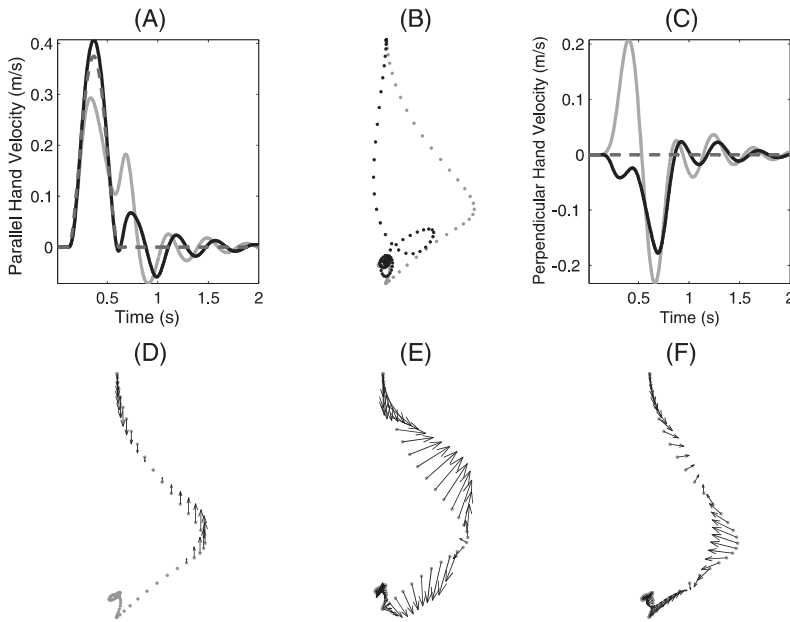


Fig. 13A–C. Simulation results for controller of Fig. 11 (switch set to 1) for movements in field B_2 for a movement in a downward direction. The inverse model correctly expects B_2 while the forward model expects B_1 . **A** Hand velocity parallel to the direction of target for the actual trajectory of the arm (gray line), estimated trajectory (black line, i.e., output of the forward model), and desired trajectory (dotted line). **B** Hand paths for actual trajectory (gray dots) and estimated trajectory (black dots). **C** Similar to **A**, except for a plot of the hand velocity perpendicular to the direction of target. **D–F** Desired, estimated and actual acceleration signals plotted as vectors at 20-ms time points on the actual hand trajectory. The largest acceleration vector in the three plots has a magnitude of 4.6 m/s^2 and all other vectors are scaled relative to that

To address this concern, we note that it has been shown that for systems with no delay in feedback, addition of an element that provides appropriate local error feedback on the plant can provide for an accurate velocity estimation from only position measurements in an inertial system (Lohmiller and Slotine 1996). Here, we demonstrate that for the system of Fig. 11, the design of the musculo-skeletal system allows the forward model to be remarkably robust to noise in velocity measurements.

We injected a random 5-Hz noise into the velocity signal received by the forward model in order to observe the behavior of the system. The magnitude of this noise was at its maximum equal to the actual velocity signal. In Fig. 14, the hand paths and hand velocity signals are plotted for the actual movement trajectory, the measured movement trajectory and the estimated movement

trajectory for each of these cases. Despite the fact that there is substantial noise in velocity measurements, the estimated velocity profile, i.e., the output of the forward model, is almost exactly the same as the actual one.

The reason for this remarkable robustness of the forward model to measurement noise is that the descending commands are not specifying a particular torque. Rather, the commands are specifying an equilibrium-like state for the distal visco-elastic system. The forward model integrates the descending neural commands over a 200-ms period given some initial state of the system. Therefore, even when the initial states are incorrect due to noise, the output from the forward model tends towards the actual state of the system because of the equilibrium properties of the system that it is being controlled.

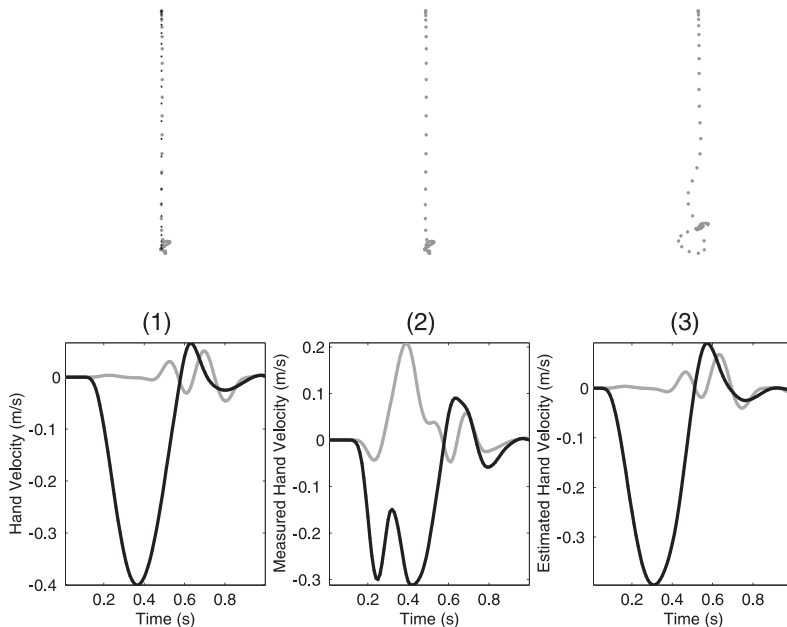


Fig. 14. Simulation results for a movement downward in a condition where measurements of velocity are very noisy. **1** Actual hand path (top row) and velocity (bottom row) of the arm. Hand velocity perpendicular to the direction of target is shown in gray. Hand velocity parallel to the direction of target is shown in black. The velocity signal is noise free. **2** Behavior of the controller when measured velocity is inaccurate. The actual hand path is very similar to the case when there was no noise in the velocity measurements. The controller is robust to measurement noise in velocity. **3** Output of the forward model during the control process. Estimated hand path and velocity are plotted. Estimated velocity is very robust to measurement noise

