

HUMAN MOTOR LEARNING IN STATIONARY AND NONSTATIONARY NOVEL DYNAMIC  
ENVIRONMENTS: PSYCHOPHYSICAL, ELECTROMYOGRAPHIC, AND COMPUTATIONAL  
VERIFICATION AND EXTENSION OF THE INVERSE MODEL HYPOTHESIS

by

Kurt Andrew Thoroughman

A dissertation submitted to The Johns Hopkins University in conformity with the requirements for the  
degree of Doctor of Philosophy

Baltimore, Maryland

December, 1999

© Kurt Andrew Thoroughman 1999

All rights reserved

## Abstract

In this thesis I investigated the computational nature of human motor learning of novel dynamics. In my experiments human subjects learned to make reaching movements while grasping a robotic arm. After learning a baseline task, subjects trained to move the robot while it produced viscous forces. The ability of humans to control their arms while interacting with environments with differing dynamic properties is computationally impressive. Recent studies have suggested that a neural implementation of an inverse model of dynamics likely underlies this ability. Psychophysical and theoretical studies have suggested that inverse models exist, are spatially focused, and could be computed using mathematical primitives termed basis functions.

I measured neuromotor output using electromyography to investigate whether muscle activations themselves supported the inverse model hypothesis. Two novel metrics of EMG suggested that subjects did learn an inverse model specific to dynamics, and that the inverse model was likely trained via an integration of feedback into feedforward control.

I also tested an assumption of previous studies, that desired trajectories remain constant in different environments. I found a small but significant difference in asymptotic behaviors in the null and force fields. Comparison of performance to simulated optimal trajectories indicated that while dynamics were considered, kinematic influences far outweighed dynamics in movement planning.

Previous investigations were also limited to study of steady-state performance after dynamics were well trained. Investigators sometimes surprised subjects with quick dynamic changes, but the effect of these changes on future performance was not studied. My subjects who trained in nonstationary dynamic environments demonstrated unlearning with every surprise dynamic change. These changes provided a unique window into movement-by-movement learning of dynamics. A state-space model fit to the adaptation patterns provided an estimate of the spatial extent of unlearning, a motor learning receptive field. This receptive field revealed that sudden changes in dynamics had opposing effects on subsequent behavior, dependent on the direction of movement in which dynamics changed. Theoretical and computational analysis showed that such a receptive field could result from broad, sparse velocity-dependent Gaussian basis functions.

## **Preface**

I thank Reza Shadmehr for his leadership, support, and generous sharing of ideas as my mentor. I also thank Maurice Smith for his constant and critically important insights and for being a true partner in my work. Thank you many times over to my family for their love and care. Thank you to the friends I have met through Hopkins, especially Drs. Larry Schramm, Amy Brisben, David Kanter, and Robert Fasciano. And final and most important thanks to my amazingly wonderful wife, Elizabeth Pickard, with whom life is extraordinary and without whom this work would have been impossible.

# Table of Contents

**Abstract ii**

**Table of Contents iv**

**Table of Figures vi**

**Table of Tables viii**

**Chapter 1 Introduction: Psychophysical, physiological, and theoretical explorations for the computational and neural bases of human motor learning 1**

**Chapter 2 EMG correlates of inverse model formation during training in stationary dynamics 7**

*2.1 Introduction 7*

*2.2 Materials and Methods 8*

*2.2.1 Experimental apparatus and protocol 8*

*2.2.2 Displacement analysis 11*

*2.2.3 EMG normalization and polar analysis 11*

*2.2.4 Computational modeling: structure of the model 12*

*2.2.5 Computational modeling: predicting rotation of directional bias 15*

*2.2.6 Computational modeling: predicting time series of appropriate activations 16*

*2.2.7 Using model predictions to remap subjects' EMG 20*

*2.2.8 Wasted contraction: a measure of co-contraction 22*

*2.2.9 Calculation of force produced by the human controller 22*

*2.2.10 Bootstrapping 23*

*2.3 EMG correlates of learning: a single direction of movement 23*

*2.4 EMG correlates of learning across all directions of movement 25*

*2.5 Rotation of angular dependence of forces and EMG 30*

*2.5.1 Simulation results 30*

*2.5.2 Rotation of angular dependence of EMG with training in B1 31*

*2.6 Activation function rotation is specific to the force field 33*

*2.7 Activation function rotation during training in B2 35*

*2.8 Anterior deltoid activation functions in the null field, B1, and B2 37*

*2.9 Reduction of wasted contraction during practice in null field and in B1 39*

*2.10 Discussion 40*

**Chapter 3 Precurvature and cost minimization 45**

*3.1 Introduction 45*

*3.2 Asymptotic behavior in null and viscous force fields 49*

*3.3 Calculation of minimum torque change trajectories 53*

*3.4 Perpendicular force profiles: actual dynamics of force field trajectories and hypothetical dynamics of null field trajectories 57*

*3.5 Torque change and jerk costs of individual movements 61*

*3.6 Neural-noise induced variance and precurvature 66*

*3.7 Discussion 67*

## **Chapter 4 Human motor learning of a nonstationary dynamic environment 71**

*4.1 Introduction 71*

*4.2 Kinematic performance during training with catch trials 74*

*4.2.1 Performance over the duration of training 74*

*4.2.2 Movement-to-movement kinematic effects of catch trials 77*

*4.3 Electromyographic correlates of catch trial-induced unlearning 81*

*4.4 Modeling force field learning and unlearning using a state-space representation with scalar input 84*

*4.5 Modeling force field learning and unlearning using a state-space representation with vector input 88*

*4.6 Theoretical relationship between map of error into subsequent performance and inverse model basis functions 94*

*4.7 Specifying basis function properties using parameters of vector-input state-space model 97*

*4.8 Modeling and simulating unlearning receptive fields using velocity-dependent Gaussian basis functions 102*

*4.9 Learning constant force fields using velocity-dependent Gaussian basis functions 106*

*4.10 Discussion 108*

## **Chapter 5 Discussion: Computational and neural bases of human motor learning 113**

### **Bibliography 118**

### **Curriculum Vitae 126**

## Table of Figures

Figure 1.1: Block diagram summarizing current theories of the computation underlying human motor control.....	2
Figure 2.1: Experimental apparatus.....	9
Figure 2.2: Data used for building a simple force activation model for the anterior deltoid and data used for testing the validity of the model .....	17
Figure 2.3: Hand paths and EMG recorded during movements.....	24
Figure 2.4: Perpendicular displacement and subjects' $a_m$ (time-averaged agonist burst EMG) during movements toward 135.....	25
Figure 2.5: Subjects' composite EMG traces and forces.....	26
Figure 2.6: The composite EMG trace appropriate for field B1 shifts with training .....	28
Figure 2.7: Rotations in muscle-tuning curves as predicted from a computational model that assumes adaptation of an IM and tuning curves as recorded from subjects' EMG.....	31
Figure 2.8: Perpendicular displacement and rotation of $a_m$ resultant vectors during movements in force field B1.....	32
Figure 2.9: Perpendicular displacement of the hand and rotation of $a_m$ resultant vectors in subjects who, after training in B1 , trained in the null field 3 min later.....	34
Figure 2.10: Perpendicular displacement and rotation of $a_m$ resultant vectors in subjects who, after training in B1, trained in B2 3 min or 6 hr later.....	36
Figure 2.11: Orientations of anterior deltoid resultant vectors during training in the null field, B1 , and B2.....	38
Figure 2.12: Wasted contraction during movements in null field and B1.....	40
Figure 3.1: Mean parallel and perpendicular displacements in all eight directions of movement .....	49
Figure 3.2: Trajectories averaged across directions.....	51
Figure 3.3: Perpendicular displacement 150 msec into the movement for null field and force field movements.....	52
Figure 3.4: Parallel velocities of minimum jerk and minimum torque change movements in the null field and in force field B1 .....	54
Figure 3.5: Perpendicular velocities of minimum jerk and minimum torque change movements in the null field and in force field B1 .....	55
Figure 3.6: Handpaths of minimum jerk and minimum torque change movements in the null field and in force field B1 .....	57
Figure 3.7: Perpendicular force produced in actual and hypothetical (null field) trajectories in force field.....	59
Figure 3.8: Derivative of perpendicular force produced in actual and hypothetical (null field) trajectories in force field .....	60
Figure 3.9: Cumulative normalized jerk during training in the force field B1 and the null field set in each direction of movement and the mean across directions.....	62
Figure 3.10: Comparing null field and third set of force field B1 movements in perpendicular and parallel velocity, and raw and normalized jerk .....	64
Figure 3.11: Cumulative normalized torque change during training in the force field B1 and the null field set in each direction of movement and the mean across directions .....	65
Figure 3.12: Mean and 95% confidence intervals of endpoints of simulated movements, produced by noisy agonist and antagonist torque actuators.....	67
Figure 4.1: Perpendicular displacement 250 msec into movement; mean of constant force group, catch trial group, and preceding or following catch trials.....	74
Figure 4.2: Perpendicular displacement 150 msec into movement; preceding or following catch trials, compared to nul field movements .....	76
Figure 4.3: Handpaths of movements toward 90 degrees in single subject and mean across subjects; null field, first exposure in force field, before, during, and after catch trial .....	78
Figure 4.4: Perpendicular displacement in catch trial movements and change in perpendicular displacement across catch trial movements .....	79
Figure 4.5: Latency of post-catch-trial increase in perpendicular displacement.....	80
Figure 4.6: Timing of B1-appropriate EMG activation during training .....	81

Figure 4.7: B1-appropriate EMG in movements before (ct-1) and after (ct+1) catch trials.....	82
Figure 4.8: Rotation of movement-initiating EMG resultant vectors from null field orientations, before, during, and after catch trials .....	83
Figure 4.9: Perpendicular displacement 250 msec into movements toward 90 degrees, fitted by scalar-input and vector-input state-space models.....	86
Figure 4.10: Parameters of the scalar-input state space model.....	87
Figure 4.11: State-space model vectors b in each modeled direction of movement .....	91
Figure 4.12: The mean across directions of the state-space parameter vector b .....	93
Figure 4.13: The mapping of torque error into torque compensation .....	101
Figure 4.14: Estimating receptive fields using snapshot evaluations of basis functions.....	103
Figure 4.15: Simulation of catch trial learning using inverse model with Gaussian bases .....	104
Figure 4.16: Response of simulated learning using inverse model with Gaussian basis functions ...	107

## Table of Tables

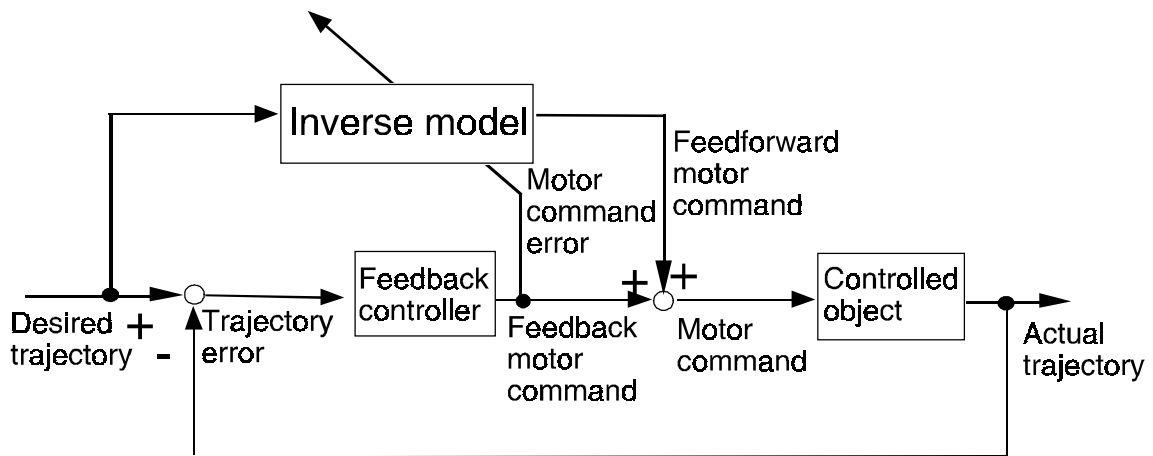
<b>Table 2.1: Computational representation of an internal model in terms of expected changes in the angular orientation of EMG.....</b>	<b>16</b>
<b>Table 2.2: Values of fitted parameters <math>p</math> in Equation 2.7.....</b>	<b>19</b>
<b>Table 2.3: Values of the vector operator <math>v</math>.....</b>	<b>20</b>
<b>Table 4.1: Correlation between simulation and actual data, of force field movement perpendicular displacement and receptive fields, as a function of the simulation's basis function width.....</b>	<b>107</b>



## **Chapter 1 Introduction: Psychophysical, physiological, and theoretical explorations for the computational and neural bases of human motor learning**

People are typically unimpressed with their ability to not only move their arms in almost any desired trajectory, but also to skillfully move both familiar and novel objects which have different dynamic properties. The neuroscientific, computational, and robotic communities were also not appropriately impressed with this human skill, until robotic control researchers attempted to control multijoint manipulanda (robotic arms). They discovered that even the simplest control of multijoint arms is fraught with complexity. Since the dynamics of multijoint systems are coupled; dynamic commands between all joints must therefore be coordinated (Atkeson, 1989). The successful completion of movement requires a transformation between many coordinate systems, including spaces of vision, joints, and the hand (Bizzi et al, 1991; Wolpert, 1997). From a control perspective, perhaps the worst complexity of multijoint arm movement lies in its strong nonlinearities. Only recently have theorists discovered complex mathematical routines to perform stable adaptive control of nonlinear dynamic systems (Narendra and Annaswamy, 1989; Slotine and Li, 1991). In the late 1980s, researchers questioned whether human nervous systems could perform the complex computations required to generate motor commands informed by precise kinematics and dynamics (Hildreth and Hollerbach, 1987). In the next twelve years, investigators have theoretically, psychophysically, and physiologically searched for the neural computational basis of human motor learning.

Figure 1.1 summarizes several current theories regarding the human motor control and learning of arm movements (taken from Wolpert et al, 1998). The “controlled object” refers to the excitation-contraction coupling of muscles, the dynamics of the arm, and any interaction from the environment. The collective system of muscles and manipulated dynamics is often referred to as the “plant.” The plant then transforms a neural motor command into the trajectory (position and velocity) of the arm.



**Figure 1.1: Block diagram summarizing current theories of the computation underlying human motor control. Taken from Wolpert et al (1998).**

One technique humans could utilize to lessen computational complexity would be to rely primarily on feedback control. Feedback control of the human arm produces torque with stiffness and viscosity, in response to differences between desired and actual hand position and velocity. A feedback control system would track desired trajectories robustly with little computational cost (Arbib, 1995). Visual and proprioceptive signals provide inputs to short- and long-loop feedback signals that could drive feedback control. A very simple strategy, then, would be to generate a desired trajectory that slightly leads the actual trajectory, with the feedback controller pulling the arm along. This simple control forms the core of an early theory termed the equilibrium point hypothesis (Flash, 1987). Both sensory systems have inherent delays, however, of up to 100 msec. In spite of these delays, humans are able to make rapid, accurate movements (Miall, 1995). Furthermore, when learned dynamics are unexpectedly removed, subjects generate errors in movement that reject the notion that movement is generated through stiffness and viscosity alone (Lackner and DiZio, 1994; Shadmehr and Mussa-Ivaldi, 1994). Humans can, however, respond to unexpected dynamics and correct their trajectories mid-movement (Gottlieb, 1996; Sainberg et al, 1999). Human motor control must therefore combine feedback control, responding to information during a movement, and feedforward control, generating commands before sensory feedback is available.

In order to generate feedforward commands, the human CNS must predict the dynamic demands of the movement. The prediction of dynamics could have relatively low computational cost if the mapping

between desired trajectory and required motor commands were stored in a lookup table. The continuous range of positions and velocities accessible by multiple joints, however, would demand a lookup table far too large to be physiologically plausible (Mussa-Ivaldi, 1999). The full spectrum of motor behaviors, then, likely stems from a functional, systematic mapping of desired movement states into motor commands, reflecting the dynamic demands of particular movements. This hypothesized unit of neural computation inverts the transformation carried out by the plant; hence this unit is termed an inverse dynamic model, or simply an inverse model. Such a systematic mapping would, unlike a lookup table representation, continuously interpolate between and extrapolate from experienced movement states.

Recent psychophysical studies have investigated the spatial and temporal properties of the internal model by training subjects in novel dynamic environments. The kinematic features of movements produced by subjects who train in viscous force field correlate well with a biomechanical simulation of movements generated by a velocity-based inverse model. When trained subjects move in novel regions of the workspace, dynamics seemed to generalize in angular, rather than Cartesian, coordinates (Shadmehr and Mussa-Ivaldi, 1994). Subjects generalized well across different speeds and velocities in the same direction as dynamic training (Goodbody and Wolpert, 1998), but generalization trailed off rapidly when subjects are tested in different movement directions (Gandolfo et al, 1996). Dynamic information seemed to generalize in terms of velocities even when the actual structure of learned dynamics depends on time, not state (Conditt and Mussa-Ivaldi, 1999). The adaptation of an inverse model to one dynamic environment also seems to be susceptible to anterograde and retrograde interference: exposure to a second anti-correlated environment is harder to learn and damages previous learning more prominently if the two environments are closely spaced in time (Brashers-Krug et al, 1996; Shadmehr and Brashers-Krug, 1997). All of these studies measured kinematics of performance to infer properties of the inverse model.

Movement kinematics, however, are a product of the plant's transformation of the actual output of the inverse model, the feedforward motor command. This motor command can be measured using surface electromyography (EMG) (Basmajian and De Luca, 1985; De Luca, 1997). EMG traces measure the output of motor neurons that drive muscle contractions. Quantification of neural output has elucidated how the CNS generates appropriate muscle forces in the multi-joint arm in many directions of free movement

(Flanders, 1991; Karst and Hasan, 1991) or isometric force resistance (Flanders and Soechting, 1990). Recording of EMG has also quantified changes in motor output while subjects learn to make single-joint movements against novel loads (Gottlieb, 1994, 1996). Single-joint movements, however, do not excite dynamics of the complexity that inspired the inverse model hypothesis. A study of EMG during dynamic learning of multi-joint movements would directly test the predictions of inverse model existence stemming from theory and psychophysics.

The inferences regarding inverse model properties in previous studies were also limited by the assumptions of investigators regarding learning and training stationarity. Metrics employed in previous studies focused on learning over the course of many minutes of training. These methods elucidated the steady state response, but did not explain the transient response of the human motor learning system to sudden changes in dynamics. Some studies did in fact interleave sudden changes of dynamics into training in a baseline environment (Shadmehr and Mussa-Ivaldi, 1994; Goodbody and Wolpert, 1998; Sainberg et al, 1999). These investigators, however, studied only how subjects responded during the surprise dynamics, not how those surprises affected subsequent movements back in the baseline environment. Just as impulse functions reveal important characteristics in systems engineering, sudden changes in dynamic environments may create the interrogation necessary to understand how inverse models adapt.

Although previous psychophysical studies have made important inferences regarding the natural of inverse dynamic models, the computational basis of the inverse dynamic computation remains largely unknown (Mussa-Ivaldi, 1999). The mapping between desired trajectories and motor commands appropriate for particular dynamics could be framed as fitting an unknown function between kinematic input and neuromuscular output. A computationally efficient method used for function fitting represent the unknown function as a sum of coefficients multiplying regular functions termed basis functions (Poggio and Girosi, 1990). These basis functions are typically radial, peaking at a center and decaying with distance away from the center. The variance of the function across the input space is therefore represented by different coefficients multiplying the basis functions located at different centers. The basis function representation is identical to the input/output mapping generated by layered neural networks (Poggio and Girosi, 1990). Presuming a neural network architecture of dynamic learning, basis functions of inverse models would

depend on the desired trajectory of each movement. Learning novel dynamics would require the update of the weights that multiply those basis functions (Jordan and Rumelhart, 1992; Sanner and Slotine, 1995; Sanner and Kosha, 1999). The response of the inverse model to sudden, frequent changes in dynamics could cause varied weight changes; the effect of these weight changes on subsequent motor output and behavior could reveal important properties of the hypothesized inverse model basis functions.

The functions of both feedback and feedforward control depend on desired trajectories as input (Figure 1.1). Movement goals are specified in a discrete number of dimensions, whereas trajectories are specified as a continuously (or at least densely represented) function of time (Wolpert, 1997). Human motor control therefore requires a transformation between movement goals and desired trajectories. Since the goal-to-trajectory mapping is many-to-one, the human CNS must use some other constraints to specify trajectories. A hypothesized class of constraints posits that humans generate movements that optimize performance, that minimize some undesirable characteristic of movement. Two undesirable costs incurred in movement are movement jerk, the third time derivative of position (Flash and Hogan, 1985), and the time derivative of torque change (Kawato et al, 1987). Minimizing jerk and torque change would maximize smoothness in, respectively, kinematics and dynamics. Natural movement profiles seemed to be explained well by both hypotheses (Desmurget et al, 1998), although specific tasks seem to be fit better by minimizing torque change (Osu et al, 1997). The altered behavior of subjects who are given perturbed visual feedback, however, indicates that minimizing visually perceived jerk is of primary concern (Wolpert et al, 1995b; Ghahramani et al, 1996). An alternate hypothesis, recently proposed, posits that humans generate desired trajectories in order to minimize error created by noise in the neural motor command (Harris and Wolpert, 1998). Quantification of theoretical and actual costs incurred in kinematics, dynamics, and accuracy during movements in a novel dynamic environment might provide insight regarding the relative relevancy of these costs to movement planning.

In this thesis I investigated the properties of the hypothesized inverse model of dynamics. Human subjects trained in novel dynamic environments, produced by a robotic arm. In Chapters 2 and 3, all subjects practiced movements in a consistently applied viscous force field. In Chapter 2, EMG is used to record motor neuronal output from subjects learning the force field. I developed novel metrics of EMG to

facilitate a comparison of neural output to an anthropomorphic biomechanical model which learned via updating an inverse model. The electromyographic metrics confirmed the existence of an inverse model, revealed how feedback control could instruct the feedforward controller, and indicated how learning an internal model could correlate with changing motor cortical activity.

In Chapter 3 I examined the asymptotic behavior of these subjects to investigate the relative importance of neural noise and kinematic and dynamic optimization in movement planning. I discovered that trajectories converge upon a more curved trajectory in the force field than without the force field, suggesting that humans do at least consider dynamic optimization in planning movements. Computational consideration of jerk, torque change, and neural noise in the force field, coupled with the actual costs incurred by subjects during movements, suggest that subjects first and foremost consider kinematic optimization during movement planning.

In Chapter 4 I examined the kinematic and electromyographic response of subjects who trained in a different paradigm, during which surprise movements without the force field were interspersed between force field movements. These “catch trial” movements induced subsequent reversals of improvements in movement accuracy and changes in muscle activation, indicating short-term unlearning of dynamics. A state-space model quantified this unlearning and uncovered its unexpected spatial extent. The parameters of the state space fit defined receptive fields of motor learning. The receptive fields, when considered in the neural network implementation of inverse models, revealed important characteristics of the basis functions underlying dynamic motor learning.

Parts of this dissertation have been published (Thoroughman and Shadmehr, 1999). Other parts of this dissertation rely on retrospective analysis of previously trained subjects. Kinematic parameters of these subjects’ movements were previously published (Shadmehr and Brashers-Krug, 1997). The analysis performed here focused on the effect of catch trial performance on subsequent movements in the force field, which constitutes novel and substantial work beyond the analyses performed by the previous paper’s authors.

## **Chapter 2 EMG correlates of inverse model formation during training in stationary dynamics**

### ***2.1 Introduction***

People learn to move novel objects along desired trajectories, in any direction, by simply practicing the task a few times. This adaptation is remarkable because of the computational complexity inherent to learning dynamics (Atkeson, 1989). Previous studies have hypothesized that adaptation of a neural internal model (IM), transforming desired trajectories of the hand into appropriate muscle activations, likely underlies this ability (Jordan, 1995; Wolpert et al., 1995). Aftereffects, errors that people make when learned dynamics are unexpectedly changed, suggest that IMs are built gradually with practice (Shadmehr and Mussa-Ivaldi, 1994), that learning one IM can interfere with the learning of a second IM (Brashers-Krug et al., 1996), and that the interference fades over the course of hours (Shadmehr and Brashers-Krug, 1997). The evidence for the formation of IMs, however, comes mainly from psychophysics. A measure of neural output, such as electromyography (EMG), could provide insight into the neural basis of the formation of IMs.

Previous work has reported EMG changes during learning of single-joint movements with different loads (Corcos et al., 1993; Gottlieb, 1994). When subjects make elbow movements against an unexpected load, EMG patterns up until 200 msec into the movement remain unchanged compared with unperturbed trials (Gottlieb, 1996). This delayed response to the perturbation reflects the latency of short- and long-loop reflexes that use proprioceptive information to produce an error–feedback action (Marsden et al., 1978). After subjects practice moving the novel load, EMG differs from unperturbed trials from the beginning of movement (Gottlieb, 1994). This change in the movement-initiating EMG reflects the adaptation of descending control commands. In computational studies, the changes in descending commands are attributable to adaptation of an IM (Wada and Kawato, 1993; Miall and Wolpert, 1996; Barto et al., 1998; Bhushan and Shadmehr, 1999). An elegant idea is that adaptation may be driven by error–feedback motor responses generated by reflex circuits (Kawato et al., 1987; Stroeve, 1997). In other words, the delayed, reflex-based error feedback might serve as a “blueprint” for how the CNS needs to change descending commands. To test this idea, the computational concept of an IM for multijoint movements needs to be

described in a way that its formation could be quantitatively tested by changes in EMG. Here I asked whether changes in EMG correlate with the formation of an IM, and whether the error–feedback response drives this learning.


We measured EMG as subjects learned to reach while grasping a manipulandum that created novel forces. I used a biomechanical model to transform EMG into a composite trace, corresponding to activation that specifically compensated for the imposed field. I found that with training, activation of the appropriate musculature gradually shifted from a delayed error– feedback response to a predictive feedforward response. To quantify how these changes corresponded to the formation of an IM, I analyzed the directional tuning functions of movement-initiating EMG. These spatial functions relate EMG to movement direction in hand-centered coordinates (Flanders and Soechting, 1990; Sergio and Kalaska, 1998). Previous reports have demonstrated that reaching movements or isometric force production generate EMGs with broadly tuned directional biases (Flanders and Soechting, 1990; Flanders, 1991; Karst and Hasan, 1991). This modeling revealed that learning of an IM should result in specific rotations of muscles’ tuning curves. I found that with practice, the directional patterns of activation gradually rotated; the final magnitude of the rotations corresponded well with the predicted change. Temporal features and angular orientation are invariant to the overall magnitude of the EMG; therefore, across subjects and across recording sessions, despite removal and reapplication of surface electrodes, these metrics robustly quantified the formation of internal models.

## ***2.2 Materials and Methods***

### **2.2.1 Experimental apparatus and protocol**

Twenty-four right-handed subjects (12 women and 12 men), students in the Johns Hopkins School of Medicine, provided written consent to participate in this study. With their right hand, subjects grasped the end effector (handle) of a 2 degree-of-freedom (df) manipulandum mounted in the horizontal plane (Shadmehr and Brashers-Krug, 1997). Subjects sat in front of the manipulandum, with the right arm supported by a sling. Sensors on the manipulandum accurately measured joint position and velocity





**Figure 2.1: Experimental apparatus. Subjects make reaching movements while grasping a 2 df manipulum (robotic arm). The manipulum could be programmed to produce a force field.**

(details provided by Shadmehr and Brashers-Krug, 1997). Two motors, mounted on the base of the manipulum, could independently produce torque on the proximal and distal joints of the robotic arm. A computer monitor mounted above the manipulum displayed a cursor representing hand position and boxes representing targets. Using pediatric cardiac electrodes, EMG was recorded from four muscles: biceps, triceps lateralis/ longus, anterior deltoid, and posterior deltoid. Before the application of the electrodes, the skin was cleaned and abraded using a preparation gel (NuPrep; Weaver, Aurora, CO). EMG signals were amplified by Grass (Quincy, MA) AC amplifiers (model 8A5) with 60 Hz notch filters. The amplified signal was bandpass-filtered (17–530 Hz), processed through a root-mean-square circuit with a 25 msec integration window (De Luca, 1997), and digitized. Hand position, hand velocity, and processed EMG were recorded at 100 Hz. The processed and digitized EMG formed the signals analyzed below.

