ADAPTIVE CONTROL OF GOAL-ORIENTED HUMAN ARM MOVEMENT

by

Mark R. Hinder
M.Sci. (Hons) University of Birmingham, UK, 1999

THESIS SUBMITTED IN PARTIAL FULFILMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE

in the
School of Kinesiology

SIMON FRASER UNIVERSITY
June 2003

Approval

Name: **Mark Hinder**

Degree: Master of Science

Title of Thesis: Adaptive Control of Goal-Oriented Human Arm Movement

Examining Committee:

Chair: **Dr. David Goodman**
Professor

Dr. Ted Milner
Senior Supervisor
Professor, School of Kinesiology

Dr. Steve Robinovitch
Supervisor
Assistant Professor, School of Kinesiology

Dr. Antony Hodgson
External Examiner
Associate Professor, Mechanical Engineering
University of British Columbia

Date Approved: 26th June 2003

PARTIAL COPYRIGHT LICENCE

Title of Thesis/Project/Extended Essay

ADAPTIVE CONTROL OF GOAL-ORIENTED HUMAN ARM MOVEMENT

Author:

(signature)
Mark Hinder

(name)
12 August 2003

(date)

Abstract

Experiment 1 directly tested a fundamental prediction of the equilibrium point hypothesis (EPH) of movement control, namely, equifinality. Subjects learned wrist flexion movements while a motor provided assistance in proportion to the instantaneous velocity. Subjects stopped short of the target on the trials where the magnitude of the assistance was randomly decreased, i.e., equifinality was not achieved. This is contrary to the EPH, although such effects are entirely consistent with predictions based on the formation of an internal dynamics model.

In experiment 2 subjects performed 0.25 m horizontal reaches while exposed to a position-dependent parabolic field (PF), located in a region extending to 0.1 m from the start position. Following extensive learning of the PF, the field was randomly perturbed to 2/3 or 3/2 times the original strength for 5 trials, a total of 24 times. Mechanical channel trials were used to test for evidence of internal model formation. Internal model formation was evident as early as the second or third PF trial. Subjects produced lateral force to counteract the field, although it did not perfectly match the PF force, resulting in bowed trajectories. As learning progressed, we found significant decreases in the EMG in all 6 shoulder and elbow muscles once subjects had passed through the force field. We suggest subjects modified their internal model to more accurately compensate for the PF, such that the cocontraction could be reduced.

Subjects quickly readjusted to the PF following each ΔPF set, suggesting retention of the motor commands required in the PF. On the first ΔPF trial of each set, subjects were displaced in the direction of the change in force. On the second trial, trajectories were substantially straighter, but the displacement was not reduced significantly more in the 3^{rd}-5^{th} trials. Subjects adapted the internal model to more accurately compensate for higher strength ΔPF, while simultaneously employing higher levels of cocontraction. In the lower strength ΔPF results suggest a reduction in the level of cocontraction. We found no consolidation of learning between ΔPF sets, i.e., performance in the ΔPF did not improve as subjects were exposed to more sets.

Acknowledgements

I would like to thank Dr. Ted Milner for his encouragement and attention to detail, ensuring that research is completed to the highest scientific standards. I am indebted to Mitsuo Kawato and the Computational Neuroscience Laboratories at Advance Telecommunications Research International, Kyoto, Japan, for allowing me the opportunity to work as an intern researcher at ATR, where I was able to plan and conduct experiment 2. I would like to thank Dave Franklin for assistance in conducting experiments at ATR, and providing an earlier version of Figure 8. I also thank Mitsuo Kawato for critiquing an earlier version of the text relating to experiment 1. Thanks to Amanda for her great friendship, encouragement and understanding. Thanks to my parents and brother for support from afar, and for coming to visit me here in Canada (what a good excuse for a holiday!).

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada, SFU graduate scholarships and the Gordon Diewert Memorial Scholarship.

Table of Contents

List Of Figures

List of Tables

Specific Aims of the Research

Research in the area of motor control, specifically investigating how humans learn to make accurate, goal-oriented movements has progressed a long way in the last 30 years, although many questions, as yet, remain unanswered. In this study, we aim to determine the mechanisms used to control voluntary arm movement, and how these are adapted during learning in a novel environment.

Research has focused on three possible mechanisms that could be utilised in the CNS. An internal model (IM) (Shadmehr and Mussi-Ivaldi, 1994; Lackner and Dizio, 1994; Gomi and Kawato, 1996; Gandolfo et al, 1996; Goodbody and Wolpert, 1998; Throughman and Shadmehr, 2000; Scheidt et al, 2000, 2001; Takahashi et al 2001) is a transformation of a particular dynamic process, which either occurs in a logical (forward), or inverse, direction. Forward models represent transformations from cause to effect while inverse models represent the reverse transformation. An IM uses a complex set of calculations requiring information on the desired task trajectory and information stored in the brain about properties of the limbs and environment dynamics.

An alternative control theory is the Equilibrium Point Hypotheses (EPH). The EPH is predicated on the fact that the length and stiffness of muscles can be controlled centrally by the central nervous system (CNS), establishing a balance of force in the springs. (Feldman 1966a, b; Feldman and Astrayan, 1965). Impedance Control, which involves specifically controlling the stiffness of individual muscles, was proposed as an optimal way for the CNS to compensate for unstable environmental forces (Hogan, 1985). By increasing the stiffness of the human limb in specific directions, the displacement due to a predictable perturbing force could be reduced along the force direction without requiring higher stiffness in other directions.

In this study we investigate how the CNS controls human goal-oriented movement of upper limbs. Subjects are asked to learn specific single and multi-joint movement tasks whilst exposed to forces of different natures. Using a well-known technique of

unexpectedly perturbing the forces in which subjects move their limbs, and analyzing the associated dynamic, kinematic and electromyograhic data we aim to distinguish between control theories.

Experiment 1 was designed to test the theory of equifinality, to determine whether subjects form an IM or use EP control. Subjects made goal-oriented wrist flexion movements under the assistance of a velocity dependent torque, without visual guidance. The level of this assisting torque was unexpectedly altered in 10% of trials. A direct test of the EPH was achieved by considering the final position of the wrist in each trial. Under EPH control, equifinality would be expected, as the conditions at the initial and final positions remained unaltered. Under IM control, we hypothesized a undershoot proportional to the level of reduction of the assisting load.

Based on the evidence from experiment 1, and considerable published research (Shadmehr and Mussi-Ivaldi, 1994; Lackner and Dizio, 1994; Gomi and Kawato, 1996; Gandolfo et al, 1996; Goodbody and Wolpert, 1998; Throughman and Shadmehr, 2000; Scheidt et al, 2000, 2001; Takahashi et al 2001), learning appears to involve formation of an IM after prolonged exposure to a novel force field. Experiment 2 investigated the mechanisms involved in the early periods of learning. Here we considered the kinematic and EMG patterns associated with learning a double joint arm movement in the horizontal plane. Subjects learnt to make straight-line movements in a novel field, before the strength of this field is either increased or decreased. The learning of the two perturbed-strength fields was interrupted by numerous trials in the original learnt field. In this manner we aimed to repeat the early learning period and answer specific questions as to how the CNS develops an IM when exposed to novel environments. We hypothesized that during early learning subjects may use impedance control, in the form of cocontraction of antagonistic muscles, to reduce displacement of the hand in trials before the IM was accurate enough to compensate for the external forces. We hoped to determine how the relative reliance on IM formation and impedance control changed during extended learning of the original field, and in early learning of the perturbed field.

We hypothesized that when subjects adapted to the novel field the IM would be become more accurate at compensating for the applied field, with an associated drop in the level of cocontraction. We aimed to find out how quickly the IM begins to develop when subjects encounter the perturbed field.

The objectives of this work are therefore twofold. Firstly, we aimed to elucidate whether EP or IM control is utilised by the CNS. Single joint movements are sufficiently simple that predictions can be made to distinguish the contrasting control theories. Secondly, a multi-joint task was used to present the CNS with a more complex control problem than a single joint task, thereby providing greater insight into the general principles of adaptation and motor learning. It is hoped that the two experiments within this thesis will lead to a clearer understanding of which control mechanisms the CNS uses in goal-oriented movements, and how these develop as learning occurs in novel environments.

Introduction

To improve the way in which we carry out everyday tasks we must learn to adapt our performance in tasks which, at first, we struggle to complete accurately. Many species show the ability for motor learning, which is essential in situations where the organism's environment or task may change.

Movement provides the only way in which we can actively interact with the environment. Without the ability to adapt to changes occurring in growth and exposure to new environmental conditions humans would undoubtedly struggle to survive for long. The process of adapting to these ever-changing conditions over a series of movements is termed motor adaptation.

How do we learn to compensate for the forces that we encounter in everyday tasks, such as turning a door handle, opening a jar or hammering a nail? How may the learning mechanisms differ in adapting to situations which are either stable or, by their very nature, unstable?

It is known that we learn by adapting both the neural networks and structural anatomy of our bodies. We are able to either adapt the controller i.e., the brain, or the body, or both, which will in general lead to a better process of motor learning (Kawato 1999) Both sensory and cognitive processes are indispensable in the process of motor learning. The sensory information that we gain from interacting with the environment is used in determining future motor output. This can be viewed as kinematic and dynamic transformations, which consider the transformation of the system co-ordinates and forces respectively.

An important question to ask is how does the Central Nervous System enable this motor adaptation to occur quickly and accurately such that motor tasks can be completed as desired. The answer to this question is not entirely known, although over the last 30 years several theories have been proposed within the field of neuromuscular control. Three

prominent theoretical control schemes are Internal Model Formation, Equilibrium Point Control and Impedance Control. These control principles aim to explain compensation for environmental forces that act on the limb during movement. Internal models and impedance control can either be utilised individually, or can act together depending on the nature of the applied environmental forces.

One way to elucidate which of the possible schemes is adopted by the CNS to control human movement is to study goal-oriented movement tasks. By asking subjects to learn a specific movement of upper or lower limbs it is possible to distinguish between control theories by studying the effects of perturbing the movement on the associated dynamic, kinematic and electromyograhic data.

The answer to this question might have applications in a wide variety of areas, including neuromuscular rehabilitation, design and control of humanoid robots and design and control of prosthetic limbs to produce motion which more closely resembles natural human movement.

Motor Control Theories

Internal Model Formation

Much recent work has provided evidence that the CNS learns, and stores, internal models of sensorimotor transformations. An internal model (IM) is a neural network within the cerebellar region of the brain that mimics the input-output (sensorimotor) transformation of a particular dynamic process in the external world (Kawato, 1999). Internal models occur as forward, or inverse, processes. Forward models represent transformations from cause to effect while inverse models represent the reverse transformation. A forward internal model is a model, or representation, within the brain, which enables the prediction of a sensory consequence of a specific motor command, or action via efference copy of the motor command. An inverse internal model is a model that predicts, or estimates, the feedforward motor command required to produce a desired consequence

based on the desired trajectory information for the task to be undertaken and information stored in the brain about properties of the limbs and environment dynamics.

The task of bringing the hand, or any other limb, to a target in three-dimensional space involves a series of both sensory and motor commands. The visual and sensory information available to us is translated by the CNS, which calculates the motor commands required to achieve the task. Many natural movements, such as reaching and grasping tasks, occur at a speed that precludes the use of pure feedback control. These movements generally last for 150-600 ms. The feedback loops relaying visual information on arm movements to the brain, which responds with a command, range between 100-250 ms. This is far too slow for the response to keep the arm on the desired trajectory. Even a spinal reflex loop, such as a stretch reflex, has a delay of 20-50 ms, which represents a large proportion of a ballistic arm movement. The movement is either partially, or completely finished before on-line correction can be made for any trajectory errors caused by perturbations or variability in the motor command.

Because of the inherent delays in either the spinal reflex feedback loop, or voluntary responses via higher centres of the CNS, the feedback gains cannot be set too high. A high gain would cause a large amount of instability in the system. We must therefore develop some type of inverse internal model, where the commands are executed in a feedforward manner, if we are to learn to undertake the task more accurately and achieve the desired goal. Learning can therefore be described as the amalgamation of both forward and inverse internal models such that a wide variety of tasks can be achieved quickly and accurately.

Initially a task may be undertaken using motor commands generated in a feedforward manner from previous experiences in a similar situation. We may draw on an existing internal model, which may or may not be appropriate to the new situation. If these commands are, in fact, incorrect and produce error in the hand path, some on-line (feedback) adjustment may occur in that trial using the afferent copy of the error signal.

The motor commands will then be adjusted in a feedforward manner for the subsequent trial to compensate for the task dynamics being different from the expected dynamics.

The scientific literature highlights numerous situations that give evidence for the use of IM's in dynamic tasks. Studies have drawn upon behavioural, neurophysiological and imaging data to form structured plans of the ways in which the internal models functions. Shadmehr and Mussi-Ivaldi (1994) studied two-joint reaches in an externally applied curl force field. The presence of after-effects when the field was removed and subjects again moved in a null field was used as the basis of an argument that suggested that humans developed an IM to account for the predictable forces applied during a reaching task. The random removal of the force led to aftereffects seen as hand paths which were mirror-symmetric to the hand paths upon original exposure to the force field. The IM was used to repeat the motor command on a feed-forward control basis, and hence was attempting to compensate for a force that was no longer present.

Lackner and DiZio (1994) conducted a study in which humans made goal-oriented reaching tasks in a rotating room. The rotation required subjects to adapt their reaching to account for the Coriolis force, a force that is proportional to the cross product of the angular velocity of the rotation and the linear velocity of the arm during the reach. In the first of two experiments the subjects moved to an end point that they contacted, allowing sensory evaluation of their performance. Even in a darkened room subjects were able to re-adjust their trajectories, once rotation began, to compensate for the Coriolis force, producing straight paths that ended at the target. Once the rotation ceased, hand paths were significantly curved in the opposite direction compared to initial trials during the rotation. In a second experiment, where no sensory feedback was available at the target, subjects were still able to adapt to the Coriolis force by straightening their hand paths, but transient errors in the final position were present. End point errors of the opposite sign i.e., aftereffects were present when subjects returned to a non-rotating situation. These aftereffects are characteristic of IM development.

Gomi and Kawato (1996) measured human arm stiffness by perturbing reaching movements in 8 different directions. The equilibrium point trajectory was calculated using these stiffness estimations, the actual trajectory and computed muscle torques. The complexity of the equilibrium point trajectories compared to the simplicity of the actual trajectories, was presented as evidence against a simple rate-controlled shift in equilibrium position, but consistent with IM control.

Gandolfo et al (1996) investigated the ability to generalize reaching tasks to different areas of the workspace, and to different hand orientations. They investigated whether motor adaptation to externally applied perturbations could be generalized to different regions of the workspace, or whether the adaptation was specific to the positions within the workspace that had been visited. The effect of hand orientation on task success was also investigated. Subjects were unable to generalize their movements throughout the workspace. After learning the reaching task in the presence of a force field, aftereffects were seen upon removal of this field suggesting an internal model is formed. The IM produces the required endpoint forces to exactly compensate for the field, rather than simply increasing the overall stiffness of the limb to reduce perturbing effects. When the task was undertaken in different areas of the workspace, the magnitude of the aftereffect was seen to decay as distance from the direction in which the task was learnt increased. This suggests that a map is formed between the required somatosensory input and the environmental forces, which is local to the area in which the task was learnt. Changing the orientation of the hand by 90 degrees, such that the palm was either horizontal or vertical, also affected performance. Upon changing orientation, but keeping the task dynamics the same, repetition of the motor command via the IM occurs. This realises endpoint errors, as the two orientations lead to different joint torques for the same commands to the muscles. Adaptation occurs when the brain realises the commands need to be different under the two orientations. In conclusion, the results together suggest that an internal model of the external perturbing forces is formed in the intrinsic co-ordinates of muscles and joints.

Goodbody and Wolpert (1998) considered the generalization of amplitude and temporal aspects of motor learning. Movements of twice the amplitude or half the duration of the learnt movement were considered. It was found that substantial generalization occurred, suggesting that the internal model developed was a non-local control process. The generalization is best described as a linear extrapolation of the force field whereby the control process i.e., the internal model, learns the relationship between the velocity and force and can map this to new states in the faster and longer amplitude trials. This result is in keeping with the well known training strategies of first practising a new skill slowly, such as a golf swing, before implementing it at correct speed. The same is true for size generalization where young children are taught to form large letters when learning to write. In time they are able to re-create the letters on a reduced scale.

Taking the results from the these generalization studies (Gandolfo et al; 1996, Goodbody and Wolpert; 1998) together it can be concluded that the control mechanism is more powerful at generalizing movements of the same direction which are scaled in time or amplitude, compared to movements in directions which are translated or rotated from the original direction.

Throughman and Shadmehr (2000) considered the motor adaptation process from a computational viewpoint. The internal model was considered as a sensorimotor map which is used to transform the desired arm trajectories onto the required muscular force via a flexible set of motor primitives. Comparisons were made between theoretical representations of how humans can adapt motor primitives and human performance in goal-oriented reaching tasks. The force vector, calculated from force requirements in the dynamical task, can be represented by a vector of scalar-valued primitives, dependent upon the desired position, velocity and acceleration, multiplied by a weighting factor, although in this study a simpler relationship, whereby the primitives were only dependent on desired velocity, was used. The internal model of task dynamics is learnt by modification of the weighting matrix, such that the square of the difference between the actual and desired trajectory is minimised. Assuming that modifications to the weighting

matrix are made in incremental adjustments, preceding each movement using error feedback, the change in internal model is dependent on the experienced error and the mutual projection between evaluations of the primitives, but is independent of the weighting factor. Given the primitives used in this model are dependent only upon velocity the change in the internal model should be proportional to the error experienced for movements to the same target. When the movements aim to locate different targets the change in command will also depend upon the breadth of the receptive fields of the primitives (Throughman and Shadmehr, 2000) in relation to the change in orientation.

Measuring error as displacement from a line joining the start and target, trials where subjects moved under a viscous force field were compared to trials in which the field was randomly removed. After learning, perturbed trials led to error in the opposite direction to which the field had initially caused displacement. Trials immediately after the perturbation had significantly higher error than the trial immediately before the perturbation, highlighting the partial unlearning of the field predicated by the force-primitive equation. As predicted, the level of unlearning preceding perturbations was well correlated to the relative size of the error in the perturbed trials. Previous work had also predicted that spatial tuning of the EMG would undergo a specific rotation during learning (Throughman and Shadmehr 1999). The unlearning effect was visible as a counter rotation between trials prior to, and immediately following, the perturbed trial. Recovery of the learnt trajectory occurred within three trials. The shape of the primitives was then tested by considering the temporal dynamics within, and across, movement directions. Using a series of random movement directions, movement errors were fitted to a representation of the internal model which was dependent on the actual error signal and whether the force field was present. Their initial model, which incorrectly assumed that errors occurring in moving to one target did not affect the internal models for other directions, was expanded to a model including eight dimensional vectors representing the movement directions. It was found, as one would expect, that movement error in a specific direction maximally influenced the internal model in that direction. The influence of error in one direction decayed in neighbouring directions, except for angular

distances larger than 90 degrees where the error appeared to destructively interfere with the generation of the internal model. This suggests that a linear basis for encoding velocity is not possible. A controller was simulated, attached to a biomechanical model of the arm, which, like Purkinje cells within the cerebellum, encoded velocity as a Gaussian distribution of a certain width, specified by the standard deviation. Modelling with a wide width yielded results very similar to actual subject performance, including the aforementioned negative sensitivity at large angles. Results suggest that the brain composes motor commands for the desired task through an adaptive combination of motor primitives. The primitives are broadly tuned to the arm velocity during the reach as a Gaussian distribution with a preferred movement direction.

