

MECHANISMS OF ADAPTIVE MOTOR CONTROL

by

DAVID W. FRANKLIN

B.Sc. (Honours), Simon Fraser University, 1995

M.Sc., Simon Fraser University, 2000

DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in the School

of

Kinesiology

© David W. Franklin 2004
SIMON FRASER UNIVERSITY
September 2004

All rights reserved. This work may not be reproduced in whole or in part, by
photocopy or other means, without permission of the author.

APPROVAL

NAME: David W. Franklin
DEGREE: Doctor of Philosophy
TITLE OF THESIS: Mechanisms of Adaptive Motor Control

EXAMINING COMMITTEE:

Chair: Dr. David Goodman
Professor

Dr. Theodore Milner
Professor, School of Kinesiology
Senior Supervisor

Dr. Parveen Bawa
Professor, School of Kinesiology
Supervisor

Dr. Mitsuo Kawato
Director, ATR Computational Neuroscience Labs
ATR, Kyoto Japan
Professor, Nara Institute of Science and Technology
Supervisor

Dr. Richard Vaughan
Assistant Professor, Computing Science
Internal Examiner

Dr. Neville Hogan
Professor, Mechanical Engineering
Massachusetts Institute of Technology
External Examiner

Date Approved: 1st September, 2004

SIMON FRASER UNIVERSITY



PARTIAL COPYRIGHT LICENCE

The author, whose copyright is declared on the title page of this work, has granted to Simon Fraser University the right to lend this thesis, project or extended essay to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users.

The author has further granted permission to Simon Fraser University to keep or make a digital copy for use in its circulating collection.

The author has further agreed that permission for multiple copying of this work for scholarly purposes may be granted by either the author or the Dean of Graduate Studies.

It is understood that copying or publication of this work for financial gain shall not be allowed without the author's written permission.

Permission for public performance, or limited permission for private scholarly use, of any multimedia materials forming part of this work, may have been granted by the author. This information may be found on the separately catalogued multimedia material and in the signed Partial Copyright Licence.

The original Partial Copyright Licence attesting to these terms, and signed by this author, may be found in the original bound copy of this work, retained in the Simon Fraser University Archive.

W. A. C. Bennett Library
Simon Fraser University
Burnaby, BC, Canada

SIMON FRASER UNIVERSITY



Ethics Approval

The author, whose name appears on the title page of this work, has obtained human research ethics approval from the Simon Fraser University Office of Research Ethics for the research described in this work, or has conducted the research as a co-investigator of a project, or member of a course, approved by the Ethics Office.

A copy of the human research ethics approval letter has been filed at the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for ethics approval and letter of approval is filed with the Office of Research Ethics. Inquiries may be directed to that Office.

W. A. C. Bennett Library
Simon Fraser University
Burnaby, BC, Canada

ABSTRACT

To manipulate objects or use tools we must compensate for any forces arising from interaction with the physical environment. It has been suggested that this compensation is achieved by learning an inverse model of the dynamics. However, many interactions involve not only compensation for forces, but also compensation for unstable or unpredictable environments. The adaptation to both stable and unstable environments was investigated in order to characterize the changes which occur during and after adaptation and examine the mechanism of this adaptation process. Subjects made point to point reaching movements of a robotic manipulandum which generated either a stable or unstable interaction with the limb. Changes in kinematics, dynamics, muscle activity and stiffness were examined to try to elucidate the mechanism of adaptation. In stable dynamics a change in endpoint stiffness occurred that was directly related to changes in net joint torque necessary to compensate for the dynamics. However in unstable dynamics there was a selective change in endpoint stiffness without any modification of net joint torque. Despite this difference in final stiffness adaptation to the two force fields, the indices of learning suggested that the CNS used both an impedance controller and inverse dynamics model during adaptation to both force fields. The adaptation to unstable dynamics was further examined by looking at various levels of instability. The results confirmed that the endpoint stiffness was selectively adapted to stabilize the instability of the environment while reducing the metabolic cost. Finally, I examined the neural mechanisms by which motor learning takes place by analyzing the trial-by-

trial changes in muscle activation during adaptation to stable and unstable dynamics. The results are interpreted in terms of a computationally simple, robust learning mechanism which learns to adapt to both stable and unstable dynamics. This novel learning algorithm can explain the changes in stiffness, trajectory, force, and muscle activity through a mechanism which simultaneously learns both the necessary changes in joint torques and the optimal endpoint stiffness.

ACKNOWLEDGMENT

I especially thank my supervisors and collaborators on this and related works: Dr Ted Milner, Dr. Mitsuo Kawato, Dr. Rieko Osu and Dr. Etienne Burdet. I would also like to thank Dr. Parveen Bawa for her support and encouragement during this PhD. I thank the Co-op students who have traveled to work with me in Japan assisting with some of the work incorporated into this thesis (Udell So) as well as working on other projects, allowing me to focus more of my time to this thesis (Frances Leung, Bernard Ng and Gary Liaw). Of course almost none of this research could have been performed without the wonderful support staff at ATR, especially Toshinori Yoshioka who is responsible for the PFM.

I would like to thank my parents, Derek and York Franklin, for their support. I also thank all of the rest of my family and friends for their help in keeping me sane throughout this process. In particular, I thank Sarah for always finding a place for me to stay whenever I arrived from Japan back to Vancouver. Finally I especially thank Sae Masuda for her incredible support and love.

This research was supported in part by the National Institute of Information and Communications Technology of Japan, the Natural Sciences and Engineering Research Council of Canada, the Human Frontier Science Program, and the Advanced Telecommunications Research Institute (ATR) where all of the experiments were performed.

TABLE OF CONTENTS

Approval	ii
Abstract	iii
Acknowledgements	v
Table of Contents	vi
List of Tables	vii
List of Figures	viii
Background	1
Introduction	17
Specific Aims	19
I: Methods	22
II: Functional Significance of Stiffness in Adaptation of Multijoint Arm Movements to Stable and Unstable Dynamics	32
Introduction	32
Methods	35
Results	49
Discussion	66
III: Adaptation to Stable and Unstable Dynamics Achieved by Combined Impedance Control and Inverse Dynamics Model	71
Introduction	71
Methods	74
Results	86
Discussion	103
Appendix A	113
IV: Impedance Control Balances Stability with Metabolically Costly Muscle Activation	115
Introduction	115
Methods	117
Results	123
Discussion	137
V: New principles for a universal framework of motor learning	145
Introduction	145
Methods	147
Results	155
Discussion	163
Conclusion	173
References	182

LIST OF TABLES

Table 3.1	Summary of the least square fit to the adaptation of handpath error and EMG during learning in the VF and DF	100
------------------	--	-----

LIST OF FIGURES

Figure 1.1	The PFM and experimental setup	24
Figure 1.2	Electrode placement used in the experiments	26
Figure 2.1	Movements in the force fields	51
Figure 2.2	Force profiles after adaptation to the force fields	52
Figure 2.3	Mean endpoint force after adaptation to the force fields	53
Figure 2.4	Joint torque profiles after adaptation to the force fields	54
Figure 2.5	Stiffness geometry after adaptation to different force fields	55
Figure 2.6	Relations between reconstructed and measured changes in force for displacements in the VF and DF	58
Figure 2.7	The stiffness geometry predicted by modeling joint stiffness as a linear function of joint torque	61
Figure 2.8	Changes in joint stiffness associated with the change in endpoint stiffness	63
Figure 2.9	Rectified, averaged and smoothed surface EMG during the DF after effect trials	65
Figure 3.1	Experimental setup to study the adaptation to stable and unstable dynamics	75
Figure 3.2	Timing of corrective (reflexive and voluntary) EMG in the VF and DF	85
Figure 3.3	Movements in the VF and DF	87
Figure 3.4	Change in shoulder and elbow torque during learning in the VF and DF	89
Figure 3.5	Early changes in shoulder and elbow torque during learning in the DF	93
Figure 3.6	Muscle activity after adaptation to the VF and DF	94

Figure 3.7	Model fits to activity of the triceps long head during learning in the VF for various intervals	96
Figure 3.8	Evolution of the EMG for six arm muscles during learning in the VF	99
Figure 3.9	Evolution of the EMG for six arm muscles during learning in the DF	102
Figure 3.10	Summary of the time course of events taking place during learning in the VF and DF	105
Figure 4.1	Unstable force field initially amplifies the variability of trajectories	124
Figure 4.2	Hand path error is reduced during adaptation to the DF	127
Figure 4.3	Endpoint forces were similar across all force field strengths after learning	130
Figure 4.4	Endpoint stiffness scales with strength of the unstable force field	131
Figure 4.5	Endpoint stiffness reduced in y-direction with sufficient training	133
Figure 4.6	The x- and y- components of endpoint stiffness together with the net stiffness of the interaction between the limb and the force field	134
Figure 4.7	Joint stiffness increases independently of joint torque for adaptation to unstable dynamics	136
Figure 5.1	Changes in muscle activity during learning of a novel skill	157
Figure 5.2	Determining the time of feedback onset during perturbed trials	158
Figure 5.3	Feedback responses in the shoulder muscles for three time intervals	160
Figure 5.4	The change in the feedforward command is dependent upon the feedback of the previous trial	161
Figure 5.5	Change in feedforward activity shown relative to the signed kinematic error on the previous trial during learning	162

BACKGROUND

Humans have excellent abilities for motor adaptation. We can learn new skills, adapt to new surroundings, and interact with new objects. In order to control movement, we need to learn a set of motor commands which will produce the necessary muscle activation, and hence muscle forces, to propel the limb in the desired direction at the desired speed. This means that we need to specify a set of time-varying motor commands which will produce the desired movement. This could involve simply controlling the trajectory of the arm or also exerting force against some object, such as when lifting a glass of water to our mouth for a drink.

Goal-directed movements of single joints are usually performed with a bi- or tri-phasic pattern of muscle activity (Brown and Cooke 1981a; Hallett et al. 1975). The movement is performed by alternating bursts of activity in the agonist and antagonist muscles. This pattern of muscle activity has been shown to be generated by central feedforward commands (Hallett et al. 1975; Rothwell et al. 1982). However, feedback from the periphery during the movement can modify the pattern of activity. This occurs if movements are blocked (Angel 1977), disturbed by torque pulses (Cooke 1980), or subjected to modified afferent feedback by tendon vibration (Capaday and Cooke 1983). Therefore, although the fundamental pattern of activation is produced by centrally generated commands (feedforward activity), it is also modulated through afferent feedback.

In order move faster or farther, the size and timing of the tri-phasic muscle pattern is modified. The first burst of activity (in the agonist) accelerates the limb towards the target. It is increased in size to increase the speed or amplitude of the movement (Brown and Cooke 1981a; Hallett and Marsden 1979) and its duration is decreased as the speed of movement increases (Berardelli et al. 1984; Brown and Cooke 1984; Mustard and Lee 1987). The second burst of activity (in the antagonist) decelerates the limb (Brown and Cooke 1990), whereas the third burst (in the agonist) is related to the maintaining the limb at the target position (Ghez and Martin 1982; Hannaford and Stark 1985). All of these components can be modified to change the amplitude, duration, or symmetry of the velocity profile (Brown and Cooke 1981b; Brown and Cooke 1990; Cooke and Brown 1990; Cooke and Brown 1994; Hallett and Marsden 1979).

One of the first studies examining the changes in muscle activity during learning was for a complex movement of the thumb (Person 1958). In this study, the initial movements included co-contraction from which a pattern of reciprocal activation of the agonist and antagonist muscles gradually emerged. Darling and Cooke (Darling and Cooke 1987a; Darling and Cooke 1987b) had subjects learn to make fast targeted movements with their elbow. As the subjects learned to make faster movements, the variability of their movements decreased (Darling and Cooke 1987a). This was achieved by producing larger agonist and antagonist muscle bursts (Darling and Cooke 1987b) to the increase the acceleration necessary to increase movement speed. The variability of the muscle activity and kinematics decreased with practice. This appeared to also coincide with a gradual decrease in the antagonist

muscle activity. It appears, therefore, that initial movements may have excessive co-contraction which is gradually decreased as the subjects practice.

Multi-joint Movements

Normally, movements are produced by the coordinated action of several joints. Even the movement of a single joint such as lifting the hand and forearm with the elbow, will require compensatory muscle actions at the shoulder in order to offset the interaction forces which are produced when the arm begins to move. For example, if the forearm is lifted upwards in front of the body, this will produce an extensor torque at the shoulder which would need to be resisted by shoulder flexor muscles in order to maintain the shoulder posture. This becomes even more critical in fast multi-joint movements where the interaction torques can dominate the dynamics of the movements (Hollerbach and Flash 1982). Although it was at first thought that the passive properties of muscles could compensate for these forces to reduce the complexity of the computations performed by the central nervous system (CNS) (Flash 1987), the impedance of the limbs was found to be insufficient for this purpose (Gomi and Kawato 1996; Katayama and Kawato 1993). Instead, the CNS must compensate for these forces by learning a pattern of muscle activation which produces muscle forces to deal with the time varying interactions torques.

Evidence that the brain compensates for interaction torques was found by examining patients with neuropathologies that left them without peripheral proprioceptive information. The reaching trajectories of these subjects were consistent with the idea that they did not compensate for the interaction torques

(Sainburg et al. 1995; Sainburg et al. 1993). This also demonstrates that proprioceptive information is essential for the adaptation process. Further evidence that the brain compensates for interaction torques was obtained from the patterns of muscle activity of normal subjects during movements in which either the elbow or shoulder remained stationary (Gribble and Ostry 1999). This means that our brain needs to be able to shape patterns of activation in order to achieve the smooth kinematics (Atkeson and Hollerbach 1985) which characterize our movements. In order to understand how this takes place, researchers began to study how the kinematics, dynamics and muscle activation change during adaptation to novel dynamics.

Learning and Motor Adaptation

In motor adaptation studies, subjects were required to make simple point to point movements in different directions to a target. Shadmehr and Mussa-Ivaldi (1994) investigated movements in a curl force field. This force field is dependent on the velocity of the movement, initially causing large deviations in trajectories. However, after several hundred learning trials, the subjects were capable of consistently moving along a straight path to the target. When the force field was then unexpectedly removed, the subjects' trajectories were again disturbed, but this time in the opposite direction compared to the original disturbed trajectories. These trajectories are referred to as after-effects. This study showed that the subjects were able to change the joint torques during the movement to adapt to the force field and

did not simply stiffen the joints by co-contraction in order to adapt. If that had been the case, then removal of the force field should have resulted in straight movements. It also indicates that the learned compensation is feedforward in origin as reverting to the original dynamics, which had no perturbing effect, now disturbs the limb trajectories. At the same time, another group of researchers were studying adaptation to Coriolis forces (Lackner and Dizio 1994). Subjects were placed in a rotating room and asked to produce point to point movements. The Coriolis force initially disturbed the limb trajectories but subjects were able to adapt to the novel environment. The interpretation of these studies was that subjects had developed an internal model of the new dynamics and combined this with a previously learned internal model of the original null force field. Internal models are neural representations of the relationship between motor commands and their effect on movement of the limbs (Kawato 1999). The learned internal model produces the feedforward control of the movement, while a viscoelastic or feedback system corrects for errors (Shadmehr and Mussa-Ivaldi 1994).

What is learned?

What exactly is learned when subjects adapt to the dynamics of the environment? Subjects could either be learning the necessary endpoint force to resist the force field or the joint torques which give rise to this endpoint force. Shadmehr and Mussa-Ivaldi (1994) investigated this by first having subjects learn a force field in one part of the workspace. Then they had the subjects move in another part of the workspace that required a different arm configuration. The force field in this part of

the workspace was either had the same characteristics as the previous force field in terms of endpoint force (and therefore different in terms of joint torque) or had the same characteristics in terms of joint torque (and therefore different in terms of endpoint force). The subjects' movements were perturbed less in the new force field when the joint torque needed to compensate for the force field was the same as what they had previously learned. The subjects, therefore, had a better internal representation of the force field in terms of joint torque than endpoint force (Shadmehr and Mussa-Ivaldi 1994). A similar result was reported for a virtual object manipulation task even when no movement was required (Mah and Mussa-Ivaldi 2003).

Several studies have attempted to examine what parameters the motor control system relies on when adapting to the environment. For example, when we interact with objects in the environment, time varying forces are applied to our limbs. Subjects could index these forces for compensation in terms of time, sequence or state. However, studies have shown that subjects are not able to learn either a time dependent force field (Conditt and Mussa-Ivaldi 1999; Karniel and Mussa-Ivaldi 2003) or a sequence dependent force field (Karniel and Mussa-Ivaldi 2003) well. Only when the time-varying forces could be indexed in terms of state variables such as position or velocity was there generalization of the force field that could be demonstrated by after-effects. If a time-dependent force field was created which was similar to a force field described by state variables, then subjects incorrectly learned the state-variable dependency. For movements in directions different from the ones learned, there were large kinematic errors (Conditt and Mussa-Ivaldi 1999).

Therefore, it appears that the neural structures responsible for motor learning rely primarily on state-dependent information.

When subjects are learning a force field, the inclusion of catch trials (null force field trials) induces short-term unlearning of the force field (Thoroughman and Shadmehr 2000). The effects can be seen both in the trajectory and the EMG of force field trials following the catch trial, but they gradually disappear. In contrast, a virtual wall has no effect on the movements (Scheidt et al. 2000). When subjects were unknowingly constrained to move within a double-sided virtual wall (mechanical channel) after learning a force field, they pushed against the channel wall for up to a hundred trials after the force field had been replaced by the channel and produced large after-effects even after 25 trials. In this case, a slow de-adaptation occurs, suggesting that a slow reduction of effort or loss of internal force calibration likely occurred. These studies illustrate that the presence of kinematic error drives the learning process even if the kinematic error is produced by a null force field (Thoroughman and Shadmehr 2000) or is smaller than 1 cm (Scheidt et al. 2001). However, in the absence of kinematic error, the internal model is maintained even when not required (Scheidt et al. 2000).

When the force field produces consistent forces on the limb throughout learning for each movement direction, then the learning proceeds until the movements are fairly straight, similar to null field movements. However, if the strength of the force field varies from trial to trial, then to what does the internal model adapt? Over all, subjects appear to adapt roughly to the mean value of the force field (Scheidt et al. 2001). However, the most important influence on the adaptive force produced a given

trial was the strength of the force field on the previous trial. In fact, using a model which was based only on state information from the last few trials, Scheidt et al. (2001) were able to account for about 85% of the variance during the adaptation process and reproduce the size of after-effects on catch trials. Therefore, while the internal model adapts to the mean strength of a variable strength force field, it is highly influenced by the last few trials. This could explain the unlearning effect produced by the inclusion of catch trials (Thoroughman and Shadmehr 2000).

Generalization

It is apparent that the nervous system is able to learn time varying motor commands to compensate for the complex forces produced by interaction forces due to the motion of limb segments or interaction forces produced by the environment. However, does the CNS need to explicitly learn the appropriate compensation for every possible movement or are we able to generalize what we have previously learned and apply it to similar movements or to different regions of the workspace. This question of the ability of the CNS to generalize motor learning to movements in other directions was investigated by Gandolfo et al. (Gandolfo et al. 1996). The subjects trained in only two directions separated by 45 degrees. After learning was complete, after-effects were investigated for movements in a variety of movement directions. In the region between the directions used in training, and for neighboring directions, after-effects were evident. However, the magnitude of the after-effects decreased smoothly as the movement directions were displaced farther from the training directions. This study illustrates that the brain is able to generalize the

learned motor adaptation to directions different from those which were initially practiced. This is similar to visuomotor adaptation which smoothly generalizes between two learned movement directions with different visual distortions, but decays away from the trained directions of movement (Ghahramani et al. 1996). (Conditt et al. 1997) examined generalization to movements not previously practiced in the force field. In particular, after subjects had trained by performing point to point movements, they attempted to draw circles in the NF. They produced misshapen circles in the NF which were similar to the after-effects of subjects trained to draw circles in the force field. Furthermore, they were able to draw round circles when asked to do so in the learned force field. This study is significant because it demonstrates that the CNS is not learning the force fields by simple rote memorization. Instead it learns it in such a way that it can generalize to parts of the state space which it has not directly experienced. This generalization does not extend far beyond the range of movements which have been experienced, but it does generalize within this range (Conditt et al. 1997).

After subjects had learned a velocity dependent force field, Goodbody and Wolpert (Goodbody and Wolpert 1998) examined movements of half the duration and twice the amplitude. They found that subjects were able to use their previously learned model and scaled it linearly with the velocity in order to extrapolate beyond the velocity used in training. This linear extrapolation, however, levels off at high speeds. These studies demonstrate that the CNS is able to learn the mapping between the limb states and the forces that were required in order to generalize to regions of the workspace and state space.

How is this model of the dynamics learned and generalized to the environment?

By carefully examining the changes in forces and kinematics during adaptation to a velocity-dependent force field Thoroughman and Shadmehr (Thoroughman and Shadmehr 2000) demonstrated that a movement towards a target in one direction affects what has been learned (the internal model) in all other directions. This was done by analyzing how movement in one direction affected the kinematic error of the following movement as a function of its direction. There was a positive effect for directions close to the original movement direction, which was similar to what would have occurred in the original direction if the movement had been repeated. As the following movement direction deviated more from the original direction the effect decreased. For directions which differed by more than 90 degrees from the original movement direction, there was a negative effect. This meant that what was learned during the original movement degraded performance for directions opposite to the original target movement. This suggests that the dynamics of this force field across the state space was represented by broadly tuned Gaussian-like basis functions. This was also supported by further studies with other force fields (Donchin et al. 2003), although they found that the performance for the movements exactly 180 degrees away from the original movement direction improved. However, they also found that performance became worse for movements in perpendicular directions.

Consolidation

After subjects have initially learned a novel force field, what determines how well the memory is retained? In a series of experiments, Shadmehr and his colleagues

have investigated the consolidation of recent motor learning into long term memory. They showed that if a second force field was learned immediately after the first field, then it destroyed the memory of the initial force field when tested on the next day (Brashers-Krug et al. 1996). However, if the second force field was learned at least 4 hours after the first force field, then the memory of the first force field was retained on the second day of experiments and could be retained as long as five months. This suggests that it requires some time for the 'fragile' motor memory to be consolidated into a more robust long term memory. One particularly interesting result was related to the size of the after-effects during the learning of the second field (Shadmehr and Brashers-Krug 1997). If the second field was learned immediately after the first field then the size of the after-effects was reduced compared to when the second field was learned 5.5 hours later. At the earlier time, the new internal model appears to use the same memory space as the previously learned model causing interference with the new learning. Instead, when the second field is learned at a later time, the results indicated that subjects were using a memory space which was akin to 'tabula rasa.' However, if the second field is learned with a different arm configuration then the learning of the second field does not interfere (Gandolfo et al. 1996).

Changes in the muscle activity during learning

What changes in muscle activation occur during the adaptation to novel dynamics? When subjects were presented with novel dynamics, they were initially disturbed by the force fields (Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994). This is associated with a large change in the muscle activity delayed with

respect to the onset of the movement (Thoroughman and Shadmehr 1999). During early learning, extensive co-contraction of arm muscles was also seen. As learning progresses, and the kinematic disturbance decreases, this change in muscle activity shifted forwards in time. This shift in time indicates that the initial feedback is gradually reduced and changed to a feedforward command (Thoroughman and Shadmehr 1999). At the same time, they found that the co-contraction gradually decreased as learning progressed. This was similar to what was seen in single joint studies (Milner and Cloutier 1993) and novel multi-joint movements (Osu et al. 2002; Osu et al. 2002; Person 1958). The changes in muscle activity were interpreted in terms of changes in the preferred direction of muscles (Shadmehr and Moussavi 2000; Thoroughman and Shadmehr 1999) during the learning of an internal model. However, while the changes could indicate a change in the preferred direction of a muscle, they could simply result from the necessary change in muscle activation to compensate for the external force. Whether the adaptation process occurs through the change in the preferred direction of motor cortex neurons controlling the particular muscles involved in the task is not clear.

Control of Impedance

It is clear that motor adaptation changes the joint torque in order to compensate for the external dynamics of a novel environment. However, while muscles produce force roughly in proportion to their activation, they also increase the impedance of the limb. By co-contracting two or more antagonist muscles the

impedance of a joint or limb can even be modulated independently of the joint torque or endpoint force. The impedance of muscles, their resistance to external disturbances, is composed of two properties: stiffness and damping, which together are often referred to as viscoelasticity (Gasser and Hill 1924; Huxley and Simmons 1971). However, the viscoelasticity of muscle is non-linear and depends on a number of factors such as activation, velocity, length, and prior history (Kearney and Hunter 1990) as well as the reflexive gains of sensory receptors (Rack 1981).

Many muscles produce torque about a joint. All the muscles that have actions about a particular axis of rotation of the joint will contribute to the viscoelasticity about that axis. Joint stiffness has been shown to increase linearly with joint torque under isometric conditions (Cannon and Zahalak 1982); (Hunter and Kearney 1982; Weiss et al. 1988). Joint viscosity also increases linearly with joint torque under isometric conditions (Hunter and Kearney 1982; Weiss et al. 1988). The moment of inertia about a single joint remains constant with respect to joint torque (Hunter and Kearney 1982)

Joint elasticity is non-linear in response to displacements. Stiffness is highest with small displacements and decreases exponentially as displacement size increases (Kearney and Hunter 1982; MacKay et al. 1986). This occurs because cross-bridge bonds are broken as a muscle is stretched (Huxley and Simmons 1971).

At the joint level, muscles can have antagonistic functions. Two active muscles producing torques in opposing directions (co-contraction) may produce no net torque. However, the impedance of the joint is the sum of the impedance of all

muscles (Hogan 1984). Co-contraction allows the impedance of a joint to vary independently of joint torque (Hogan 1984; Milner et al. 1995).

Multi-Joint Mechanical Impedance

The multi-joint impedance of a limb was first examined by Mussa-Ivaldi, Hogan and Bizzi (1985), who developed a method for determining the magnitude of the passive stiffness at the hand, in different postures in the workspace. They perturbed the hand of the subject in eight directions and measured the force once the hand was at rest in the new posture. The change in force in response to the displacement not only had a component opposite to the direction of displacement but also had a component along the perpendicular direction. This was used to estimate the endpoint stiffness of the hand, which was represented as an ellipse after removing the effects of non-conservative forces. Similar to single-joint stiffness, the endpoint stiffness decreases with increasing perturbation displacement (Shadmehr et al. 1993).

The postural behavior of the stiffness has strong directional character (anisotropy) and varies in a regular way with arm configuration (Flash and Mussa-Ivaldi 1990; Mussa-Ivaldi et al. 1985; Tsuji et al. 1995). The endpoint viscosity also varies with posture (Dolan et al. 1993; Tsuji et al. 1995). The viscosity was rotated counter-clockwise slightly (5°) with respect to stiffness and changed similarly to stiffness with changes in posture.

Several studies have examined multi-joint viscoelasticity during force control tasks in the arm (Franklin and Milner 2003; Perreault et al. 2004) and the finger (Milner and Franklin 1998). Single joint stiffness during multi-joint force control has

been shown to increase linearly with joint torque for both the elbow and shoulder. The double joint stiffness increased linearly with elbow torque. Damping also roughly followed this pattern of stiffness, increasing with joint torque, although the correlation coefficients were small (Gomi and Osu 1998). However, in a recent experiment, it appears that endpoint damping is not modulated with endpoint stiffness (Perreault et al. 2004). Instead while stiffness increased linearly with generated force, the damping increased in a non-linear fashion.

Previous work of Gomi and Osu (1998) found that joint stiffness increased linearly with joint torque in a isometric force control task. In single joint studies, it has been shown that the subjects co-contract their muscles when it is necessary to increase the stability without changing the joint torque. This is also possible in the multijoint case, where co-contraction of all muscles can result in an increase in the size of the endpoint stiffness (Mussa-Ivaldi et al. 1985). In contrast to force control tasks where the environment stabilizes the arm or joint, stiffness is modified to compensate for load stability in position control tasks of single joints (Akazawa et al. 1983; Doemges and Rack 1992a; Doemges and Rack 1992b). This is achieved due both to increased muscle co-contraction and reflex gain. More recently it has been shown that in the multi-joint arm, a similar change in the components of the stiffness matrix occurs to maintain stability under position control tasks (Franklin and Milner 2003) suggesting that the endpoint stiffness of the arm is modified independently of joint torque to achieve stability. Recently, it has been shown that subjects are able to learn to selectively increase the endpoint stiffness of their arm in the direction of an instability without changing the endpoint force (Burdet et al. 2001) finally

demonstrating that humans are able to selectively control the impedance of their limbs (Hogan 1984).

Control of Impedance and Force

The previous results indicate that when adapting our environment we are capable of adapting both the force we apply to it and the impedance of our limbs. The ability to control both impedance and force at the same time was nicely demonstrated by asking subjects to make movements in randomly varying environments (Takahashi et al. 2001). In this study, the subjects both increased impedance to resist the random variations in the force field as evidenced by a decrease in the after-effect size and formed an internal model of the mean force field as evidenced by the decrease in trajectory error. The increased impedance would be used to stabilize the limb against instabilities in the environment whereas the change in force could be used to change the trajectory. For example, when sculptors are working on their carvings, they need to produce a force with their carving tools to extract the material. At the same time, they must stiffen up their limbs in order to both stabilize the inherent instability of tool use and resist any sudden changes in material consistency. Failure to do so could cause their carving implements to slip and damage the sculpture. Similar results would occur for almost all tool use tasks (Rancourt and Hogan 2001).

INTRODUCTION

The overall goal of this study is to investigate the adaptive abilities of the motor control system of humans and to propose a mechanism by which this might be performed. Humans, more so than most animals, skillfully manipulate our environment using tools. Our ability to use tools is one of our defining characteristics, and has shaped the evolution of our species. While this ability is not limited to humans, we alone, appear to have the amazing range and control to adapt our motor system to so many different tools, so many different environmental dynamics, and so many novel tasks. In order to perform even apparently simple actions with tools, we must not only control the time-varying forces at the endpoint of our limbs, but stabilize the entire linkage with the appropriate impedance. Furthermore, we often manage to learn these tasks while limiting energy consumption. The first time we perform an action, we generally monitor the entire movement continuously. However, with practice, the action becomes more and more automatic, requiring less and less correction. This is called adaptation or learning. Using common everyday tasks is difficult when studying adaptation because all subjects have varying levels of experience with these tasks. Therefore, we examine this process using dynamical environments generated with a robot.

Extensive work has been done looking at adaptation to stable dynamic environments. However, little is known about how we adapt to unstable dynamics. In

this study I will compare and contrast the learning and adaptation to both stable and unstable dynamics. I will examine changes in the trajectories, endpoint forces, joint torques, electrical activity of muscles and endpoint stiffness during and after the adaptation processes. By doing this I hope to characterize the abilities of the motor control system for adaptation. Finally, I would like to propose a possible mechanism which is responsible for both types of adaptation.

SPECIFIC AIMS

1. Stiffness changes associated with adaptation to stable and unstable environments.

Burdet et al. (2001) examined stiffness adaptation to movements in an unstable environment. They showed that humans are able to selectively change the endpoint stiffness of their arm to compensate for directional instabilities in the environment. I examine this adaptation further and compared the stiffness adaptation in the unstable interaction with stiffness adaptation in a stable interaction. In particular, I test if the stiffness changes in the two environments could be produced by changes in the joint torques required to produce the movements in the two tasks or if they were related to a pure co-contraction of antagonist muscles. This would indicate if the limb impedance was being controlled directly by the CNS or if it was a by-product of the changes in endpoint force and joint torques.

2. Changes in muscle activity during adaptation to stable and unstable environments.

While measuring endpoint stiffness after adaptation to the dynamics can tell us the state of the final adaptation, it provides little information about the way in which the CNS may perform this adaptation. In particular, I test if the adaptation to novel dynamics occurred with both an inverse dynamics model and an impedance controller even if only one of these was required for final adaptation.

In order to examine this, I record the electrical activity of six muscles before, during and after adaptation to stable and unstable dynamics.

3. Selective control of endpoint stiffness in an unstable environment. Burdet et al. (2001) demonstrated that the endpoint stiffness of the arm is selectively increased in the direction of instability during movements in an unstable environment. However, only a change in rotation of the ellipse was found as only one level of instability was examined. While the net stiffness (arm-environment) after adaptation was the same as in the null force field, it is still unclear if this means that the CNS is adaptively controlling the net stiffness of the arm in interaction with the environment or not. In order to examine this, subjects adapted to several levels of instability of the environment, ranging from less than their arm stiffness (a stable interaction) to greater than their arm stiffness (an unstable interaction). I test if the CNS is controlling the level of endpoint stiffness to adapt to the level of environmental instability.
4. Mechanisms of adaptation. The three previous experiments examine the types of changes that occur during and after adaptation to both stable and unstable dynamics. In this portion of the thesis, I examine the trial by trial changes in the EMG during the adaptation process. By separating the feedforward and feedback components of the motor command, I can investigate how the feedforward command changes between each trial. Based on the results of these experiments I propose a physiologically plausible learning mechanism which can adapt to both stable and unstable environments.

The first three of these studies have now been published. They appear as:

Franklin, D. W., Burdet, E., Osu, R., Kawato, M., and Milner, T. E. Functional significance of stiffness in adaptation of multijoint arm movements to stable and unstable dynamics. *Exp Brain Res* 151: 145-57, 2003.

Franklin, D. W., Osu, R., Burdet, E., Kawato, M., and Milner, T. E. Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J Neurophysiol* 90: 3270-82, 2003.

Franklin, D. W., So, U., Kawato, M., and Milner, T. E. Impedance Control Balances Stability with Metabolically Costly Muscle Activation. *J Neurophysiol* (June 16, 2004). doi:10.1152/jn.00364.2004.

These three papers are reproduced in this thesis by permission of Springer and the American Physiological Society.

I: METHODS

The experiments in this thesis all use similar methodology. The details of the techniques used are explained in this section in more detail than that found for each experiment.

Apparatus

Subjects sat in a chair and moved the parallel-link direct drive air-magnet floating manipulandum (PFM) in a series of forward reaching movements performed in the horizontal plane (Figure 1.1). The PFM itself is a horizontal planar manipulator (robotic interface), powered by two DC direct-drive motors controlled at 2 kHz by a digital signal processor as has previously been explained in (Gomi and Kawato 1996, 1997). The thin links of the robot, which are attached to the subject by way of the handle, are driven by the wide links. These wide links are in turn driven by two electric motors located under the metal table. The handle itself is supported by a friction-free air magnet floating mechanism which supports some of the weight of the arm and prevents the handle from tilting and bending the thin links of the robotic interface. The subject's shoulders were held against the back of the chair by means of a shoulder harness. The right forearm was securely coupled to the PFM using a rigid custom molded thermoplastic cuff. The cuff functioned as a rigid connection between the subject's arm and the robotic interface. It also immobilized the wrist joint, permitting movement of only the shoulder and elbow joints. The subject's right

forearm rested on a support beam projecting from the handle of the PFM. Motion was, therefore, limited to a single degree of freedom at the shoulder and at the elbow. The subject's hand position was determined using optical encoders which sensed motor shaft angle (409600 pulse/rev) and the force exerted by the hand was measured using a force sensor (Nitta Corp. No. 328) located between the handle and the manipulandum (resolution 0.059 N). Position and force were sampled at 500 Hz. The minimum mechanical rigidity of the PFM within the normal working area is 3.28×10^4 N/m when it is under control. The maximum speed, acceleration and force are 4 m/s, 50 m/s^2 , and 150 N respectively.

Reaching movements were performed from a start point ($[x, y] = [0, 0.31] \text{ m}$) to a target located at ($[0, 0.56] \text{ m}$) relative to the subjects shoulder ($[0, 0]$) for a total movement length of 0.25 m. The start and target circles along with the instantaneous hand position were projected onto opaque plexiglass plates covering the workspace. These plates also prevented the subject from viewing his/her arm.




Figure 1.1: The PFM and experimental setup. The PFM is a planar two link robotic interface which can exert forces on the hand during horizontal point-to-point arm movements performed by the subject. The hand was linked by means of a stiff thermoplastic cuff to a handle at the end of the robot. Subjects make reaching movements from a start circle to a target circle, both 2.5 cm in diameter. The center of the start circle was located 31 cm in front of the shoulder, i.e. at (0, 0.31) m relative to the shoulder, while the center of the target circle was 56 cm in front of the shoulder, i.e. at (0, 0.56) m.

Electromyography

Surface electromyographic (EMG) activity of six arm muscles was recorded using pairs of disposable silver-silver chloride surface electrodes. The electrode locations were chosen to maximize the signal from a particular muscle while avoiding cross-talk from other muscles. The skin was cleansed with alcohol and prepared by rubbing in electrode paste. This was removed with a dry cloth and pre-gelled electrodes were then attached to the skin with tape. The spacing between the electrodes of each pair was approximately 2 cm. The impedance of each electrode pair was tested to ensure that it was below 10 k Ω .

The activity of two monoarticular shoulder muscles, pectoralis major and posterior deltoid, two biarticular muscles, biceps brachii and long head of the triceps, and two monoarticular elbow muscles, brachioradialis and lateral head of the triceps, was recorded. The EMG signals were analog filtered at 25 Hz (high pass) and 1.0 kHz (low pass) using a Nihon Kohden amplifier (MME-3132) and then sampled at 2.0 kHz. EMG was aligned on the movement onset and averaged over twenty trials to visually compare changes that had occurred during learning.