Subjects made 10 cm point-to-point movements toward 8 mm square targets represented by boxes on the computer monitor. Subjects made movements in four directions (0, 45, 90, and 135°) away from the center of the work space and four directions (180, 225, 270, and 315°) back to the center; the order of the outward directions was determined pseudorandomly. If subjects completed the movement in  $500 \pm 50$  msec, the box “exploded,” and the computer generated a pleasing sound. If the cursor reached the box too slowly (in

$\geq 550$  msec), the target filled in blue; if the cursor reached the box too quickly (in  $\leq 450$  msec), the target filled in red. The only instruction provided was to explode as many targets as possible. I provided no instructions regarding straightness of movements or smoothness of trajectories. Subjects made reaching movements in two or three different dynamic environments. One, termed the “null field,” is the condition in which the torque motors do not create any forces; in the null field the subjects encounter only the inertial dynamics of the manipulandum. In the second and third environments, B1 and B2, torque motors produce an additional force as described by the equation:

$$\vec{F} = B\vec{x}, \quad (2.1)$$

where  $\vec{F}$  is the force produced at the end effector by the robot’s motors,  $\vec{x}$  is the instantaneous velocity vector of the hand, and B is a viscosity matrix. In the first environment, the null field, B0 equals [0 0; 0 0] N sec m<sup>-1</sup>. In the second environment, B1 equals [0 13; -13 0] N sec m<sup>-1</sup>; in the third, B2 = -B1. These “force fields” exert a force proportional in strength to the instantaneous speed of the hand, in the direction perpendicular to the instantaneous velocity vector.

Subjects’ training was divided into sets of 192 movements. Each set was followed by a 3 min rest period. On a first day of training, all 24 subjects completed three sets of movements in the null field. On a second day, all subjects completed one set in the null field and then three sets in the force field B1. Here I report data only from the second day.

To determine the whether observed changes in EMG were specific to B1, three groups of eight subjects each completed two additional sets of movements in other environments. One group completed these two sets in the null field after the standard 3 min rest. A second group trained in B2 3 min after B1, and a third group trained in B2 6 hr after B1; EMG in these groups also explored the neural basis of anterograde interference (Brashers-Krug, Shadmehr, and Bizzi, 1996). During the 6 hr period, EMG electrodes were removed, and subjects left the laboratory. After their return, electrodes were reapplied on approximately the same arm positions as the earlier session, as marked by the earlier abrasion of the skin.

## 2.2.2 Displacement analysis

To facilitate averaging across movements, all movement data were temporally aligned on the basis of the movement onset, which was defined as when the hand speed first crossed a threshold (0.03 m/sec).

Perpendicular velocity was defined as the component of the velocity vector that pointed perpendicular to the line connecting the initial hand position and target position. Perpendicular displacement was computed by summing perpendicular velocity over the first 250 msec of the movement. Curvature in null field movements was often nonzero. To focus on recovery from the force field, in each direction I subtracted away the average perpendicular displacement generated in the null field (first set, day 2).

## 2.2.3 EMG normalization and polar analysis

Each subject's (bandpassed, RMS'd) EMG traces were normalized based on movement-initiating EMG (as defined below), recorded during the initial null field set (first set of movements in day 2). For each muscle  $m$  and in each direction of movement, I calculated a scalar activation  $a_m$  by averaging the normalized EMG trace  $A_m(t)$  over time interval between 50 msec before and 100 msec after the onset of movement (capital letters refer to full time series; small letters refer to scalar measures) and then averaging over all movements made in that direction during the initial null field set. I constructed an 8 x 1 vector  $\vec{a}_m$  for each muscle, each element representing a single direction. I then adjusted the scale of each muscle's EMG trace in each movement, producing the normalized trace  $A^n$ , using the equation:

$$A_{k,m}^n(t) = 50 * \frac{A_{k,m}(t) - \min(\vec{a}_m)}{\max(\vec{a}_m) - \min(\vec{a}_m)} + 25, \quad (2.2)$$

where  $k$  is an index of movement number, and  $m$  is an index of muscle. Within each subject, therefore, the EMG values recorded during all movements (in the null field and in force fields) were normalized using the same linear transformation. This transformation facilitated comparisons of activation across subjects but preserved the training-induced changes of every subject's EMG.

I used a form of data analysis termed polar analysis to evaluate changes in EMG. I began the analysis by computing the movement-initiating activation for each subject, binned within each direction across eight

movements. I then multiplied each scalar by the unit vector pointing in the direction of movement, so that these vectors, when plotted together, display  $a_m$  as radii pointing in the direction of movement corresponding to that activity. These polar plots summarized the function mapping target direction into initial EMG activity.

To determine the directional bias of each activation function, in each movement bin I added the eight  $a_m$  vectors together to form one resultant vector. This resultant vector had the same direction and is proportional in magnitude to the mean of the eight vectors (if a unit mass would have been placed on the end point of each vector, the mean of the vectors would have been the center of mass of the collection of points). By calculating the orientation of this resultant vector, I detected changes in the angular dependence of  $a_m$  as subjects learned dynamic environments.

## 2.2.4 Computational modeling: structure of the model

The purpose of the computational modeling was to predict the change in the pattern of joint torques that should result if an adaptive control system learned to completely compensate for the dynamics of the force field. To calculate the changes in joint torques attributable to adaptation, we simulated the dynamics of the four-link system of Figure 2.1 through the following coupled differential equations:

$$I_r(\vec{p})\ddot{\vec{p}} + G_r(\vec{p}, \dot{\vec{p}})\dot{\vec{p}} = M(\dot{\vec{p}}) + J_r^T \vec{F}, \quad (2.3)$$

$$I_s(\vec{q})\ddot{\vec{q}} + G_s(\vec{q}, \dot{\vec{q}})\dot{\vec{q}} = \vec{C}(\vec{q}, \dot{\vec{q}}, \vec{q}^*(t)) - J_s^T \vec{F}, \quad (2.4)$$

where I and G were inertial and coriolis/centripetal matrix functions, M was the torque field produced by the robot's motors, i.e., the environment,  $\vec{F}$  was the force vector measured at the handle of the robot,  $\vec{C}$  was the adaptive controller implemented by the motor system of the subject,  $\vec{q}^*(t)$  was the reference trajectory planned by the motor control system of the subject, J was the Jacobian matrix describing the differential transformation of coordinates from end point to joints,  $\vec{q}$  and  $\vec{p}$  were column vectors

representing joint positions of the subjects and the robot (Figure 2.1), and the subscripts s and r denoted subject or robot matrices of parameters.

The robot's passive mechanical dynamics,  $I_r$  and  $G_r$ , as well as a typical subject's passive dynamics of the arm,  $I_s$  and  $G_s$ , were previously estimated (Shadmehr and Brashers-Krug, 1997). A key design question is with respect to the function C. Simulations have previously suggested that a reasonable lumped model of the subject's biomechanical controller in the case of these targeted movements is as follows (Shadmehr and Brashers-Krug, 1997):

$$\vec{C} = \hat{I}_s(\vec{q}^*)\ddot{\vec{q}}^* + \hat{G}_s(\vec{q}^*, \dot{\vec{q}}^*)\dot{\vec{q}}^* - \hat{M}(\dot{\vec{q}}^*) - K(\vec{q} - \vec{q}^*) - V(\dot{\vec{q}} - \dot{\vec{q}}^*). \quad (2.5)$$

K is a linear estimate of subject's joint stiffness and was set to a constant  $K = [15, 6; 6, 16] \text{ N m}^{-1} \text{ rad}^{-1}$  (Mussa-Ivaldi et al., 1985). V is a linear estimate of joint viscosity and was set at  $V = [2.25, 0.9; 0.9, 2.4] \text{ N m}^{-1} \text{ sec}^{-1} \text{ rad}^{-1}$ . This controller is dependent on a model of the subject's passive dynamics, represented by  $\hat{I}_s$  and  $\hat{G}_s$ , and model of the force field environment,  $\hat{M}(\dot{\vec{q}}^*)$ . I assume the desired trajectory,  $\dot{\vec{q}}^*$ , was a minimum jerk motion of the hand to the target (Flash and Hogan, 1985) with a period of 0.5 sec.

Previous simulations have suggested that such a controller produced trajectories in the above differential equations that reasonably agreed with performance of subjects in the field (Shadmehr and Brashers-Krug, 1997).

In previous work, Shadmehr and Brashers-Krug (1997) simulated Equations 2.3 and 2.4 for a configuration of the arm in which the hand made the first movement starting from  $x = -0.1$  and  $y = 0.45$  m, based on the coordinate system shown in Figure 2.1. This is the work space for which the EMG data in this thesis were acquired. The force field produced by the robot was always of the form  $\vec{F} = B\vec{x}$ . This field can be written in joint coordinates of the robot or the subject, i.e., the M variable in Equation 3, using simple kinematic transformations (Shadmehr and Mussa-Ivaldi, 1994). I considered two conditions. In the first condition, the field produced by the robot was  $B = B_0$ , i.e., the null field. In the second condition, I had  $B = B_1$ . I performed a numerical simulation of Equations 2.3 and 2.4 for each condition under the

assumption that the controller had a perfect model of the field. Movements were simulated to the same eight directions as in the psychophysical studies.

The torques produced by the controller of Equation 2.5 on the shoulder and elbow joints were  $\tau_s(t,\theta)$  and  $\tau_e(t,\theta)$ , where  $t$  is time into the movement, and  $\theta$  is the target direction. The calculated elbow and shoulder torques were parceled into muscle force by solving the set of equations,

$$\vec{\tau} = J_m^T \vec{F}$$

$$F_b F_t = k_{bt} \tag{2.6}$$

$$F_a F_p = k_{ap},$$

where  $\vec{F} \equiv [F_b, F_t, F_a, F_p]^T$  is a vector of muscle forces,  $J_m^T$  is the transpose of the differential transformation from muscle lengths to joint angles, i.e.,  $d\lambda/d\theta$ , the Jacobian representing muscle moment arms, and  $k_{bt}$  and  $k_{ap}$  are constants of reciprocal activation. The reciprocal activation constants were derived through analysis of the actual null field EMG data. From the recorded EMG data,  $a_m$  was calculated for each muscle for each direction (this represented the average EMG from  $t = -50$  to  $t = 100$  msec into the movement). Next, a resultant vector was calculated for each muscle. The ratio of the length of the resultant vector to the mean of the  $a_m$  function for each muscle was calculated. This represented the strength of the directional bias of each muscle's EMG. I observed that selection of reciprocal activation constants  $k_{bt}$  and  $k_{ap}$  determined the strength of the directional bias of muscle forces in the model (also calculated over  $t = -50$  to  $t = 100$  msec). I adjusted the activation constants in the model so that it produced muscle forces that matched the strength of the directional biases measured in the  $a_m$  data. I arrived at  $k_{bt} = 1 \times 10^{-3} \text{ N}^2$  and  $k_{ap} = 3 \times 10^{-3} \text{ N}^2$ .

There is no general consensus regarding the magnitude of muscle moment arms (i.e., matrix  $J_m$  in Eq. 6) for the human arm in the horizontal plane. Here I considered six different models (Table 1). Four of the matrices (models A–D) were derived presuming different levels of effective biarticulation of biceps and

Model	Moment arm matrices (cm) $J_m^T = \left( \frac{d\lambda}{d\theta} \right)^T$	Expected rotation in the muscle's EMG function when field B1 is learned (°)			
		Biceps	Triceps	Ant del	Post del
A	$\begin{bmatrix} 4 & -4 & 0 & 0 \\ 0 & 0 & 4 & -4 \end{bmatrix}$	-26.6	-26.5	-14.1	-13.6
B	$\begin{bmatrix} 4 & -4 & 2 & -2 \\ 0 & 0 & 4 & -4 \end{bmatrix}$	-26.6	-26.5	-13.2	-13.0
C	$\begin{bmatrix} 4 & -4 & 4 & -4 \\ 0 & 0 & 4 & -4 \end{bmatrix}$	-26.6	-26.6	-13.7	-13.6
D	$\begin{bmatrix} 4 & -4 & 4 & -4 \\ 0 & 0 & 2 & -2 \end{bmatrix}$	-26.6	-26.6	-19.7	-20.2
E	$\begin{bmatrix} 4.58 & -0.70 & 2.37 & -2.74 \\ 0 & 0 & 4.67 & -1.48 \end{bmatrix}$	-26.6	-26.5	-13.1	-16.6
F	$\begin{bmatrix} 3.16 & -6.07 & 2.20 & -3.43 \\ 0 & 0 & 4.84 & -2.21 \end{bmatrix}$	-26.6	-26.5	-12.7	-15.2

**Table 2.1. Computational representation of an internal model in terms of expected changes in the angular orientation of EMG.**

triceps: model A presumed no biarticulation, and models B–D presumed that the moment arms of biceps and triceps across the shoulder were one-half, the same, and twice the moment arms of the anterior and posterior deltoids, respectively. The fifth model (E) was derived from cadaver data (Wood et al., 1989) recorded with the shoulder adducted and forearm horizontal and extending forward, and the sixth model (F) was model E transformed through a rigid body rotation to the horizontal plane of the shoulders.

## 2.2.5 Computational modeling: predicting rotation of directional bias

For each moment arm model, the simulation produced a time series of muscle forces appropriate to make movements in the null field. In each direction of movement  $i$ , I averaged over the first 150 msec of

predicted force production (corresponding to the interval 50 msec before through 100 msec after movement initiation) to determine for each muscle a scalar variable  $f_m(i)$ . I performed polar analysis (see above) on these scalars to determine a resultant vector to summarize, in null field movements, the angular dependence of  $f_m$  on movement direction. I then performed the same analysis on the forces appropriate to move in the force field B1. From these two polar analyses I determined, for each muscle, the difference between the orientations of  $f_m$  resultant vectors in the null field and in force field B1.

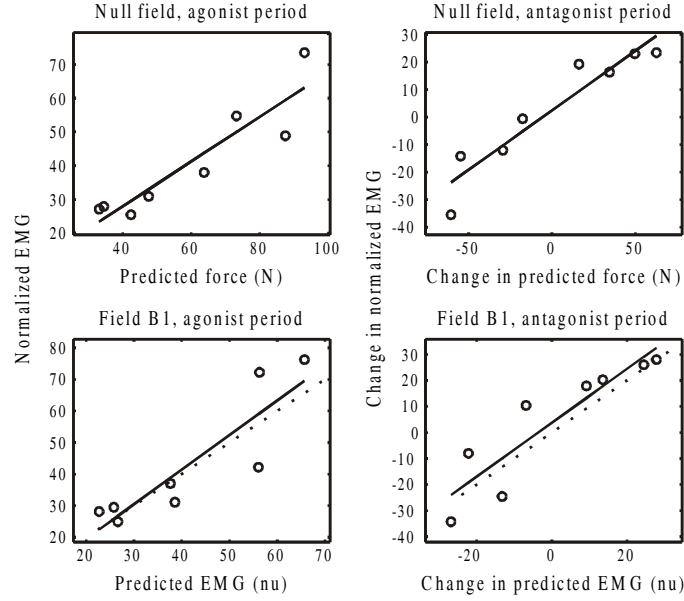
The results for each model are shown in Table 1. I found that changing moment arm matrices tended to change the orientation of preferred directions in both null field and B1 but did not tend to change the angle between them. All models predicted approximately the same levels of rotation in the tuning functions of biceps and triceps. For the anterior and posterior deltoids, predicted rotations were very similar for all models except that of model D. The results in Table 1 show how much rotation of the spatial activation function, as represented in hand-centered coordinates, should occur as a subject learns an internal model for field B1.

## 2.2.6 Computational modeling: predicting time series of appropriate activations

Studies that have recorded EMG during single-joint movements have reported increases or decreases in the activation of individual muscles during learning. This analysis is appropriate for tasks with a single degree of freedom, but here we needed to consider movements with 2 df, actuated by multiple muscles, moving in several directions. To facilitate analysis of EMG across directions of movement, we derived an inverse muscle model, whose parameters were determined by subjects' EMG in the null field. The model was validated by predicting EMG patterns in the force field. These predictions were then used to remap EMG into composite traces based on their functional relevance.

To determine the muscle activations appropriate for this task, I first used Equations 2.3 through 2.6 and muscle model C to calculate the time series of muscle forces necessary to make a minimum jerk reaching movement,  $F_{i,m,n}(t)$ , where  $i$  indicates movement direction,  $m$  muscle, and  $n$  force field ( $n = 0$  signifies the null field,  $n = 1$  the force field B1). Next, I transformed the modeled muscle force  $F_{i,m,n}(t)$  into appropriate





**Figure 2.2: Data used for building a simple force activation model for the anterior deltoid and data used for testing the validity of the model. Top plots: The anterior deltoid force predicted by the inertial model is fitted to the mean EMG recorded from subjects during null field movements. Each point is data for movement to a given direction. Each circle is the average normalized EMG recorded during that period in this muscle from all subjects. The line is the best linear fit ( $p < 0.001$  for each figure), that represents the muscle model. Bottom plots: Validating the muscle model. The linear fit from the null field was used to transform the predicted forces necessary for movements in field B1 to EMG units. The x-axis is the predicted average EMG in anterior deltoid. The circles are the mean recorded EMG in all subjects after full adaptation. The solid line is the muscle model's predictions. The dotted line is the best possible linear fit of the validation data.**

muscle activations. To estimate the relationship between predicted muscle force and recorded EMG  $[A_{i,m,n}(t)]$ , I found a piecewise linear fit between the force predictions in the null field and the recorded EMG from subjects, using the following analyses. In the null field, for each muscle  $m$  and in each direction  $i$ , I calculated the average predicted force and recorded EMG during the agonist period, from 50 msec before through 100 msec after the onset of the movement ( $f_m^{ag}$  and  $a_m^{ag}$ ), and during the antagonist period, from 200 msec through 300 msec into the movement ( $f_m^{ant}$  and  $a_m^{ant}$ ). I then fit the null field forces across all eight directions to the recorded null field EMG using the equations:

$$\vec{a}_m^{ag} = p_m^1 \vec{f}_m^{ag} + p_m^2 \quad (2.7)$$

$$\vec{a}_m^{ant} - \vec{a}_m^{ag} = p_m^3 (\vec{f}_m^{ant} - \vec{f}_m^{ag}) + p_m^4,$$

Parameter	Biceps	Triceps	Ant del	Post del
$p^1$ (nu/N)	1.58	1.70	0.66	0.85
$p^2$ (nu)	-0.10	-0.14	0.015	-0.034
$p^3$ (nu/N)	0.75	0.71	0.43	0.60
$p^4$ (nu)	0.036	0.022	0.024	0.045

**Table 2.2: Values of fitted parameters  $p$  in Equation 2.7. nu: normalized units of EMG.**

where  $\vec{a}$  and  $\vec{f}$  are  $8 \times 1$  vectors reflecting recorded EMG and predicted force in all eight directions, and  $p$  are parameters of the fit.

I calculated the parameters to find the least-squared fit of the null field data and found that the fit was significant for all four muscles, during both the agonist and antagonist bursts (for all eight fits,  $r > 0.88$ ;  $p < 1 \times 10^{-3}$ ; the top two plots in Figure 2.2 show the fits for anterior deltoid in the null field). To validate the fits, I used the same parameters to transform  $f_{i,m,1}$ , the predicted force for field B1, into predicted EMG. The transformed forces fit the mean EMG recorded from the subjects rather well [for one of the fits (triceps during the antagonist burst),  $r = 0.73$ ;  $p < 0.04$ ; for all other fits,  $r > 0.80$ ;  $p < 0.015$ ; the bottom two plots in Figure 2.2 show, for anterior deltoid, the relationship between transformed force and EMG in the field B1 ]. The parameters used for the fits are given in Table 2. This rather simple muscle model, therefore, was validated on data not used to fit the parameters.

For each movement direction, each muscle, and both dynamic environments, I used the parameters  $p$  to map predicted force ( $f$ ) into predicted EMG ( $e$ ; note that  $a$  refers to actual activations) at baseline attributable to the reciprocal activation (base), during movement initiation (ag, from -50 through 100 msec), and midmovement (ant, 200–300 msec). I then compared these activation scalars to determine, in the null field, what increases or decreases from baseline were appropriate for initiating movements. I also

Predicted force direction	Direction (°)	Biceps	Triceps	Ant del	Post del
Null field, excited (0,+)	0	0	0.536	0	0.845
	45	0	0.971	0	0.241
	90	0	0.859	0.513	0
	135	0	0.070	0.998	0
	180	0.604	0	0.797	0
	225	0.975	0	0.222	0
	270	0.813	0	0	0.583
	315	0.063	0	0	0.998
Null field, inhibited (0,-)	0	0.669	0	0.744	0
	45	0.965	0	0.265	0
	90	0.761	0	0	0.649
	135	0.080	0	0	0.997
	180	0	0.605	0	0.796
	225	0	0.944	0	0.331
	270	0	0.876	0.482	0
	315	0	0.135	0.991	0
Difference between null and B1, excited (1,+)	0	0	0.921	0.389	0
	45	0	0.182	0.983	0
	90	0.869	0	0.495	0
	135	0.997	0	0.074	0
	180	0.937	0	0	0.350
	225	0.396	0	0	0.919
	270	0	0.819	0	0.574
	315	0	0.9991	0.042	0
Difference between null and B1, inhibited (1,-)	0	0.863	0	0	0.505
	45	0.027	0	0	0.9996
	90	0	0.825	0	0.565
	135	0	0.988	0	0.154
	180	0	0.947	0.322	0
	225	0	0.563	0.826	0
	270	0.814	0	0.581	0
	315	0.9995	0	0	0.031

**Table 2.3: Values of the vector operator  $v$ , as defined in Equation 2.9.**

determined what increases or decreases from null field activations were appropriate to counter the novel forces produced by field B1 midway through the movement (when field-induced forces were maximal). I summarized these predictions in a three-dimensional object  $x(i, m, n)$ :

$$x(i, m, 0) = e_{m,0}^{ag}(i) - e_m^{base}(i); \quad x(i, m, 1) = e_{m,1}^{ant}(i) - e_{m,0}^{ant}(i). \quad (2.8)$$

For each movement direction  $i$  and force field  $n$ ,  $x(i, m, n)$  for two of the muscles were positive; the size of these activations predicted the proportion that each of those muscles participated as an agonist in making a movement in the null field (for  $n = 0$ ) or as an agonist for resisting the forces of field B1 (for  $n = 1$ ). On the other hand,  $x(i, m, n)$  was negative for the remaining two muscles. These negative values predicted that the two remaining muscles produced a force opposite that of the imposed field and should be inhibited. I constructed 1 x 4 vector operators  $v$ ,

$$v_{i,n,+}(m) = \begin{cases} x(i, m, n) & \text{if } x(i, m, n) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2.9)$$

$$v_{i,n,-}(m) = \begin{cases} -x(i, m, n) & \text{if } x(i, m, n) < 0 \\ 0 & \text{otherwise} \end{cases},$$

to identify functional agonists (+) and antagonists (-) in a given force field. The magnitude of each vector  $v$  was normalized to 1. The vector values are shown in Table 3 for the null field and for field B1. Each vector operator predicts, in the four-dimensional space of muscle activation, in what direction activation should increase or decrease (at the peak of activity) to effectively make movements in a given field. To give an example, consider a movement toward a target at  $90^\circ$ . In the null field ( $n = 0$ ), the vector operator  $v_{90,0,+}$  predicts that to initiate an accurate movement, the triceps and anterior deltoid should be activated above their baseline levels. Furthermore, if one considered the triceps and anterior deltoid activations to form an  $(x, y)$  plane, the model predicts that the average activity over the agonist period would point in the direction of the unit vector  $(0.859, 0.513)$ . To move in the field B1 toward  $90^\circ$ ,  $v_{90,1,+}$  predicts that biceps and anterior deltoid should be activated more strongly than in the null field. Midway through the movement (200–300 msec into the movement), the unit vector pointing in the direction of the most appropriate increase in activation would equal  $(0.869, 0.495)$ .

### 2.2.7 Using model predictions to remap subjects' EMG

I used the vector operators (Table 2.3) to remap EMG recorded from the four muscles into composite traces,  $T$ :

$$T_{i,n,\pm}(t) = \vec{v}_{i,n,\pm} \cdot \vec{A}_i(t), \quad (2.10)$$

where  $\vec{A}_i \equiv [A_{b,i}(t), A_{t,i}(t), A_{a,i}(t), A_{p,i}(t)]$  is a concatenation of EMG from the four muscles in a movement toward direction  $i$ . For example, the composite trace  $T_{i,1,+}$  reveals, at each time point, the projection of the four-dimensional EMG vector (one dimension for each muscle) onto the direction of activation that performs work against the force field B1. Whereas the original EMG traces were strongly direction-dependent, the composite traces reveal the direction-independent muscular activity that compensates for a given field. The composite traces therefore can be compared and averaged across directions to understand adaptation of neural output across the work space.

The transformation in Equation 2.10 remapped the full time series of EMG traces using the predicted optimal activation over specific time intervals (the average over the agonist period for null field data, over midmovement for the increase appropriate for field B1). These remappings retained the natural temporal structure of the EMG traces for each muscle; the activation in each trace at each time point was weighed equally. When the remapping of the increase appropriate for B1 was instead determined by predictions during other time intervals, the direction of appropriate activation changed very little over the range from 50 msec before through 300 msec after the onset of movement (averaged across directions of movement, change in unit vector direction  $\leq 15^\circ$ ). The unit vectors  $v_{i,+}$  in Table 3, therefore, accurately reflect the predicted direction of increased activity throughout the movement.

The composite trace  $T_{i,1,+}$  summarized EMG activity that I estimated as effectively countering the forces of field B1. This trace was averaged across directions and binned over 64 movements (three bins per target set of 192 movements). The averaged and binned traces were then fit with fourth-order polynomials. Of primary interest was the question of how this function evolved as subjects practiced in the force field. A major component of this change was an apparent shift in the peak of the function. Bootstrapping techniques (see below) were used to estimate 95% confidence intervals of timing of the peak.

### 2.2.8 Wasted contraction: a measure of co-contraction

Previous studies have measured “co-contraction” through the addition of the EMG traces of opposing muscles (Milner and Cloutier, 1993). If the activation of a single muscle increases, however, the sum of activation of this muscle and its antagonist would increase. The co-contraction metric would then increase, even though only one muscle’s activation increased. To avoid this ambiguity, I used a new measure termed “wasted contraction,” which used the remapped composite traces to determine the amount of contraction that is cancelled by the opposing groups and therefore does not lead to effective force production. For each subject, I first remapped the (previously scaled) EMG of each movement into composite traces (Eq. 10). At each sampling point, I then determined, in each agonist–antagonist pair of composite traces, which trace was smaller. Because this activation was cancelled by the larger, opposing activation, the time series of this smaller activity was termed wasted contraction. This amount was then subtracted from the larger trace, creating a time series of “effective contraction.” The wasted contractions were averaged across force directions ( $n \in 0, 1$  from above), averaged across movement direction, and binned over 32 movements (six bins per set). Each subject’s wasted contraction was then scaled by the maximum effective contraction measured over bins in the null field and in B1. The wasted contraction traces were then averaged across subjects. This average wasted contraction was fit with a fourth-order polynomial; the maximum of this polynomial was calculated and compared during different stages of training. The 95% confidence intervals for the change in maximum wasted contraction were calculated by bootstrapping (see below).

### 2.2.9 Calculation of force produced by the human controller

I also used the structure and parameters of the model of the human–robot interaction to estimate the amount of force subjects create during movements. Using Equations 2.3 and 2.4 and using estimates of the human and robot dynamical parameters, I transformed individual subjects’ hand position and velocity into estimated torque produced by the subjects’ muscles (the acceleration terms in these equations were estimated by filtering the velocity data through a second-order Savitsky–Golay [differentiating] filter). I then multiplied this torque by  $J_s^{-T}$ , the transposed inverse of the subject’s hand Jacobian, to determine the force produced at the hand attributable to joint torques. I remapped the forces produced in each direction of movement into the components parallel and perpendicular to the target direction. I averaged across

directions and across subjects, fit each force trace with a fourth-order polynomial, and determined the timing of the maximum of each trace. Confidence intervals were determined by bootstrapping.