Scheidt et al (2000) performed an experiment in which subjects learnt goal-directed movements of the arm, with the idea of investigating the relative role of kinematic and dynamic errors in recovering from an adapted force field state. Once subjects had learnt to compensate for the applied external force the force was removed and aftereffects were studied in two contrasting ways. Either a mechanical channel was implemented to restrict subjects to a straight-line path, eliminating any kinematic feedback as to the incorrect hand path, or trials were conducted in the normal null field. It was found that recovery from the novel field was much slower in the trials constrained by the channel where kinematic information was unavailable. In these trials subjects persisted in generating large forces against the channel, which were unnecessary in the null field. The rate of reduction of this force was significantly slower than when kinematic error information was available. The results show that both kinematic and dynamic information influence the motor adaptation process, and that kinematic factors are dominant in instigating a rapid return to a null field situation after exposure to the novel force field.

Scheidt et al (2001) consider the effects of a randomly changing force field. Unpredictable perturbations were made to the amplitude (gain) of the viscous force field in which subjects made goal-oriented movements in the horizontal plane. Double-joint reaches were made whilst the subject grasped a robot manipulator. Previous studies that

applied perturbations to the hand looked at perturbations that were fixed, and repeated throughout the course of the experiment. This study was undertaken to more accurately emulate the situation that may occur in nature where perturbations are random and do not have a fixed or repeated structure. The authors use perturbations due to wind or diffraction of light rays through water as two examples. The way in which the subjects adapted to these changes was analyzed, primarily by looking at the maximum deviation of the hand from a straight line joining the start position and the target. Results indicate that instead of counteracting the mean field strength of individual trials, subjects used information from a limited number of preceding trials. A model utilizing error feedback from a single previous trial was developed and found to be an accurate representation of how subjects developed their control strategy. This method enables near-optimal performance by minimizing the mean square variance of deviation. The study highlights that the neural networks involved in motor learning only require short-term memory of a very limited number of preceding trials. Representations of trials occurring before this were shown to be unimportant in generating the optimal motor response in the aforementioned task.

Takahashi et al (2001) also looked at the effect of an unpredictable field. Subjects were exposed to two fields, a "mean field", with a constant strength and a "noise field", where the strength varied unpredictably according to a normal distribution about the same mean strength as in the "mean field". It was found that the "noise field" did not affect the subject's performance relative to their performance in the "mean field", as measured by the rate of adaptation and the position error about the target. This suggests that both internal model and impedance control strategies are utilised by subjects. The general co-contraction stiffened the limb and led to decreased aftereffects on removal of the "noise field". Internal model control in the absence of any cocontraction, would work well when the "noise field" strength is similar to the "mean field" strength but when the strength of the "noise field" is not close to the expected mean strength, poor performance would occur, characterized by large errors in the trajectory. Impedance control (see Impedance Control Review) is a "back-up" strategy of benefit in the "noise field". Hand path error

on trials in the "noise field" where the strength is substantially different from the mean will be smaller, improving performance.

More recently Franklin and colleagues (Franklin et al, 2003a, b) and Osu (Osu et al 2002) have also considered the use of both internal models and impedance control when learning to compensate for both stable and unstable environments. Regardless of whether the dynamics subjects were exposed to were stable i.e., where an internal model could be used to compensate for the predictable forces, or unstable i.e., a field with unpredictable forces, Franklin (Franklin et al, 2003a, b) found that an increase in endpoint stiffness served to reduce hand-path error. Osu et al (2002) reported that impedance control could be used in conjunction with internal model learning. They reported that cocontraction contributed mostly in early learning and internal models contributed more towards the end of learning. The CNS appears to be able to regulate the stiffness enabling stability to be achieved during the whole learning process.

Popescu and Rymer (2000) considered how applying small, barely noticeable, force pulses at the beginning, middle or end of the motion altered the endpoints of reaching movements. By applying small perturbations the experimenters were able to avoid startling the subjects into an undesired voluntary reaction. Subjects made 10 cm reaching movements across their body in a two dimensional plane. Pulses were applied in a direction perpendicular to the motion of the hand causing increased elbow flexion (towards the body) or increased elbow extension (away from the body). The significant difference in endpoint error, predominantly in the elbow extension, relative to the unperturbed trials, did not significantly decrease with extensive practice. This persistence of endpoint error, i.e., equifinality was not observed, and the non-symmetric effect in the flexion direction, where endpoint error was not affected, were taken as evidence of inconsistency with the EPH, supporting the idea that an internal model is developed. Although the analysis failed to consider the effect the orientation of the two-dimensional endpoint stiffness ellipse of the arm would have on the relative error in the flexion and extension directions, persistence of the endpoint error is still evidence against EP control.

Equilibrium Point Hypotheses (EPH)

The spring like behaviour of muscles has long been recognized as a key element in the control of limb movement. Equilibrium Point models (or mass-spring control models) of human movement rely on this stiffness, and are predicated on the fact that these muscular springs whose length and stiffness can be controlled centrally by the central nervous system (CNS), establishing a balance of force in the springs (Asatryan and Feldman 1965, Feldman 1966a, b, Bizzi and Polit 1978, 1979, Bizzi et al 1982)

The equilibrium point (EP) is described as the point where torque produced by the agonist and antagonist muscles are equal and opposite, producing a stationary limb posture. By activating the flexor or extensor muscles to varying degrees, via a specific motor command, the EP can be shifted. Adding an external load can alter the position of the EP by contributing additional torques to those produced by the muscles. The level of activation in the muscles may be altered to compensate for the externally applied torque, such that the EP is kept in the same position. An example of this would be filling a glass of water. The hand holds the glass in a stationary position by activating the elbow flexor muscles (and perhaps extensor muscles to some level in a co-contraction capacity to increase stability) to compensate for the torque produced by the glass. As the glass is filled the flexor torque will be increased to allow for the additional torque due to the liquid such that the glass is kept still i.e., the EP remains stationary. If the muscles did not compensate as the glass was filled the elbow would straighten, leading to lowering of the glass and ultimately spilling the water. If a limb is displaced from an original EP by an external disturbance, a force is generated within the stretched muscles to restore the original EP as long as the muscular activities are not voluntarily altered as a consequence of the perturbation.

The Lambda (λ) model is one version of the EPH postulated by Feldman in 1965 (Feldman 1966a, b, Astrayan and Feldman 1965) upon which most research has been focused. The alpha model, proposed by Bizzi and his co-workers (Mussa-Ivaldi, Hogan, & Bizzi, 1985), is an alternative equilibrium point control model, although this model has

not received as much attention as Feldman's model. In the λ-model of motor control, motion of a joint arises through a reciprocal (R) command which alters the relative activation of agonist and antagonist muscles, to produce a shift in the equilibrium between the muscle force and external load. The new equilibrium position is stabilized by the stiffness of the joint, which can be enhanced by co-contracting agonist and antagonist muscles, termed the co-contraction (C) command (Feldman, 1986). The R and C commands are referred to as Invariant Characteristics (IC's) and specify a positional frame of reference for activation of flexor and extensor motorneurons. The CNS need only specify the rate of the equilibrium shift and the joint stiffness.

The R command specifies the referent joint angle (R) at which the transition between net flexor and net extensor active torque occurs during changes in the joint angle elicited by an external force. Neurophysiologically, the R command is associated with reciprocal changes in the membrane potentials of agonist and antagonist motor neurons (Levin, Feldman, Milner, & Lamarre, 1992). The C command defines an angular range where the flexor and extensor muscles are simultaneously active (C>0) or silent (C=0) (St-Onge, Adamovich, & Feldman, 1997) i.e., C is equal to zero if there is no co-contraction, and stabilisation of the EP occurs solely due to joint stiffness and the spring-like properties of the muscle. The C command provides increased movement stability and effective energy dissipation, reducing oscillation at the end of movement. The higher the level of C, the higher the co-contraction of the flexor and extensor muscles and the more effectively oscillations will be damped out. The C command is associated with simultaneous depolarisation of the agonist and antagonist motor neurons. (Levin, Feldman, Milner, & Lamarre, 1992)

λ is defined as the activation length of the muscle. In the λ model a muscle is silent i.e., it does not produce any force if its length, x, is shorter than the activation length λ. Both flexor and extensor muscles have an activation length, λ_{fl} and λ_{ex}, which is controlled by the CNS. The range of λ is very important to the model such that relaxation is provided at all muscle lengths. The maximum and minimum values of λ must be respectively larger

and smaller than the maximum and minimum muscle lengths, represented by x. Breaking this condition would lead to absence of force producing capability at short lengths (hypotonia) or spasticity when it is impossible to relax the muscle at long lengths. Stretching a muscle to a length greater than the activation length will elicit a stretch-reflex force to be generated in that muscle to re-establish the EP.

Changes to R and C elicit different changes in the joint invariant characteristics. A change in R produces unidirectional changes in the λ_{fl} and λ_{ex} curves, such that the gradient of the joint invariant characteristic is unchanged. The activity of one of the muscles will be increased whilst the other will be decreased. Position of the joint is altered whilst joint stiffness is kept constant. A change in the C command elicits a contra-directional change in the λ_{fl} and λ_{ex} curves. Either activity is increased or decreased in both muscle groups simultaneously. This causes a change in the gradient of the joint invariant characteristic, affecting joint stiffness without a shift in the position.

Weeks and colleagues (Weeks et al, 1996) considered adaptation of movement to changes in load. Two different strategies were postulated as possible mechanisms when subjects encounter unexpected changes in the load. Under the recurrent strategy the system adapts to the new load condition by adopting a new set of R and C commands within the single trial, such that the target position would be located accurately on the second trial of the new load condition, without the need for secondary corrective movements, assuming an accurate estimation of the new control commands. Under the invariant strategy the original control variables are kept resulting in endpoint errors on the perturbation trial, and any subsequent trials under the new load condition, but an accurate result on returning to the original load. In the study the subjects made elbow flexion movements. After null field trials an assisting, opposing or alternating assisting and then opposing spring-like force was introduced for 5 trials, before the null field was re-introduced. Subjects were instructed to correct any positional error and reach the target as soon as possible. Secondary movements were generally made on the first of the perturbed trials, but were not sufficient to eliminate endpoint errors. On the second

perturbed trial subjects could generally locate the target in a single movement. Weeks et al (1996) research concluded that the recurrent strategy is dominant over the invariant strategy, in agreement with the EP control, or λ model.

Recent expansion of the λ model allowed a description of the way in which the R command changes between the initial and final EP's, with two conflicting ideas appearing. One suggests shifts are monotonic ramp changes (Feldman, Adamovich, & Levin, 1995; St-Onge et al, 1997) ending near the peak velocity of movement, well before the final position is achieved. The other predicts that they are nonmonotonic, but actually "N-shaped" and proceed until the end of movement (Latash & Gottlieb, 1991). Gribble and Ostry (1998) asked the question whether complex control signals, such as the inverse dynamics calculation undertaken in any inverse internal model, are necessary in controlling human arm movement. Using a two-joint model of the arm, as specified under the λ model of the EPH, the authors address claims that the complex equilibrium trajectories are required to explain non-monotonic shifts in the equilibrium point. Under EP control it is proposed that smooth EP trajectories can be achieved via simple monotonic shifts in the EP. This notion has been challenged by Latash and Gottlieb (1991), who reported that complex, nonmonotonic EP shifts are required to produce the torques necessary to produce rapid, single joint movements and Gomi and Kawato (1996), who showed complex calculations must be undertaken to compensate for the dynamics of the multijoint arm. Gribble and Ostry (1998) state that simplifications made in the force generating model of Gomi and Kawato (1996), specifically the assumption that joint torques vary linearly with the difference between actual and equilibrium position, led to the large differences seen between the actual and the calculated EP trajectory. This simplification led to a situation in which the hypothesized EP trajectory must first lead, and then lag, the actual limb position such that acceleration and deceleration could be produced.

Gribble and Ostry (1998) developed simulations using a more explicit muscle model, instead of a simple motor to produce torque. With a constant level of cocontraction and a

constant rate of equilibrium shift, in the presence of an external torque they found that the actual trajectories recorded by Gomi and Kawato (1996) as well as rapid, single joint movements (Latash and Gottlieb, 1991) could be accurately simulated. The resulting simulated trajectories were, indeed, more complex than simple bell shaped velocity profiles, and resembled the actual trajectories in the empirical studies (Gomi and Kawato, 1996; Latash and Gottlieb, 1991). This showed that although the hand path trajectories were complex they could be reproduced from simple monotonic control signals under EP control. Although this appears to support the EPH, the authors point out that even though simple control signals, such as those under the λ-model, can control simple movements, movements with complex kinematics, control of movements at various speeds, or movements requiring compensation for external loads may be controlled using more complex control signals. The authors also indicate that the nature of the neuromuscular plant used in any simulation i.e., the muscular model, greatly influences the inferred trajectories, even if the control signals are themselves simple

Gribble and Ostry (2000) also investigated the control of arm movements when movement–dependent loads were present. A position control model based on EP control, predicting shifts in the equilibrium point in direct proportion to the relative area between the desired and actual hand position, was used in the analysis of subject's performance. Simulations were made of movements made with internally generated loads, i.e., joint interaction torques, and in externally applied velocity force fields. Using a simple linear compensation method the model predicts trajectories very similar to the actual experimentally recorded trajectories with either internal or external loads. Complex coordinate transformations or inverse dynamic calculations were not required, suggesting that an iterative, position-control model, such as the EPH, could feasibly govern human motor tasks, rather than the force control (internal) models suggested in a number of recent studies (Lackner and Dizio, 1994; Shadmehr and Mussi-Ivaldi, 1994; Conditt et al 1997, Gomi and Kawato, 1996). Although the position-controller suggested in this study is not incompatible with the idea of internal models anticipating the consequences of force in a feedforward manner, it has shown that position control is also a possibility.

According to the λ-model (Feldman, 1986), increased muscle activation through stretch reflex pathways should compensate for small, transient perturbations or loads that disappear once movement ceases, thereby preserving the final equilibrium position. For example, if a load, which depends only on velocity, is unexpectedly increased, the velocity of a joint will be reduced and the disparity between the joint angle and the pre-programmed EP will increase relative to movements which were not perturbed in this way. Consequently, the difference between joint angle and stretch reflex threshold (λ) increases, which should result in an increase in agonist muscle activity that remains until the joint reaches the originally programmed final equilibrium position.

Impedance Control

The value of internal model learning is limited if the forces that the limb is exposed to are unpredictable or inherently unstable. Such forces arise on a day-to-day basis (Colgate and Hogan 1988; McIntyre et al, 1996) and the CNS needs to be able to deal with them in such a way that tasks can still be completed accurately and efficiently. The forces that are generated when a limb interacts with an external perturbation causing a displacement play a key role in the stability of the interaction.

What control strategies can be adopted by the CNS to ensure that the forces occurring due to the interaction do not cause the system to become unstable? In terms of human movement control, stability is often defined by the Lyapunov stability. This states that a dynamical system is stable when small perturbations of the state converge to zero and is unstable when a small perturbation produces divergence. (Vidasager, 1978) An example of an unstable equilibrium would be balancing on one leg. A small push from any direction will cause instability. Once we leave the stable position the error accumulates and it is then very difficult to regain a stable posture, culminating in losing our balance or falling over.

Hogan (1985) proposed that Impedance Control, specifically controlling the stiffness of individual muscles, could be implemented by the CNS as an optimal way to compensate

for unstable environmental forces. By increasing the stiffness of the human limb (for example the shoulder-elbow system) in specific directions, the displacement due to a predictable perturbing force could be reduced in the force direction without requiring higher stiffness in other directions. Impedance Control is possible as the spring-like properties of the two-joint system have a directional property. A displacement in a particular direction may lead to a restoring force along that direction, or at some angle to it, depending on the direction of the displacement.

The two dimensional stiffness of the human arm (shoulder-elbow system in a horizontal plane), K_e, is given by

$$\Delta F = -K_e \Delta r$$

where

$$K_e = \begin{bmatrix} k_{xx} & k_{xy} \\ k_{yx} & k_{yy} \end{bmatrix}, \Delta r = \begin{bmatrix} \Delta x & \Delta y \end{bmatrix}^T$$

K_e can be represented as the sum of symmetric and anti-symmetric matrices. The symmetric component represents forces that have no curl component and can therefore be derived from a potential function. The anti-symmetric component represents the forces that have a non-zero curl component. Comparison of the relative magnitude of the symmetric and anti-symmetric parts provides a measure of the spring-like nature of the system. Mussi-Ivaldi et al (1985) found that for small displacements the anti-symmetric component was indeed small, relative to the symmetric component, and hence a multiarticular system can be deemed predominantly 'spring-like'.

The magnitude, or area, of the endpoint stiffness ellipse represents the overall endpoint stiffness, with the major and minor axis lengths proportional to the maximum and minimum eigenvalues of K_e. Mathematically, the orientation of the major axis is along

what is known as the principal eigenvector. A displacement along the direction of the maximum or minimum eigenvalue will result in a restoring force along that direction. A displacement in any other direction will lead to a restoring force at some angle to that direction. Specifically, the restoring force will occur along a direction perpendicular to the tangent of the ellipse at the point of displacement. The orientation of the ellipse is determined by the direction of the maximum eigenvalue. In the very specific case that the maximum and minimum eigenvalues are equal the ellipse becomes a circle. In this situation, the restoring force will always be oriented along the direction of the displacement.

Along the direction of the major and minor axes, the restoring force is determined by the stiffness, which is directly proportional to the axis length. In any other direction the displacement is determined by the vector sum of the displacements along the major and minor axes, and the resultant force is determined by the forces along the axes. Given that the stiffness of individual muscles increases as the muscular activation increases, it is possible to co-activate agonist and antagonist muscle groups spanning a single joint to increase the stiffness, and therefore, in most situations, increase the stability of that joint to any disturbances it may encounter. Applying this principle to a multi-joint system, with the added directionality effect of the stiffness, Hogan (1985) postulated that the CNS could use Impedance Control as a strategy to control posture. He theorized that the CNS could alter the stiffness in preferred directions by activating specific muscles to reduce the effect of perturbations along those directions. Verifying impedance control is not as simple as it may first appear, as it must be shown that the changes in stiffness orientation, shape and size are independent of the force produced by the hand.

Milner and Cloutier (1993) considered the stability at the target following wrist flexion movements under a destabilizing, negative viscosity load. Subjects initially used more cocontraction than necessary to stabilize the wrist in the target position for the required duration. Although cocontraction globally increased the endpoint stiffness, helping to reduce the amplitude of the ensuing oscillations, this was not always a helpful strategy

since it tended to increase reflex gain. Subjects had no difficulty in producing the correct amount of torque to locate the target, but found the task of quickly damping oscillations within the 3 ° target window more difficult. With practice they reduced both the level of cocontraction and the amplitude of the oscillations. It appears that with practise subjects were able to reduce cocontraction such that they generated "optimal impedance", providing the required stability at the target without over-increasing the reflex gain or causing unnecessary metabolic cost in the muscles.