Electrodes were oriented along the muscle fiber direction and placed over the center of the muscle to maximize the signal (Figure 1.2). The subject's arm was positioned in the horizontal plane as it would be in the experiments with the elbow flexed about 90 degrees. The subject contracted each muscle in turn after being instructed by the experimenter. The pectoralis major electrode site was located approximately 5 cm proximal and 5 cm inferior to the glenohumeral joint. The posterior deltoid electrodes were placed on the posterior surface, approximately at the glenohumeral joint. The biceps brachii electrodes were placed over the middle of the two heads of the biceps muscle, approximately 10 cm proximal to the elbow joint. The triceps longus electrodes were placed approximately 15 cm distal to the shoulder joint, along the inferior-posterior edge of the arm. The brachio-radialis electrodes were placed about 5 cm distal to the elbow joint along the superior-medial edge. Finally the triceps lateralis electrodes were placed about 10 cm proximal to the elbow joint along the superior-posterior edge of the arm. The electrodes were positioned on an angle relative to the long axis of the upper arm along the direction of the muscle fibers.

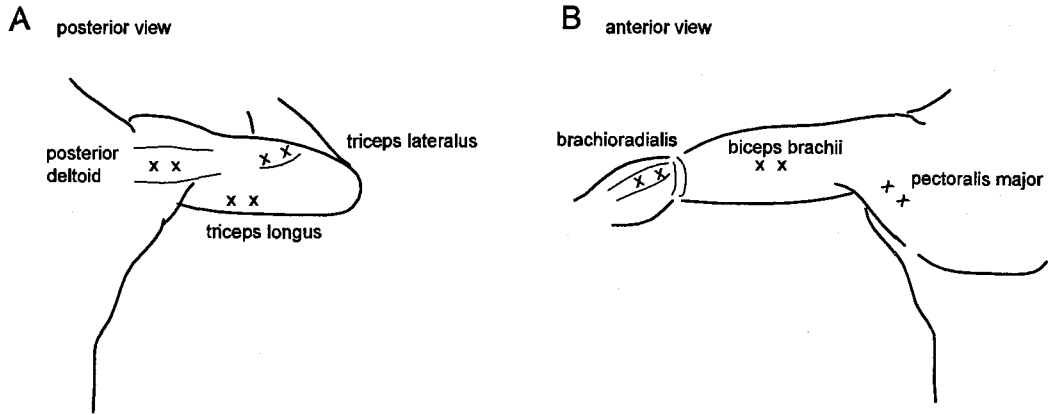


Figure 1.2: Electrode placement used in the experiments. A. Posterior view of subject. Shown are the placements for extensor muscles: posterior deltoid, triceps longus and triceps lateralis. B. Anterior view of subject. Shown are the flexor muscles: pectoralis major, biceps brachii, and brachioradialis.

Force Fields

The experiment examined stiffness and EMG adaptation in two force fields: a velocity dependent force field (VF), producing a stable interaction with the arm, and a position dependent (divergent) force field (DF), producing an unstable interaction. Results were compared to those in a null field (NF). The force (F_x, F_y) (in N) exerted on the hand by the robotic interface in the VF was computed as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \chi \begin{bmatrix} 13 & -18 \\ 18 & 13 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (1.1)$$

where (\dot{x}, \dot{y}) is the hand velocity (m/s) and the scaling factor, χ , was adjusted to the subject's strength ($2/3 \leq \chi \leq 1$). In particular, χ was 2/3 for women, 15/18 for most men and 1 for the largest men. The DF produced a negative elastic force perpendicular to

the target direction with a value of zero along the y -axis, i.e., no force was exerted when trajectories followed the y -axis, but the hand was pushed away whenever it deviated from the y -axis. The DF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \quad (1.2)$$

where the x -component of the hand position was measured relative to the shoulder joint. $\beta > 0$ (N/m) was adjusted for each subject so that it was larger than the stiffness of the arm measured in NF movements so as to produce an unstable interaction. Both force fields were inactivated once the subject reached the target position.

Learning

All subjects practiced making movements in the NF on at least one day prior to the experiment. These training trials were used to accustom the subjects to the equipment and to the movement speed and accuracy requirements. Experiments were normally specified in terms of number of successful movements in order to maintain subject motivation to completing the task requirements. Successful trials were those which ended inside the 2.5 cm diameter target window within the prescribed time (0.6 ± 0.1 s). All movements during learning trials were recorded whether successful or not. Movements were self-paced so subjects were able to rest between movements if they wished. On each trial, the start circle, the target circle and the instantaneous hand position, represented by a 0.5 cm diameter cursor, were projected onto an opaque

horizontal surface covering the arm. Positioning the hand in the start circle initiated a sequence of 3 beeps at 500 ms intervals. The subject was instructed to begin movement on the third beep and complete it on a fourth beep, 600 ms later. Two additional beeps followed at 500 ms intervals to indicate the target hold time. Feedback of the performance indicating final hand position (OK, OUT) and movement duration (OK, LONG, SHORT) was given after each trial by displaying a message on a monitor in front of the subjects. Duration was considered OK if it was 600 ± 100 ms.

Impedance estimation

We measured stiffness in VF and DF movements after extensive learning, as well as in NF movements. The basic principles of trajectory prediction, used in the stiffness measurement, were originally devised by Franklin and Milner (1997) and are described in detail by Burdet et al. (2000). The method for estimating stiffness was based on a method originally employed by Bennett (1993) which uses a controlled displacement relative to a prediction of the undisturbed trajectory. Bennett used the mean of past trajectories as a prediction of the current trajectory. However, this becomes much less accurate for the multi-joint movements. The trajectory variation from one trial to the next is quite large in multi-joint movements because they are less constrained.

The prediction algorithm proceeds as follows. Initially, 10 unperturbed movements were recorded. The velocity profiles of these movements were low-pass filtered at 50 Hz, scaled to produce the same amplitude movement (in the main

movement direction) and truncated at both ends using a 0.03 m/s velocity threshold. The profiles were then averaged. However, to avoid distortion of the template due to differences in the times of the peak velocity, the mean was computed independently for the acceleration and deceleration phases of the movement. A mean velocity profile was computed for both the x - and y -axes. These mean velocity profiles were then scaled over a range of 11 amplitudes and 11 time shifts to generate 121 candidate velocity profiles to be used for trajectory prediction. On each trial, the candidate best matching the current trajectory up to the time of perturbation onset was selected and used as the prediction. The best match was determined by minimizing the following recursive function after the velocity of the movement had crossed the 0.03 m/s threshold:

$$d_{i,k} = (\mathbf{v}_k - \mathbf{w}_{i,k})^2 + \alpha d_{i,k-1} \quad (1.3)$$

where \mathbf{v}_k is the velocity of the current time sample k , $\mathbf{w}_{i,k}$ is the i^{th} velocity candidate at time k and $\alpha = 0.94$ is a forgetting factor which limits matching to the last 100 ms. The movement was decomposed parallel and perpendicular to the main movement direction and candidate velocity profiles were selected independently for the x - and y -directions.

The mean velocity profile was updated after each unperturbed trial during a session. If the latest unperturbed movement was within 1.05 times the standard deviation of the currently used ten trials for the template, then it was included in the

ten stored trials while the oldest trial was removed. If it was outside this range then the template was not changed.

Prior to each stiffness measurement session, the subjects were retrained in the force field to ensure that they had readapted to the field (NF: 40 trials; DF, VF: 80 trials). For three subjects, one hundred and sixty movements were then performed in each force field, of which eighty were randomly selected for stiffness measurement. Half of these were randomly selected for stiffness measurement. On these trials, displacements were introduced at the midpoint of the movement in one of 8 directions chosen randomly from the set $\{0^\circ, 45^\circ, 90^\circ, 135^\circ, 180^\circ, 225^\circ, 270^\circ, 315^\circ\}$. The PFM briefly displaced the hand by a constant distance from the prediction of the undisturbed trajectory (Burdet et al. 2000). This displacement had an amplitude of 8 mm and lasted 300 ms. The perturbation profile consisted of a 100 ms smooth ramp away from the current trajectory, a 100 ms hold portion and 100 ms smooth ramp back to the predicted trajectory. In order to minimize vibration, a sixth order polynomial with zero velocity and acceleration at beginning and end and zero end jerk was used as the servo command for the ramp portions of the displacement. During the hold phase, the hand was moved with the predicted velocity of the unperturbed movement. Assuming that the prediction were perfect there would be no difference in velocity between the perturbed and unperturbed trajectories, eliminating any contribution of damping or acceleration to the change in measured endpoint force. Although the prediction is not perfect, our results indicate that the errors are small and that the average prediction over several trials is very close to the average of the actual trajectory (Burdet et al. 2000). Therefore, we can be quite confident that

forces due to damping or acceleration did not introduce error in the stiffness measurements. Using the average force and displacement during a 60 ms interval towards the end of the hold portion of the perturbation window, an estimate of the 2x2 endpoint stiffness matrix (**K**) was obtained by linear regression of the mean change in hand force and the mean change in position, as represented by the equation:

$$\begin{bmatrix} \Delta F_x \\ \Delta F_y \end{bmatrix} = \mathbf{K} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \quad (1.4)$$

The stiffness in different directions was represented in terms of an ellipse by plotting the elastic force produced by a unit displacement (Mussa-Ivaldi et al. 1985).

II: FUNCTIONAL SIGNIFICANCE OF STIFFNESS IN ADAPTATION OF MULTIJOINT ARM MOVEMENTS TO STABLE AND UNSTABLE DYNAMICS

INTRODUCTION

We constantly interact with the world around us; moving and manipulating objects in our environment. Tasks such as opening a door, which involve a stable interaction with the environment, are relatively simple to learn, because similar motor commands will result in similar movements, so the dynamics can be easily identified and compensated. For example, we may be surprised the first time we open a door with high friction, but will be able to open the door smoothly on the second or third trial. Unstable tasks are more difficult to learn, because they are affected by different initial conditions, neuro-motor noise (Schmidt et al. 1979; Slifkin and Newell 1999) or any small external perturbation, which can lead to unpredictable, inconsistent and unsuccessful performance. For example, during sculpting, material irregularities can displace the chisel to the left or to the right of the intended path, and it requires extensive practice for a sculptor to acquire the skill necessary to compensate for such instability.

Many tasks that humans perform, particularly those involving tool use, are inherently unstable (Rancourt and Hogan 2001). However, while the tasks may be unstable, the mechanical impedance of the musculoskeletal system, which is

stabilizing in nature, counteracts the instability. Ultimately, it is the interaction between our limbs and the environment that determines whether or not a movement will be stable. Coupled stability of the limb and the environment is a prerequisite for successful actions, as it provides robustness to motor output variability and perturbations from the environment. Limb impedance can be modified by changing the force at the hand since muscle stiffness inherently scales with activation level (Hunter and Kearney 1982). Modification of impedance geometry, e.g., shape and/or orientation of endpoint stiffness, has been reported (Gomi and Osu 1998; Lacquaniti et al. 1993; McIntyre et al. 1996) but is generally associated with a change in applied force, or a change in the limb configuration. Such inherent modulation of impedance with applied force often cannot be used to achieve stability, e.g., when the task constrains the direction or magnitude of the hand force. Similarly, task demands may limit the variation in the endpoint impedance that could be achieved by changing the limb configuration. To deal with such situations, the CNS must be able to modify the size, shape or orientation of the impedance independently of the force applied by the hand. This could be achieved by co-contracting specific groups of muscles to select the geometry of the endpoint stiffness (Hogan 1985). We have recently demonstrated that the CNS can optimize the magnitude, shape and orientation of the endpoint stiffness of the arm to compensate for environmental instabilities (Burdet et al. 2001).

If the CNS is concerned about minimizing metabolic cost then we would expect that the endpoint stiffness, in tasks which are mechanically stable, would be directly related to the minimum muscle torque necessary to move the arm and compensate for any external dynamics. Specifically, if metabolic energy is to be

minimized, then muscles will be recruited with the minimal activation so as to produce the desired net joint torques. Accordingly, we would expect a relatively linear relation between joint torque and joint stiffness similar to that found under stable isometric conditions (Gomi and Osu 1998; Perreault et al. 2001) as muscle and joint stiffness have been found to increase monotonically with muscle force and joint torque, respectively (Carter et al. 1993; Kirsch et al. 1994; Milner et al. 1995). It should, therefore, be possible to predict the joint stiffness from the net torque generated by muscles during performance of the task. This is in contrast to tasks which require greater stability than would be conferred by the minimum muscle activation necessary to move the arm and compensate for any external dynamics. Under such unstable conditions, co-contraction of opposing muscles may be required above and beyond that ordinarily occurring with the movement. In such cases, the joint stiffness should be higher than that predicted by the net joint torque.

The present study was undertaken to analyze and compare adaptation to environments producing stable or unstable interactions with the arm from the perspective of changes in limb mechanics. Previously, we measured the endpoint stiffness of the arm after adaptation to mechanical instability and found a selective increase in stiffness in the direction of instability, which appeared to compensate precisely for the instability (Burdet et al. 2001). We extend that work in the current paper by presenting new data and analysis of stiffness after adaptation to instability, as well as examining stiffness after adaptation to stable dynamics. Adaptation to stable dynamics, which has been examined previously by looking at kinematics and EMG (Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999), has

never been characterized in terms of the change in endpoint stiffness. In order to properly compare and contrast the adaptation to stable and unstable environments we rigorously modeled the trial to trial variation in the force response to displacement. We investigated whether the change in stiffness observed after adaptation to the stable or unstable environment was controlled independently of the change in joint torque needed to compensate for the new dynamics. In an unstable environment, the change in stiffness was independent of the change in joint torque. However, in a stable environment, we found that changes in endpoint stiffness were well correlated with changes in joint torque. Therefore, while the change in endpoint stiffness in the unstable environment was directly controlled and could be attributed to adaptive co-contraction of antagonist muscle pairs, the change in endpoint stiffness in the stable force field reflected only the adaptive change in joint torque.

METHODS

Six healthy individuals participated in the study (20-34 years of age; 2 females and 4 males). The institutional ethics committee approved the experiments and the subjects gave informed consent prior to participation.

Apparatus

Subjects sat in a chair and moved the parallel-link direct drive air-magnet floating manipulandum (PFM) in a series of forward reaching movements performed in the horizontal plane. Their shoulders were held against the back of the chair by

means of a shoulder harness. The right forearm was securely coupled to the PFM using a rigid custom molded thermoplastic cuff. The cuff immobilized the wrist joint, permitting movement of only the shoulder and elbow joints. The subjects' right forearm rested on a support beam projecting from the handle of the PFM. Motion was, therefore, limited to a single degree of freedom at the shoulder and at the elbow. The manipulandum and setup are described in detail elsewhere (Gomi and Kawato 1996; Gomi and Kawato 1997). In brief, the subjects' hand position was measured with joint-position sensors (409600 pulse/rev) and the force exerted by the hand was measured using a force sensor located between the handle and the manipulandum (resolution 0.059 N). Position and force were sampled at 500 Hz.

Subjects performed reaching movements from a start circle located 0.31 m in front of the shoulder to a target circle located 0.56 m in front of the shoulder (total distance 0.25 m). The start and target circles along with the instantaneous hand position were projected onto an opaque horizontal surface located directly overtop of the subject's arm. This surface also removed the subject's arm from view.

Force Fields

The experiment examined stiffness and EMG adaptation in two force fields: a velocity dependent force field (VF), producing a stable interaction with the arm, and a position dependent (divergent) force field (DF), producing an unstable interaction. Results were compared to those in a null field (NF). The force (F_x , F_y) (in N) exerted on the hand by the robotic interface in the VF was computed as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \chi \begin{bmatrix} 13 & -18 \\ 18 & 13 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (2.1)$$

where (\dot{x}, \dot{y}) is the hand velocity (m/s) and the scaling factor, χ , was adjusted to the subject's strength ($2/3 \leq \chi \leq 1$). In particular, χ was $2/3$ for women, $15/18$ for most men and 1 for the largest men. The DF produced a negative elastic force perpendicular to the target direction with a value of zero along the y -axis, i.e., no force was exerted when trajectories followed the y -axis, but the hand was pushed away whenever it deviated from the y -axis. The DF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \quad (2.2)$$

where the x -component of the hand position was measured relative to the shoulder joint. $\beta > 0$ (N/m) was adjusted for each subject so that it was larger than the stiffness of the arm measured in NF movements so as to produce an unstable interaction. Both force fields were inactivated once the subject reached the target position.

Learning

All subjects practiced making movements in the NF on at least one day prior to the experiment. These training trials were used to accustom the subjects to the equipment and to the movement speed and accuracy requirements. Subjects were randomly assigned to one of two groups. Group 1 initially performed the experiments

with the DF, and then proceeded to the VF, whereas group 2 adapted to the fields in reverse order. Stiffness was normally measured the day after adaptation to a particular force field once sufficient practice in the field had taken place (see impedance estimation below).

Subjects first practiced in the NF until they had achieved 50 successful trials. Successful trials were those which ended inside the 2.5 cm diameter target window within the prescribed time (0.6 ± 0.1 s). All movements were recorded whether successful or not. Movements were self-paced so subjects were able to rest between movements if they wished. At the completion of 50 successful trials, the force field was activated. No information was given to the subjects as to when the force field trials would begin. Subjects then practiced in the force field until achieving 75 successful trials. They took a short break and then performed 100 more movements, 20 of which were random trials in the NF. The NF trials were called after effects and were recorded to confirm that subjects adapted to the force field.

Impedance estimation

We measured stiffness in VF and DF movements after extensive learning, as well as in NF movements. The method is described in detail by (Burdet et al. 2000). Prior to each stiffness measurement session, the subjects were retrained in the force field to ensure that they had readapted to the field (NF: 40 trials; DF, VF: 80 trials). For three subjects, one hundred and sixty movements were then performed in each force field, of which eighty were randomly selected for stiffness measurement. For one subject, in the case of all three force fields, and for two other subjects, in the case

of two of the force fields, only eighty movements were performed. Half of these were randomly selected for stiffness measurement (as in (Burdet et al. 2001)). On these trials, displacements were introduced at the midpoint of the movement in one of 8 directions chosen randomly from the set $\{0^\circ, 45^\circ, 90^\circ, 135^\circ, 180^\circ, 225^\circ, 270^\circ, 315^\circ\}$. The PFM briefly displaced the hand by a constant distance from a prediction of the undisturbed trajectory (Burdet et al. 2000). This displacement had an amplitude of 8 mm and lasted 300 ms. This was composed of a 100 ms ramp away from the current trajectory, a 100 ms hold portion and 100 ms ramp back towards the predicted trajectory. During the hold phase of the perturbation, the hand was displaced with the predicted velocity of the unperturbed movement. Assuming that the prediction is perfect there would be no difference in velocity between the perturbed and unperturbed trajectories, eliminating any contribution of damping to the change in measured endpoint force. Although the prediction is not perfect, our results indicate that the errors are small and that the average prediction over several trials is very close to the average of the actual trajectory (Burdet et al. 2000). Therefore, we can be quite confident that forces due to damping did not introduce error in the stiffness measurements. Using the average force and displacement during a 60 ms interval towards the end of the hold portion of the perturbation window, an estimate of the 2x2 endpoint stiffness matrix (\mathbf{K}) was obtained by linear regression of the mean change in hand force and the mean change in position, as represented by the equation:

$$\begin{bmatrix} \Delta F_x \\ \Delta F_y \end{bmatrix} = \mathbf{K} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \quad (2.3)$$

The stiffness in different directions was represented in terms of an ellipse by plotting the elastic force produced by a unit displacement (Mussa-Ivaldi et al. 1985).

In the perturbation trials in the VF, the force field was activated prior to and after the perturbation to ensure that the same trajectory and motor commands were used in the perturbation and non-perturbation trials. In the trials where the hand was displaced in the DF, the force field was not activated prior to the displacement in order to avoid amplification of trajectory deviations that would contribute to error in the stiffness estimates. Since subjects could not detect the absence of the force field on these trials we were able to rule out the possibility that changes in stiffness were reactive to the environment. The joint stiffness (\mathbf{R}) was calculated from the endpoint stiffness (\mathbf{K}) using the relation:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{ss} & \mathbf{R}_{se} \\ \mathbf{R}_{es} & \mathbf{R}_{ee} \end{bmatrix} = \mathbf{J}^T \mathbf{K} \mathbf{J} + \frac{\partial \mathbf{J}^T}{\partial \boldsymbol{\theta}} \mathbf{F} \quad (2.4)$$

where \mathbf{J} is the Jacobian, a matrix which represents the geometric transformation of small changes in joint angles to small changes in endpoint position, and \mathbf{F} is the endpoint force (McIntyre et al. 1996).

Stiffness dependence on joint torque

Previously, it had been shown that joint stiffness linearly increases with joint torque during postural tasks (Gomi and Osu 1998). Recently however, we

demonstrated that the endpoint stiffness after adaptation to an unstable force field (DF) was selectively increased along the x -axis without any corresponding mean change in net joint torque (Burdet et al. 2001). However, even if the mean change in joint torque is zero, this does not mean that the stiffness could not be produced by variations in the joint torque from trial to trial which sum to zero over the entire experiment. Previously, we looked for a simple correlation between the actual endpoint force in the x -direction and the expected force based on the measured stiffness and the size of the displacement. A more rigorous approach has been taken in the current work, which considers the forces and stiffness in both the x and y -directions. We wish to examine whether trial to trial variation in net muscle torque is correlated with the variation in measured stiffness. In the previous approach we calculated muscle torque from the endpoint force recorded during the perturbation. This value includes not only the force which would be produced by muscles during a normal movement but also the elastic force resisting the perturbation. Consequently, part of the measured force would have been correlated *a priori* with the expected change in force calculated from the measured stiffness. To avoid this confounding effect we have estimated the force necessary to compensate for the force field using our trajectory prediction algorithm (Burdet et al. 2000). For each trial, we use the trajectory in the first portion of the movement along with information from previous trials to predict what the trajectory would have been without the perturbation. From this trajectory we could determine the force that the force field would have applied to the subject to obtain a more accurate estimate of the net muscle torque at each joint. In this way, it is possible to investigate whether or not the changes in stiffness in the

VF and DF are dependent or independent of any change in joint torque while avoiding the confounding effects of the perturbations.

The following formulas use three different symbols to denote differences in variables: d , δ and Δ . For example, when referring to torque, d refers to the change in joint torque, $d\tau$, produced by displacing the hand, i.e., the torque resisting the displacement; δ refers to the difference between the resisting torque, $d\tau$, in one force field and $d\tau$ in a different force field; and Δ refers to the difference between the joint torque required to move the arm in one force field and the joint torque required for the same movement in a different force field.

When the arm is constrained to move in the horizontal plane at shoulder height, as in this experiment, it can be modeled as a two degree-of-freedom mechanical system. The associated 2x2 joint stiffness matrix \mathbf{R} is defined by:

$$\begin{bmatrix} d\tau_s \\ d\tau_e \end{bmatrix} = \begin{bmatrix} R_{ss} & R_{se} \\ R_{es} & R_{ee} \end{bmatrix} \begin{bmatrix} d\theta_s \\ d\theta_e \end{bmatrix} \quad (2.5)$$

where the change in joint torque $d\tau = [d\tau_s \ d\tau_e]^T$ is the elastic resistance to a differential displacement $d\theta = [d\theta_s \ d\theta_e]^T$. The stiffness \mathbf{R} can be estimated during arm movements by measuring the restoring force in response to small perturbations relative to a prediction of the undisturbed trajectory (Burdet et al. 2000). As we wish to examine the change in stiffness that occurs during the adaptation to the force fields, the change in torque associated with joint stiffness in the NF

$$\begin{bmatrix} d\tau_s^{NF} \\ d\tau_e^{NF} \end{bmatrix} = \begin{bmatrix} R_{ss}^{NF} & R_{se}^{NF} \\ R_{es}^{NF} & R_{ee}^{NF} \end{bmatrix} \begin{bmatrix} d\theta_s \\ d\theta_e \end{bmatrix}$$

is subtracted from the change in torque associated with joint stiffness in the force field

$$\begin{bmatrix} d\tau_s^{FF} \\ d\tau_e^{FF} \end{bmatrix} = \begin{bmatrix} R_{ss}^{FF} & R_{se}^{FF} \\ R_{es}^{FF} & R_{ee}^{FF} \end{bmatrix} \begin{bmatrix} d\theta_s \\ d\theta_e \end{bmatrix}$$

resulting in

$$\begin{bmatrix} \delta\tau_s \\ \delta\tau_e \end{bmatrix} = \begin{bmatrix} \Delta R_{ss} & \Delta R_{se} \\ \Delta R_{es} & \Delta R_{ee} \end{bmatrix} \begin{bmatrix} d\theta_s \\ d\theta_e \end{bmatrix} \quad (2.6)$$

where

$$\begin{bmatrix} \Delta R_{ss} & \Delta R_{se} \\ \Delta R_{es} & \Delta R_{ee} \end{bmatrix} = \begin{bmatrix} R_{ss}^{FF} - R_{ss}^{NF} & R_{se}^{FF} - R_{se}^{NF} \\ R_{es}^{FF} - R_{es}^{NF} & R_{ee}^{FF} - R_{ee}^{NF} \end{bmatrix},$$

$$\begin{bmatrix} \delta\tau_s \\ \delta\tau_e \end{bmatrix} = \begin{bmatrix} d\tau_s^{FF} - d\tau_s^{NF} \\ d\tau_e^{FF} - d\tau_e^{NF} \end{bmatrix}$$

and the change in joint angle $d\theta$ is assumed to be the same in all force fields.

Joint stiffness has been shown to be well correlated with joint torque for single joints such as the ankle (Hunter and Kearney 1982), elbow (Cannon and Zahalak 1982), wrist (Milner et al. 1995) and interphalangeal joints of the thumb (Akazawa et al. 1983) and index finger (Carter et al. 1990) during maintenance of posture, and for the elbow during movement (Bennett 1993). Similar correlations between joint stiffness and joint torque have been found during maintenance of multi-joint posture of the arm (Gomi and Osu 1998; McIntyre et al. 1996). Gomi and Osu (1998) showed that: i) the shoulder stiffness R_{ss} was well correlated with shoulder torque $|\tau_s|$ and weakly correlated with the elbow torque $|\tau_e|$; ii) the cross-joint stiffness elements R_{se} and R_{es} are of similar magnitude, well correlated with $|\tau_e|$ but poorly correlated with $|\tau_s|$; and iii) the elbow stiffness is well correlated with $|\tau_e|$ and weakly with $|\tau_s|$. However, as the biarticular muscles should contribute to all four elements of the joint stiffness matrix, and as biarticular muscle activation may be much higher during movement as compared to isometric tasks (Karst and Hasan 1991; Tax et al. 1990a; Tax et al. 1990b; Tax et al. 1989; van Groeningen and Erkelens 1994), the shoulder stiffness term may vary as a function of both $|\tau_s|$ and $|\tau_e|$. The presence of the elbow torque term in the expression for shoulder stiffness reflects the contribution of the biarticular muscles which have been shown to be activated primarily to produce elbow torque (Buchanan et al. 1986; Flanders and Soechting 1990; Gomi and Osu 1998; Osu and Gomi 1999; Wadman et al. 1980). In particular, we assumed that the dependence of joint stiffness on joint torque could be represented as:

$$\mathbf{R}^{NF} = \begin{bmatrix} \alpha_1 |\tau_s^{NF}| + \alpha_2 |\tau_e^{NF}| & \alpha_3 |\tau_e^{NF}| \\ \alpha_3 |\tau_e^{NF}| & \alpha_4 |\tau_e^{NF}| \end{bmatrix} + \mathbf{R}_p, \quad \alpha_1 \dots \alpha_4 \text{ constants} \quad (2.7)$$

and

$$\mathbf{R}^{FF} = \begin{bmatrix} \alpha_1 |\tau_s^{FF}| + \alpha_2 |\tau_e^{FF}| & \alpha_3 |\tau_e^{FF}| \\ \alpha_3 |\tau_e^{FF}| & \alpha_4 |\tau_e^{FF}| \end{bmatrix} + \mathbf{R}_p, \quad \alpha_1 \dots \alpha_4 \text{ constants} \quad (2.8)$$

where \mathbf{R}_p is the passive stiffness. Subtracting equation (2.7) from equation (2.8)

yields:

$$\Delta \mathbf{R} = \begin{bmatrix} \alpha_1 \Delta |\tau_s| + \alpha_2 \Delta |\tau_e| & \alpha_3 \Delta |\tau_e| \\ \alpha_3 \Delta |\tau_e| & \alpha_4 \Delta |\tau_e| \end{bmatrix}, \quad \alpha_1 \dots \alpha_4 \text{ constants} \quad (2.9)$$

where

$$\begin{bmatrix} \Delta |\tau_s| \\ \Delta |\tau_e| \end{bmatrix} = \begin{bmatrix} |\tau_s^{FF}| - |\tau_s^{NF}| \\ |\tau_e^{FF}| - |\tau_e^{NF}| \end{bmatrix}$$

Equation (2.6) was transformed from joint co-ordinates to endpoint co-ordinates to obtain the relation:

$$\Delta \mathbf{F} = (\mathbf{J}^{-1})^T \Delta \mathbf{R} \begin{bmatrix} d\theta_s \\ d\theta_e \end{bmatrix} \quad (2.10)$$

As the trajectories of movements in the VF and DF after learning were similar to movements in the NF and as $\mathbf{J}(\theta)$ changed little in the region of stiffness measurement, we assumed that $\mathbf{J}(\theta)$ was constant. To examine the correlation between the change in stiffness relative to the NF and the corresponding change in joint torque using individual trials, (2.9) can be substituted into (2.10) and rearranged as:

$$\Delta \mathbf{F} = (\mathbf{J}^{-1})^T \begin{bmatrix} d\theta_s \Delta|\tau_s| & d\theta_s \Delta|\tau_e| & d\theta_e \Delta|\tau_e| & 0 \\ 0 & 0 & d\theta_s \Delta|\tau_e| & d\theta_e \Delta|\tau_e| \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} \quad (2.11)$$

where $\mathbf{J}(\theta_0)$ is calculated from the joint angles of each subject at the midpoint of the movements θ_0 . The predicted trajectory was used in combination with the force field equation to calculate the endpoint force applied to the hand by the force field as described above. From the endpoint force and kinematics, the joint torques were then computed using inverse dynamics as explained in next section. In this way, the difference in joint torque ($\Delta|\tau|$) for force field movements, relative to NF movements could be estimated. On the other side of the equation, the actual difference in force ($\Delta \mathbf{F}$) produced by the perturbation in the force field relative to the NF was calculated

directly from the difference in measured restoring force ($d\mathbf{F}$) in the force field and the NF:

$$\Delta\mathbf{F} = \begin{bmatrix} \Delta F_x \\ \Delta F_y \end{bmatrix} = \begin{bmatrix} dF_x^{FF} - dF_x^{NF} \\ dF_y^{FF} - dF_y^{NF} \end{bmatrix} \quad (2.12)$$

To determine whether the stiffness was being controlled independently of the joint torque necessary to cancel the force exerted by either force field, the parameters α_i in equation (2.11) were determined by linear regression for each subject using the trials in which stiffness was measured. An R^2 value close to 1 would indicate that endpoint stiffness could be explained by the joint torque required to perform the task. An R^2 value close to zero would indicate that overall, from trial to trial, endpoint stiffness was controlled independently of joint torque.

Torque estimation

Time varying muscle torque at the shoulder and elbow was computed using the equations of motion for a two-link planar arm (c.f. (Hollerbach and Flash 1982).

The joint torque was calculated as:

$$\begin{aligned} \tau_s &= \ddot{\theta}_s (2X \cos(\theta_e) + Y + Z) + \ddot{\theta}_e (X \cos(\theta_e) + Y) - \dot{\theta}_e^2 X \sin(\theta_e) \\ &\quad - 2\dot{\theta}_s \dot{\theta}_e X \sin(\theta_e) - (l_1 \sin(\theta_s) + l_2 \sin(\theta_s + \theta_e)) F_x \\ &\quad + (l_1 \cos(\theta_s) + l_2 \cos(\theta_s + \theta_e)) F_y \\ \tau_e &= \ddot{\theta}_e Y + \ddot{\theta}_s (X \cos(\theta_e) + Y) + \dot{\theta}_s^2 X \sin(\theta_e) \\ &\quad - l_2 \sin(\theta_s + \theta_e) F_x + l_2 \cos(\theta_s + \theta_e) F_y \end{aligned} \quad (2.13)$$

where:

$$\begin{aligned}X &= m_2 l_1 c_{m2} + m_c l_c c_{mc} \\Y &= I_2 + m_2 c_{m2}^2 + I_c + m_c c_{md}^2 \\Z &= I_1 + m_1 c_{m1}^2 + (m_2 + m_c) l_1^2\end{aligned}$$

τ is joint torque, θ is joint angle (defined according to the convention of Mussa-Ivaldi et al. (1985)), I is moment of inertia about the center of mass (c_m) of the segment, l is segment length and m is segment mass. The subscript s refers to the shoulder joint (or the upper arm), e to the elbow (or the lower arm) and c to the wrist cuff. The mass and inertia of the subject's arm segments were estimated from the weight and segment lengths of each subject based on anthropometrical scaling relations (Winters 1990).

EMG measurement of DF adaptation

To examine the muscle activity after learning in the DF, surface EMG was recorded in the NF, after learning in the DF and during after effect trials in the DF for four of the six subjects. All EMG was recorded on the same day with the same electrode placement. Activity was recorded from six muscles producing torque at the shoulder and elbow joints. The muscles included two monoarticular shoulder muscles: the pectoralis major and the posterior deltoid; two biarticular muscles: the biceps brachii and the long head of the triceps; and two monoarticular elbow muscles: the brachioradialis, and the lateral head of the triceps. The EMG was recorded by using pairs of disposable silver-silver chloride surface electrodes in a bipolar configuration with a separation distance of approximately 2 cm. The skin was

thoroughly cleaned with alcohol and prepared by rubbing in electrode paste. Excess paste was wiped from the skin prior to attaching the electrodes. The resistance of each electrode pair was tested to ensure that it was less than 10 k Ω . EMG signals were filtered at 25 Hz (high pass) and 1 kHz (low pass) and sampled at 2 kHz.

EMG during the DF after effect trials was compared to EMG during NF movements and EMG after complete adaptation to the DF. The rectified EMG was integrated over the entire movement, from 100 ms prior to movement onset until 800 ms after movement onset. Twenty trials were used from each condition and the data for all subjects were used in an ANOVA with subjects as a random variable. A post-hoc test was then performed on the three conditions using Scheffe's test with a significance level of 0.05. The rectified EMG was smoothed using a 125 point moving average and averaged over twenty trials for purposes of illustration.

RESULTS

The two force fields investigated in this study produced distinctive perturbations of the trajectories prior to adaptation. Initial trials in the stable VF perturbed the subjects to the left (Fig. 2.1B), but they quickly learned to compensate for the force field and soon began to make straighter movements. By the 5th trial, subjects usually succeeded in moving to the final target in a single movement. After learning, the trajectories were relatively straight and consistently reached the final target position. Movements in the unstable DF were initially perturbed either to the right or the left (Fig. 2.1C), depending on the initial deviation in the path. However,

subjects were able to adapt to the force field, successfully completing the task on most trials and exhibiting straight trajectories to the final target. After learning, trajectories in both the VF and DF were similar to NF trajectories. The recorded after-effects in the VF were similar to those reported in previous studies using similar force fields (Shadmehr and Mussa-Ivaldi 1994), confirming that subjects adapted to the endpoint forces imposed by the force field. This has been previously suggested to indicate the formation of an internal model of the environmental dynamics (Shadmehr and Mussa-Ivaldi, 1994). In the DF, after-effect trajectories deviated less from the straight-line trajectory than did NF movements (Burdet et al. 2001).

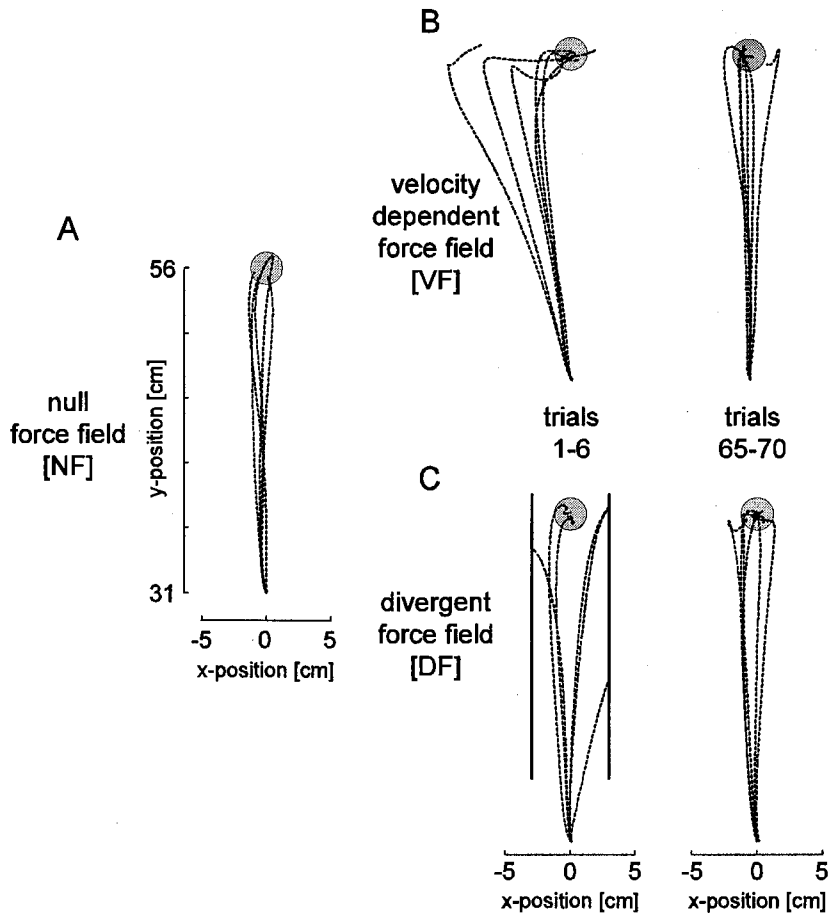


Figure 2.1 Movements in the force fields. A: in the NF. B: in the VF. C: in the DF. Movements are shown for the initial movements in the force field (trials 1-6) and the late portion of learning (trials 65-70). The solid lines on either side of trials 1-6 for the DF indicate the safety zone, outside of which the field was turned off for safety reasons.