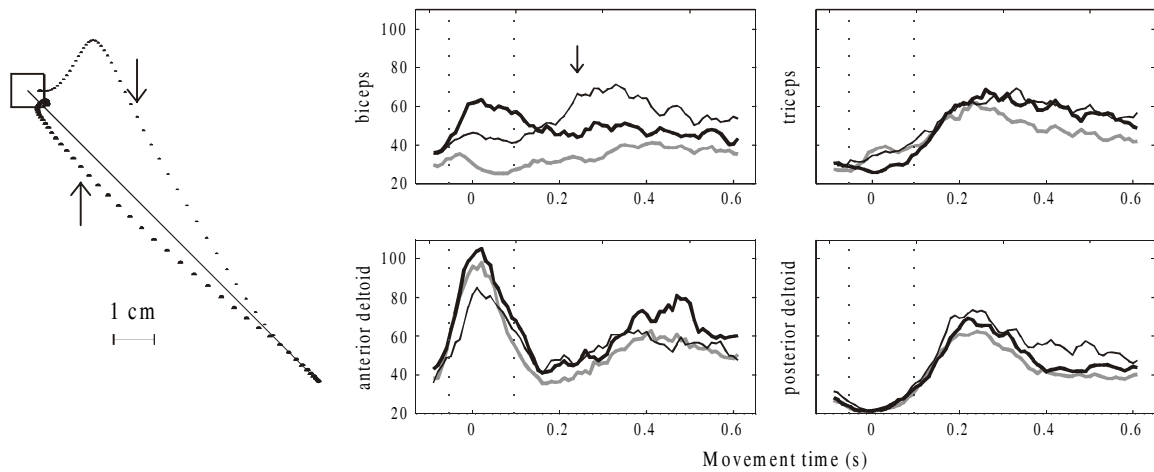
### 2.2.10 Bootstrapping

Confidence intervals for the rotation of EMG resultant vectors, the timing of the first peak of the B1-appropriate EMG trace, and the maximum level of wasted contraction were determined by using bootstrapping (Fisher, 1993; Welsh, 1996). The original pool of data from 24 subjects was resampled, with replacement, 1000 times to create 1000 new sets of 24 data points (for rotation) or 24 time series (of EMG). Each resampled set was averaged; for time series, each series was fitted with a fourth-order polynomial, and then the first peak (for B1-appropriate EMG and forces) or the maximum (for wasted contraction) was calculated from this fit. The resulting set of 1000 statistics was sorted by rank; the 95% confidence interval was determined by the 26th and 975th elements of this vector.

## ***2.3 EMG correlates of learning: a single direction of movement***

Let us initially consider movements made toward a single direction. When subjects made movements in the null field at 135°, their EMG traces (Figure 2.3, gray lines) revealed that the anterior deltoid was active early to initiate the movement, whereas triceps and posterior deltoid were active later to brake the movement. Early in training in force field B1, when subjects attempted to make movements toward 135° they generated inaccurate hand paths that diverged from the target direction (Figure 2.3, top trajectory). Midway through a typical movement, the activity in biceps (Figure 2.3, thin black line) increased; the hand path was then corrected toward the target. The timing of the increased biceps activation suggested that this muscle was activated in response to the robot forces that extended the elbow during the reach. The increased activation of biceps was possibly an error–feedback response mediated via a stretch reflex mechanism. In contrast to these field-induced changes in the biceps, the time series of activations for anterior and posterior deltoids in the force field closely resembled the activations recorded in these muscles in the null field.


As subjects trained in field B1, the movements became straighter (Figure 2.3). The EMG from these improved movements featured larger biceps activity very early, preceding the movement itself (Figure 2.3,



**Figure 2.3: Hand paths and EMG recorded during movements. Left, Hand paths of a typical subject in her first (top) and last (bottom) movements toward a target at 135° in the force field B<sub>1</sub>. Dots are at 10 msec intervals. The arrow indicates 250 msec into the movement. Remaining plots, Normalized EMG in the null field (all 24 movements toward 135°, gray line), during initial stages of training in the force field (first 8 movements toward 135°, thin black line), and late in training in the force field (last 8 movements, thick black line). Movement begins at  $t = 0$ . Each EMG trace represents data averaged across all 24 subjects. The vertical dotted lines delimit the 150 msec interval over which EMG is averaged to calculate the scalar variable  $a_m$ , representing time-averaged agonist burst EMG.**

thick black line). The biceps then became less active later in the movement, during the same period in which it was active when subjects first trained in B<sub>1</sub>. Therefore, a major component of the adaptation process was a modification of the EMG patterns necessary to initiate the movement. To summarize the changes in early EMG in movements toward 135°, I averaged EMG over the time interval from 50 msec before to 100 msec after the onset of movement. This scalar measure of activation was termed  $a_m$ . Figure 2.4 shows the evolution of  $a_m$  as subjects moved in the null field and then trained in B<sub>1</sub>. The figure also displays perpendicular displacement, a measure of movement error that is computed 250 msec into the movement. These data show that as subjects trained in B<sub>1</sub>, the curvature (error) of movements became significantly smaller ( $p < 5 \times 10^{-8}$ ),  $a_m$  in biceps significantly increased ( $p < 1 \times 10^{-4}$ ), and  $a_m$  in triceps significantly decreased ( $p < 0.002$ ) compared with their null field activation. Neither anterior nor posterior deltoid activation significantly changed with training in B<sub>1</sub> (anterior,  $p > 0.15$ ; posterior,  $p > 0.35$ ). In summary, the changes in EMG traces measured in movements toward 135° suggested that when subjects were first exposed to B<sub>1</sub>, they generated additional muscle activations late in the movement to correct for



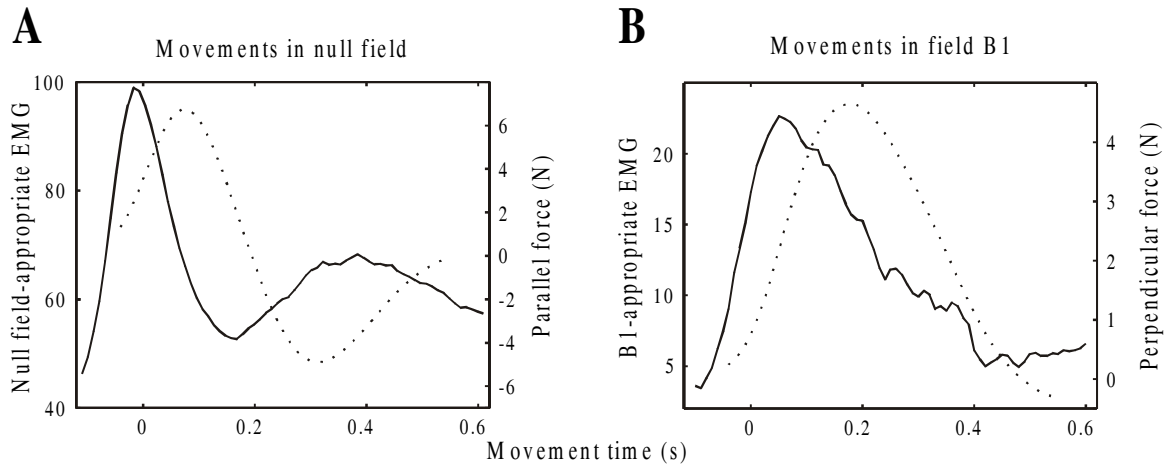


**Figure 2.4: Perpendicular displacement (*top*) and subjects'  $\alpha_m$  (time-averaged agonist burst EMG) during movements toward  $135^\circ$ . Subjects completed one set of movements in the well learned null field environment and then three sets in the force field B1. Perpendicular displacement is measured 250 msec into the movement, averaged across directions and subjects; these data reflect change from average perpendicular displacement in null field movements. *x*-axis **tick marks** indicate breaks between sets of movements. Each **point** includes data from eight movements (3 data points per set). Error bars reflect 95% confidence intervals of the mean.**

movement curvature. With further training, subjects learned to generate this added activation early in the movement, predicting rather than responding to the force field.

## **2.4 EMG correlates of learning across all directions of movement**

To quantify the correlation between changes in EMG and the formation of an internal model, I wanted to extend analysis from a single direction to many directions of movement. This extension is hampered by the fact that each muscle makes different contributions to movements in different directions. I circumvented this obstacle by using the predictions of the simulated controller. I considered the time pattern of EMG in the muscles for a given movement as a path in a four-dimensional space, each axis representing a muscle. The model also predicted paths in this space: one path for moving in a null field and one for moving in a force field. The difference between these two predicted paths indicated activity in that space which should be increased, peaking midway through the movement, to move accurately in B1. In each movement direction, the increased activity projected into a single component of muscle activity (Eq. 2.9, Table 3; see Materials and Methods). The discovery of appropriate activations allowed us to project the EMG recorded



**Figure 2.5: Subjects' composite EMG traces (solid lines) and forces (dotted lines). A, Null field-appropriate EMG and force in the direction parallel to the target, averaged over all subjects and all movements in the initial null field set (192 movements). B, B1-appropriate EMG trace and force perpendicular to the direction of target, averaged over all subjects and over all movements during the final set of training in force field B1. For the time series of EMG and force, maximum range of 95% confidence intervals of EMG and force were: null field-appropriate EMG,  $\pm 5.3$  units; parallel force,  $\pm 0.28$  N; B1 -appropriate EMG,  $\pm 6.2$  units; perpendicular force,  $\pm 0.13$  N.**

from each subject, during each movement in the field B1, into functionally appropriate composite traces (Eq. 2.10).

We calculated the null field- and B1-appropriate EMG traces during the first null field set and the third set in field B1, respectively. The resulting traces were averaged over all movement directions and all subjects to arrive at a single trace. The results are shown in Figure 2.5 (solid lines). The lefthand plot (Figure 2.5A) shows the average null field-appropriate EMG trace produced by subjects as they reached in the null field. The peak of this trace was at 213.3 msec, i.e., before the start of the movement [95% confidence interval of the mean (CIM), (217.5, 28.2) msec]. In comparison, the average force that was produced by the subjects (Figure 2.5A, dotted line) had its peak 68.4 msec into the movement [95% CIM, (66.1, 70.6) msec].



The composite EMG trace shown in Figure 2.5A represents the component of activation appropriate for moving the hand toward a target, i.e., the agonist burst of activity. A second composite EMG trace (data not shown) produces force in the direction opposite the movement direction. This composite trace (defined by  $v_{0,-}$  in Table 3) peaked 231 msec into the movement [95% CIM, (222, 240) msec]. This peak of activity generated the burst of force in the negative direction in Figure 2.5A, which reached a minimum (a

maximum in the negative force direction) 316 msec into the movement [95% CIM, (313, 319) msec]. As expected, the temporal profile of the force trace was essentially a delayed version of the transformed EMG traces.

This close correspondence was maintained when I compared the B1-appropriate EMG with the pattern of forces that subjects had produced perpendicular to the direction of target during the training. Figure 2.5B shows the B1-appropriate EMG and the corresponding forces, both averaged across subjects during the last set of training in the force field. With training, subjects had learned to alter their EMG early into the movement [Figure 2.5B, peak of the solid line at 78 msec; 95% CIM, (67, 91) msec] so that the muscles produced a bell-shaped force pattern perpendicular to the direction of the target, effectively canceling the force field.

The relationship between the composite EMG traces and the subject-generated forces confirms the aptness of the transformation of EMG into composite traces. The vector operators in Table 3, created from the predicted activations of the biomechanical model, have elicited the components of EMG appropriate for moving in the null field and resisting the additional force produced by field B1. I will now examine the component appropriate for field B1 throughout the training in the force field, to determine how the neural activation of the appropriate musculature evolves with learning.

Within subjects, we averaged the B1-appropriate EMG across directions and across bins of 64 movements (three bins per set, nine bins spanning training in B1). The mean trace for each subject for each bin was averaged across subjects. The resulting traces are shown in Figure 2.6. Because the imposed force field was a function of hand velocity, and hand velocity was generally unimodal, we expected the EMG that effectively countered this force to be also unimodal. Indeed, we found that the B1-appropriate EMG was always unimodal (single-peaked). Furthermore, because of force in the negative direction in Figure 2.5A, which reached a minimum (a maximum in the negative force direction) 316 msec into the movement [95% CIM, (313, 319) msec]. As expected, the the response to the imposed force field could initially be only through a delayed error feedback system (e.g., through spinal reflexes), we expected the initial B1-appropriate EMG to lag the imposed forces. This was also observed.



**Figure 2.6:** The composite EMG trace appropriate for field B1 shifts with training. Top: the binned and averaged increase in the B1-appropriate EMG trace. Each trace is averaged over 64 movements and over all 24 subjects. The progression of plots from top to bottom show EMG traces from early to late in training. The dashed lines in the top and bottom axes represent the magnitude of the force created by the viscous field B1, which remained consistent throughout training in B1. While early in training the peak of the field-appropriate EMG lagged the imposed force, with training it preceded the imposed force. The timing of the peak was estimated by a fit to a fourth order polynomial. The best estimate of the peak and its change with training is marked by the line that crosses the EMG traces from top to bottom. In the bottom figure, the timing of the peak of the EMG traces is plotted vs. movement number. The dotted lines connecting data points represent the 3 min breaks between sets. Error bars reflect 95% confidence intervals of the mean (across all 24 subjects), as determined by bootstrapping. The dotted horizontal line represents the timing of the peak of the force created by the force field.

With practice in the force field, the peak of the B1-appropriate EMG trace shifted to occur earlier into the movement. To estimate the peak location of each EMG trace, the data were fitted to a fourth-order polynomial. I found that when subjects started to train in B1, the peak of the B1-appropriate EMG occurred 246 msec into the movement [95% CIM, (215, 274) msec]. This peak arrived 48 msec after the peak of the force created by the force field B1; the peak of the robot-imposed force coincided with the peak tangential velocity of the hand and on average occurred at 195 msec into the movement throughout training in B1. By the end of the third set of training in B1, subjects had modified their motor commands such that the peak of the B1-appropriate EMG occurred just 67 msec into the movement [95% CIM, (48, 86) msec; in t test of time shift,  $p < 1 \times 10^{-4}$ ]. Therefore, when subjects first encountered B1, they activated the appropriate muscles midway through the movement. The timing of this activation, coupled with the marked curvature of these movements, suggested that subjects activated these muscles through an error-driven feedback mechanism. With training in B1, subjects learned to activate appropriate muscles earlier in the movement, in the period during which control was purely feedforward. The changes in timing of the B1-appropriate EMG activity revealed that subjects adapted their neuromotor output throughout all three sets of movements (Figure 2.6, bottom). This measure was much more sensitive than the kinematic measure of perpendicular displacement, which showed no learning during the second or third sets (compare Figure 2.8).


With training, subjects generated B1-appropriate EMG earlier in the movement. During the 3 min break between sets, however, the subjects partially unlearned this adaptation and began the next set with activation of B1-appropriate muscles shifted later than at the end of the previous set (Figure 2.6, bottom). In between the second and third sets of training in B1, this shift forward in time, i.e., the “forgetting,” was significant (mean shift, 31 msec;  $p < 0.05$ ). The shift of the B1-appropriate activation backward in time during the second and third sets and forward in time during breaks between sets indicate that the timing of this component of activation is a valid metric of learning that reveals more about the state of the adaptation than do kinematic measures alone.

## **2.5 Rotation of angular dependence of forces and EMG**

### **2.5.1 Simulation results**

These results suggested that subjects initially responded to B1 by generating appropriate EMG quite late in the movement, perhaps because of an error–feedback mechanism, but with training they learned to activate the same muscles earlier in the movement. This transition suggested that reactive responses early in training effected proactive behavior late in training. I then focused attention on the subjects' initial, purely feedforward, responses to the task: how do humans generate movement-initiating muscle activations based on the visual cue of a target, a desired end point of a movement? How does this neural computation, mapping stimulus to motor neuronal activation, change with learning? To guide investigation of feedforward muscle activation, I revisited the predicted forces generated by the ideal controller. I summarized each muscle's force production in both the null field and in B1 by averaging over the first 150 msec of the time series of force to create a scalar  $f_m$ . I used polar analysis (see Materials and Methods) to visualize the angular dependence of  $f_m$  (Figure 2.7, top). I then calculated resultant vectors (radial vectors in Figure 2.7, top), which indicate the directions of maximal force production and summarize the dependence of force on the direction of movement. Because I predicted muscle activations using a linear transformation of muscle force, the same polar analysis reveals the predicted directional bias of the muscle activations.

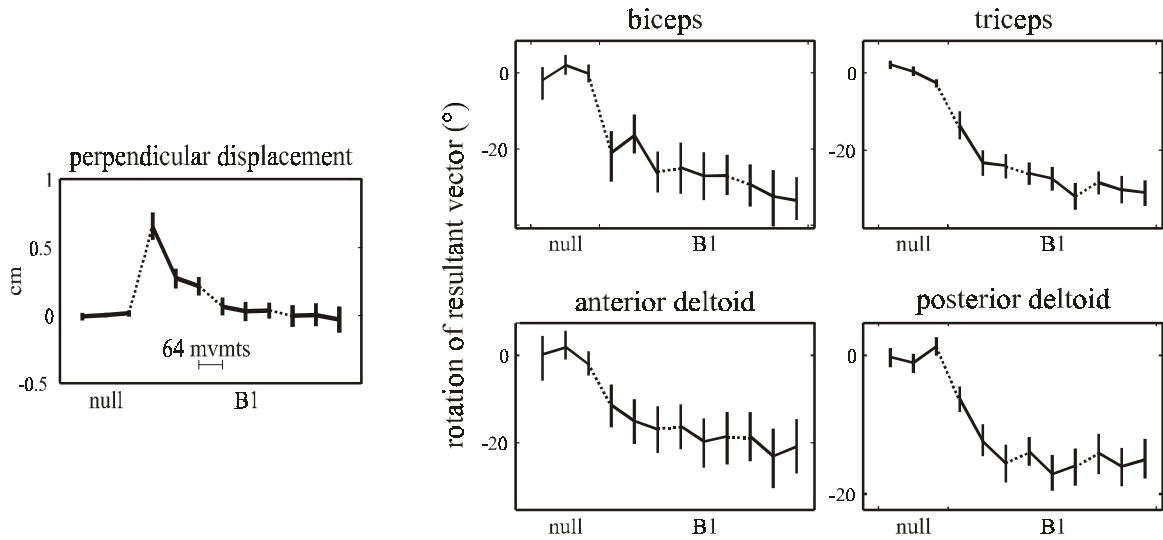
The simulation results revealed that by learning to compensate for B1, the resultant vector of each muscle should rotate by a specific amount (Figure 2.7, top). I found that the degree of rotation was similar over a wide range of moment arm models (Table 1) and remained constant over all levels of co-contraction. (For model C, the directions of the resultant vectors were biceps,  $-130.7^\circ$  for null and  $-157.3^\circ$  for B1 ; triceps,  $49.3^\circ$  and  $22.7^\circ$ ; anterior deltoid,  $150.0^\circ$  and  $136.3^\circ$ ; posterior deltoid,  $-29.6^\circ$  and  $-43.2^\circ$ .) In learning B1, the simulation predicted that the resultant vectors should rotate clockwise by  $26.6^\circ$  for biceps and triceps and  $13\text{--}20^\circ$  for anterior and posterior deltoid, depending on the moment arm model used (Table 1). These observations suggested that humans might calculate a function mapping desired end point location to initial alpha-motor neuronal output, and that when subjects learn to move in B1, the activation function rotates to produce appropriate muscle forces to accurately initiate movements.



**Figure 2.7: Rotations in muscle-tuning curves as predicted from a computational model that assumes adaptation of an IM and tuning curves as recorded from subjects' EMG. *Top*, Movement-initiating muscle forces as predicted by the model for movements in the null field (*gray*) and in B1 (*black*). Polar plots display predicted movement-initiating force ( $f_m$ ) for each direction of movement. The radii of the scale circles (centered at the origin of each plot) are 20 N. The *thick vector* represents a tuning curve's resultant vector (preferred direction). The model predicts that the resultant vector should rotate between the null field and B1. The *bars* below each plot reveal the rescaling of the data (using Eq. 2.7 and the parameters in Table 2) into predicted muscle activations, in normalized units (*nu*) of EMG. *Bottom*, Tuning functions representing subjects' movement initiating EMG ( $a_m$ ) during training in null field (*gray lines*), early force field (*thin black lines*), and late force field (*thick black lines*). Radii are resultant vectors (preferred directions). The error bars around individual data points reflect 95% confidence intervals of the mean activation.**

### 2.5.2 Rotation of angular dependence of EMG with training in B1

Armed with the result that the simulated controller's muscle activation resultant vectors rotated between the null field and B1, I analyzed the EMG generated at the beginning of movements in both the well-learned null field and in B1. Just as in analysis of movements made toward  $135^\circ$ , I averaged each EMG trace over the period from 50 msec before through 100 msec after the beginning of the movement to create the scalar  $a_m$ . I then averaged these scalars over bins of eight movements within directions and then averaged across all subjects. I compared  $a_m$  generated at three stages of training: at the end of the set in the null field; at the beginning of the first set of training in B1; and at the end of the third set of training in B1. The results are shown in Figure 2.7, bottom row.



**Figure 2.8:** *Left, Perpendicular displacement of the hand 250 msec into the movement, averaged across movement directions, during movements in null field and force field B1. Remaining plots, Rotation of  $a_m$  resultant vectors during movements in force field B1 with respect to null field. Each point contains data from 64 movements. Each point is mean and 95% confidence interval (across subjects). Dotted lines represent 3 min breaks in training.*

In the null field (Figure 2.7, gray lines), each muscle was strongly activated in certain “preferred” directions and much less activated in the opposing directions. This directional bias to the activation function was summarized by using polar analysis to construct a resultant vector for each muscle. The resultant vectors pointed toward the direction (for biceps, 2130°; triceps, 50°; anterior deltoid, 127°; posterior deltoid, 217°) of maximal activity presuming a sinusoidal fit of the data.

During the first 64 movements in field B1 (Figure 2.7, thin black lines), subjects began to generate movement-initiating muscle activations that differed from those generated in the null field. The resultant vectors of  $a_m$  rotated clockwise from the directions recorded in the null field. Through the end of training in field B1, the resultant vectors summarizing  $a_m$  continued to rotate in a clockwise direction (Fig. 2.7, bottom). To quantify this adaptation across subjects, we began by determining the orientation of individual subjects’ resultant vectors during null field and B1 training. We rotated each subject’s activation functions so that the mean null field resultant vectors pointed toward 0°. We then calculated the mean and 95% confidence intervals of the orientations of the  $a_m$  vectors during training in the null field and B1. The results are shown in Figure 2.8.




As subjects made reaching movements in the null field, the orientation of the EMG resultant vectors remained constant, and the curvature of movements remained small. In the first bin of 64 movements in B1, subjects' movements were markedly curved, with mean perpendicular displacement, averaged across directions, of 0.65 cm [Figure 2.8; 95% CIM, (0.55, 0.75) cm]. By the end of the third set of training, subjects learned to make much straighter movements [mean perpendicular displacement, 20.03 cm; 95% CIM, (20.13, 0.07) cm; in t test of curvature reduction,  $p < 1 \times 10^{-8}$ ]. The  $a_m$  resultant vectors revealed that subjects began to adapt their initial muscle activations in the first 2 min of training. In the first bin of 64 movements in B1, the orientations of the vectors summarizing all four muscles had significantly rotated from the average null field orientations ( $p < 0.002$  for anterior deltoid;  $p < 1 \times 10^{-5}$  for other muscles). Although these movements were still markedly curved, subjects had begun to adapt their neuromotor output.

As subjects continued to train in B1, the orientation of the  $a_m$  resultant vectors continued to rotate farther away from the null field orientation. By the end of the third set of training in B1, the final directions of the  $a_m$  resultant vectors were significantly different from the orientations early in training in B1 ( $p < 0.05$  for anterior deltoid;  $p < 2 \times 10^{-5}$  for other muscles). The magnitudes of rotation of the  $a_m$  resultant vectors over the three sets of training in B1 were  $-33.7^\circ$  for biceps,  $-28.0^\circ$  for triceps,  $-19.1^\circ$  for anterior deltoid, and  $-15.6^\circ$  for posterior deltoid. These observed values were similar to the predictions made by the simulated controller: within  $7^\circ$  for the flexors and within  $2^\circ$  for the extensors. This result demonstrates a technique whereby a computational model predicts the change that should occur in the activations of muscles if the brain is learning an internal model specific to a force field. The results show that the evolution of the changes in muscle activations early into the movement are reasonably close to the model's predictions.

## ***2.6 Activation function rotation is specific to the force field***

I have suggested that subjects learned to rotate their initial muscle activations to accurately reach in the novel force field B1. An alternate interpretation of the above data would be that the observed rotation was a function of either an increased familiarity with the task itself, regardless of the force field, or that the observed changes were a function of fatigue. To test these alternate hypotheses, I presented one group of eight subjects with the null field after having just completed B1 training. When subjects were returned to



**Figure 2.9: Perpendicular displacement of the hand and rotation of  $a_m$  resultant vectors in subjects who, after training in B1 , trained in the null field 3 min later. Each *point* contains data from 64 movements. The first three data points are from the third set of movements in B1 . Error bars reflect 95% confidence intervals of the mean (across subjects).**

the null field (Figure 2.9), their hand trajectories were initially curved, in the direction opposite the error initially generated in B1 [mean perpendicular displacement, 20.57 cm; 95% CIM, (20.71, 20.43)]; this curvature reveals aftereffects lasting minutes after training in B1. After two sets (350 movements) of training, however, the trajectories again resembled the straight movements made in that day's first null field set [mean perpendicular displacement, 0.02 cm; 95% CIM, (-0.04, 0.08)].

As subjects made straighter movements, they generated movement-initiating muscle activations that reverted back to their original dependence on the direction of the target. Over the last two sets of training in the null field, the orientation of the  $a_m$  resultant vectors all rotated counterclockwise, back to the initial null field directions [amount of rotation: mean (95% CIM) of the amount of rotation: biceps, 26.8° (14.2°, 40.6°); triceps, 23.6° (16.8°, 29.3°); anterior deltoid, 12.3° (4.7°, 19.1°); posterior deltoid, 17.7° (15.1°,

20.4°); for all four muscles,  $p < 0.002$ ]. The average resultant vectors in the last set of movements pointed in directions statistically indistinguishable from the vectors summarizing the  $a_m$  generated in the day's first null field set (for all four muscles,  $p > 0.35$ ).


The return of the orientation of the  $a_m$  resultant vectors to the original direction indicated that the rotation observed when subjects learned B1 was not attributable to an increasing familiarity with the general task or to fatigue but reflected a force field-specific change in the neural command sent to the muscles.

## ***2.7 Activation function rotation during training in B2***

I have suggested that the subjects may have learned to make movements in B1 by rotating the function mapping desired end point to early muscle activation. To determine whether a different rotation of this activation function could underlie learning of other dynamic environments, we trained two groups of eight subjects each in B2, a force field anticorrelated to B1. One group of eight subjects completed two sets in B2 3 min after the three sets in B1; the other group of eight subjects took a 6 hr break between the two force fields. Previous work (Brashers-Krug, Shadmehr, and Bizzi, 1996) has indicated that at 3 min, but not at 6 hr, learning field B1 interferes with subjects' ability to learn the anti-correlated B2. In each of the subjects in these groups, we measured the  $a_m$  generated for each movement, binned the data across eight movements, and constructed resultant vectors. We then determined the orientation of these vectors compared with each subject's initial null field orientation.

The 6 hr group had their electrodes removed after the first training session and reapplied for the second training session. Because the preferred direction of EMG is independent of the total signal strength, the resultant vector orientations calculated from the two sessions are directly comparable.

In the first 64 movements in B2, both groups of subjects made markedly curved movements [Figure 2.10; across both groups, mean of perpendicular displacement, -1.64 cm; 95% CIM, (-1.45, -1.84) cm]. After two sets of training, both groups of subjects made much straighter movements [mean, -0.23 cm, 95% CIM, (-0.12, -0.34) cm]. The 3 min group, however, generated movements that were more curved than did the 6 hr group. A two-way ANOVA (two groups 3 six bins) revealed that the group factor is significant ( $F_{(1,84)} =$



**Figure 2.10: Perpendicular displacement and rotation of  $a_m$  resultant vectors in subjects who, after training in B1, trained in B2 3 min (o) or 6 hr (●) later. Each point contains data from 64 movements. Error bars reflect 95% confidence intervals of the mean (across subjects). The first three data points are from the third set of movements in B1.**

40.82;  $p < 1 \times 10^{-8}$ ) and that in all but the last bin of movements, the larger amount of curvature in the 3 min group is significant ( $p < 0.05$ , Tukey test). The 3 min group initially had more difficulty than the 6 hr group in making accurate movements; this difference remained statistically significant through 10 min of training in B2.

As both groups of subjects reduced the curvature of their movements, subjects rotated their activation functions back to and beyond the orientations appropriate for the null field. By the second set of training in B2, all four  $a_m$  resultant vectors pointed in directions that were rotated clockwise from the activity appropriate for B1; three of the four vectors pointed in directions significantly more clockwise than the directions appropriate for the null field [orientation of  $a_m$  vectors in the second set of B2, mean (95% CIM): biceps,  $11.8^\circ$  ( $6.2^\circ$ ,  $17.3^\circ$ );  $p < 1 \times 10^{-4}$ ; triceps,  $16.9^\circ$  ( $11.4^\circ$ ,  $22.2^\circ$ );  $p < 1 \times 10^{-4}$ ; anterior deltoid,  $4.1^\circ$  ( $-9.9^\circ$ ,  $16.3^\circ$ );  $p < 0.5$ ; posterior deltoid,  $22.2^\circ$  ( $17.3^\circ$ ,  $26.8^\circ$ );  $p < 1 \times 10^{-4}$ ].