4 Rates of adaptation of the internal models

While certain features of the biological controller, e.g., behavior about the segmentation points, suggest that generation of descending commands relies on a combination of forward and inverse models (Fig. 11), we have yet to determine how these models might change during a practice session with a novel force field. Of particular concern is the question of whether the data from our 16 subjects who practiced in novel force fields allows us to differentiate between the rate of adaptation of the models. In other words, during practice in a force field, how fast do each of models adapt? Is there any evidence that the two models adapt at different rates?

4.1 Learning of field B_1

To illustrate our approach, we begin with an extreme example. Consider the behavior of the controller in Fig. 11 (switch set to 1) when the arm is initially exposed to field B_1 . How much improvement in performance can we expect if the inverse model completely adapts to B_1 but the forward model does not? How much improvement in performance can we expect if only the forward model adapts? To answer these questions, we assumed a decaying exponential change in either the forward or the

inverse models and plotted key parameters of the performance of the system for the sequence of targets that were presented to our subjects (Fig. 15). There are two lines in each sub-figure, one corresponding to the change in a particular parameter as the forward model adapts (black line), and the other corresponding to the change as the inverse model adapts (gray line). We initially assumed an exponential learning rate with a time constant of 0.02/movement. This implies that by the 50th movement, each model accounts for 63% of the dynamics of the force field. The two cases clearly predict different adaptation curves for most movement parameters. When the forward model is adapting but the inverse model is not, performance improves dramatically. In contrast, when only the inverse model is adapting, there are much smaller improvements in performance. Therefore, the performance of the system is highly dependent on the rate of adaptation of the forward model, and much less so on the rate of adaptation of the inverse model.

Now consider the possibility that during practice, both the inverse and the forward models are adapting but with possibly different rates. Assume that these rate are r_{fm} for the forward model and r_{im} for the inverse model.

$$FM(n) = FM(0) + \Delta\{FM\}(1 - e^{-nr_{fm}})$$

$$IM(n) = IM(0) + \Delta\{IM\}(1 - e^{-nr_{im}})$$

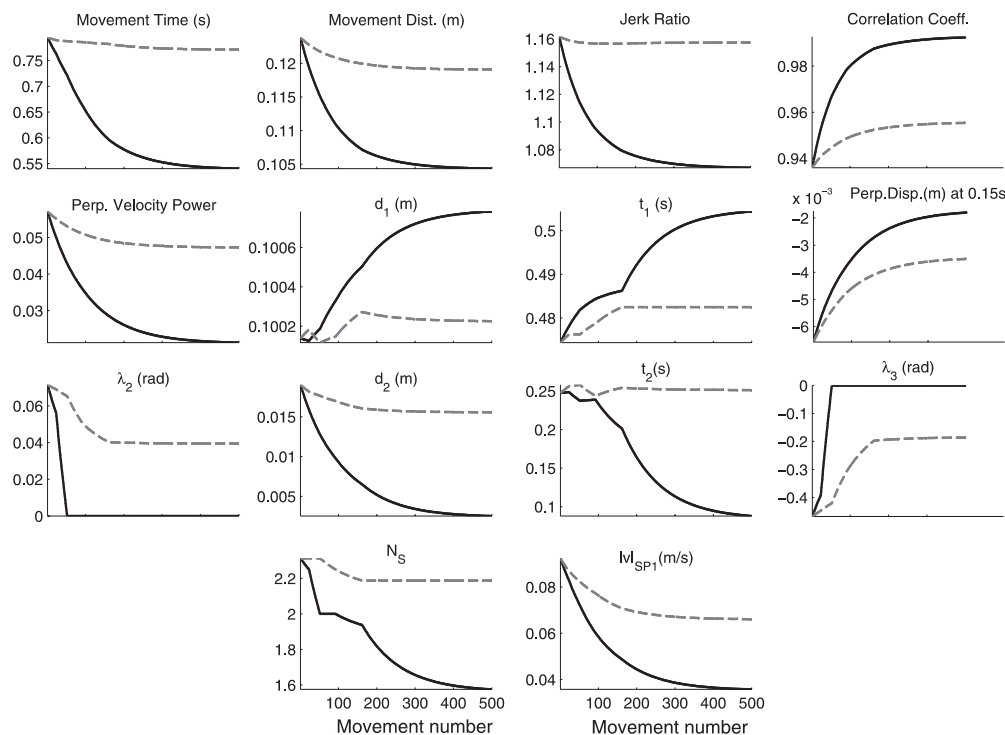


Fig. 15. Adaptation curves for movement parameters during learning of field B_1 in two cases – (1) *dotted line*, only the inverse model adapts exponentially to B_1 at a rate of $r_{im} = 0.02$, $r_{fm} = 0$; (2) *solid line*, only the forward model adapts to B_1 at a rate of $r_{im} = 0$, $r_{fm} = 0.02$. The desired trajectory was a 10-cm minimum jerk motion performed in 0.5 s. Jerk ratio is the ratio of cumulated squared jerk in a movement with respect to the minimum jerk possible for a movement of the same peak speed. The correlation coefficient is with respect to the minimum jerk motion. The perpendicular distance refers to the distance of the hand from the min jerk motion at 150 ms into the movement. Perp. Power refers to the power in the frequency spectrum of the velocity of hand along a direction perpendicular to the direction of target. d_i , t_i and λ_i , refer to the distance to, time to, and angle at the i th segmentation point. N_s is the number of segmentation points, and $|v|_{sp1}$ is the hand speed at the first segmentation point