Endpoint force after learning

The similarity of trajectories in the VF and the DF to those in the NF indicates that the subjects learned to compensate for the force fields. To compare how compensation for the dynamics was achieved in the two cases, we measured the endpoint force after learning. Figure 2.2 shows the endpoint force of subject 1 in the NF, VF and DF averaged over ten movements with standard deviations. Subjects adapted to the VF by counteracting the force produced by the field. The force was

more positive in the x -direction and more negative in the y -direction than in the NF. During the early stages of learning in the DF, subjects experienced large forces which could be either positive or negative depending on the initial movement direction. However, once subjects had adapted to this field, similar endpoint forces were recorded in the DF as had been recorded in the NF.




Figure 2.2 Force profiles after adaptation to the force fields **A**: Mean endpoint force profiles for the x - and **B**: y - directions during the movement after learning in the NF (grey), VF (dotted) and DF (solid black). Standard deviations are shown for movements based on the last twenty successful trials. Data are shown for one subject.

Figure 2.3 shows mean endpoint force and standard deviations at the midpoint of the trajectory (time of stiffness measurement) in the three force fields for six subjects. The endpoint force after learning in the VF was significantly larger (positive) in the x -direction and significantly smaller (more negative) in the y -direction compared to the endpoint force in the NF. For no subject was the x - or y -force at the midpoint of the movement trajectory in the DF significantly different from that in the NF (Fig. 2.3). In contrast to the VF, adaptation in the DF did not require a change in the endpoint force.




Figure 2.3 Mean endpoint force after adaptation to the force fields. **A:** in the x -direction. **B:** in the y -direction. The mean x - and y -forces and standard deviations at the midpoint of the movements in the NF (squares), VF (diamonds) and DF (stars) for all 6 subjects. The time of measurement corresponds to the time at which stiffness was estimated in the movements.

In order to further examine the difference in adaptation that was required in these two force fields we computed the joint torques using inverse dynamics (Fig. 2.4). In comparison to the NF, the adaptation to the VF required a larger extensor torque at the shoulder throughout the entire movement and a small extensor torque at the elbow late in the movement to compensate for the force field. Since no change in endpoint force occurred in the DF, the mean joint torque after adaptation was not different from the NF. There was, however, a larger variation in the joint torque in the DF than the NF because force was applied to the hand whenever the trajectory deviated from the y -axis. Because this force acted to increase the deviation, variation in the trajectory was also greater in the DF than the NF.




Figure 2.4 Joint torque profiles after adaptation to the force fields **A**: shoulder torque. **B**: elbow torque. The shoulder and elbow joint torque profiles are shown for the NF (grey), VF (dotted) and DF (solid black). Adaptation to the VF required an extensor torque at the shoulder whereas no changes in joint torque were observed in the DF. Data are shown for the same subject as Fig. 2.2.

The endpoint stiffness changes with adaptation to force fields

The endpoint stiffness of the arm, measured after adaptation to the NF, VF and DF is shown in Figure 2.5. The ellipses represent stiffness at the midpoint of the movement. In both the VF and DF, the endpoint stiffness was modified in shape and orientation compared to stiffness in the NF. In the VF, the endpoint stiffness increased along the direction of force compensation (Fig. 2.3), although this occurs only for specific force directions (c.f. Gomi and Osu, 1998). In the DF the stiffness increased dramatically in the direction of the instability (indicated by small black arrows for subject S6), but there was relatively little change in the movement direction. The increase in stiffness in the direction of instability was achieved without a change in the endpoint force (Fig. 2.3). This illustrates the ability to control the endpoint stiffness of the arm independently of endpoint force as previously described by Burdet et al. (2001).

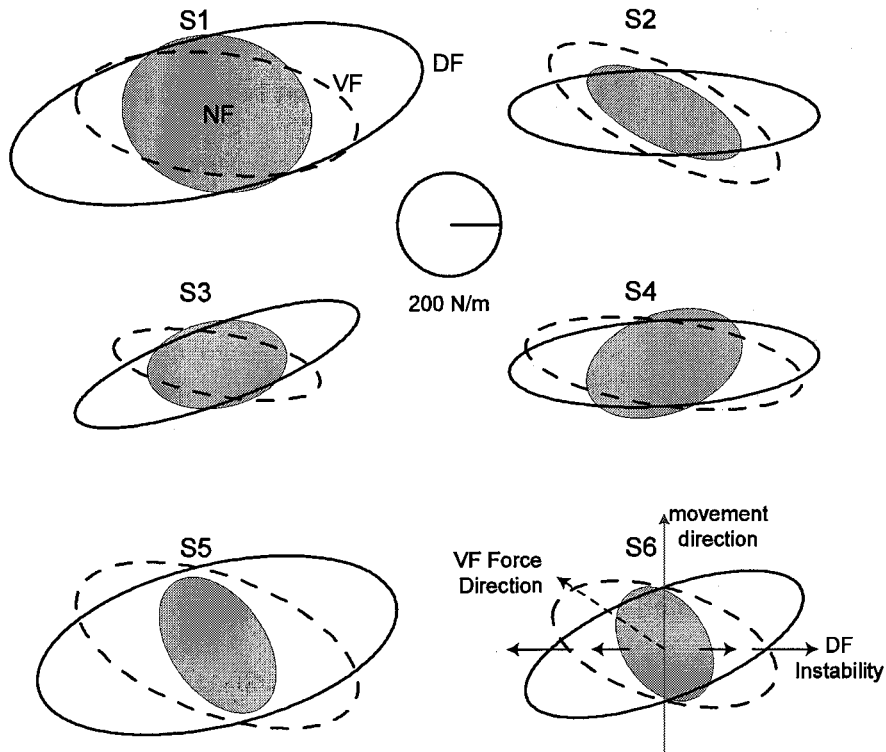


Figure 2.5 Stiffness geometry after adaptation to different force fields. Stiffness ellipses are shown for six subjects in the NF (light grey filled ellipse), VF (dotted ellipse) and DF (dark solid ellipse). For the ellipses of subject 6, a grey arrow illustrates movement direction, and dotted and solid arrows show the direction of the external force in the VF and direction of instability in the DF, respectively.

The endpoint stiffness can be decomposed into two components: a symmetric component (conservative) and an antisymmetric component (rotational) (Mussa-Ivaldi et al. 1985). Neither during estimation nor visualization procedures, did we constrain the endpoint stiffness to be symmetric. For comparison to previous work, the curl components of the endpoint stiffness were estimated as in Mussa-Ivaldi et al. (1985). In particular, we estimated Z_{mean} which is the square root of the ratio of the determinant of the antisymmetric stiffness to the determinant of the symmetric stiffness. Expressed as a percentage, we found the following Z_{mean} values (mean \pm standard deviation of six subjects) for stiffness under the three conditions, NF: $23.0 \pm$

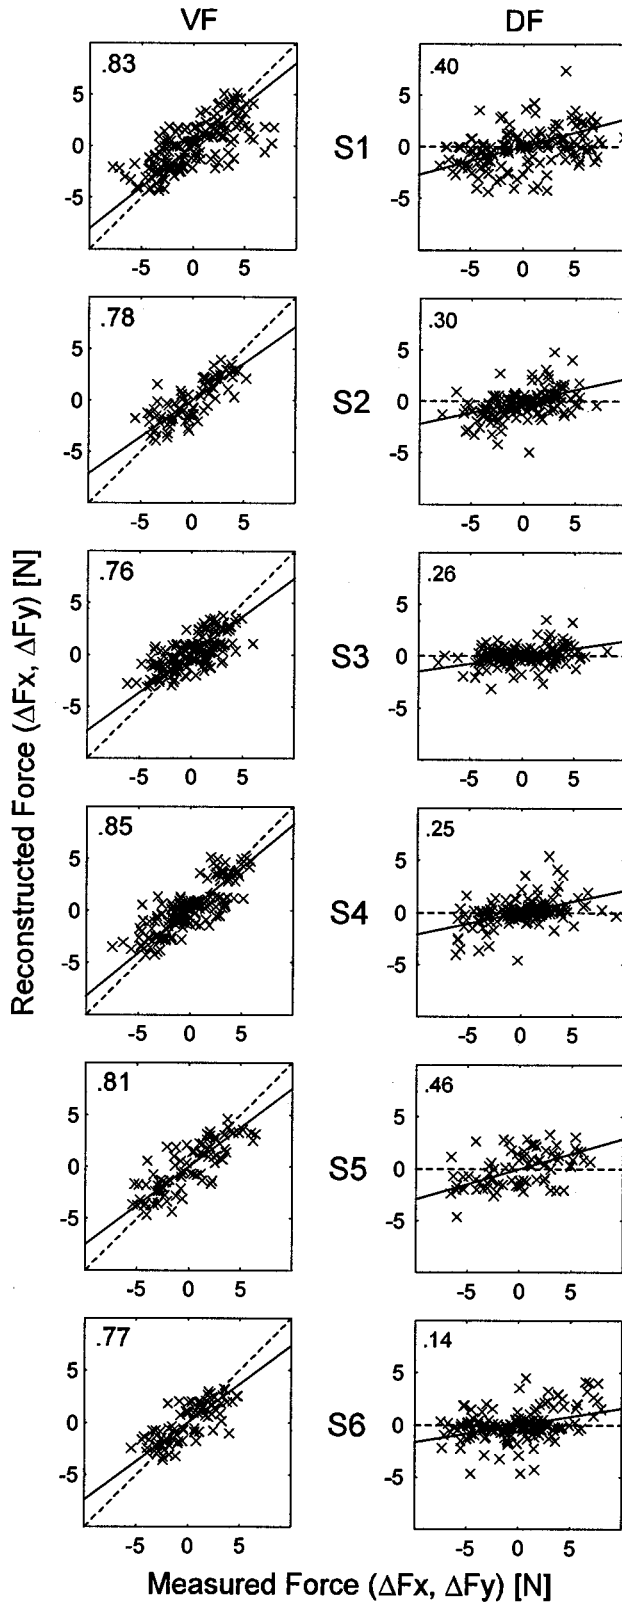
10.8; VF: 15.5 ± 11.1 ; DF: 24.6 ± 12.9 . In general, these are larger than those found under isometric conditions at rest (Mussa-Ivaldi et al., 1985). Differences in Z_{mean} across force fields was tested using an ANOVA with subjects as a random effect. There was no significant difference in the contribution of the antisymmetric component to the endpoint stiffness after adaptation to the DF compared to the NF ($p=0.83$) or after adaptation to the VF compared to the NF ($p=0.28$).

Stiffness in VF can be predicted only from change in joint torque

Endpoint stiffness increased in specific directions both for movements in the VF and the DF. However, the reasons for the increase in stiffness were likely different. To address this issue we performed an analysis to determine whether or not stiffness was being controlled independently of the joint torque necessary to compensate for the forces applied by the force field. We used the linear relation represented by equation (2.11) to predict the stiffness and so determine whether the increased stiffness could be explained by increased joint torque relative to NF movements. The change in endpoint force produced by the perturbation used to measure stiffness was estimated for each trial based on the prediction of the unperturbed trajectory and the field strength. Figure 2.6 compares the measured and estimated changes in endpoint force relative to endpoint force in the NF (ΔF_x and ΔF_y). The parameters α_i were estimated from the data obtained in the VF (left) and DF (right) for each subject using trials where stiffness was measured. From the results it is evident that the stiffness in the VF was relatively well correlated with the applied force, whereas in the DF the correlation was generally poor. The difference in slope

for the VF and DF relations demonstrates clearly that stiffness changed for different reasons. For the VF, the slope was close to one indicating that the difference in joint torque accounts well for the difference in resistive force due to the difference in stiffness relative to the NF. In contrast, for the DF the slope was close to zero for all subjects indicating that the difference in resistive force was independent of the difference in joint torque.

Figure 2.6 Relations between reconstructed and measured changes in force for displacements in the VF (left) and DF (right). The measured change in force is calculated from the recorded force data during the experiments (Equation 2.13). The reconstructed force is based on estimates of the joint torque, the displacement size and the estimates of α (Equation 2.12). Changes of endpoint force in the x - and y -direction relative to force in the NF (ΔF_x and ΔF_y) are plotted for each subject. R^2 values for the relation, calculated using data from both the x - and y -directions, are placed in the top left corner of each panel. The slope of each relation is shown with the solid line. For reference purposes, lines of slope one (VF) and slope zero (DF) are overlaid on each plot with dotted lines.



To further investigate how much of the difference in endpoint stiffness could be explained by the difference in joint torque we used equation (2.9) to compute the difference in joint stiffness in the force field relative to the NF, which was then added to the NF stiffness. The parameters α_i were estimated for each subject using the respective data for the VF and DF. The stiffness was computed using force data from individual trials and averaged. The resulting stiffness ellipses in Figure 2.7 illustrate the stiffness geometry that would have resulted if the stiffness scaled directly with the calculated joint torques. The modeled stiffness for movements in the VF shows an increase in stiffness along the same direction as the observed stiffness, of almost identical size and shape (Fig. 2.7A). In the DF, the modeled stiffness is essentially the same as the NF stiffness so the model fails to reproduce the observed increase in stiffness along the direction of instability (Fig. 2.7B). This indicates that the increased stiffness in the DF was not due to any increase in net joint torque accompanying adaptation to the DF and must, therefore, have been independently modulated due to co-contraction of antagonist muscles or contribution from reflexes.




Figure 2.7 The stiffness geometry predicted by modeling joint stiffness as a linear function of joint torque (equation (2.10)) (broken ellipse) with the stiffness measured in the NF (grey) and force field (solid) superimposed. A: VF and B: DF. Stiffness was computed for individual trials and averaged, using the force data for each subject. The ellipses represent the stiffness averaged across all six subjects.

Joint stiffness changes under VF and DF

The joint stiffness was calculated from the endpoint stiffness matrix (McIntyre et al. 1996). The elements of the joint stiffness matrix in the NF are shown in Figure 2.8A. The relative values for all subjects were similar in the NF. After adaptation to the VF, the shoulder joint stiffness (R_{ss}) and cross-joint stiffness (R_{se} and R_{es}) increased by 36.7 Nm/rad, 11.7 Nm/rad and 13.8 Nm/rad, respectively, while the elbow joint stiffness (R_{ee}) increased by only 2.9 Nm/rad (Fig. 2.8B). The greatest changes were in the shoulder joint and cross-joint stiffness terms. This indicates that the major contribution came from the shoulder extensor muscles, followed by the biarticular extensor muscle.

The changes in joint stiffness terms, seen after adaptation to the DF, present a different pattern (Fig. 2.8C). In this case, the increases in the shoulder joint stiffness term (82.1 Nm/rad) and cross-joint stiffness terms (49.5 and 60.8 Nm/rad) were more comparable, with a somewhat smaller increase in single joint elbow stiffness (39.3

Nm/rad). As the biarticular muscles must contribute to the single joint elbow and shoulder stiffness, as well as to the cross-joint stiffness terms it seems that the single joint muscles likely contribute little to the increased joint stiffness.

One noteworthy point is that in the NF the cross-joint stiffness term, R_{se} , is larger than the elbow stiffness, R_{ee} , for all five subjects. The other cross-joint term, R_{es} , is almost the same as R_{ee} . After adaptation either to the VF or the DF, the cross-joint stiffness terms increased by more than the single joint elbow stiffness term. This finding could be explained by biarticular muscles contributing more to cross-joint stiffness than to single joint elbow stiffness in this condition. The actual contribution to single joint stiffness relative to shoulder stiffness or cross-joint stiffness will be dependent on the ratios of the moment arms of the biarticular muscles at the elbow and shoulder joints.

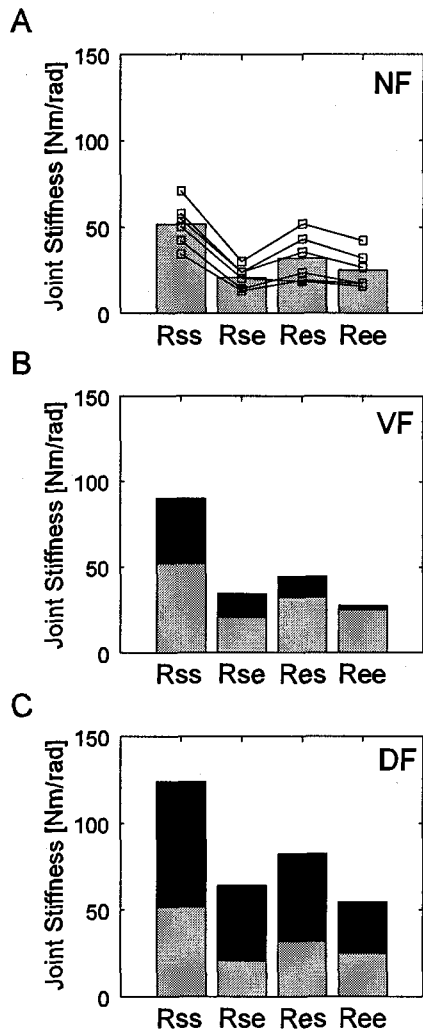


Figure 2.8 Changes in joint stiffness associated with the change in endpoint stiffness. **A:** Joint stiffness in the NF. The bar graph shows the mean stiffness for each term in the matrix, whereas the joined lines show the individual values for five subjects. **B:** Mean joint stiffness after adaptation to the VF (black bars) relative to the joint stiffness in the NF (grey bars) (same as in A). **C:** Mean joint stiffness after adaptation to the DF (black bars) relative to the joint stiffness in the NF (grey bars).

Stiffness measurements in the DF were performed with the force field turned off prior to the applied perturbation. In order to determine if this could affect the value of the stiffness measurement and to confirm that the stiffness adaptation is preprogrammed; the EMG after adaptation was compared to that during after effects.

Figure 2.9 compares the EMG during after effect trials of the DF with that during NF

trials. The EMG was much higher during DF trials compared to NF trials ($p < 0.001$ for all six muscles), illustrating the large increase in co-contraction in antagonist muscle pairs, and corresponding to the adapted impedance. Similarly, the after effect EMG was much higher than NF trials ($p < 0.001$ for all six muscles) although these movements were also performed in the NF. Furthermore, the EMG during after effect trials was very similar in magnitude and shape to DF trials. For three muscles there was no significant difference in the magnitude between after effect and DF trials: posterior deltoid ($p = 0.17$), pectoralis major ($p = 0.25$) and brachioradialis ($p = 0.96$); whereas for the other three muscles the DF EMG was slightly higher than the EMG during after effects: triceps long head ($p = 0.017$), biceps brachii ($p = 0.005$) and triceps lateral head ($p = 0.022$). Although the DF was off during the after effect trials, the subjects were not aware on which trials the DF was on or on which it was off, and assumed that it was always on. Thus, the EMG during the after effect trials reveals a preprogrammed motor command to compensate for the DF. The impedance was predictively and not reactively controlled. Similarly, this indicates that having the force field off prior to the stiffness measurement should not have affected the stiffness measurement. The small differences in EMG between DF and after effect trials, occurring in the later half of the movement, likely represent the neural feedback response to perturbations caused by the DF. Such perturbations would not have occurred during trials in which stiffness was being measured, regardless of whether the DF had been activated or inactivated prior to the interval of stiffness measurement. Because the differences in EMG were small compared to the overall magnitude of the EMG and because their onset tended to be later than the onset of

stiffness measurement, we can be quite certain that the muscles were in the same mechanical state on normal DF trials and on DF trials when stiffness was being measured.

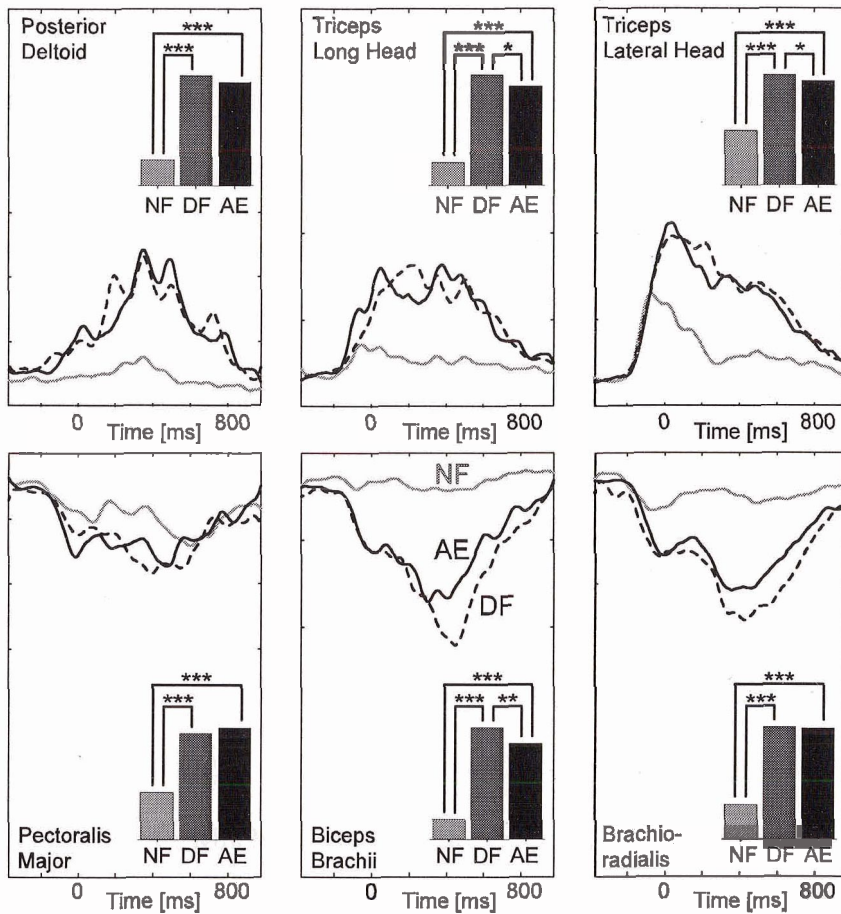


Figure 2.9 Rectified, averaged and smoothed surface EMG during the DF after effect trials for subject 1. Black dashed curves denote EMG in the DF-on trials of the after effect experiment (DF). Black solid curves denote EMG in the NF trials of the after effect experiment (AE). Gray curves denote EMG in NF trials, where the DF was never presented (NF). In each plot the mean integrated EMG activity over the entire movement for all subjects is shown with bar graphs. The asterisks indicate significant differences between the three conditions using Scheffe's post hoc comparison (* $p < 0.05$, ** $p < 0.01$ and *** $p < 0.001$).

DISCUSSION

This study compared changes in endpoint impedance after adaptation to stable and unstable environments in order to further understand the processes of learning and adaptation. The endpoint stiffness of the arm was modified during adaptation to both environments, but for different reasons. In the stable environment (VF), the change in stiffness was due to changes in the net joint torques needed to counteract the large environmental force. In contrast, the change in stiffness after adaptation to the unstable environment (DF) far exceeded what would have been predicted from the change in joint torques necessary to cancel the environmental force. Furthermore, the change in stiffness was uncorrelated with the change in joint torque. Instead, there was a directionally selective increase in stiffness which was likely brought about by biased co-contraction of biarticular muscles.

Stiffness in stable dynamics is side effect of change in joint torque

Previous studies which investigated the adaptation of arm movements to novel dynamics did not measure how endpoint stiffness changed after learning, although changes in the EMG of selected muscles were reported (Thoroughman and Shadmehr 1999). While these changes in muscle activity patterns likely corresponded to changes in endpoint stiffness the latter has not been quantified. This is the first quantitative investigation of the change in endpoint stiffness following learning of an internal model.

After adaptation to the VF, the stiffness ellipses became elongated along the direction of compensation for the force field. Such a change in endpoint stiffness with force direction has been reported previously for static postures (Gomi and Osu 1998; McIntyre et al. 1996; Perreault et al. 2001). The observed change in stiffness orientation is similar to that reported by (Gomi and Osu 1998) for the corresponding force direction. We found that the difference in the elastic force, resisting displacement of the hand, relative to the NF was well correlated with the estimated difference in force based on an assumed linear relation between the difference in joint stiffness and the difference in joint torque relative to the NF. Furthermore, the same linear relation generated an excellent prediction of the stiffness measured in the VF. This is strong evidence that the change in stiffness observed in the VF was simply a byproduct of the change in endpoint force required to adapt to the force field.

Stiffness in unstable dynamics is directly controlled

When subjects were faced with the unstable dynamics of the DF, they responded by selectively adapting their stiffness. This provided necessary stability to perform the task successfully. This impedance control was performed by co-contracting antagonist muscle pairs in order to increase stiffness without changing the net joint torque. Such impedance control was first postulated by Hogan (Hogan 1985). While EMG measurements during movements to a single point from various starting positions provided some limited evidence that impedance may vary depending on past history (Gribble and Ostry 1998), no conclusive evidence for the existence of impedance control had been found until recently (Burdet et al. 2001).

While reflexes could have contributed to the measured stiffness it is unlikely that they were responsible for the adaptation which permitted stable movement in the DF. Reflexes would only act following displacement of the hand from the planned trajectory and would occur too late to counteract the DF because its effect would be increasing throughout the delay period. Since the trajectories in the DF after learning were similar to those in the NF, with the hand moving smoothly to the final target, there is no evidence that the hand was sufficiently disturbed to evoke reflex responses.

Because measurements of stiffness are based on multiple trials they contain no information about trial-by-trial variations. Although unlikely, it is possible that the observed increase in stiffness was linked to the force which the subject exerted to counteract the force applied by the force field. This would vary from trial to trial due to variations in the trajectory. If this had been the case, then the difference in stiffness, relative to the NF, should have given a change in force which was correlated with the difference in measured endpoint force. When this was tested we found that the two were essentially independent. Similarly, stiffness ellipses generated from presumed linear relations between joint torque and joint stiffness failed to reproduce the measured endpoint stiffness in the DF. From this, it is clear that the change in endpoint stiffness in the DF with respect to the NF could not be attributed to compensation for the small forces exerted by the force field. Therefore, we can conclude that after adaptation to the unstable DF, the stiffness was controlled independently of the endpoint force and was selectively tuned to the instability of the environment.

Change in stiffness is a byproduct of development of an inverse model

As the CNS learns to adapt to the environment, the activity of the muscles is modified in order to counteract the environmental dynamics. This is presumed to occur during formation of an inverse model of the dynamics (Conditt et al. 1997; Flanagan et al. 2001). Because muscle stiffness increases linearly with muscle activation and force (Hunter and Kearney 1982; Kirsch et al. 1994), any change in the activation produced by the development of an internal model should produce a corresponding change in stiffness. We found that the change in limb impedance was well explained by the change in the joint torque that occurred during the learning in the VF. The only other study which has attempted to address changes in the impedance of the arm after adaptation to novel force fields is the recent work of Wang et al (2001). They proposed that an inverse model would not have any effect on the impedance of a limb and specifically would not change the intrinsic stiffness or the short latency reflex response to a perturbation. The restoring force due to large perturbations during movements was attributed to the long latency reflex and it was presumed that the gain of this reflex was changed due to a modification of the sensory-motor feedback control pathway (Wang et al. 2001). This conclusion is based upon the assumption that the use of an inverse model to adapt to the environment would not affect the intrinsic or reflexive stiffness of a muscle. However, modifying the feedforward command to muscles will automatically change both the intrinsic and reflexive components of the muscle impedance. As the descending drive to muscles (activation) is increased the muscles force response to a disturbance (impedance)

increases. This has been shown at both the muscle level (Rack and Westbury 1969) and joint level (Carter et al. 1993; Gottlieb and Agarwal 1988; Hunter and Kearney 1982; Weiss et al. 1988). Similarly, previous work has also shown that both the short and long latency reflexes are directly influenced by the level of muscle activation (Capaday et al. 1994; Dufresne et al. 1978; Marsden et al. 1976; Matthews 1986; Smeets and Erkelens 1991; Stein et al. 1995). These results indicate that any change in the activation level of a muscle due to a modification in the descending command, such as would be produced by an inverse model would affect both the intrinsic and reflexive stiffness of the muscles. This is also seen in the multi-joint arm where a change in the joint torque has a direct effect on the joint stiffness and, therefore, the endpoint stiffness (Gomi and Osu 1998; Perreault et al. 2001). A key issue is the question of whether the change in impedance following adaptation to novel dynamics is due to a change in intrinsic and reflexive stiffness which occurs as a direct effect of a change in motor command, or if it is due to the modification of the sensory-motor feedback pathway. The results of our analysis on adaptation to the VF clearly indicate that the change in impedance scales with the change in joint torque, exactly as in isometric multi-joint tasks (Gomi and Osu 1998; Perreault et al. 2001) where no change to the sensory-motor feedback pathways is required. This indicates that development of an inverse model can directly affect the limb impedance. This must be taken into account when using the force responses to perturbations to infer the properties of the motor control system.

III: ADAPTATION TO STABLE AND UNSTABLE DYNAMICS ACHIEVED BY COMBINED IMPEDANCE CONTROL AND INVERSE DYNAMICS MODEL

INTRODUCTION

Humans have exceptional abilities to move and interact with objects in the environment. When faced with novel tasks, they adapt to environmental disturbances in a way that indicates a fundamental knowledge of the mechanics of the external world (Conditt et al. 1997; Flanagan et al. 2001; Flanagan and Wing 1997; Krakauer et al. 1999; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999). Studies of individuals performing goal-directed movements in novel mechanical environments have shown that the CNS acquires internal models of the external world (Kawato 1999). However, most adaptation studies have employed paradigms that involve stable interactions with the environment whereas many tasks that humans perform, particularly those involving tool use, are inherently unstable (Rancourt and Hogan 2001).

Adaptation to perturbations that do not compromise mechanical stability appears to involve the acquisition of an inverse dynamics model through feedback error learning (Kawato et al. 1987). However, conventional feedback error learning does not address the issue of modifying mechanical impedance to counteract mechanical instability, although this has been repeatedly observed (Akazawa et al.

1983; Burdet et al. 2001; De Serres and Milner 1991; Milner 2002a; Milner and Cloutier 1993; Milner and Cloutier 1998; Milner et al. 1995).

When the dynamics do not induce mechanical instability, changes in muscle activation patterns closely follow adaptive changes in joint torques, although there is excess activation, particularly in the early stages of learning (Thoroughman and Shadmehr 1999). Agonist-antagonist muscle co-contraction was found to increase upon exposure to the novel dynamics and then decrease as learning progressed. The ability to co-contrast specific groups of muscles would permit selective changes to the geometry of the endpoint stiffness (Hogan 1985). Recently, we were able to demonstrate that the magnitude, shape and orientation of the endpoint stiffness of the arm can be controlled in a predictive way to compensate for environmental instabilities (Burdet et al. 2001).

Previous research has suggested the existence of two separate motor control mechanisms: inverse dynamics models and impedance control. An inverse dynamics model is a controller which computes feedforward commands of the net joint torques for movement, based on the estimated effects of internal and external dynamics. An impedance controller, in contrast, modifies the impedance of the limb by co-contraction of agonist and antagonist muscles without changing net joint torque. These two controllers could operate either independently or together during adaptation. While there is no direct evidence for (or against) separate brain mechanisms being implicated in these two types of control, they perform different functions and, therefore, can be considered as separate controllers. Different features of our visual experience, for example, colour, contour orientation, motion, and retinal

disparities, are not independently encoded at the sensory level, but are represented separately in the visual cortex (Ts'o and Roe 1994). Functionally, these features of the visual world provide very different information about our environment and so the brain processes each separately. It seems possible, therefore, that functionally distinct features of motor output might be represented and processed in a similarly independent fashion by the motor system. Some evidence already exists that control of reciprocal activation and co-contraction occurs in separate areas of the cortex (Humphrey and Reed 1983). Recently, both Takahashi et al. (Takahashi et al. 2001) and Osu et al. (Osu et al. 2002) have provided evidence that inverse dynamics models and impedance control operate as separate mechanisms for motor control.

The present study investigated adaptation to novel force fields, in which trajectories were initially stable or unstable, to compare features of inverse dynamics model formation and impedance control. Adaptation to the force field, in which hand trajectories were stable, could be achieved by simply modifying joint torque. An inverse dynamics model alone was sufficient to compensate for the perturbing effects of this force field (Franklin et al. 2003a). Adaptation to the other force field, in which hand trajectories were unstable, required an increase in the endpoint impedance of the arm, but no change in net joint torque, i.e., only impedance control (Burdet et al. 2001). However, we hypothesized that both of these processes are generally active during learning. To test this hypothesis, we analyzed the time course of changes in muscle activation patterns, as well as joint torques, during the learning of novel force fields. In particular, we were interested in comparing the extent to which the CNS used generalized muscle co-contraction in the early stages of learning to increase

endpoint stiffness and how patterns of muscle activity were later refined. Two parallel processes were identified from the evolution profile of the muscle activation patterns. One was an activation process involved in increasing the endpoint stiffness of the arm by means of muscle co-contraction during the early stages of learning. The other was a deactivation process, which led to gradual reduction in muscle activity as learning progressed.

METHODS

Five healthy individuals participated in the entire study (20-34 years of age; 1 female and 4 males; all right-handed). The institutional ethics committee approved the experiments and the subjects gave informed consent prior to participation.

Apparatus

Subjects sat in a chair and moved the parallel-link direct drive air-magnet floating manipulandum (PFM) (Fig. 3.1) in a series of forward reaching movements performed in the horizontal plane. Their shoulders were held against the back of the chair by means of a shoulder harness. The right forearm was securely coupled to the PFM using a rigid custom moulded thermoplastic cuff. The cuff immobilized the wrist joint, permitting movement of only the shoulder and elbow joints. The subjects' right forearm rested on a support beam projecting from the handle of the PFM. Motion was, therefore, limited to a single degree of freedom at the shoulder and at the

elbow. The manipulandum and setup are described in detail elsewhere (Gomi and Kawato 1996; Gomi and Kawato 1997).

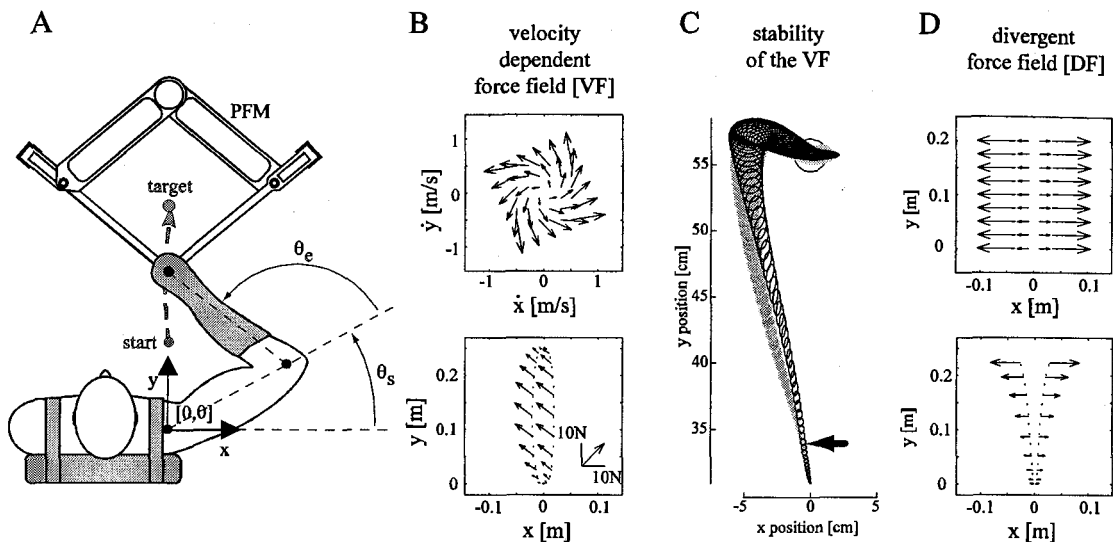


Figure 3.1 A: Experimental setup to study the adaptation to stable and unstable dynamics. Subjects were seated in a chair with their shoulders restrained by a harness, and their hand and forearm firmly attached to the parallel-link direct drive air-magnet floating manipulandum (PFM) with a thermoplastic splint. Reaching movements were performed from a start point ($[x, y] = [0, 0.31]$ m) to a target located at ($[0, 0.56]$ m) relative to the subjects shoulder ($[0, 0]$) for a total movement length of 0.25 m. The conventions for shoulder angle (θ_s) and elbow angle (θ_e) are shown. B: Force in the velocity dependent force field (VF) plotted as a function of hand velocity (top). The vectors indicate the direction and magnitude of the forces. In the VF, the force vector rotates and increases in magnitude with velocity. Force vectors when hand trajectories (dotted curves) are slightly deviated from a straight line along the y -axis to the left and right (bottom). In the VF, the applied force vectors are similar despite trajectory variations. C: Stability of the VF environment. Before effect trials (random VF trials in a majority of NF trials) with (grey filled circles) and without (black circles) a force perturbation (15 N for 25 ms) were recorded. The onset of the force perturbation (100 ms after start) is shown by the black arrow. Both series of trajectories converge to the same path demonstrating the stability of the interaction between the arm and the VF. D: Force in the divergent force field (DF) plotted as a function of the hand position (top). The force vector increases in magnitude with the distance from the y -axis. Bottom: Force vectors when hand trajectories (dotted curves) are slightly deviated from a straight line along the y -axis to the left and right. The DF amplifies trajectory variations.