Recently several studies have identified situations in which impedance control occurs (Burdet et al, 1999, 2001; Franklin, 2000). When exposed to a predictable velocity dependent curl field subjects develop an internal model to reach the target accurately and accomplish the task with a straight-line path. Characteristic aftereffects are seen when the force is removed, where subjects' paths diverge from a straight line in the opposite direction to that in which they were perturbed when first exposed to the force (Shadmehr and Mussa-Ivaldi, 1994). This was not seen when subjects were exposed to a divergent field, where the perturbing force depended on the distance from a straight-line path between start and target positions. Because the force direction at the onset of movement could be unpredictable due to movement variability subjects apparently did not develop an internal model of the force field. When this field was unexpectedly removed no aftereffects were visible. In fact, the hand paths on post-perturbation trials were significantly straighter than the paths in the null field trials prior to the force field being activated. Using a technique whereby the hand is displaced by a small amount from the theoretical unperturbed path, during the middle of the movement, the endpoint stiffness was estimated from the resistance to displacements in eight separate directions. In the diverging field subjects were found to increase the endpoint stiffness of their arm in a direction parallel to the direction of the force field without increasing the stiffness in the orthogonal direction, demonstrating Impedance Control.

Control by Internal Model or EPH?

Studies of human muscular control often use learning paradigms in which subjects are asked to make task-oriented movements. Subjects learn to move to a target in a null field, where no external force is applied to the system. A novel force field, to which subjects have not previously been exposed, is then randomly, or consistently, introduced to perturb the learnt hand path trajectory. Analysis of adaptation to these forces over time, and the aftereffects that occur when the field is subsequently removed, can be measured using both kinematic and EMG measures. Patterns in learning can be used to deduce how the CNS modifies motor commands to improve performance. The results of such studies suggest three possible schemes, although not necessarily mutually exclusive, which may be employed by the CNS in controlling these goal-directed tasks: Internal Model formation, EP Control or Impedance Control. Present research continues in this area of neuromuscular control in an attempt to unambiguously determine whether or not the aforementioned mechanisms are used in particular tasks. The experiments within this thesis aim to more clearly elucidate the mechanisms that are used be the CNS in adapting to novel dynamic environments.

Experiment 1: The Case for Internal Dynamics Model versus Equilibrium Point Control in Movement

Introduction

The simplicity of the EPH (Feldman 1966a, b; Feldman and Astrayan, 1965) is very appealing given that no knowledge of any geometry of the constituent limbs or inverse modelling calculations is necessary. However, many researchers seem unconvinced that the brain can control complex movements without such information. Abundant evidence suggests that the ability to perform skilled movements in novel mechanical environments is only gradually acquired with practice (Shadmehr and Mussa-Ivaldi 1994; Conditt et al. 1997; Lackner and Dizio, 1994; Gomi and Kawato, 1996; Gandolfo et al, 1996; Goodbody and Wolpert, 1998; Scheidt et al 2000, 2001; Takahashi et al 2001; Thoroughman and Shadmehr 2000). These studies provide convincing evidence that motor learning involves formation of an internal model of task dynamics as skill is acquired.

Proponents of the EPH have argued that these results are not incompatible with the EPH and that there is no reason to abandon it. Feldman et al. (1998) argued that Lackner and Dizio's 1994 experiment entailed a change in the motor command, violating a necessary condition for equifinality under the EPH. Gomi and Kawato's results were dismissed because their estimated equilibrium trajectory did not arrive at the final position until sometime after the actual movement had already ceased. Feldman et al. (1998) argued that this was not physically possible and attributed it to simplified assumptions about the force-generating mechanism (Gomi and Kawato, 1996).

DiZio and Lackner (2001) investigated whether the subjects in their previous study (Lackner and DiZio; 1994) could have reacted, as Feldman had suggested (Feldman et al, 1998) to the Coriolis force due to the rotation. Five labyrinthine-defective subjects, unable to sense body rotation in the dark, were compared to five age-matched normal healthy subjects. Similarly to the previous study subjects made reaching movements to a

just-extinguished target. Both groups initially showed large trajectory and endpoint errors in the rotating room trials. Normal subjects adapted both trajectory and endpoint within 11 trials, but the labyrinthine-defective subjects were unable to reduce their endpoint errors, even after 40 trials. The specific dynamics of the labyrinthine-defective subjects' hand paths, namely higher elevation, denied them somatosensory information about the accuracy of the endpoint. Aftereffects showed mirror-symmetric trajectories for both groups, although aftereffects of endpoint were much more pronounced for normal subjects, highlighting the lack of adaptation by labyrinthine-defective subjects. These results provide evidence for IM control. Characteristic aftereffects of adaptation and the need for somatosensory information to reduce endpoint errors would not be predicted under EP control.

Although Popescu and Rymer (2000), showed persistence of endpoint error, reported as evidence against the EPH, the study has several shortcomings. First, the stiffness of the arm was relatively low, favouring residual position error. Second, perturbations during small amplitude movements have been shown to have a greater effect on endpoint error than during large amplitude movements (Sanes and Evarts, 1983; Sanes 1986). Third, the demonstrated lack of equifinality only for the extension perturbations is equivocal.

In addition to the aforementioned shortcomings of previous tests of the EPH, no test so far has been designed to compare alternative hypotheses, which predict distinct outcomes leading to the validity of the results being questioned. We designed an experiment for which formation of an IDM predicts that final position error will increase with the magnitude of an unexpected change in the environmental force. The EPH, in contrast, predicts that final position should not change.

The environmental force was realized as an assisting torque, proportional to the velocity of wrist movement. Under the IDM framework, an unexpected reduction in the assisting torque would cause the movement to fall short of the intended goal by an amount directly proportional to the amount of the reduction. Under the EPH framework, regardless of the

level of reduction in the assisting torque, and regardless of whether changes in the equilibrium point occur monotonically or in a more complex manner, the same final position should be achieved, albeit more slowly.

Our experimental design, unlike that of Gomi and Kawato (1996) requires no estimation of an equilibrium trajectory. It also addresses another issue raised by Feldman et al. (1998). Specifically, the property of equifinality is based on the premise that the pattern of central commands is not altered by a perturbation, something which Lackner and Dizio (1994) and Popescu and Rymer (2000) did not demonstrate. To deal with this issue, we recorded the activity of all of the primary muscles contributing to the movement to demonstrate maintenance of the same feedforward commands under all force perturbation conditions. Furthermore, to address the claim of Feldman et al. (1998) that the force perturbations in Lackner and Dizio's (1994) study, and possibly also in Popescu and Rymer's (2000), may have been lower than a threshold, below which the EPH demonstrates no corrective muscle activation, we imposed perturbations of 4 different strengths, which included load changes of more than 100%. The 30° flexion movements of the present study can be considered comparable to the large amplitude movements in Sanes' studies (Sanes and Evarts 1983, Sanes 1986), and the co-contraction of wrist flexor and extensor muscles required to perform the task successfully produced high stiffness (Milner and Cloutier, 1998). In this manner we could test whether the EPH would hold for high stiffness, large amplitude movements, as previously suggested (Bizzi et al 1978, Bizzi and Polit 1978) but only be prone to failure for low amplitude, low stiffness movements (Popescu and Rymer 2000, Sanes and Evarts 1983, Sanes 1986).

Methods

Experimental Procedure

Six healthy, right-handed subjects (three men and three women aged 21-32) took part in the study. All subjects gave their informed, and written, consent to the protocol, which was approved by the Simon Fraser University Ethics Review Committee.

The apparatus and data acquisition system is the same as that described by Milner and Cloutier (1998). Briefly, a torque motor (PMI U16M4) with a fixture for securing the hand was used to apply an assisting torque to the wrist during flexion movements. Position, velocity and torque were measured together with surface EMG. Active bipolar surface electrodes (Butterworth 4^{th} order bandpass 20-500 Hz) were used to record activity from flexor carpi radialis (FCR), flexor carpi ulnaris (FCU), extensor carpi radialis (ECR) and extensor carpi ulnaris (ECU). All signals were sampled at 2000 Hz.

Subjects sat 0.70 m from a computer screen on which they viewed a display depicting wrist position as a vertical cursor. A 1 ° start window, centred on the subject's relaxed wrist position (0 °), and a 3 ° target window, centred at 30 ° of flexion, were displayed as shaded regions above and below the position cursor, a distance of 0.18 m apart on the screen. The full screen represented 42 ° and had a resolution of $1/15^{th}$ of a degree of wrist motion. Subjects could easily detect movement of $1/5^{th}$ of a degree. The cursor was active only between trials to allow subjects to reposition the wrist in the start window. A box placed over the hand and apparatus prevented the subject from viewing the hand. Subjects were instructed to move the wrist from the start window to the centre of the target window at a prescribed peak velocity of 300 $°\cdot s^{-1}$ ±10%. No visual feedback of wrist motion was available during the movement. The final position of the cursor and peak velocity during the movement were displayed only at the completion of the data acquisition, 1.5 s after movement onset. A horizontal bar at the bottom of the display screen reflected whether subjects had moved too slowly, too quickly, or in the desired range of peak velocity of 300 $°s^{-1}$ ± 10%. Subjects could make adjustments to the velocity on subsequent trials using this information. Although subjects were aiming to match the required peak velocity, we decided, prior to analysis of the data, that all trials were included in the analysis, regardless of the actual peak velocity.

Movement was assisted by the torque motor in proportion to angular velocity, ω. The control law used was $\tau = B\omega$. The velocity feedback gain, B, was positive, such that the torque, τ, assisted movement to the target. This produced negative damping of the torque

motor, which resulted in underdamped oscillation of the wrist about the final position. The value of B was selected individually for each subject to ensure that the oscillation had completely subsided several hundred milliseconds before the end of data acquisition. Values ranged from 0.009-0.011 $Nm \cdot s \cdot deg^{-1}$.

Subjects performed 500-700 practice trials in sessions conducted prior to the day of the experiment to become familiar with the apparatus and to learn task dynamics. All practice trials were performed with the assisting torque. The experimental session began with an additional 100 practice trials with the assisting torque, followed by 5 sets of 100 trials in which the assisting torque was occasionally and unexpectedly reduced. There was a break between each set to prevent fatigue. Since trials were self-initiated subjects could rest within sets, if desired.

On 10% of trials, selected randomly, the value of B was reduced prior to movement onset, without warning. The amount of the reduction was randomly chosen as 0%, 25%, 50%, 75% or 100%. Each reduced value of B occurred 10 times. No feedback of wrist position or peak velocity was given on these trials, which are hereafter referred to as perturbation trials. Subjects were instructed to produce a single movement and not to react voluntarily to any changes in the assisting torque. Trials in which voluntary reaction was evident, i.e., as one or more late corrective movements, where zero crossings were present in the velocity trace, were excluded from later analysis (see below). This was detected in 3 perturbation trials for one subject and in 4 perturbation trials for another. In the statistical analysis, the data for these trials were replaced by the average of the remaining trials corresponding to the same perturbation condition.

One subject was consistently unable to suppress corrective reactions to changes in load during the perturbation trials. Late corrective movements towards the target were detected in 24 of the 50 perturbed trials. Because the subject was not able to adhere to the instructions, his data was not included in subsequent analysis.

Analysis

This study enabled us to test the two very distinct, hypotheses predicted under EPH and IDM control. Under the IDM framework, learning to move with the assisting torque would consist of a reduction in the joint torque produced by agonist muscles to accelerate the wrist to the required velocity and an increase in the torque produced by antagonist muscles to decelerate the limb and accurately locate the target. Under the EPH framework, learning would involve decreasing the rate of the shift in equilibrium position, which would now be augmented by the assisting torque. Initial overshoot of the final position, followed by underdamped oscillations about the target, would be predicted under both EP and IDM control.

An IDM model predicts that an unexpected reduction in the assisting torque will cause the movement to fall short of the intended goal by an amount which increases with the magnitude of the torque reduction. Assuming that the feedforward motor command is not altered, the speed of the movement will be reduced in relation to the amount by which the assisting torque is reduced. The slower the movement, the closer to the start position that the extensor torque impulse will begin and consequently the farther short of the target that the movement will stop. Oscillation will still occur about the new final position if the assisting torque is not reduced to zero. Under the EPH, movement kinematics will be affected in a similar way when the assisting torque is reduced, i.e., the movement speed will be reduced. Assuming that the pre-programmed change in equilibrium position remains unaltered, the flexor muscles will be farther from their stretch reflex threshold (λ) than before so their activation will increase, compensating for the reduction in the motor torque and driving the wrist to the *same* final position. The velocity profile may be skewed, reflecting a longer deceleration phase since flexor muscle activation will be reduced as the wrist approaches the programmed final equilibrium position. Depending on the level of torque reduction, some overshoot and oscillation about the target would be expected if the wrist-motor system remains underdamped. The predicted result of removing 100% of the assisting torque is compared for IDM control and EP control in Figure 1.

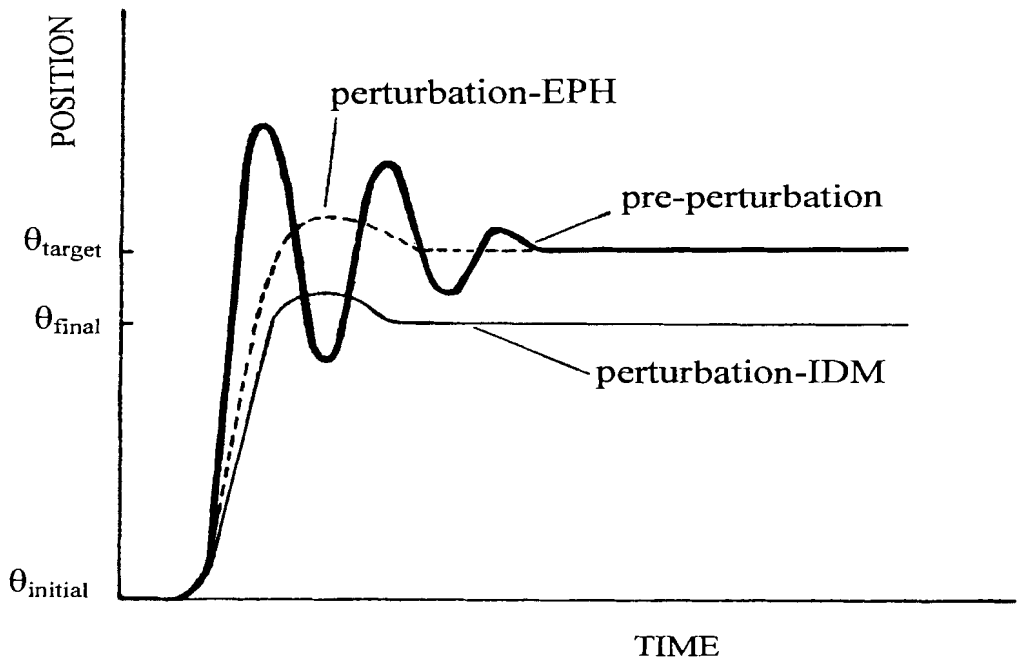

Figure 1. Theoretical predictions of the movement kinematics under IDM and EPH control.
Movement under IDM and EPH result in identical movements after extensive learning of the control (pre-perturbation) movements (thick line) from $\theta_{initial}$ to θ_{target}. On a 100% torque-removed perturbation trial, stretch reflex activity creates extra flexor torque such that the same final position, θ_{target}, is achieved under the EPH (dashed line). In contrast, repetition of the feedforward motor command under IDM control results in an undershoot of the target (thin line) of θ_{target}-θ_{final}, directly proportional to the magnitude of the torque change.

Position errors were analyzed by calculating the difference in final wrist position between perturbed trials and control trials immediately preceding (pre-perturbation) or following (post-perturbation) the perturbed trials. Final position was calculated as the average position during the final 100 ms of the data acquisition period, i.e., 1.4-1.5 s after movement onset. We confirmed that all motion had ceased before this time. Two-way repeated measures ANOVA and linear regression were used to establish the statistical significance of the differences. An *a priori* significance level of 0.05 was used in the ANOVA. The repeated measures assumption of sphericity (Howell, 1997) was tested using Mauchly's test. Sphericity, ε, is a measure as of how similar all the eigenvalues of the orthonormally transformed variance-covariance matrix are to unity (1). When violations of the assumption were detected, i.e., $\varepsilon<1.0$, sphericity-corrected repeated measures analysis of variance tests were performed. Greenhouse-Geisser corrections

were implemented when $\varepsilon<0.7$, while Huynh-Feldt corrections were used when $0.7<\varepsilon<1.0$ to downward-adjust the degrees of freedom (Howell, 1997). Although the p-values of the results change due to the correction factors, the significance of the results were not affected at the *a priori* level.

To determine whether the motor commands to muscles were modified when the assisting torque was reduced, the cumulative EMG on pre-perturbation and perturbation trials was compared for each muscle. The EMG was rectified and smoothed using a 10-point moving average. The cumulative EMG at any point in time was calculated by summing the rectified EMG up to that point, starting 250 ms prior to movement onset. The mean cumulative EMG was calculated for each perturbation condition by averaging values for each of the 10 trials. Mean cumulative EMG for perturbation and pre-perturbation trials was compared using a non-parametric Kolmogorov-Smirnov test (Sokal and Rohlf 1981) with an *a priori* significance level of 0.05. Using this test, the cumulative EMG for the two conditions could be compared up to any selected time to determine whether it differed significantly at any prior time.

The nature of the environmental force allowed us to analyse stretch reflex activity during the underdamped oscillations about the final position. We compared stretch reflex gain and threshold in pre-perturbation and post-perturbation trials. In both pre- and post-perturbation trials the wrist oscillated about the final position. These oscillations stretch and shorten the wrist muscles, producing phasic modulation of the EMG that can be attributed to the stretch reflex (Milner 2002). If the stretch reflex was altered in any way during the random perturbation we would expect to see a change in the stretch reflex response in the post-perturbation trial. Although this is an indirect means of addressing the question of changes in reflex threshold (λ), the method can be applied to all levels of torque reduction, whereas direct analysis could only have been undertaken for trials corresponding to 0% and 25% reduction in the assisting torque, since with higher torque reductions there was no oscillation about the final position. EMG integrated over a period of 50 ms, beginning 20 ms after onset of muscle stretch, as indicated by zero crossings of

wrist velocity, was used as a measure of the monosynaptic stretch reflex. The amplitude of the wrist displacement, relative to its final position, was used as a measure of the magnitude of the muscle stretch. All pre- and post-perturbation trials were analysed. The 50 pre-perturbation trials for the 5 levels of torque reduction formed one set of data on which linear regression between cumulative EMG and displacement amplitude was performed for each subject. Similarly the 50 post-perturbation trials formed a second set of data on which linear regression was performed for each subject. By conducting a test of parallelism, an inference could be made about differences in stretch reflex gains, while testing for common intercepts provided information about differences in stretch reflex thresholds.

Results

With extensive practice subjects adapted to the assisting torque and were able to perform accurate wrist flexion movements. The mean final position of pre-perturbation trials was 29.7 ° (s.d. 2.23), almost at the centre of the target window, while the mean peak velocity was 334 $°s^{-1}$ (s.d. 49.4), approximately at the upper bound of the velocity target window. This was achieved without visual feedback during the movement. No secondary corrective movements were present in any of the trials included in the analysis.

The pre-perturbation trials (full assisting torque) are compared with the perturbation trials in which the assisting torque was reduced by 100% in Figure 2. There was a prominent overshoot and an underdamped oscillation of the wrist prior to stabilization at the final equilibrium position in pre-perturbation trials due to the assistance and negative damping of the torque motor. The positive velocity feedback control law, governing the behaviour of the torque motor, can be readily identified by comparing the recorded velocity and torque signals. The average final position was very close to the centre of the target window (30 °), indicating that subjects accurately compensated for the dynamics of the torque motor. Muscle activity was characterized by prominent bursts in the flexor carpi radialis (FCR) and flexor carpi ulnaris (FCU) muscles to initiate movement. Bursts in the

extensor carpi radialis (ECR) and extensor carpi ulnaris (ECU) muscles occurred shortly after movement onset to decelerate the wrist. There were later bursts which reflected stretch reflex responses during the oscillations about the final position.