Although both groups produced muscle activations that rotated from the null field activity, the initial responses of the 3 min group were less adapted to B2 than were the initial responses of the 6 hr group. The orientations of the  $a_m$  resultant vectors were compared over the three bins spanning the first set of training in B2 (Figure 2.10, open circles represent the 3 min group  $a_m$ , filled circles represent the 6 hr group  $a_m$ ). I performed a two-way ANOVA (two groups x three bins) on the  $a_m$  vector orientations of biceps, triceps, and posterior deltoid. The orientations of the  $a_m$  vectors of the 3 min group had rotated significantly less than did the 6 hr group (for biceps,  $F_{(1,42)} = 5.3$ ;  $p < 0.03$ ; for triceps and posterior deltoid,  $F_{(1,42)} = 5.2$ ;  $p < 0.03$ ).

In the second set of training in B2 (i.e., after 192 movements), the 3 min group continued to lag behind the 6 hr group in their adaptation of their activation of their posterior deltoid. A second round of two-way ANOVA on the  $a_m$  resultant vectors revealed that the posterior deltoid activation of the 3 min group was significantly closer to the orientation of the null field activation ( $F_{(1,42)} = 11.25$ ;  $p < 0.002$ ). By the second set of training in B2, the two groups had begun to generate activations of biceps and triceps whose resultant vectors pointed in similar directions (for biceps,  $F_{(1,42)} = 0.02$ ; for triceps,  $F_{(1,42)} = 0.43$ ; for both,  $p > 0.5$ ). Although both groups learned equally well how to appropriately control the biceps and triceps, subjects in the 3 min group continued to have added difficulty controlling their posterior deltoid appropriately.

Even though the two groups train the same amount in both force fields, the subjects with 3 min separating the two fields had more difficulty generating muscle activations appropriate to B2. The lagging of the 3 min group suggests that the counterclockwise rotation of the activation function, that is, the internal model appropriate for B1, had lingered in these subjects and had hampered their ability to adapt to B2.

## ***2.8 Anterior deltoid activation functions in the null field, B1, and B2***

Compared with each group's orientation in the null field, the 3 min and the 6 hr groups have different rotations of their anterior deltoid activation functions in the last set of training in B1 (Figure 2.10). The difference between the two groups stems not from differing orientations in B1 but from differing initial orientations of anterior deltoid activation functions during the null field (Figure 2.11). Across the 24 subjects in all three groups, the variability of individual subject's activation function orientation in the null



**Figure 2.11: Orientations of anterior deltoid resultant vectors during training in the null field, B1, and B2. Whereas the previous Figures 2.8 through 2.10 represent the rotation of  $a_m$  resultant vectors with respect to null field orientations, this figure represents the actual orientation during each stage of training. The symbols indicate the mean orientations in subjects who had a 3 min (○) or 6 hr (●) separation between B1 and B2. Error bars reflect 95% confidence intervals of the mean (across subjects).**

field was almost three times larger for anterior than for posterior deltoid (size of 95% confidence intervals of the mean: biceps, 15.0°; triceps, 14.3°; anterior deltoid, 18.7°; posterior deltoid, 6.4°). Although the two groups of subjects had the same amount of training in the null field, the two groups generated anterior deltoid effective fields that pointed in different directions (3 min group mean, 113.7°; 6 hr group mean, 133.8°; in two-way ANOVA, two groups x three bins of movements in null field set,  $F_{(1,42)} = 6.69$ ;  $p < 0.02$ ). The differences in activations led to different movement profiles in the null field (in comparing mean perpendicular velocity traces in movements made toward 135° and 180°, mean correlation between groups,  $r = 0.04$ ). Once the two groups were well trained in B1, the anterior deltoid activation functions then pointed in similar directions (3 min group mean, 96.9°; 6 hr group mean, 100.0°;  $F_{(1,42)} = 0.04$ ;  $p > 0.8$ ), and the movement profiles of the two groups converged (135° and 180° perpendicular velocity correlation between groups,  $r = 0.87$ ).

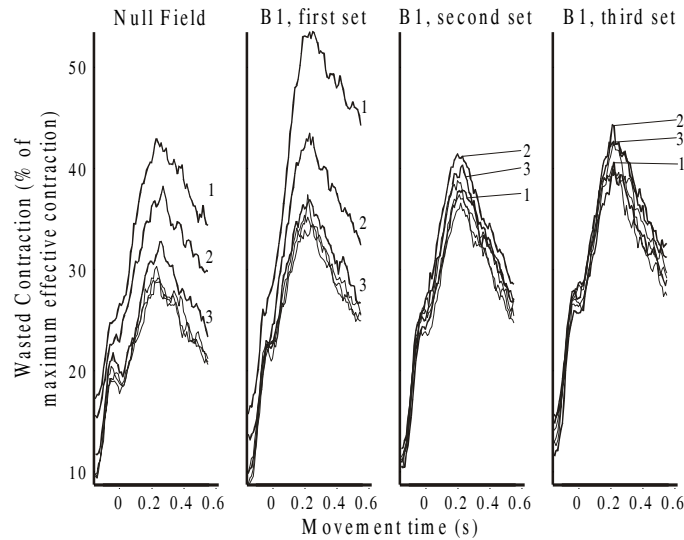
When the two groups trained in their first sets of movements in B2, the mean of the orientations of the 3 min group's anterior deltoid activation functions rotated less toward the counterclockwise direction than did the activation functions of the 6 hr group, but the difference between the groups is marginally significant (3 min group mean, 104.8°; 6 hr group mean, 125.4°;  $F_{(1,42)} = 4.02$ ;  $p < 0.055$ ). These results

suggest that despite the high subject-to-subject variation in anterior deltoid activation function orientation in the null field, the actual orientations of the anterior deltoid activation functions remained a good measure of internal model building during the learning of novel dynamics.

## ***2.9 Reduction of wasted contraction during practice in null field and in B1***

It has been shown that as subjects learn to move their wrist against novel forces, they decrease the coincident activation of flexors and extensors, thereby decreasing the stiffness of the joint (Milner and Cloutier, 1993). To quantify how co-contraction of muscles changed during learning of multijoint reaching movements, I defined a measure of wasted contraction. I partitioned subjects' EMG into composite traces using the simulation's predictions of appropriate muscle activations in the null field and field B1 (Table 3). For each movement, I used these coefficients to partition the four muscle traces into four direction-independent composite traces. Among these four composite traces, two pairs of traces produced forces that opposed each other. In each opposing pair, at each moment in time, the weaker of the two composite activations was cancelled by the stronger. I compared these activations to determine, in each movement, time series of wasted contraction and effective contraction. The level of wasted contraction decreased with practice, both during training in the null field and during learning of B1 (Figure 2.12). I averaged the time series of wasted and effective contractions in each subject across bins of 32 movements, scaled the wasted contraction as a percentage of maximum effective contraction, and then averaged these scaled wasted contractions across subjects. The peak of wasted contraction occurred ~200 msec into the movement for all bins spanning both fields. The magnitude of the peak of wasted contraction was reduced by one-third over the course of a training set in both the null field (from 43 to 29%;  $p < 0.005$ ) and the first set in B1 (from 54 to 35%;  $p < 0.005$ ). Subjects did not decrease their wasted contraction over training in the second or third sets in B1 ( $p > 0.5$ ).

During training in the null field, subjects also decreased their wasted contraction generated during the beginning of the movement, the interval over which the  $a_m$  above was calculated (average level in first bin of null, 24.5%; in last bin of null, 18.3%;  $p < 0.03$ ). The mean level of early wasted contraction also changed in the first set of B1 but not significantly ( $p > 0.2$ ). Whether the robot produced a viscous force field, the maximum levels of wasted contraction occurred consistently at the midway point of the



**Figure 2.12: Wasted contraction during movements in null field and B1. These plots show the wasted contraction (as a percent of effective contraction), averaged across subjects, as a function of time into each movement. Each line represents a bin of 32 movements. The leftmost plot contains traces from null field movements; the next three plots contain traces from the three sets of training in B1. The numbers inside the plot label the wasted contraction from the first, second, and third bins of movements in each set.**

movement. Subjects reduced their peak wasted contraction when they first learned the novel force field B1, but also, previously, when they again familiarized themselves with the well learned null field. This suggested that subjects may use wasted contraction to increase stiffness both when first regaining familiarity moving in previously learned environments and when first learning novel environments.

## **2.10 Discussion**

Subjects learned to reach in a novel viscous force field. The convergence of movements onto previous (desired) trajectories, and the existence of aftereffects when learned forces are removed, have suggested that the adaptation of an internal model underlies this learning (Shadmehr and Mussa-Ivaldi, 1994; Shadmehr and Brashers-Krug, 1997). The internal model hypothesis posits that the CNS effectively computes inverse dynamics, transforming desired trajectories into appropriate muscle activations. Here I used a simulation to predict that formation of the IM should accompany specific rotations in the directional bias of muscle activation functions. I found that the actual rotation of subjects' spatial EMG tuning curves closely matched the predictions. The simulation also allowed us to estimate the component of muscle



activation functions that produced a force that countered those in the force field. Using this transformation, I found that early in training EMG changes were driven primarily by a delayed error–feedback response. This feedback response may have formed the template for the eventual learned feedforward response.

In generating reaching movements, the CNS combines two elements of control: feedforward elements, which generate neural commands based on information available before the movement (e.g., desired trajectory); and feedback elements, which generate neural commands based on delayed visual and proprioceptive signals received during the movement. In recent computational models (Jordan, 1995; Miall, 1995; Bhushan and Shadmehr, 1999), it is assumed that feedback control provides an estimate of error, i.e., the disagreement between the actual and expected (or desired) sensory signals, and generates a delayed motor response that attempts to correct that error. Here I found that when subjects first moved in the field, their field-appropriate EMG was generated quite late in the movement, after the movement had significantly diverged from the desired trajectory. The timing of this activation relative to the perturbing force suggested that the EMG was generated by error-driven feedback control. During training, the shape of this field-appropriate EMG remained essentially invariant, but its peak gradually shifted earlier and earlier into the movement. The temporal shift culminated in a modification of the EMG pattern very early into the movement, before proprioceptive or visual information was available. This modified EMG pattern effectively predicted and canceled the upcoming pattern of imposed forces. Therefore, improvements in performance occurred as the CNS gradually incorporated the error–feedback response of the system into the feedforward command that initiated the movement.

Previous computational studies (Kawato et al., 1987; Stroeve, 1997) have shown that the motor response generated by an error–feedback system may be used to learn inverse dynamics of reaching movements. These studies suggested that a copy of the error–feedback response may be sent to the brain and then used to train the IM in a supervised learning paradigm. The results support this general framework by showing that learning may entail an incorporation of the error–feedback response into the feedforward commands through a gradual temporal shifting process.

Using this computational approach, I initially predicted the shape of the directionally tuned muscle activation functions for null field movements, with no external perturbing forces. I found that the predicted functions closely matched the measured EMG functions for the four muscles examined (Figure 2.7). I summarized the directional bias of EMG functions by calculating the orientation of these functions' resultant vectors (preferred directions). The resultant vectors remained stable as well-trained subjects performed reaching movements in the null field. However, as training began in a force field, the preferred direction of each muscle's EMG function began to rotate and eventually reached an asymptote. The internal model theory had predicted a specific rotation in each muscle's EMG function; that prediction was confirmed. The preferred directions returned to initial conditions when the force field was removed and reversed their direction of rotation when the direction of imposed forces was reversed. These results, therefore, signified an experience-dependent adaptation of the IM. These methods linked the computational concept of learning an internal model with changes in the directional tuning of muscle activation functions. After training in field B1, some subjects trained in the anticorrelated force field B2 either 3 min or 6 hr later. I had reported before that the 3 min group has much more difficulty in both their initial response and eventual adaptation to the second field (Brashers-Krug et al., 1996; Shadmehr and Brashers-Krug, 1997). Here I found that in field B2, not only are the movements made by the 3 min group more disturbed than the 6 hr group, but the preferred directions of biceps, triceps, and posterior deltoid are oriented in directions more appropriate for B1 (less appropriate for B2 ) than the 6 hr group. The lingering rotations of the preferred directions provide physiological evidence that, 3 min after training, the internal model appropriate for the field just learned is still active and biasing the subjects' ability to learn the second force field. After 6 hr have passed, the internal model no longer lingers, and subjects can learn the newest dynamics with relative ease. The results also demonstrate the robustness of the methodology in that I could compare changes in muscle activation functions despite removal and reapplication of surface electrodes in the 6 hr group.

Previous studies that have recorded EMG changes during motor learning have focused on single-joint movements and have reported increases or decreases in the magnitudes of EMG traces to demonstrate changes in neural output (Corcos et al., 1993; Gottlieb, 1994, 1996). Quantitative comparisons of these changes in magnitude with the concept of learning of an internal model are difficult because of the intrinsic

variability in the strength of the signal and the difficulty in relating the absolute magnitude of EMG and generation of force. In contrast, the methods used here focus on the directional and temporal nature of the data for multijoint movements. The directional bias and the timing of the muscle activations are independent of the overall signal strength; this independence facilitates the analysis of adaptation across subjects and the quantitative comparison of EMG and computational models.

The orientation of muscle activation functions are reminiscent of the preferred direction of primary motor cortical neurons (Georgopoulos et al., 1982, 1986) in that both summarize the dependence of neuronal activity on the direction of movement. In the motor cortex, the preferred direction of some cells changes when movements are made in different postures (Sergio and Kalaska, 1997; Scott and Kalaska, 1997). A recent study has shown that the rotation of preferred direction of some cells was similar to the rotation of the directional bias of muscle activation functions (Sergio and Kalaska, 1997). Therefore, the rotation in the muscle activation functions that occurs with learning of an internal model might echo a similar change in the preferred direction of certain cortical cells. Indeed, preliminary results (Benda et al., 1997; Li et al., 1998) suggest that the tuning of select cells in a monkey's primary motor cortex rotates when the animal learns a force field similar to the one used here. Previous work has also shown that in making single-joint movements, subjects tend to reduce the level of co-contraction of flexors and extensors as they learn to move a novel load (Milner and Cloutier, 1993). Here I quantified co-contraction because subjects could respond to novel dynamics by simply stiffening the limb through co-contraction, rendering the arm less vulnerable to perturbations (Hogan, 1984). Therefore, one might predict that during learning of a novel task, co-activation levels should start high but decline as the internal model is acquired. Here I found that co-contraction levels (termed wasted contraction to denote the EMG that was directly cancelled by opposing activation) did significantly increase from null field to when subjects were initially exposed to the force field. With learning in the force field, the co-contraction levels declined rapidly during the initial 96 movements but remained stable with further training. The same reduction in wasted contraction, however, was also observed when subjects moved in that day's first set in the null field, a well-learned dynamic environment. These results suggest that the CNS increases the limb's stiffness to reduce the effect of "unmodeled dynamics" early in the learning process and reduces limb stiffness as dynamics are learned. The increased wasted contraction in the null field extends this idea to well learned dynamics; higher

stiffness may facilitate the recollection of an internal model that is appropriate but stored in long-term memory.

Linking IM formation to rotation of the muscle-tuning functions raises two interesting possibilities regarding the computational organization of the internal model. First, I had previously reported that learning of a force field at a given arm configuration resulted in generalization at a new arm configuration (Shadmehr and Mussa-Ivaldi, 1994). Because learning a field at a given arm configuration results in a rotation of the muscle-tuning functions, it is possible that the magnitude of rotation is conserved when the arm must make movements at a new configuration. This would result in generalization of the force field in joint coordinates. Second, I have observed that certain force fields are much easier to learn than others. Because learning of each field is coupled to a specific rotation in the muscle-tuning functions, it is possible that the degree of difficulty in learning a field relates to how much each tuning function needs to rotate to represent that field. Further research is needed to examine whether the link between internal models and muscle tuning functions can be used to understand representation of IMs by the CNS.

## Chapter 3 Precurvature and cost minimization

### 3.1 Introduction

A fundamental assumption of many current theories of human motor learning is that subjects move in both perturbed and unperturbed environments using the same desired trajectories. In their seminal paper in 1994, Shadmehr and Mussa-Ivaldi quantified the learning of the viscous force field by how well movements made in the force field correlated with movements made in the initial unperturbed null field. This comparison explicitly depends on the assumption that subjects intend to move in the same trajectories in the force field as they did in the null field. Other investigators (including this one, earlier in this thesis) measure other movement parameters, such as movement length, magnitude of curvature (Goodbody and Wolpert, 1998), or area of curvature (Sainberg et al, 1999), and implicitly assume that the goal of learning is to replicate the null field behavior as quantified by these metrics.

This assumption could be validated by the veracity of the minimum jerk hypothesis. To generate a reaching movement, humans must first transform initial and desired final states into a full trajectory (Hildreth and Hollerbach, 1987; Wolpert, 1997). The initial and final states cover relatively few degrees of freedom, describable by two simple state vectors. The full hand trajectory, however, needs to be described throughout the movement, requiring state vectors to be represented at many points in time. To make the transition from the low- to the high-dimensional state, some theories posit that humans strive to minimize some cost functional throughout the movement. One of these theories, the minimum jerk hypothesis (Hogan, 1984; Flash and Hogan, 1985), propounds that people generate movements in an effort to be maximally smooth in hand coordinates. A maximally smooth movement minimizes the third derivative of hand position, called jerk. In order to produce a maximally smooth movement, humans would need to produce movements that minimize the cost function

$$J = \int_{t_i}^{t_f} \left\| \frac{d^3 x}{dt^3} \right\|^2 dt. \quad (3.1)$$

This cost function depends solely on the kinematic features of movement, and solely on the extrinsic manifestation of the motor output. Changes in dynamic task demands, therefore, should not alter desired trajectories if minimizing jerk fully informs movement planning.

A competing hypothesis suggests that instead of generating maximally smooth hand trajectories, the CNS produces movements which generate maximally smooth time series of joint torque. The minimum torque change hypothesis (Uno et al, 1989), like the minimum jerk hypothesis, argues that people map boundary conditions into full trajectories by minimizing a cost function. This cost function, however, integrates the time derivative of torque made throughout the movement:

$$J = \int_{t_i}^{t_f} \left\| \frac{d\tau}{dt} \right\|^2 dt . \quad (3.2)$$

While minimum jerk trajectories can be analytically produced from start and end points and movement time, minimum torque change trajectories can only be determined numerically (Jordan, 1995).

The typical kinematic features of short, point-to-point movements can be reasonably well explained by both the minimum jerk and the minimum torque change hypotheses. Free movements in space tend to be gently curved and smooth in position, and symmetric and bell-shaped in velocity (Atkeson and Hollerbach, 1985). Both minimum jerk (Flash, 1987) and minimum torque change (Uno et al, 1989) optimizations produce symmetric, bell-shaped velocity profiles and smooth movements. The minimum torque change hypothesis posits that the gentle curvature in movements stems from the inherently curved desired trajectories that result from minimizing an intrinsic (angular) cost (Osu et al, 1997). Minimum jerk movements are maximally straight, so any curvature in movements would at first seem to indicate that subjects do not minimize jerk. The hypothesis could still hold, however, if people misperceive their gently curved movements to be straight. A quantification of subject's perception of their own movements indeed indicated that people perceived their curved movements to be straight (Wolpert et al, 1994). Since both hypotheses explain short point-to-point movements well, the behavior of humans in other tasks needed to be measured in order to differentiate between jerk and torque change minimization.

Human performance in movements along long trajectories, through intermediate points, and around obstacles has suggested that dynamic information influences trajectory formation. In the original formulation of the minimum torque change hypothesis, Uno et al (1989) examined subjects who made 30 to 40 cm long movements in non-horizontal planes. These movements were markedly curved, as were more complex movements made through via points. A simulated controller that generated movements which minimized torque change predicted these curvatures well. In a separate group of experiments (Sabes and Jordan, 1997; Sabes et al, 1998) subjects were trained to avoid virtually projected obstacles. The handpaths of subjects changed with the angular orientation of start and end points, suggesting that subjects incorporate inertial properties of their arm into movement planning in order to steer clearly away from the obstacle. Differing task specifications seems to change the amount of curvature in movements (Desmurget et al, 1998), perhaps indicating that subjects attempt to optimize torque or some other intrinsic variable when planning movements.

The results of studies in which visual information is altered, however, indicate subjects likely optimize kinematic variables when planning movements. Subjects were trained to move their hands while a cursor represented hand position. Subjects were provided either full visual feedback which was perturbed to display fictitious curvature (Wolpert et al, 1995b; Goodbody and Wolpert, 1999), or visual feedback only at the end of the movement which exaggerated end point error (Ghahramani et al, 1996; Vetter et al, 1999). When straight movements were displayed to be curved, subjects adjusted their movements to appear straight, therefore producing trajectories that were actually curved. When end point visual perturbations were localized, adapted performance generalized in a pattern resembling a spherical eye-centered coordinate system (Vetter et al, 1999), further indicating the importance of visual information. The conservation of visually measured kinematic movement features implies that kinematics are relatively more important in movement planning than dynamics. The examination of subjects who learn with consistent kinematics but varying dynamics could provide important supplementary information: just how much dynamic cost will humans incur to conserve kinematics?


A third optimization hypothesis posits that subjects generate desired trajectories to minimize the variance of the movement endpoint (Harris and Wolpert, 1998). The simulation of Harris and Wolpert postulates

that neural commands to muscles are noisy, with variance that increases with the mean signal. Humans might then vary the intended trajectories in order to minimize the effect of the neural noise on the accuracy of movement. Although the minimum variance hypothesis was tested mainly on movements of the eye and a single arm joint, this optimization does generate bell-shaped velocities and appropriately gently curved handpaths.

These previous studies weigh the relative importance of jerk, torque change, and endpoint variance by measuring the curvature of movement, then comparing curvature to simulated trajectories that minimize kinematic, dynamic, or precision cost. The resemblance of actual performance to hypothetical ideals does provide inferences and circumstantial evidence regarding the importance of jerk and torque change. The absolute minimization of one cost, however, may produce unacceptable penalties in other costs. Myopia on one specific cost may ignore movement plans that satisfy reasonable cost limits in jerk, torque change, and variance. A more direct test of these hypotheses would be to calculate the jerk and torque change costs actually incurred during movement, along with the endpoint variance associated with actual trajectories. A calculation of the costs accrued during movement, and the reduction of these costs with learning, would directly quantify the relative importance of optimizing smooth kinematics, smooth dynamics, or endpoint accuracy.

Here I examined the kinematics and dynamics of movements made by subjects who train in the viscous force field. I found in asymptotic performance, subjects produced trajectories that differed slightly but significantly from trajectories produced in the null field. I then calculated the hypothetical trajectories that would minimize torque change in the viscous force field. Whereas minimum jerk and minimum torque change trajectories are similar for many natural movements, jerk and torque change minimization predict sharply different trajectories in the force field. The slight change in actual subjects' movements in the force field agrees with the direction, but not the magnitude, of the gross curvature in minimum torque change movements. Lastly I estimated how neural noise-induced endpoint variances, cumulative normalized jerk, and cumulative normalized torque change costs accrued changed with training.





**Figure 3.1: Mean parallel and perpendicular displacements in all eight directions of movement, averaged over the initial null field set (dots), the first bin of eight movements in B1 (gray), and the last three bins of eight movements in B1 (black lines). Subplot headings indicate the direction of movement; the subplot located 90 degrees from the center, for instance, displays mean trajectories of movements toward the target at 90 degrees. These data are plotted as rotated versions of Cartesian x- and y-position. I use the sign convention that positive perpendicular displacement reflects curvature induced by the force field, which is the reverse of the natural abscissa. Note, therefore, that these abscissae have negative values atop positive values.**

### **3.2 Asymptotic behavior in null and viscous force fields**

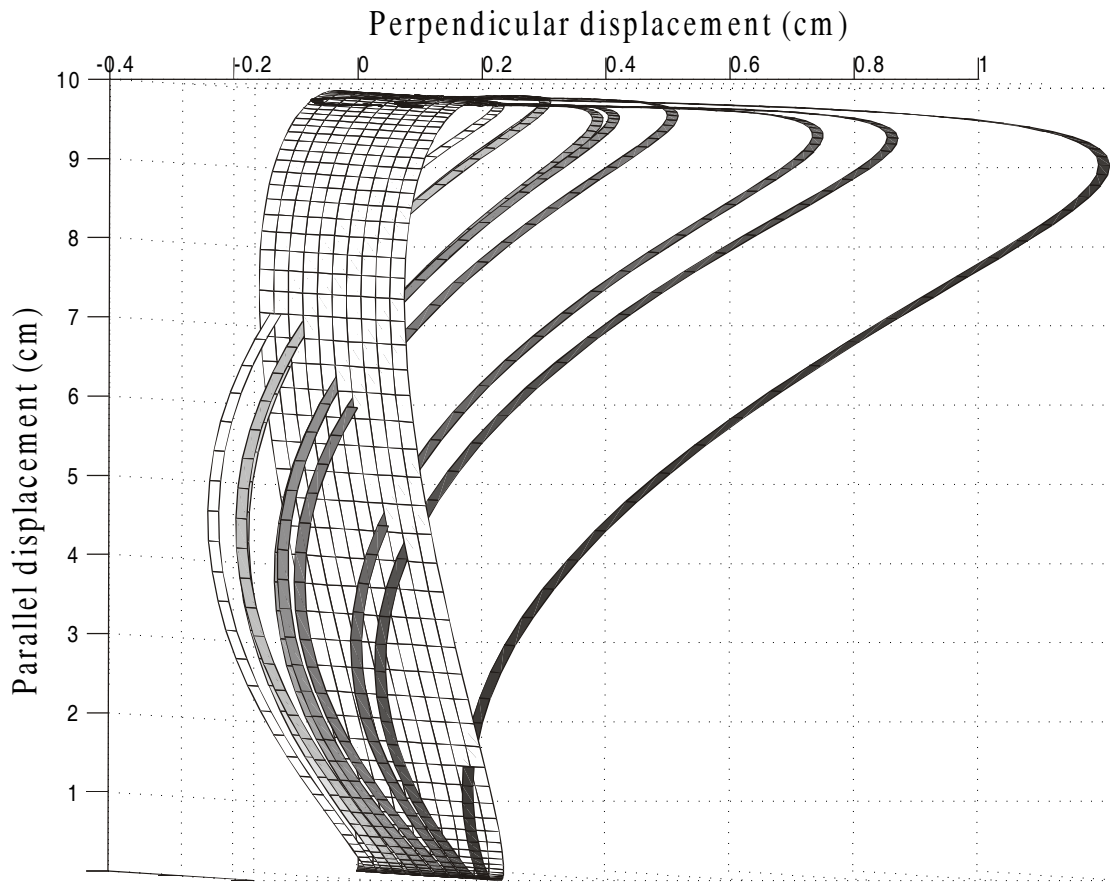
To investigate the accuracy of the assumption that subjects move in the null and force fields with the same desired trajectories, at first I simply looked at the trajectories that constant force field subjects generated in the null and force fields. Within directions, the null trajectories were averaged across the initial null field set (the first set of training on Day 2); the force field movements were averaged across eight movement

bins. I rotated each hand path from the original x and y coordinates into parallel and perpendicular displacement.

In each direction of movement, the mean null field trajectories trace a path toward the target that deviates from the straight line path by a few millimeters (Figure 3.1, dots). Although this curvature is small, the difference from the straight path is significant in most directions of movement (more details of this curvature will be provided below). When subjects first moved in the novel viscous force field, movements were strongly curved (Figure 3.1, gray lines). All of these curvatures pointed in the clockwise direction, as the force field pushes subjects perpendicular to velocity in the clockwise direction. As subjects trained in the field, movements did get substantially straighter. In movements toward 90 degrees, the mean trajectory late in training deviated from the straight line path by a maximum of less than 0.5 cm, whereas the first bin of movements had a peak deviation of almost 2 cm.

The reduction of the curvature induced by the force field, however, was coincident with a small but significant increase in curvature in the opposite direction. In all directions of movement, as a function of training, the handpaths early in movements late in training (Figure 3.1, black lines) curved away from the straight line path in the clockwise direction. This early curvature deviates not just from the straight line path, but also from the null field trajectories. In some directions, this additional curvature appeared as early as the second bin of movements, three to four minutes after the start of training. The “precurvature” of the mean trajectories deviated from the null field trajectories by only a few millimeters in each direction. This small deviation, however, clearly separates asymptotic behavior in the force field from asymptotic behavior in the null field (Figure 3.2). If one presumes that subjects converge upon desired behaviors, these small precurvatures indicate a change in desired trajectories with a change in dynamics.