$FM(n)$ and $IM(n)$ are the adaptation states of the two models at movement number n in the force field and represent the time course of adaptation for the two models. $FM(0)$ and $IM(0)$ are the initial states of the forward and inverse models at the beginning of the force field training. $\Delta\{FM\}, \Delta\{IM\}$ are the difference in the initial state of the models and the force field being learned. The equations are obtained by considering a rate of learning of the models that is proportional to the difference in the model and the field at any instant of time. When subjects begin the training in field B_1 , they initially expect the null field. Hence, the initial states of both the forward and inverse models are set to the dynamics appropriate for the null field, i.e., the coupled dynamics of the human arm and the robot with no forces being produced by the robot's motors. This is referred to $FM(0) = IM(0) = \text{null}$. Furthermore, $\Delta\{FM\} = \Delta\{IM\} = B_1$. Here, we tried to find a best estimate of the two learning rates by comparing the performance of the model system at a given rate of adaptation to the performance of the 16 subjects. We considered six values for $r_{im} = \{0.0003, 0.003, 0.01, 0.03, 0.1, 0.3\}$, and five values for $r_{fm} = \{0.003, 0.01, 0.03, 0.1, 0.3\}$, for a total of 30 combinations.

To compare performance of the controller at a given rates of adaptation with the performance of our subjects, we initially quantified each movement of each subject with 16 parameters. These parameters were movement time, movement distance (total length of a movement), peak speed, jerk ratio (the ratio of the cumulative squared jerk of a movement with respect to the cumulative squared jerk for a minimum jerk movement of the same peak speed), correlation with a minimum jerk movement, perpendicular displacement of the hand from a straight line to the target at 150 ms into the movement, power in the frequency spectrum of the velocity of the hand along a vector perpendicular to the direction of target, d_i (distance to the i th segmentation point), t_i (time to the i th segmentation point), λ_i (angle of the hand trajectory about the i th segmentation point, see Fig. 5), N_s (number of segmentation points) $|v|_{sp1}$ (speed at the first segmentation point). Let us refer to these parameters with variable $p = 1 \dots m$, $m = 16$.

These movement parameters were quantified for each of the 576 movements of each subject. Let us refer to the movements with variable $n = 1 \dots 576$. The resulting "learning curves" for all subjects are shown as average \pm SD in Fig. 17.

We next quantified the same movement parameters p for a control system with particular learning rates r_{im} and r_{fm} . To exactly simulate experimental conditions faced by our subjects, we used the same sequence of targets (movement directions) that were experienced by the subjects. This procedure was repeated for all 30 combinations of r_{im} and r_{fm} . Let us label each pair of rates by the variable $q = 1 \dots 30$. The next step was to find the one pair of rates that resulted in a control system that had movement characteristics that most resembled data of our subjects. To do this, an error measure e was defined for each movement parameter p , as measured over n movements for adaptation rates q :

$$e_{pq} = \frac{\sum_n (|y_{pqn} - \mu_{pn}| + \sigma_{pn})}{\sum_q \sum_n (|y_{pqn} - \mu_{pn}| + \sigma_{pn})} \quad (4)$$

where, y_{pqn} is the value of the simulated movement parameter p at the n th movement corresponding the adaptation rate q . μ_{pn} is the mean of the parameter value for our 16 subjects at movement n and σ_{pn} is the corresponding standard deviation. The numerator in the equation is almost equal to the average area between the simulated and experimental adaptation curves. This error is normalized by the denominator which is the sum of the errors for the 30 different rates of adaptation for the inverse and forward models, making it independent of parameter units and values. To combine the information from the different movement parameters p , the errors were summed together to give a net error E_q for a particular pair of adaptation rates q ,

$$E_q = \frac{\sum_p e_{pq}}{m} \quad (5)$$

This net error is plotted as a function of r_{fm} and r_{im} in Fig. 16. The region surrounded by the thick black line outlines the minima. The error is at its lowest for $r_{fm} = 0.01$. This means that by the 100th movement, a typical subject's




Fig. 16A,B. Normalized error in matching performance of the adaptive controller of Fig. 11 (switch set to 1) with that of 16 subjects that practiced in field B_1 for 572 movements. **A** The error is plotted as a function of rates of adaptation of the forward model r_{fm} and inverse model r_{im} . **B** The region of minimum error is highlighted by the *thick black line* in the two-dimensional projection. Note that while the error surface is sharply defined in terms of changes in the forward model, it is fairly flat to variations in the learning rate of the inverse model

forward model accounted for 63% of the dynamics of the force field. This value lies at the bottom of a sharply defined region, suggesting a high degree of sensitivity, and therefore confidence, in estimating the rate of adaptation of the forward model. In contrast, the rate of adaptation for the inverse model cannot be precisely estimated because the minimum lies in a fairly shallow valley. There are no significant differences in the error measure whether the rate of learning of the inverse model is $r_{im} = 0.003$ or $r_{im} = 0.03$. The reason for the shallowness of the valley for r_{im} is the much weaker dependence of movement parameters on the rate of adaptation of the inverse model.