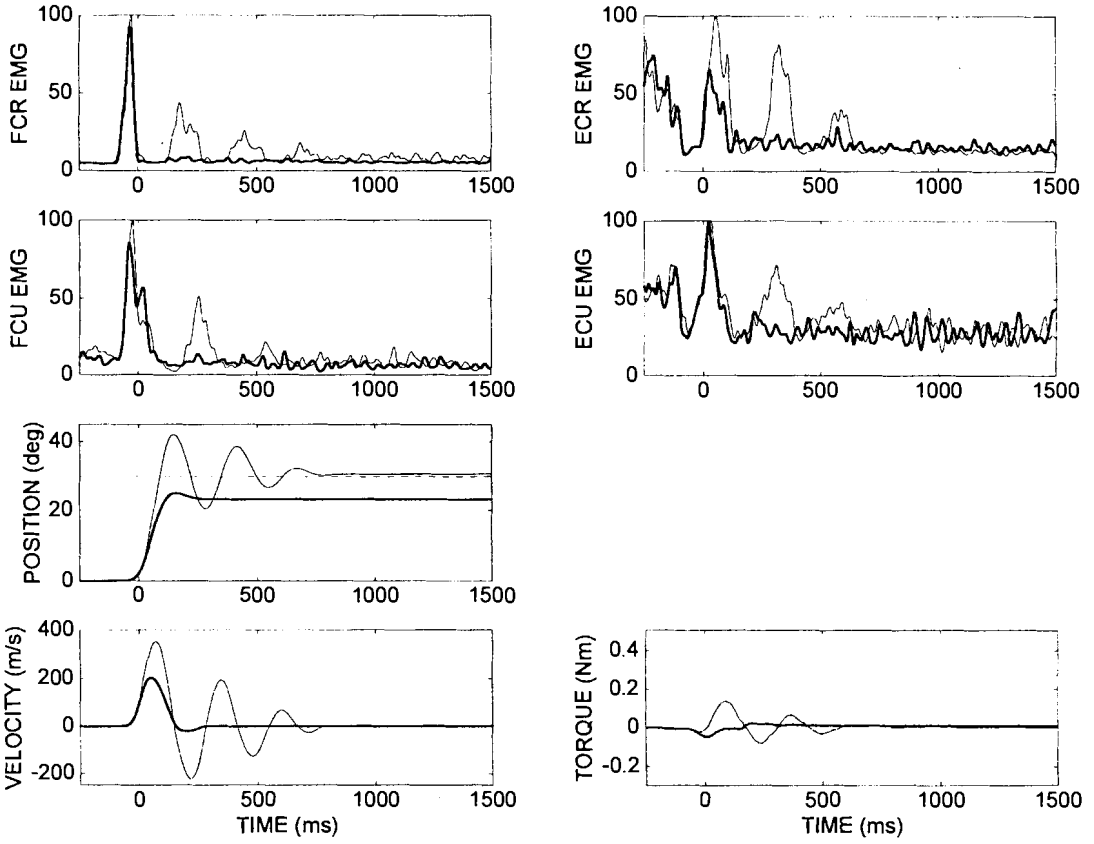

Figure 2. Kinematic and EMG data for the pre-perturbation (thin lines) and 100% removed perturbation trials (thick lines).

The key feature of the perturbation trials i.e., reduced assisting torque, was the prominent error in final position. The wrist always stopped short of the target. Without the assisting torque there was much less overshoot in the position and the wrist stabilized quickly at its final position. The peak velocity was much lower than the target value since there was no assisting torque. The initial EMG bursts in both flexor and extensor muscles were almost identical to those recorded in the pre-perturbation trials, although later bursts were generally undetectable due to the absence of oscillation.

Mean position traces for the perturbation trial at each level of torque reduction are shown in Figure 3. As a higher percentage of the assisting torque is removed, the amount of oscillation about the final position is significantly reduced, reflecting a reduction in the negative damping of the wrist-motor system. To test whether the final position error was

related to the percentage of the torque that had been removed, we calculated the difference in final position between each pre-perturbation trial and the subsequent perturbation trial. The differences were averaged for each perturbation condition. The average over all subjects is displayed as a linear regression in Figure 4. It is clear that the error increased linearly with the percentage of assisting torque removed. Regression lines fitted to the data of individual subjects were highly linear with coefficients of variation (R^2) between 0.74 and 0.94.

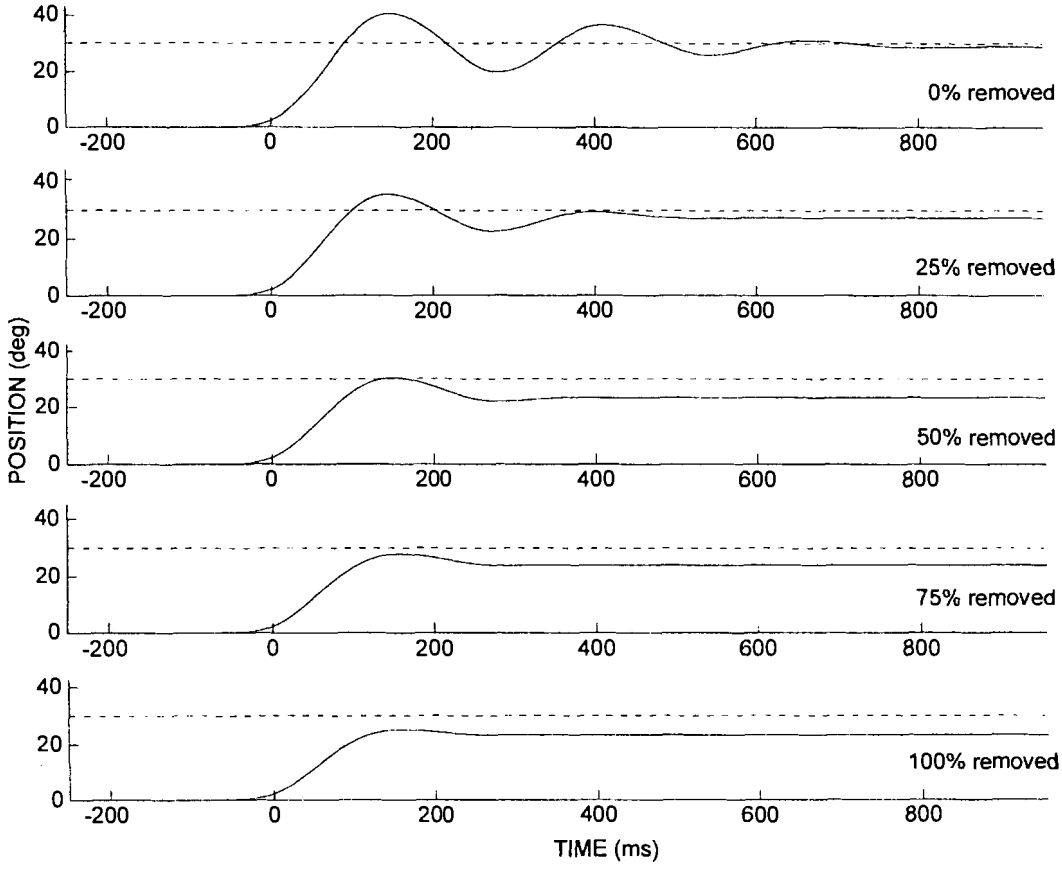

Figure 3. Position traces for the control perturbation trials (0% torque reduction) and the 4 levels of torque reduction.
Endpoint error increases and the amplitude and number of oscillations about the target decrease as more of the assisting load is removed. Data is averaged over the 10 trials in each condition for subject DM.

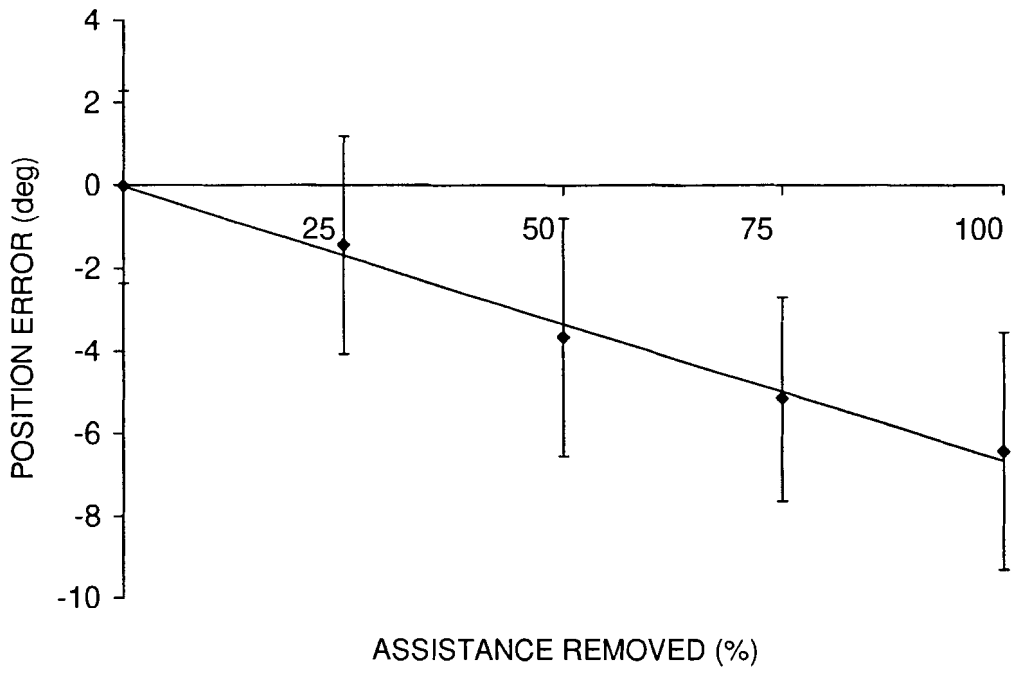

Figure 4. Line fit of the difference in final position between in the perturbation trials and corresponding pre-perturbation control trials, as a function of the percentage reduction in assisting torque.
The data points represent the averages for the 5 subjects. Error bars represent the standard deviation of all trials for each perturbation condition.

Two-way repeated measures ANOVA was used to compare final position in pre-perturbation and perturbation trials. The main effect of reducing the assisting torque was highly significant ($F=15.472$, $p=0.003$), showing that subjects stopped shorter of the target on the perturbation trials than the pre-perturbation trials. The trial main effect was not significant ($F=0.469$, $p=0.648$) i.e., as subjects were exposed to more and more perturbation trials, performance in the post-perturbation trial did not alter, relative to early exposures. The interaction between reduced assisting torque and trial number was not significant either ($F=1.131$, $p=0.377$).

ANOVA to compare pre- and post-perturbation trials, i.e., trials immediately following the perturbation trials, did not reveal any difference in final position (mean –0.009 °, s.d. 2.576 °). There was no significant effect of the percentage of assisting torque removed in

the intervening perturbed trial ($F=1.965$, $p=0.195$) or of the trial number ($F=0.507$, $p=0.667$), nor was there any significant interaction between the two ($F=1.179$, $p=0.358$). Figure 5 shows the pre- and post-perturbation trials corresponding to the 100% torque removed perturbation. The amplitude of the overshoot and the number of oscillations were the same in both cases as the torque motor provided full assistance on these trials. In fact, the kinematic traces correspond so precisely that the two traces are completely coincident for the duration of the movement. Furthermore, a two-way repeated measures ANOVA revealed no significant difference between the peak velocity in pre- and post-perturbation trials over all 5 subjects ($F=2.920$, $p=0.160$), nor was there any significant difference in peak velocity between trials in the 5 different load conditions ($F=1.939$, $p=0.178$). Stretch reflex activity was evident in all muscles during oscillations about the final position and the reflex bursts were modulated with the oscillation amplitude in a similar manner for pre- and post-perturbation trials.

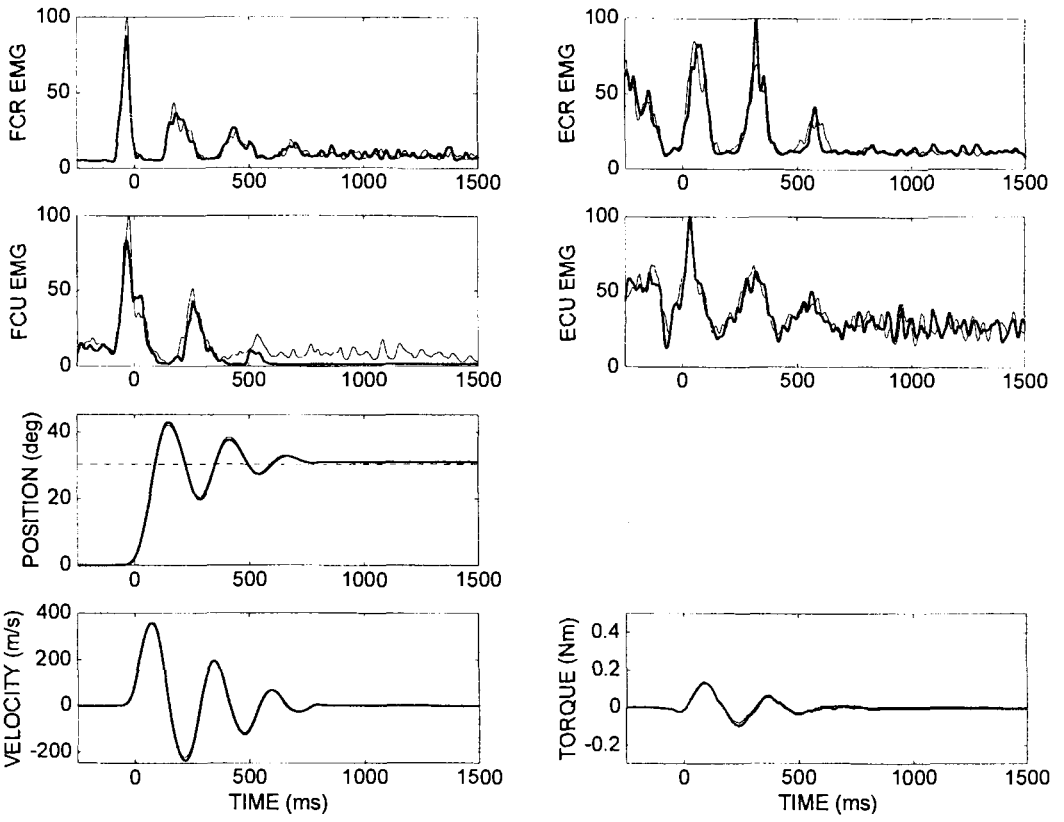

Figure 5. Comparisons of EMG (normalized to peak), position and velocity for pre-perturbation (thick lines) and 100% assistance removed post-perturbation trials (thin lines).
Data were averaged for the 10 trials under each condition for subject DM. EMG bursts during the period of stabilization are phase-locked to wrist displacement. The burst amplitudes can be seen to decrease as oscillation amplitude decreases. Note the distinct similarity in EMG and kinematic traces, and that the target is accurately located in both pre- and post- perturbation trials.

The cumulative EMG for one subject is shown in Figure 6 for pre-perturbation trials and perturbation trials, where the assisting torque was reduced by 100%. There were no significant differences in the cumulative EMG ($p<0.05$) for any muscle for this subject or for any of the other four subjects, who did not make corrective movements, prior to the time at which stretch reflex responses occurred in the pre-perturbation trials during oscillations about the final position. Furthermore, regardless of the percentage of assisting torque reduction, no significant difference in the cumulative EMG during this period was found for any of these subjects. This indicates that subjects issued the same pattern of feedforward motor commands to wrist muscles in perturbation trials as in pre-perturbation trials.

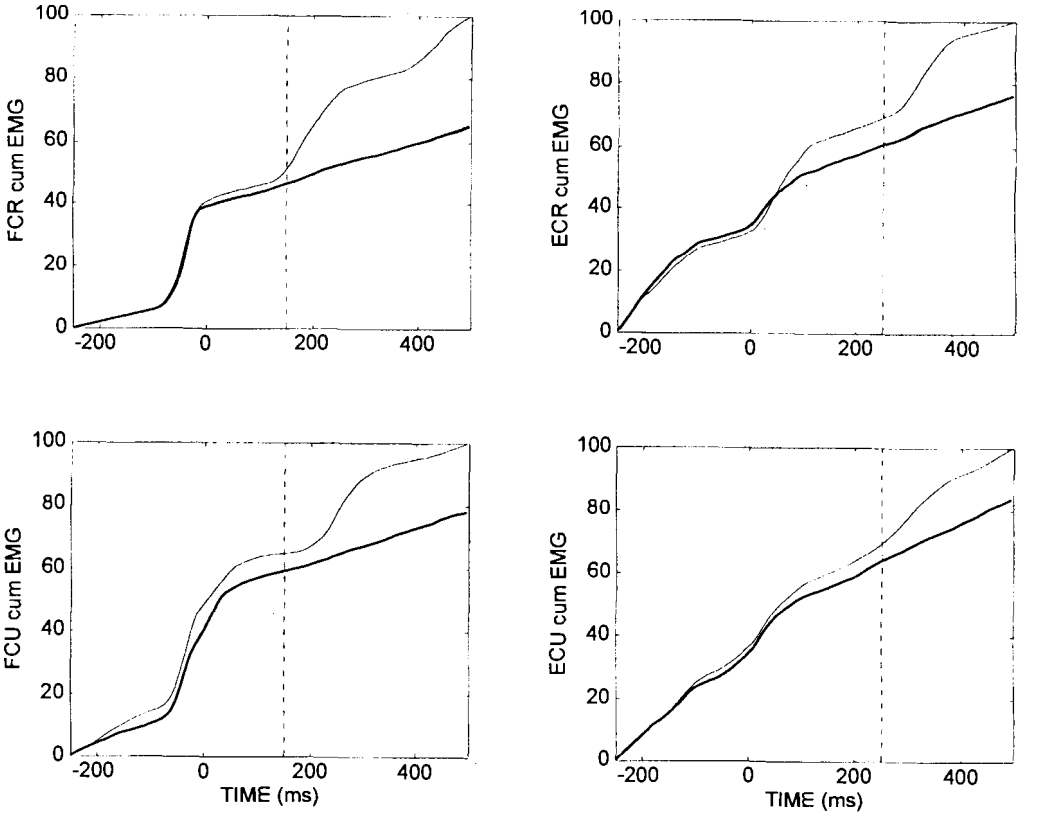

Figure 6. Normalized cumulative EMG signals for the four muscles (top left, FCR; bottom left, FCU; top right, ECR; bottom right, ECU), beginning 250 ms before movement onset.
Data shown are for the perturbation trials with assisting torque reduced by 100% (thick lines) and the corresponding pre-perturbation trials (thin lines), averaged over the 10 trials for subject DM. The vertical lines on the flexor and extensor plots represent the time at which the first oscillation occurred in the pre-perturbation trials, causing muscle stretch in the flexors or extensors, respectively. Differences in cumulative EMG were not statistically significant from start of data collection until the time of oscillation onset.