Force Fields

The experiment examined trajectory and EMG adaptation in two force fields (Fig 3.1 B,D): a velocity dependent force field (VF), in which hand trajectories were stable from the outset and a position dependent (divergent) force field (DF), in which hand trajectories were initially unstable. Results were compared to those in a null field (NF). Details of the implementation and protocol have been described elsewhere (Burdet et al. 2001). Only a general overview is given here. The force (F_x, F_y) (in N) exerted on the hand by the robotic interface in the VF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \chi \begin{bmatrix} 13 & -18 \\ 18 & 13 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (2.1)$$

where (\dot{x}, \dot{y}) is the hand velocity (m/s) and the scaling factor, χ , was adjusted to the subject's strength ($2/3 \leq \chi \leq 1$). The stability of the initial trajectories was tested by applying small perturbations during before effect trials and comparing the resulting trajectories (Fig. 3.1C). While subjects were performing movements in the NF the force field was switched to the VF on random trials (before effects). On some of these trials a brief triangular force pulse (25 ms duration) with amplitude of 15 N was applied to the hand 100 ms after movement onset. Both the perturbed and unperturbed before effect trajectories are perturbed by the force field. However, all trajectories converge illustrating that they are stable. This result is consistent with the defining characteristic of Lyapunov stability that the addition of a small perturbation does not produce divergent behaviour.

The DF produced a negative elastic force perpendicular to the target direction with a value of zero along the y -axis, i.e., no force was exerted when trajectories followed the y -axis, but the hand was pushed away whenever it deviated from the y -axis. The DF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \quad (2.2)$$

where the x -component of the hand position was measured relative to the shoulder joint. $\beta = (300 \text{ to } 500) \text{ (N/m)}$ was adjusted for each subject so that it was larger than the stiffness of the arm measured in NF movements so as to produce instability. For safety reasons, the DF force field was inactivated if the subjects' trajectory deviated more than 3 cm from the y axis. Both force fields were inactivated once the subject reached the target position.

Learning

All subjects practiced in the NF on at least one day prior to the experiment. These training trials were used to accustom the subjects to the equipment and to the movement speed and accuracy requirements. Subjects were randomly assigned to one of two groups. Group 1 adapted to the DF on one day and adapted to the VF on another day, whereas group 2 adapted to the fields in reverse order.

Subjects first practiced in the NF until they had achieved 50 successful trials. Successful trials were those which ended inside a 2.5 cm diameter target window

within the prescribed time (0.6 ± 0.1 s). All movements were recorded whether successful or not. The movement distance was 0.25 meters. Movements were self-paced so subjects were able to rest between movements if they wished. At the completion of 50 successful trials, the force field was activated. No information was given to the subjects as to when the force field trials would begin. Subjects then practiced in the force field until achieving 75 successful trials. They took a short break and then performed 100 more movements, 20 of which were random trials in the NF. The NF trials were called after effects and were recorded to confirm that subjects had adapted to the force field.

Hand path errors

The adaptation to the force fields was quantified by calculating the error relative to a straight line joining the centers of the start and target circles. The absolute hand path error

$$S(|e_x|) = \int_{t_0}^f |x(t)| |\dot{y}(t)| dt \quad (2.3)$$

represents the area between the actual movement path and the straight line. The signed hand path error, defined as

$$S(e_x) = \int_{t_0}^f x(t) |\dot{y}(t)| dt \quad (2.4)$$

is a measure of the mean directional extent by which the path deviates from the straight line. Hand path errors were calculated from the start time, t_o (75 ms prior to crossing a hand velocity threshold of 0.05 m/s), to the termination time, t_f (when curvature exceeded 0.07 mm^{-1}) (Pollick and Ishimura 1996).

The handpath error for each subject was fit with an exponential curve using a least square error method. This fitted error was expressed as:

$$S(t) = Ae^{-\frac{t}{\tau}} + C \quad (2.5)$$

where A is the gain of the exponential process, τ is the time constant, C is the constant error and t refers to the trial number.

Torque estimation

Time varying muscle torque at the shoulder and elbow was computed using the equations of motion for a two-link planar arm (c.f. (Hollerbach and Flash 1982)). However, the equations used also include the contributions to joint torque from external forces applied at the hand. The joint torque was calculated as:

$$\begin{aligned} \tau_s &= \ddot{\theta}_s (2X \cos(\theta_e) + Y + Z) + \ddot{\theta}_e (X \cos(\theta_e) + Y) - \dot{\theta}_e^2 X \sin(\theta_e) \\ &\quad - 2\dot{\theta}_s \dot{\theta}_e X \sin(\theta_e) - (l_1 \sin(\theta_s) + l_2 \sin(\theta_s + \theta_e)) F_x \\ &\quad + (l_1 \cos(\theta_s) + l_2 \cos(\theta_s + \theta_e)) F_y \\ \tau_e &= \ddot{\theta}_e Y + \ddot{\theta}_s (X \cos(\theta_e) + Y) + \dot{\theta}_s^2 X \sin(\theta_e) \\ &\quad - l_2 \sin(\theta_s + \theta_e) F_x + l_2 \cos(\theta_s + \theta_e) F_y \end{aligned} \quad (2.6)$$

where:

$$\begin{aligned} X &= m_2 l_1 c_{m2} + m_c l_c c_{mc} \\ Y &= I_2 + m_2 c_{m2}^2 + I_c + m_c c_{mc}^2 \\ Z &= I_1 + m_1 c_{m1}^2 + (m_2 + m_c) l_1^2 \end{aligned}$$

τ is joint torque, θ is joint angle (defined as in Fig. 3.1), I is moment of inertia about the center of mass (c_m) of the segment, l is segment length and m is segment mass.

The subscript 1 refers to the upper arm, 2 to the forearm, s to the shoulder, e to the elbow and c to the wrist cuff. The mass and inertia of the subject's arm segments were estimated from the weight and segment lengths of each subject based on anthropometric scaling relations (Winters 1990).

Joint torque during learning was further quantified by estimating the variation in torque on a particular trial relative to the joint torque after adaptation to the force field. The learned joint torque profile was calculated using the mean of the last twenty successful trials during the learning. The absolute torque error ($|\tau_{error}|$) represents the total difference between the joint torque on a particular trial and the mean final joint torque both in terms of extent and timing. It is calculated as:

$$|\tau_{i,error}^j| = \int_{t_0}^{t_f} |\tau_i^j(t) - \tau_m^j(t)| dt \quad (2.7)$$

where τ_i is the joint torque on a given trial, τ_m is the mean joint torque after adaptation, the superscript j refers to either the shoulder or elbow joint, t_0 is the time

of movement onset, and t_f is the end of the time of interest. The torque error was estimated over 1000 ms. The absolute torque error was fit with an exponential curve using least squares similar to the handpath errors (Equation (2.5)).

In order to examine the development of the torque during the first few trials in the DF we used two measures of torque change relative to the torque in the NF. The first measure, absolute torque development, was calculated as in equation (2.7) but with τ_m referring to the mean joint torque in the NF. A second measure termed signed torque development ($\tau_{develop}$) was calculated as:

$$\tau_{i,develop}^j = \int_{t=t_0}^{t_f} (\tau_i^j(t) - \tau_m^j(t)) dt \quad (2.8)$$

where τ_m again refers to the mean joint torque in the NF. Both the absolute and signed torque development were calculated over the initial 400 ms to avoid the influence of corrective movements or the safety boundary.

Electromyography

Surface electromyographic (EMG) activity of six arm muscles was recorded using pairs of silver-silver chloride surface electrodes during the learning sessions. The electrode locations were chosen to maximize the signal from a particular muscle while avoiding cross-talk from other muscles. The skin was cleansed with alcohol and prepared by rubbing in electrode paste. This was removed with a dry cloth and pre-gelled electrodes were then attached to the skin with tape. The spacing between the

electrodes of each pair was approximately 2 cm. The impedance of each electrode pair was tested to ensure that it was below 10 k Ω .

The activity of two monoarticular shoulder muscles, pectoralis major and posterior deltoid, two biarticular muscles, biceps brachii and long head of the triceps, and two monoarticular elbow muscles, brachioradialis and lateral head of the triceps, was recorded. The EMG signals were analog filtered at 25 Hz (high pass) and 1.0 kHz (low pass) using a Nihon Kohden amplifier (MME-3132) and then sampled at 2.0 kHz. All comparisons between force field EMG and NF EMG involved data recorded on the same day without removal of the electrodes. EMG was aligned on the movement onset and averaged over twenty trials to visually compare changes that had occurred during learning.

To quantify changes during learning, the rms (root-mean-square) value of the EMG was calculated from 100 ms before movement onset until 350 ms after movement onset to include all the early movement associated activity. The muscle activity during this time will include both the feedforward command and the reflex responses. In all cases, the recorded EMG data were characterized by an initial increase in activity followed by a more gradual decrease when examined over all of the trials. Both processes appeared to be exponential. In order to compare the time course of the changes in EMG with the time course of other kinematic and dynamic parameters, we fit these data with a model. Our model expressed EMG as a function of trial number (t) using a double exponential with four free parameters:

$$EMG(t) = A \left(1 - e^{-\frac{t}{\tau_1}} \right) - B \left(1 - e^{-\frac{t}{\tau_2}} \right) + C \quad (2.9)$$

A and B represent the gains of the two exponential processes, τ_1 and τ_2 represent the time constants, C is EMG at time zero (fixed as the mean of the EMG in the NF) and t is the trial number. One process (the first term) represents an exponential increase in EMG in response to exposure to the force field. It is bounded by the maximal level of activation that is possible ($A+C$), which has been shown to be significantly lower in co-contraction than reciprocal activation (Milner et al. 1995). The other process (the second term) represents an exponential decrease in EMG. For comparisons across subjects, the EMG was normalized such that the mean NF EMG (C) was equal to a value of one. This model corresponds well with the recent work of Osu et al. (2002) which describes EMG activity during motor learning as a gradually decreasing function with occasional increases or decreases that are related to the error in previous trials. To provide further support for the choice of this model, Akaike's Information Criterion (AIC) (Akaike 1974) was calculated for this and two simpler models (Appendix A).

The time course of the EMG was compared to the time course of the handpath error and absolute joint torque error. The handpath error was calculated by one of two methods. In the DF, the absolute handpath error $S(|e_x|)$, a measure of the area between the actual movement path and the straight line joining the start and end targets, was used. In the VF, the signed handpath error $S(e_x)$, a measure of the mean directional extent by which the path deviates from the straight line, was used.

Although different representations of the error were used for the DF and VF, we could have used absolute handpath error for both DF and VF without changing any of the results or conclusions. The handpath error and joint error for each subject were fit with an exponential curve using a least square error method. The time constants of these changes were then compared to the time constants of the change in EMG during the adaptation process using t-tests with a significance level of 0.05. These statistics were performed on the reciprocal of the time constants, i.e., the rates, to reduce the variance of processes that exhibited little change over time (long time constants). All time constants were significantly different from zero using a t-test at the 0.05 level ($p < 0.00001$).

Timing of Reflex Responses

In order to find an appropriate interval for EMG analysis, the onset of voluntary responses representing feedback correction in the VF and DF was examined using before effects. Movements were performed in the NF. On randomly selected trials this was changed to the force field to elicit reflex and voluntary reactions to the imposed force field. A total of 80 NF and 20 force field trials were recorded for each force field. A comparison of the before effect EMG to that of the NF trials allowed us to determine the onset of the corrective responses (Figure 3.2). These generally occurred more than 200 ms after the onset of movement except in the case of the posterior deltoid muscle for which the responses were faster. In this case, a small early reflex response could be seen as early as 150ms after the onset of movement.

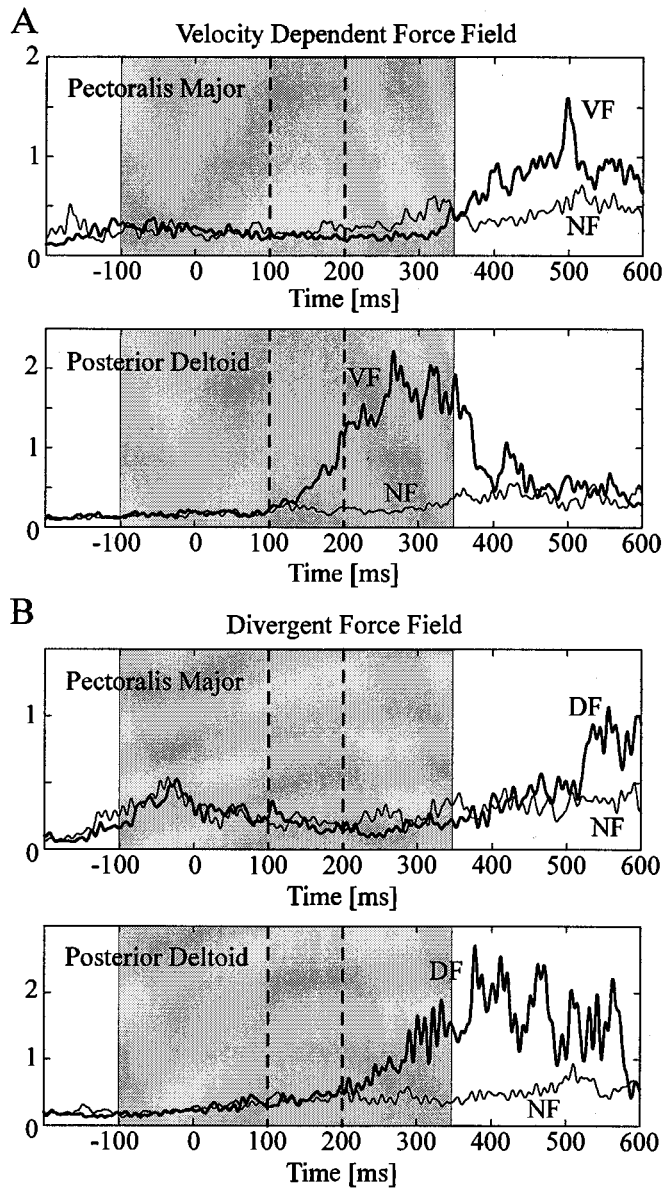


Figure 3.2 Timing of corrective (reflexive and voluntary) EMG in the VF (A) and DF (B). On random trials in a majority of NF movements, either the VF or DF (in separate experiments) was applied to elicit the initial corrective response to the field without learning (before effects). The timing of the corrective responses could then be seen by comparing the force field EMG (thick line) to the NF EMG (thin line). Each line represents the mean of ten trials which have been rectified and low pass (zero phase) filtered with a 5th order 100 Hz cutoff Butterworth filter. Traces are shown for the muscle with the fastest response (posterior deltoid) and its antagonist muscle (pectoralis major). Except for the posterior deltoid muscle in the VF, no corrective EMG was seen prior to 200 ms, and in most muscles no activity was seen prior to 350 ms.

RESULTS

The two force fields investigated in this study produced distinctive perturbations of the trajectories prior to adaptation. Initial trials in the VF consistently perturbed the trajectories to the left (Fig. 3.3A), but subjects quickly learned to compensate for the force field and soon began to make straighter movements. By the 25th trial, the trajectories were relatively straight and consistently reached the final target position. The trajectories of later trials were similar. The mean signed hand-path error illustrates the large reduction in trajectory error over the first ten movements (Fig. 3.3B). The signed hand-path error was reduced almost to zero, with relatively little variability, by about trial 30. The absolute handpath error changed in a similar fashion. Movements in the DF were at first perturbed either to the right or the left (Fig. 3.3C), depending on the initial deviation in the path. However, again subjects were able to adapt to the force field. Following the 25th trial, subjects were able to successfully complete the task on most trials, exhibiting straight trajectories to the final target. After learning, trajectories in both the VF and DF were similar to NF trajectories. The mean signed error across subjects and its standard deviation achieved their minimum value around trial 30 in the DF (Fig. 3.3D), similar to the VF. The trajectory of the initial trials in the DF was not perturbed as much as in the VF because of the safety zone. Nevertheless, the mean absolute error decreased only gradually over the first 75 trials. In both the VF and the DF, the signed error tended to be negative, as overall, subjects made movements which were slightly biased to the left. This is consistent with NF movements which were also slightly biased to the left.

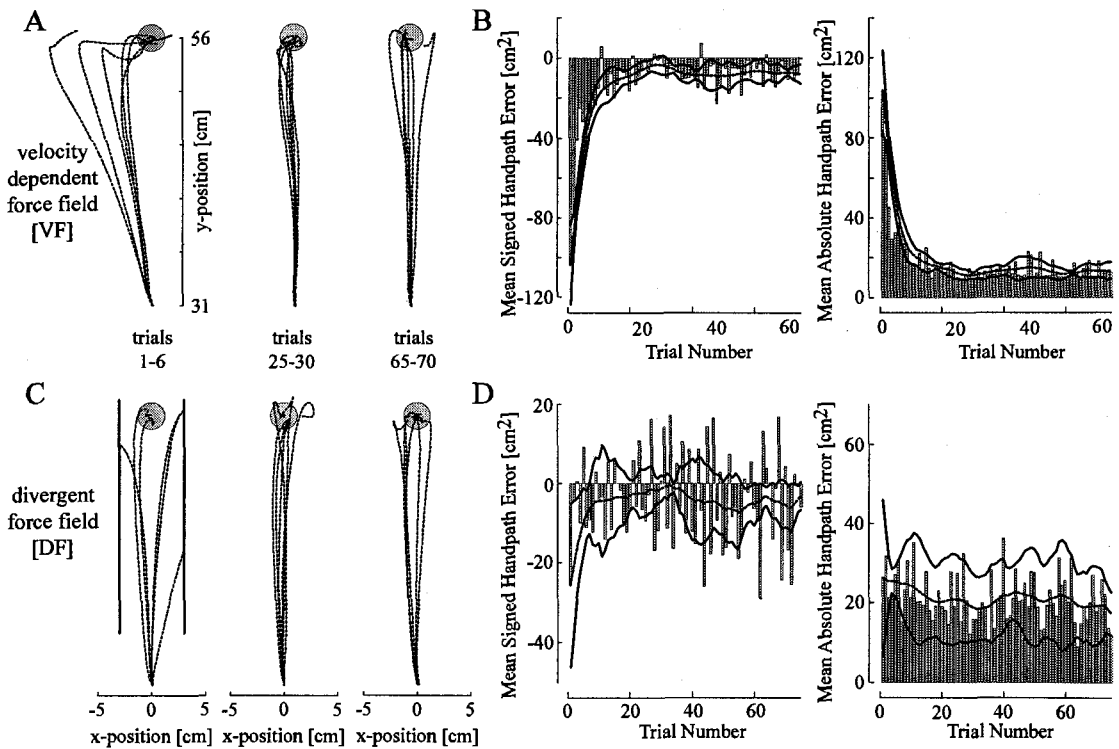
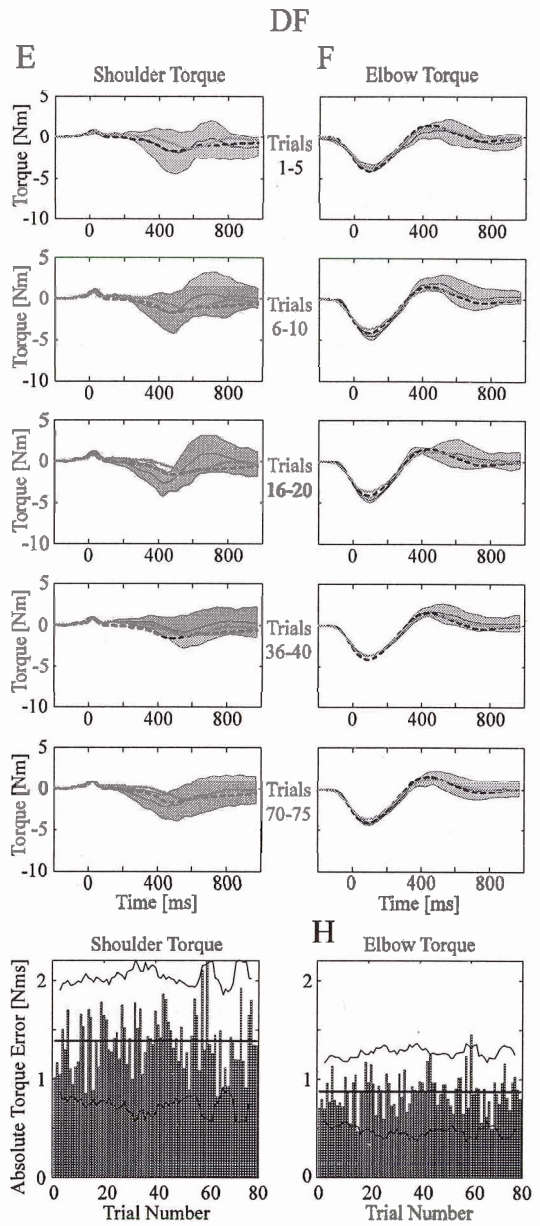
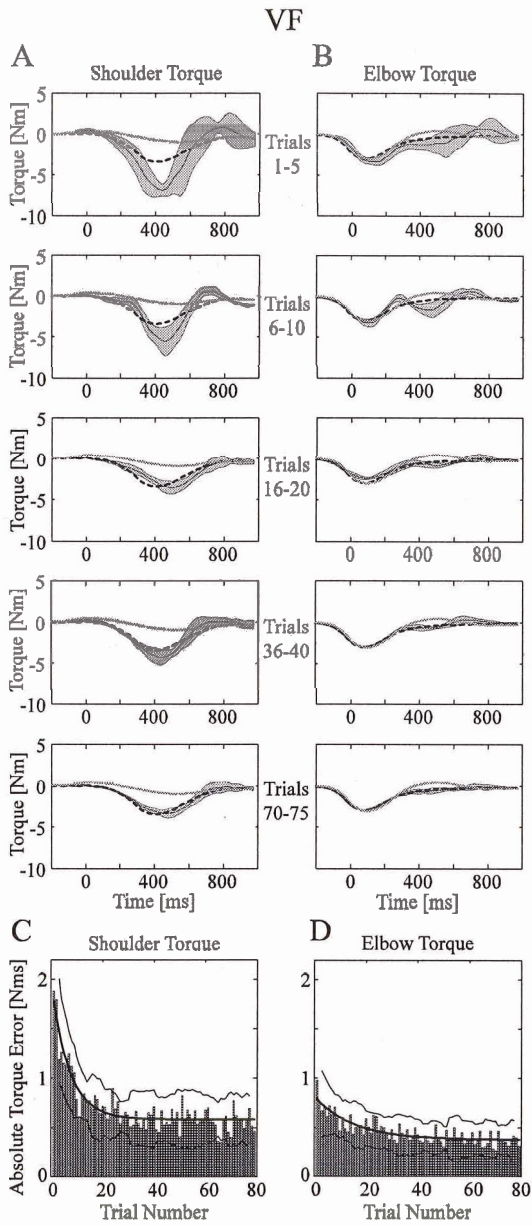


Figure 3.3. Movements in the VF (A) and DF (C). Movements are shown for the initial movements in the force field (trials 1-6), the early portion of learning (trials 25-30) and the late portion of learning (trials 65-70). The black lines on either side of trials 1-6 for the DF indicate the safety zone, outside of which the field was turned off for safety reasons. B,D: The bar graphs represent signed hand-path error (left) and absolute handpath error (right) averaged over all subjects during the first 75 trials of learning in the VF (B) and DF (D). The middle dotted line indicates the mean values smoothed with an 8 point moving average. The standard deviations about this mean are also shown using solid lines (5 point moving average).

During learning, subjects adapted to the novel forces applied by the VF by modifying the joint torques (Fig. 3.4 A-D). After adaptation, the shoulder joint torque became an extensor torque throughout the entire movement and increased to several times the value in the NF trials. At the elbow, only a small change in torque in the extensor direction, during the second half of the movement, was required. In order to

adapt to the VF, subjects gradually changed both the amplitude and shape of their joint torque profile until they had compensated for the force field's effects. Early in learning, the torque varied both above and below the final adaptation profile. In particular, early in the trial the torque tended to overshoot the final torque profile whereas later in the movement the opposite occurred. The early effect is due to the force applied by the force field and the corresponding changes in torque produced by changes in hand trajectory. The later effect is likely the result of reflexive feedback and voluntary correction in response to the disturbance. As subjects gained more experience performing movements in this force field, the amplitude and shape of the torque profile were adjusted until trajectories were no longer disturbed by the force field. The variability in the joint torque was large among the early trials (shaded grey areas, Fig. 3.4A,B). As subjects modified the torque profile they also reduced the trial to trial variability in the joint torque. The absolute torque error of all subjects decreased quickly during learning with most of the adaptation occurring in the first 30-40 trials (Fig. 3.4 C,D). However, subjects continued to decrease joint torque errors by small amounts as their performance became more skilled and less variable. While subjects were able to reduce their hand-path trajectory errors close to the final level by approximately the 20th trial during learning, changes in the joint torque profiles indicate that further updating and refining of the feedforward commands was still taking place throughout the learning process.

Figure 3.4. Change in shoulder and elbow torque during learning in the VF (left) and DF (right). VF (left). A: Mean shoulder torque during 20 NF movements (grey dotted line), during 20 movements after adaptation to the VF (black dotted line) and during the adaptation process (5 trials) (solid grey line) along with standard deviations (light grey area). Individual plots are shown for five separate sets of trials during the learning process. Data are shown for one subject. B: Mean elbow torque during learning in the VF. Data are plotted as in A. C: Absolute shoulder torque error averaged over all subjects during learning. Absolute shoulder torque is the summation of the difference between the torque on a given trial and the mean of twenty successful trials after full adaptation to the force field. D: Absolute elbow torque error during learning averaged over all subjects. DF (right). E: Mean shoulder torque during 20 NF movements (grey dotted line), during 20 movements after adaptation to the DF (black dotted line) and during the adaptation process (5 trials) (solid grey line) along with standard deviations (light grey area). Individual plots are shown for five separate sets of trials during the learning process. F: Mean elbow torque during learning in the DF. Data are plotted as in E. G: Absolute shoulder torque error averaged over all subjects during learning. H: Absolute elbow torque error during learning averaged over all subjects.



Early trials in the DF were characterized by unstable trajectories that varied either to the left or the right of the straight line joining the initial and final targets. Adaptation to this force field did not require a change in the net joint torques. The joint torque after adaptation to the force field was similar to that in the NF (Fig. 3.4 E,F). The moving average of the joint torque (over five consecutive trials) during the learning process was also quite similar to that in the NF, although the standard deviation remained high throughout the learning process. The computed joint torque took into account the measured hand force so the effect of any force applied to the hand by the DF due to small fluctuations in the handpath was included. Unlike the adaptation to the VF, the absolute torque error remained fairly constant (Fig. 3.4 G,H).

When subjects were initially presented with the DF, their trajectories deviated to the left or the right of the target and they rarely completed the movement successfully to the target. After learning, however, subjects were able to produce successful movements to the target. Movement along the target trajectory in the DF did not require any change in the joint torque compared to the NF. However, during the first few movements the joint torque varied greatly from trial to trial (Fig. 3.5). In particular, it tended to vary from an extensor torque to a flexor torque and back again. This alternating pattern, which can also be seen in the signed torque development plots (Fig. 3.5 C,D), did not necessarily occur on successive trials but was consistently seen on some time scale for every subject. The mean absolute torque development across all subjects increased progressively for the first six trials (Fig. 3.5 E,F). This occurred both for shoulder and elbow torque. An ANOVA was performed

comparing the first 3 trials (trials 1-3) to the next three trials (trials 4-6) with subjects as a random variable. The second three trials were found to be significantly larger at the 0.05 level than the first three trials for both the shoulder torque ($p=0.033$) and the elbow torque (0.001). This indicates that subjects initially responded to the disturbing effects of the force field by trying to modify the net joint torques to reduce the error. However, changing the net joint torques modifies the hand-path trajectory, which generally causes larger perturbing forces to be applied by the DF.

Adaptation in the VF required a modification of the net joint torques at the shoulder and elbow. Specifically, adaptation produced a change in net extensor joint torque at the shoulder throughout the movement and a small extensor moment at the elbow late in the movement. Therefore, we would expect to see increased activity in muscles contributing to these joint torques, particularly the posterior deltoid and long head of the triceps. After adaptation to the DF, the joint torques were not different from those in the NF. Instead, as we have shown previously, subjects modified the endpoint impedance of the limb (Burdet et al. 2001). We, therefore, expected to see increased activity in one or more muscle pairs contributing to increased co-contraction. The expected changes in muscle activation patterns were confirmed from the EMG after learning (Fig. 3.6). In the VF, the EMG increased predominantly in the posterior deltoid and long head of the triceps muscles, the two muscles contributing the extensor torque at the shoulder needed to compensate for the force field. In contrast, in the DF the muscle activity increased in both muscles of all antagonist pairs. Since the endpoint forces and joint torques were the same in the DF and the NF, this increase in muscle activation represented balanced co-contraction.

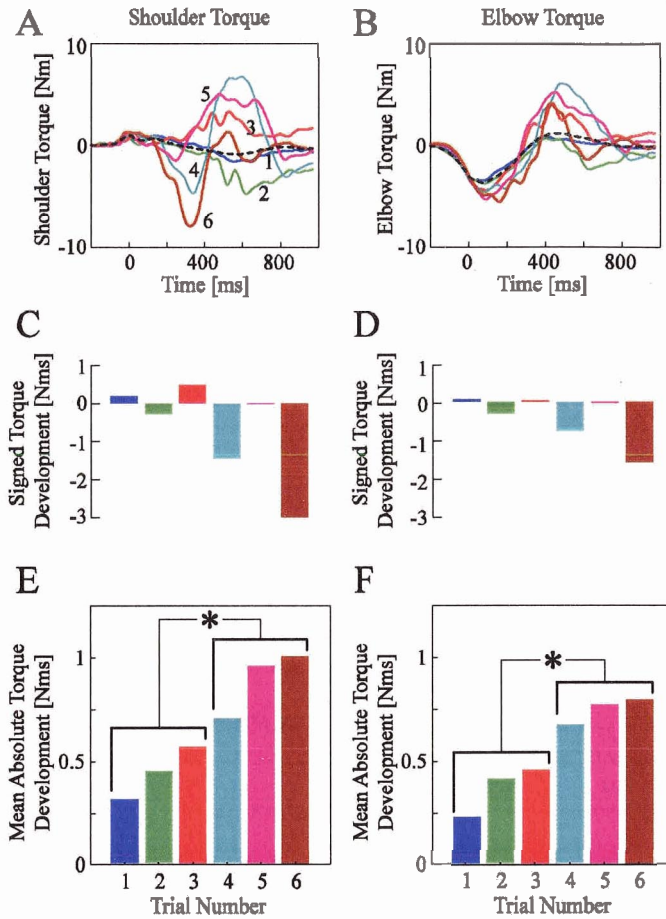


Figure 3.5 Early changes in shoulder and elbow torque during learning in the DF. A,B: Estimated shoulder torque (A) and elbow torque (B) during the first six trials of learning in the DF for one subject (coloured lines). The level of joint torque during movements in the NF is shown with a black dotted line (mean of 20 successful trials). C,D: The signed torque error for the trials shown in A and B, respectively. Signed torque development is the sum of the difference between the torque on a given trial and the mean torque in the NF over the first 400 ms after movement onset. E,F: The absolute torque development for the first six trials averaged across all subjects. Absolute torque development is the sum of the absolute difference between the torque on a given trial and the mean torque in the NF (first 400 ms after movement onset).

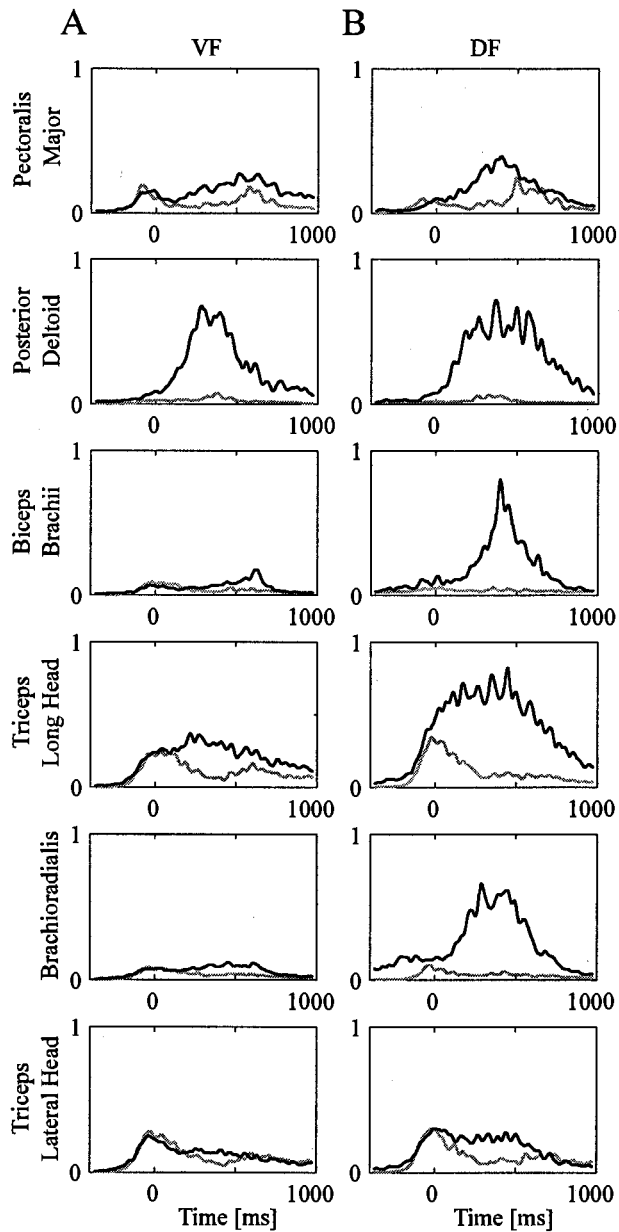


Figure 3.6 Muscle activity after adaptation to the VF (A) and DF (B) shown for one subject. The EMG activity of six arm muscles is shown for the NF (grey), compared to either the VF (A) or DF (B) (solid black). The EMG was smoothed using a 75-point (37 ms) smoothing routine and averaged over twenty successful trials. The EMG is expressed in arbitrary units where the mean NF EMG level for each muscle is equal in A and B.

After subjects experienced the VF for the first time, the muscle activity increased dramatically over the next few trials particularly for the posterior deltoid and long head of the triceps. This initial increase in EMG activity gradually diminished as learning proceeded. The change in EMG activity as learning progressed was fit with the double exponential function of Equation (3.9), which represents EMG as a function of trial number. In the following sections we present detailed analysis of the EMG from the interval of -100 to 350 ms from the start of movement. However, other intervals were also examined ([-100 to 200 ms], [250-350 ms], [-100 to 600 ms]) and all gave similar results. To illustrate the effect of the interval chosen for EMG analysis, the fitted curves for EMG activity using eight different time intervals were calculated for the long head of the triceps (Fig. 3.7). Although the rms EMG varies depending on the interval chosen, the overall trend for the way in which the activity changed on successive trials and the rate constants for the change in activity varied little. From this figure we can see that the choice of interval is not critical. Our analysis using the interval -100 to 350 ms included both pre-programmed feedforward motor commands and reflex responses to the perturbations. However, based on before effect EMG (Fig. 3.2), we determined that the earliest response occurred at about 150 ms after movement start for the long head of the triceps and posterior deltoid, but for most muscles, no significant change in activity was seen prior to 350 ms. Analysis over this shorter interval (-100 to 200 ms) for the muscles that exhibited early responses gave similar results to the longer interval.

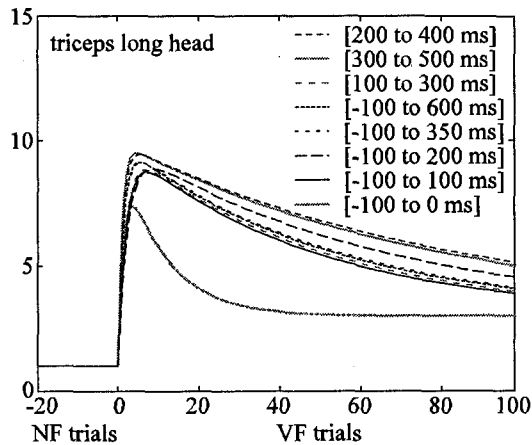


Figure 3.7 Model fits to activity of the triceps long head during learning in the VF for various intervals. Fits for 8 different time intervals are shown (Equation 3.9) to illustrate that the parameters of the model are relatively insensitive to the time interval chosen. The parameters A, B and the NF level have been normalized across all intervals so that the time constants can be examined.