To visualize how well the rate of adaptations of the inverse and forward models accounted for the pattern of learning in our subjects, we compared changes in performance of the model with that of the subjects. In Fig. 17 we have the changes in various movement parameters for the 16 subjects (mean \pm SD) as they practice in the force field. In this figure we also have the changes in movement parameters of the adaptive controller for a particular combination of adaptation rates of the forward and inverse models, $r_{fm} = 0.01$, $r_{im} = 0.01$. The performance of the adapting model controller accurately captures the trajectory of changes in the performances of our subjects. The simulation results even mimic the set structure, which is due to the sequence of movement directions for the experimental data.

In summary, theoretical results suggest that significant improvements in performance can be achieved with adaptation of only the forward model. In contrast, adaptation of only the inverse model results in much more modest improvements in performance. While this might suggest that a reasonable policy for an adaptive

controller might be to allocate resources mostly to learning of the forward model, the actual performance of our subjects does not provide convincing evidence to support this hypothesis. The psychophysical data suggests that the inverse and the forward models adapt at fairly comparable rates, accounting for approximately 63% of the dynamics of the field by the 100th movement. We note, however, that our estimation of the rate of adaptation in the inverse model is much more tentative than that of the forward model because of the relative insensitivity of the movement parameters (for the task studied here) to changes in the inverse model.

4.2 Learning of field B_2

A fundamental observation in the way humans learn force fields is that the ability of subjects to learn a counter example (field B_2) of a previously learned field (B_1) depends on the time that has passed since adaptation to that field (Brashers-Krug et al. 1996). During this period, significant changes appear to occur in the functional properties of the motor memory (Shadmehr and Brashers-Krug 1997), changes which coincide with shifts in the neural correlates of the memory (Shadmehr and Holcomb 1997). Here, we applied the above computational framework and quantified the rate of adaptation of the forward and inverse models under two conditions: first, when subjects were exposed to B_2 5 min after completing 572 targets in B_1 ; second, when another group of subject was exposed to B_2 6 h after completion of practice in B_1 .

The estimation procedure was as described above, except that at the start of practice in field B_2 we assumed

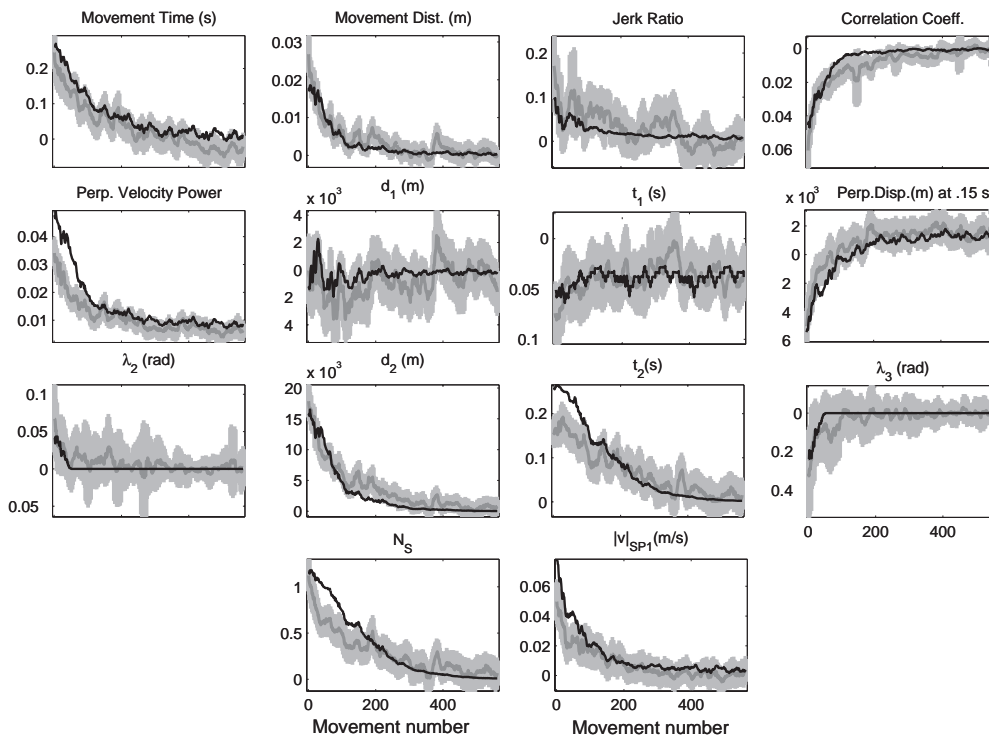


Fig. 17. Changes in movement parameters during learning of field B_1 for the adapting controller (black) with adaptation constants $r_{fm} = 0.01$, $r_{im} = 0.01$, and for experimental data from 16 subjects (gray) plotted as the mean and standard deviation. All values are represented as change in the movement parameter with respect to values recorded after subjects/models had adapted to movements in the null field. Movement parameters are as in Fig. 15




Fig. 18A,B. Normalized error in matching performance of the adaptive controller of Fig. 11 (switch set to 1) with that of 16 subjects that learned field B_2 after training in B_1 . **A1** and **A2** Error as a function of adaptation rates in the forward and inverse models for subjects that trained in B_2 at 5 min after completion of training in B_1 . The position with minimum error value is at $r_{im} = 0.0028$ and $r_{fm} = 0.015$. That is, the forward model appeared to adapt about 5.4 times faster than the inverse model. **B1** and **B2** Subjects that trained in B_2 at 6 h after B_1 . The position with minimum error value is at $r_{im} = 0.004$ and $r_{fm} = 0.020$. Again, the forward model appeared to adapt about five times faster than the inverse model. Furthermore, at 6 h both the forward and inverse models appear to be adapting significantly faster than at 5 min

that the forward and inverse models had fully adapted to B_1 . This was based on the result that the best estimate of the adaptation rate in field B_1 was 0.01 for both models, i.e., by the 300th target, both the inverse and forward models could account for 95% of the dynamics of the force field (subjects received 572 targets).