Earlier in this work I characterized the erroneous curvature of movements, induced by unmodeled robot dynamics, with the metric of perpendicular displacement 250 msec into the movement. The possible precurvature of constant force field subjects, however, seems to occur very early in the movement; by 250 msec, trajectories seem to be heading back to (perhaps even crossing) the straight line connecting movement start and target. The perpendicular displacements of the constant force field subjects were



**Figure 3.2: Trajectories averaged across directions. The thick ribbon represents the mean trajectory in the null field set, averaged across all directions, movements, and subjects. Thin ribbons represent averages over bins of eight movements in each direction, then averaged over directions and subjects. Progressing from the front (black) to the back (white) signifies the progression of training in B1.**

therefore measured at 150 msec to test whether the precurvature of the group mean trajectory was discernable in individual subjects.

In certain directions of movement, there was a clear precurvature of individual subjects' movements late in training in the force field (Figure 3.3). In two directions of movement, toward 180 and 225 degrees, although the mean perpendicular displacement was more clockwise than null field trajectories, the 95% confidence intervals of the means overlapped. In the other six directions, however, the magnitude of



**Figure 3.3: Perpendicular displacement 150 msec into the movement for null field (n) and force field (B1) movements. Each data point in the force field represents the mean over an eight-movement bin. Error bars represent 95% confidence intervals of the mean.**


curvature is significantly larger at the end of force field training than during the null field set. The center plot, showing the mean behavior across all directions, shows the clearest separation between the confidence intervals of the mean. On average, in the third set of training, the magnitude of this precurvature 150 msec into the movement is 1.3 mm ( $p < 1e-5$ ). The perpendicular displacements early in the movement, then, indicate a small but significant difference between asymptotic behavior in the force field and in the null field.

### ***3.3 Calculation of minimum torque change trajectories***

The significant precurvatures suggest that the desired trajectory of these subjects is slightly different late in force field training than it was in the null field. The minimum jerk hypothesis demands that movements are planned in an attempt to minimize a purely kinematic cost function. Since the kinematic feedback of the task remains constant throughout training, significant precurvatures indicate that subjects cannot generate movements based on kinematic information alone. Although the kinematics of the task stays constant, the addition of the force field grossly alters the dynamics of the task. The torque change cost function, then, could be minimized by different hand trajectories in the null field and in the force field. A potential motivation behind precurvatures, therefore, could be an attempt to minimize some combination of the kinematic cost function of squared jerk and the dynamic cost function of squared torque change.


To investigate this possibility, I utilized the biomechanical controller in order to calculate trajectories that minimized mean squared torque change. To calculate the minimum torque change trajectory, in each direction of movement, I first parameterized hand trajectories using five floating knots spaced equally in time throughout the movement (as in Harris and Wolpert, 1998). The knots specified the position, velocity, and acceleration of the hand in Cartesian coordinates. The full vector time series of hand trajectory was therefore parameterized by a 30-element vector (2 dimensions x 3 state variables x 5 knot locations). The start and end of the movement (set at 600 msec) served as two fixed knots. The full time series of position, velocity, and acceleration was then specified by six quintics in time, each satisfying the boundary conditions imposed by each knot. I then used the biomechanical model to determine the dynamics of the quintic spline trajectory. The dynamic demands of the trajectory finally specified what torques the anthropomorphic controller would need to exert to produce the trajectory. The numerical minimization therefore consisted of varying the trajectory, as parameterized by the floating knots, in order to minimize the cumulative squared torque change.

The parallel velocities of movements of minimum jerk, minimum torque change in the null field, and minimum torque change in the force field are all similar (Figure 3.4). The minimum jerk velocity profile in each direction displays the prototypical bell shape, reflecting the analytical fourth-order polynomial relation between time and speed. The minimum jerk bells are perfectly symmetric about the central time point of



**Figure 3.4: Parallel velocities of minimum jerk (gray lines) and minimum torque change movements in the null field (thin black lines) and in B1 (thick black lines).**

the movement. The parallel velocity profile minimizing jerk in the null field has a slightly faster rise and fall of velocity, coupled with a slightly broader peak. Since torque has components dependent on velocity and acceleration, one would expect a minimum torque change trajectory to have some aspects of minimum acceleration movement profiles combined with minimum jerk. The minimum acceleration velocity profile rises very quickly and is very broad and flat midway through the movement. The minimum torque velocity profile in the null field does indeed then hint at some subtle minimum acceleration contribution. The force



**Figure 3.5: Perpendicular velocities of minimum jerk (gray lines) and minimum torque change movements in the null field (thin black lines) and in B1 (thick black lines).**

field minimum torque change trajectory also rises and falls faster and has a broader peak, but the difference between this profile and the minimum jerk velocity is even subtler. The velocity profiles of minimum torque change movements in individual directions are slightly asymmetric, with the peak arriving tens of milliseconds before or after the central time point, but the mean parallel velocity traces (averaged across directions) are symmetric.


The perpendicular velocities of movements that minimize jerk and that minimize torque change in the null field are also quite similar (Figure 3.5). The handpath of the minimum jerk movement, not surprisingly,

moves in a perfectly straight line between the start and end of the movement. The handpath that minimizes torque change in the null field is slightly curved, most noticeably in movements toward 135 and 315 degrees. In all directions of movement, the small curvature in the minimum torque change trajectories points in the same Cartesian direction for movements that oppose each other. The inverse dynamics of the human-robot interaction indicate that torque depends on functions of position multiplying joint acceleration and on Coriolis terms that depend on the square of joint velocity. Since both joint acceleration and joint velocity reverse sign between movements of opposite direction, only the Coriolis term would have the same sign in both outward and inward movements. This small curvature is therefore likely attributable to Coriolis interactions.

The perpendicular velocities of trajectories that minimize jerk and that minimize torque change in the null field are markedly different, however, from the trajectory that minimizes torque change in the force field (Figure 3.5). The handpath that minimizes torque change in the null field is grossly curved in all directions of movement. The magnitude of the curvature, as characterized by perpendicular displacement 150 msec into the movement, ranges from 0.5 cm toward 45 degrees to 1.6 cm toward 315 degrees. In trajectories toward 135 and 315 degrees, the initial movement angle differs from the target location by more than 45 degrees (Figure 3.6). The mean perpendicular displacement across directions is 1.0 cm, an order of magnitude larger than the precurvature observed in subjects (1.3 mm). The increase in curvature produces a torque change cost that is 37% smaller than the torque change associated with moving in minimum jerk (straight line) trajectory. This additional curvature, however, increases the cumulative jerk of the movement by 103%.

The large difference between actual and hypothetical minimum torque change trajectories, specifically the curvature magnitude, indicates that the amount that subjects weigh the torque change cost is likely small compared with the cost associated with jerk. All of the additional curvature generated in the force-field minimum-torque-change trajectories, however, points counterclockwise, in the same direction as the





**Figure 3.6: Handpaths of minimum jerk (gray lines) and minimum torque change movements in the null field (thin black lines) and in B1 (thick black lines).**

precurvature. This agreement suggests that precurvature in subjects might be producing cumulative torque change that is smaller than if they moved in the force field with straighter handpaths.


### ***3.4 Perpendicular force profiles: actual dynamics of force field trajectories and hypothetical dynamics of null field trajectories***

As a first step in investigating the possible reduction of torque change cost through precurvature, I studied the dynamic component of the group mean trajectories produced in the task. Since the force field produced forces perpendicular to velocity, I first examined the perpendicular forces associated with different

trajectories. Velocity traces were first averaged across bins of eight movements within directions (three bins per direction per set), then averaged across subjects. The group mean trajectories were processed through the biomechanical model to determine the dynamic output of the human controller. Since the viscous field produces force perpendicular to velocity, I isolated the component of human-generated force at the hand which pointed perpendicular to the movement direction. Both the force trace and the time derivative of force were averaged over time to determine the total motor output and to determine a dynamic cost function associated with the force field.


In order to determine the significance of the changes in trajectory between the null and force fields, in the calculation of perpendicular force, I processed the null field trajectories as if they were generated in the viscous force field. This calculation determined what motor output would have been produced and what dynamic cost would have been incurred if subjects adopted the same trajectories in the force field as they did in the null field. Comparing the actual dynamic cost incurred late in training with the hypothetical “null field trajectory” cost determines what, if any, advantage in terms of perpendicular force results from precurvature.

As suggested by my earlier EMG timing findings, during the course of training, subjects learn to generate perpendicular force earlier into the movement. In Figure 3.7, the thin black trace shows the increase in perpendicular force produced during the first 64 movements (eight movements per direction) made in the force field. Subjects produce a unimodal perpendicular force profile, mimicking the velocity profile, but the peak of perpendicular displacement arrives 100 msec after the peak in hand speed and thus after the peak in field-induced force. By the end of training in the force field, subjects produce the perpendicular force trace shown by the thick black line. This force trace is also unimodally bell-shaped, but peaks earlier, 190 msec into the movement, coincident with the peak in field-induced force. The peak height, however, is significantly shorter than the peak height of the force trace created early in training. The perpendicular force created early in training peaks at 4.7 N; late in training it peaks out at only 4.2 N ( $p < 5e-4$ ).



**Figure 3.7: Perpendicular force produced, averaged over the first 64 movements (thin black line) and the last full set of 192 movements (thick black line) of training in force field B1. Also shown is the perpendicular force that would have been necessary if the null field trajectories were produced in force field B1 (gray line). All force profiles averaged across directions and across subjects.**

Although perpendicular force production begins much earlier later in training than at the beginning of training (just as would be expected), the mean perpendicular force produced during the movement is also smaller late in training (mean difference across subjects, 0.2 N,  $p < 5e-4$ ). With training, then, subjects learn to produce force earlier and produce accurate movements while producing less total perpendicular force.



**Figure 3.8: Derivative of perpendicular force produced, averaged over the first 64 movements (thin black line) and the last full set of 192 movements (thick black line) of training in force field B1. Also shown is the perpendicular force that would have been necessary if the null field trajectories were produced in force field B1 (gray line). All force profiles averaged across directions and across subjects.**

The perpendicular force produced during actual movements in the third set are significantly different than the force profile that would have created the null field trajectories in the force field (Figure 3.7, gray line).


The peak of the two force traces occur at roughly the same time; on average, the actual force trace produced late in training peaks at 199 msec, the hypothetical “null” profile at 184 msec. The force profile necessary to produce the null field trajectory in the force field, however, is narrower and taller than the force profile actually produced. The force necessary to produce the null field trajectory peaks at 4.7 N,

significantly more than the actual force output ( $p < 5e-6$ ). Although the magnitude of “null trajectory” perpendicular force output is clearly larger than the force actually produced around the peak of force production, the actual force profile is larger early and late in the movement. Because of the breadth of the actual force profile, the mean force produced by subjects up to 500 msec into the movement is actually slightly larger than the “null trajectory” force trace (mean difference, 0.05 N); across subjects, this difference is marginally significant ( $p < 0.07$ ). Deviation from the null field trajectory, therefore, does not lessen the total integrated force necessary to move in the force field.

The deviation from the null trajectory does, however, lessen the time derivative of force over the course of the movement (Figure 3.8). The derivative of both actual and “null trajectory” force traces, as determined by a 130-msec wide Savitsky-Golay filter, shows that both the rise and fall of the perpendicular force has a smaller absolute value for the trajectories actually produced than for the hypothetical “null trajectory” traces. During the first 500 msec of movement, to generate the null trajectory in the force field, subjects would have had to average an absolute value derivative of force of  $16.7 \text{ N sec}^{-1}$ . The mean actual trajectories in the third set of force field training required only  $14.5 \text{ N sec}^{-1}$ ; this difference was significant ( $p < 1e-4$ ). When movements are combined into bins of eight movements, then, the perpendicular force traces indicate that subjects might be adopting the slightly precurved trajectories in order to lessen the torque change cost associated with hand force pointing perpendicular to movement direction by 13%.

### ***3.5 Torque change and jerk costs of individual movements***

Perpendicular force and its derivative summarizes well the changes in dynamics associated with the force field, and therefore serves as an intuitive proxy for torque and torque change. This single component of hand force, however, cannot fully encompass the torque change. Since cumulative jerk and torque change costs depend on third time derivatives of displacement, the smoothing created by averaging of movements over bins also underestimates dynamic and kinematic costs. Finally, both the analytical solution to minimize jerk and the numerical solution to minimize torque change sensitively rely on the peak speed of movement, which varies across trials. Therefore, although the analysis of binned trajectories provides a good thumbnail of kinematic and dynamic cost, a movement-by-movement analysis of jerk and torque change is necessary to quantify the relative contributions of these costs to movement generation.




**Figure 3.9: Cumulative normalized jerk during training in the force field B1 and the null field set (n) in each direction of movement and the mean across directions. Each data point in individual direction is averaged over 8 movements; in the mean plot, over 64 movements. Although the null field training precedes the force field training, the null field data is shown at the end to facilitate comparison with normalized jerk at the end of training in field B1.**

To calculate the cumulative jerk of every individual movement, each velocity vector trace was processed through a third order Savitsky-Golay filter, then summed over the entire movement data, from 100 msec before the start to the very end of the movement. To calculate the torque change of each movement, necessary joint torques were calculated using the biomechanical model, then differentiated, again with a Savitsky-Golay filter. As with the binned data, the torques calculated from null field trajectories were

processed as if the movements were made in the force field. Each jerk and torque change score was then normalized using the minimum values possible for movements of the same peak speed as the movement. Minimum jerk trajectories were determined analytically. Minimum torque change trajectories were determined for 21 peak speeds, in all eight directions of movement. Both minimum jerk and minimum torque change trajectories were resampled at 100 Hz; each resampled trajectory was then processed as was the actual movements to calculate jerk and torque changes. These steps were taken to replicate the sampling and filtering bias inherent in processing actual movements. The minimum possible torque change in each direction at each speed was numerically found by varying the knot locations, as described above. This simulation indicated a linear relation between peak speed and calculated torque change in all eight directions of movement. The minimum torque change possible for each actual movement, then, was found by interpolating between the torque changes found for the 21 calculated speeds. Finally, the cumulative jerk and torque change from each movement was normalized by the minimum possible values. Each movement was therefore scored by normalized jerk and normalized torque change. These normalized scores were finally averaged across bins of eight movements, and compared across subjects.

Although the observation of precurvature early in movements suggests that subjects would have a higher jerk in force field than catch trial movements, the actual data shows no difference in jerk between null and late force field trajectories (Figure 3.9). In all eight directions of movement, jerk seems to decrease from the first bin of movements to the last bin. In no direction, however, is there a clear separation of jerk between the third set of movements in the force field and in the null field movements. The slight precurvature should generate slightly more jerk than the less curved null field movements, but the normalized jerk scores in each of the directions do not indicate any increased jerk.


Revisiting the binned and averaged trajectories and introducing the normalization of jerk scores by peak velocity provides insight regarding the absence of jerk increase in force field movements. The perpendicular velocity of the average movement in the third set of force field training is indeed larger early in the movement than the perpendicular velocity of the null field movement (Figure 3.10a). The peak of parallel velocity of the force field movements, however, is also slightly larger than the parallel velocity in the null field (Figure 3.10b). Although the percent change in perpendicular velocity is clearly more than



**Figure 3.10: Comparing null field (gray) and third set of force field B1 (black) movements in perpendicular (A) and parallel (B) velocity, and raw (C) and normalized (D) jerk. Trajectories averaged over the entire set, all directions of movement, and all subjects.**

the percent change in parallel velocity, the absolute change in parallel velocity in m/sec is actually larger. The raw calculation of jerk in the force field movements does indicate that higher jerk is incurred in the curved, force field movements than in the straight, null field movements (Figure 3.10c). A fair quantification of jerk, however, requires normalization based on peak speed. When the two time series of jerk are normalized by the minimum jerk possible for the two different peak speeds, the early curvature no longer adds more jerk cost (Figure 3.10d). In fact, the small curvature in null field movements actually adds more normalized jerk than the larger curvature in the force field movements, since the force field movements are faster. Against intuition, the small increase in curvature actually does not increase the jerk





**Figure 3.11: Cumulative normalized torque change during training in the force field B1 and the null field set (n) in each direction of movement and the mean across directions. Each data point in individual direction is averaged over 8 movements; in the mean plot, over 64 movements. Although the null field training precedes the force field training, the null field data is shown at the end to facilitate comparison with normalized jerk at the end of training in field B1.**


of even the stereotypical binned trajectories. The lack of increased normalized jerk in the movement-to-movement analyses of force field movements, therefore, is no longer surprising.

Unlike the cumulative jerk score, the cumulative torque change does significantly differ between late force field and “null field” movements (Figure 3.11). With training, the binned normalized torque change scores decrease. By the third set of training, in four directions of movement (90, 135, 270, and 315 degrees), the

normalized torque change was significantly smaller than it would have been had subjects instead moved in the straighter, null field trajectories ( $p < 0.01$ ). In a fifth direction (180 degrees) the difference was marginally significant ( $p < 0.06$ ). When the torque change scores were averaged across directions, the difference between actual force field and hypothetical “null trajectory” torque change costs is strongly significant ( $p < 1e-4$ ). By precurving their trajectories by an average of 1.3 mm, subjects reduce their total torque change cost by 6%. Subjects therefore significantly reduce their cumulative torque change cost without increasing their cumulative jerk cost.

### ***3.6 Neural-noise induced variance and precurvature***

Harris and Wolpert (1998) suggest that desired trajectories are determined not by optimizing kinematic or dynamic smoothness, but by an attempt to maximize movement endpoint accuracy in the face of noisy neural commands. To estimate the relevance of neural noise on trajectory formation, I added noisy torque generators to the biomechanical model used above. A Simulink/Matlab implementation of the biomechanical model (provided by Maurice Smith) simulated a realistic anthropomorphic controller moving the dynamics of the manipulandum. The simulation transformed desired trajectories into feedforward motor commands using an inverse dynamic model. These feedforward commands, along with stiffness and viscosity feedback (Shadmehr and Brashers-Krug, 1997), drove the simulated plant to produce movement. Controller-produced torque at the elbow and shoulder was split into agonist and antagonist torque using Equation 2.6. Noise was added to each agonist and antagonist torque. The noise was white with zero mean and standard deviation set to 25% of the total (feedforward plus feedback) deterministic torque, the maximum of the range of noise magnitudes considered by Harris and Wolpert (1998). The accuracy of the simulated controller in its attempt to overcome torque noise was quantified by running the simulation 1000 times, then measuring the standard deviation across iterations of position immediately after the desired movement ending time (600 msec). All simulated movements were toward the target at 90 degrees. To quantify the effect of precurvature on endpoint variance, I drove the noisy model with nine different desired trajectories. I parameterized trajectory space with  $\lambda$ , where  $\lambda = 0$  corresponded to minimum jerk and  $\lambda = 1$  corresponded to minimum torque change. I then varied  $\lambda$  between  $-1$  and  $1$ , with steps of 0.25, to generate nine trajectories with different curvatures.



**Figure 3.12: Standard deviation and 95% confidence intervals of standard deviation of endpoints of simulated movements, produced by noisy agonist and antagonist torque actuators. The curvature of movements were parameterized with  $\lambda$ , with  $\lambda$  equaling 0 and 1 indicated minimum jerk and minimum torque change trajectories.**

I found that torque noise did produce some endpoint variance, but that precurvature only increased this variance (Figure 3.12). Although endpoint variance was smallest for straight movements, the increase in variance caused by curvature added a very small (<100 micron), statistically insignificant amount of standard deviation. The hypothesis of Wolpert and Harris (1998) suggested that minimizing variance motivated movement planning. I found little changes in variance with large changes in desired curvature; furthermore, since straight movements generated the smallest curvature, the goal of minimizing endpoint variance could not account for precurvature generated by subjects.

### **3.7 Discussion**

When subjects train in a consistent viscous force field, they learned to generate movements that were much straighter than their initial error. Coincident with the decrease in force field induced curvature, however,

subjects slightly increase the curvature during the beginning of movement. The increase in curvature is opposite the direction of the initial erroneous curvature. Although the magnitude of this "precurvature" is small, across subjects it is highly significant. Since well-learned force field movements are significantly more curved than catch trial movements, a first-order analysis of precurvature would suggest that subjects cannot only minimize jerk to plan movements. The magnitude and direction of curvature has previously been the benchmark for verifying or falsifying optimization hypotheses (Uno et al, 1989; Wolpert et al, 1994; Nakano et al, 1999). The change in asymptotically generated trajectories with learned dynamics, then, indicates that jerk minimization cannot by itself motivate movement planning.

Varying simulated movement trajectories revealed that in order to minimize cumulative torque change, subjects would need to produce grossly curved movements. The minimum jerk and minimum torque change solutions to natural movement tasks produce very similar trajectories (Desmurget et al, 1998). Minimum torque change trajectories in the viscous force field, however, are very different from minimum torque trajectories. Initial directions of minimum torque change movements can point as much as 50 degrees away from the target location, while minimum jerk trajectories are perfectly straight. Such gross curvature could not be confused with the millimeter-scale errors in human estimation of straightness (Wolpert et al, 1994). Although the minimum torque change trajectories are also grossly different from movements actually made late in force field training, the small precurvatures in each direction of movement point in the same direction as the increased curvature in minimum torque change trajectories. The correlation in curvature direction, but not in magnitude, suggests that subjects may combine a small regard for torque change into their plan to minimize jerk.

Previous studies differentiating kinematic from dynamic planning simply compared the curvature of actual to predicted movements. While these comparisons provided important inferences regarding movement planning, curvature alone does not quantify the costs associated with movement. Here I quantified the kinematic and dynamic costs of the actual movements produced by subjects. I first focused on perpendicular force output, as the viscous force field pushed subjects in the direction perpendicular to instantaneous velocity. Trajectories actually generated late in force field required a smaller peak of perpendicular force than if subjects would have generated their null field trajectories in the force field.

Since the actual perpendicular force profile was also broader than the “null field” force, the cumulative perpendicular force exerted was actually slightly greater than if subjects would have moved in the original null field trajectories. The broader and shorter force profile, however, associated with less cumulative force change during the movement. Since force change is closely related to torque change, the shorter and broader forces produced by subjects suggests that deviation from null field behavior could be driven by an attempt to lessen torque change cost.

To directly evaluate the relative costs of jerk and torque change, I calculated the actual kinematic and dynamic costs incurred during individual movements in the force field. Since torque change costs change with dynamics, I processed null field movements as if they were made in the force field, in order to determine if subjects changed their incurred dynamic cost by altering trajectories from their null field behavior. As subjects trained in the force field, they did reduce their torque change cost: the cumulative torque change incurred was in fact significantly less than if subjects would have generated null field trajectories (Figure 3.13a). Precurvature does seem to lessen the dynamic cost of movement in the force field.

Surprisingly, however, the additional curvature does not significantly increase the normalized cumulative jerk during the movement (Figure 3.13b). Any additional curvature does increase the raw jerk of a movement. Since jerk depends so sensitively on movement velocity, however, a fair quantification of cumulative jerk must normalize each movement’s jerk score by its peak velocity. Movements in the force field are slightly faster than in the null field; the normalized cumulative jerk cost of these faster movements is not significantly different than the null field movements. Although precurvature does indicate that subjects may consider torque change costs in movement planning, the dynamic cost savings does not seem to add any more kinematic cost.

A third optimization theory recently developed posits that subjects plan movements to minimize the variance of endpoints as caused by noisy neuromuscular commands. The proponents of the theory (Harris and Wolpert, 1998) varied trajectories in order to minimize endpoint variance, and discovered that minimum variance trajectories grossly resemble natural movements. Our realistically stiff and viscous

biomechanical controller, however, produced quite accurate movements even at the maximum estimate of torque noise. Furthermore, the precurving of movements very slightly increased the endpoint variance of movements. Precurvature, then, cannot result from an attempt to minimize endpoint variance.

Previous attempts to support or reject the minimum jerk and minimum torque change hypotheses were hampered by the high correlation between both theories' optimal trajectories in most tasks (Desmurget et al, 1998). The viscous force field, however, produces unique dynamics. The minimum torque change hypothesis predicts very curved movements, very different from the perfectly straight minimum jerk trajectories. Although subjects precurve their movements late in force field training, overall trajectories are very straight and the additional curvature does not significantly raise normalized jerk above null field levels. Subjects therefore seem to care primarily about minimizing jerk, as predicted by visual perturbation experiments (Wolpert et al, 1995; Ghahramani et al, 1996; Goodbody and Wolpert, 1999). Inside the space of movements with acceptably smooth kinematics, subjects discovered a way to lessen their torque change cost. Since precurved movements actually require slightly more force, the consideration of dynamics in movement planning seems to concern torque change, not torque itself. The planning of reaching movements then likely involves a primary requirement to move straight, followed by a secondary goal of reducing cumulative torque change.

## **Chapter 4 Human motor learning of a nonstationary dynamic environment**

### ***4.1 Introduction***

Robotic and computational theory has indicated the necessity of an accurate inverse model of dynamics to move the human arm in desired trajectories. The human motor learning of perturbing forces indicates that human subjects can adapt their internal model of inverse dynamics (Shadmehr and Mussa-Ivaldi, 1994; Lackner and DiZio, 1994). In studies of kinematic learning where the mapping between the visual representation of the limb and its proprioceptively sensed position was altered, investigators tested for the motor adaptation to kinematic perturbations by observing behavior after the perturbation was removed (Kitazawa et al, 1997). Subjects produced errors during the unperturbed period that opposed the direction of initial error during the perturbations. These “aftereffects” were deemed to be a hallmark of specific motor, rather than cognitive, learning.

Inspired by the kinematic aftereffects, investigators altered the mechanical dynamics of the limb and measured subjects’ behavior in three stages. Subjects first trained making movements in natural, unperturbed dynamics. Next subjects learned to train in novel but systematic dynamics. Finally, the novel dynamics were removed to test for possible aftereffects. Akin to kinematic aftereffects, the first discoveries of these dynamic aftereffects provided an important hallmark specifying computation of inverse dynamics (Shadmehr and Mussa-Ivaldi, 1994; Lackner and DiZio, 1994). The magnitude and direction of aftereffects suggested that inverse models are learned in joint space (Shadmehr and Mussa-Ivaldi, 1994), learned dynamics was directionally local (Gandolfo et al, 1996) but generalize across speeds and movement length (Goodbody and Wolpert, 1998), and learned dynamics may not have an explicit temporal component (Conditt and Mussa-Ivaldi, 1999). The temporal evolution of aftereffects generated by several groups of subjects also suggested that different short-term and long-term motor memories likely underlie human behavior (Shadmehr and Brashers-Krug, 1997).

In order to quantify the state of the inverse model during training, some of investigators interspersed “catch trial” movements during novel force training (Shadmehr and Mussa-Ivaldi, 1994; Shadmehr and Brashers-Krug, 1997). In these catch trials the novel forces were unexpectedly removed, resulting in aftereffects. Other investigators interspersed the usual novel dynamics with other novel dynamics, in order to test the generalization across movement parameters (Goodbody and Wolpert, 1998) or to investigate the interaction between feedforward, feedback, and postural controllers during movements in unexpected dynamics (Sainburg et al, 1999). In these studies the interspersing of surprise movements is intended to interrogate the feedforward control system while it is learning the usually applied dynamics.

These occasional, sudden changes in dynamics, however, introduce a nonstationarity to the training paradigm. A trademark of motor learning is that testing and training cannot be separated. Although investigators introduce sudden dynamic changes to interrogate the human motor learning system, these changes very likely also train the inverse model. Each catch trial, then, could induce an unlearning of the learned dynamics in the usual environment. When a block of learning one dynamic environment is followed closely in time by a block of another environment, the second training interferes with the memory of the first dynamics (Brashers-Krug et al, 1995; Brashers-Krug et al, 1996). On a shorter time scale, catch trials may have a similar disruptive effect on the very learning that catch trials were intended to measure.


Although catch trials may disrupt the human motor learning system, this very disruption may reveal important characteristics of that system. Two classic methods of system identification quantify the steady state and the impulse responses of unknown systems. The differential responses to consistent, long-term input and short bursts of greatly changing input provides complementary, sometimes supplementary, information regarding the system’s inner workings. In neural network implementations of inverse model learning, the calculation of the inverse model output has depended on weights multiplying basis functions which depend on the desired trajectory of the movement (Kawato et al, 1987; Schaal and Atkeson, 1998; Sanner and Kosha, 1999). These architectures imply motor learning through weight adaptation. The weight adaptation that minimizes squared error depends on both movement error and the basis function generated for that movement. Since the torque output from the inverse model depends on the updated weights multiplying the basis functions, the relation between error and subsequent performance may



provide important information regarding the basis functions that underlie inverse dynamic models. The occasional but sudden changes in input created by catch trials supplies the necessary errors to investigate the properties of these basis functions.