The evolution of muscle activity with practice in the VF is shown for one subject in Fig. 3.8. There was an increase in the EMG for all six muscles during the first few trials. A t-test was performed comparing the rms EMG of trials 2-21 to 20 trials in the NF for each muscle for each subject. Twenty-seven of the thirty comparisons showed significant increases in the rms EMG at the 0.01 significance level indicating that the activity of all muscles increased early in learning. The peak EMG activity appeared to occur between the 3rd and 10th trial. With further practice, the EMG gradually decreased in all muscles to an asymptotic level. Similar to Fig. 3.6, the final rms EMG after adaptation was approximately the same as that of the NF for all muscles except those producing extensor torque at the shoulder. These muscles, the posterior deltoid and the long head of the triceps, compensated for the force applied by the VF. The actual difference after learning was tested using the same technique as for early learning. A t-test was used to compare the rms EMG of

the final 20 trials in the VF to twenty trials in the NF for each muscle for each subject with the significance level set at 0.01. For the posterior deltoid and long head of the triceps, all comparisons (10/10) were significantly larger after learning compared to the NF. In the case of the lateral triceps, four out of five comparisons were significantly larger after learning compared to in the NF. However, in the other three muscles, only eight out of the fifteen comparisons were significantly larger in the force field.

The variation in EMG activity over trials was best fit by a double exponential process (equation(2.9)). The results for all subjects are presented in Table 3.1. The activation and deactivation rates for each muscle were compared to the rate of reduction of the signed handpath error. In the VF, subjects quickly adapted to the disturbing effects of the force field and reduced their handpath error by two orders of magnitude within about ten trials. Activation either led or paralleled handpath error reduction and was characterized by a relatively short time constant, τ_1 . The deactivation time constant, τ_2 , was much longer. For three muscles (posterior deltoid, long head of the triceps and brachioradialis), activation occurred more rapidly than handpath error reduction ($p=0.012$). Both the posterior deltoid and long head of the triceps contribute to the necessary increase in shoulder extensor torque. The activity of the brachioradialis acts to counteract the extensor torque at the elbow resulting from the increase in activity of the long head of triceps. The activity of the remaining muscles increased more slowly, at a rate not significantly different from the rate of handpath error reduction ($p=0.45$). The activation of both sets of muscle groups was also faster than the reduction in absolute torque error at either the shoulder ($p=0.008$;

$p=0.026$) or elbow joint ($p=0.008$; $p=0.016$). Deactivation occurred at a similar rate in all muscles and proceeded much more slowly than either activation ($p=0.007$), handpath error reduction ($p<0.0001$), reduction in absolute shoulder torque error ($p=0.012$) or reduction in absolute elbow torque error ($p<0.0001$). The fact that the reduction in handpath error occurred faster than muscle deactivation, indicates that subjects only gradually reduced their reliance on increased impedance to resist the force field even after handpath error had been reduced to NF values. The largest increase in the activation process (A in eqn. 3.9) was seen for the posterior deltoid, followed by the long head of the triceps. However, all muscles showed a significant increase in muscle activation. The final level of adaptation, expressed as an increase from NF values ($A-B$ in eqn. 3.9), indicates that the activation of most muscles, posterior deltoid and long head of the triceps being the exceptions, was reduced to near NF levels. The rates of the decrease in absolute shoulder and elbow torque errors were not found to be significantly different from each other ($p=0.4062$). The rate of decrease in the absolute elbow torque error was found to be significantly slower than that of handpath error ($p=0.041$). However, the shoulder torque error was not found to be significantly different from the handpath error at a level of 0.05 ($p=0.063$).

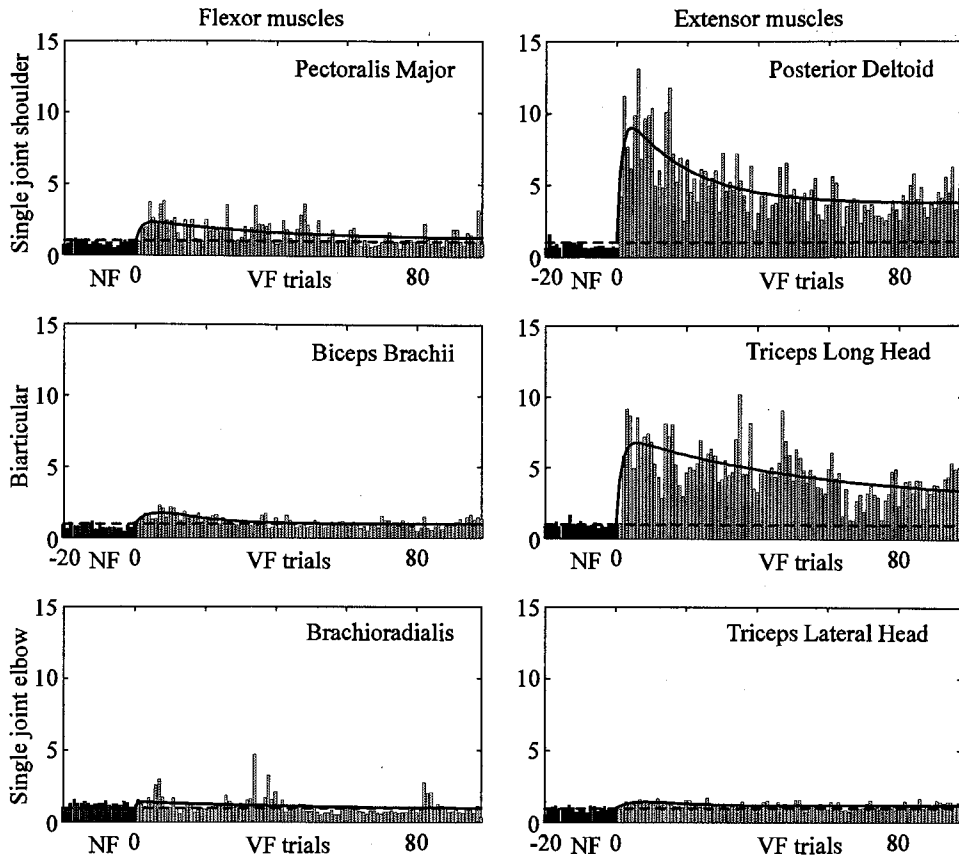


Figure 3.8 Evolution of the EMG for six arm muscles during learning in the VF for one subject. Subjects initially performed movements in the NF. The EMG of the last 20 trials of the NF is depicted for each muscle by the dark grey bars to the left of each graph. A baseline was determined from the mean NF rms EMG (dotted line) and was set to a value of one for each muscle. All trials recorded in the VF are depicted by light grey bars. The best-fit double exponential curve (using least squares error) is shown (solid line). All values are rms EMG, expressed relative to NF values, from 100 ms prior to movement initiation until 350 ms after movement initiation.

TABLE 3.1 Summary of the least square fit to the adaptation of handpath error and EMG during learning in the VF and DF

VF Adaptation

Handpath Error		τ_{\pm} SD			
		3.32 ± 2.21			
Torque Error		τ_{\pm} SD			
Shoulder		6.85 ± 14.51			
Elbow		9.78 ± 11.39			
EMG	A ± SD	τ_1 ± SD	B ± SD	τ_2 ± SD	A-B ± SD
Pectoralis Major	1.08 ± 1.00	2.89 ± 3.91	0.97 ± 0.95	22.27 ± 11.85	0.11 ± 0.22
Posterior Deltoid	8.17 ± 6.14	1.49 ± 2.10	5.90 ± 3.44	17.69 ± 5.92	2.27 ± 2.71
Biceps Brachii	1.84 ± 1.59	5.42 ± 3.55	1.74 ± 1.31	21.20 ± 17.50	0.11 ± 0.35
Triceps Longus	3.75 ± 2.24	1.74 ± 1.66	3.31 ± 1.88	27.16 ± 19.55	0.44 ± 0.70
Brachioradialis	1.51 ± 1.03	0.96 ± 0.87	1.32 ± 0.86	29.16 ± 25.39	0.20 ± 0.20
Triceps					
Lateralus	1.39 ± 1.01	3.20 ± 3.00	1.30 ± 0.92	22.57 ± 19.95	0.09 ± 0.12

DF Adaptation

Handpath Error		τ_{\pm} SD			
		13.65 ± 15.55			
EMG	A ± SD	τ_1 ± SD	B ± SD	τ_2 ± SD	A-B ± SD
Pectoralis Major	2.06 ± 3.08	6.51 ± 5.69	1.46 ± 2.15	26.68 ± 18.32	0.60 ± 0.93
Posterior Deltoid	4.14 ± 2.17	8.09 ± 5.06	3.39 ± 1.72	58.40 ± 54.57	0.725 ± 0.70
Biceps Brachii	7.40 ± 7.93	13.75 ± 13.53	6.35 ± 6.15	73.60 ± 66.55	1.05 ± 1.95
Triceps Longus	5.92 ± 7.36	17.53 ± 22.15	4.53 ± 5.20	73.45 ± 93.11	1.39 ± 2.31
Brachioradialis	5.21 ± 6.34	6.43 ± 1.72	4.61 ± 6.67	65.02 ± 67.98	0.60 ± 0.67
Triceps Lateralas	2.90 ± 1.07	8.92 ± 5.83	1.98 ± 0.71	21.06 ± 13.08	0.92 ± 0.49

Values for the gain of the two exponential functions (A , B and $A-B$) are expressed relative to the NF value (normalized to 1). The mean and standard deviation (SD) over all subjects are shown for each muscle and for the handpath error.

During adaptation to the DF, there was also an initial increase in EMG activity of all muscles, which then gradually declined. The initial increase in EMG activity reached its peak after 20 to 40 trials. A t-test was performed comparing the rms EMG of trials 11-30 to 20 trials in the NF for each muscle for each subject. Twenty-seven of the thirty comparisons showed significant increases in the rms EMG at the 0.01 significance level indicating that all muscles increased activity early in learning. The EMG then gradually declined to an asymptotic level. Again, a t-test was performed comparing the rms EMG of the last 20 trials in the DF to 20 trials in the NF for each muscle for each subject. Twenty-six of the thirty comparisons showed a significant increase in the rms EMG at the 0.01 significance level after adaptation. The initial increase in activation, which represented co-contraction, occurred at a similar rate in all muscles. The ensuing deactivation proceeded much more slowly. This suggests that the selective control of endpoint impedance was a slow process, which likely involved tuning of the relative activation of muscle pairs. The double exponential model of equation (2.9) was again able to accurately capture the variation in EMG activity over the period of adaptation to the DF (Fig. 3.9).

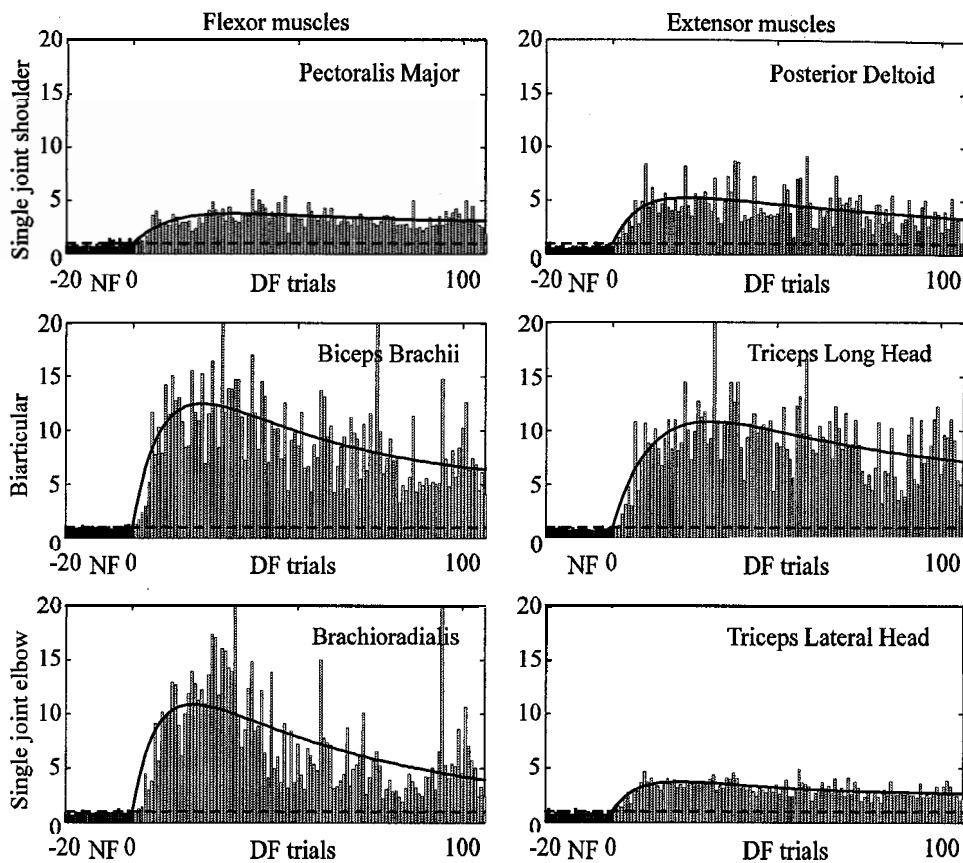


Figure 3.9 Evolution of the EMG for six arm muscles during learning in the DF for one subject. Subjects initially performed movements in the NF. The EMG of the last 20 trials is depicted for each muscle by the dark grey bars to the left of each graph. A baseline was determined from the mean NF rms EMG (dotted line) and was set to a value of one for each muscle. All trials recorded in the DF are depicted by light grey bars. The best-fit double exponential curve (using least squares error) is shown (solid line). All values are rms EMG, expressed relative to NF values, from 100 ms prior to movement initiation until 350 ms after movement initiation.

The activation and deactivation time courses in the DF for each muscle were compared to the time course of the handpath error reduction. Activation proceeded at a similar rate to handpath error reduction ($p=0.59$), whereas deactivation proceeded much more slowly ($p<0.0001$). The time courses were similar for the six muscles. As in the case of the VF, activation was significantly faster than deactivation ($p<0.0001$).

Values for all parameters are shown in Table 3.1. The gain of the activation process

(*A*) is larger for the biarticular muscles than the single joint elbow or shoulder muscles. Interestingly, the deactivation process (*B*) also is largest for the biarticular muscles, even though the final level after adaptation (expressed as an increase from NF values) (*A-B*) shows the largest increase for the biarticular muscles.

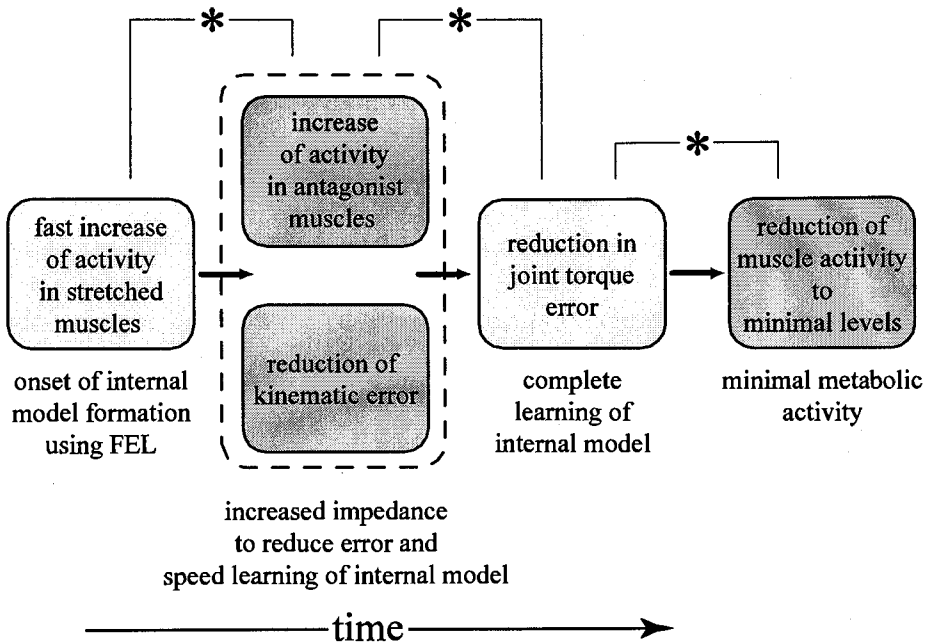
DISCUSSION

This study compared the learning processes during adaptation to a force field in which hand trajectories were stable and one in which they were initially unstable. The force fields were designed so that adaptation could be achieved solely by a change in the net joint torques in the stable case, in contrast to the unstable case, which required only a change in the limb impedance. Initial trajectories were consistently displaced in the same direction in the stable case (VF) but subjects quickly reduced their handpath error and learned to produce straight movements. The modification of joint torques occurred more slowly. The adaptation in muscle activity as learning progressed was well described by concurrent activation and deactivation processes. Learning was characterized by a rapid increase in activation of the muscles needed to compensate for the environmental force together with generalized co-contraction, which lagged slightly behind. This generalized co-contraction was later reduced, but much more slowly than the reduction of hand-path error. In the unstable case (DF), initial trajectories were displaced in both directions, usually ending outside a safety zone. With practice, subjects were able to reduce the handpath error and learned to produce straight movements to the final target location. The mean joint

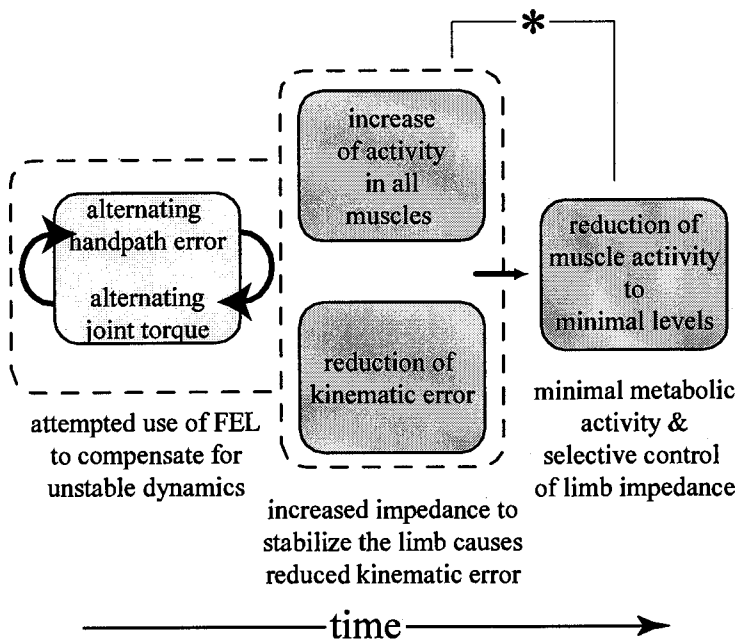
torques changed relatively little during the adaptation process. However, the early trials were characterized by alternating errors in joint torque from trial to trial and increasing absolute joint torque. Again, the modification of muscle activity was well described by a concurrent activation and deactivation process. Differential co-contraction occurred rapidly and matched the rate of hand-path error reduction. The later reduction in muscle activity occurred slowly, gradually falling to a level where endpoint stiffness would guarantee a normal safety margin for stability (Burdet et al. 2001). These processes are summarized in Fig. 3.10.

Figure 3.10 Summary of the time course of events taking place during learning in the VF (A) and DF (B). Asterisks indicate a significance difference in time constants at the 0.05 level. A: VF. Initially, fast increases are seen in the stretched muscles suggesting quick onset of internal dynamics model formation. This is followed by the increase in activity of antagonist muscle (onset of impedance controller) and reduction in kinematic error. Later, the joint torque error is reduced, which would indicate the complete learning of the IDM. Finally muscle co-contraction is reduced to minimal levels by the impedance controller. B: DF. Early trials in the DF are characterized by alternation in the direction of handpath error and joint torque which indicates that development of an IDM using feedback error learning is being attempted. Also, during the early trials, increased co-contraction occurs in all muscles (impedance control) causing a reduction in kinematic error. Finally, the muscle activity is reduced to minimize metabolic activity while ensuring the stability of the movements. This results in a selective control of the endpoint impedance.

A velocity dependent force field



B divergent force field



The VF is similar to previously studied dynamical environments in which it has been suggested that inverse dynamics models are formed (Conditt et al. 1997; Flanagan et al. 2001; Flanagan and Wing 1997; Krakauer et al. 1999; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999). Early in learning, the activation of three muscles (posterior deltoid, long head of the triceps and brachioradialis) increased rapidly. The increased activity may have comprised reflex activity arising from muscle stretch, voluntarily activation during the movement to correct for the initial displacement caused by the force field and predictive activation to counteract the expected force. The posterior deltoid and long head of the triceps contributed to the increase in shoulder extensor torque necessary to counteract the VF while the brachioradialis would have counteracted the extensor torque at the elbow resulting from the increase in activity of the long head of triceps. This adaptation is consistent with the development of an inverse dynamics model by means of feedback error learning.

In addition, muscle activity increased due to agonist-antagonist co-contraction, which can be inferred from the increase in the activity of the remaining muscles. This co-contraction acts to increase the endpoint stiffness of the arm. Increased co-contraction during the learning of novel force fields is certainly not a new finding. Thoroughman and Shadmehr (1999) have also shown increased co-contraction during adaptation to stable dynamics, which they termed wasted contraction. However, we suggest that this increased contraction of antagonist muscles is not wasted but performs an essential role in the learning of novel dynamics

by stabilizing the limb, resisting the disturbing effects of the force field, and allowing for the refinement of the feedforward inverse dynamics model (Osu et al. 2002). With improvement in the inverse dynamics model, co-contraction (limb stiffness) was reduced. This was associated with a slow deactivation process, which resulted in the activity of most muscles returning close to NF levels.

Learning in the DF

When subjects initially performed movements in the DF, their movements varied to either side of the mid-line from one trial to the next. This is evident from the joint torques, which alternated between extensor and flexor moments. The torque, relative to the NF, increased on each of the first six trials. This suggests that the subjects may have been attempting to incorporate the error information from the previous trial into the feedforward command for the next movement as would occur during feedback error learning of an inverse dynamics model. In the case of stable dynamics, this method works well and allows for quick adaptation. In contrast, this mechanism alone will not succeed when dynamics are unstable (Burdet et al. 2001). If error information from the first trial was used to update the feedforward command for the next movement, that movement would be made to the opposite side of the mid-line of the force field since subjects would produce a force opposite to direction of the previous disturbance. The resulting error would be larger and opposite in sign to the previous error. As this process continued, joint torque would tend to increase and alternate between extensor and flexor moments, although the degree to which this

would occur would depend on parameters such as the learning factor (c.f. (Scheidt et al. 2001)) and the magnitude of motor noise.

However, early in the learning period there was an increase in agonist-antagonist co-contraction, which can be interpreted as an increase in endpoint stiffness. The handpath error was reduced at about the same rate as the increase in stiffness. Furthermore, the results indicated that the faster the stiffness increased, the faster the handpath error was reduced. This provides evidence that the increased stiffness directly contributed to the reduction in handpath error. This is not unexpected, since increased endpoint stiffness counteracts the instability of the DF.

As the subject became more successful in counteracting the instability of the DF, the EMG was gradually reduced. This reduction in superfluous co-contraction would reduce metabolic energy requirements and possibly also reduce variability in motor output that tends to increase with muscle activity (Clancy and Hogan 1995; Harris and Wolpert 1998; van Galen and van Huygevoort 2000). We suggest that an interplay between two competitive processes: increased muscle activation in response to trajectory errors and a drive to minimize muscle activation necessary to perform the task results in a selective control of endpoint impedance as seen after learning in the DF (Burdet et al. 2001; Franklin et al. 2003a). The final level of muscle activity, relative to the NF, is expressed as $A-B$ (Table 3.1). The largest values are found for the biarticular muscles (long head of the triceps and biceps brachii), indicating that these muscles may play the most prominent role in the selective control of the endpoint impedance in the DF.

Impedance Controller

It is clear that the CNS is able to control the impedance of the limb. Increased muscle activity and stiffness has been seen during adaptation to many types of environments (De Serres and Milner 1991; Jones and Hunter 1992; Milner and Cloutier 1993; Takahashi et al. 2001; Thoroughman and Shadmehr 1999). More recently, we have shown that this impedance can be directionally tuned to the environment (Burdet et al. 2001). However, it is not yet clear how the CNS produces optimally oriented limb impedance. We propose the existence of an impedance controller which employs a fast activation process in response to error signals and a slower deactivation process. This controller will initially increase the impedance of the limb during adaptation to any novel dynamics. It will also attempt to minimize the activation of all muscles. By employing these two opposing processes, the initial high impedance will eventually be reduced to a minimal level necessary for stability.

It appears that the CNS engages a similar process of impedance control during the initial period of adaptation whether trajectories are inherently stable or unstable. When an additional force must be applied to the environment, as in the case of the VF, there is a natural increase in limb impedance due to muscle activation (Franklin et al. 2003a). Even though stability may be guaranteed by this naturally occurring impedance, the CNS chooses to augment it by co-contraction during the earliest phase of learning. The superfluous impedance is later eliminated once the inverse dynamics model had been acquired. However, when the naturally occurring impedance is insufficient to provide stability, as in the case of the DF, the impedance controller generates a global increase in stiffness, which is selectively reduced as the optimal

stiffness geometry is determined (Franklin et al. 2003a). Takahashi et al. (Takahashi et al. 2001) have also suggested that impedance control can coexist with the formation of inverse dynamics models for control. Robotic implementations for learning novel tasks often use high impedance as a method of achieving faster learning (Katayama et al. 1998; Sanger 1994). The increased stiffness can reduce the disturbing effects of the novel dynamics to provide better tracking of a desired trajectory during the early stages of learning. The dynamics during learning are then closer to the desired final dynamics, which will increase the speed of learning.

Learning rates for impedance control and inverse dynamics model formation

The inverse dynamics model learning and impedance learning occur simultaneously, although the former appears to proceed at a faster rate than the latter. Several lines of evidence support this conclusion. Learning in the VF, ultimately realized by inverse dynamics model formation, is faster than learning in the DF, achieved by impedance control. Deactivation of all muscles (late phase of impedance learning) is slower than acquisition of accurate torque profiles (late phase of inverse dynamics model learning) in the VF. Activation of all muscles in the DF (early phase of impedance learning) is slower than activation of the posterior deltoid and the triceps long head in the VF (early phase of inverse dynamics model learning). In summary, the early phase of inverse dynamics model learning is almost simultaneous or slightly faster than the early phase of impedance learning, which is much faster than the late phase of inverse dynamics model learning, which, in turn, is slightly

faster than the late phase of impedance learning. The interrelation of these two learning processes suggests that they could be unified in a single model.

Unified model for motor learning

The following natural extension of feedback error learning could coherently unify the two learning processes, and at least qualitatively reproduce our results, as well as other recent data on motor learning (Burdet et al. 2001; Osu et al. 2002). First, centrally generated feedforward motor commands would comprise both a reciprocal component for agonist and antagonist muscle pairs (difference in muscle activation similar to net joint torque) and a co-activation component (summation of agonist and antagonist muscle activation similar to joint stiffness). This concept originated with Feldman and has been elaborated by his collaborators (Feldman 1980a; Feldman 1980b; Gribble et al. 1998; Latash 1992; Levin et al. 1992). Second, the feedforward co-activation signal to antagonists should increase on trials following perturbation of the hand path during early learning even when only agonist muscles are stretched. This hypothesis is required to account for activation of pectoralis major up to the 5th trial in the VF (Fig. 3.8) despite the fact that it was not stretched, as all the trajectories deviated to the left (Fig. 3.3A, trials 1-6). This mechanism also contributes to the co-activation of all muscles in the DF because trajectories deviate to left on some trials and to the right on others. Third, the feedforward co-activation signal decays with a large time constant as manifested by the deactivation time constants of all muscles. We are currently developing methods to quantitatively evaluate this conceptual model both experimentally and by means of computer simulations.

APPENDIX A

The change of EMG during the learning in the force fields was characterized by an initial increase and a gradual decrease. We therefore modeled this as a double exponential process with two functions, one contributing to each of these changes. In order to assess the accuracy of this model we compared the residuals of three related models using Akaike's Information Criterion (Akaike 1974). The EMG at time zero was set equal to the mean EMG in the NF prior to presentation of the force field. In the first model, EMG was expressed as a linear function of trial number. In the second model, EMG was expressed as an exponential function of trial number, as used previously to characterize handpath error (Burdet et al. 2001; Flanagan et al. 1999) and suggested by the results of Thoroughman and Shadmehr (1999). The third model expressed EMG as a double exponential function of trial number with four free parameters (Equation (2.9)).

The three models were fit to the rms EMG of consecutive trials during learning in both the VF and DF. Because increasing the number of parameters of a model can over fit the data without adding any information, the residuals were used in Akaike's Information Criterion (AIC) (Akaike 1974) to determine if the added parameters of the double exponential function were justified by explaining more of the variation in the data. To calculate a single representative value of AIC for each model, the data from six muscles and five subjects was combined. Because the EMG magnitude varied among muscles and subjects, the data was first normalized by dividing the residual by the sum of the fitted EMG data for each muscle and subject.

The residuals, now expressed as a function of how well they fit the data, were then combined to calculate a single AIC value:

$$AIC = N \left(\log 2\pi \left(\frac{1}{sub} \sum_{i=1}^{sub} \frac{1}{mus} \sum_{j=1}^{mus} \hat{\sigma}_{i,j}^2 \right) + 1 \right) + 2(m+1) \quad (A1)$$

where N is the total number of data points, m is the number of fitted parameters in the equation multiplied by the number of muscles and subjects, and $\hat{\sigma}_{i,j}^2$ is the integrated squared error in EMG modeling for the j^{th} muscle of the i^{th} subject. This was performed separately for VF and DF data to determine which model fit the data for learning in each force field most appropriately.

The AIC values were calculated for each model in both the VF and the DF. In both fields, the double exponential model fit the data better. In the VF, the lowest Akaike's Information Criterion (AIC) was obtained for the double exponential model (42696). This was 50 less than AIC for the single exponential model and 2023 less than AIC for the linear model. A difference greater than 2 is usually considered statistically significant (Sakamoto et al 1986). The results for the DF were similar, with the lowest AIC (45374) for the double exponential model, which was 209 less than AIC of the single exponential model and 1521 less than AIC for the linear model. This supports our choice of the double exponential model, comprising an activation and a deactivation process, to describe adaptation to both the VF and DF dynamics.

IV: IMPEDANCE CONTROL BALANCES STABILITY WITH METABOLICALLY COSTLY MUSCLE ADAPTATION

INTRODUCTION

Humans constantly interact with their environment in everyday life. For example, to drink from a cup, a person must successfully reach for the cup, pick it up and then bring it to his/her mouth without spilling the contents. In order to successfully complete this activity, the person must be able to compensate for the forces exerted on the arm by the movement of the cup and its contents. It is widely accepted that the human central nervous system (CNS) learns about the dynamics of the physical world and, in particular, learns to compensate for externally imposed forces on the arm during movements (Conditt et al. 1997; Krakauer et al. 1999; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994).

Previous studies have provided evidence that the CNS learns an internal model of the interaction dynamics when movements are performed in novel mechanical environments (Kawato 1999). That is, the CNS obtains a neural representation of the relationship between motor command and actual movement. However, these studies focused primarily on stable interactions with the environment.

Relatively few studies have investigated adaptation to unstable interactions. Unstable interactions are as important as stable interactions since they routinely occur in daily activities that involve the manipulation of tools and utensils (Rancourt and

Hogan 2001). For example, the action of an artist using a hammer and chisel to make a sculpture is inherently unstable. Small forces in the direction parallel to the surface of the material being sculpted can cause the chisel to slip. The chisel must be struck squarely with the head of the hammer or the resulting force may be misdirected and mar the artist's work. However, with practice, humans can acquire the skill to compensate for such unstable interactions.

The viscoelastic properties of muscle play an important role in motor control as they respond instantaneously to disturbances and stabilize movements. The greater the arm's viscoelastic impedance, the more it resists disturbances that perturb it away from its equilibrium position or intended trajectory. The ability to control viscoelastic impedance is particularly important for stabilizing movements in unstable environments (Burdet et al. 2001) or unpredictable situations (Takahashi et al. 2001).

Hogan (1985) hypothesized that impedance might be selectively controlled by the CNS and we recently verified that hypothesis. We demonstrated that the endpoint stiffness of the arm could be tuned to the direction of an instability (Burdet et al. 2001; Franklin et al. 2003a). Such selective change in stiffness orientation, toward the direction of an instability is likely achieved by changes in the feedforward muscle activation, i.e. by controlling the activation level of particular muscle pairs. To investigate the sophistication of impedance control by the CNS we varied the level of instability of a force field and measured the endpoint stiffness of the arm after adaptation. Based upon our previous experiments, we expected a gradual increase in endpoint stiffness in the direction of the instability with little or no change in the perpendicular direction. We also anticipated that the net stiffness of the arm and the

environment would remain constant if the CNS attempted to maintain a specific margin of limb stability while interacting with the environment. Such regulation of endpoint stiffness would produce a gradual elongation of the stiffness ellipse in the direction of the instability as the degree of instability increased.

METHODS

Subjects

Five healthy, right-handed subjects participated in the study (4 male and 1 female). The experiments were approved by the institutional ethics committee and subjects gave informed consent.

Apparatus

Subjects were seated with their shoulders restrained against the back of a chair by a shoulder harness. A custom-molded rigid thermoplastic cuff was securely fastened around the subjects' right wrist and forearm, immobilizing the wrist joint. Only the shoulder and elbow joints remained free to move in the horizontal plane. The subjects' forearm was secured to a support beam in the horizontal plane and the cuff and beam were coupled to the handle of the parallel-link direct drive air-magnet floating manipulandum (PFM). Movement was thus restricted to a single degree of freedom in each joint in the horizontal plane. Our coordinate system was positive to the right (x-axis) and forward (y-axis) relative to the shoulder.

The PFM was powered by two DC direct-drive motors controlled at 2 kHz and the subjects' hand position was measured using optical joint position sensors (409,600 pulse/rev). The force applied by subjects at the handle of the PFM was measured using a six-axis force-torque sensor (Nitta Corp. No. 328) with a resolution of 0.06 N. The handle of the PFM (subjects' hand position) was supported by a frictionless air-magnet floating mechanism. The PFM was controlled by a digital signal processor (0.5 ms/cycle) to reduce the effect of the PFM's dynamics on the subjects' hand. Detailed descriptions of the PFM and controller have previously been published (Gomi and Kawato 1996; Gomi and Kawato 1997).

The subjects' view of the PFM-cuff-forearm coupling was blocked by a table placed above the PFM. The start point, end point and current hand position were displayed on the surface of the table with a projector mounted in the ceiling above the PFM. The start point and end point were respectively located 31 cm and 56 cm directly in front of the subjects' shoulder joint. A computer monitor placed behind the PFM displayed visual feedback information on the acceptance of a trial based on the timing and final location.

Procedure

Subjects produced point to point movements with their arm in a null field (NF) and in a position dependent (divergent) force field (DF) of several different strengths. The divergent force fields added negative stiffness to the arm causing a destabilizing or unstable interaction between the robotic interface and the subjects' arm. The DF produced negative elastic force perpendicular to the target direction. If

the subject made a perfectly straight movement along the y-axis from start to end point, zero perpendicular force was produced. However, if the subject deviated from the y-axis, a perpendicular x-force was introduced, increasing with deviation. Noise due to variability in the descending motor command would normally cause the initial movement direction to vary from trial to trial. The DF had an amplifying effect on this motor output variability. Even relatively small deviations in the initial trajectory, to one side or the other, caused the arm to deviate farther and farther from the y-axis as the perpendicular force increased. The DF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \quad (3.1)$$

where β (N/m) was chosen from the set {200, 300, 400, 500}. Values of β both larger and smaller than each subject's measured NF stiffness were used such that the interaction ranged from slight destabilization to instability. Only the two strongest subjects experienced the 500 strength force field. x is the lateral component of the subjects' hand position relative to the shoulder and F_x and F_y are the forces generated by the PFM in the x- and y-directions respectively. Each subject learned the force fields in a different order, and no two force field strengths were presented on a single day. Usually many days separated the presentation of two strengths of the force field. Previous work has shown that subjects are able to adapt to this unstable force field (Burdet et al. 2001; Franklin et al. 2003a; Franklin et al. 2003b; Osu et al. 2003).

A safety boundary (deviation of y-axis greater than 0.05 m) was implemented beyond which the force field reverted to a NF.

The experiment had two parts for each force field: learning and stiffness estimation. The stiffness measurements were taken either on the same day as learning or one day after based on the availability of the equipment. In the learning phase, subjects first performed 20 successful movements in the NF. A successful trial was one in which movement terminated within the allocated time (0.6 ± 0.1 s) and within the 2.5 cm diameter end target. After 20 successful NF trials, the DF was activated, although the subjects were given no warning. Subjects then practiced in the force field until 100 successful trials had been completed. All trials were recorded whether they were successful or not.

After learning, the endpoint stiffness was measured using controlled displacements as in (Burdet et al. 2001; Burdet et al. 2001; Burdet et al. 2000; Franklin et al. 2003a). Subjects first completed 40 successful movements in the force field. Then an additional 160 trials were performed, 80 of which were randomly selected for stiffness estimation. For each of these 80 trials, the PFM was programmed to briefly displace the subjects' hand at the midpoint of the movement in one of 8 randomly chosen directions. The PFM briefly perturbed the subjects' hand by a constant distance then returned the hand to its predicted unperturbed trajectory. Full details of the stiffness estimation method appear in (Burdet et al. 2000).

Learning

To determine whether learning had occurred, hand-path error E , representing the area between the actual trajectory and the straight line joining the start and end targets, was calculated as in (Osu et al. 2003). The calculation was performed from time 0 (75 ms before hand-velocity crossed a threshold of 0.05 ms^{-1}) to time T (the termination time when curvature exceeded 0.07 m^{-1} (Pollick and Ishimura 1996)). Hand-path error was calculated for all practice trials. If the subjects' hand deviated from the y -axis by more than 0.05 m at any time during a particular trial, crossing the safety boundary, the x -axis position was considered to remain at 0.05 m for the rest of the trial. This method of calculation was used to ensure that the error was accurately represented even though the force field was shut off if subjects crossed the safety boundary. An analysis of variance (ANOVA) was performed to examine whether learning took place. The hand-path error on the first 10 and last 10 trials in the DF force fields was compared across all subjects (random effect) and field levels. A similar analysis was performed using the last 10 trials in the NF prior to the start of the DF and comparing them to the last 10 trials in the DF. This was done to compare performance at the end of practice to performance in the NF field.