We simulated performance of the system (Fig. 11, switch set to 1) in field B_2 with the following rates of adaptation: $r_{fm} = (0.008, 0.01, 0.015, 0.025, 0.04)$, and $r_{im} = (0.001, 0.003, 0.005, 0.008, 0.01)$. An error measure, as defined in Eq. (4), was estimated for each movement parameter and adaptation rate. We used Eq. (5) to arrive at a normalized error measure for all parameters at a given adaptation rate. The results are plotted in Fig. 18. We found that in learning of field B_2 at 5 min after completion of practice in B_1 , the minimum in this error function was at $r_{im} = 0.0028$ and $r_{fm} = 0.015$. This suggested that the forward model was adapting at a rate that was approximately 5.4 times that of the inverse model.

When the procedure was repeated for the data of the 6-h group, the normalized error was minimized when $r_{im} = 0.004$ and $r_{fm} = 0.020$. Remarkably, in this group, the forward model also adapted at approximately 5.0 times that of the inverse model. To our knowledge, this is the first evidence that in learning to control the arm, the human adaptive controller may rapidly learn a forward model but acquire an inverse model at a much more slower rate.

The results also show that when subjects were presented with field B_2 , the rates of adaptation in both the inverse and the forward models were faster in the 6-h group than in the 5-min group. This suggests that as the memory of B_1 consolidated, certain computational re-

sources that might have been used in adaptation of the internal models once again became available, resulting in more rapid rates of adaptation at 6 h as compared to 5 min.

5 Discussion

Here, we approached the problem of controlling a time-delayed mechanical system similar to the human arm. The task that we considered was reaching movements in novel force fields. We demonstrated that essential characteristics of the arm's trajectory could not be accounted for if the supra-spinal controller was an open loop system composed of a model of the inverse dynamics of the arm (Fig. 4). These characteristics were related to the frequency response of the system to a perturbation and behavior about segmentation points, i.e., points where both the derivative of hand speed and the direction of hand velocity changed rapidly.

We next considered the design of the supra-spinal controller with a forward model (also known as an observer). The forward model predicted the position and velocity of the arm from a copy of descending motor commands and the latest sensory feedback from the moving arm. Descending neural commands were generated through a comparison of the predicted state of the system and its desired state. It was found that psychophysical data could not be suitably modeled with this system. First, because of small delays inherent in the computations of the forward model, gains of the feedback loop that generated an error signal could not be set large enough to effectively provide an inverse model of the dynamics of the distal system. Second, the gains were

further limited because of the instability that resulted when there was unmodeled dynamics in the distal system, for example, when the arm was coupled to a novel force field. This suggested that it was unlikely that the biological controller generated descending motor commands solely through the use of a forward model.

We next considered the possibility that descending commands were generated through a control architecture that included both an inverse and a forward model. In this scenario, the forward model did not provide the driving input to the system, but an input only when the feedforward pathway (which consisted of the inverse model) could not precisely account for dynamics of the distal system. In this architecture, the essential idea was that the entire control system, including the set points for the spinal reflex loops and state-dependent maps in the inverse model, relied on an estimate of the current state, i.e., the output of the forward model, rather than the desired state trajectory for the system. We showed that using this architecture the behavior of the human arm in a force field could be almost precisely accounted for. The response of the system in a novel force field had dynamics which included segmentation points. The behavior about these points were remarkably similar to those we had observed in our subjects. This suggested that the trajectory of the human arm about its segmentation points were due to the action of a supra-spinal feedback system that used a forward model which, in turn, provided an error signal that, through an inverse model, resulted in modification of descending commands. These results were found to be robust to changes in model parameters, including inertial properties of the limb, muscle visco-elastic properties, joint stiffness, and gains of the spinal and supra-spinal control loops.

A number of previous theoretical studies have represented the architecture of the biological controller via either an adaptive inverse model based system (Kawato 1989; Shadmehr 1990; Katayama and Kawato 1993; Barto et al. 1998; Schweighofer et al. 1998), or an adaptive forward model based system (Miall et al. 1993; Darlot et al. 1996). Our results show that an architecture that is more parsimonious with our experimental data is one where both systems play a role in the generation of descending motor commands.

Why should the CNS have both a forward and an inverse model of dynamics of a system? If we assume that motor memory is a collection of internal models, then an important problem facing the controller is that of selecting an appropriate model when the hand comes into contact with the environment. If a collection of both forward and inverse models are present, then the behavior of the arm in the environment could be immediately compared to the outputs of a collection of forward models (Wolpert and Kawato 1998). The model that most closely predicted the behavior of the arm could then be used for identification of the appropriate inverse model. In contrast, if motor memory consisted of only inverse models, each would have to be tried out by actually controlling the arm using that model and measuring its performance. It appears that a controller

that used both models would have a distinct advantage in unstructured environments.

5.1 Rates of learning of the two models

From an adaptive control perspective, an interesting prediction can be made regarding the behavior of the proposed control system of Fig. 11. When this controller needs to make reaching movements in a novel force field, performance is much more sensitive to changes in the forward model than in the inverse model (Fig. 15). If there were limited resources available for adaptation, it would be more important to quickly adapt the forward model than to adapt the inverse model (obviously, adaptation in both models is required to completely eliminate errors). Does this actually happen when subjects are learning force fields? We found that when the dynamics of the arm changed from the null field to B_1 , subjects appeared to learn a forward model at a rate that was approximately the same as the rate of learning in their inverse model (Fig. 16). However, when the dynamics changed from B_1 to B_2 (a magnitude of change that resulted in significantly more errors in performance, possibly straining learning resources), there was a tendency for a more rapid adaptation of the forward model than the inverse model (Fig. 18). In this situation, the forward model appeared to adapt at approximately five times the rate of the inverse model.