This chapter will explore the effect of catch trials on the learning of the viscous force field. I quantified the kinematic ramifications of force field removal from movement-to-movement and across the training session. Electromyographic traces indicated how neural output changed with each catch trial, and how the overall training in nonstationary dynamics differentiated subjects' neural response from subjects who trained in the constant force field. I developed a state-space model to identify the effect of force field undulations on performance, and to measure the extent to which the effects generalized across movement direction. I developed the theoretical relation between dynamic error and subsequent dynamic output of the inverse model. The kinematic parameters elicited from the state-space model were processed through the biomechanical model to determine errors in torque space. These errors were finally examined, using the theoretical result, to discover the spatial extent of basis functions underlying the inverse model. This spatial extent provides insight into the relative contributions of previously hypothesized basis functions to the human motor system's response to sudden dynamic changes.

To determine the effect of nonstationarity on motor learning, I examined the behavior of 40 subjects who trained in the viscous force field that was occasionally turned off in catch trials. Catch trials occurred during 33 of the 192 movements in each force field set. The catch trials arrived pseudorandomly throughout the set, but during the same movements in every set in every subject. Data from 27 of the 40 subjects were retrospectively analyzed, as they were trained for previous studies (Shadmehr and Brashers-Krug, 1997). In the other 13 subjects, EMG was recorded, processed, and analyzed as in Chapter 2.



**Figure 4.1. Perpendicular displacement 250 msec into movement, during null field and force field training; means over constant force field (<const>) and catch trial (<ctg>) groups, and the subset of fielded movements before (ct-1) or after (ct+1) catch trials.**

## **4.2 Kinematic performance during training with catch trials**

### **4.2.1 Performance over the duration of training**


Previous investigators have trained subjects in force fields with catch trials, presuming that occasional removal of the forces did not affect the long-term performance of subjects. To test this presumption, I calculated perpendicular displacement generated by the 40 subjects in the catch trial group during movements made in the force field, excluding catch trial movements. I binned perpendicular displacements across directions and over bins of 64 movements. The catch trial subjects' movement curvature was then compared to the 24 other subjects who trained in a constant force field, without catch trials.

Throughout training in the force field, the subjects who trained with catch trials produced significantly more curved trajectories than did subjects who constantly trained in the force field (Figure 4.1, solid lines). During the initial null field set, both sets of subjects produced very similar, small curvatures in their

movements. In the very first bin of 64 movements, during the first two minutes of training in the force field, constant force field subjects produce an average perpendicular displacement of 0.65 cm, whereas subjects who trained with catch trials produce 0.87 cm of perpendicular displacement. Although both groups of subjects improve over the next two minutes of training, the group without catch trials improves more over their initial two minutes' performance than the catch trial group improves. This gap becomes wider and wider until the end of the three sets of training. Although the catch trial group experiences catch trials only one movement out of six, the occasional removal of the catch trials significantly affects the performance of the catch trial subjects throughout the 20 minutes of training.

The catch trials affect the overall performance during training; do they also affect performance in the shorter term? Over the course of training, do the catch trials generally hinder the performance of all force field movements, or is there a specific effect after each catch trial? To address this question I examined specific subsets of force field movements. Inside each direction of movement, I identified catch trial movements, then examined force field movements that immediately preceded and immediately followed each catch trial. Note that consecutive movements in each direction of movement are always interleaved with at least one, usually several, movements to other directions. For this analysis (and all analyses below until Section 4.5) I indexed movements within each direction. A movement that I identify as preceding or following a catch trial movement are the movements that occur most recently before or soonest after the catch trial, in the same direction as the catch trial. I am therefore investigating how catch trials could possibly affect force field movements in the same direction as the catch trials.

In each of the 64 movement bins (8 movements in each of 8 directions), I binned together the perpendicular displacements of the subsets of force field movements that immediately preceded or immediately followed catch trials. During the first 64 movements, the perpendicular displacements of movements preceding and following catch trials are quite similar (Figure 4.1, dotted lines). Starting with the next 64 movements, however, the curvature of movements preceding catch trials is noticeably, if not significantly, different from the mean curvature in fielded movements. These valleys of curvature, however, do not approach the performance of the constant force field subjects. During the third set of training, movements made just before catch trials average 0.30 cm of perpendicular displacement, 0.42 cm more than the average



**Figure 4.2. Perpendicular displacement 150 msec into movement, as a function of movement direction, for null field movements (solid) and movements preceding (lower dotted line) and following (top dotted line) catch trials.**

curvature during their movements in the initial null field set. These movements have 0.12 cm less curvature, on average, than the total mean for force field movements, a significant difference ( $p < 0.03$ ). Even just before catch trials, however, subjects with catch trials move with much more curvature, with 0.49 cm more perpendicular displacement, than do constant force field subjects moving in the third set of training ( $p < 1 \times 10^{-6}$ ).


The perpendicular displacement data, averaged across bins of 8 movements per direction and averaged across directions, suggests that catch trials have two effects on the performance of human subjects. Each catch trial might induce a jump in curvature, generating a short term unlearning of the force field. This short-term unlearning, followed by a recovery to the next movement just preceding a catch trial, seems to produce an oscillation in kinematic performance. The cumulative effect of all the catch trials seems to elevate the asymptotic level of curvature higher than the final curvature of the constant force field subjects. Even the bottom of the curvature envelope sits clearly above the performance subjects who train in the constant force field.

The asymptotic behavior of subjects who trained in the constant force field was characterized by not only reduced curvature in the direction of perturbing forces, but a slight early precurvature in the opposite (counterclockwise) direction. I therefore also examined the perpendicular displacements of movements made by catch trial subjects, measured 150 msec into the movement, to determine whether these subjects exhibited any precurvature. I examined again the subset of movements that immediately preceded and followed catch trial movements. The early perpendicular displacement was averaged within each direction of movement, across the second and third sets of training. These displacements were then compared to the mean null-field perpendicular displacement generated by the 24 subjects in the constant force field group.

In none of the eight directions of movement did the early perpendicular displacement of movements preceding catch trials significantly differ from null curvature (Figure 4.2; solid lines, null; upper dotted lines, after catch trials; lower dotted lines, before catch trials). In six out of the eight directions of movement, the mean early curvature was actually greater in the force field movements than in the null field. The mean perpendicular displacement in movements preceding catch trials was actually larger than in the null field, but by only 0.03 cm and insignificantly ( $p > 0.9$ ). There was no evidence, therefore, that the catch trial subjects adopted the precurvature that the constant force field subjects did. The catch trials, therefore, hamper the ability of subjects to fully reduce the perturbations caused by the force field, but also seem to prevent the adoption of precurvature to lessen the dynamic cost of moving in the force field.

#### 4.2.2 Movement-to-movement kinematic effects of catch trials


I next explored whether the possible unlearning due to catch trials could be identified in individual movements. Figure 4.3a shows the hand trajectories during five movements of a single subject toward a single direction, 90 degrees. The very first movement in the viscous force field is toward this direction. When this subject first moved in the novel force field (trajectory labeled “1st”), her handpath was markedly curved compared to the null field trajectory (“n”). By the midpoint of her second set of training, during the 263rd movement overall and 38th movement in this direction, the handpath is much straighter (“ct-1”). The next movement in this direction (the 283rd movement overall) is a catch trial; this trajectory (“ct”) is quite curved, but in the opposite direction of the initial error (“1st”). The next movement toward 90



**Figure 4.3. Handpaths of movements toward 90 degrees in single subject (A) and mean across subjects (B); mean in null field (n), first movement in B1 (1st), and, mid-training, movements just before (ct-1) and after (ct+1) a catch trial (ct). Error bars in B indicate 95% confidence intervals of the mean of parallel and perpendicular displacement 250 msec into movement.**

degrees (287th movement overall; 40th movement toward 90 degrees) is more curved than the movement toward 90 degrees preceding the catch trial.

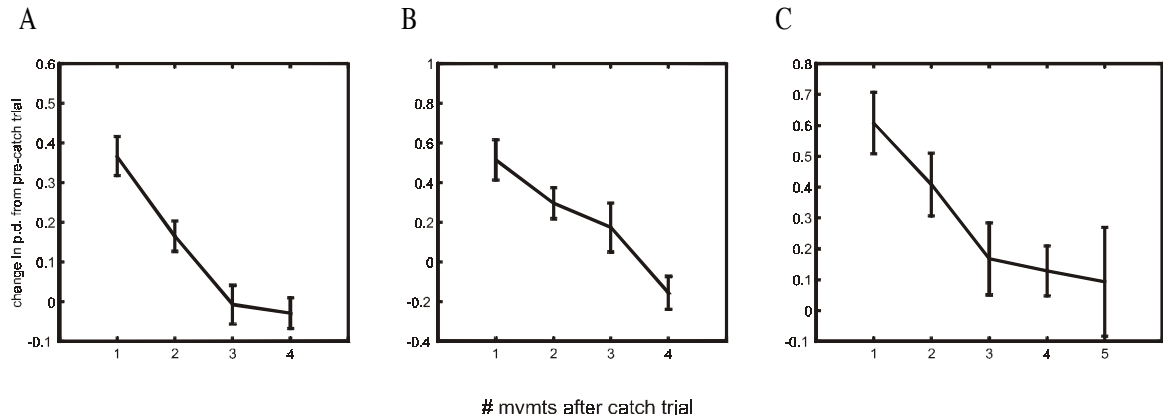
Across all subjects who trained with catch trials, this particular movement toward 90 degrees had greater curvature than the movement toward 90 degrees preceding the catch trial (Figure 4.3b). From the 1st movement in the force field to the 38th, subjects' perpendicular displacement decreased from 2.38 cm to 0.45 cm (95% confidence intervals of the mean (CIM), [2.15, 2.61] cm for the 1st movement and [0.30,



**Figure 4.4. Perpendicular displacements in catch trial movements (A) and change in perpendicular displacements across catch trial movements (B) in all directions of movement. Error bars reflect 95% confidence intervals of the mean.**

0.60] cm for the 38th movement). In the next movement toward 90 degrees, a catch trial, subjects produce a large negative perpendicular displacement (mean  $-1.35$  cm, 95% CIM [1.14, 1.56] cm). In the next movement toward 90 degrees, the perpendicular displacement significantly increased to 1.22 cm (95% CIM of the increase, [0.54, 1.00] cm). Before this particular catch trial, after 37 movements toward 90 degrees and 262 movements overall, subjects had eliminated 81% of their perpendicular displacement. After the catch trial toward 90 degrees, subjects increased their curvature substantially, losing over half of the improvement in their curvature. These trajectories suggest that the catch trial induced a short-term unlearning of the force field that causes the increase in curvature.

In all directions, catch trials caused an increase in perpendicular displacement in subsequent (force-field) movements in the same direction. Within each direction of movement, within each subject, we averaged perpendicular displacements across catch trials occurring during the second and third sets of training in the force field. We also calculated the increase in curvature between the movements preceding and following the catch trial, in the same direction as the catch trial. Figure 4.4a shows the mean and 95% CIM for the perpendicular displacements of catch trial movements. In all eight directions, perpendicular displacement was significantly less than zero; the mean magnitude of the curvature ranged from less than 1 cm for catch trials toward -180 degrees to more than 1.5 cm for catch trials toward 45 degrees, 90 degrees, and 135




**Figure 4.5. Latency of post-catch-trial increase in perpendicular displacement: mean across directions (A) and in movements toward 90 (B) and 135(C) degrees.**

degrees. Figure 4.4b shows the mean and confidence intervals of the mean for the change in curvature between force field movements that sandwich catch trials. In all eight directions, perpendicular displacement increased after a catch trial. The mean jump in perpendicular displacement ranged from 0.2 cm in movements toward -180 degrees to greater than 0.5 cm in movements toward 135 degrees. The spread of after-effect magnitude in catch trials and of catch-trial induced jumps in curvature suggests that more unlearning occurs in directions which have large after-effects.

The increase in curvature after catch trials decayed quickly during subsequent force field movements. I next examined consecutive force field movements following catch trials, once again focusing on the order of movements only within direction. All force field movements that occurred two, three, or four movements after the most recent catch trial movement were grouped together. All such movements during the second and third sets of training were combined within subjects. The typical jump in perpendicular displacement after catch trials, averaged across directions, is 0.35 cm with respect to the movement before the catch trial (Figure 4.5a). In the next force field movement, the additional curvature has been cut in half. No increased curvature is discernable in the third or fourth movements after the catch trial. In two specific directions of movement, increases in curvature lingered for three movements (90 degrees, Figure 4.5b) or four movements (135 degrees, Figure 4.5c), but did not linger beyond that. The kinematic manifestation of catch-trial induced unlearning seems to be short-term, lasting only two to four movements after the catch trial.






**Figure 4.6. Timing of B1-appropriate EMG activation during training. A: Evolution of B1-appropriate EMG over the course of training. From top to bottom, subplots represent mean B1-appropriate activation averaged over bins of 64 movements (eight in each direction) and over subjects. The dotted lines in the first and last subplots represent the force produced by the field. The dotted line traversing the subplots identifies the location of the activation peak, as quantified by a fourth-order polynomial fit. B: Timing of peak EMG as a function of training. Each data point includes 64 movements in the force field. Error bars represent 95% confidence intervals of the mean.**

### **4.3 Electromyographic correlates of catch trial-induced unlearning**


To determine the neurophysiological basis for the differences in behavior induced by catch trials, I analyzed the electromyographic traces generated by catch trial subjects and compared their metrics of learning with constant force field subjects. I began by analyzing the B1 appropriate component of activation. As with the constant force field subjects, I binned the B1 appropriate EMG across eight movements per direction, then averaged across movement directions (Figure 4.6a). At the beginning of training in the force field, the catch trial subjects, like the constant force field subjects, generated appropriate activation which peaked 200 milliseconds into the movement. Since the force induced by the field peaked at 190 milliseconds and there was a delay of 80 – 100 msec between recorded EMG and force production (from Chapter 2), this activation came too late to compensate completely for the force field.



**Figure 4.7. B1-appropriate EMG in movements before (ct-1) and after (ct+1) catch trials.**

Between the first and second bins of movements, the appropriate activation significantly shifted to earlier in the movement, to 145 milliseconds. However, throughout the rest of training the appropriate activation peak remained at 145 milliseconds (Figure 4.6b). Constant force field subjects moved their peak of activation all the way back to 70 milliseconds – significantly earlier than the catch trial group. The shift in the timing of the peak of appropriate activation in constant force field subjects (Section 2.4) marked a transfer of dependence on feedback control to feedforward control. The later activation of the appropriate musculature indicates that throughout training catch trial subjects may rely more heavily on feedback control.

The kinematic features of movements made by catch trial subjects shows that, although movements improved over the course of training, catch trials induced a short term unlearning of the force field. To focus on the neurophysiological basis of this short term unlearning I binned movements in the same direction as catch trials but arriving one movement before or after the catch trial (Figure 4.7). Force field appropriate EMG peaked 5.3 msec later after the catch trial than before the catch trial (162.8 vs. 157.5 msec), but this shift was not significant ( $p>0.3$ ). The integrated activation under these traces, however, are significantly different. In particular, the activation occurring during the agonist burst (-50 to 100



**Figure 4.8: Rotation of movement-initiating EMG resultant vectors from null field orientations, before (ct-1), during (ct), and after (ct+1) catch trials. Error bars reflect 95% confidence intervals of the mean.**

milliseconds) is significantly larger before the catch trial (mean increase 2.7 normalized EMG units,  $p < 0.03$ ). This suggests that although the timing of the activation does not change across the catch trial, the ability to produce appropriate feed forward control may have been partially unlearned.

To quantify this possible unlearning, I constructed resultant vectors of movement initiating EMG produced by each muscle in each direction of movement. As with the constant force field subjects the rotation of resultant vectors away from their initial null field orientation quantifies the learning of the inverse model. In movements preceding catch trials, the EMG resultant vectors have all rotated significantly away from their null field orientations (Figure 4.8). The magnitudes of all four muscles' rotations were not significantly different from the orientations of the constant force field group. During catch trials, as previously noted, subjects generated large curvature in the opposite direction of initial error. Despite this kinematic aberration, the orientation of the movement initiating activation resultant vectors remained unchanged from the previous movement. This result confirms the assumption that what I term the agonist

period of activation represents feedforward control that is unaltered by catch trials. In movements following catch trials, each resultant vector points in a direction less appropriate for the force field than before the catch trials. This change in orientation is significant in all four muscles. Analysis of the force field appropriate EMG traces and resultant vectors of movement initiating activation indicates that while catch trial subjects rely on feedback control throughout training, catch trials induce specific unlearning of appropriate feed forward control, indicating short term unlearning of the inverse model.

#### ***4.4 Modeling force field learning and unlearning using a state-space representation with scalar input***

My analysis of movements has revealed that subjects learn to reduce the large curvature produced in initial exposure to the force field; that during catch trials, subjects produce curvature in the opposite direction of initial error; and that force field movements subsequent to catch trials in the same direction exhibit increased curvature. We used these three characteristics as a guide to build descriptive state-space models of learning and unlearning the force field, one model per direction of movement. These models fit the relationship between the input, the application or removal of the force field; and the output, curvature of movements as measured by perpendicular displacement 250 msec into the movement.

Each direction of movement was modeled independently with a different state-space model. Each model was described by the pair of equations:

$$x_{n+1} = ax_n + bu_n \quad (4.1)$$


$$y_n = x_n + du_n, \quad (4.2)$$

where  $y$  represented perpendicular displacement and  $x$  represented the the contribution to perpendicular displacement generated by the feedforward output of the CNS. The state of the system,  $x$ , at time step  $n+1$  depended only upon the last state ( $x_n$ ) and the last input value ( $u_n$ ). The parameter  $a$  determined how the previous feedforward input affects the current feedforward input. For values of  $a$  between 0 and 1, this component of the state equation describes an exponential decay; if the system were not driven by input, the

solution to the discrete system  $x_{n+1} = ax_n$  would be a pure exponential decay with a time constant of  $\frac{-1}{\ln(a)}$  movements. The input,  $u$ , was a binary measure of the force output of the robot;  $u$  was equal to 1 for movements in the force field, -1 for catch trial movements (for computational purposes, the full input time series was mean corrected). The parameter  $b$  quantified the difference in adaptation following force field or following catch trial movements. Note that, as in the kinematic and EMG analysis, the characterization of movements as “following catch trial movements” refers to the order of movements within the modeled direction, and ignores movements that occur in other directions; a vector input model below will consider the influence of catch trials across directions. The adaptation of the state, representing the feedforward control of movement, depended solely on information from the previous time step; the system had no explicit memory of earlier states. The state adaptation also depended solely on previous state and the binary measure of whether, during the previous movement, the force field was on; it did not depend on the curvature of the movement itself.

The second equation (4.2) mapped state and robot dynamics into curvature. Since the state represented the portion of curvature that resulted from feedforward control, the state is mapped directly into output. The second term describes the relationship between robot dynamics and kinematic output. The parameter  $d$  quantified the difference in perpendicular displacement between force field and null field (catch trial) conditions. This parameter multiplied the same binary input  $u$  as in the state equation. The two equations of the state-space formulation, then, model an adaptive controller and its interaction and adaptation when controlling a system with alternating dynamics.


In each direction of movement, the relationship between perpendicular displacements and the binary inputs of robot dynamics was fit with a state-space model. All fits were calculated using Matlab and its System Identification Toolbox. Inputs to the fit were perpendicular displacements  $y$  and the pattern of force field movements and catch trials as flagged by  $u$ . The fit used Equations 4.1 and 4.2 to produce the hidden variable  $x$  and estimate parameters  $a$ ,  $b$ ,  $d$ , and the initial value of  $x$ , termed  $x_0$ . The original data pool of 40 subjects was resampled with replacement 1000 times. Each resampled data set was averaged across subjects, then fit to generate 1000 models. Since any model representing the large negative curvatures of



**Figure 4.9: Perpendicular displacements 250 msec into movements toward 90 degrees, fitted by scalar-input (thin line) and vector-input (thick line) state-space models, as a function of movement number. Dots represent average across subjects, mean corrected (for computational purposes).**

catch trials would correlate well with the entire series of perpendicular displacements, the quality of each fit was quantified by the correlation between the predictions and the actual perpendicular displacements of the force field movements. The 1000 estimates of each parameter were sorted by rank; the 26<sup>th</sup> and 975<sup>th</sup> elements of the sorted parameters provided a bootstrap estimate of the 95% confidence interval of the mean.

The scalar input state-space model provided a significant and reasonably good fit of the series of perpendicular displacement in all eight directions of movement. Figure 4.9 shows the mean perpendicular displacements for movements made toward 90 degrees. Both the plotted data and all data fit with the models were mean corrected. The best scalar-input state-space model fit to this data is displayed with the



**Figure 4.10: Parameters of the scalar-input state space model for all eight directions of movement. “Mean” represents the mean subtracted out from the perpendicular displacement data before processing; “ $r^2$ ” represents the correlation between the model fit and the actual displacements during force field movements. Error bars represent 95% confidence intervals of the mean as determined by bootstrapping.**

thin line. The model output represents well three key features of the variation in perpendicular displacements: the initial large error decays quickly, large negative curvatures are generated by catch trials, and movements following catch trials have larger positive curvatures. The correlation ( $r^2$ ) between the force field movements and the model output equaled 0.82 (95% CIM, [0.77, 0.85]).

The parameters of the best fitting scalar input state-space models reveal that directions with larger aftereffects have larger catch-trial induced unlearning. Figure 4.10 shows the parameters, subtracted

means, and qualities of fit for scalar state-space models fit to the eight directions of movement. The variance of  $a$  across directions shows that error generated upon initial exposure to the force field decays more quickly in certain directions than others. The relationship between  $b$  and  $d$  relates the kinematic effects of unlearning to the magnitude of the catch trial error. These parameters are highly correlated ( $r^2 = 0.81$ ). This correlation suggests that when perpendicular displacement in the catch trial is large, unlearning due to the catch trial is also large. Large catch trial errors, then, could induce large changes in the weights underlying inverse model formation, producing larger unlearning of the appropriate force field.

#### ***4.5 Modeling force field learning and unlearning using a state-space representation with vector input***

In the scalar input formulation of the state-space model,  $x_{n+1}$  is influenced only by its prior state and by whether the previous movement was a force field or a catch trial movement. Movements in each target direction were modeled separately; each model considers as input only recent dynamic history in the modeled direction. These models fit well the initial learning of the force field, the negative curvature due to catch trials, the jump in positive curvature after catch trials, and the slow decay in curvature after catch trials. This model, however, allows only for a very stereotypical response to the catch trials. The actual perpendicular displacement produced by subjects reveals some variance in the magnitude of negative curvature in catch trials and the magnitude of the increase in curvature after catch trials. In addition, in between the jump in curvature and the next catch trial, the mean perpendicular displacement often does not decay monotonically, but rather has some bumps and valleys.

Whereas the scalar input formulation incorporates the pattern of catch trials and force field movements in the modeled direction, these movements are interleaved between force field movements and catch trials in all the other directions. To investigate whether the movement-to-movement variance in catch trial and force field curvature could be attributed to experience in these other directions, I developed a second class of state-space model which, for every time step, inputs a vector which represents immediate history of forces in all directions of movement. In the previous model, the scalar input  $u$  was equal to 1 for force field movements and  $-1$  for catch trial movements. Since the state equation relates feedforward control at time  $n+1$  to state and input at time  $n$ , the scalar  $b$  indicated the differential influence of force field and catch trial



movements on subsequent feedforward control. I generalized this concept to include an eight-dimensional vector input  $\vec{u}$ , one dimension per direction of movement:

$$x_{n+1} = ax_n + \vec{b} \cdot \vec{u}_n. \quad (4.3)$$

Since this input vector will differentiate the feedforward control in the modeled direction  $i$  between time steps  $n$  and  $n+1$ ,  $\vec{u}$  at time step  $n$  actually indicated the existence of catch trial or force field movements between movements  $n$  and  $n+1$  in direction  $i$ . If no movement occurs in direction  $j$  between movements  $n$  and  $n+1$  in direction  $i$ , then the  $j$ th element of  $\vec{u}$  at time step  $n$  equals 0. If there are one or more movements in direction  $j$  between movements  $n$  and  $n+1$  in direction  $i$ , only the movement that occurs the latest (the closest in time to movement  $n+1$  in direction  $i$ ) counts. If this movement is a force field movement, then the  $j$ th component equals 1; if it is a catch trial movement, then the  $j$ th component equals -1. The parameter vector  $\vec{b}$  therefore maps recent dynamic history in all directions of movement into subsequent dynamic behavior in the modeled direction.

The vector-input state space model was fitted to the mean perpendicular displacement across subjects. The thick line in Figure 4.9 compares the fit of the scalar and vector input state space models to the mean (across subjects) perpendicular displacement of all movements made toward 90 degrees. The vector input model, like the scalar input model, fits well the initial reduction of curvature, the large negative curvature due to catch trials, and the increase in curvature after catch trials. In addition, the vector-input model seems to capture more of the smaller variation in the curvature of force field movements. For instance, between the catch trials in movements 46 and 55, the vector model seems to capture more of the bimodal sequence. In addition, the magnitude of the post-catch-trial increase in perpendicular displacement varies in the model, capturing some of the actual variance in the data.


In fitting the model to the force field movements made in all three sets toward 90 degrees, the vector input model, not surprisingly, accounted for more of the variance in perpendicular displacement than did the scalar input model ( $r^2 = 0.91$  for the vector model, [0.87 0.93] 95% CIM). The improvement in fit was more impressive, however, for the force field movements made during the second and third sets of training

( $r^2 = 0.68$  for the scalar model, [0.55 0.75] 95% CIM;  $r^2 = 0.87$  for the vector model, [0.79 0.91] 95% CIM).

The models of all directions of movement share this characteristic, fitting the second and third sets of training particularly well (mean of  $r^2$  across directions in scalar model = 0.60, [0.49 0.69] 95% CIM; in vector model,  $r^2 = 0.81$ , [0.72 0.87] 95% CIM). The expansion of the input states seems to fit the initial recovery from the first exposure to the field about as well as the scalar input model. The improvement in fit seems concentrated on curvatures generated during the second and third sets, when the overall envelope of performance has reached an asymptote and the variance in performance seems to be dominated by catch trials.

The improvement in the quality of fit is somewhat unremarkable, as we are increasing the number of parameters of each state space model from four ( $a, b, d, x_0$ ) to eleven. To investigate whether the expansion to vector input actually revealed a meaningful dynamic of learning, I bootstrapped the curvature data across subjects to produce 1000 samples. I averaged over each of these resamplings, then fit each mean with the vector state space model. The resulting statistics of the parameters would determine whether the interactions in the model were significant. In particular, I focused on the statistical significance of the bootstrapped values for the vector parameter  $\vec{b}$ , which in each modeled direction fit the differential effects of catch trials and force field movements in all directions. Components of  $b$  which, across the bootstrapping, were significantly different from zero would indicate that the pattern of forces in those directions have a significant effect on feedforward control, rejecting the null hypothesis that the parameters merely fit each resampling without any other significance.

In Figure 4.11, the mean and 95% CIM of the values for  $\vec{b}$  are plotted for each modeled direction. The location of the subplot indicates the direction of movement modeled. Within each subplot, "0" indicates the same direction of movement as the direction modeled; for example, in the subplot corresponding to 90 degrees, the data point at "0" represents the value of  $\vec{b}$  for 90 degrees, quantifying the differential effect of catch trials and force field movements toward 90 degrees on movements made toward 90 degrees. The x-location of the remaining data reflect how catch trials and force field movements made in other directions affect each modeled direction, as a function of each direction's distance away from the modeled



**Figure 4.11: State-space model parameter vectors  $\vec{b}$ , in each modeled direction of movement. The x-axis represents the difference in angle between the input and output directions. Negative angle differences indicate that the input direction is farther clockwise than the input direction. Error bars represent 95% confidence intervals of the mean as determined by bootstrapping.**

direction. In the subplot for 90 degrees, for instance, the data at  $x$  equal to -45 and 45 degrees correspond to the components of  $\vec{b}$  for movements made toward 45 and 135 degrees, respectively. In movements toward 90 degrees, then, the unlearning induced by catch trials occurring at 45 or 135 degrees is much smaller than unlearning induced by catch trials occurring at 90 degrees.