Endpoint force of the arm was recorded during all of the experiments. After adaptation the mean force applied to the hand should not have been different from that in the NF, regardless of the strength of the DF. To confirm this, an ANOVA was performed across all force field strengths, with subjects as a random effect, comparing the x - and y -forces in the 20 NF trials prior to the onset of the DF to the last 20 successful trials during practice in the DF. Time varying muscle torque at the

shoulder and elbow was estimated using the equations of motion for a two-link planar arm. Full details appear in (Franklin et al. 2003b).

Stiffness estimation

The endpoint stiffness of the arm was measured after learning the NF as well as after learning in each DF. Full details of the method and analysis procedure are found in (Burdet et al. 2000; Franklin et al. 2003a). Basically, a ramp up/hold/ramp down servo-controlled displacement was applied to the hand near the midpoint of the movement. Using the average force and displacement during a 60 ms interval toward the end of the hold period, an estimate of the 2x2 endpoint stiffness matrix (\mathbf{K}) was obtained by linear regression as represented by the equation:

$$\begin{bmatrix} \Delta F_x \\ \Delta F_y \end{bmatrix} = \mathbf{K} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} K_{xx} & K_{xy} \\ K_{yx} & K_{yy} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \quad (3.2)$$

where ΔF_x , ΔF_y , Δx , Δy represent the mean change in endpoint force in the x- and y-directions and the mean change in displacement in the x- and y-directions, respectively. The stiffness in different directions was represented in terms of an ellipse by plotting the elastic force produced by a unit displacement (Mussa-Ivaldi et al. 1985). This was done using the singular value decomposition method (Gomi and Osu 1998).

The joint stiffness (\mathbf{R}) was calculated from the endpoint stiffness (\mathbf{K}) using the relation:

$$\mathbf{R} = \begin{bmatrix} R_{ss} & R_{se} \\ R_{es} & R_{ee} \end{bmatrix} = \mathbf{J}^T \mathbf{KJ} - \frac{\partial \mathbf{J}^T}{\partial \theta} \mathbf{F} \quad (3.3)$$

where \mathbf{J} represents the Jacobian transformation matrix from endpoint coordinates to joint coordinates, and the last term represents the change in endpoint force due purely to the change in the geometry of the arm produced by the displacement (Franklin and Milner 2003). The Jacobian and force matrices were fixed using the mean values for position and force during the measurement interval for each subject.

RESULTS

Changes in Kinematics during Practice

In the NF, subjects performed smooth accurate movements to the target (Fig. 4.1). Small variations in the duration and path occurred from trial to trial, likely due to factors such as motor noise (Clancy and Hogan 1995; Harris and Wolpert 1998; Jones et al. 2002; Van Beers et al. 2004; van Galen and van Huygevoort 2000) and temporal deformation of the motor command (Morishige et al. in press). However, these were insignificant compared to trajectory deviations produced when the DF was first activated. Subjects' movements were markedly perturbed either to the right or the left of their normal paths and many crossed the safety boundary. This was true even for low levels of instability, which only marginally destabilized the interaction between the arm and the PFM. The destabilizing properties of the DF caused an

amplification of the trial to trial variability seen in the NF. As learning progressed, subjects' performance improved and relatively straight movements to the target were achieved on most trials. They were able to adapt the control of their arm to counteract the destabilizing effect of the DF. By the end of the training period the trajectories were relatively straight and generally successfully reached the end target, similar to NF trials. Subjects required an average of 114 ± 10 trials to achieve 100 successful trials.

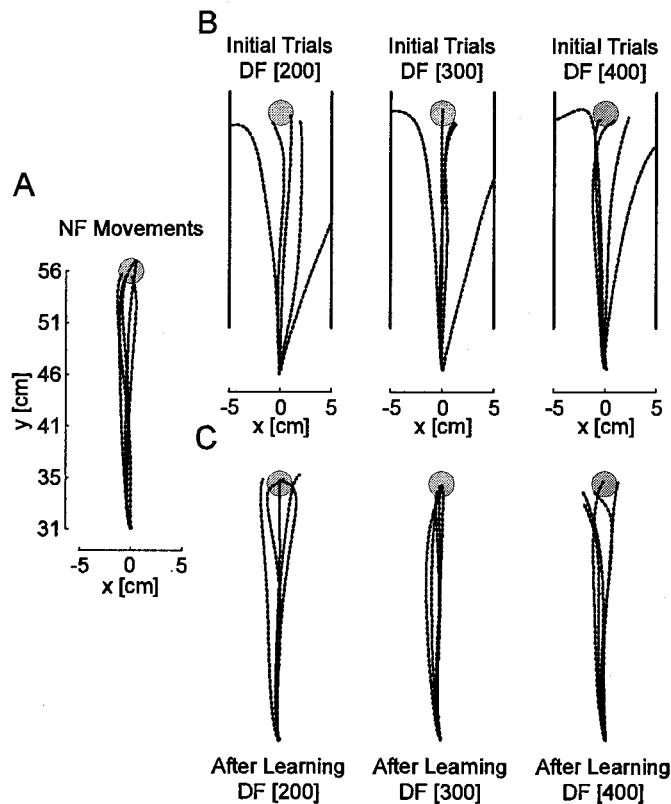







Figure 4.1 Unstable force field initially amplifies the variability of trajectories. A. Movements in the NF. B. The initial five trials in the DF fields. The safety boundary is shown by the solid black lines to either side of the trajectories, outside of which the force field was inactivated for safety reasons. C. The final five movements in the DF. The force fields for this subject were 200, 300 and 400 N/m.

The handpath error was calculated for all levels of instability (NF and DF). Early trials in the DF had large handpath errors, but these were gradually reduced during practice (Fig. 4.2). An analysis of variance (ANOVA) was performed to examine whether performance improved. The hand-path error on the first 10 trials and the last 10 trials in the DF (learning) was compared across all subjects (random effect) and field strength. A significant main effect of practice was found ($F=9.754$, $p=.029$) indicating that subjects were able to significantly reduce their hand-path error by the end of practice in all fields. A significant interaction effect between practice and field strength was also found ($F=4.053$, $p=.045$) indicating that the different force field strengths caused different amounts of error in subjects' movements. This likely arises because the different field strengths would amplify the variability in the trajectories to different degrees, causing the magnitude of the error at the beginning of practice to depend of the strength of the force field. In order to confirm this, the handpath error for the first ten trials in the force fields was compared using an ANOVA (Fig. 4.2 D). The field strength had a significant effect on the size of the error ($F=3.031$, $p=0.031$). Further support was provided by Scheffe's post hoc multiple comparison test, which indicated that the handpath error in the lowest and highest strength fields were significantly different. However, in all cases subjects adapted to the changed dynamics and the hand path error values were reduced. Not surprisingly, the success rate paralleled the hand path error. When the hand path error was large the success rate was low. The relatively low initial success rate at the highest force field strength is similar to that found by Milner (Milner 2002b), who investigated the ability of subjects to maintain a stable posture in the DF. As subjects

adapted to the force field and reduced the hand path error, the success rate also improved, particularly for the highest force field strength.

To determine the extent to which the errors were reduced, the first 10 NF trials before the activation of the DF and the last 10 practice trials in the DF were examined using an ANOVA with main effect of field strength and with subjects as a random effect. No significant difference was found between NF and DF hand-path errors ($F=2.365$, $p=.197$). This result indicates that hand-path error in the DF's was reduced to a level similar to hand-path error in the NF by the end of practice. Furthermore, no effect of field strength was found ($F=1.974$, $p=.189$) suggesting that error was reduced to levels similar to that in the NF for all force field strengths. This indicates that subjects were able to stabilize movements and achieve movement trajectories similar to those in a stable, free environment by the end of practice across all levels of instability.

Figure 4.2 Hand path error is reduced during adaptation to the DF. A-C. Error in the NF for 20 trials prior to the onset of the DF, and 70 practice trials in the DF are plotted for the 200, 300 and 400 N/m fields. Values shown are the average of all five subjects. Early trials in the DF have large errors but these were reduced quickly as subjects adapted to the environmental dynamics. D. Initial hand path error (over first 10 movements) increased as the field strength increased. Hand path error in the strongest field was significantly different from that in the weakest field ($p = 0.03$). The black line indicates the exponential best fit to the data as described previously (Osu et al. 2003). E. Success rate in DF. The percent of successful trials (not exiting the safety boundary) are shown for the first (grey) and last (black) ten learning trials at each field strength. Error bars represent the standard error of the mean. The success rate in the early trials depended on the field strength ($F = 8.282$, $p = 0.002$) where the 500 field was significantly less than the 200 and 300 fields (Scheffe's post hoc test). After learning, the success rate on the last ten trials was not significantly different across any of the force fields ($F = 2.517$, $p = 0.104$).



Endpoint Forces after Practice

At the beginning of the learning session in the DF, subjects experienced relatively large forces in the x-direction, which were positive or negative, depending on the initial direction of the trajectory. Once adaptation had occurred and subjects were able to make relatively straight movements, the x- and y-forces measured in the DF did not differ significantly from those in the NF (Fig. 4.3). An ANOVA compared the x- and y-forces in the 20 NF trials prior to activation of the DF to the corresponding forces in the last 20 successful trials of practice in the DF (practice main effect), across all levels of instability (field strength effect) with subjects as a random effect. The results from the ANOVA found no significant differences for the forces in the DF compared to the forces in the NF for either the x- or y-directions ($F=1.639$, $p=.266$ and $F=6.010$, $p=.067$, respectively). Similarly, no effect was found for level of instability in either the x- or y-force directions ($F=1.898$, $p=.200$ and $F=.493$, $p=.696$, respectively). Even for different levels of instability, the mean forces produced by subjects remained similar to those in the NF once learning had taken place. Variance in endpoint force in the DF was greater than in the NF because similar trial-to-trial variation in path was accompanied by larger forces in the DF than in the NF. Nevertheless, adaptation to the DF required no change in mean endpoint force compared to the NF.

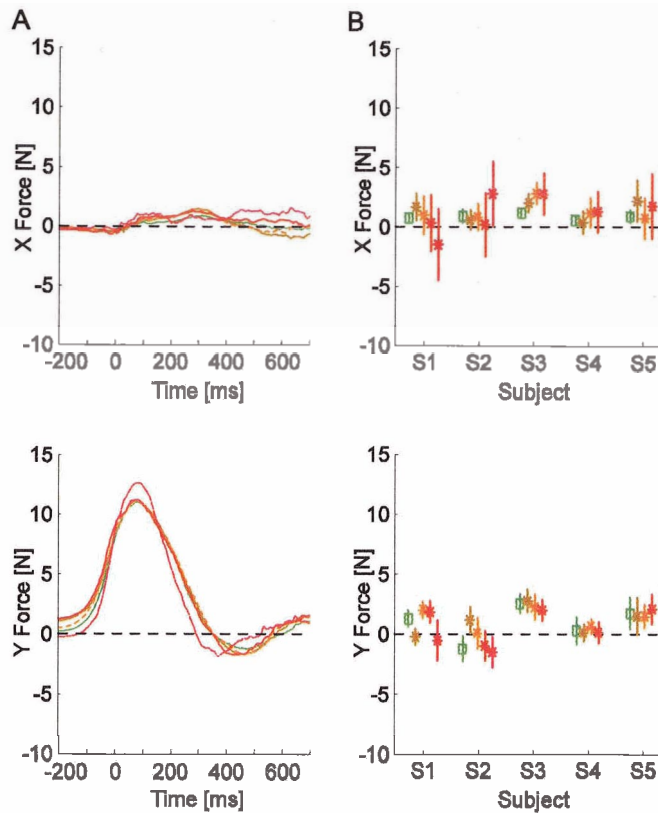


Figure 4.3 Endpoint forces were similar across all force field strengths after learning. A. Mean endpoint force profiles for the x - (top) and y -axes (bottom) during movements after learning in the NF (green), 200 DF (brown), 300 DF (light orange dashed), 400 DF (dark orange) and 500 DF (red dotted). Only subjects 1 and 2 performed trials in the 500 DF. B. Mean and standard deviation x - (top) and y -force (bottom) at the midpoint of movements after adaptation for all force fields. The midpoint of movements refers to the time at which stiffness was estimated.

Endpoint Stiffness

After adaptation to each force field, each subjects' endpoint stiffness was measured at the midpoint of the movement during movements in the force field (Fig. 4.4). Compared to the NF stiffness ellipses, the DF stiffness ellipses were more anisotropic or elongated. In particular, the dimension of the stiffness ellipses remained similar to those in the NF for the y -direction, whereas in the x -direction the dimension of the stiffness ellipses increased as field strength (level of instability)

increased. Although endpoint forces remained unchanged, subjects were able to modify their endpoint stiffness, indicating that the CNS was controlling stiffness via impedance control rather than as a side effect of changes in endpoint force (Burdet et al. 2001; Franklin et al. 2003a; Jones and Hunter 1993). This specific scaling of stiffness in the x-direction (direction of instability) with level of instability indicates that stability was not achieved by simply co-contracting all the muscles of the arm equally. If subjects had adapted to a force field by using such generalized muscle co-contraction, the resulting stiffness ellipse would have been modified mainly in size rather than in orientation and shape.




Figure 4.4 Endpoint stiffness scales with strength of the unstable force field. Endpoint stiffness ellipses after adaptation are shown for all five subjects in the NF (green line), 200 DF (brown line), 300 DF (light orange dashed line), 400 DF (dark orange line) and 500 DF (red dotted line).

When the experiments were initially conducted with two of the subjects (S3 & S4), the stiffness ellipses were found to have increased in both the x - and y -directions after adaptation to several levels of the force field. This differed from the results for

the other subjects and for the results of previous experiments (Burdet et al. 2001; Franklin et al. 2003a). We hypothesized that these subjects were either unable to control stiffness independently along the two directions unlike other subjects or they required more training to learn the adaptation strategy used by the other subjects. To test these hypotheses, we had the subjects perform an additional practice session followed by stiffness measurement for each force field. Subjects performed another 200 practice trials and then stiffness was again measured. These subjects, therefore, had practiced at least 500 trials in the force field altogether before the second stiffness measurement. The stiffness measurements before and after the second practice session are compared in Figure 4.5. After more extensive practice the endpoint stiffness was smaller in both the x- and y-directions than measurements made earlier ($p=0.004$; $p=0.011$, respectively). The reduction $((\text{original } K - \text{new } K) / \text{original } K)$ along the x-axis was small (12.6%) compared to the reduction along the y-axis (49.4%).

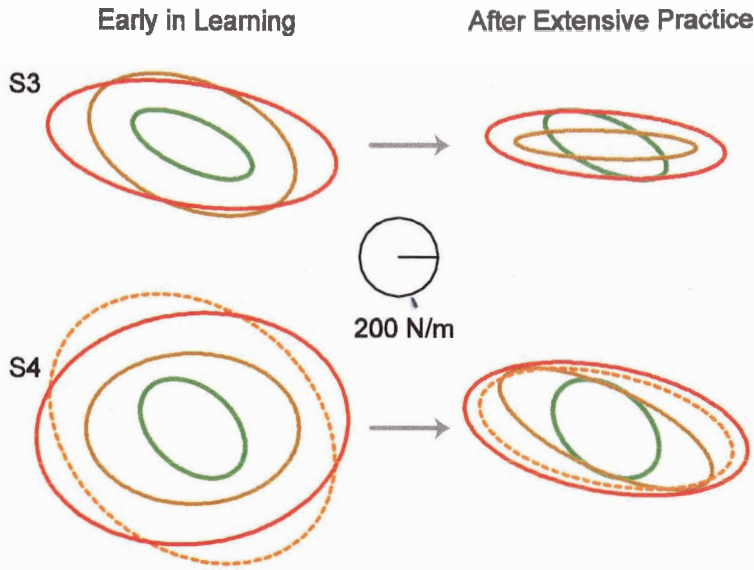


Figure 4.5 Endpoint stiffness reduced in y -direction with sufficient training. Endpoint stiffness ellipses are shown for two subjects after the initial learning session (left) and after further training on a subsequent day (right). Ellipses are shown using the same color code as in Fig 4.3. The NF ellipses are also shown for comparison.

To further investigate the effect of the level of instability (field strength) on the endpoint stiffness of the limb, the x - and y - (diagonal) components of the stiffness matrix \mathbf{K} for each subject and field strength were examined (Fig. 4.6). The x -stiffness in the DF increased progressively with level of DF instability. An ANOVA was performed on the x -component of stiffness for field strength, with the NF being classified as 0. A highly significant effect was found ($F=35.015$, $p<.001$), supporting this observation. In contrast, no significant differences (ANOVA) in stiffness in the y -direction were found with level of instability ($F=.362$, $p=.832$). Further support was provided by Scheffé's post hoc multiple comparison test, performed on the y -component of stiffness for each field strength, which showed no significant differences between any two levels and found only one homogenous subset. The same test performed on the x -component of stiffness found 4 homogenous subsets at

a significance level of $\alpha=0.05$. This means that the x-component of stiffness was selectively increased as field strength was increased without significantly changing the y-component of stiffness.

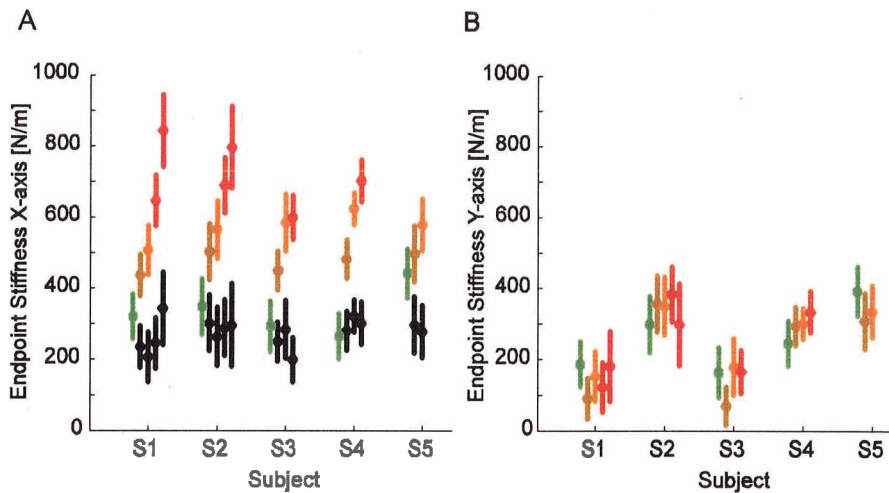


Figure 4.6 The x- and y- components of endpoint stiffness together with the net stiffness of the interaction between the limb and the force field. A. The x-component of endpoint stiffness after adaptation for all subjects in the NF (green), 200 DF (brown), 300 DF (light orange), 400 DF (dark orange) and 500 DF (red). The net stiffness (stiffness of the interaction between the limb and DF) after learning is shown in black. Vertical bars show 90% confidence intervals. B. The y-component of endpoint stiffness in the same force fields.

The net stiffness is a representation of the combined stiffness of the arm and the environment in the x-direction and is calculated as the difference between the measured arm stiffness and the force field stiffness. The net stiffness determines the overall stability of the interaction between the arm and the environment. For each force field strength, the net stiffness required to stabilize movements remained constant, at a level similar to the NF (Fig. 4.6A, black). An ANOVA performed on the net stiffness in each field across all force levels and subjects indicated no

significant difference ($F=2.176$, $p=.129$). A post hoc test also found only one homogenous subset and confirmed that there was no significant difference between net stiffness for any two field strengths, including the NF. The results indicate that in adapting to the unstable dynamics of this study, the impedance controller selectively increased stiffness in the direction of instability while maintaining stiffness in the stable direction at a level similar to that in the NF. The arm stiffness in the unstable direction was modified in a manner such that net stiffness remained at a level similar to that in the NF. This indicates that the CNS attempted to maintain a specific level of stability when adapting to unstable dynamics of different strengths.

Joint Stiffness

The joint stiffness was estimated from the measurements of endpoint stiffness. The shoulder and elbow joint stiffness (R_{ss} , R_{ee}) values were plotted against shoulder and elbow joint torque, respectively, for all subjects at all force field strengths (Fig 4.7). A linear regression was performed and no significant correlation was found between shoulder joint stiffness and shoulder joint torque ($r^2 = 0.002$). Similarly, no significant relationships were found for either of the cross joint stiffness terms with either elbow or shoulder joint torque. The elbow joint stiffness was correlated with the elbow joint torque ($r^2 = 0.32$). However, the slope of this relation was much greater than that found in previous studies under isometric conditions (Franklin and Milner 2003; Gomi and Osu 1998).

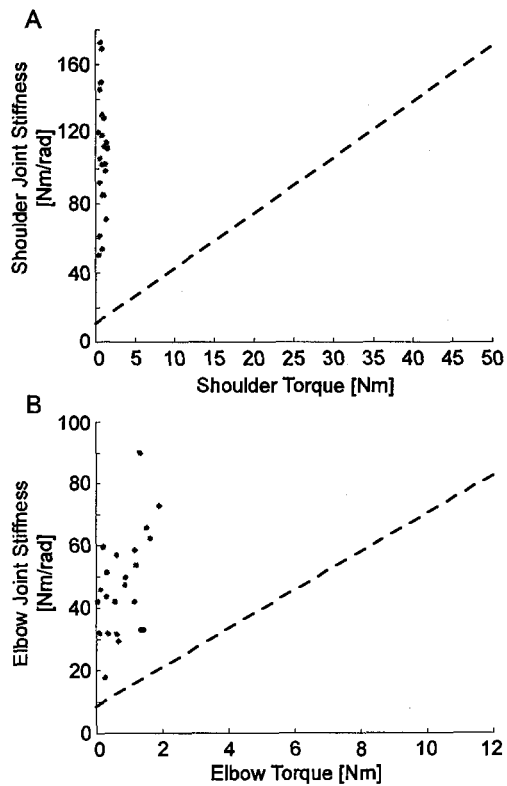


Figure 4.7 Joint stiffness increases independently of joint torque for adaptation to unstable dynamics. A. Shoulder joint stiffness R_{ss} plotted against shoulder joint torque τ_s for all subjects in all force fields. The dashed line was produced using the mean slope and intercept data from the work of Gomi and Osu (1998). The line represents the relationship between shoulder joint stiffness and shoulder torque found during stable posture [$R_{ss} = 3.18 * \tau_s + 10.80$]. B. Elbow joint stiffness R_{ee} plotted against τ_e for all subjects in all force fields. The dashed line again represents the relationship between elbow joint stiffness and elbow torque from Gomi and Osu (1998) [$R_{ee} = 6.18 * \tau_e + 8.67$].

DISCUSSION

We have investigated the adaptation of limb impedance to different levels of instability perpendicular to the direction of movement. These destabilizing force fields initially caused the subjects' trajectories to deviate to either side of the straight line between start and end targets. With sufficient practice the subjects were able to adapt to the force fields and perform movements to the target similar to those in the null force field. After adaptation, there was little difference between the forces applied to the hand for any of the force fields compared to the null force field. Therefore, any changes in the stiffness were not produced simply by a change in the joint torque. The endpoint stiffness of the limb increased with the strength of the force field. In particular, the stiffness in the x-direction increased in proportion to the increasing instability of the force field, which was oriented in that direction, whereas the stiffness in the y-direction did not significantly change with the strength of the force field. The overall stiffness in the x-direction, encompassing the stiffness of the limb and that of the environment (net stiffness), was maintained at a constant level across all force field strengths.

Interaction between motor noise and instability

When we make repeated movements, there are small variations in the trajectories from trial to trial. These trajectory variations are produced by factors such as motor noise (Clancy and Hogan 1995; Harris and Wolpert 1998; Jones et al. 2002; Van Beers et al. 2004; van Galen and van Huygevoort 2000) and temporal

deformation of the motor command (Morishige et al. in press). These variations may generally be small enough not to interfere with the performance of most activities. For movements in an unstable environment, however, the effect of motor noise is critical. The variability in trajectory produced by the motor noise is amplified by the instability of the force field. This produces a diverging trajectory to either one side or the other of the desired path. In order to compensate for this effect, the CNS must counteract the instability of the force field by increasing the stability of the limb or by reducing the variability to zero. It seems that eliminating the motor noise is not an option, so instead the CNS increases the stiffness of the limbs in order to stabilize the total system of the limb and its environment. This is similar to the strategy used to achieve higher accuracy for a given movement speed or to achieve the same accuracy for faster movements (Grey 1997; Gribble et al. 2003).

Stiffness change is independent of joint torque

As muscles are activated, the stiffness increases along with the increasing muscle force (Kirsch et al. 1994). This translates into an increase in joint stiffness with increasing joint torque (Carter et al. 1990; Gottlieb and Agarwal 1988; Hunter and Kearney 1982; Milner et al. 1995; Weiss et al. 1988). Similarly, the joint stiffness terms have been found to increase linearly with the joint torque under stable conditions for a multi-joint limb (Gomi and Osu 1998; Perreault et al. 2001; Perreault et al. 2002). This is in sharp contrast to the results reported here, where the joint stiffness increased with little change in the joint torques. Although the stiffness of the elbow joint was correlated with the joint torque, the slope of the relationship showed

that only a small portion of the change in stiffness could have been produced by the change in torque. By co-contracting antagonist muscles the joint stiffness can be increased without changing the net joint torques (Hogan 1984). This increased co-contraction and joint stiffness allows the limb to counteract the instability in the environment such that unperturbed movements are possible (De Serres and Milner 1991; Milner 2002a).

Our previous work, examining this adaptation to unstable dynamics carefully, modeled the changes in force and stiffness on a trial by trial basis and found that the change in stiffness could not be explained by the direct relationship between the changes in muscle stiffness and muscle force (Franklin et al. 2003a). However, as only a single field strength was examined, the variation in joint stiffness was relatively small. Consequently, it was not possible to correlate joint stiffness with joint torque. In this experiment, there was a large variation in joint stiffness but we found almost no correlation between joint stiffness and joint torque, confirming that the stiffness change was controlled independently of joint torque. This supports our hypothesis that the CNS explicitly controls the impedance of the limb by learning an internal model related to the stability of the external environment (Franklin et al. 2003b; Osu et al. 2003).

Net stiffness of the interaction is regulated

The net stiffness is the combined stiffness of the environment and the limb. In our case, it represents the stability of the interaction between the man and machine. Burdet et al. (2001) found that the net stiffness after adaptation was positive,

indicating a stable interaction, and that it was similar to the limb stiffness in the null force field. In this experiment, the net stiffness of the combined arm and environment was found not to vary with the level of instability. This has important ramifications for the understanding of motor control. The neuromuscular system must be able to judge the instability of the environment and adapt the stiffness of the limb in order to maintain this constant net stiffness. The idea that we are able to judge the stiffness of the environment is not new, other researchers have found evidence that humans are able to judge their environmental impedance (Jones and Hunter 1990; Jones and Hunter 1993). This may occur through the estimation of both changes in environmental force and its effect on limb kinematics. However, when moving, kinematic error alone may be sufficient to learn to counteract the effect of instability in the environment. We propose that the CNS uses the kinematic error on a given trial to update an impedance model of the environment for modifying feedforward motor commands for the next trial. Using error information from movements during repeated trials, we can slowly build up an internal representation of the stability of a task performed in a specific environment or with a specific tool and compensate for any instability. Such changes in the feedforward control would be intimately related to the development of an inverse dynamics model to compensate for necessary changes in the external force (Franklin et al. 2003a).

Control of the Endpoint Stiffness Characteristics

This study confirms the findings of our previous work (Burdet et al. 2001; Franklin et al. 2003a) and generalizes them for different levels of instability. In

particular, we were able to demonstrate that the stiffness of the limb is not just selectively increased in the direction of the instability but that it increases in proportion to the level of instability. This provides clear evidence that the CNS is able to control the impedance of the limb as proposed by Hogan (1985).

Studies of multi-joint impedance under static conditions have generally found limited control over the orientation and shape of the stiffness ellipse (Gomi and Osu 1998; Perreault et al. 2002). However, in these studies subjects were instructed to modulate their stiffness based on the feedback of EMG or measured endpoint stiffness. In contrast, during movement in an unstable environment we have shown that the stiffness of the limb can be rotated and elongated selectively in the direction of the instability. The question remains whether this is particular to movement as opposed to posture or whether it occurs as a result of the feedback and learning signals which are used to guide the subjects. In our experiment, subjects must increase the impedance of the limb above that of the environment in order to succeed at the task. They are also required to practice extensively. They, therefore, receive two key signals which may allow the CNS to learn this control of the endpoint stiffness. The first is proprioceptive sensory information from the periphery which can be used to evaluate the instability of the environment. The second is the cumulative effect of co-contraction. High levels of co-contraction allow subjects to succeed at the task but can also lead to fatigue. Subjects performing the task only once may increase stiffness far above the minimum that is required. However, subjects performing the task repeatedly will find it useful to learn to minimize muscle activity. This view is supported by the finding that there was less co-contraction

during a force control task than during a less stable position control task (Franklin and Milner 2003).

However, there are other possible explanations that could account for differences in impedance control between posture and movement. In recent experiments examining load-sensitivity of neurons in primary cortex, it appears that independent populations of neurons respond to loads during posture (static loading) and movement (dynamic loading) (Scott 2004). Different excitatory and inhibitory projections of two such populations of cortical-motor neurons might result in differences in the ability to regulate co-contraction and stiffness. Nevertheless, there are theoretical limits to the modification of stiffness geometry by the CNS, and we are currently investigating how the CNS works within these constraints when the interaction dynamics are unstable.

Metabolic cost is reduced

To adapt to the DF, the CNS selectively modified stiffness such that a specific level of stability was maintained in the direction of the instability without change in the perpendicular direction. Under static conditions without training, Mussa-Ivaldi et al. (1985) found a global increase in stiffness in response to sinusoidal perturbations applied in specific directions. The strategy seen here however, to maintain a specific stability margin for all levels of force field instability, was energy efficient as any additional stiffness in either the stable or unstable directions would have required more muscle co-activation and thus, more expenditure of metabolic energy. Increased muscle activity is generally related to increased metabolic cost (Foley and Meyer

1993; Hogan et al. 1996; Sih and Stuhmiller 2003). A non-specific increase in stiffness magnitude would require co-activation of all pairs of antagonist muscles in the arm, which would be more costly in terms of energy expenditure than selective co-activation. The selective increase in stiffness along the direction of instability indicates that the CNS does not increase the activation of all muscles equally, but attempts to find a balance when modifying arm impedance. It faces the optimization problem of maintaining stability while minimizing metabolic cost.

A greater reduction in the metabolic cost might be achieved by relying on feedback control, either reflexive or voluntary, rather than feedforward control. However, there are several problems with such a strategy. First, feedback correction operates with a significant delay. During this time, the arm would be pushed farther away from the desired trajectory by the force field and the force applied to the arm would increase. Not only would the error be large, but the force needed to correct the error would be difficult to estimate because of the instability. In addition, the metabolic cost of initial force production is higher than that required to maintain the force level (Russ et al. 2002). Consequently, a series of brief contractions, e.g. 250 ms, could have a higher metabolic cost than a single contraction of longer total duration (Hogan et al. 1998). Therefore, it may be more efficient metabolically to co-contract muscles continuously to produce a stable interaction with the environment throughout the entire movement rather than reciprocally activating muscles and reacting to perturbations after they occur. Furthermore, such feedforward control can ensure that perturbations do not occur, which is a prerequisite for skilled tool use that could not be guaranteed if the CNS relied on delayed feedback control.

While this paper attempts to explain the results in terms of achieving stability with minimal metabolic cost, there are also other possible explanations why the muscle activity is being reduced. In particular, other parameters such as net muscle generated stress at a joint or a measure of the effort or the neural activity could easily replace our proposed metabolic cost. The current work cannot distinguish between any of these related possibilities. However, we feel that the assumption that the CNS balances stability and metabolic cost is the most intuitive interpretation, and the most understandable in terms of the evolutionary pressures on the development of motor learning.

V: NEW PRINCIPLES FOR A UNIVERSAL FRAMEWORK OF MOTOR LEARNING

INTRODUCTION

In order to perform even simple movements, we must control time-varying forces applied with our limbs. Evidence from many adaptation studies have illustrated that the central nervous system (CNS) learns to compensate for the dynamics of the environment by forming an internal representation of the external dynamics (Condit et al. 1997; Flanagan et al. 2001; Imamizu et al. 2000; Krakauer et al. 1999; Lacquaniti et al. 1992; Li et al. 2001; McIntyre et al. 2001; Shadmehr and Holcomb 1997; Shadmehr and Mussa-Ivaldi 1994). It has been proposed that this adaptation occurs using a feedback error learning mechanism (Kawato et al. 1987) where the change in the joint torque required to adapt the movement to the environment is produced by a shift in the command based on the size of error (Bhushan and Shadmehr 1999; Donchin et al. 2003; Gribble and Ostry 2000; Sanner and Kosha 1999; Schweighofer et al. 1998). While this mechanism was consistent with the learning of stable environmental dynamics, recent studies investigating unpredictable (Scheidt et al. 2001; Takahashi et al. 2001) or unstable environments (Burdet et al. 2001; Franklin et al. 2003a; Franklin et al. 2003b) have illustrated that this mechanism cannot explain (Burdet et al. 2001; Osu et al. 2003) the true range of human adaptation which encompasses not only stable tasks but unstable tasks such as tool use (Rancourt and Hogan 2001). These previous mechanisms cannot explain the

learning of impedance or changes from trial to trial in muscle activation. Here we investigate trial-by-trial changes in the motor command during motor learning and propose a mechanism of adaptation which can explain adaptation to both stable and unstable environments.

While stable environments produce predictable, repeatable dynamics from trial to trial, unstable or unpredictable environments produce dynamics which will vary from trial to trial based upon the environment and initial conditions. In these environments, subjects increase the impedance of the limbs in order to limit the disturbing effects of the dynamics and compensate for the instability (Burdet et al. 2001; Takahashi et al. 2001). The subjects learn to selectively control the impedance of the arm in a way that decreases the metabolic cost of maintaining stability (Burdet et al. 2001; Franklin et al. 2003a; Franklin et al. 2004). Whether the environment interacts in a stable or unstable manner with the arm, similar patterns of muscle activation are observed if the environment initially disturbs the normal trajectory (Franklin et al. 2003b). Subjects appear to use a strategy of increasing co-contraction along with a change in joint torque acting to reduce the perturbing effects of the environmental dynamics. Gradually, as the subjects learn to compensate for the environmental dynamics, excess co-contraction is reduced and subjects are able to compensate for the new dynamics with less metabolic cost (Franklin et al. 2003a; Franklin et al. 2003b; Franklin et al. 2004; Osu et al. 2003). So far, no mechanism of motor learning has been proposed which can deal with both stable and unstable dynamics nor do we have a mechanism that would allow the mechanical impedance of the arm to adapt selectively to the characteristics of the environmental stability. In

order to elucidate the adaptation mechanisms, we examined the trial-by-trial changes in muscle activation as subjects learned to compensate for the perturbing effects of different force fields.

METHODS

A total of nine healthy individuals participated in this study (20-34 years of age; 4 females and 5 males; all right-handed), although only one subject participated in all experiments. The institutional ethics committee approved the experiments and the subjects gave informed consent prior to participation.

Apparatus

Subjects sat in a chair and moved the parallel-link direct drive air-magnet floating manipulandum (PFM) in a series of forward reaching movements performed in the horizontal plane. Their shoulders were held against the back of the chair by means of a shoulder harness. The right forearm was securely coupled to the PFM using a rigid custom molded thermoplastic cuff. The cuff immobilized the wrist joint, permitting movement of only the shoulder and elbow joints. The subjects' right forearm rested on a support beam projecting from the handle of the PFM. Motion was, therefore, limited to a single degree of freedom at the shoulder and at the elbow. The manipulandum and setup are described in detail elsewhere (Gomi and Kawato 1996; Gomi and Kawato 1997).

Force Fields

The experiment examined trajectory and EMG adaptation in three force fields: a velocity dependent force field (VF), in which hand trajectories were stable from the outset and two position dependent (divergent) force fields (DF and 45DF), in which hand trajectories were initially unstable. Results were compared to those in a null field (NF). Details of the implementation and protocol have been described elsewhere (Burdet et al. 2001; Franklin et al. 2003b). The force (F_x, F_y) (in N) exerted on the hand by the robotic interface in the VF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \chi \begin{bmatrix} 13 & -18 \\ 18 & 13 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (4.1)$$

where (\dot{x}, \dot{y}) is the hand velocity (m/s) and the scaling factor, χ , was adjusted to the subject's strength ($2/3 \leq \chi \leq 1$). While subjects were performing movements in the NF the force field was switched to the VF on random trials (before effects).

The DF produced a negative elastic force perpendicular to the target direction with a value of zero along the y -axis, i.e., no force was exerted when trajectories followed the y -axis, but the hand was pushed away whenever it deviated from the y -axis. The DF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (4.2)$$

where the x -component of the hand position was measured relative to the shoulder joint. $\beta = (300 \text{ to } 500) \text{ (N/m)}$ was adjusted for each subject so that it was larger than the stiffness of the arm measured in NF movements so as to produce instability. For safety reasons, the DF force field was inactivated if the subjects' trajectory deviated more than 5 cm from the y -axis.