This can only be taken as preliminary evidence because, for the reaching movements considered here, the ability to estimate the rate of adaptation in the inverse model was hampered by the relative insensitivity of movement parameters to changes in this part of the adaptive controller. Nevertheless, it is intriguing that the data from our subjects shows that the forward model adapts more rapidly than the inverse model when there are very large errors in performance. This is in agreement with a crucial prediction of the theory that suggests that the CNS may rapidly learn a forward model in order to use the acquired knowledge for off-line learning of an inverse model (Miall and Wolpert 1996). Further experiments are necessary to determine whether the adaptive state of the inverse model changes during this off-line period. However, we note that, in the current experiment, we found evidence in support of the idea that passage of time during an off-line period affected states of the models; the rates of learning in both models substantially accelerated when the memory of field B_1 had been allowed to consolidate before B_2 was introduced.

5.2 Noise sensitivity of the forward model

From a control perspective, the use of a forward model can lead to serious problems because estimating the current state of a non-linear system from delayed sensory feedback depends strongly on the quality of the sensory information. Actions taken based on erroneous sensory data might easily destabilize the

system. In effect, the further in time we need to estimate a non-linear process, the more sensitive we might be to its initial conditions. However, here we found the surprising result that the forward model was, in fact, quite accurate even when there were larger errors in state measurements (of velocity). This, we hypothesize, is because of the nature of the descending commands to the distal system. Whereas in a robotic system the descending commands might represent a torque about a joint, in the biological arm, the descending commands loosely correspond to force fields, i.e., torques that are position and velocity dependent. These neural commands drive the system toward an implicit equilibrium position and velocity (Mussa-Ivaldi and Giszter 1992). If a forward model can accurately describe the dynamics of the arm, then a copy of the neural commands will similarly drive the estimated dynamics of the system toward this equilibrium position and velocity, providing robustness to errors in measurements of state. Therefore, it appears that a forward model of the biological arm can function adequately despite significant noise in proprioceptive velocity sensors.

5.3 Neural representation of internal models

The current report suggests that acquiring a motor skill likely involves learning two different kinds of computational models, and that the states of these models may be influenced by the passage of time. What regions of the brain might be involved in the implementation of the forward and inverse models? Many reports have speculated that there may be a role for the cerebellum in representing the forward (Miall et al. 1993), inverse (Kawato and Gomi 1992; Shidara et al. 1993; Houk and Wise 1995; Barto et al. 1998; Schweighofer et al. 1998), or both models (Wada and Kawato 1993; Miall and Wolpert 1996; Wolpert et al. 1998b). The cerebellum of the electric fish provides an excellent example of the neural implementation of a forward model (Bell et al. 1997). In primates, there are also some experimental data to support a role for the cerebellum in representation of the inverse and forward models. For example, simple spike activity of Purkinje cells during generation of eye and arm movements can be closely fitted to an inverse dynamics representation of the movements of the eye (Gomi et al. 1998) and the hand (Ebner and Fu 1997). Alternatively, if the cerebellum is involved in representation of a forward model, then during the delay period before initiation of a movement, simple spike activity may represent the predicted sensory outcome of a movement (Miall et al. 1998). If this predicted outcome differs from the desired behavior, complex spikes representing predicted error should be detected soon afterwards. In fact, there is evidence for this temporal order in the firing activity of some Purkinje cells (Miall et al. 1998).

Lesion studies in humans, however, are not currently in agreement with the view that the forward model is represented in the cerebellum. To show this, in Fig. 11 consider a switch that would allow for cutting off of

descending commands, N_c , to the spinal cord. In this situation, the supra-spinal controller can generate the neural commands (output of the inverse model) and observe its sensory consequences (output of the forward model) without actually making the movement. This is one way by which a forward model can be used to mentally simulate a movement. In fact, humans are able to accurately predict the consequences of imagined movements of the hand (Sirigu et al. 1996). Damage to the parietal cortex adversely affects this ability (Sirigu et al. 1996), while damage to the motor cortex (Sirigu et al. 1996) and the basal ganglia (Dominey et al. 1995) does not. This would suggest that the parietal cortex must be involved in some aspect of using the output, or storing the contents, of the forward model. Indeed, a recent report of a parietal patient with deficits in maintaining state estimates of his arm supports this view (Wolpert et al. 1998a). Accordingly, if motor imagery is affected in the parietal patients because a pathway for interpreting the output of the forward model has been lost, then perhaps the forward model resides in the cerebellum and its output is provided to the parietal cortex via thalamic pathways. However, a recent report has shown that patients with cerebellar lesions can predict consequences of imagined arm movements as accurately as normal subjects (Kagerer et al. 1998). Therefore, this casts some doubt on the idea that, in humans, the forward model is stored in the cerebellum. Current lesion studies only implicate the parietal cortex in the neural machinery that might represent or use the forward model.