The preceding analysis might be interpreted as an incomplete characterization of the central motor commands because it does not address the state of the stretch reflex, which is integral to the λ-model. Under Feldman's λ-model the stretch reflex threshold is crucial in determining muscle activation to correct for any error in final position. If a muscle, either flexor or extensor, is stretched to a length longer than at the corresponding threshold angle, λ_{flexor} or $\lambda_{extensor}$, stretch reflex activation will produce muscular force to move the joint until the muscles and the load are in equilibrium. To test whether the state of the stretch reflex was altered when the assisting torque was reduced, we compared the

slope and intercept of the relation between integrated EMG and wrist displacement during oscillation about the final position for pre- and post-perturbation trials (see Methods). There were generally 2 cycles of underdamped oscillation about the final position for each trial, such that each regression was performed on approximately 100 data points. Because subjects consistently undershot the target when the motor torque was reduced an increase in flexor muscle activation would have been required to correct the error, so only the reflex response of flexor muscles was analyzed. According to the λ-model, an increase in reflex activation should have occurred because the flexor muscles would have been at a longer length than that specified by the central command. Under the λ-model only if the central command were altered so as to increase the stretch reflex threshold could it be claimed that an error would occur due to inadequate reflex activation. For completeness, we also considered the effect of the perturbation trial on the stretch reflex gain, as measured by the slope of the regression lines. Both stretch reflex gain and threshold are central in the α-model of EP control (Bizzi et al. 1982).

Representative relations between integrated EMG and wrist displacement are shown in Figure 7. Table 1 lists the intercepts and slopes, together with the significance levels for the 5 subjects for both FCR and FCU.

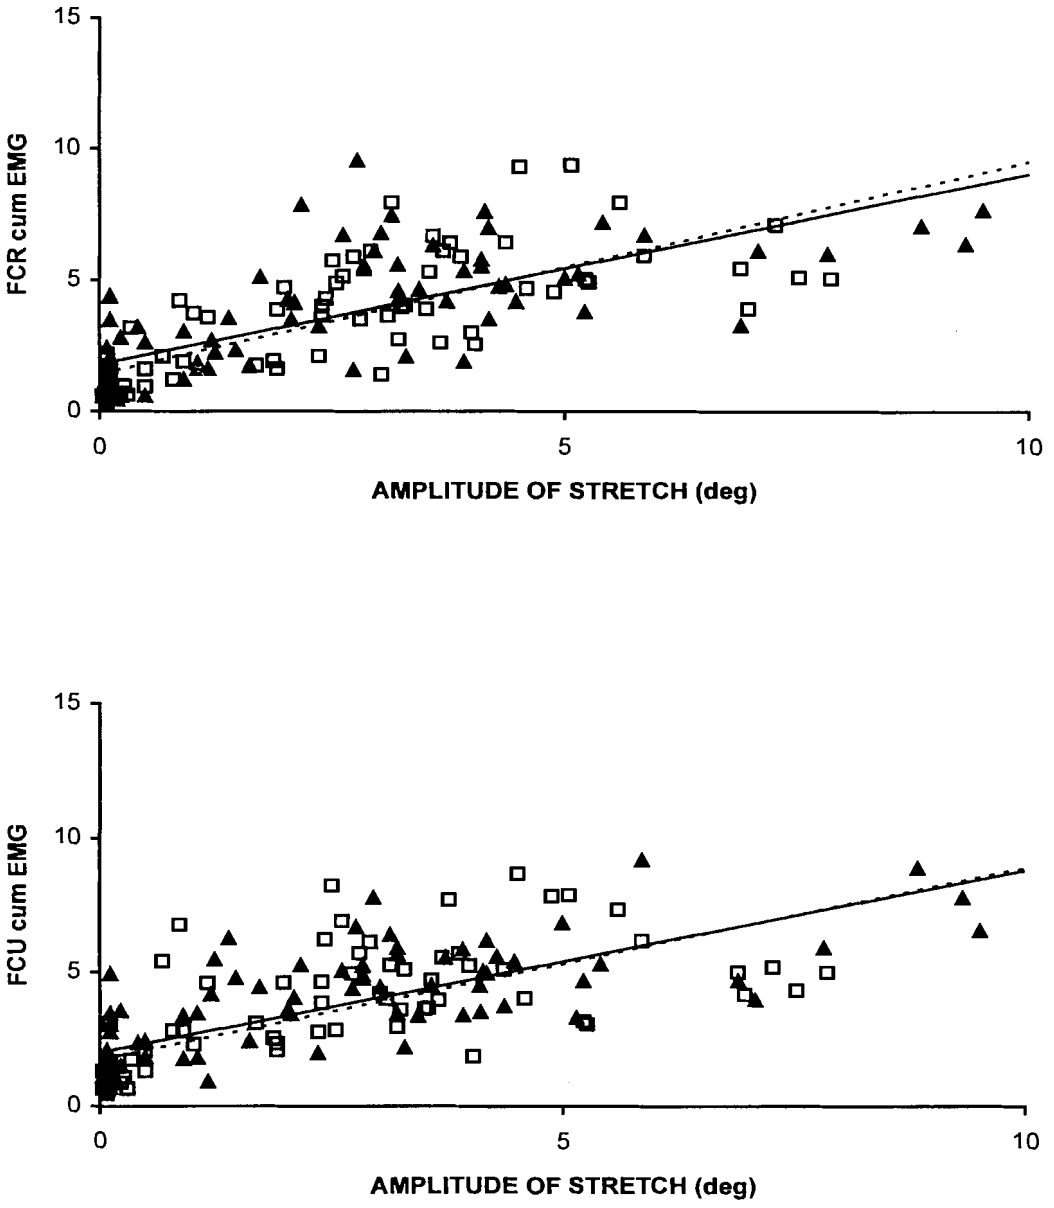

Figure 7. Linear regression between cumulative EMG and wrist displacement.
The 10 pre-perturbation trials (▲, solid line) and 10 post-perturbation trials (□, dashed line) for all 5 levels of torque reduction on the intervening perturbation trial were combined. The data for FCR are shown in the upper panel and for FCU in the lower panel for subject AN.

FCR	Pre-perturbation trials		Post-perturbation trials		p-values	
Subject	Slope	Intercept	Slope	Intercept	Slope	Intercept
mh	0.4807	3.9190	0.5017	3.7350	0.6763	0.5967
dm	0.3889	3.4141	0.3962	3.4389	0.8850	0.9898
an	0.7294	1.7823	0.8129	1.4225	0.3962	0.8484
al	0.4854	3.4876	0.5326	2.7879	0.6286	0.8522
bm	0.3259	4.9161	0.2998	4.6700	0.6079	0.9451

FCU	Pre-perturbation trials		Post-perturbation trials		p-values	
Subject	Slope	Intercept	Slope	Intercept	Slope	Intercept
mh	0.4985	3.8100	0.3877	4.4300	** 0.0105	* 0.0272
dm	0.8223	2.3927	0.8733	2.9735	0.8811	0.9656
an	0.6788	2.0127	0.7166	1.7443	0.7144	0.9025
al	0.7570	2.8660	0.6165	2.8601	0.2009	0.9988
bm	0.8857	4.1880	0.5914	5.8438	** 0.0218	0.8323

Table 1. Slope and intercept values for FCR and FCU for all 5 subjects.
Significant differences between pre- and post-perturbation trials are shown with a * (significant at p=0.05) or ** (significant at p=0.05 after multiple-comparison adjustment).

Correlation coefficients for the 20 regressions (pre-perturbation and post-perturbation for both flexor muscles, for each of the 5 subjects) ranged from 0.55-0.83 (R^2 values 0.30-0.74), with an average of 0.71. Student's paired *t*-tests revealed that for 9 of 10 comparisons (FCR and FCU for 5 subjects) there was no significant difference between the intercepts of the linear regressions for the pre- and post-perturbation trials (2-tailed test, *a priori* alpha 0.05, P> 0.60) When a family-wise adjustment was made in the critical *t* value to account for multiple tests, i.e., testing for both slopes and intercepts using the same data set, none of the 10 comparisons produced a statistically significant difference.

Student's paired *t*-tests for the slope (gradient) of the FCR regressions showed that in 8 out of 10 tests the slope in pre- and post-perturbation trials was not significantly different (2-tailed test, *a priori* alpha 0.05, P>0.20). The 2 cases where significant differences were found showed a decrease in the slope in post-perturbation trials relative pre-perturbation trials. Both of these tests remained significant when a family-wise adjustment was made in the critical *t* value to account for multiple tests.

Discussion

The results of our study clearly demonstrate that final equilibrium position is not preserved when the magnitude of an assisting torque, proportional to velocity, is unexpectedly reduced. This obviously violates the EPH. However, the strong correlation between the error in final position and the percentage by which the assisting torque was reduced matches the prediction based on the formation of an IDM of the specific task dynamics.

A fundamental prediction of the EPH is that final position should not be affected by a velocity dependent perturbation if the pattern of central commands to agonist and antagonist muscles does not change (Feldman et al. 1998). In the λ-model, the equilibrium point specified by the λ state, shifts from start to final position at some rate that determines the speed of the movement. An assisting torque that is proportional to velocity will affect the speed of the movement, but will not affect the final position specified by the λ state. Note that λ is not synonymous with final equilibrium position, which depends on load. The greater the load, the farther the final equilibrium position will be from λ because muscle torque must equal load torque. Once λ is fixed muscle torque increases by moving the joint angle farther from λ. In the case where a velocity dependent load is added, the added load disappears as movement is halted leaving the final equilibrium position unaltered. Our results unequivocally demonstrate that not only did subjects stop short of the target when the assisting torque was reduced, but the position error increased with the reduction in assisting torque. End point error was as large as 6 °, which represents 20% of the movement amplitude. Although we cannot definitively state that subjects could perceive this error, it is more than 2.5 times greater than the standard deviation in final position on unperturbed trials. The relative error, as a proportion of movement distance, is the same whether the error is expressed in terms of joint position or end-point (hand) position. Therefore, it is unlikely that errors as large as 6 ° could be considered as subthreshold for correction by the EPH. End point errors could not be attributed to a change in the pattern of feedforward central commands since there

was no difference in the cumulative EMG of any muscle, until later in the movement when oscillations ensued about the final position.

As well as showing maintenance of the feedforward motor commands for all levels of torque reduction, we investigated changes in the stretch reflex during oscillations about the final position. We found no evidence of an increase in stretch reflex threshold (higher intercept) between pre- and post-perturbation trials for any subject in either the FCR or FCU. Similarly, when we considered the reflex gain, 8 out of 10 comparisons showed no evidence of a change in the reflex gain. The average final position on pre-perturbation trials was 29.7 ° (s.d. 2.23 °), almost at the centre of the target window. A reduction in stretch reflex gain would, of course, produce lower muscular activity to correct for any residual end point error, relative to the programmed equilibrium position. Assuming that the programmed final equilibrium position was at the centre of the target window i.e., 30 ° then the residual error would have been 0.3 °. Given that in the perturbation trial the final commanded equilibrium point defined by the R command remains unaltered, a change in the reflex gain would magnify the small residual error observed in pre-perturbation trials. A reduction in reflex gain of 50% on the perturbation trial relative to the pre-perturbation trial, much larger than our results suggest, would increase the final position error from 0.3 ° to 0.6 °. This is much smaller than the errors which we observed, indicating that observed changes in reflex gain could not be responsible for final position errors on the perturbation trials.

It might be argued that when the movement slowed, friction in the apparatus caused motion to stop and that the driving torque arising from the position error was insufficient to overcome the friction. This is unlikely for several reasons. First, the friction torque of the apparatus is small (0.05 Nm). Second, under three of the perturbation conditions where endpoint errors occurred (25%, 50% and 75% reduction in assisting torque) the assisting torque from the motor would have counteracted the friction torque. Furthermore, because of the negative damping effect of this assisting torque, there was often a transient overshoot of the target position even when position finally stabilized

short of the target (Figure 3). Third, wrist stiffness tends to be greater than 0.1 Nm/deg after adaptation to this type of assisting torque (Milner and Cloutier 1998). Consequently, the final position error should have been reduced to less than 0.5 °, according to the EPH, even without any compensation for friction. The fact that the average final position error ranged from about 1.5 ° to 6 ° belies such an explanation.

The results can be readily explained if the CNS employed an IDM of the dynamics of the assisting torque to control movement (Shadmehr and Mussa-Ivaldi 1994; Conditt et al. 1997; Thoroughman and Shadmehr 2000; Scheidt et al. 2000). In this case, we assume that subjects form an IDM of the relation between velocity and assisting torque through practice. The CNS learns the muscle activation patterns to produce the torque needed to reach the target position. Specifically, it learns the requisite wrist flexor torque to accelerate toward the target at the specified speed and the requisite extensor torque to counteract the assisting torque during deceleration so as to stop in the target window. It also produces sufficient co-contraction of agonist and antagonist muscles to ensure stability (Milner, 2002). Because the assisting torque was reduced on perturbation trials, the flexor torque required to reach the target was increased while the extensor torque required to stop the movement was reduced. However, the CNS issued the same feedforward motor commands on all trials, as judged by the similarity in cumulative EMG up to the first velocity zero crossing. This is precisely as expected from a feedforward control scheme using an IDM and would result in the observed final position errors. There would be a deficiency in flexor torque and an excess extensor torque, equal to the reduction in assisting torque. Consequently, the movement would stop short of the target by an amount that increased with the percentage by which the assisting torque had been reduced.

Several previous studies have yielded evidence which casts doubt on the principle of equifinality. However, in none of them was equifinality completely violated nor was an alternative control hypothesis tested. Rothwell et al. (1982) found undershooting of the target position under conditions where muscle was areflexive. Since this only occurred

during faster movements and did not occur when stretch reflexes were intact, it could not be considered as evidence against the λ-model of equilibrium point control. Similarly, Sanes and Evarts (1983) and Sanes (1986) found that unexpected changes in viscous loading produced final position errors, but only for small amplitude movements. Consequently, these errors may have been below some hypothesized threshold for stretch reflex compensation, as suggested by Feldman et al. (1998). This criticism, and the possibility of a change in motor command, weighs against the findings of Lackner and Dizio (1994) as providing definitive evidence against the EPH. The results of Popescu and Rymer (2000) were equivocal. They did not show that equifinality was consistently violated. Furthermore, their experimental paradigm involved relatively small amplitude movements unopposed by any external load, where stiffness would have been low, a situation in which final position error would be expected given the elastic nature of the internal load (muscle and tendon).

Our test of the EPH is stronger than in any previous study. We created a situation where high joint stiffness was required and trained subjects to perform movements of moderate amplitude and speed. Furthermore, we rigorously examined the possibility of a change in motor command both from the perspective of change in feedforward command and change in stretch reflex threshold or gain. There should be no doubt that if the EPH was valid then equifinality should have held under our experimental conditions. The final position errors that we found were inconsistent with the EPH, but entirely consistent with *a priori* predictions based on feedforward control by an IDM. In addition, the similarity of the cumulative EMG on pre-perturbation and perturbation trials supports the use of such feedforward control. Comparison of the final position on pre-perturbation and post-perturbation trials provided more supportive evidence. Not only did we find no difference in final position, but also there was no effect of the magnitude of the intervening perturbation trial. This indicates that the feedforward motor commands were robust and could not be altered by occasional, unexpected changes in the relation between assisting torque and velocity, when subjects were aware that that they any changes in load were only random trials, i.e., they were not trying to adapt to the dynamics of the perturbation

trial. Although we have not tried to show that the EPH fails universally since we examined only the case of wrist flexion under an assistive load, our results are in clear contradiction to the EPH. Furthermore, we have demonstrated that the results are compatible with an alternative movement control strategy, which postulates formation of an IDM.

We have presented evidence both from movement kinematics and patterns of muscle activation that provides strong support for the theory of IDM formation, but contradicts the EPH. This conclusion is not restricted to the λ-model, but applies generally to all formulations of the EPH. However appealing the EPH may be, we have shown that one of its fundamental predictions, equifinality, does not hold under certain conditions. We conclude that learning leads to the formation of an IDM of the load dynamics by the CNS, which is used for feedforward control of motor commands to muscles.

In experiment 2 we take the result from experiment 1 and consider the kinematic and EMG patterns when a subject learns a double joint task. Four main questions will be addressed with the aim of elucidating the mechanisms employed during learning to move in a novel force field. 1) What are the relative contributions of IM formation and impedance control in early adaptation to a novel force field 2) How does the level of cocontraction vary with repeated exposure to the field? 3) How quickly can subjects adapt their IM when the expected environmental force is perturbed? 4) How does the level of cocontraction vary when perturbed field strengths are initially encountered and subsequently learnt?

Experiment 2. Adaptation to Novel Dynamics- an Investigation into the Short- and Long-Term Stages of Learning

Introduction

This experiment was based on the premise that, as experiment 1 and considerable published research suggests (Shadmehr and Mussi-Ivaldi, 1994; Lackner and Dizio, 1994; Gomi and Kawato, 1996; Gandolfo et al, 1996; Goodbody and Wolpert, 1998; Throughman and Shadmehr, 2000; Scheidt et al, 2000, 2001; Takahashi et al 2001), subjects learn an internal model (IM) after prolonged exposure to a novel force field which compensates for all predictable external forces. How we learn IM's for novel force fields, and the way in which the IM's are adapted when the force field is perturbed were considered in this experiment.

Whilst developing an IM of the task dynamics it is possible that subjects may use impedance control, in the form of cocontraction of antagonistic muscles, to reduce the hand displacement in trials before the IM is accurate enough to compensate for the external forces. Prior research suggests that impedance control, in terms of cocontraction, is utilized when subjects are exposed to destabilizing force fields (Milner and De Serres, 1991; Milner and Cloutier 1998, Burdet et al 2001). It has also been reported that the level of cocontraction in a movement is adapted to the amount of stability in the task (Milner 2002).

In this study subjects learnt to make accurate movements in a field that was subsequently perturbed, at random, to a lower or higher strength for a number of trials. By asking subjects to adapt to the perturbations we investigated the relative roles of IM formation and impedance control (cocontraction) in subsequent adaptation.

We hypothesized that stretch reflex EMG, resulting from trajectory error in perturbed trials, would provide an indirect estimate of the afferent error signal available to the brain. How this feedback is incorporated into a feedforward model on the subsequent

trials was investigated. A model was developed to describe the process, by which changes to the IM or the mechanical impedance, through impedance control, occur during learning.

Methods

Equipment and Set-Up

Multi-joint movement of the elbow and shoulder was studied using the Parallel-Link Direct-Drive Air and Magnet Floating Manipulandum (PFM) at ATR, Kyoto, Japan. Two-dimensional movement occurs in the horizontal plane. Details of its design and operation have been described previously (Gomi and Kawato, 1997)

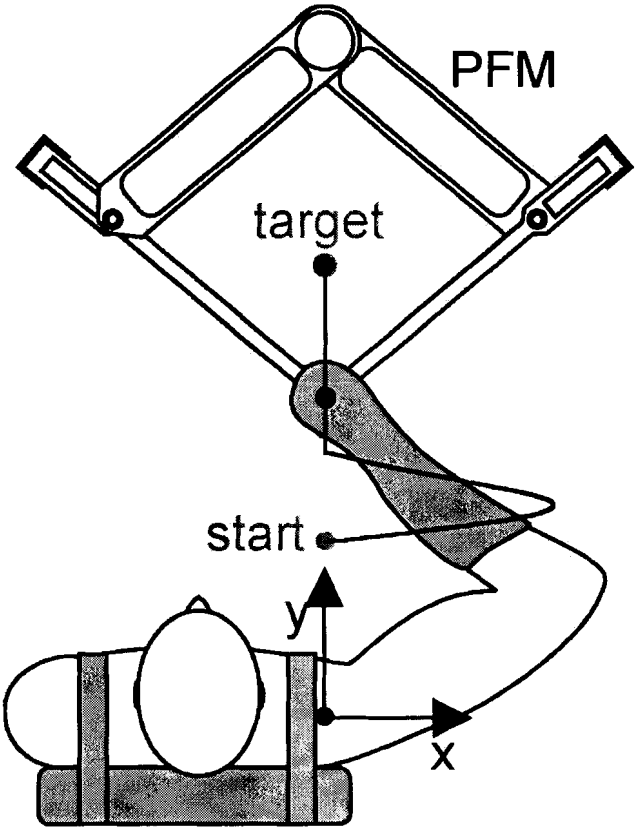

Figure 8. Plan View of the PFM.
Subjects made reaching movements of 0.25 m, from the start position, located 0.31 m from the shoulder. The parabolic force profile is shown schematically. The force acts over the first 0.10 m of the movement, with a peak x force of 6 N occurring 0.05 m from the start position. (Adapted with the permission of David Franklin.)