The 45DF produced a negative elastic force rotated 45 degrees from the target direction. Similar to the DF, the force field exerted no force on the hand along the y -axis. The 45DF was implemented as:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta & 0 \\ \beta & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \quad (4.3)$$

$\beta = (300 \text{ to } 500) \text{ (N/m)}$ was the same value as used in the DF. Again, for safety reasons, the DF was inactivated if the subjects' trajectory deviated more than 5 cm from the y -axis. All force fields were inactivated once the subject reached the target position.

Learning

All subjects practiced in the NF on at least one day prior to the experiment. These training trials were used to accustom the subjects to the equipment and to the movement speed and accuracy requirements. Experiments in each force field were conducted on separate days, usually several days apart.

Subjects first practiced in the NF until they had achieved 50 successful trials. Successful trials were those which ended inside a 2.5 cm diameter target window within the prescribed time (0.6 ± 0.1 s). All movements were recorded whether successful or not. The movement distance was 0.25 meters. Movements were self-paced so subjects were able to rest between movements if they wished. At the completion of 50 successful trials, the force field was activated. No information was given to the subjects as to when the force field trials would begin. Subjects then practiced in the force field until achieving 75 successful trials.

Hand path errors

The adaptation to the force fields was quantified by calculating the error relative to a straight line joining the centers of the start and target circles. The signed hand path error, defined as

$$S(e_x) = \int_{x=t_0}^f x(t) |\dot{y}(t)| dt \quad (4.4)$$

is a measure of the mean directional extent by which the path deviates from the straight line. Hand path errors were calculated from the start time, t_0 (75 ms prior to crossing a hand velocity threshold of 0.05 m/s), to the termination time, t_f (when curvature exceeded 0.07 mm^{-1}) (Pollick and Ishimura 1996).

Electromyography

Surface electromyographic (EMG) activity of six arm muscles was recorded using pairs of silver-silver chloride surface electrodes during the learning sessions. The electrode locations were chosen to maximize the signal from a particular muscle while avoiding cross-talk from other muscles. The skin was cleansed with alcohol and prepared by rubbing in electrode paste. This was removed with a dry cloth and pre-gelled electrodes were then attached to the skin with tape. The spacing between the electrodes of each pair was approximately 2 cm. The impedance of each electrode pair was tested to ensure that it was below 10 k Ω .

The activity of two monoarticular shoulder muscles, pectoralis major and posterior deltoid, two biarticular muscles, biceps brachii and long head of the triceps, and two monoarticular elbow muscles, brachioradialis and lateral head of the triceps, was recorded. The EMG signals were analog filtered at 25 Hz (high pass) and 1.0 kHz (low pass) using a Nihon Kohden amplifier (MME-3132) and then sampled at 2.0 kHz. All comparisons between force field EMG and NF EMG involved data recorded on the same day without removal of the electrodes. During all analysis the EMG was aligned on movement onset. In order to quantify changes in components of the EMG, all analysis was performed using the integrated rectified value of the EMG over various time intervals relative to the time of movement onset.

Timing of Reflex Responses

After learning of a particular force field had been completed, subjects de-adapted by performing 30 to 40 trials in the NF and then took a short break. More

movements were then performed in the NF. On randomly selected trials the NF was unexpectedly changed to the force field to elicit reflex and voluntary corrective responses to the perturbing effect of the force field. A total of 80 NF and 20 force field trials were recorded for each force field. The force field trials were called before effects (BE). A comparison of the BE EMG to that of the NF trials allowed us to determine the latency of the corrective responses.

During the BE trials, 20% of the trials included force field trials which disturbed the trajectory of the subjects. These trials could have influenced the muscle activity of the intervening NF movements and consequently modified the feedforward activity in the BE trials. An increase in feedforward co-contraction during NF trials was found during intermittent learning of a destabilizing velocity-dependent force field (Milner and Franklin, unpublished).

If the muscle activity was modified in the BE conditions compared to the NF conditions, then there could be differences between the effect of the force field on BE trials and on the first force field trial during learning. In particular, the trajectory disturbance on the BE trials might not be as great as that on the first trial during learning if the stiffness of the arm increased due to feedforward co-contraction. This could influence the latency and size of the feedback responses estimated using the BE trials. In order to test for this, the muscle activity on the NF trials preceding each BE trial was compared to NF trials recorded prior to any experiments. The integrated muscle activity was estimated from 100 ms prior to movement onset until 130 ms after movement onset, corresponding to the latency of the observed feedback responses (see Results section). An ANOVA was then used to test if the muscle

activity was different between these two conditions. The ratio of the muscle activity on the NF trials, which were interspersed between BE trials, divided by that of the original NF trials was also estimated. No significant differences were found for any of the six muscles at the 0.05 significance level (pectoralis major: $p = 0.33$, ratio = 93%; posterior deltoid: $p = 0.61$, ratio = 102%; biceps brachii: $p = 0.14$, ratio = 94%; triceps longus: $p = 0.83$, ratio = 97%; brachioradialis: $p = 0.95$, ratio = 99%; triceps lateralis: $p = 0.18$, ratio = 97%). Similarly, the ratio is close to 100% in all cases, demonstrating that the muscle activity was not increased on the BE trials. This indicates that the muscle activity and, therefore, the limb stiffness would have been similar both at the start of learning and during the BE trials. This would result in similar perturbations being applied to the limbs. We can, therefore, assume that the BE trials reproduced the feedback latency and response amplitude which would be seen during the first trials of learning.

Separation of feedback and feedforward commands

The muscle activity during a movement arises from two processes: the feedforward command and the feedback command. In order to examine how the feedforward command is modified, it was necessary to separate the contributions of these two processes. From analysis of the BE trials, the feedback activity of the posterior deltoid muscle was determined to exhibit a roughly linear increase with kinematic error. The feedforward activity on the first trial in the force field was assumed to be equal to the mean of the preceding ten NF trials. The feedback activity on this trial was then estimated as the difference in integrated EMG between the first

force field trial and the NF trials. All EMG was rectified and then integrated from -100 ms before movement to 600 ms after the start of the movement. For learning in the VF, feedback gain was defined as the ratio of this feedback activity to the kinematic error on the first trial. The feedback on each trial throughout the learning session was then estimated as the kinematic error on that trial multiplied by this gain. In particular, the feedback (FB) on each trial (i) was estimated as:

$$FB_i = S(e_x)_i \cdot \frac{\sum_{t=-100}^{600} EMG_1(t) - \sum_{t=-100}^{600} EMG_{NF}(t)}{S(e_x)_1} \quad (4.5)$$

where EMG_1 indicates the rectified muscle activity on either the first trial, EMG_{NF} indicates the mean rectified muscle activity of ten trials in the NF, and $S(e_x)$ refers to the signed error on each trial. Our estimate of the feedback activity was then subtracted from the actual integrated EMG to give an estimate of the feedforward EMG (FF). In particular, the feedforward EMG on each trial (i) was estimated as:

$$FF_i = \sum_{t=-100}^{600} EMG_i - FB_i \quad (4.6)$$

This was done only for the learning in the VF because the force field elicited repeatable perturbations of the shoulder joint. Analysis was carried out for the posterior deltoid, which acted as an agonist to counteract the VF. The change in the feedforward command was estimated by subtracting the feedforward activity of the

previous trial from the feedforward activity of the current trial. The cross-correlation function between change in feedforward activity and feedback activity was then computed.

A second method was also used to estimate the feedforward component of the muscle activity which involved constraining the analysis to the early portion of the movement. As mentioned previously, the BE trials were used to determine the latency of feedback activity. The integrated muscle activity prior to the onset of feedback was then used to directly estimate the early portion of the feedforward command. The change in the feedforward command was then estimated for the posterior deltoid and pectoralis major and related to the kinematic error of the previous trial. Fifteen to twenty-five trials from the beginning of the learning sessions for the three different force fields (VF, DF and 45DF) were selected for each subject. The force field EMG for each subject was scaled by the corresponding NF EMG and combined for all subjects. The change in the feedforward command was calculated and the value entered into one of ten groups, based on the size of the trajectory error on the previous trial. The change in feedforward activity for each group was tested for a significant difference from zero using a t-test at an alpha of 0.05.

RESULTS

When we first make movements in novel dynamics, our limbs are perturbed away from our planned trajectory. This elicits a reflex response, and on a longer scale, a voluntary correction, in order to restore the arm to the desired location or to

complete the task appropriately. Gradually, as the subjects learn, this feedback is replaced by a feedforward control and the subjects are no longer disturbed by the external dynamics. The muscle activity in the VF was compared in the first few movements in the novel dynamics to determine how it changed (Fig. 5.1). The activity in the first movement included large reflex responses occurring at a delay of approximately 200 ms after the start of movement in a stretched muscle and at a longer delay in its antagonist. However, on the second movement, a large increase in the muscle activity occurred prior to the reflex latency. Since the trajectory error decreased on the second movement, we can assume that error information from the feedback activity of the first movement had been used to change the feedforward activity. Furthermore, the antagonist muscle, which was not stretched by the perturbing effect of the novel dynamics, but did show a reflex response, also had increased feedforward activity in the following trials. There was gradual straightening of the path over the first few trials in the force field accompanied by changes in the profile of muscle activity. As the feedforward command was adapted to match the external dynamics, there was a gradual decay of the feedforward muscle activity as illustrated by the reduced muscle activity on trial 85.

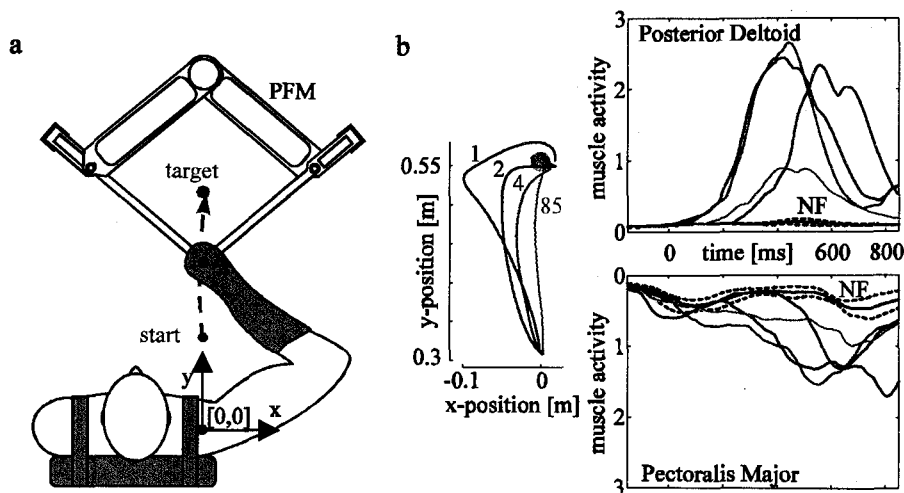


Figure 5.1 Changes in muscle activity during learning of a novel skill. A, The PFM exerts forces on the hand during horizontal point-to-point reaching movements. B, Initial movements in the VF colored according to the trial number. Hand trajectories (left). EMG activity in the posterior deltoid and pectoralis major (right). The NF activity is shown in black (solid: mean of 20 movements; dashed: standard deviation). The activity in the first movement (blue) contains large feedback responses delayed about 200 ms from the start of movement in the stretched posterior deltoid and slightly longer in the antagonist (pectoralis major). By the second movement (red), a large increase in the activity occurs prior to the previous trial's feedback, suggesting a feedforward pathway. A similar response increase was also seen in the shortened pectoralis major. The shape of this feedforward command is similar to the feedback of the first movement, suggesting that this initial error signal is driving the learning of both the agonist and antagonist muscles. The response is similar in the fourth movement, although small refinements of muscle activity likely result in the improved trajectory compared to the second trial. Gradually, as the feedforward command matches the external dynamics, there is a decrease of the muscle activity.

Feedback

The feedback muscle activity was examined by combining data from the BE trials of all force fields and subjects. In BE trials, the subjects believe that they are moving in the NF. Therefore, the feedforward command was assumed to be equal to the mean EMG in the NF. The feedback was calculated as the difference between these two values. The latency of the feedback response was determined by testing if

the BE EMG was significantly greater than the NF EMG over 10 ms intervals from 100 ms prior to the start of the movement (Fig 5.2). Because the magnitude and latency of the feedback command could vary with the size of the kinematic error in the BE, the signed error was used as an estimate of the total kinematic error on a particular trial. An ANOVA with main effects of force field (BE or NF), signed error, and subjects was used to test if the BE was different from the NF activity for each time interval. Significance differences were examined at the alpha 0.01 level.

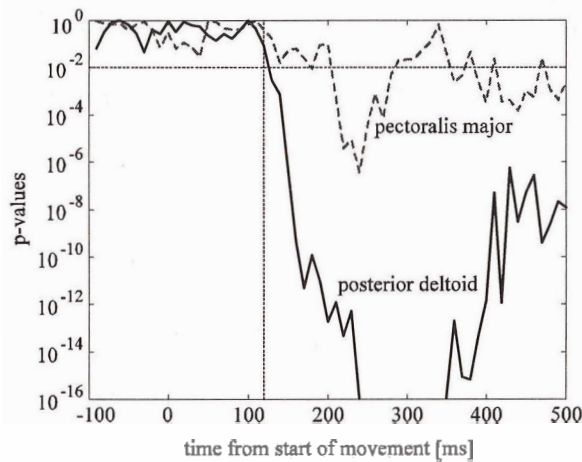


Figure 5.2 Determining the time of feedback onset during perturbed trials. Muscle activity in BE trials was compared to that in NF movements. For each 10 ms period from -100 ms prior to the start of the movement, until 500 ms after the start of movement the BE integrated EMG was compared to the NF integrated EMG over the same interval using ANOVAs with subjects as a random effect. The p-values of these comparisons are plotted on a log scale against time interval to determine the onset of the feedback response. No significant differences were seen prior to 130 ms (dotted vertical line) after the start of the movement. The horizontal line is located at $p = 0.01$. The earliest feedback responses were seen in the posterior deltoid.

The BE activity was not significantly different from that of the NF trials prior to 130 ms after the start of the movement for any of the six muscles (Fig 5.2). The earliest significant difference was found in the posterior deltoid muscle at 130-140 ms

after the start of the movement. The muscle activity was examined over three separate intervals: before feedback (-100 to 130 ms); early feedback (150 to 250 ms); and late feedback (350 to 450 ms) relative to the onset of movement using an ANOVA to test if it was different from NF activity (Fig 5.3). All differences were considered significant at the $p < 0.05$ level. There was no significant difference for the before feedback interval [-100 to 130 ms] ($p=0.99$ pectoralis major, $p=0.28$ posterior deltoid) across all of the force fields. The posterior deltoid, stretched in BE trials, exhibited increased feedback activity for both the early and late intervals ($p < 0.00001$, $p < 0.00001$). In the pectoralis major, shortened during the BE trials, two responses were found: initial inhibition ($p < 0.00001$) followed by later excitation ($p < 0.00001$).

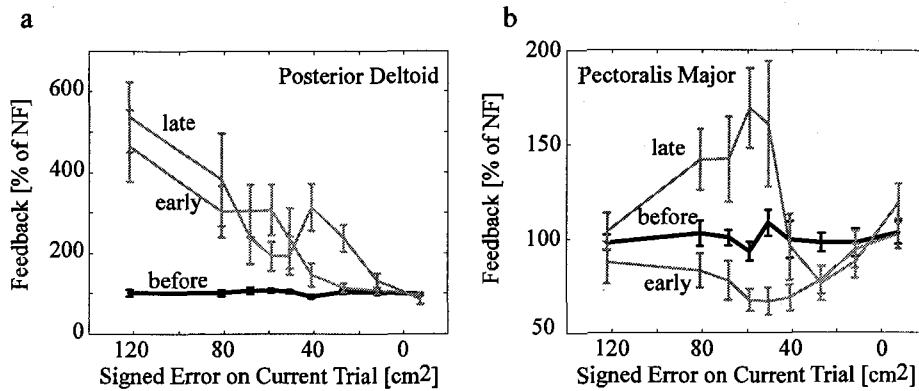


Figure 5.3 Feedback responses in the shoulder muscles for three different time intervals. Results are shown for both the pectoralis major (A) and posterior deltoid (B). For each interval: before feedback [-100 to 130 ms], early feedback [150 to 250 ms] and late feedback [350 to 450 ms], the integrated muscle activity was determined and grouped depending on the signed kinematic error. An ANOVA was used to examine differences between the activity of the BE trials with a main effect of signed error and random effect of subjects. Black indicates no significant difference and grey indicates significant difference from the NF. The mean values \pm sem are shown as a percentage of the muscle activity in the NF movements in the same interval. No difference was found before feedback. However, significant differences were found for the other two intervals. The posterior deltoid shows an increase in the feedback response as the kinematic error increases for both the early and late feedback intervals. However, the pectoralis major illustrates two distinct responses. Over the early interval the activity is reduced (inhibited) but during the later interval the feedback activation increases. This excitation increases as kinematic error increases up to a limit after which it is reduced. For the largest kinematic errors little difference was seen.

Separation of feedback and feedforward commands

The feedback activity of the posterior deltoid muscle exhibited a roughly linear increase with kinematic error (Fig. 5.3A). By taking advantage of this, we estimated the feedback component from the kinematic error. This allowed us to separate the feedback from the feedforward components of the EMG for the posterior deltoid during learning in the VF. Once the feedforward command was estimated, the change in the feedforward command was calculated by subtracting the feedforward command on the previous trial from the feedforward command on the current trial.

The change in this feedforward component from one trial to the next was then cross-correlated with the feedback during the learning resulting in a single significant peak for the preceding trial (Fig 5.4). This indicates that the change in the feedforward activity depends uniquely on the feedback activity of the previous trial.




Figure 5.4 The change in the feedforward command is dependent upon the feedback of the previous trial. Changes in the feedforward and feedback commands of the posterior deltoid muscle were estimated during learning in the VF. The change in feedforward activity of the posterior deltoid was cross-correlated with the feedback activity. A single peak occurs for the previous trial (-1) indicating that the change in the feedforward activity was dependent only on the feedback from the previous trial. The horizontal lines indicate two standard errors.

Change in feedforward command

The way in which this preceding trial influenced the change in muscle activity was then examined. In order to determine the effect of the positional error information on the feedforward command, the change in the feedforward command from one trial to the next was computed and plotted against the signed error of the previous trial (Fig. 5.5). Feedback activity was excluded by calculating the change in the feedforward component between 100 ms prior to and 130 ms after movement initiation. The change in the feedforward command was determined for the initial

learning trials in several force fields. The change in the feedforward command was separated into ten groups according to the size of the trajectory error on the previous trial. The change in feedforward activity for each group was tested for a significant difference from zero using a t-test with significance at an alpha of 0.05. We found that for both the posterior deltoid and pectoralis muscle, the feedforward command on the next trial was significantly higher if the muscle had been stretched. However, the command was also higher if its antagonist muscle had been stretched on the previous trial. Finally, if there was little kinematic error, then the feedforward command was actually decreased on the next trial.

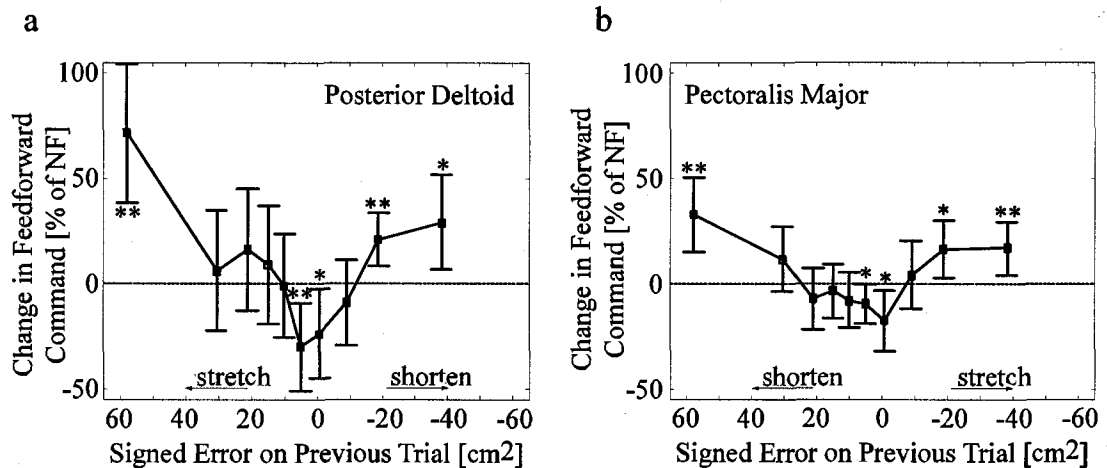


Figure 5.5 Change in feedforward activity shown relative to the signed kinematic error on the previous trial during learning. Data is shown for the pectoralis major (A) and posterior deltoid (B). The feedforward command was calculated over the -100 to 130 ms interval relative to the start of movement before any feedback activity could contribute to the measurement. The mean values \pm sem are shown as a percentage of the muscle activity in the NF movements in the same interval. Values from all subjects and all force fields were combined and then separated into 10 groups, according to the kinematic error of the previous trial. For each group, a t-test was performed to examine if the change in the feedforward activity was significantly different from zero. The symbols * ($p < 0.05$) and ** ($p < 0.01$) indicate that the change in feedforward activity was significantly different from zero. Both muscles showed similar responses: if the muscle was either stretched or shortened on the previous trial the feedforward command increased on the subsequent trial. However, if the kinematic error was close to zero on the previous trial, then the feedforward command of both muscles was reduced.

DISCUSSION

This research examined the motor learning process in detail. Muscle activity was examined during learning in three force fields and in BE trials. Feedback activity from trajectory errors under these force field conditions did not occur until at least 130 ms from the start of the movement. The earliest response of feedback was excitation of the stretched muscle and inhibition of the shortened muscle. However, the activity later increased in both stretched and shortened muscles. When the muscle activity was separated into feedforward and feedback activity, the change in the feedforward component was significantly correlated with the trajectory error of the previous trial. The feedforward activity was shown to increase if the perturbing effect of the force field produced a trajectory error large enough to either stretch or shorten a muscle on the previous trial, but decrease if the error was close to zero, i.e. if the perturbing effect of the force field was below some threshold.

Feedback components of the muscle activation

In order to examine the adaptation process, the muscle activity was separated into feedback and feedforward components. The feedback was examined by recording BE trials in the force fields (Osu et al. 2003). In these trials, the feedforward component is equal to the muscle activity in the NF which allows us to directly estimate the entire feedback component. The earliest onset time of the feedback component was estimated as 130 ms after the start of the movement. This is long compared to the time normally associated with reflex responses to perturbations

in the arm muscles which have a latency of 30-40 ms (Lacquaniti and Soechting 1984; Lacquaniti and Soechting 1986a; Sinkjaer and Hayashi 1989). However, in such experiments, the reflex response was elicited by a fast controlled perturbation. In the present study, the force field only gradually applied a force as the movement proceeded. The trajectory only deviated gradually from the mean trajectory of NF movements during the first 100 ms. Based on the latency of the feedback onset, it was only after 90-100 ms that the trajectory deviated sufficiently to elicit a reflex response.

The feedback activity of a stretched muscle increased as the muscle was stretched more. This is consistent with the results of other studies (Lee and Tatton 1982; Smeets and Erkelens 1991; Stein and Kearney 1995). For muscles that were shortened by the effect of the force field, two responses were found. At a shorter latency the muscle was either inhibited or defacilated. However, at a longer latency it was actually excited. Single joint studies have often shown that the antagonist of a stretched muscle is inhibited after the limb is perturbed (Gielen et al. 1988). Under the multi-joint condition, other responses can occur. Lacquaniti and his colleagues (Lacquaniti and Soechting 1984; Lacquaniti and Soechting 1986a; Lacquaniti and Soechting 1986b; Soechting and Lacquaniti 1988) in a series of elegant studies showed that in a multijoint limb, a shortened muscle, which is inhibited at a short latency, can be excited at a longer latency. These later responses were organized in such a way that the overall goal of the movement was more closely achieved. This was later confirmed for both single and multi-joint movements (Koshland and Hasan 2000). While the long latency response attempts to return the perturbed movement to

the intended trajectory, it does not necessarily directly oppose the perturbation (Gielen et al. 1988; Koshland and Hasan 2000). The short monosynaptic reflexes appear to produce relatively invariant hardwired responses to the disturbances, whereas the later responses, possibly mediated through the cortex (Matthews 1991), produce much more sophisticated responses to the perturbation. These responses can include excitation of a shortened muscle such as seen in this experiment.

Feedforward components of the muscle activation

The feedback component of the muscle activity was found to increase approximately linearly with the kinematic error on a particular trial. Due to this feature, the feedback component was estimated during learning based on a linear reconstruction using the feedback component of muscle activity on the first trial and the kinematic error of all of the trials. This allowed us to compute the feedforward component of the muscle activity and examine how it was changed as learning progressed. We hypothesized that the feedback activity was used to change the feedforward activity. To test this hypothesis we computed the cross-correlation function of the estimated change in feedforward activity with the estimated feedback activity. The change in the feedforward activity was found to depend only on the feedback activity of the preceding trial. This suggests that the feedforward motor command is modified uniquely by the error of the preceding trial.

In this analysis, there are several important assumptions. First, that the feedback is linearly dependent of the kinematic error, and secondly, that I am able to estimate the feedback component of the muscle activity simply by scaling the

kinematic error by a constant gain factor determined by the first trial in the force field. There are two problems with this approach. First, is that the behavior of the reflex responses was determined by the before effects. From postural balances studies it is apparent that the long latency reflexes can be adapted (reduced or enhanced) with practice if they do contribute to keeping the body upright (Nashner 1976). The means that that the reflex gain may have been modified during before effect trials or more possibly during learning trials, contributing to an error in the estimates. Secondly, the reflex responses increase with the underlying muscle activation (Carter et al. 1990; Kearney and Hunter 1983; Marsden et al. 1976; Stein et al. 1995). Therefore, during learning we should expect larger reflex responses as the feedforward EMG increases. Both of these will contribute to errors in the feedback estimates and, therefore, to errors in the feedforward command estimates as the feedback was assumed to increase purely with the kinematic error. If we could characterize how these factors influence the reflex response, then we could include a more realistic feedback component in our model. The reflex response could be expressed both as a function of the feedforward command on the current trial, to model the increase with muscle activation, and as a function of perturbation number (or trial number), to model the possible reflex adaptation. The dependence on activation level could be expressed as a linear function as the reflex response is linear up to about 50% maximum voluntary contraction (Carter et al. 1990; Stein et al. 1995). On the other hand, modeling of reflex adaptation could be more complex as the process would be asymptotic and might vary from condition to condition. However, if these contributing factors could be accurately incorporated into the model of feedback, the estimate of the

feedforward components would also be more accurate, which should be corroborated by higher cross-correlation values.

In order to exclude the possibility that feedback activity contaminated the estimates of feedforward activity, the feedforward activity was examined only over the earliest portion of the movement when no feedback was present. The change in the feedforward activity was compared to the size of the kinematic error on the previous trial. The feedforward component of the muscle activity increased after the muscle was stretched on the previous trial. It also increased if the muscle was shortened on the previous trial. However, if the kinematic error on the previous trial was very close to zero, then the feedforward component of muscle activation decreased on the subsequent trial.

The response to error during learning, therefore, is not a simple response to the kinematic error. Some models of adaptation propose that the internal model is updated in a direct way relative to the direction of the visual or kinematic error, or based upon the minimization of feedback torque (Bhushan and Shadmehr 1999; Donchin et al. 2003; Kawato et al. 1987; Schweighofer et al. 1998). While this may be true in visual adaptation studies, it does not appear to be the case during adaptation to novel dynamics. These algorithms are incapable of explaining adaptation to unstable dynamics (Burdet et al. 2001; Osu et al. 2003). As discussed above, the longer latency feedback components of the previous trial are related to the kinematic error. By incorporating the feedback command into the feedforward command on the subsequent trial, the central nervous system could utilize this signal to drive the learning process.

A proposed learning mechanism

Based upon the analysis of the adaptation data, we propose a model of adaptation with three major features. First, the feedback command of the previous trial is incorporated into the current feedforward command with a phase advance. Second, the feedforward co-activation signal to all muscles increases on trials following perturbation of the hand path during early learning even when muscles are not stretched. Third, the feedforward activation signal decays slowly. Together these three simple features appear to produce a powerful algorithm capable of adaptation to a variety of environmental dynamics, which may be sufficient to capture the fundamental principles of human adaptation.

This model leads to feedforward motor commands which comprise both a reciprocal component for agonist and antagonist muscle pairs (difference in muscle activation similar to net joint torque) and a co-activation component (summation of agonist and antagonist muscle activation similar to joint stiffness). Many studies have proposed this mixture of reciprocal activation and co-contraction in the motor command (Feldman 1980a; Feldman 1980b; Gribble et al. 1998; Latash 1992) and evidence exists from physiological studies (Levin et al. 1992). As we can both co-contraction and independently activate our antagonist muscle pairs, the final command to the muscles can be comprised of both co-contraction and reciprocal activation. The real question is to what extent are these commands separate in the representation of the command by the central nervous system. Some evidence suggests that these commands are represented separately (Humphrey and Reed 1983). Depending on the characteristics of the force field, the two components could be differentially affected

during adaptation to produce either a change in the endpoint force or a change in the endpoint impedance to achieve successful movements (Franklin et al. 2003a).

Online feedback is used to limit the effects of the disturbance and allow the movement to proceed to the desired location. It contains extremely useful information about the novel dynamics. The late portions of the feedback command already contain information about the temporal profile of muscle activity needed to achieve the goal of the movement (Gielen et al. 1988; Koshland and Hasan 2000). However, in order to be used as a feedforward command to counteract the disturbance of the limb it must be phase advanced on the following trial. This could be done in the cerebellum in the following manner. The sensory feedback, along with the change in the cortical commands is used to update the internal model. This error signal would result in complex spikes in the Purkinje cells produced by climbing fiber input. The phase advance could be created by the temporal specificity of the long-term depression in the parallel fiber Purkinje synapses (Chen and Thompson 1995). They found that the strongest conditioning was produced when the climbing fiber input was delayed by up to 250 ms from the firing of the parallel fibers. In the cross-correlation, the change of the feedforward command uniquely depended on the feedback command of the last trial. This is evidence that the feedback is incorporated into the subsequent movement. In addition, the profile of the feedback as a function of the kinematic error (Fig. 5.5) is similar to the shape of the late component of the feedback (Fig. 5.3) when looking at similar quantities of kinematic error.

The feedforward co-activation signal to all muscles should increase on trials following perturbation of the hand path during early learning. This hypothesis is

required, for example, to account for activation of pectoralis major up to the fifth trial in the VF despite the fact that it was not stretched (Franklin et al. 2003b). This mechanism also contributes to the co-activation of all muscles in the DF because trajectories deviate to the left on some trials and to the right on others. Gradually the co-activation builds up until stability is achieved, at which point subjects are able to achieve normal trajectories. The results of our analysis showed that muscle activity was increased on the subsequent trial even if the muscle had been shortened on the previous trial. This was seen for both the posterior deltoid and the pectoralis major. This leads to co-contraction at the beginning of learning. When both the stretched and shortened muscles increase activity, the joints become stiffer and resist any perturbing effects of the novel dynamics. Such co-contraction has been observed at the onset of learning during all studies of adaptation to novel force fields (Franklin et al. 2003b; Thoroughman and Shadmehr 1999).

Gradual deactivation of the feedforward co-activation signal has been shown during adaptation to different force fields (Franklin et al. 2003b). All muscles exhibited slow deactivation time constants throughout the learning process compared to changes in kinematic error and joint torque. Evidence for this process was also found in the trial by trial analysis of the early learning trials (Fig. 5.5). When the error was close to zero, the feedforward command was reduced on the subsequent trial. Co-contraction increases due to repeated errors early in learning. The co-contraction initially limits the trajectory errors. However, as an internal model of the dynamics is learned, the level of co-contraction will be more than necessary for the task, which is metabolically wasteful. Co-contraction is gradually reduced until a minimal level of

muscle activity is achieved. This decay in the feedforward command could be responsible for the decrease in endpoint stiffness after extensive learning (Franklin et al. 2004). The combination of increased activation of particular muscle pairs in response to muscle stretch with the gradual decay of feedforward muscle activation could be used to drive the learning of the selective control of endpoint impedance seen during adaptation to unstable dynamics (Burdet et al. 2001; Franklin et al. 2003a; Franklin et al. 2004).

With this simple control mechanism using the efferent and afferent information from the muscles as a learning signal, the CNS can quickly develop a model of the external environment. No explicit calculation of inverse dynamics or impedance needs to be performed. With antagonistic muscle activity increasing in response to disturbances, joint torques and limb impedance are modified until kinematic error is reduced to some threshold. The decrease in motor command when error is small, explains muscle deactivation during learning (Franklin et al. 2003b) such that metabolic cost is limited. This clever brain mechanism provides a very simple learning algorithm with powerful capabilities of adaptation.

An important caveat to this claim is that no attempt has been made to model changes in the temporal pattern of muscle activation. The analysis in this thesis deals purely with the integrated EMG over particular time intervals. The same integrated EMG could result in completely different movements depending on its temporal profile. Although this issue has not been addressed in this thesis, some preliminary analysis has demonstrated that the proposed learning mechanism is capable of

reproducing changes in the temporal profile of muscle activation during learning (Franklin, PhD proposal; Ng and Franklin, unpublished observations).

Another issue which has not been addressed is the role of voluntary online correction in learning. The relative contributions of feedback and voluntary activation to the learning signal have not been considered. Simulations have been carried out which demonstrate that a simple feedback signal is sufficient for the learning to occur (Franklin, Burdet, Tee, Osu, Chew, Milner and Kawato, unpublished observations). It appears that the same muscles activated by long latency reflexes are activated by later voluntary commands (Gielen et al. 1988) to correct for trajectory errors (Gielen et al. 1988; Koshland and Hasan 2000; Smeets and Erkelens 1991). Therefore, either or both of these signals could be used in updating the feedforward command.

Presumably, if the learning occurs through the long-term depression in the parallel fiber Purkinje synapses (Chen and Thompson 1995), then the temporal specificity of the long term depression may act as a filter to constrain the learning signal to particular components of the modified efferent command or afferent information at particular latencies. The strongest reinforcement occurs when the climbing fiber input is delayed by up to 250 ms from the firing of the parallel fibers (Chen and Thompson 1995). Given the delays both in neural transmission and force generation in the muscle (25-40 ms) (Ito et al. 2004; Norman and Komi 1979), combined with the fact that the onset of reflexes during disturbed movements occur 130 ms (Fig 2) to 200 ms (Shapiro et al. 2004; Shapiro et al. 2002) after movement onset, the long latency reflex component is the most likely candidate for the learning signal.

CONCLUSION

In the first experiment, I examined stiffness changes which occur after adaptation to stable and unstable interactions with the environment. The changes in endpoint stiffness after adaptation to the stable environment were well correlated with the changes in joint torques. In contrast, changes in endpoint stiffness in the unstable environment could not be explained by changes in the joint torque. Therefore in the unstable environment stiffness had been adaptively learned to stabilize the arm. These results indicate that the internal model produced during learning can comprise both an inverse dynamics model and an impedance controller.

In the second experiment, I again had subjects adapt to the same stable and unstable environments. The adaptation process itself was examined by looking at the changes in the trajectory, joint torques, and electrical activity of the muscles. In both the stable and the unstable interactions, subjects co-contracted during the early stages of the adaptation. In the stable condition, this was gradually reduced as the inverse dynamics model was learned. In the unstable condition, this was also reduced to a final level which was associated with the selective change in the endpoint stiffness reported in earlier work. Early trials in the unstable force field also indicated that the subjects were likely attempting to learn an inverse dynamics model as well as increasing their impedance. The results show that both the impedance controller and inverse dynamics models were being learnt during the early stages of learning, but that finally only one of the two was functioning after complete adaptation.

In the third experiment, I examined the changes in the endpoint impedance with different levels of instability. I found that as the instability in the x -direction increased the endpoint stiffness of the limb increased in this direction. In contrast I found no changes in the endpoint stiffness in the y -direction after adaptation. When this stiffness change was further examined I found that the combined stiffness of the arm and environment remained constant for all levels of instability with a value similar to that in the null field. This suggests that the CNS maintains a specific margin of stability. By maintaining this level of stability the CNS is able to balance the stability of the limb against metabolic cost.

The previous results suggested that the CNS learns both the inverse dynamics model (IDM) and selectively controls the endpoint impedance of the limb (impedance controller). The evidence suggests that for both stable and unstable environments the brain is attempting to learn both the IDM and impedance controller. The fourth study examined the trial by trial changes in the muscle activity during learning. The change in the feedforward command was dependent only on the feedback from the previous trial. When this effect was analyzed more closely, several features stood out. If a muscle was stretched on the previous trial, then the feedforward muscle activation was increased on the subsequent trial. This also occurred if the muscle was shortened. If the error was close to zero, then the feedforward muscle activation was actually decreased.

These results suggest a mechanism of adaptation which could explain the results of all of the studies. The feedforward motor command would be comprised of both reciprocal activation to learn the net joint torques and co-activation to learn the

necessary endpoint stiffness. The feedforward command would be increased if the muscle was unexpectedly shortened or lengthened, contributing to the learning of these two components. However, when the trajectory was close to the desired path, the motor command would be reduced to limit the metabolic cost of the movement. By combining all of these processes, the CNS can quickly develop a model of the external environment without any explicit calculations of inverse dynamics or impedance. Instead, the model in the brain is updated by efferent and afferent information on the previous trial to gradually learn the correct time dependent muscle activations to perform tasks in both stable and unstable environments.