Our observation in this report has been that learning of reaching movements likely involves adaptation of two distinct internal models, possibly at different rates. Let us consider the possibility that one of these models, perhaps the inverse model, resides in the cerebellum, while the other does not. Note that if this is the case, when there is damage to the neural representation of the inverse model, some improvement in performance will still take place through adaptation of the forward model. However, if the damage to the neural representation of the inverse model is temporary, when it comes back on-line the performance of the system would suddenly decline, requiring relearning. In fact, it has been observed that cats can learn to move a manipulandum and retain this motor memory despite inactivation of their cerebellar nuclei (Wang et al. 1998). However, they have to relearn the task when the cerebellum comes back on-line. It is possible that this occurs because of a mismatch between an adapted forward model that might reside outside the cerebellum and the inverse model in the cerebellum.

A glance at the design of the forward model suggests that a number of functions must be closely synchronized in order for the forward model to learn the dynamics of a force field. Perhaps the most important function is to have at any time t a copy of the efferent commands from time $t - t_0$ until t . This would allow the forward model to estimate the current position of the arm by integrating the descending commands from an initial condition set by the sensory measurements taken at time $t - t_0$. In effect, there should be a place where a copy of the

descending commands is held for up to 200 ms. While short-term retention of sensorimotor information is often associated with the dorsolateral prefrontal cortex (Fuster 1997), there is no evidence to suggest that damage to this part of the brain results in a movement disorder. The neural basis of this very short-term motor memory remains to be explored.

We have demonstrated here that the human motor-control system has an architecture that likely includes both forward and an inverse models. When there are large errors in performance, improvements during training occur because of rapid adaptation in the forward model but a much slower rate of adaptation in the inverse model.

Acknowledgements. This work was part of a Masters Thesis submitted by N.B. to the Biomedical Engineering Department at Johns Hopkins University. The thesis, which includes more complete details of the model parameters and further experimental results, is available from www.bme.jhu.edu/~reza/nbthesis.pdf. We are very grateful to Kurt Thoroughman who provided us with the human subject data used in this report. Our work has been greatly enriched by our interactions with Maurice Smith, Kurt Thoroughman, and Steve Wise. This work was funded in part by grants from the U.S. Office of Naval Research, the Whitaker Foundation, and a Multi-University Research Initiative from the U.S. Department of Defense.

References

- Astrom KJ, Wittenmark B (1984) Computer controlled systems. Prentice-Hall, Englewood Cliffs, NJ
- Atkeson CG (1989) Learning arm kinematics and dynamics. *Annu Rev Neurosci* 12:157–183
- Barto AG, Fagg AH, Sitkoff N, Houk JC (in press) A cerebellar model of timing and prediction in the control of reaching. *Neural Comput*
- Bell CC (1981) An efference copy which is modified by reafferent input. *Science* 214:450–453
- Bell CC, Han VZ, Sugwara Y, Grant K (1997) Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Science* 287:278–281
- Blakemore SJ, Goodbody SJ, Wolpert DM (1998) Predicting the consequences of our own actions: the role of sensorimotor context estimation. *J Neurosci* 18:7511–7518
- Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. *Nature* 382:252–255
- Cordo P, Carlton L, Bevan L, Carlton M, Kerr GK (1994) Proprioceptive coordination of movement sequences: role of velocity and position information. *J Neurophysiol* 71:1848–1861
- Darlot C, Zupan L, Etard O, Denise P, Maruani A (1996) Computation of inverse dynamics for the control of movements. *Biol Cybern* 75:173–186
- Dominey P, Decety J, Broussolle E, Chazot G, Jeannerod M (1995) Motor imagery of a lateralized sequential task is asymmetrically slowed in hemi-Parkinson's patients. *Neuropsychologia* 33:727–741
- Ebner TJ, Fu, Q (1997) What features of visually guided arm movements are encoded in simple spike discharge of cerebellar Purkinje cells? *Prog Brain Res* 114:431–447
- Flanagan JR, Wing AM (1997) The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J Neurosci* 17:1519–1528
- Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703
- Fuster JM (1997) The prefrontal cortex: anatomy, physiology, and neuropsychology of the frontal lobe. Lippincott-Raven, Philadelphia
- Gielen CCAM, Houk JC (1987) A model of the motor servo: Incorporating nonlinear spindle receptor and muscle mechanical properties. *Biol Cybern* 57:217–231
- Gomi H, Kawato M (1992) The cerebellum and VOR/OKR learning models. *Trends Neurosci* 15:445–453
- Gomi H, Kawato M (1996) Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272:117–120
- Gomi H, Shidara M, Takemura A, Inoue Y, Kawano K, Kawato M (1998) Temporal firing patterns of Purkinje cells in the cerebellar ventral paraflocculus during ocular following responses in monkeys I. Simple spikes. *J Neurophysiol* 80: 818–831
- Gribble PL, Ostry DJ, Sanguineti V, LaBoissiere R (1998) Are complex control signals required for human arm movement? *J Neurophysiol* 79:1409–1424
- Hanneton S, Berthoz A, Droulez J, Slotine JJE (1997) Does the brain use sliding variables for the control of movements? *Biol Cybern* 77:381–393
- Houk JC, Wise SP (1995) Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cerebral Cortex* 5:95–110
- Johansson RS, Westling G (1984) Roles of glabrous skin receptors and sensorimotor memory in automatic-control of precision grip when lifting rougher or more slippery objects. *Exp Brain Res* 56:550–564
- Jordan MI, Rumelhart DE (1992) Forward models: supervised learning with a distal teacher. *Cogn Sci* 16:307–354
- Jordan MI, Flash T, Arnon Y (1994) A model of the learning of arm trajectories from spatial deviations. *J Cogn Neurosci* 6:359–376
- Kagerer FA, Bracha V, Wunderlich DA, Stelmach GE, Bloedel JR (1998) Ataxia reflected in the simulated movements of patients with cerebellar lesions. *Exp Brain Res* 121:125–134
- Karniel A, Inbar GF (1997) A model for learning human reaching movements. *Biol Cybern* 77:173–183
- Katayama M, Kawato M (1993) Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. *Biol Cybern* 69:353–362
- Kawato M (1989) Adaptation and learning in control of voluntary movement by the central nervous system. *Adv Robot* 3:229–249
- Kawato M, Gomi H (1992) A computational model of four regions of the cerebellum based on feedback-error learning. *Biol Cybern* 68:95–103
- Krylow AM, Sandercock TG, Rymer WZ (1995) Muscle models. In: Arbib, MA (ed) *The handbook of brain theory and neural networks*. MIT Press, Cambridge, mass, pp 609–613
- Lee RG, Tatton WG (1975) Motor responses to sudden limb displacements in primates with specific CNS lesions and in human patients with motor system disorders. *Can J Neurol Sci* 2:285–293
- Lohmiller W, Slotine JJE (1996) On metric observers for nonlinear systems. *Proc IEEE Int Conf Control App* 320–326
- Massaquoi SG, Slotine, J-JE (1996) The intermediate cerebellum may function as a wave-variable processor. *Neurosci Lett* 215:60–64
- Miall RC, Wolpert DM (1996) Forward models for physiological motor control. *Neural Netw*, 9:1265–1279
- Miall RC, Weir DJ, Wolpert DM, Stein JF (1993) Is the cerebellum a Smith predictor? *J Mot Behav* 25:203–216
- Miall RC, Keating JG, Malkmus M, Thach WT (1998) Simple spike activity predicts occurrence of complex spikes in cerebellar Purkinje cells. *Nat Neurosci* 1:13–15
- Montgomery JC, Bodznick D (1994) An adaptive filter that cancels self-induced noise in the electrosensory and lateral line mechanosensory systems of the fish. *Neurosci Lett* 174:145–148