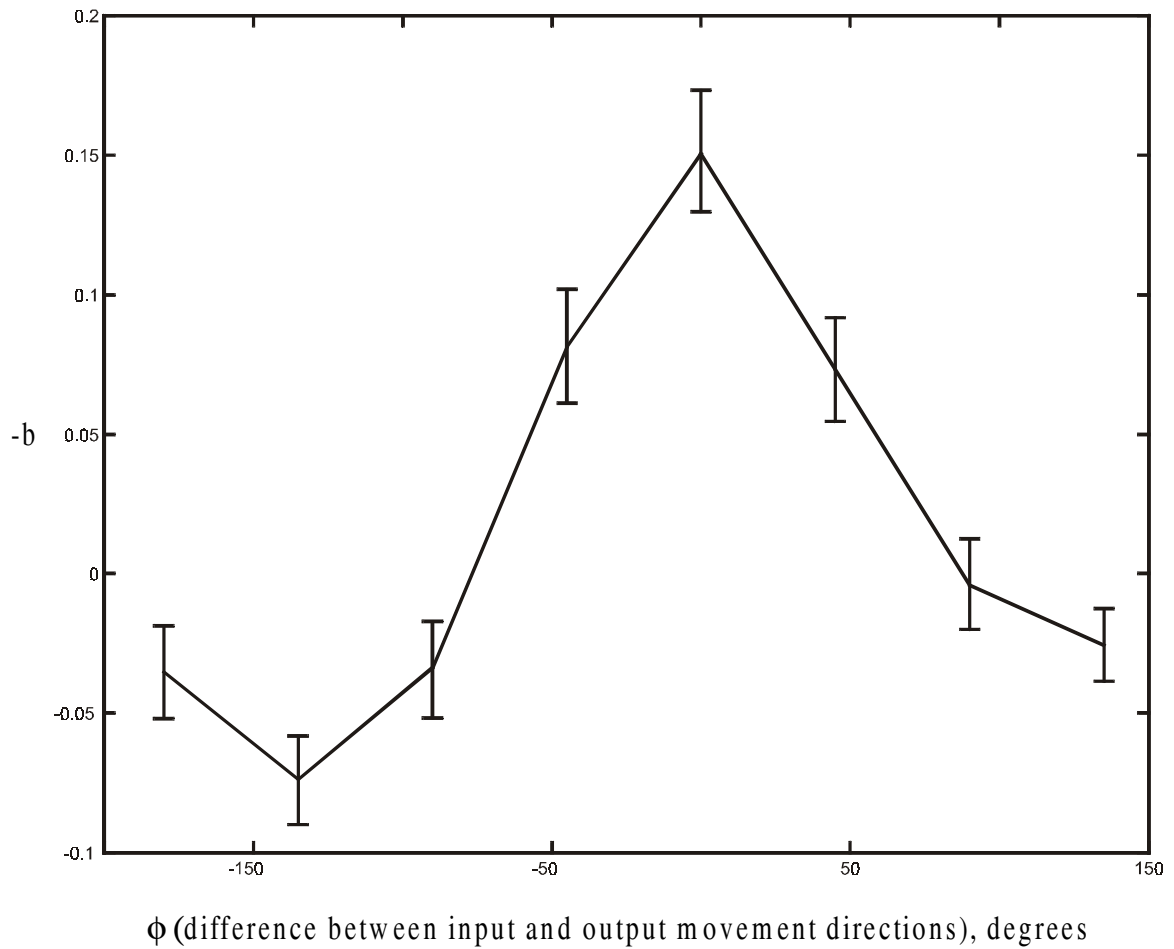
The bootstrapping revealed that in each modeled direction, the  $\vec{b}$  parameters were significantly different from zero in multiple input directions. In each direction of movement, the self-influence of catch trials was significant; this was expected, since this component corresponds to the influence of catch trials modeled in

the scalar input formulation. The value of the self-influential component of  $\vec{b}$ , averaged across directions, is 0.15. The difference between the input values ( $\vec{u}$ ) corresponding to catch trials and force field movements equals 2. Since the feedforward control component is mapped directly into perpendicular displacement (Equation 4.3), this result indicates that on average, a catch trial induces 0.3 cm more curvature in the same direction's next movement than does a force field movement.

Also in each direction of movement, there was at least one direction of movement, 45 degrees away (either clockwise or counterclockwise), which also exerted a significant influence on the feedforward control of movements. In seven out of eight directions of movement, this second direction's influence on the modeled direction was not significantly different than the self-influence. In the eighth direction, toward 90 degrees, there was a sharp drop between self-influence and the influence of 45 degrees away; the influence of movements made toward 135 degrees was small, but significant. This result indicates that other directions of movement can have a positive effect on the modeled direction.

Much more surprisingly, in each modeled direction, at least one other direction of movement has a component of  $\vec{b}$  which is significantly less than zero. In the model for 90 degrees, for instance, the influence of catch trials and force field movements made toward -135, -90, and -45 degrees are all significantly smaller than zero. A negative component of  $\vec{b}$  for an input direction  $\theta$  indicates that, following catch trials in that input direction, perpendicular displacement in the output direction will actually decrease. The sign change of  $\vec{b}$  across input directions indicates that the differential effect of catch trials and force field movements can actually have opposite effects on feedforward control, depending on the direction of input.

To determine the typical response to force field movements and catch trials across the workspace, in each of our bootstrapped data sets I averaged the parameter vectors  $\vec{b}$  across directions of movement, then determined the confidence intervals across the bootstrap. The mean and confidence intervals of the generalized parameter vector  $\vec{b}$  reveal three key features of the typical response of feedforward control to force feedback (Figure 4.12). First, the self-influential effect of force field input is the strongest, producing



**Figure 4.12: The mean across directions of the state-space parameter vector  $\mathbf{b}$ , as a function of the difference in angle between input and output. Error bars represent 95% confidence intervals of the mean as determined by bootstrapping.**

on average a 0.3 cm difference in curvature between force field and catch trial input. Second, the effect of neighboring inputs is half as strong as the self-influential effect, indicating a full-width half-maximum of 90 degrees. Thirdly, and most surprisingly, four out of the eight directions show a significant effect of force field input which opposes the self-influential effect. The magnitude of this opposing influence is 0.15 cm, also half the effect of self-influence.

#### **4.6 Theoretical relationship between map of error into subsequent performance and inverse model basis functions**

Surprisingly, the parameter vector  $\vec{b}$  indicates that force field and catch trial exemplars have different, even opposite, effects on subsequent movements depending on the direction of the exemplar movement. This finding strongly suggests that learning a representation of dynamics in some directions of movement can adversely affect dynamic representations in other directions. To test this hypothesis, I examined extant theories describing the representation and learning of inverse models using basis functions. I then transformed these equations into a formulation that made possible the calculation of the projection of these basis functions across directions of movement.

The neural implementation of the inverse model was modeled presuming the architecture of a layered feedforward neural network. The input of this network depended on desired trajectory; the output equaled the feedforward torque command. A particular subset of this architecture presumes that learning of the inverse model takes place at the output layer. The adapting weights,  $w$ , would therefore relate to the torque output from the inverse model ( $\hat{\tau}$ ) using the equation

$$\hat{\tau} = w^T g(q^*, \dot{q}^*, \ddot{q}^*), \quad (4.4)$$

where  $g$  is a concatenation of basis functions of the inverse model, depending on desired position, velocity, and acceleration. In all equations, hats indicate the output of the forward model (the hat signifying “estimate of”). Since torque is two-dimensional, while  $g$  has  $n \times 1$  elements,  $w^T$  has  $2 \times n$  elements. Since desired position, velocity, and acceleration depends on the direction of movement, the values of these basis functions change with movement direction. Since learning takes place by updating weights, however, this architecture postulates that the basis functions of the inverse model do not change with learning.

The inverse model updates its weights based on the error created during movement. To directly update the inverse model, the weight adaptation should aim to minimize a cost function of error in torque,

$$J = \frac{1}{2}(\hat{\tau} - \tau^*)^T (\hat{\tau} - \tau^*). \quad (4.5)$$

where  $\tau^*$  represents the appropriate feedforward torque which would produce the desired trajectory  $q^*$ .

Using a gradient descent technique, the least-mean-squared learning rule to try to minimize J would be

$$w_{n+1} = w_n - \eta \nabla_w J_n, \quad (4.6)$$

where  $\eta$  is a small constant. Note that rather than calculating the actual derivative  $\nabla_w J$ , the system uses the online approximation  $\nabla_w J_n$ . Presume this system makes a movement, indexed 1, which is imperfect.

The change in weights that would result based on Equation 4.6 would be

$$\Delta w_1 = -\eta \left. \frac{d\hat{\tau}}{dw^T} \right|_1 \tilde{\tau}_1. \quad (4.7)$$

In (4.7) and in all other equations, the tilde indicates error, hence  $\tilde{\tau}$  represents torque error. The vertical

bar and subscript in  $\left. \frac{d\hat{\tau}}{dw^T} \right|_1$  indicates that the derivative is evaluated along the trajectory in movement 1.

This learning rule is common to three different previous studies that examined how the human nervous system could learn an inverse model. The direct inverse modeling technique used by Kawato et al (1987) directly multiplies error in torque by the basis functions to determine weight changes. Sanner and Kosha (1999) adapt weights by multiplying the inverse model basis function by an error state combining position and velocity. The same error state, however, is used to compute feedback torque, which is equal in magnitude and opposite in direction to error in torque. The weight update rule of Sanner and Kosha differs from (4.5) only in that  $\eta$  is a small positive definite matrix, rather than a small scalar. The distal supervised learning algorithm of Jordan and Rumelhart (1992) calculated weight change with the equation

$$\Delta w = -\eta \frac{d\tau}{dw^T} \frac{d\hat{q}}{d\tau^T} \ddot{q}. \quad (4.8)$$

The last two terms of this equation, however, uses the derivative of the forward model to transform distal error into torque error. All three formulations, then, rely on the same least-mean-squared technique to map errors in torque into appropriate weight changes.

The weight change responding to error in movement 1 would then alter the system's performance on a subsequent movement indexed 2. The torque generated by the inverse model ( $\hat{\tau}$ ) during movement 2 can be simplified using a Taylor expansion:

$$\hat{\tau}_2(w + \Delta w_1) = \hat{\tau}_2(w) + \left. \frac{\partial \hat{\tau}}{\partial w^T} \right|_2 \Delta w_1 + O(\eta^2). \quad (4.9)$$

A sufficiently small  $\eta$  permits the negligence of terms second order and higher in the Taylor expansion.

Let the change in torque generated in movement 2 caused by the learning from movement 1 to be termed  $\Delta_1 \tau_2$ , such that

$$\Delta_1 \hat{\tau}_2 \equiv \hat{\tau}_2(w + \Delta w_1) - \hat{\tau}_2(w), \quad (4.10)$$

then, using (4.7) and (4.9),

$$\Delta_1 \hat{\tau}_2 = -\eta \left( \frac{\partial \tau}{\partial w^T} \right)^T \Big|_2 \left( \frac{\partial \tau}{\partial w^T} \right) \Big|_1 \tilde{\tau}_1. \quad (4.11)$$

Presuming that all the adapted weights multiply the inputs to the output layer of the inverse model (Equation 4.3), the change in torque in movement 2 due to learning from movement 1 becomes

$$\Delta_1 \hat{\tau}_2 = -\eta g_2^T g_1 \tilde{\tau}_1. \quad (4.12)$$



where  $g_i \equiv g(q_i^*, \dot{q}_i^*, \ddot{q}_i^*)$ .

Up to this point, I have considered the effect of error in one movement on the subsequent movement. The state-space parameter vector  $\vec{b}$ , however, estimates how the effect of one movement on the next changes with the direction of the first movement. Consider therefore directions of movement  $i$  and  $j$ . If a force field movement in direction  $i$  is indexed “if” and a catch trial movement is indexed “ic,” then

$$\Delta_{if} \hat{\tau}_j = -\eta g_j^T g_i \tilde{\tau}_{if} \quad (4.13)$$

and

$$\Delta_{ic} \hat{\tau}_j = -\eta g_j^T g_i \tilde{\tau}_{ic}. \quad (4.14)$$

Note that since basis functions depend on only desired trajectories, the evaluations of basis functions are equal in catch trials and force field movements. Subtracting these two equations, and defining

$\Delta_{id} \hat{\tau}_j \equiv \Delta_{if} \hat{\tau}_j - \Delta_{ic} \hat{\tau}_j$  and  $\tilde{\tau}_{id} \equiv \tilde{\tau}_{if} - \tilde{\tau}_{ic}$  (where “d” indicates difference), then

$$\Delta_{id} \hat{\tau}_j = -\eta g_j^T g_i \hat{\tau}_{id}. \quad (4.15)$$

Equation (4.15) therefore indicates that errors in torque in direction 1 are mapped into changes in torque in direction 2 by the projection between the basis functions generated for the two movement directions. This formulation now can be compared to the spatial dependence of  $\vec{b}$  to determine unlearning receptive fields in terms of torque. These receptive fields will then reveal features of the basis functions  $g$ .

#### **4.7 Specifying basis function properties using parameters of vector-input state-space model**

The state-space model considers the relative effects of catch trials and force field movements on subsequent behavior. Since catch trials induce large errors, and Equation 4.15 relates errors to subsequent performance via inverse model basis functions, the model parameters could provide information about the basis

functions. The state-space model, however, fit the input in terms of the binary presence of the force field, and fit the output in terms of perpendicular displacement 250 msec into the movement. Equation 4.15 relates errors in torque to subsequent production of torque. The biomechanical simulation of movement in the force field, however, estimates the torque requirements of movements made in the null field and in the force field. The biomechanical controller and the inherent nature of force field and catch trial movements can be utilized to use the state-space model parameters to investigate the nature of inverse model basis functions.

To quantify the torque error in Equation 4.15, first assume that the proper additional torque required to move in the force field is represented by  $\tau^*$ . In force field movements, the error in torque  $\tilde{\tau}$  equals the difference between the appropriate torque and the additional output of the inverse model ( $\tau^* - \hat{\tau}$ ). In catch trial movements, no force field is applied, so any torque generated in addition to the null field requirements is in error ( $\tilde{\tau} = -\hat{\tau}$ ). The difference between torque error in the force field and in catch trials precisely equals the additional torque necessary to move in the force field ( $\tau^*$ ). All  $\tilde{\tau}_{id}$ 's, therefore, simply equaled the corresponding force-field appropriate additional torque  $\tau_i^*$ .

To quantify the change in torque output on the left hand side of Equation 4.15 required the estimation of the relation between torque and positional error. I therefore simulated many trajectories, holding the parallel component of velocity fixed and varying the perpendicular component. I found that the relationship between perpendicular displacement and torque was almost perfectly linear. Remember that each direction of movement  $j$  was fit with a separate model and therefore was fit with a separate vector  $\vec{b}_j$ . The quantity  $2b_j(i)$  indicates the effect that recent dynamic input in direction  $i$  has on subsequent kinematics in direction  $j$ . The linear mapping of this kinematic error into dynamic error therefore requires only the correct constant multiplying  $2b_j(i)$ .

The appropriate constant for this linear mapping can be discovered by looking to catch trial movements. The difference in perpendicular displacement between force field and catch trial movements in direction  $j$

equals  $2d_j$  (from Equation 4.2). The difference in error between force field and catch trial movements in direction j is  $\tau_j^*$ , as discussed above. The proper ratio between torque and perpendicular displacement errors in direction i, therefore, is  $\tau_j^*/2d_j$ . The change in torque on the left hand side of Equation 4.15 can finally be expressed as

$$\Delta_i \hat{\tau}_j = \tau_j^* \cdot b_j(i) / d_j. \quad (4.16)$$

Transforming Equation 4.15 into known quantities generates

$$\tau_j^* \cdot b_j(i) / d_j = -\eta g_j^T g_i \tau_i^*. \quad (4.17)$$

The product of the basis functions can now be estimated using the torque predictions of the biomechanical model and the parameters of the vector input state-space models.

Both the input and output torques were converted to scalar measures by integrating over the first 200 msec of simulated torque production. As a final simplifying assumption, I focused on the change of the basis function projections as a function of the angle between movements i and j, so if  $\phi = \theta_j - \theta_i$ , then

$$G(\phi) \equiv g(\theta_j)^T g(\theta_i). \quad (4.18)$$

The magnitude of the desired positions, velocities, and accelerations are consistent in all directions of our task; these desired trajectories differ only in their angular information. This task specification makes reasonable the focus on the angular dependence of basis function projections.

To estimate the mutual basis function projections  $G$  for a particular difference in input and output directions  $\phi$ , I first calculated the output (left hand side) of Equation 4.17 for all eight instances (indexed by i) of the input and output angles differing by  $\phi$ :

$$O_{\phi}(i) = b_{\theta_i+\phi}(\theta_i) / d_{\theta_i+\phi} \tau_{\theta_i+\phi}^* \quad (4.19)$$

I also calculated the corresponding torque error inputs:

$$I_{\phi}(i) = \tau_{\theta_i+\phi}^* \quad (4.20)$$

Since the individual basis functions are scalar valued,  $G(\phi)$  is a scalar. Torques, however, are two-element vectors, so Equation 4.17 is of the form  $v_2 = s v_1$ . I estimated  $s$  by determining the ratio between

the projection of  $v_2$  into  $v_1$  with the length of  $v_1$  ( $\hat{s} \equiv \frac{v_2 \cdot v_1}{v_1 \cdot v_1}$ ). For each input-output pair, I estimated

$-\eta G(\phi)$  using the equation:


$$E_{\phi}(i) = \frac{O_{\phi}(i) \cdot I_{\phi}(i)}{I_{\phi}(i) \cdot I_{\phi}(i)} \quad (4.21)$$

An average over each estimate provided the best overall estimate of  $-\eta G(\phi)$ :

$$-\eta G(\phi) \cong \sum_i E_{\phi}(i) / 8 \quad (4.22)$$

From my bootstrapping of the initial perpendicular displacement data, I produced 1000 data sets resampled, with replacement, from the original pool of 40 subjects. From these resamplings I found 1000 parameters of the state-space model fits for each modeled direction  $i$ , including  $d_i$  and  $\vec{b}_i$ . I therefore produced 1000 estimates for  $-\eta G(\phi)$  for each angle difference  $\phi$  from which I estimated confidence intervals.

The mean values of  $-\eta G$  reveal that the receptive field of torque learning strongly peaks around  $\phi = 0$ , but has significant positive components in directions far away from  $\phi = 0$  (Figure 4.13). Echoing the central peak of the receptive field of angular kinematic error, the torque learning is twice as strong if torque input



**Figure 4.13: The mapping of torque error into torque compensation, as a function of the difference between input and output movement directions. Dotted lines indicate reduced confidence in values at  $\phi = -90$  and  $90$  degrees (see text). Error bars represent 95% confidence intervals of the mean as determined by bootstrapping.**

and output are in the same direction than if the input direction is  $-45$  or  $45$  degrees away. At  $90$  degrees, the center of the bootstrap reveals a small positive effect that is marginally significant; at  $-90$  degrees, the mean eigenvalues are essentially zero. The torque learning at  $135$ ,  $-180$ , and  $-135$  degrees, however, all have significant positive values, the mean of which are 25 – 35% of the peak in torque learning transfer. Note that Equation 4.15 indicates that  $-\eta G$  maps torque error into subsequent torque output, whereas these eigenvalues are of  $\eta G$ .

Recall that the torques used in Equation 4.17 are the appropriate torques (additional to the null field) necessary to move in the force field. The two-dimensional torque appropriate to move in one direction is roughly orthogonal to the torque appropriate to move in a direction  $90$  degrees away. This orthogonality renders troublesome the estimates of the basis function projections for  $\phi$  equal to  $-90$  and  $90$  degrees. A fairer estimate of the relation between  $\eta G$  and  $\phi$ , then, might ignore these points, as illustrated by the dotted lines in Figure 4.13.

To better understand the consequences of the values of  $\eta G$  as a function of  $\phi$ , consider a specific movement that generates a specific torque error  $\tilde{\tau}$  early in the movement. The eigenvalues at  $\phi = 0^\circ$  indicates that in the next movement in the same direction torque would be augmented in the opposite direction of  $\tilde{\tau}$ , with 16% of the magnitude of  $\tilde{\tau}$ . In the next movement in the opposite direction ( $\phi = 180^\circ$ ), torque would also be augmented in the opposite direction of  $\tilde{\tau}$ , but with 4% of  $\tilde{\tau}$ 's magnitude. For all seven values of  $\phi$  which have significant values of  $\eta G$ , subsequent torque compensation is always in the opposite direction of the torque error input.


#### ***4.8 Modeling and simulating unlearning receptive fields using velocity-dependent Gaussian basis functions***

I next explored whether learning inverse models using Gaussian basis functions could replicate the catch-trial unlearning receptive field measured in subjects. I considered the class of inverse models that estimated dynamics by multiplying adjustable weights by velocity-based Gaussians

$$g(k) = \exp(-\|\dot{x} - c_k\|^2 / (2\sigma^2)) \quad (4.23)$$

The centers of the basis functions were fixed in a grid spanning hand velocity space from  $-40$  to  $40$  cm/s. I varied  $\sigma$  to investigate how the width of the basis function altered the learning and unlearning of force fields. As I changed basis function width, I also changed the density of the basis functions across the velocity range: the centers of the basis functions  $c_k$  were always spaced  $1 \sigma$  apart.


Using Equation 4.17 and evaluating each basis function at a specific instant in time, the basis function width predicts the receptive field parameterized by  $\vec{b}$ . The evaluation of wide basis functions substantially overlapped each other in movements toward different directions (Figure 4.14A). This overlap induced error in one direction of movement to propagate into compensation in other directions; when this overlap spills into opposing directions, the reversal of force error parlayed into an increase in curvature in subsequent movements. A calculation of the receptive field of force feedback using basis functions evaluated 100 msec into movements reveals substantial negative components of the receptive field for



**Figure 4.14: Estimating receptive fields using snapshot evaluations of basis functions. A: Evaluation of Gaussian basis functions with  $\sigma = 0.1$  m/s, 100 msec into the movement. Position of subplot reflects direction of movement. Remaining figures: estimation of receptive field ( $b(\phi)$ ) with varying basis function widths when bases are evaluated at 100 msec (B) or 200 msec (C) into the movement.**


$|\phi| > 90$  degrees only if basis functions are sufficiently wide (Figure 4.14B). Receptive fields show very little negative-going effects, however, when basis functions are calculated 200 msec into the movement (Figure 4.14C). Since the estimates of receptive fields depended heavily on when into the movement basis functions were evaluated, I built a simulation that generated full movements using the basis function

A



B

Biomechanical simulation: receptive fields of catch-trial induced unlearning



**Figure 4.15: Simulation of catch trial learning using inverse model with Gaussian bases. A: Actual perpendicular displacement (\*) with simulated perpendicular displacements of models using bases with  $\sigma = 0.04$  (thin) or  $0.12$  (thick) m/s. B: Receptive fields of simulations with different widths of Gaussians.**

representation of external forces. The simulated controller knew perfectly well the dynamics of its own arm; the passive dynamics of the robot were omitted from this simulation. The simulated controller made movements in the same 192 movement target set, with catch trials, that subjects experienced. The forces



generated by the controller were determined by an inverse model, with  $\hat{F} = w^T g$ , and by stiff and viscous elements. After each movement, weights were updated by  $\Delta w = -\eta g \tilde{F}$ , with  $\eta = 0.025$ . The sequences of perpendicular displacements created by 3 sets of simulation were modeled by the state-space formulation to determine sensitivities to dynamic input (the catch-trial receptive fields) and to determine how the sensitivities changed with basis function width.

Although simulated controllers using narrow and wide basis functions unlearned with catch trials, only wide basis functions featured receptive fields with appropriately negative components. For instance, in simulated movements toward 270 degrees, an inverse model in which  $\sigma = 0.04$  m/s adapted to reduce initial error, unlearned with catch trials, and smoothly recovered after catch trials (Figure 4.15A, thin line). An inverse model with wide basis functions ( $\sigma = 0.12$  m/s) also reduced initial error and unlearned with catch trials, but dynamic input from other directions of movement induced variance other than smooth recovery after catch trials. During fielded movements in the 2<sup>nd</sup> and 3<sup>rd</sup> sets of movement, when variance in performance is dominated by catch trial-induced unlearning, the correlation of simulation to data jumped in this direction of movement from  $r^2 = 0.48$  to  $r^2 = 0.80$ . Across directions of movement, the correlation between simulated and actual perpendicular displacements was highest for  $\sigma = 0.10$  and  $\sigma = 0.12$  m/s.

Wide basis functions also replicated the dependence of the catch-trial receptive field to input direction (Figure 4.15B). When the simulated inverse model was driven by narrow basis functions, the state-space parameter  $\vec{b}$  was strongly peaked at  $\phi = 0$ , dropped off rapidly, and remained at zero for  $\phi$  far from zero. With increased width of basis functions, the peak of sensitivity broadened, and negative values appeared for  $\phi$  far away from zero. Surprisingly, just as in the state-space model fit to actual perpendicular displacements (Figure 4.12), the receptive field of the simulated controller featured a trough which was deepest at  $\phi = -135$  degrees, not  $\phi = -180$  degrees as the algebraic calculation predicted (Figure 4.14B). The simulation's receptive field most closely approximated the actual findings when  $\sigma = 0.12$  or  $0.16$  m/s.


Width of Gaussian in simulation (m/s)	Simulation-to-data correlation of mean perpendicular displacement in force field movements, 2 <sup>nd</sup> and 3 <sup>rd</sup> sets ( $r^2$ )	Simulation-to-data correlation of mean state-space estimate of receptive field parameter “b” ( $r^2$ )
0 (producing $b_{(\phi=0)}=1$ , $b_{(\phi \neq 0)}=0$ )	----	0.707
0.04	0.582	0.769
0.05	0.630	0.787
0.08	0.723	0.876
0.10	0.749	0.920
0.12	0.743	0.939
0.16	0.701	0.935
0.20	0.654	0.913

**Table 4.1: Correlation between simulation and actual data, of force field movement perpendicular displacement and receptive fields, as a function of the simulation’s basis function width.**

In these simulations, the width of the Gaussians ( $\sigma$ ) also described the spacing of the basis functions. When I doubled the number of basis functions such that the spacing between basis function centers was  $\sigma/2$ , the simulation-to-data correlation of both perpendicular displacement and receptive field went down (average  $r^2$  reduced by 0.058 for the receptive field and 0.118 for perpendicular displacement). The increase in density seems to lessen the interference caused by broad basis functions. This finding indicates that basis functions need to be both broad and sparse in order to mimic subject performance and estimated catch-trial receptive fields.

#### ***4.9 Learning constant force fields using velocity-dependent Gaussian basis functions***

Finally I removed the catch trials from the training of the simulation in order to investigate how varying basis function widths altered adaptation to the constant force field. The perpendicular displacements of the first movements made in the force field (toward 90 degrees) are equally bad (Figure 4.16). Inverse models dependent on wide basis functions, however, recovered from the initial error more quickly than narrow basis functions. The overlap between evaluations of basis functions in neighboring directions seems to aid initial adaptation to the force field. Curvature of movements generated by narrow bases decays directly to zero. Very surprisingly, by the eighth movement toward the target at 90 degrees and by the 25<sup>th</sup> movement overall, movements generated by wide basis functions begin to generate curvature in the opposite direction of the initial error. This behavior mimics the precurvature seen in subjects learning the constant force field in Chapter 3. This result introduces the intriguing possibility that precurvature results not from torque



**Figure 4.16. Response of simulated learning using inverse model with Gaussian basis functions. A through C: Trajectories of simulated controllers with varying Gaussian widths in first (A), 11<sup>th</sup> (B), and 21<sup>st</sup> (C) movements toward 90 degrees. D: Perpendicular displacement of simulated movements made toward 90 degrees, in first two sets of simulated movements. E: Perpendicular displacement of simulated movements in all directions, smoothed with a 13-movement wide window.**

change considerations, but as a behavior emergent from wide basis functions. Whereas the precurvature in the simulation is transient and disappears in the second and third sets of training, actual precurvature increased throughout the three sets of training. Despite this inconsistency between simulated and actual behavior, the precurving of movements by the simulated controller suggests that wide basis functions could describe the computational basis of adaptation to both stationary and nonstationary environments.

#### **4.10 Discussion**

Kinematic and electromyographic features of subjects' movements indicate that occasional catch trials do indeed disrupt the learning of the usually applied force field. Catch trials induced an increase in curvature in subsequent force field movements in all directions of movement. The increase in curvature decayed within two to four movements, but even movements immediately preceding catch trials were significantly more curved than the movements generated by subjects who trained in a constant force field. The peak of force-field appropriate EMG arrived at the same time into the movement both before and after catch trials. Since the timing of the peak activation quantifies the relative contribution of feedforward and feedback processes, the constancy of the peak timing suggests that throughout training subjects rely more heavily on feedback control than the constant force field subjects. Previous studies that included sudden changes in dynamics interspersed within dynamic force training (Shadmehr and Brashers-Krug, 1997; Goodbody and Wolpert, 1998; Sainberg et al, 1999) represented the performance during the sudden changes in dynamics as representing the activity of an inverse model which was learning the usually applied force field. The sudden changes themselves, however, generate a different response to training than stationary dynamic training.