Subjects sat in a chair with a harness to constrain the trunk so that the elbow and shoulder joints could only move in the horizontal plane. The forearm and wrist were held in a thermoplastic splint rigidly attached to the manipulandum, constraining movement to two degrees of freedom. Forces were applied to the hand by means of two torque motors driving the parallel linkage.

The chair's height was adjusted such that the arm was in approximately 80 degrees abduction. A circular cursor 0.5 cm in diameter, representing the current hand position, was initially positioned in a 2.5 cm start circle, the centre of which was located 0.31 m directly in front of the shoulder. The cursor, as well as the start and target circles were projected onto an opaque horizontal screen suspended above the arm, such that subjects could not see their arm during trials. Subjects made 0.25 m horizontal point-to-point movements away from the body along the y-axis (sagittal plane) to a 2.5 cm diameter target circle. The prescribed movement time of 600 ms was indicated by acoustic signals. Subjects initiated each trial on hearing the third beep of a series, and tried to complete the movement on the first beep of the second set of three beeps. The last two beeps of the second series indicated the time that the subject should remain in the target circle before returning towards the start position. No force acted on the hand until movement was initiated on each trial. The final hand position (OK or OUT) and movement duration (OK, LONG, or SHORT) were presented to subjects on a screen. The final position was deemed OK if the target circle was located accurately. The duration was deemed OK if it was within $\pm$ 100 ms of the desired movement time. Subjects were instructed that their goal was to produce movements which always met the OK criteria. The movement duration and final hand position feedback were presented primarily as incentive for subjects to improve performance. However, all trials, irrespective of whether the duration and hand position were deemed OK, were used in subsequent data analysis. Each trial was self-initiated by placing the cursor in the start position, enabling subjects to rest between trials, if desired.

A bipolar arrangement of silver-silver chloride disposable electrodes was used for recording surface EMG. The surface of the skin was prepared using alcohol wipes, skin abrasion and electrolyte gel to ensure an interelectrode resistance of less than 10 kΩ. The EMG signals were sampled at 2000 Hz and were filtered using a 25 Hz (high pass) and 1kHz (low pass) filter. EMG signals were acquired beginning 150 ms prior to movement onset for 1400 ms from two monoarticular elbow muscles (the lateral head of triceps and brachioradialis), two monoarticular shoulder muscles (pectoralis major and posterior deltoid) and two biarticular muscles (biceps brachii and the long head of the triceps).

A position dependent force, in the form of a parabolic force field (PF), was used during this experiment. Prior work using the PFM has shown that the application of velocity dependent force fields leads to a progressively increasing position error until near the end of trial. In such experiments subjects remained relatively close to the desired straight path for the first half of the movement, before diverging from the unperturbed path later in the trial. To investigate how error feedback related to the stretch reflex might be incorporated into the feedforward internal model it was preferable to have the displacement occurring near the beginning of the movement. This makes it easier to distinguish changes to the feedforward command from voluntary on-line corrections.

Force in the lateral (x) direction was dependent on the position of the hand along the direction of movement in the trial (y axis). The maximum PF force was dependent on the amplitude of the parabola, and the distance over which the force acts. The PF is given by

$$\begin{aligned} F_x &= -G(y_{start} - y)(y - y_{end}), & -(y_{start} - y)(y - y_{end}) \geq 0 \\ F_x &= 0, & -(y_{start} - y)(y - y_{end}) < 0 \end{aligned}$$

where F_x is the force in the x direction (positive to the right), y_{start} and y_{end} define the boundaries of the force field, y is the current location and C is the field strength. The field parameters were set such that the field was 'active' over the first 0.1 m of the 0.25 m reaches. G was set to produce a maximum force of 6 N in the PF.

Following displacement from a straight-line path when subjects unexpectedly encounter a perturbation in the PF strength, the stiffness of the arm will naturally cause the subject's hand to return towards the straight path once the perturbation is complete, i.e., over the last 0.15m of the movement. The maximum force in the x direction was chosen following pilot studies. We chose a field strength that resulted in perceivable hand perturbations, relative to the learnt path when no external forces were present (null field). This ensured that subjects had incentive to adapt to the new field, correcting both final position and straightening hand paths. The perturbed strength fields were then commanded to be 33% weaker and stronger than the chosen PF strength (see Experimental Protocol).

Experimental Protocol

Subjects who were naïve to PFM experimentation underwent a training period of two sets of 100 trials in the null field (NF) to familiarize themselves with the equipment, movement duration and general experimental procedure. This session was conducted on the morning of the actual experiment without any EMG recording. Learning of the PF consisted of 150 trials in which subjects aimed to produce straight hand paths as possible and accurately locate the target. Visual feedback was available throughout the experiment. The cursor representing hand position was visible throughout the task. The start and target circles were also always visible. The inherent latency in using the visual feedback ensured that the first 200 ms of the motion would not involve any on-line use of visual information. Feedback of final position could be used to modify performance on subsequent trials during learning.

Every 5^{th} PF trial was replaced with a channel trial (CT) which restricted subjects to a straight line between the start and target positions. The "walls" of the channel were created by commanding the motor to act like a very stiff spring, which restrained subject's lateral movement. CT's allowed us to analyze the feedforward motor command (analyzed by EMG) in the absence of any stretch reflex activity. We were also able to analyze the force produced by the hand in isolation from motor force. More information related to the PFM set-up has been previously reported (Gomi and Kawato, 1996)

Following the 150 trial PF learning period subjects had a short break, to prevent any possible effects of fatigue. Subjects were re-familiarized with the PF for 27 trials, before the strength of the PF was perturbed 24 times, for 5 consecutive trials at a time, to either a higher or lower strength. ΔPF_{low} trials constituted a reduction in the PF strength such that the maximum lateral force produced by the PFM was 33% less i.e., 4 N, whereas ΔPF_{high} trials had a maximum lateral force 33% higher than the PF i.e., 8 N, as shown in figure 9.

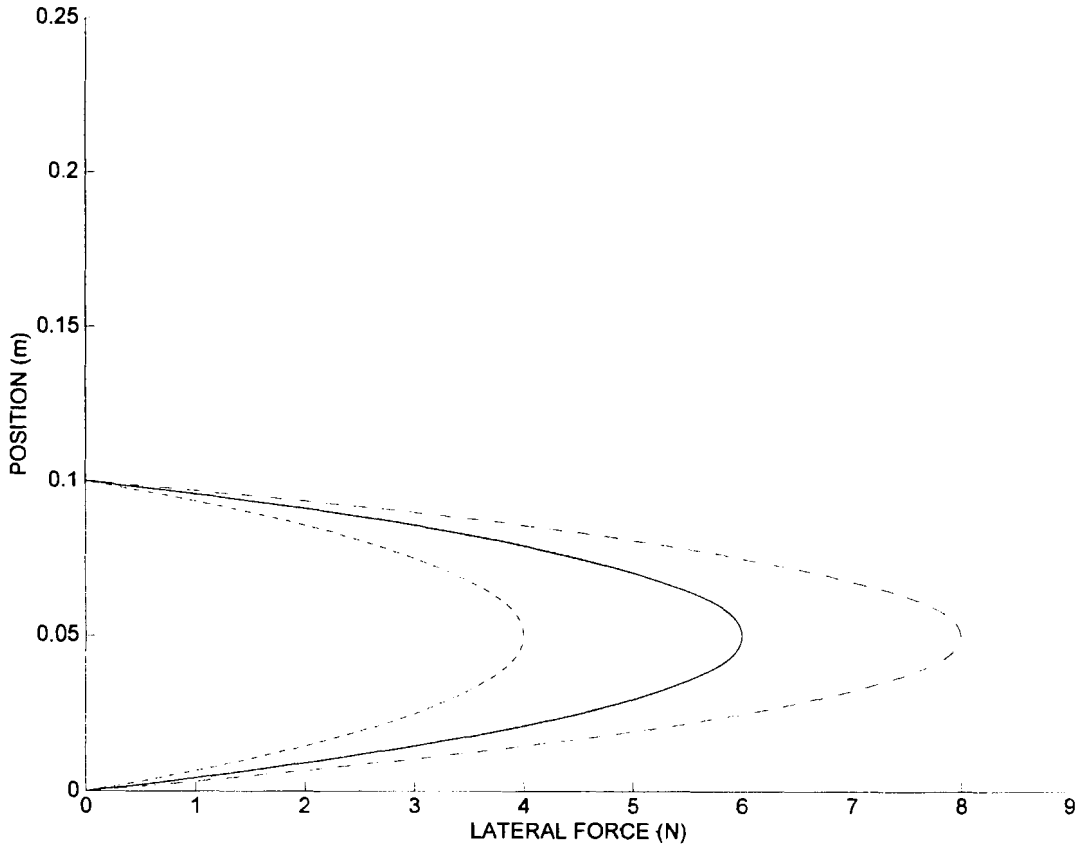

Figure 9. Force in the lateral (x) direction as described by the parabolic field (PF).
The PF was active over the first 0.1 m of the 0.25 m amplitude movements. The learnt field (PF) is shown with the solid line, the high and low strength perturbation fields are shown in the dotted and dashed lines, respectively.

Subjects performed 12 sets of 5 trials for each ΔPF, in a randomly generated order. Channel trials (CT) occasionally replaced ΔPF trials in the second, third, fourth and fifth trial of the 5 trial sequences. The trial on which the CT occurred was selected randomly

but with equal probability, such that a CT occurred in each set of 5 trials, albeit in a different position. The CT occurred 3 times in each position of the sequence i.e., 25% of the ΔPF trials in positions 2-5 were CT's. Subjects were instructed to try to adapt to the new fields and produce as straight a path to the target as possible. Between each perturbed set subjects were re-exposed to the PF for 6-8 trials, selected randomly, in which again they were instructed to make straight paths. Pilot experiments had revealed that significant hand path straightening in the ΔPF occurred within 5 trials and on returning to the learnt PF, reproduction of learnt hand paths occurred within 2-3 trials. 6-8 PF trials were used between each set of ΔPF trials to allow any cocontraction to fall associated with the change in field to drop, such that both performance and motor commands in the PF prior to any perturbations could be re-developed.

This experimental protocol allows us to consider how subjects develop IM's in the PF over a large number of trials. By exposing subjects to only 5 trials in which the field strength was altered at a time, we do not permit subjects to fully adapt to the ΔPF. thereby recreating the stages of early learning 12 times in each perturbed field. This allows a more thorough examination of the control strategies within this period. In total 315 trials were performed, in addition to the 150 trial practice session. The experiment was completed within 90 minutes.

Results-PF learning

To analyze changes in subjects' performance from trial to trial we considered a number of kinematic variables relating to hand path. These analyses allowed us to determine to what extent subjects were able to learn in the PF, and where within each trial most of the learning, in terms of straightening the hand paths, occurred. Although changes in hand path will give an insight into the learning, they do not reveal what mechanisms or processes subjects utilize to enable them to adapt to the novel dynamics. A thorough analysis of force and EMG data was also undertaken to ascertain the mechanisms subjects employed to compensate for the external force, and determine how these may have changed with repeated exposures to the dynamics of the force field.

In all analyses requiring repeated measures ANOVA, the assumption of sphericity was tested using Mauchly's test. When violations of the assumption were detected, i.e., $\varepsilon<1.0$, sphericity-corrected repeated measures analysis of variance tests were performed. Greenhouse-Geisser corrections were implemented when $\varepsilon<0.7$, while Huynh-Feldt corrections were used when $0.7<\varepsilon<1.0$ to downward-adjust the degrees of freedom (Howell 1997).

Prior to analysis of PF learning trials, the effect of imposition of CT trials was considered. Two important issues were addressed. Firstly, we determined whether the CT imposed any significant perturbation to the hand, relative to the path that would have occurred on a PF trial. Secondly, we considered the effect of the CT on subsequent PF trials. Only after these questions were addressed could we determine how best to analyze the learning.

Comparison of learnt hand paths in the NF and PF

The use of CT's as a means of investigating force relies on the presumption that the CT *does not* impose any significant perturbation to the expected hand path, i.e., CT's should restrict subjects to a path which is not significantly different from the path they would otherwise follow in the PF, and subjects should not experience forces during a CT that are significantly different from those which they would experience on a PF trial. Hand paths in a null field (NF), i.e., a field where the PFM produces no additional force, other than to compensate for the friction, and the PF field were compared for 7 of the subjects who had also participated in a recent experiment involving numerous trials conducted in a NF. CT's have previously been used successfully when implemented within other force fields, e.g., velocity dependent and divergent (curl) force fields, where the learnt path was not significantly different from that in the NF.

To determine whether hand paths in the PF were significantly different from hand paths in the NF we compared the last 8 NF trials (learnt-NF) prior to introduction of any novel field in the previous experiment to the last 8 PF trials (learnt-PF) in the present

experiment, prior to the introduction of any ΔPF trials. If there was no significant difference in the NF and PF paths, we would be able to analyze CT's in a similar manner to previous experiments.

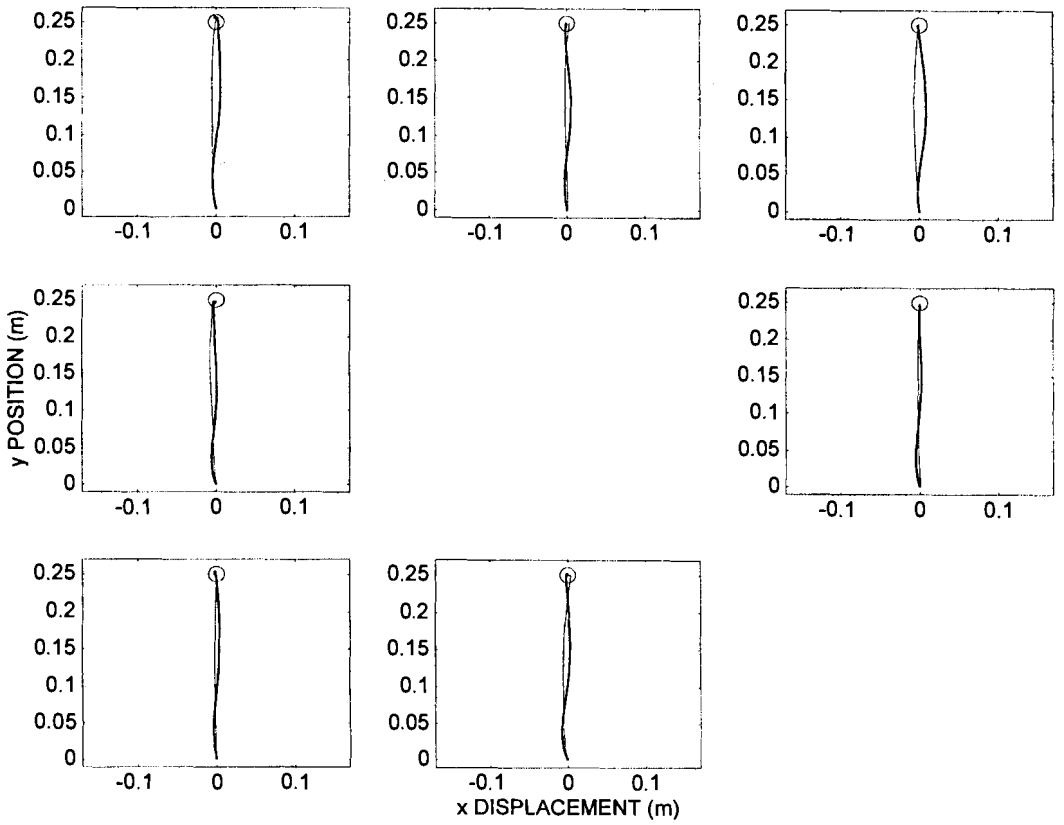

Figure 10. Learnt hand paths in the NF (thin) and PF (thick) for the 7 subjects who participated in a previous study in which trials in the NF were recorded.
Data are averaged over the last 8 trials in each case prior to exposure to other novel fields. Top left-bottom right subjects MH, US, AN, JJ, SM, DF and RO.

Curvature, described by the equation

$$C = \frac{\dot{x}\ddot{y} - \dot{y}\ddot{x}}{(\dot{x}^2 + \dot{y}^2)^{3/2}},$$

where $\dot{x}$, $\dot{y}$, $\ddot{x}$ and $\ddot{y}$ represent the instantaneous x (lateral to the direction of the reach) and y (the direction of the reach) linear velocities and accelerations, was used as a measure of path straightness. Curvature was calculated for each data point, beginning 5 mm from the start position (where velocity in the direction of reach was approximately

0.04 $m \cdot s^{-1}$), until the tangential velocity fell below 0.05 $m \cdot s^{-1}$. In this manner, we were able to calculate curvature for the majority of data points within a single reach, whilst disregarding the regions at the beginning and end of movement where very low velocity resulted in anomalously high curvature, and high curvature regions due to online, corrective movements near the target.

To characterize performance we determined the cumulative positive curvature and cumulative negative curvature in each of the learnt-NF and learnt-PF trials, for the 7 subjects for whom NF data was available. These values were subsequently divided by the number of points in each respective trial where the curvature value was either positive or negative to obtain mean positive and mean negative curvature. These two distinct measures were used in evaluating performance, realising more information than simply considering mean curvature averaged over the whole trial.

We compared the mean positive and negative curvatures in the learnt-NF and learnt-PF trials using two separate two-way, repeated measures (RM) ANOVA. Results revealed no significant difference in the mean positive curvature between the two fields (9.78 m^{-1} for the NF versus 7.60 m^{-1} for the PF; field main effect: F=0.324, p=0.590). The trial effect and interaction were also non-significant (trial main effect: F=0.468, p=0.642; interaction effect: F=0.1.039, p=0.381 using Greeenhouse-Geisser corrections for sphericity on both results), highlighting the fact that there were no significant trends amongst the 8 trials in each field. When we considered the mean negative curvature we found a significant field main effect (-20.29 m^{-1} for the NF versus –68.05 m^{-1} for the PF; F=21.205, p=0.04), confirming greater negative curvature in the PF trials compared to the NF trials. There was no significant trial or interaction effect (trial main effect: F=0.1.039, p=0.381; interaction effect: F=0.877, p=0.473 using Greeenhouse-Geisser corrections for sphericity on both results).

Taken together, these results reveal that the learnt-NF and learnt-PF paths *are* statistically different in terms of trajectory curvature. The larger negative curvature in the

PF can be wholly accounted for within the very early period of each trial, where subjects move laterally "into" the applied field, at relatively low velocity (see PF learning-kinematics).