Limitations

There are several limitations with this work. These range from the focus of the movements studied on a single direction of movement to some generalizations in the stiffness estimates. Each of these limitations will be discussed in turn.

First, only one direction of motion was examined. It is possible that we would see different effects with different directions of movement. Certainly, the adaptability of the endpoint stiffness ellipse will be highly dependent on the joint angles, particularly the elbow angle. Most force field adaptation studies have quantified learning over a larger region in the state space, using up to eight different movement directions over a greater proportion of the workspace. In this way, they were able to draw conclusions about the generalization of the learning. The present series of experiments, however, were not designed to examine generalization. Consequently, this limits the generality of our learning algorithm, which was derived on the basis of

observations related to a single movement direction. In order to generalize this model to other movement directions it would be necessary to combine it with a model of generalization such as that proposed by Shadmehr and his colleagues (Donchin et al. 2003; Donchin and Shadmehr 2002; Thoroughman and Shadmehr 2000). The advantage of using a single movement direction was that it was easier to interpret the changes in EMG from trial to trial. It also made it feasible to estimate the change in endpoint stiffness after adaptation to a particular force field as this could probably not have been completed in a single session had we attempted to measure the stiffness for several directions of motion.

Another limitation of this work is that the stiffness was measured only in the middle of the movement where the velocity was largest. In general, endpoint stiffness varies with the muscle activity and with the geometric configuration of the arm, both of which vary with time throughout the movement. Only a single estimate of stiffness was obtained and assumed to be representative of what occurred at other locations in the movement. This has implications for the conclusion that the double joint muscles are largely responsible for the selective increase in the stiffness in the unstable force field. In fact, this may only be true for the particular movement direction and configuration of the arm where the measurement was made. At other locations, other muscle pairs may be more effective at increasing the stiffness selectively in the x -direction.

The stiffness was also determined by using an average value across a range of positions. As stiffness varies with posture due to the geometric components of the stiffness matrix (Franklin and Milner 2003), the stiffness would not be constant over

the region of estimation. Similarly, only single values for change in position and endpoint force were used in estimating the stiffness. These were the average values over the region of estimation. Therefore, the stiffness is really an average value across the entire region of estimation rather than the values at a particular instant in time or workspace location. However, as we used the same estimation procedure for all force fields which were examined, and the movements were similar in all conditions, this limitation should not effect the conclusions of the study.

A related issue is that of reflex contributions to the stiffness measurement. The stiffness measurement included data recorded up to 200 ms after the onset of the perturbation. This is a sufficiently long interval that reflexes can contribute to the estimated stiffness values. Unfortunately, with the current hardware there is no way of excluding reflex contributions to the stiffness estimates, although there are techniques that could work with the appropriate hardware (Kearney et al. 1997). Clearly, the possible contribution of reflex responses to the stiffness estimates cannot be ignored. However, it is unlikely that they could have been responsible for stabilizing the arm in the unstable force fields. After learning, the hand moved smoothly to the final target, so there is no evidence that the hand was sufficiently perturbed to evoke reflex responses.

In the third study, the results are interpreted in terms of minimizing metabolic cost. This is an attractive concept which I think is important for the human body in terms of daily life. However, no measurements of metabolic cost were performed during the experiments. It is simply an interpretation of the findings in terms of what

might be a physiologically relevant quantity that the brain would be trying to minimize when it learns new tasks.

The fourth study investigated changes in the motor command with respect to the kinematic error. However the error measure which was used only quantifies errors perpendicular to the principal movement direction. If subjects were perturbed along the principal direction of movement by the force field, this error measure would not capture that error. However, movements tend to have greater variability in the direction of the movement than they do in the perpendicular direction (Messier and Kalaska 1999). This makes it difficult to estimate the actual error in this direction. Nevertheless, trajectory errors in this direction contribute to changes in motor commands but were not investigated in this thesis.

Only horizontal planar movements with a fixed wrist were studied because the robotic manipulandum is only capable of movement in this plane. This also meant that the stiffness could only be measured in two dimensions, although the actual endpoint limb stiffness would be three-dimensional. The redundant degrees of freedom of the arm were eliminated which also limits the range of strategies which the subjects could use for adaptation. For example, if the subjects were able to use all of the degrees of freedom of the shoulder, they might choose to modify their stiffness through changing the limb posture rather than through muscle co-contraction. These are issues which should be addressed in the future.

An important question is to what extent the results of these studies can be generalized to daily activities. The movements investigated in this thesis are simple point to point goal directed movements. What can studying these movements really

tell us about, for example, a hockey player's slapshot? First of all, studying these movements cannot tell us everything there is to know about such a complex movement. It does not begin to explain how a sequence of movements is coordinated to produce the shot. However, by studying a very simple movement we can perform highly detailed studies to investigate the learning mechanisms of the brain. We are able to generate virtual environments with the robotic manipulandum which none of the subjects will have experienced in real life. We are able to perturb the subject's trajectories to estimate how the limb stiffness changes after adaptation. While simplifying the task may limit what it tells us about performing complex tasks, it does provide detailed information about some important aspects of motor control.

Future Directions

Every study is a work in progress. While this thesis investigates the mechanisms of motor control, there are many important questions which it brings up. In this section I will outline a few of the important directions of this research which I intend to follow up. Some of these projects have already begun.

The next study will be to investigate the adaptation of the stiffness ellipse to other directions of instability. The unstable force fields in this study are all in the direction perpendicular to the movement. However, if we are able to selectively change the endpoint stiffness then we should expect similar changes in the orientation of the stiffness ellipse to force fields oriented in other directions. It would also be worthwhile to determine what constraints there are on the ability to selectively orient

the stiffness ellipse. In particular, is there a point at which global co-contraction replaces selective impedance control in the adaptation to environmental instability?

A related investigation is to attempt to model the intrinsic and reflex contributions to the endpoint stiffness. By using the recorded muscle activity it is possible to recreate the endpoint stiffness of the arm (Franklin and Milner 2001). This technique could be further refined to estimate reflex contributions to selective changes in the endpoint stiffness.

I plan to further investigate the trial by trial changes in muscle activity during adaptation. Whereas the current study pools together the adaptation to several different force fields from several different days, I would like to conduct a study in which the subject adapts to continually changing force fields. By modulating the direction and the gain of the force field similar to other studies (Scheidt et al. 2001; Takahashi et al. 2001), a more accurate estimate of the changes in the feedforward command could be achieved. I would like to better understand how changes in muscle activation are related to particular errors. It would also be interesting to compare the adaptation when muscle proprioception is the source of error feedback to that when vision is the source of error feedback. Using a device like the KINARM, which allows perturbations to be applied selectively to either the shoulder or the elbow joints, we could also examine possible changes in feedforward muscle activation that might be mediated through heteronymous reflex connections.

Another project is to examine how the proposed learning mechanism predicts changes in the temporal profile of muscle activity. In this thesis, a single integrated value of muscle activity was examined. However, in controlling the limbs,

it is critical that the correct temporal pattern of muscle activation is learned. Although this has not been studied much to this point, preliminary results indicate that the learning mechanism appropriately learns tri-phasic patterns of muscle activity during fast targeted movements which are modified in timing and size similar to those of human movements (Ng and Franklin, unpublished observations).

Another direction is to investigate the brain activity during the learning of stable and unstable dynamics using functional magnetic resonance imaging (fMRI). If we can develop a robotic manipulandum which is capable of being used in these environments we can begin to explore the neural substrates for adaptation. Of particular interest is comparing adaptation to stable and unstable environments. If the same internal model is used for both of these types of dynamics then similar areas of the brain may be involved in the learning and production of movements in both environments. However, it is possible that different areas of the brain are involved in the adaptation to unstable dynamics.

REFERENCES

- Akaike, H. A new look at the statistical model identification. *IEEE Trans Automat Control AC-19*: 716-723, 1974.
- Akazawa, K., Milner, T. E., and Stein, R. B. Modulation of reflex EMG and stiffness in response to stretch of human finger muscle. *J Neurophysiol* 49: 16-27, 1983.
- Angel, R. W. Antagonist muscle activity during rapid arm movements: central versus proprioceptive influences. *J Neurol Neurosurg Psychiatry* 40: 683-6, 1977.
- Atkeson, C. G. and Hollerbach, J. M. Kinematic features of unrestrained vertical arm movements. *J Neurosci* 5: 2318-30, 1985.
- Bennett, D. J. Torques generated at the human elbow joint in response to constant position errors imposed during voluntary movements. *Exp Brain Res* 95: 488-98, 1993.
- Berardelli, A., Rothwell, J. C., Day, B. L., Kachi, T., and Marsden, C. D. Duration of the first agonist EMG burst in ballistic arm movements. *Brain Res* 304: 183-7, 1984.
- Bhushan, N. and Shadmehr, R. Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol Cybern* 81: 39-60, 1999.
- Brashers-Krug, T., Shadmehr, R., and Bizzi, E. Consolidation in human motor memory. *Nature* 382: 252-5, 1996.
- Brown, S. H. and Cooke, J. D. Amplitude- and instruction-dependent modulation of movement-related electromyogram activity in humans. *J Physiol* 316: 97-107, 1981a.
- Brown, S. H. and Cooke, J. D. Responses to force perturbations preceding voluntary human arm movements. *Brain Res* 220: 350-5, 1981b.
- Brown, S. H. and Cooke, J. D. Initial agonist burst duration depends on movement amplitude. *Exp Brain Res* 55: 523-7, 1984.
- Brown, S. H. and Cooke, J. D. Movement-related phasic muscle activation. I. Relations with temporal profile of movement. *J Neurophysiol* 63: 455-64, 1990.

- Buchanan, T. S., Almdale, D. P., Lewis, J. L., and Rymer, W. Z. Characteristics of synergic relations during isometric contractions of human elbow muscles. *J Neurophysiol* 56: 1225-41, 1986.
- Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414: 446-9, 2001.
- Burdet, E., Osu, R., Franklin, D. W., Yoshioka, T., Milner, T. E., and Kawato, M. A method for measuring endpoint stiffness during multi-joint arm movements. *J Biomech* 33: 1705-9, 2000.
- Burdet, E., Tee, K. P., Chew, C. M., Franklin, D. W., Osu, R., Kawato, M., and Milner, T. E. Stability and learning in human arm movements. Proceedings of the 2001 International Conference on Computational Intelligence, Robotics and Autonomous Systems. 2001.
- Cannon, S. C. and Zahalak, G. I. The mechanical behavior of active human skeletal muscle in small oscillations. *J Biomech* 15: 111-21, 1982.
- Capaday, C. and Cooke, J. D. Vibration-induced changes in movement-related EMG activity in humans. *Exp Brain Res* 52: 139-46, 1983.
- Capaday, C., Forget, R., and Milner, T. A re-examination of the effects of instruction on the long-latency stretch reflex response of the flexor pollicis longus muscle. *Exp Brain Res* 100: 515-21, 1994.
- Carter, R. R., Crago, P. E., and Gorman, P. H. Nonlinear stretch reflex interaction during cocontraction. *J Neurophysiol* 69: 943-52, 1993.
- Carter, R. R., Crago, P. E., and Keith, M. W. Stiffness regulation by reflex action in the normal human hand. *J Neurophysiol* 64: 105-18, 1990.
- Chen, C. and Thompson, R. F. Temporal specificity of long-term depression in parallel fiber--Purkinje synapses in rat cerebellar slice. *Learn Mem* 2: 185-98, 1995.
- Clancy, E. A. and Hogan, N. Multiple site electromyograph amplitude estimation. *IEEE Trans Biomed Eng* 42: 203-11, 1995.
- Conditt, M. A., Gandolfo, F., and Mussa-Ivaldi, F. A. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78: 554-60, 1997.
- Conditt, M. A. and Mussa-Ivaldi, F. A. Central representation of time during motor learning. *Proc Natl Acad Sci U S A* 96: 11625-30, 1999.

- Cooke, J. D. The role of stretch reflexes during active movements. *Brain Res* 181: 493-7, 1980.
- Cooke, J. D. and Brown, S. H. Movement-related phasic muscle activation. II. Generation and functional role of the triphasic pattern. *J Neurophysiol* 63: 465-72, 1990.
- Cooke, J. D. and Brown, S. H. Movement-related phasic muscle activation. III. The duration of phasic agonist activity initiating movement. *Exp Brain Res* 99: 473-82, 1994.
- Darling, W. G. and Cooke, J. D. Changes in the variability of movement trajectories with practice. *J Mot Behav* 19: 291-310, 1987a.
- Darling, W. G. and Cooke, J. D. Movement Related EMGs Become More Variable During Learning of Fast Accurate Movements. *J Mot Behav* 19: 311-31, 1987b.
- De Serres, S. J. and Milner, T. E. Wrist muscle activation patterns and stiffness associated with stable and unstable mechanical loads. *Exp Brain Res* 86: 451-8, 1991.
- Doemges, F. and Rack, P. M. Changes in the stretch reflex of the human first dorsal interosseous muscle during different tasks. *J Physiol* 447: 563-73, 1992a.
- Doemges, F. and Rack, P. M. Task-dependent changes in the response of human wrist joints to mechanical disturbance. *J Physiol* 447: 575-85, 1992b.
- Dolan, J. M., Freidman, M. B., and Nagurka, M. L. Dynamic and loaded impedance components in the maintenance of human arm posture. *IEEE Trans Systems Man Cybern* 23: 698-709, 1993.
- Donchin, O., Francis, J. T., and Shadmehr, R. Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J Neurosci* 23: 9032-45, 2003.
- Donchin, O. and Shadmehr, R. Linking motor learning to function approximation: learning in an unlearnable force field. Cambridge, MA, MIT Press. 2002, 197-203.
- Dufresne, J. R., Soechting, J. F., and Terzuolo, C. A. Electromyographic response to pseudo-random torque disturbances of human forearm position. *Neuroscience* 3: 1213-26, 1978.
- Feldman, A. G. Superposition of motor programs--I. Rhythmic forearm movements in man. *Neuroscience* 5: 81-90, 1980a.

- Feldman, A. G. Superposition of motor programs--II. Rapid forearm flexion in man. *Neuroscience* 5: 91-5, 1980b.
- Flanagan, J. R., King, S., Wolpert, D. M., and Johansson, R. S. Sensorimotor prediction and memory in object manipulation. *Can J Exp Psychol* 55: 87-95, 2001.
- Flanagan, J. R., Nakano, E., Imamizu, H., Osu, R., Yoshioka, T., and Kawato, M. Composition and decomposition of internal models in motor learning under altered kinematic and dynamic environments. *J Neurosci* 19: RC34, 1999.
- Flanagan, J. R. and Wing, A. M. The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J Neurosci* 17: 1519-28, 1997.
- Flanders, M. and Soechting, J. F. Arm muscle activation for static forces in three-dimensional space. *J Neurophysiol* 64: 1818-37, 1990.
- Flash, T. The control of hand equilibrium trajectories in multi-joint arm movements. *Biol Cybern* 57: 257-74, 1987.
- Flash, T. and Mussa-Ivaldi, F. Human arm stiffness characteristics during the maintenance of posture. *Exp Brain Res* 82: 315-26, 1990.
- Foley, J. M. and Meyer, R. A. Energy cost of twitch and tetanic contractions of rat muscle estimated in situ by gated ³¹P NMR. *NMR Biomed* 6: 32-8, 1993.
- Franklin, D. W., Burdet, E., Osu, R., Kawato, M., and Milner, T. E. Functional significance of stiffness in adaptation of multijoint arm movements to stable and unstable dynamics. *Exp Brain Res* 151: 145-57, 2003a.
- Franklin, D.W. and Milner, T.E. Using invariant features of voluntary movements to predict hand trajectories, *Society for Neuroscience Abstracts*. 812.17, 1997.
- Franklin, D. W. and Milner, T. E. A technique for estimating endpoint stiffness from electromyographic activity during multijoint movements. *The International Symposium on Movement, Analysis and Modeling of Human Functions*. IMEKO. 334-339, 2001.
- Franklin, D. W. and Milner, T. E. Adaptive control of stiffness to stabilize hand position with large loads. *Exp Brain Res* 152: 211-20, 2003.
- Franklin, D. W., Osu, R., Burdet, E., Kawato, M., and Milner, T. E. Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J Neurophysiol* 90: 3270-82, 2003b.
- Franklin, D. W., So, U., Kawato, M., and Milner, T. E. Impedance Control Balances Stability with Metabolically Costly Muscle Activation. *J Neurophysiol* 2004.

- Gandolfo, F., Mussa-Ivaldi, F. A., and Bizzi, E. Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843-6, 1996.
- Gasser, H. S. and Hill, A. V. Dynamics of muscular contraction. *Proc R Soc Lond Ser B* 96: 398-437, 1924.
- Ghahramani, Z., Wolpert, D. M., and Jordan, M. I. Generalization to local remappings of the visuomotor coordinate transformation. *J Neurosci* 16: 7085-96, 1996.
- Ghez, C. and Martin, J. H. The control of rapid limb movement in the cat. III. Agonist - antagonist coupling. *Exp Brain Res* 45: 115-25, 1982.
- Gielen, C. C., Ramaekers, L., and van Zuylen, E. J. Long-latency stretch reflexes as co-ordinated functional responses in man. *J Physiol* 407: 275-92, 1988.
- Gomi, H. and Kawato Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272: 117-20, 1996.
- Gomi, H. and Kawato, M. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biol Cybern* 76: 163-71, 1997.
- Gomi, H. and Osu, R. Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *J Neurosci* 18: 8965-78, 1998.
- Goodbody, S. J. and Wolpert, D. M. Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825-38, 1998.
- Gottlieb, G. L. and Agarwal, G. C. Compliance of single joints: elastic and plastic characteristics. *J Neurophysiol* 59: 937-51, 1988.
- Grey, M. J. Viscoelasticity of the human wrist during the stabilization phase of a targeted movement. Msc. Thesis, Burnaby, Canada, Simon Fraser University. 1997.
- Gribble, P. L., Mullin, L. I., Cothros, N., and Mattar, A. Role of cocontraction in arm movement accuracy. *J Neurophysiol* 89: 2396-405, 2003.
- Gribble, P. L. and Ostry, D. J. Independent coactivation of shoulder and elbow muscles. *Exp Brain Res* 123: 355-60, 1998.
- Gribble, P. L. and Ostry, D. J. Compensation for interaction torques during single- and multijoint limb movement. *J Neurophysiol* 82: 2310-26, 1999.
- Gribble, P. L. and Ostry, D. J. Compensation for loads during arm movements using equilibrium-point control. *Exp Brain Res* 135: 474-82, 2000.

- Gribble, P. L., Ostry, D. J., Sanguineti, V., and Laboissiere, R. Are complex control signals required for human arm movement? *J Neurophysiol* 79: 1409-24, 1998.
- Hallett, M. and Marsden, C. D. Ballistic flexion movements of the human thumb. *J Physiol* 294: 33-50, 1979.
- Hallett, M., Shahani, B. T., and Young, R. R. EMG analysis of stereotyped voluntary movements in man. *J Neurol Neurosurg Psychiatry* 38: 1154-62, 1975.
- Hannaford, B. and Stark, L. Roles of the elements of the triphasic control signal. *Exp Neurol* 90: 619-34, 1985.
- Harris, C. M. and Wolpert, D. M. Signal-dependent noise determines motor planning. *Nature* 394: 780-4, 1998.
- Hogan, M. C., Ingham, E., and Kurdak, S. S. Contraction duration affects metabolic energy cost and fatigue in skeletal muscle. *Am J Physiol* 274: E397-402, 1998.
- Hogan, M. C., Kurdak, S. S., and Arthur, P. G. Effect of gradual reduction in O₂ delivery on intracellular homeostasis in contracting skeletal muscle. *J Appl Physiol* 80: 1313-21, 1996.
- Hogan, N. Adaptive control of mechanical impedance by coactivation of antagonist muscles. *IEEE Trans Auto Control* AC-29: 681-90, 1984.
- Hogan, N. The mechanics of multi-joint posture and movement control. *Biol Cybern* 52: 315-31, 1985.
- Hollerbach, M. J. and Flash, T. Dynamic interactions between limb segments during planar arm movement. *Biol Cybern* 44: 67-77, 1982.
- Humphrey, D. R. and Reed, D. J. Separate cortical systems for control of joint movement and joint stiffness: reciprocal activation and coactivation of antagonist muscles. *Adv Neurol* 39: 347-72, 1983.
- Hunter, I. W. and Kearney, R. E. Dynamics of human ankle stiffness: variation with mean ankle torque. *J Biomech* 15: 747-52, 1982.
- Huxley, A. F. and Simmons, R. M. Proposed mechanism of force generation in striated muscle. *Nature* 233: 533-8, 1971.
- Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Putz, B., Yoshioka, T., and Kawato, M. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403: 192-5, 2000.

- Ito, T., Murano, E. Z., and Gomi, H. Fast force-generation dynamics of human articular muscles. *J Appl Physiol* 96: 2318-24; discussion 2317, 2004.
- Jones, K. E., Hamilton, A. F., and Wolpert, D. M. Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88: 1533-44, 2002.
- Jones, L. A. and Hunter, I. W. A perceptual analysis of stiffness. *Exp Brain Res* 79: 150-6, 1990.
- Jones, L. A. and Hunter, I. W. Changes in pinch force with bidirectional load forces. *J Mot Behav* 24: 157-64, 1992.
- Jones, L. A. and Hunter, I. W. A perceptual analysis of viscosity. *Exp Brain Res* 94: 343-51, 1993.
- Jones, L. A., Hunter, I. W., and Irwin, R. J. Differential thresholds for limb movement measured using adaptive techniques. *Percept Psychophys* 52: 529-35, 1992.
- Karniel, A. and Mussa-Ivaldi, F. A. Sequence, time, or state representation: how does the motor control system adapt to variable environments? *Biol Cybern* 89: 10-21, 2003.
- Karst, G. M. and Hasan, Z. Timing and magnitude of electromyographic activity for two-joint arm movements in different directions. *J Neurophysiol* 66: 1594-604, 1991.
- Katayama, M., Inoue, S., and Kawato, M. A strategy of motor learning using adjustable parameters for arm movement. *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 2370-2373. 98.
- Katayama, M. and Kawato, M. Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. *Biol Cybern* 69: 353-62, 1993.
- Kawato, M. Internal models for motor control and trajectory planning. *Curr Opin Neurobiol* 9: 718-27, 1999.
- Kawato, M., Furukawa, K., and Suzuki, R. A hierarchical neural-network model for control and learning of voluntary movement. *Biol Cybern* 57: 169-85, 1987.
- Kearney, R. E. and Hunter, I. W. Dynamics of human ankle stiffness: variation with displacement amplitude. *J Biomech* 15: 753-6, 1982.
- Kearney, R. E. and Hunter, I. W. System identification of human triceps surae stretch reflex dynamics. *Exp Brain Res* 51: 117-27, 1983.

- Kearney, R. E. and Hunter, I. W. System identification of human joint dynamics. *Crit Rev Biomed Eng* 18: 55-87, 1990.
- Kearney, R. E., Stein, R. B., and Parameswaran, L. Identification of intrinsic and reflex contributions to human ankle stiffness dynamics. *IEEE Trans Biomed Eng* 44: 493-504, 1997.
- Kirsch, R. F., Boskov, D., and Rymer, W. Z. Muscle stiffness during transient and continuous movements of cat muscle: perturbation characteristics and physiological relevance. *IEEE Trans Biomed Eng* 41: 758-70, 1994.
- Koshland, G. F. and Hasan, Z. Electromyographic responses to a mechanical perturbation applied during impending arm movements in different directions: one-joint and two-joint conditions. *Exp Brain Res* 132: 485-99, 2000.
- Krakauer, J. W., Ghilardi, M. F., and Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* 2: 1026-31, 1999.
- Lackner, J. R. and Dizio, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J Neurophysiol* 72: 299-313, 1994.
- Lacquaniti, F., Borghese, N. A., and Carrozzo, M. Internal models of limb geometry in the control of hand compliance. *J Neurosci* 12: 1750-62, 1992.
- Lacquaniti, F., Carrozzo, M., and Borghese, N. A. Time-varying mechanical behavior of multijointed arm in man. *J Neurophysiol* 69: 1443-64, 1993.
- Lacquaniti, F. and Soechting, J. F. Behavior of the stretch reflex in a multi-jointed limb. *Brain Res* 311: 161-6, 1984.
- Lacquaniti, F. and Soechting, J. F. EMG responses to load perturbations of the upper limb: effect of dynamic coupling between shoulder and elbow motion. *Exp Brain Res* 61: 482-96, 1986a.
- Lacquaniti, F. and Soechting, J. F. Responses of mono- and bi-articular muscles to load perturbations of the human arm. *Exp Brain Res* 65: 135-44, 1986b.
- Latash, M. L. Independent control of joint stiffness in the framework of the equilibrium-point hypothesis. *Biol Cybern* 67: 377-84, 1992.
- Lee, R. G. and Tatton, W. G. Long latency reflexes to imposed displacements of the human wrist: dependence on duration of movement. *Exp Brain Res* 45: 207-16, 1982.
- Levin, M. F., Feldman, A. G., Milner, T. E., and Lamarre, Y. Reciprocal and coactivation commands for fast wrist movements. *Exp Brain Res* 89: 669-77, 1992.

- Li, C. S., Padoa-Schioppa, C., and Bizzi, E. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 30: 593-607, 2001.
- MacKay, W. A., Crammond, D. J., Kwan, H. C., and Murphy, J. T. Measurements of human forearm viscoelasticity. *J Biomech* 19: 231-8, 1986.
- Mah, C. D. and Mussa-Ivaldi, F. A. Generalization of object manipulation skills learned without limb motion. *J Neurosci* 23: 4821-5, 2003.
- Marsden, C. D., Merton, P. A., and Morton, H. B. Stretch reflex and servo action in a variety of human muscles. *J Physiol* 259: 531-60, 1976.
- Matthews, P. B. Observations on the automatic compensation of reflex gain on varying the pre-existing level of motor discharge in man. *J Physiol* 374: 73-90, 1986.
- Matthews, P. B. The human stretch reflex and the motor cortex. *Trends Neurosci* 14: 87-91, 1991.
- McIntyre, J., Mussa-Ivaldi, F. A., and Bizzi, E. The control of stable postures in the multijoint arm. *Exp Brain Res* 110: 248-64, 1996.
- McIntyre, J., Zago, M., Berthoz, A., and Lacquaniti, F. Does the brain model Newton's laws? *Nat Neurosci* 4: 693-4, 2001.
- Messier, J. and Kalaska, J. F. Comparison of variability of initial kinematics and endpoints of reaching movements. *Exp Brain Res* 125: 139-52, 1999.
- Milner, T. E. Adaptation to destabilizing dynamics by means of muscle cocontraction. *Exp Brain Res* 143: 406-16, 2002a.
- Milner, T. E. Contribution of geometry and joint stiffness to mechanical stability of the human arm. *Exp Brain Res* 143: 515-9, 2002b.
- Milner, T. E. and Cloutier, C. Compensation for mechanically unstable loading in voluntary wrist movement. *Exp Brain Res* 94: 522-32, 1993.
- Milner, T. E. and Cloutier, C. Damping of the wrist joint during voluntary movement. *Exp Brain Res* 122: 309-17, 1998.
- Milner, T. E., Cloutier, C., Leger, A. B., and Franklin, D. W. Inability to activate muscles maximally during cocontraction and the effect on joint stiffness. *Exp Brain Res* 107: 293-305, 1995.
- Milner, T. E. and Franklin, D. W. Characterization of multijoint finger stiffness: dependence on finger posture and force direction. *IEEE Trans Biomed Eng* 45: 1363-75, 1998.

- Morishige, K., Miyamoto, H., Osu, R., and Kawato, M. Positional variance on via-point reaching movement supports sequential trajectory planning and execution model. *IEICE Trans Inf & Syst (Japanese Edition)* in press.
- Mussa-Ivaldi, F. A., Hogan, N., and Bizzi, E. Neural, mechanical, and geometric factors subserving arm posture in humans. *J Neurosci* 5: 2732-43, 1985.
- Mustard, B. E. and Lee, R. G. Relationship between EMG patterns and kinematic properties for flexion movements at the human wrist. *Exp Brain Res* 66: 247-56, 1987.
- Nashner, L. M. Adapting reflexes controlling the human posture. *Exp Brain Res* 26: 59-72, 1976.
- Norman, R. W. and Komi, P. V. Electromechanical delay in skeletal muscle under normal movement conditions. *Acta Physiol Scand* 106: 241-8, 1979.
- Osu, R., Burdet, E., Franklin, D. W., Milner, T. E., and Kawato, M. Different mechanisms involved in adaptation to stable and unstable dynamics. *J Neurophysiol* 90: 3255-69, 2003.
- Osu, R., Franklin, D. W., Kato, H., Gomi, H., Domen, K., Yoshioka, T., and Kawato, M. Short- and Long-Term Changes in Joint Co-Contraction Associated With Motor Learning as Revealed From Surface EMG. *J Neurophysiol* 88: 991-1004, 2002.
- Osu, R. and Gomi, H. Multijoint muscle regulation mechanisms examined by measured human arm stiffness and EMG signals. *J Neurophysiol* 81: 1458-68, 1999.
- Perreault, E. J., Kirsch, R. F., and Crago, P. E. Effects of voluntary force generation on the elastic components of endpoint stiffness. *Exp Brain Res* 141: 312-23, 2001.
- Perreault, E. J., Kirsch, R. F., and Crago, P. E. Voluntary control of static endpoint stiffness during force regulation tasks. *J Neurophysiol* 87: 2808-16, 2002.
- Perreault, E. J., Kirsch, R. F., and Crago, P. E. Multijoint dynamics and postural stability of the human arm. *Exp Brain Res* 157: 507-17, 2004.
- Person, R. S. [Electromyographic investigations of coordination of the antagonistic muscles in development of motor habit.]. *Zh Vyssh Nerv Deiat Im I P Pavlova* 8: 17-27, 1958.
- Pollick, F. E. and Ishimura, G. The three-dimensional curvature of straight-ahead movements. *J Motor Behav* 28: 271-279, 1996.

- Rack, P. M. Limitations of somatosensory feedback in control of posture and movement. In Brooks, V. B. ed. *Motor Control. Handbook of Physiology, Sect. 1. The Nervous System.* American Physiological Society. 1981, 229-256.
- Rack, P. M. and Westbury, D. R. The effects of length and stimulus rate on tension in the isometric cat soleus muscle. *J Physiol* 204: 443-60, 1969.
- Rancourt, D. and Hogan, N. Stability in force-production tasks. *J Mot Behav* 33: 193-204, 2001.
- Rothwell, J. C., Traub, M. M., Day, B. L., Obeso, J. A., Thomas, P. K., and Marsden, C. D. Manual motor performance in a deafferented man. *Brain* 105 (Pt 3): 515-42, 1982.
- Russ, D. W., Elliott, M. A., Vandenborne, K., Walter, G. A., and Binder-Macleod, S. A. Metabolic costs of isometric force generation and maintenance of human skeletal muscle. *Am J Physiol Endocrinol Metab* 282: E448-57, 2002.
- Sainburg, R. L., Ghilardi, M. F., Poizner, H., and Ghez, C. Control of limb dynamics in normal subjects and patients without proprioception. *J Neurophysiol* 73: 820-35, 1995.
- Sainburg, R. L., Poizner, H., and Ghez, C. Loss of proprioception produces deficits in interjoint coordination. *J Neurophysiol* 70: 2136-47, 1993.
- Sakamoto, Y., Ishiguro, M., and Kitagawa, G. *Akaike Information Criterion Statistics.* Dordrecht, Kluwer Academic Publishers. 1986.
- Sanger, T. D. Neural network learning control of robotic manipulators using gradually increasing task difficulty. *IEEE Trans Robotics Automat* 10: 323-333, 1994.
- Sanner, R. M. and Kosha, M. A mathematical model of the adaptive control of human arm motions. *Biol Cybern* 80: 369-82, 1999.
- Scheidt, R. A., Dingwell, J. B., and Mussa-Ivaldi, F. A. Learning to move amid uncertainty. *J Neurophysiol* 86: 971-85, 2001.
- Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z., and Mussa-Ivaldi, F. A. Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84: 853-62, 2000.
- Schmidt, R. A., Zelaznik, H., Hawkins, B., Frank, J. S., and Quinn, J. T. Jr Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol Rev* 47: 415-51, 1979.

- Schweighofer, N., Arbib, M. A., and Kawato, M. Role of the cerebellum in reaching movements in humans. I. Distributed inverse dynamics control. *Eur J Neurosci* 10: 86-94, 1998.
- Scott, S.H. Neural correlates of limb mechanics and mechanical loads in primary motor cortex. In: *Proceedings of the 31st NIPS International Symposium*, Okazaki, Japan: NIPS, 2004, p.36
- Shadmehr, R. and Brashers-Krug, T. Functional stages in the formation of human long-term motor memory. *J Neurosci* 17: 409-19, 1997.
- Shadmehr, R. and Holcomb, H. H. Neural correlates of motor memory consolidation. *Science* 277: 821-5, 1997.
- Shadmehr, R. and Moussavi, Z. M. Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20: 7807-15, 2000.
- Shadmehr, R. and Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208-24, 1994.
- Shadmehr, R., Mussa-Ivaldi, F. A., and Bizzi, E. Postural force fields of the human arm and their role in generating multijoint movements. *J Neurosci* 13: 45-62, 1993.
- Shapiro, M. B., Gottlieb, G. L., and Corcos, D. M. EMG responses to an unexpected load in fast movements are delayed with an increase in the expected movement time. *J Neurophysiol* 91: 2135-47, 2004.
- Shapiro, M. B., Gottlieb, G. L., Moore, C. G., and Corcos, D. M. Electromyographic responses to an unexpected load in fast voluntary movements: descending regulation of segmental reflexes. *J Neurophysiol* 88: 1059-63, 2002.
- Sih, B. L. and Stuhmiller, J. H. The metabolic cost of force generation. *Med Sci Sports Exerc* 35: 623-9, 2003.
- Sinkjaer, T. and Hayashi, R. Regulation of wrist stiffness by the stretch reflex. *J Biomech* 22: 1133-40, 1989.
- Slifkin, A. B. and Newell, K. M. Noise, information transmission, and force variability. *J Exp Psychol Hum Percept Perform* 25: 837-51, 1999.
- Smeets, J. B. and Erkelens, C. J. Dependence of autogenic and heterogenic stretch reflexes on pre-load activity in the human arm. *J Physiol* 440: 455-65, 1991.
- Soechting, J. F. and Lacquaniti, F. Quantitative evaluation of the electromyographic responses to multidirectional load perturbations of the human arm. *J Neurophysiol* 59: 1296-313, 1988.

- Stein, R. B., Hunter, I. W., Lafontaine, S. R., and Jones, L. A. Analysis of short-latency reflexes in human elbow flexor muscles. *J Neurophysiol* 73: 1900-11, 1995.
- Stein, R. B. and Kearney, R. E. Nonlinear behavior of muscle reflexes at the human ankle joint. *J Neurophysiol* 73: 65-72, 1995.
- Takahashi, C. D., Scheidt, R. A., and Reinkensmeyer, D. J. Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J Neurophysiol* 86: 1047-51, 2001.
- Tax, A. A., Denier van der Gon, J. J., and Erkelens, C. J. Differences in coordination of elbow flexor muscles in force tasks and in movement tasks. *Exp Brain Res* 81: 567-72, 1990a.
- Tax, A. A., Denier van der Gon, J. J., Gielen, C. C., and Kleyne, M. Differences in central control of m. biceps brachii in movement tasks and force tasks. *Exp Brain Res* 79: 138-42, 1990b.
- Tax, A. A., Denier van der Gon, J. J., Gielen, C. C., and van den Tempel, C. M. Differences in the activation of m. biceps brachii in the control of slow isotonic movements and isometric contractions. *Exp Brain Res* 76: 55-63, 1989.
- Thoroughman, K. A. and Shadmehr, R. Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19: 8573-88, 1999.
- Thoroughman, K. A. and Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* 407: 742-7, 2000.
- Ts'o, D. Y. and Roe, A. W. Functional compartments in visual cortex: segregation and interaction. In Gazzaniga, M. S. ed. *The Cognitive Neurosciences*. Cambridge, MIT Press. 1994.
- Tsuji, T., Morasso, P. G., Goto, K., and Ito, K. Human hand impedance characteristics during maintained posture. *Biol Cybern* 72: 475-85, 1995.
- Van Beers, R. J., Haggard, P., and Wolpert, D. M. The role of execution noise in movement variability. *J Neurophysiol* 91: 1050-63, 2004.
- van Galen, G. P. and van Huygevoort, M. Error, stress and the role of neuromotor noise in space oriented behaviour. *Biol Psychol* 51: 151-71, 2000.
- van Groeningen, C. J. and Erkelens, C. J. Task-dependent differences between mono- and bi-articular heads of the triceps brachii muscle. *Exp Brain Res* 100: 345-52, 1994.

- Wadman, W. J., Denier van der Gon, J. J., and Derksen, R. J. A. Muscle activation patterns for fast goal-directed arm movements. *J Hum Move Stud* 6: 19-37, 1980.
- Wang, T., Dordevic, G. S., and Shadmehr, R. Learning the dynamics of reaching movements results in the modification of arm impedance and long-latency perturbation responses. *Biol Cybern* 85: 437-48, 2001.
- Weiss, P. L., Hunter, I. W., and Kearney, R. E. Human ankle joint stiffness over the full range of muscle activation levels. *J Biomech* 21: 539-44, 1988.
- Winter, D. A. *Biomechanics and Motor Control of Human Movement*. New York, John Wiley & Sons. 1990.