Although training in a nonstationary dynamic environment induces different learning and behavior across the training session, the response of the subjects to sudden changes in dynamics reveals important new details of the human motor learning system. As evidenced in movements toward 90 degrees, catch trial subjects do learn how to move substantially straighter and more accurately in the force field. With each catch trial, however, the force field-specific knowledge of dynamics is partially unlearned, and subsequent movements have significantly more curvature. As demonstrated in Chapter 2, the orientations of resultant vectors of movement-initiating EMG provide a window into the adaptation of the inverse model. Before

catch trials, the resultant vectors pointed in a direction appropriate for the force field, clockwise from null field orientations. After catch trials, however, each muscle's resultant vector rotated counterclockwise, back toward the null field orientation. The catch-trial induced rotations indicates that the sudden removal of the force field caused specific unlearning of the inverse model.

The state-space models reveal how the human adaptive motor control system could both learn in the initial exposure to the force field and unlearn after catch trial exposure. States in these models equaled the component of perpendicular displacement created by the human feedforward controller. The scalar input state-space representation depended on only four parameters: the initial state value, an exponential decay of state, a scalar input response which changed state based on the last dynamic environment, and a scalar quantifying the component of curvature attributable to the robot. This simple model fit well the initial recovery to large curvatures early in training, large negative curvatures in catch trial movements, jumps in curvature after catch trials, and recovery after jumps. The interaction between the a and b parameters even fit the high level of curvature in movements before catch trials, compared to constant force field subjects. The scalar input state-space architecture confirms that the human CNS could implement a simple adaptive control algorithm to learn and unlearn the viscous forces. The scalar input model, however, left some unexplained variance in the time series of movement curvatures. Much of this additional variance was explained by a vector-input formulation, which also considered how catch trials affect subsequent performance across directions of movement.

This vector model revealed an unexpected receptive field of learning. The parameter vector  $\vec{b}$  represented how the curvature of subsequent movements was differently affected depending on the movement direction of dynamic input. Averaged across directions, the parameter vector  $\vec{b}$  revealed that subsequent curvature was most strongly affected by recent dynamic history when the input (force field on or off) and the output (subsequent movement) were in the same direction. As observed in kinematics and the scalar input model, this element of  $\vec{b}$  indicates an increase in curvature following catch trials. Catch trials in neighboring directions affect curvature half as much, indicating a sharp peak of unlearning. When catch trials occurred  $-180$ ,  $-135$ , or  $135$  degrees away from the output movement direction, however, the elements of  $\vec{b}$  are

significantly negative. These significant negative values indicate that catch trials in these directions actually induce small reductions in curvature in the output direction. The receptive field of learning, as represented by curvature, reflects a “Mexican-hat” function, in which the response of the human adaptive controller can have opposite signs depending on the input movement direction.

Since the viscous force field always pushes subjects 90 degrees clockwise from their instantaneous velocity vector, the natural metric of motor error in our task is movement curvature, as measured by perpendicular displacement. Curvature, however, is defined in reference to movement direction. The relationships between curvature and hand kinematics, joint kinematics, force applied at the hand, and joint torque all change with movement direction. The natural coordinates of dynamic properties lie in fixed hand and joint coordinate systems that do not rotate with direction. To elicit information about human estimation of inverse dynamics, the angular error needed to be transformed to a natural dynamic metric. I therefore used the biomechanical model to transform the perpendicular state-space parameters into errors and subsequent compensations in torque.

The mapping of torque errors into torque adjustments reveals interesting properties of the human inverse model, especially if inverse dynamics are computed using basis functions. A technique to model unknown functions is to fit the input/output relationship with basis functions multiplied by scalars. Each basis function depends on input variables and spans some subset of input space. Function approximation using focused radial basis functions is mathematically equivalent to approximating input/output mappings using layers neural networks (Poggio and Girosi, 1990; Girosi et al, 1995). A particularly simple representation of inverse dynamics calculates nonlinear functions of movement states, then calculates inverse dynamics by specifying appropriate weights to multiply basis functions spanning those nonlinear functions (Equation 4.4). This simplification restricts the update of weights to the output layer. Such a simplification, however, is reasonable, as the inverse dynamics of robotic and human arms can be represented as nonlinear transformations of position, velocity, and acceleration linearly multiplying a vector representing inertial quantities (Slotine and Li, 1991; Sanner and Slotine, 1995). This simplified representation has been simulated to replicate human motor learning of viscous fields (Slotine and Koshla, 1999). The neural computation of inverse dynamics, therefore, can be reasonably considered to be represented by basis

functions (Mussa-Ivaldi, 1999). The relationship between torque error and subsequent torque compensation, then, reveals how these basis functions project onto each other as a function of movement direction (Equation 4.15).

The estimates of the mutual projection between basis functions  $G(\phi)$  showed that the inverse model basis functions have a sharp peak in correlation around zero angle difference in movement direction, but significant positive correlation with large angle differences in movement direction. The estimates were largest when  $\phi = 0^\circ$ , or when the input movement direction (where torque error occurred) equaled the output movement direction (where responsive torque compensation occurred). The estimates were 50% smaller when the input and output movement directions were 45 degrees apart, and approached zero when  $\phi = 90^\circ$ . These data points suggest a narrow overlap between the spans of basis functions. Many models of inverse model formation presume that inverse models formed when training in viscous fields realized the velocity dependence of force (Kawato et al, 1987; Shadmehr and Mussa-Ivaldi, 1994; Sanner and Kosha, 1999). Computational (Schaal and Atkeson, 1998) and psychophysical (Gandolfo et al, 1996) studies have suggested that each basis function spans a local subset of velocity space. The shape of the estimates curve between  $-90$  and  $90$  degrees would support the hypothesis that inverse model basis functions are local and depend on velocity.

Mutual projections at  $-180$ ,  $-135$ , and  $135$  degrees, however, have significant positive values. Basis functions that are local should not have significant positive correlations when they are realized at two far away, even opposing, directions of movement. This finding suggested that basis functions were wide enough to generate interference across directions, which was confirmed by simulating a controller with an inverse model with velocity-dependent basis functions. Simulations that relied on wide, sparse basis functions to estimate the viscous field mimicked both the variance across movements in subjects' performance and the angular dependence of the sensitivity to catch trials. Previous models of human motor learning presumed narrow basis functions, since wider basis functions could lead to interference across velocity space (Schaal and Atkeson, 1998). The discovery of interference across velocity space indicates that computational elements in the adaptive motor system are not fully optimized to minimize error. A

computationally efficient method of constructing a basis function representation strives to both approximate arbitrary functions and to constrain solutions to smooth forms (Poggio and Girosi, 1998). The high correlation between wide Gaussian basis functions and subject performance may indicate that the adaptive motor system is optimized to learn only dynamics with low frequency components in velocity space.



## **Chapter 5 Discussion: Computational and neural bases of human motor learning**

In this thesis, psychophysical and electromyographic metrics have been recorded from humans learning stationary and nonstationary viscous force fields. These metrics reveal that movement-to-movement and across multiple sets of training, subjects learned force-field specific dynamics which informed their feedforward control. Our subjects used this information to produce straighter, more accurate movements over the course of training. Catch-trial induced unlearning and subsequent recovery indicated that the force-field specific knowledge rose and fell with each new exemplar. The difference in performance in stationary and nonstationary environments, however, also indicates that humans optimized performance based on longer term features of dynamics. Learners of the stationary force field did not converge upon unperturbed trajectories, but slightly precurved their movements, which lessened the incurred torque change cost. Learners of the nonstationary force field peaked their force field appropriate muscle activation 80 msec later into the movement than their stationary environment counterparts. Since the timing of peak activity remained constant even across catch trials, the lateness of activation suggests a heightened dependence on feedback control in order to produce good movements in the force field and not-so-good movements during catch trials. The asymptotic performance of both sets of subjects suggests both specific learning of forces, generating predictions of the dynamics of each movement, and broader adaptation of strategies requiring accommodation over many movements.

The inverse dynamic model, providing feedforward control, is likely trained by the output of the feedback controller. Three theories compete to explain how inverse model weights could be updated, based on distal error in acceleration (Jordan and Rumelhart, 1992), sliding variables of position and velocity error (Sanner and Kosha, 1999), or directly from the feedback error command (Kawato et al, 1987). The previous two models presume stiff and viscous error feedback, however, and presume feedback torque is proportional in magnitude and opposite in direction to distal error. All three models therefore specify feedback control that is available to the CNS as a training signal. The electromyographic composite traces most appropriate for force field movements are consistently shaped, but shifts to earlier into the movement as subjects train. The neural output of subjects who learn stationary dynamics suggests a simple learning algorithm: in an

attempt to minimize future error, the previously generated feedback torque simply augments subsequent feedforward torque.

The transition from feedback to feedforward torque, on a shorter time scale, founds the quantification of receptive fields and basis functions of short-term dynamic learning. During catch trials, the sudden removal of learned viscous forces produces kinematic and torque error in the direction opposing initial error caused by the forces. The torque error is corrected by feedback torque that actually points in the same direction as force-field appropriate feedforward torque. Weights of the inverse model adjust in response to this odd exemplar, producing increased curvature in the next movement in the force field in that same direction. The same weights adjusted by the catch trial feedback, however, could inform movements in other directions as well. The dependence of the feedback-induced torque change describes receptive fields of motor learning, where here receptive fields connote the full transformation between learning and behavior. Previously, motor learning receptive fields have been shown to have positive transfer across a narrow range of movement directions (Gandolfo et al, 1996) and across different workspaces, with information transferring in joint coordinates (Shadmehr and Mussa-Ivaldi, 1994). These studies relied on a steady-state approximation to learning, however, and could not investigate the fine structure of learning in response to single movements. In the space of movement curvature, catch trial induced unlearning reveals a narrow band of positive transfer, but also a reciprocal band of negative transfer. The adaptation of movement accuracy due to catch trials actually reverses depending on the location of dynamic input. In torque space, when input and output movement directions are far away, feedback torque direction is conserved in subsequent augmentation of feedforward torque. This conservation is inappropriate, however, since the direction of field generated force reverses in opposing directions.

The negative transfer of learning in certain directions of movement indicates unexpected properties of the basis functions underlying the calculation of inverse dynamics. Presuming a neural network architecture of inverse modeling, adapting weights multiply basis functions which depend on desired trajectory parameters and therefore change with direction of movement. Radial basis function representations approximate well arbitrary functions (Poggio and Girosi, 1990), but are especially warranted for inverse dynamics since they can be represented by a vector on inertial parameters multiplying known nonlinearities of movement states

(Slotine and Li, 1991; Sanner and Kosha, 1999). The least-mean-squared learning rule of inverse dynamics specifies a weight change opposite in sign and proportional to the product of the basis function and the error in torque. The augmentation in torque in a subsequent movement equals that realization of the basis function times the weight change. The mapping of inputted torque error to outputted torque compensation, as quantified in the receptive field, therefore reveals the projection between the basis functions calculated in the input and output movement directions. The receptive field in torque space suggests that the basis functions have a sharp positive peak, but also significant positive components in directions where velocity and acceleration have reversed signs. Previous hypotheses have suggested that inverse dynamic basis functions could be sharply focused in velocity space (Gandolfo et al, 1996; Schaal and Atkeson, 1998) The basis functions suggested by receptive fields require a wider basis function spanning larger regions of velocity space, possibly producing interference.

The results of this thesis could indicate the involvement of more than simply motor cortical cells in motor learning. A first approximation of my results could be explained by the firing patterns in the motor cortex. Rotation of movement-initiating muscle activation correlates well with the predictions of our anthropomorphic biomechanical controller. Since changes in EMG correlate well with changes in motor cortical activity (Sergio and Kalaska, 1997), the rotations in muscle activation could indicate a rotation in the directional dependence of neurons in the primary motor cortex (Georgopoulos et al, 1982). Preliminary studies have indeed measured such a rotation of motor cortical preferred direction in monkeys learning a viscous force field (Benda et al, 1997; Li et al, 1998).

The receptive field suggested by catch-trial induced unlearning, however, indicates a neuronal substrate which is broadly tuned to velocities. To account for the negative-going components of the receptive field which approach one half of the peak strength, basis functions need to have sizeable components across all directions of movement. The tuning of cerebellar cells to velocity has been measured as much broader than motor cortical cells (Fortier et al, 1993). Just as the receptive fields of neurons in the visual system subtend larger angles in higher centers (Ahissar and Hochstein, 1997), neurons in higher motor cortical areas and the basal ganglia may also span larger velocity spaces than do motor cortical cells. When humans learn a visual discrimination task, the grain of the relevant information determined the level at which plasticity

occurred (Karni, 1996; Ahissar and Hochstein, 1997). A motor task featuring high frequency information could reveal one of two possibilities. If humans cannot learn fields with high frequencies, then perhaps plasticity is only possible in higher centers and the link between motor cortex and muscles is relatively hard-wired. On the other hand, a higher frequency task could uniquely spark plasticity in the more finely tuned motor cortex, which could be verified by testing for generalization outside the trained workspace (Shadmehr and Mussa-Ivaldi, 1994).

My work has focused on an inverse model of dynamics as the sole driver of feedforward control. A second computational model, termed a forward model, calculates the forward process of the plant. A forward model transforms an efferent copy of descending commands into a prediction of state, or position and velocity of the arm (Miall and Wolpert, 1996). Previous work has established the importance of forward models in canceling self-induced afference, preparing other systems for the consequences of movement, and combating neural delays in order to stabilize control (Miall et al, 1993; Wolpert et al., 1995a; Flanagan and Wing, 1997). An accurate forward model could even help train an inverse model (Jordan and Rumelhart, 1992). A novel control architecture has hypothesized that forward models could go beyond complementing the inverse model to actually supplementing information early in the movement, generating commands that are incorporated into the movement-initiating feedforward control signal. Kinematic movement features near the end of movement support such an involvement from the forward model (Bhushan and Shadmehr, 1999). The work presented here can only determine the receptive fields and basis functions associated with the computation of inverse dynamics. These experiments do not determine whether the forces produced result from a true inverse model or a forward model in a feedback loop that effectively inverts the plant. Further theoretical and psychophysical work is necessary to determine whether the adaptation and computation presented here lies in a forward model, and inverse model, or some temporally changing combination of both.

I have presented here evidence that humans adapt an inverse model of dynamics while training in a novel viscous force field. The aim of this adaptation is not to replicate unperturbed behavior, as well-learned movements of learners of the stationary force field significantly curve their movements to lessen dynamic cost. While precurvature rules out a strict interpretation of the minimum jerk hypothesis and necessitates

the consideration of dynamics in movement planning, asymptotic behavior differs markedly from minimum torque movements, indicating that kinematic smoothness likely outweighs dynamic smoothness. The modeling of inverse dynamics makes possible the execution of much straighter and more accurate movements than during initial force field exposure. Groups of subjects who train in stationary and nonstationary dynamics both learn to rotate muscle activation functions, which likely indicates a force-field specific adaptation of sharply-tuned velocity-dependent basis functions. The short-term unlearning induced by catch trials constrains the inverse dynamic basis functions to broadly tune and sparsely represent velocity space. Further studies of behaving animals and humans are necessary to determine the importance of these neural substrates on the full repertoire of biological motor control and learning.

## Bibliography

- Ahissar, M. and Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, 387:401-406.
- Arbib, M. A. (1995). *The Handbook of Brain Theory and Neural Networks*. MIT Press, Cambridge, MA.
- Atkeson, C. G. (1989). Learning arm kinematics and dynamics. *Ann Rev Neurosci*, 12:157-183.
- Atkeson, C. G. and Hollerbach, J. M. (1985). Kinematic features of unrestrained vertical arm movements. *J Neurosci*, 5:2318-2330.
- Barto, A. G., Fagg, A. H., Sitkoff, N., and Houk, J. C. (1998). A cerebellar model of timing and prediction in the control of reaching. *Neural Computation*, 11:565-594.
- Basmajian, J. V. and De Luca, C. J. (1985). *Muscles alive: their functions revealed by electromyography*. Williams & Wilkins, Baltimore.
- Benda, B. J., Gandolfo, F., Li, C. S. R., Tresch, M. C., DiLorenzo, D., and Bizzi, E. (1997). Neuronal activities in M1 of a macaque monkey during reaching movements in a viscous force field. *Soc Neurosci Abs*, 23:1556.
- Bhushan, N. and Shadmehr, R. (1999). Computational architecture of human adaptive control during learning of reaching movements in force fields. *Biol Cybern*, 81:39-60.
- Bizzi, E., Mussa-Ivaldi, F., and Giszter, S. (1991) Computations underlying the executive of movement: a biological perspective. *Science*, 253:287-291.
- Brashers-Krug T, Shadmehr R, and Todorov E. (1995) Catastrophic interference in human motor learning. In: *Advances in neural information processing systems*, Vol 7 (Tesauro G, Touretzky DS, Leen TK, eds), pp 19-26. Cambridge, MA: MIT.

Brashers-Krug, T., Shadmehr, R., and Bizzi, E. (1996). Consolidation in human motor memory. *Nature*, 382:252-255.

Conditt, M. A. and Mussa-Ivaldi, F. A. (1999). Central representation of time during motor learning. *Proc Natl Acad Sci*, in press.

Corcos, D. M., Jaric, S., Agarwal, G. C., and Gottlieb, G. (1993). Principles for learning single-joint movements. i. enhanced performance by practice. *Experimental Brain Research*, 94(3):499-513.

De Luca, C. J. (1997). The use of surface electromyography in biomechanics. *J Appl Biomech*, 13:135-163.

Desmurget, M., Pelisson, D, Rossetti, Y, and Prablanc, C. (1998). From eye to hand: planning goal-directed movements. *Neurosci and Biobehav Rev*, 6:761-788.

Fisher, N. I. (1993). *Statistical analysis of circular data*. Cambridge University Press, Cambridge.

Flanagan, J.R. and Wing, A.M. (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

Flanders, M. (1991). Temporal patterns of muscle activation for arm movements in three-dimensional space. *J Neurosci*, 11:2680-93.

Flanders, M. and Soechting, J. F. (1990). Arm muscle activation for static forces in three-dimensional space. *J Neurophys*, 64:1818-37.

Flash, T. and Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci*, 5:1688-1703.

Flash, T (1987). The control of hand equilibrium trajectories in multi-joint arm movements. *Biol Cybern*, 57:257-274.

- Fortier P. A., Smith A. M., Kalaska J. F. (1996). Comparison of cerebellar and motor cortex activity during reaching: directional tuning and response variability. *J Neurophysiol*, 69:1136-1149.
- Gandolfo, F., Mussa-Ivaldi, F., and Bizzi, E. (1996). Motor learning by field approximation. *Proc Natl Acad Sci*, 93:3843-3846.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci*, 2(11):1527-37.
- Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233(4771):1416-1419.
- Ghahramani, Z., Wolpert, D. M., and Jordan, M. I. (1996). Generalization to local remappings of the visuomotor coordinate transformation. *J Neurosci*, 16:7085-7096.
- Girosi, F., Jones, M., and Poggio, T. (1995) Regularization theory and neural networks architectures. *Neural Computation*, 7:219-269.
- Giszter, S., Mussa-Ivaldi, F., and Bizzi, E. (1994) Linear combinations of primitives in vertebrate motor control. *Proc Natl Acad Sci*, 91:7534-7538.
- Goodbody, S. J. and Wolpert, D. M. (1998). Temporal and amplitude generalization in motor learning. *J Neurophysiol*, 79:1825-1838.
- Gottlieb, G. L. (1994). The generation of the efferent command and the importance of joint compliance in fast elbow movements. *Exp Brain Res*, 97:545-550.
- Gottlieb, G. L. (1996). On the voluntary movement of compliant (inertial-viscoelastic) loads by parcellated control mechanisms. *J Neurophysiol*, 76:3207-3229.



- Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394:780-784.
- Hildreth, Ellen C., and Hollerbach, J.M. (1987). "Artificial Intelligence: computational approach to vision and motor control," in: *Handbook of Physiology, Section 1: The Nervous System, Volume V: Higher Functions of the Brain, Part II*, F. Plum, ed., American Physiological Society, Bethesda, Maryland, pp. 605-642.
- Hogan, N. (1984). Adaptive control of mechanical impedance by coactivation of antagonist muscles. *IEEE Trans Automatic Control*, AC-29:681-90.
- Jordan, M. I. (1995). Computational motor control. In Gazzaniga, M. S., editor, *The Cognitive Neurosciences*, pages 597-609. MIT Press, Boston.
- Jordan, M. I. and Rumelhart, D. E. (1992). Forward models: supervised learning with a distal teacher. *Cognitive Sci*, 16:307-354.
- Karst, G. M. and Hasan, Z. (1991). Initiation rules for planar, two-joint arm movements: Agonist selection for movements throughout the work space. *J Neurophysiol*, 66:1579-1593.
- Karni, A. (1996). The acquisition of perceptual and motor skills: a memory system in the adult human cortex. *Brain Res Cogn Brain Res*, 5:39-48.
- Kawato, M., Furukawa, K., and Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. *Biol Cybern*, 57:169-185.
- Kitazawa S, Kimura T, and Uka T. (1997). Prism adaptation of reaching movements: specificity for the velocity of reaching. *J Neurosci*, 17:1481-1492.
- Kudo, K. and Ohtsuki, T. (1998). Functional modification of agonist-antagonist electromyographic activity for rapid movement inhibition. *Experimental Brain Research*, 122:23-30.

Lackner, J. R., and DiZio, P. (1994). Rapid adaptation to coriolis force perturbations of arm trajectory. *J Neurophysiol*, 72:299-313.

Li, C. S. R., Benda, B. J., Padoa Schioppa, C., Gandolfo, F., DiLorenzo, D., and Bizzi, E. (1998). Neuronal plasticity in the motor cortex of monkeys learning to adapt to a viscous force field. *Soc Neurosci Abs*, 24:405.

Marsden, C. D., Merton, P. A., Morton, H. B., Adam, J. E. R., and Hallett, M. (1978). Automatic and voluntary responses to muscle stretch in man. *Prog Clin Neurophysiol*, 4:167-77.

Miall, R. C. (1995). Motor control, biological and theoretical. In Arbib, M. A., editor, *The Handbook of Brain Theory and Neural Networks*, pages 597-600. MIT Press, Cambridge, MA.

Miall, R. C., Weir, D. J., Wolpert, D. M., and Stein, J. F. (1993) Is the cerebellum a Smith predictor? *J Mot Behav*, 25:203-216

Miall, R. C. and Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9:1265-79.

Milner, T. E. and Cloutier, C. (1993). Compensation for mechanically unstable loading in voluntary wrist movement. *Exp Brain Res*, 94:522-532.

Mussa-Ivaldi, F. A (1999). Modular features of motor control and learning. *Curr Opin Neurosci*, in press.

Mussa-Ivaldi, F. A., Hogan, N., and Bizzi, E. (1985). Neural, mechanical and geometric factors subserving arm posture in humans. *J Neurosci*, 5:2732-2743.

Nakano E, Imamizu H, Osu R, Uno Y, Gomi H, Yoshioka T, and Kawato M (1999) Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. *J Neurophysiol*, 81:2140-2155.

Narendra, K. S. and Annaswamy, A. M. (1989). Stable adaptive systems. Prentice Hall, Englewood Cliffs, NJ.

Osu, R., Uno, Y., Koike, Y., and Kawato, M. (1997). Possible explanations for trajectory curvature in multijoint arm movements. *J Exp Psych: Hum Percep and Perform*, 3:890-913.

Poggio, T. and Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science*, 247:978-982.

Poggio, T. and Girosi, F. (1998). A sparse representation for function approximation. *Neural Computation*, 10:1445-1454.

Sabes, P. N., and Jordan, M. I (1997). Obstacle Avoidance and a Perturbation Sensitivity Model for Motor Planning. *J. Neurosci*, 17:7119-7128

Sabes, P. N., Jordan, M. I., and Wolpert, D. M. (1998). The Role of Inertial Sensitivity in Motor Planning. *J. Neurosci*, 18:5948-5957

Sainburg, R. L., Ghez C., and Kalakanis D. (1999). Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. *J Neurophysiol*, 81:1045-1056.

Sanner, R. M. and Kosha, M. (1999). A mathematical model of the adaptive control of human arm motions. *Biol Cybern*, 80:369-382.

Sanner R. M. and Slotine J.-J.E. (1995). Stable adaptive control of robot manipulators using "neural" networks. *Neural Computation*, 7:753-788

Schaal, S. and Atkeson, C. (1998). Constructive incremental learning from only local information. *Neural Computation*, 10:2047-2084.

Scott, S. H. and Kalaska, J. F. (1997). Reaching movements with similar hand paths but different arm orientation: I. activity of individual cells in motor cortex. *J Neurophysiol*, 77:826-852.

Sergio, L. E. and Kalaska, J. F. (1997). Systematic changes in directional tuning of motor cortex cell activity with hand location in the workspace during generation of static isometric forces in constant spatial directions. *J Neurophysiol*, 78:1170-74.

Sergio, L. E. and Kalaska, J. F. (1998). Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *J Neurophysiol*, 80:1577-1583.

Shadmehr, R. and Brashers-Krug, T. (1997). Functional stages in the formation of human long-term motor memory. *J Neurosci*, 17:409-419.

Shadmehr, R. and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *J Neurosci*, 14(5):3208-3224.

Slotine, J.-J. E., and Li, W. (1991). *Applied nonlinear control*. Prentice Hall, Englewood Cliffs, NJ.

Stroeve, S. (1997). A learning feedback and feedforward neuromuscular control model for two degrees of freedom human arm movements. *Hum Mvmt Sci*, 16:621-651.

Thoroughman, K. A. and Shadmehr, R. (1999). Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci*, 19:8573-8588.

Tresch, M., Saltiel, P., and Bizzi, E. (1999). The construction of movement by the spinal cord. *Nature Neurosci*, 2:162-167.

Uno, Y., Kawato, M., and Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement: A minimum torque change model. *Biol Cybern*, 61:89-101.

Vetter P., Goodbody S. J., Wolpert D. M. (1999) Evidence for an eye-centered spherical representation of the visuomotor map. *J Neurophysiol*, 81:935-939.

Wada, Y. and Kawato, M. (1993). A neural network model for arm trajectory formation using forward and inverse dynamics models. *Neural Networks*, 6:919-32.

Welsh, A. H. (1996). *Aspects of Statistical Inference*. Wiley and Sons, New York.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1994) Perceptual distortion contributes to the curvature of human reaching movements. *Exp Brain Res*, 98:153-156.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995a). An internal model for sensorimotor integration. *Science*, 269:1880-82.

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

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cog Sci*, 1:209-216.

Wolpert, D. M., Miall, R.C., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cog Sci*, 2:338-347.

Wood, J. E., Meek, S. G., and Jacobsen, S. C. (1989). Quantitation of human shoulder anatomy for prosthetic arm control - ii. anatomy matrices. *J Biomech*, 22:309-25.

## Curriculum Vitae

### Kurt Andrew Thoroughman

Born 31 January 1972 in Kalamazoo, Michigan

#### Education

1993 – 1999 Johns Hopkins University (JHU), School of Medicine (SOM)

Doctoral Program, Department of Biomedical Engineering

1989 – 1993 University of Chicago, Chicago, Illinois

BA in Physics, General and Departmental Honors

1985 – 1989 Central High School, Kalamazoo, Michigan

#### Teaching Experience

1997, 1998 Teaching Assistant, Neuroscience, JHU SOM

1996, 1997 Instructor, Introduction to Neuroscience, CTY, JHU

1995, 1997 Teaching Assistant, Department of Biomedical Engineering, JHU SOM

1994 Instructor, Fast-Paced Physics, Center for Talented Youth (CTY), JHU

1992 – 1993 Teaching Assistant, Department of Mathematics, University of Chicago

#### Invited Talks

2000 “Short-term unlearning reveals the computational basis of human motor learning,” at Neural Control of Movement

1997 “Human motor memory processes,” at Human Cognition and How it Fails,

**Publications and Abstracts**

- 1999                    Thoroughman KA and Shadmehr R. Electromyographic correlates of learning an internal model of reaching movements. *Journal of Neuroscience*, 19:8573-8588.
- 1999                    Shadmehr R and Thoroughman KA. Learning and memory formation of arm movements. In Biomechanics and Neural Control of Posture and Movement. JM Winters and PE Crago, (eds), New York: Springer-Verlag (in press).
- 1999                    Thoroughman KA and Shadmehr R. EMG correlates of learning and unlearning of nonstationary environments. *Soc Neurosci Abstr* 29:2178.
- 1998                    Thoroughman KA and Shadmehr R. Changes in EMG during learning of reaching movements. *Soc Neurosci Abstr* 24:419.
- 1997                    Thoroughman KA and Shadmehr R. Influence of after-effects on learning of an internal model of reaching movements. *Soc Neurosci Abstr* 23:203.
- 1997                    Thoroughman KA, Weiler MA, Lahti AC, Tamminga CA, and Holcomb HH. Implicit motor learning and memory in schizophrenia. *International Congress on Schizophrenia Research*.
- 1996                    Thoroughman KA and Shadmehr R. Forward interference in motor learning. *Soc Neurosci Abstr* 22:897.