Although the differences in curvature suggest that there are physical differences between the NF and PF hand paths, this does not automatically require us to amend analyses that have been used in previous studies to analyze CT's if the differences are sufficiently small. In particular, if the force applied to the channel wall, due to the channel straightening the PF hand path, was no larger for our PF, than for force fields in previous studies, then the analysis would still be valid. This force would arise primarily from the inability to accelerate mass in the lateral direction. Any force that would normally be developed to accelerate mass laterally would be experienced as a constraint force produced by the channel wall, which resisted lateral acceleration. We considered the maximum lateral acceleration *away* from the start-target line, for both the learnt-NF and learnt-PF trials, in the period in which the PF would have been active, i.e., over the first 0.1 m of the reach. A two-way RM ANOVA revealed that the lateral acceleration in the 8 learnt-PF trials was larger than the lateral acceleration in the 8 learnt-NF trials ($F=14.168$, $p=0.009$). The maximum lateral acceleration in the PF trials, averaged over the 8 trials and 7 subjects, was 1.71 times the magnitude of that in the NF trials (12.93 ms^{-2} and 22.48 ms^{-2} in the NF and PF, respectively).

Lateral forces recorded in NF trials, over the first 0.1 m of the trial, were generally less than 0.5 N, and did not exceed 0.7 N in either direction. Since lateral acceleration on PF trials was approximately 1.7 times that of NF trials this force should not have been greater than 1.2 N on PF trials. However, because muscle shortening velocity may have been slowed slightly by the channel constraint, the force applied to the channel wall on CT's, which replaced PF trials, may have been slightly higher. We note that this estimate does not include the force that would have been used to accelerate the human arm.

To estimate the force that would have been applied to the channel wall as a result of preventing the subject's arm from accelerating laterally we first computed the joint torques needed to accelerate the arm. We used the joint kinematics corresponding to the learnt-PF paths for each subject to compute the shoulder and elbow joint torques from inverse dynamics. We transformed joint torques (τ) into the x and y components of force (F), using the following relation.

$$\boldsymbol{F} = (\boldsymbol{J}^{T})^{-1}\tau$$

J is the Jacobian, a two-dimensional matrix that relates the joint angles and segment lengths. It was found that the peak lateral force (x component) would be 2-3 N, in the leftward direction. A profile plot revealed that this force would peak very shortly after the onset of movement.

In a similar manner, we estimated the force that would have been applied to a channel wall as a result of preventing the subject's arm from accelerating laterally in NF trials. Since the transformation from the lateral acceleration to force used in accelerating the arm inertia is non-linear, this force could not be assumed to be 1.71 times less than the force we determined for the PF trials. We estimated that the force on NF trials would have been, in general, slightly lower than on PF trials, with peak values between 1.2-3 N. We combined the estimated forces due to preventing both the lateral acceleration of the PFM and the arm. In the NF, this force would have been around 1.9-3.7 N, whereas in the PF a somewhat higher force of 3.2-4.2 N would have been expected.

To complete the kinematic comparisons between NF and PF trials, we compared maximum deviation from a straight line joining the start and target positions. Figure 10 suggests that maximum deviation from a straight line was similar in NF and learnt-PF trials. This was tested statistically using a two-way RM ANOVA, for the 8 NF trials and 8 learnt-PF trials. We actually found maximum deviation in the learnt-PF trials was larger than that in NF trials (field type main effect $F=13.95$, $p=0.01$), though the difference between the mean maximum deviation in the two field types was very small. Mean

maximum deviation in the NF was 8.05 ± 1 mm, while mean maximum deviation in the PF was 10.4 ± 1 mm.

In summary, we have found that there are indeed significant differences between the learnt-NF and learnt-PF paths, in terms of curvatures, lateral accelerations and maximum deviations from a straight line joining the start and target positions. Of most importance in the analysis of channel trials was our estimation of the perturbing effect due to the channel straightening the learnt hand path. We estimated that forces applied to the channel wall, caused by the channel preventing lateral acceleration of both the PFM and the human arm were somewhat higher for PF trials than for NF trials. However, our force estimates are similar in magnitude to the lateral forces recorded by Scheidt et al (2000), when channel trials replaced null field trials. Peak forces recorded during null field trials in their experiment were generally less than 5 N, but peak forces recorded on channel trials following null field trials were generally larger, suggesting that they measured additional force due to the constraint on lateral acceleration of the human arm and the manipulandum. Given that our channel force estimates are of a similar magnitude as those reported by Scheidt, we can assume that the perturbing effect of the channel in our experiment was no greater than in their experiment. Therefore, we should be able to draw equally valid conclusions about the nature of feedforward commands from analysis of the force profiles during CT's.

Did CT's induce a stretch reflex?

Another important question when determining if the CT's significantly perturb subjects is whether the change in path in CT's, relative to the expected path, i.e. the learnt-PF path, caused perturbations which were sufficiently large, i.e., of sufficiently high velocity, to evoke a stretch reflex response. Such a reflex response would limit the use of EMG data in CT's as a measure of subject's feedforward motor command, issued on the premise that they would encounter the PF, since the EMG would then also include a reflex component which could not be separated from the feedforward component.

Shoulder and elbow joint angles were calculated throughout the movement in CT's and learnt-PF trials, using inverse kinematics from the end point position of the limb. Angles were calculated for the hand path averaged over the last 5 CT's during the 150 trial learning period and for the learnt-PF path. We calculated the difference in both shoulder and elbow angle between the CT and learnt-PF trials, at each time point throughout the trial, before differentiating with respect to time to determine the rate of change of the difference in joint angle. These "velocities" were compared to joint perturbation velocities used in studies that had successfully evoked significant stretch reflex responses.

We found that the peak rate of change of difference in shoulder angle did not exceed 25 $°s^{-1}$ across the 9 subjects, while peak rate of change of the difference in elbow angle was less than $25°s^{-1}$ in 8 out of the 9 subjects, and did not exceed 50 $°s^{-1}$. Nakazawa et al. (2001) used elbow velocities in the range 50-300 $°s^{-1}$ to evoke stretch reflexes in the elbow flexor and extensor muscles. Their results indicate that stretch reflex responses at 50 $°s^{-1}$ were relatively small (<5% MVC), whereas higher velocities elicited responses as high as 35% MVC. The peak rate of change of the difference in joint angle between the 4^{th} PF learning trial, i.e., the trial immediately prior to the first CT trial in the learning period, and the following CT did not exceed 50 $°s^{-1}$ for the shoulder or 110 $°s^{-1}$ for the elbow, across the 9 subjects. When we considered the 9^{th} trial and the following CT, i.e., the second channel trial in the PF learning period, peak rate of change of the difference in joint angles did not exceed 50 $°s^{-1}$ or 90 $°s^{-1}$ for the shoulder and elbow. In the 14^{th} trial and 3^{rd} CT, peak rate of change of the difference in joint angle did not exceed 20 $°s^{-1}$ for the shoulder and was generally less than 50 $°s^{-1}$ at the elbow. The elbow joint of one subject (RO) was perturbed by the CT at a higher rate, i.e., around 90 $°s^{-1}$. This was because this subject's PF paths were less straight than that of other subjects during the early stages of learning.

Since Nakazawa et al. (2001) suggest a 80-100 $°s^{-1}$ perturbation velocity evokes a stretch reflex of the order 12% of MVC, we assume that some stretch reflex activity was evoked

in the first 2 CT's of the learning set, but thereafter the rate of change of the difference in joint angle suggests that little, if any, reflex activity would have been evoked. We also note that the peak values of the rate of change of the difference in joint angles occurred later than 200 ms into the movement. Given that we shall only use the initial part of the movement in CT's to investigate feedforward motor commands, stretch reflex responses this late in the trial would have little or no effect on our results.

Despite the statistical differences in curvature, acceleration and maximum deviation between the NF and PF trials, we believe that CT's will still be a valid means of investigating feedforward commands to compensate for the PF, at least within the region where the PF was active. Any force caused by their perturbing effect should only amount to a small proportion of the total force recorded in such trials. The force recorded on the channels will provide a direct measure of the level of force that subjects actively produce to compensate for the PF. In addition, we have shown that the change in joint velocity at both the shoulder and the elbow is much smaller than that used in previous experiments to induce stretch reflex responses in muscles. This allows us to use EMG data in the initial part of CT's as a measure of the feedforward motor command, in the absence of any stretch reflex activity.

The effect of CT's on subsequent PF trials

We continued by considering the effect of CT's on subsequent PF trials. Due to the very nature of the CT, no error feedback is available on these trials, i.e., subjects are unable to judge how accurate the movement resulting from their issued motor command would have been in the PF, relative to a straight line (or their learnt path). Based on the work of Scheidt et al. (2000), this lack of error feedback would be expected to adversely affect performance on PF trials that immediately follow CT's (post-CT trial). This decrement in performance would be visible as higher hand path error in the post-CT trial, compared to the preceding PF trial (pre-CT trial). We hypothesized that early in the learning period repeated exposure to the PF would lead to progressively lower trajectory error as subjects adapted to the PF. Therefore, if we found higher, or even no change in the hand path

errors in post-CT trials compared to pre-CT trials, we would conclude that CT's adversely affect learning. Later in learning, when the feedforward motor command should have been more accurate we hypothesized that the lack of error feedback during the CT's would have had less of an effect on post-CT trials, although this may still have been significant. If CT's were found to significantly affect the performance in the subsequent PF trial, it would be necessary to take this into account in analysis of the progressive learning within the PF.

We considered the absolute hand path error between the actual hand path and a straight line between the start and target positions as a measure of hand path error, and referred to this as absolute hand path error. Absolute hand path error was an applicable measure in this situation as we were interested in the overall effect of the CT, and less concerned about the point within the trial where error occurred.

A two-way RM ANOVA over the 9 subjects revealed a significant increase in absolute hand path error in post-CT trials relative to the corresponding pre-CT trials ($F=5.549$, $p=0.046$). The trial number within the PF learning set had a significant effect ($F=2.562$, $p=0.048$), suggesting, as expected, that subjects overall performance, averaged across pre- and post-CT trials, improved throughout learning. Somewhat surprisingly, we found a non-significant interaction between the trial position and trial number ($F=1.211$, $p=0.326$), suggesting that even though subjects compensated more accurately for the PF later during the learning period, the lack of error feedback in CT's had a similar adverse affect on performance (Figure 11A).

A similar two-way RM ANOVA (Figure 11B) showed no significant difference in performance between pre-CT trials and the second trials after imposition of a CT, referred to as post-CT+1 trials ($F=0.640$, $p=0.447$). This highlights that error feedback available in post-CT trials had been successfully used in modifying the feedforward command issued on the post-CT+1 trials, such that performance improved to a level similar to that in the pre-CT trial. This recovery of performance on the post-CT+1 trials

was also evident in a two-way RM ANOVA, which showed a significant reduction in absolute hand path error in post-CT+1 trials relative to post-CT trials (F=5.95, p=0.041).

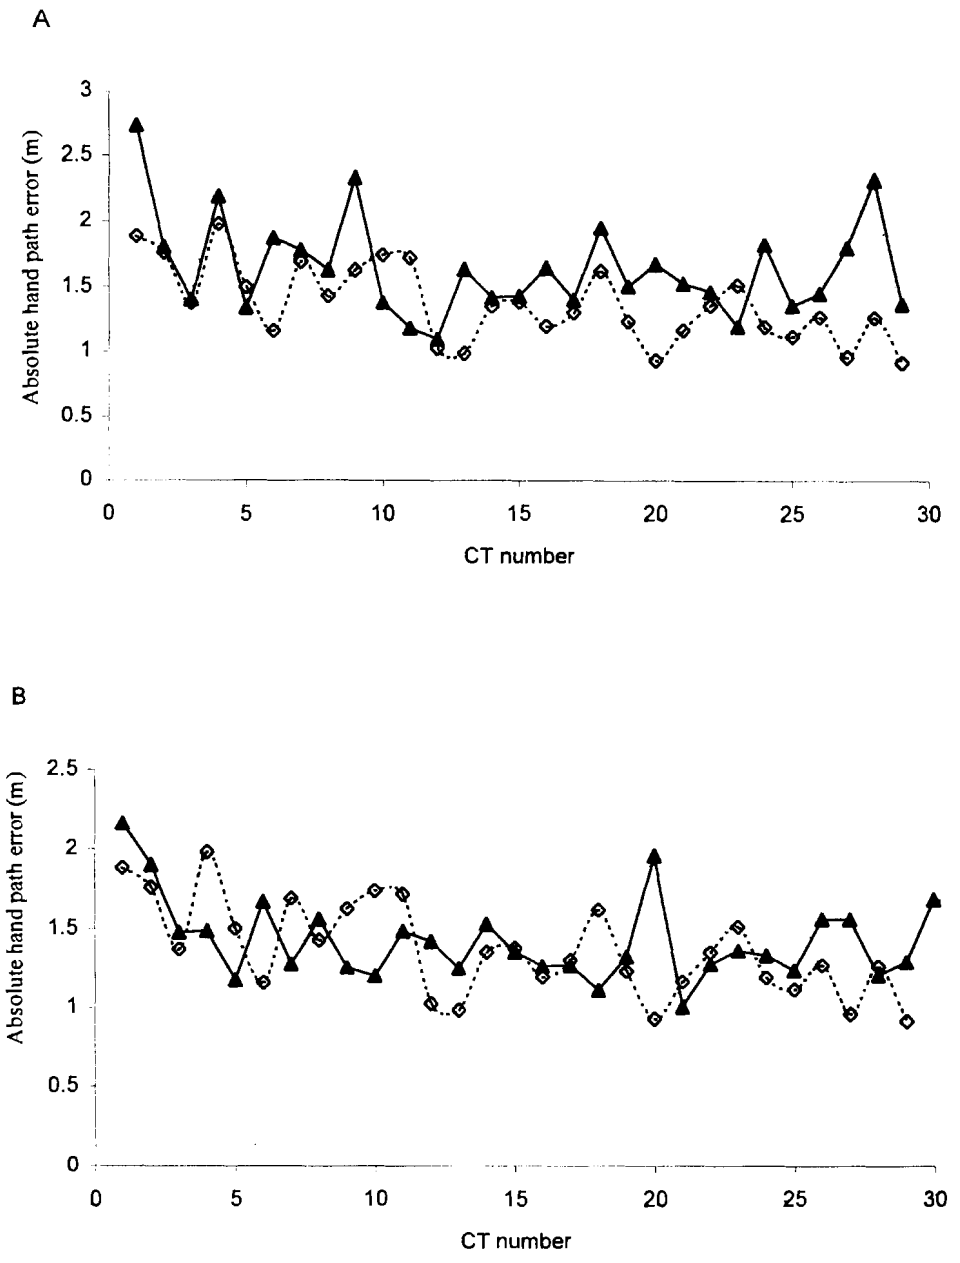

Figure 11. The effect CT's on absolute hand path errors in subsequent PF trials.
A) Absolute hand path error in the pre-CT (dotted line, ◊) and post-CT trials (solid line, ▲). Performance worsened in the post-CT trial relative to the corresponding pre-CT trial (p=0.046) B) Absolute hand path error in the pre-CT (dotted line, ◊) and post-CT+1 trials (solid line, ♦). There was no significant difference in the performance between the 2 sets (p=0.447). Both plots show data averaged over the 9 subjects.

It is clear that imposition of a CT must reduce the rate of learning of the PF. We have shown that no improvement of performance occurs between pre-CT trials and post-CT+1 trials Consequently, learning of the PF field was likely slower than if no CT's had been imposed. Following a CT, which occurred every 5^{th} trial, subjects would only have been able to improve performance during the last 2 trials of the 4 PF trials between each CT, i.e., in only 61 trials out of the 150 trial set would subjects have been able to use error information to incrementally improve PF performance.

In spite of this effect, we decided that learning of the PF could be analysed by considering all PF trials, i.e., 120 trials. We noted that although the "unlearning" effect seen in post-CT trials was statistically significant, it is small in comparison to the overall learning effect within the PF learning period. Furthermore, despite the fact that subjects learned to compensate for the PF, significantly straightening trajectories as learning progressed, there was considerable variability in the PF hand paths, even towards the end of learning. This is apparent in trials not affected by CT's. The performance decrement caused by a CT is not the only factor that may cause worsening of performance on a specific trial, relative to the previous trial. Variability of performance due to motor output variability, for example, can also cause fluctuation in the error between trials.

PF learning- kinematics

Large lateral displacements occurred for all subjects when the PF was initially encountered. This first trial can be considered as the initial stimulus for learning, and separate from the learning per se. Subjects had no knowledge of the force field they were about to encounter and hence could not prepare an anticipatory adaptive response. The maximum rightward displacements on these initial trials were 40-70 mm, and occurred between 0.15 m and 0.20 m from the start position. Subjects generally overcompensated for this displacement in the latter portion of this first trial, demonstrating online corrections to the left before locating the target (Figure 12). On the second trial for most subjects, and by the third trials for all subjects, i.e., the first or second trial in which error feedback had been available to adjust the feedforward command, subjects initially moved

to the *left* of the straight line to the target. Thus, subjects moved into the field, initially slightly overcompensating for the field before being displaced to the right, but by a smaller amount than on the first trial. This leftward movement is indicative of internal model formation, in which a feedforward command resulted in a force that compensated for the PF. The leftward movement may have been a strategy to reduce error later in the movement.

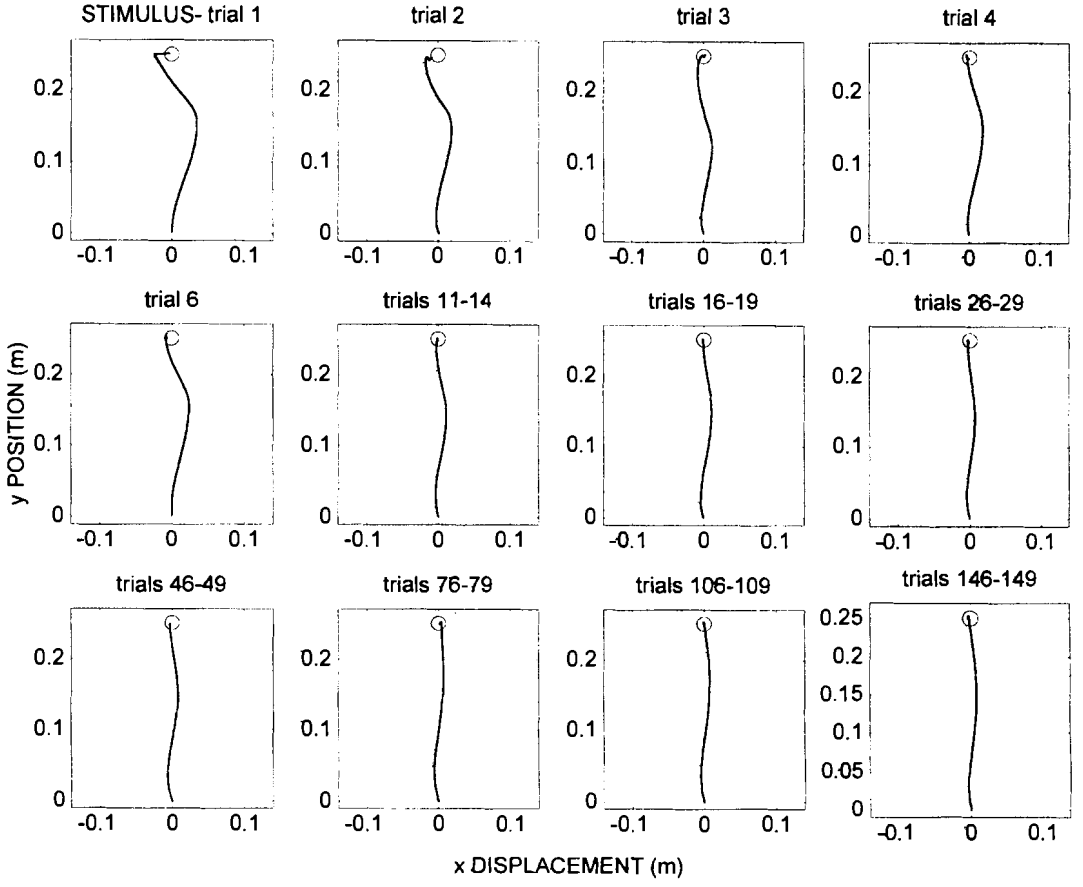

Figure 12. Development of hand paths in the PF field.
Data shown is for the first 5 PF trials and the average of 4 consecutive trials at subsequent points during the 150 trial learning set. Data is shown for subject US. and is representative of all subjects. Note that trial number 5 was a CT, hence trials 1, 2, 3, 4 and 6 represent the first 5 PF trials. The target is represented as the circle, which is drawn to scale within each panel

The initial leftward displacement was characteristic throughout the learning. Hand paths up to the PF boundary, at 0.1 m, changed little once the very large displacement on the

first trial had been reduced (the hand path during this portion of the movement is almost identical between the 2nd trial and trials 146-149 in Figure 12). The most significant path straightening occurred during the latter period of the movement, between 0.1 m and 0.25 m, following removal of the PF. After approximately 50-70 trials, the hand path beyond 0.1 m appeared significantly straighter, while the characteristic leftward displacement in the first 0.1 m remained similar in magnitude to the early trials.