- Murray WM, Delp SL, Buchanan TS (1995) Variation of muscle moment arms with elbow and forearm position. *J Biomech* 28:513–525
- Mussa-Ivaldi FA, Giszter, SF (1992) Vector field approximation: a computational paradigm for motor control and learning. *Biol Cybern* 67:491–500
- Mussa-Ivaldi FA, Hogan N, Bizzi E (1985) Neural, mechanical and geometric factors subserving arm posture in humans. *J Neurosci* 5:2732–2743
- Niemeyer G, Slotine JJE (1991) Stable adaptive teleoperation. *IEEE J Oceanic Eng* 16:152–162
- Ronco E (In press) Open-loop intermittent feedback optimal control: a probable human motor control strategy. *IEEE Trans Man Machine Cybern*
- Rothwell JC (1990) Long latency reflexes of human arm muscles in health and disease. In: Rossini, PM, Mauguiere F (ed) *New trends and advanced techniques in clinical neurophysiology*. Elsevier, Amsterdam, pp 251–263
- Sanes JN, Shadmehr R (1995) Sense of muscular effort and somesthetic afferent information in humans. *Can J Physiol Pharmacol* 73:223–233
- Schweighofer N, Arbib MA, Kawato M (1998) Role of the cerebellum in reaching movements in humans. I. Distributed inverse dynamics control. *Eur J Neurosci* 10:86–94
- Shadmehr R (1990) Learning virtual equilibrium trajectories for control of a robot arm. *Neural Comput* 2:436–446
- Shadmehr R, Arbib MA (1992) A mathematical analysis of the force-stiffness characteristics of muscles in control of a single joint system. *Biol Cybern* 66:463–477
- Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. *J Neurosci* 17:409–419
- Shadmehr R, Holcomb HH (1997) Neural correlates of motor memory consolidation. *Science* 277:821–825
- Shadmehr R, Holcomb HH (1999) Inhibitory control of motor memories: a PET study. *Exp Brain Res* (in press)
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224
- Shidara M, Kawano K, Gomi H, Kawato M (1993) Inverse dynamics model eye movement control by purkinje cells in the cerebellum. *Nature* 365:50–52
- Sirigu A, Duhamel J-R, Cohen L, Pillon B, Dubois B, Agid Y (1996) The mental representation of hand movements after parietal cortex damage. *Science* 273:1564–1568
- Sperry RW (1950) Neural basis of spontaneous optokinetic response produced by visual inversion. *J Comp Physiol Psychol* 43:482–489
- Stroeve S (1997) A learning feedback and feedforward neuromuscular control model for two degrees of freedom human arm movements. *Hum Move Sci* 16:621–651
- Wada Y, Kawato M (1993) A neural network model for arm trajectory formation using forward and inverse dynamics models. *Neural Netw* 6:919–932
- Wang JJ, Shimansky Y, Bracha V, Bloedel JR (1998) Effects of cerebellar nuclear inactivation on the learning of a complex forelimb movement in cats. *J Neurophysiol* 79:2447–2459
- Wolpert DM, Miall RC, Kerr GK, Stein JF (1993) Ocular limit cycles induced by delayed retinal feedback. *Exp Brain Res* 96:173–180
- Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. *Neural Netw* 11:1317–1329
- Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. *Science* 269:1880–1882
- Wolpert DM, Goodbody SJ, Husain M (1998a) Maintaining internal representations: the role of the human superior parietal lobe. *Nat Neurosci* 1:529–533
- Wolpert DM, Miall RC, Kawato M (1998b) Internal models in the cerebellum. *Trends Cogn Sci*. 2:338–347
- Zajac FE (1989) Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Crit Rev Biomed Eng* 17:359–411