We think that subjects may have been attempting to straighten hand paths primarily in the later portion of the movement. We determined that a suitable way to quantitatively analyze hand path error was to consider two separate periods within a trial. We designated movement until the field boundary as the early period, and from the field boundary until the end of the movement as the late period. In this way we could determine if, as hand paths in figure 12 suggest, the most significant learning occurred in the late period of the trials. We initially considered four different measures of hand path error. Maximum deviation measured the largest lateral distance between the hand path and the straight-line joining the start and target positions. Absolute hand path error was calculated by summing the absolute lateral distance between the hand path and a straight line over time. Mean positive and negative curvature were calculated as in the previous NF-PF comparisons. In all cases, the end of the movement was taken to be 0.23 m, 20 mm short of the target, or in the case of curvature, the point at which the velocity fell below 0.05 ms^{-1} (this was generally around 0.23 m amplitude, but avoided analysis of any trials where velocity dropped below this threshold earlier in the trial). In this manner, we excluded any submovements performed near the target.

There was considerable trial-to-trial variability in both the absolute hand path error and maximum deviation in the early period. Once the error observed on the first trial, i.e., the stimulus trial, had been reduced, which occurred during the first 3-4 trials, little or no decrease in error occurred over the remaining trials. There was no trend in the data reflecting a progressive reduction in the hand path error in the early period of the trials. In the late period of the trial, both the maximum displacement and the absolute hand path

error decreased rapidly in the first few trials and continued to be reduced over approximately the first 45 trials as seen in Figure 13. This result confirms that learning, as defined as reduction in maximum deviation or absolute hand path error, did predominantly occur in the late, as opposed to the early, period of movement trials.

Note that this analysis was carried out on the 120 PF trials in the learning period, i.e., the 30 CT's were not included in the analysis. PF trials have been renumbered 1-120, but one should remember that the 120 trials do in fact represent 120 trials from a 150 trial regimen. Later in the learning, trial-to-trial variability dominated any further error reduction. Both absolute hand path error and maximum deviation measures displayed very similar trends, signifying a high correlation between the two measures, suggesting either could be used to analyze overall performance, in terms of trajectory error.

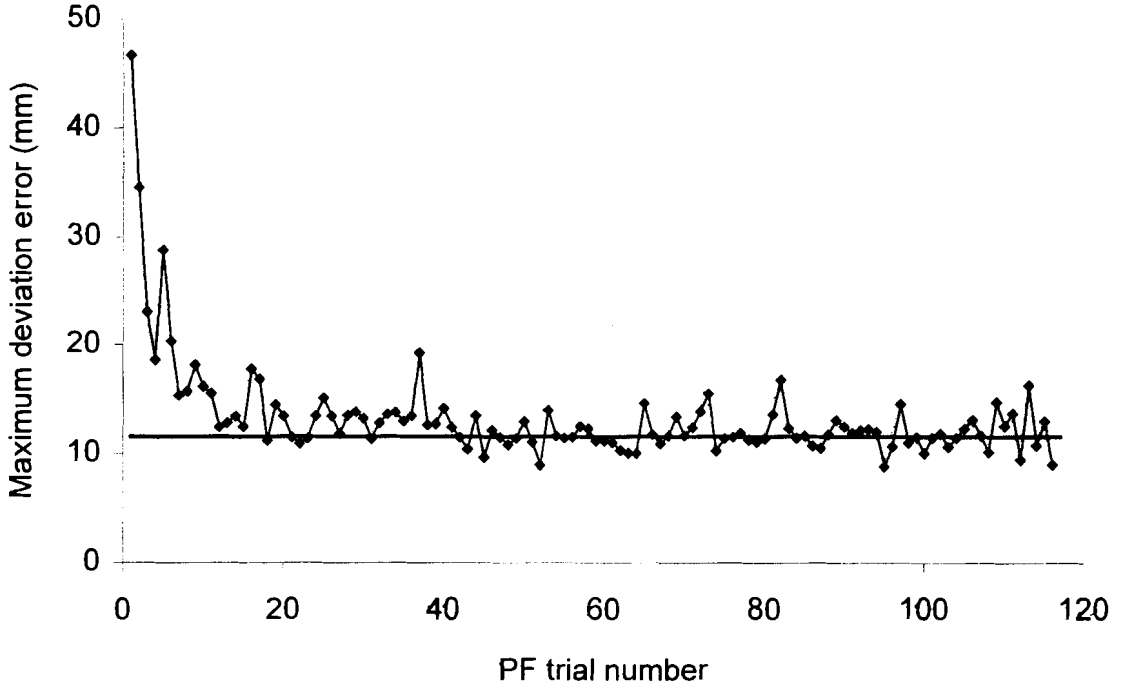

Figure 13. Maximum deviation error in the late stage of the trial.
Data are averaged over the 9 subjects. The horizontal line indicates the mean error over the last four trials.

To take into account the trial-to-trial variability in assessing error reduction, each of the PF learning trials was compared against the last 4 PF trials, for each subject. In this manner, we considered the straightness of movements during learning relative to the

adapted state, without overlooking the obvious variability in hand paths. A performance-related score was formulated for each PF trial, based on the straightness of the movement, relative to the final 4 trials, designated as learnt-PF trials. A score of 2 was assigned if the maximum deviation in a trial was larger than that of all 4 learnt-PF trials. A score of 1 was assigned if the maximum deviation was larger than that of 3 of the 4 learnt-PF trials, a score of 0 was assigned to a trial which was straighter than 2 and less straight than 2 learnt-PF trials. Scores of –1 or –2 were assigned to a trial which was straighter, i.e., had lower maximum deviation than 3, or all 4, learnt-PF trials, respectively. Scores averaged over the 9 subjects are shown in Figure 14.

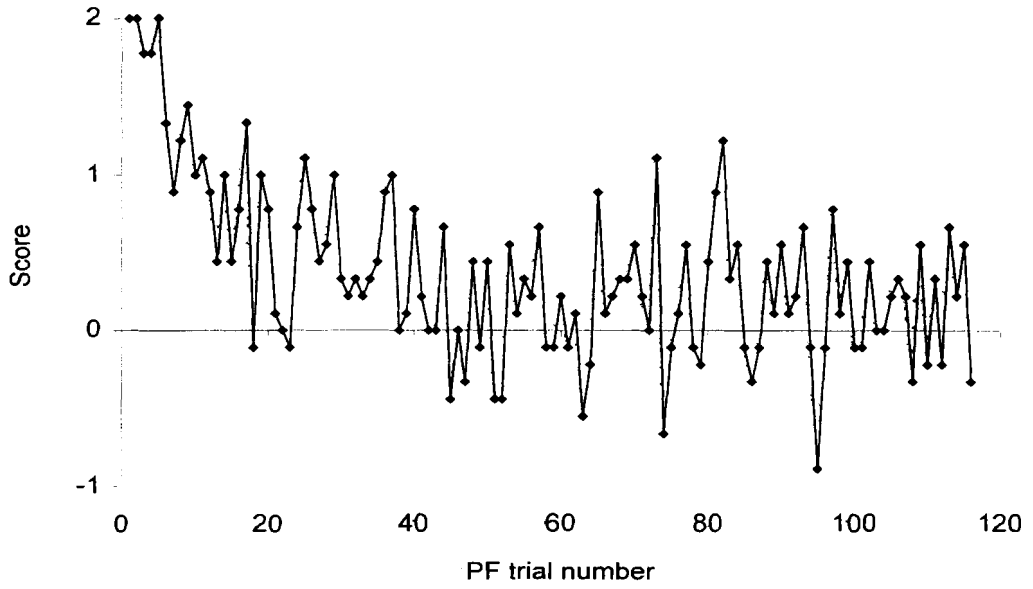

Figure 14. Performance scores using maximum deviation.
Data are averaged over the 9 subjects. The score was calculated by comparing the maximum deviation in each trial to that in each of the last four trials.

A clear improvement in the performance is seen between early trials, where a score of 2 reveals that the trial was less straight than all 4 of the learnt-PF trials, to later trials where the score is generally between 0 and 1, i.e., the trial was only straighter than 1 or 2 of the final 4 trials. Student's t-tests (2 tailed, *a priori* α=0.05) were carried out to determine whether the score for each trial was significantly different than zero. The first trial whose score was not significantly different than zero (non-significant t value) was the 14th trial. Following this trial, 69 of the remaining 102 trials (66%) had performance scores not

significantly different from zero ($p>0.05$), whereas 35 trials (34%) had performance scores significantly larger than zero ($p<0.05$). No trial had a score that was significantly less than zero. This highlights the inherent variability in performance throughout learning, and indicates that at no point in the learning, on average, was any particular trial straighter than 3 of the final 4 trials.

After the 45th trial there does not appear to be any trend for a reduction in the score. A linear regression on the first 45 trials reveals that the intercept (1.505) and gradient (-0.034) are both significantly different from zero ($p<0.0005$ in both cases), confirming general improvement in performance throughout these trials. The correlation coefficient (R-squared) of 0.49 reveals a reasonably good fit, suggesting that the improvement in performance overrides any trial-to-trial variability. A similar regression for the trials following the 45th trial reveals no significant trend in the data (the intercept and gradient are both not significantly different from zero; $p=0.20$, $p=0.69$, respectively), suggesting no further improvement in the performance. Furthermore, the very low correlation coefficient (R-squared=0.0025) confirms the overriding effect of the variability in performance over the later periods of learning. Subjects appear to progressively improve performance in the PF over the initial 45 trials. After this point, factors such as motor output variability result in significant variability in performance, with no further progression in learning, at least in terms of hand path error reduction.

Evolution of hand paths within the PF learning period was also analysed using the mean positive and mean negative curvature. We used mean positive and mean negative curvature to provide a more accurate insight into how learning progressed. Measures such as absolute hand path error and maximum deviation only provide limited information about learning. The curvature measures reveal more information regarding where in the hand path the most significant learning effects occurred. Figure 15 shows curvature values for 10 NF trials, as a reference, and the 120 PF trials. Note, as previously, the 30 CT's within the PF learning period are not included in this analysis.

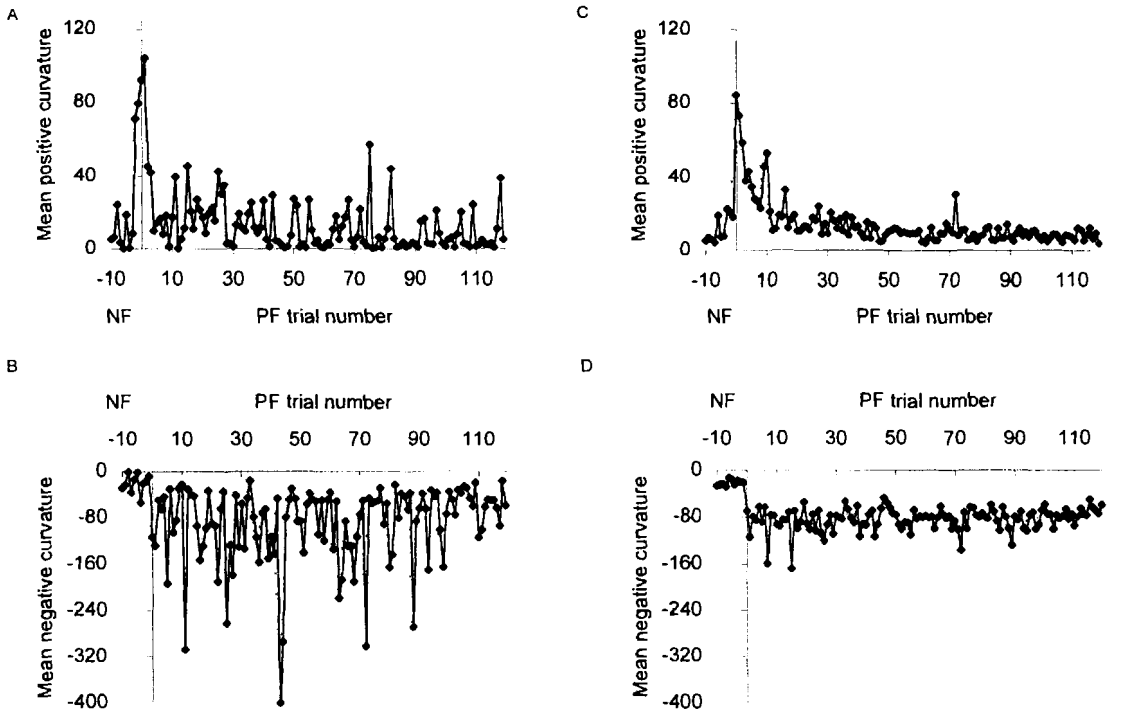

Figure 15. Mean positive and negative curvature.
A) and C) Mean positive curvature for a representative individual subject SM, and averaged over the 9 subjects. B) and D) Mean negative curvature for SM and averaged over the 9 subjects. 10 learnt-NF and 120 PF learning trials are shown. Following initial exposure to the PF, positive and negative curvature are higher than NF baseline values. Positive curvature is quickly reduced towards NF values during learning, whereas negative curvature remains elevated compared to the NF for the duration of learning.

The trends that we noted in the maximum deviation are equally evident in the curvature. On the first PF trial, subjects are perturbed to the right, resulting in high negative curvature followed by high positive curvature as they return towards the target later in the movement. As learning progressed we noted significant straightening of paths, i.e., a reduction in maximum deviation and absolute hand path error, during the latter portion of the movement. This straightening is also evident in the positive curvature, which returned towards NF levels over the first 45 trials of the learning period. The reduction of maximum deviation over the late period of the trial and the reduction in mean positive curvature occur with a very similar time course. This suggests a high correlation between the 2 measures. We noted previously that subjects appear to move into the field during

the early period of the trial, resulting in little reduction of maximum deviation or absolute hand path error in this period. Negative curvature was associated with the movement into the field and subsequent return towards the centreline at the field boundary (0.1 m) and changed little throughout learning (see Figure 15B and D), whereas positive curvature, more prevalent later during a trial, was reduced significantly during learning (Figure 15A and C).

Performance scores were also calculated for each trial in terms of mean positive and mean negative curvature, to account for the variability, which was also evident in these measures. We considered whether the mean positive or mean negative curvature in each learning trial was higher or lower than the respective curvature in the final 4 trials. This gave two scores, one for mean positive and one for mean negative curvature (Figure 16).

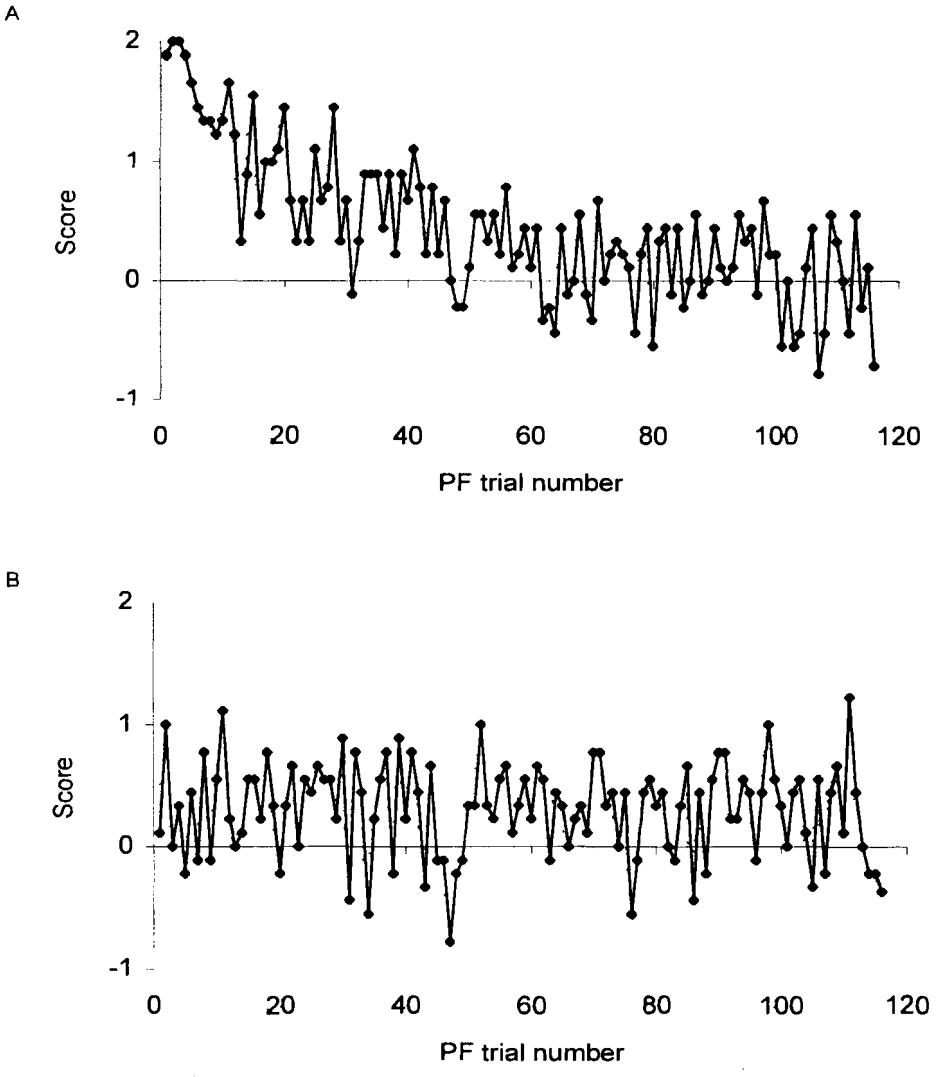

Figure 16. Performance scores for curvature.
Data are averaged over the 9 subjects, during PF learning. A) Mean performance score for positive curvature. B) Mean performance score for negative curvature. Scores were calculated by comparing the positive or negative curvature for each trial to the respective curvature of each of the last four trials.

Mean performance scores for positive curvature (Figure 16A) clearly decreased as learning progressed, whereas no trend was seen in mean negative curvature (Figure 16B). Initial trials had significantly higher positive curvature than the final 4 PF learning trials, reflected as scores of 1-2 over the first 15 trials. Mean positive curvature decreased over a similar number of trials as performance scores for maximum deviation (Figure 14). Both Figure 14 and Figure 16A indicate that significant reduction in the score occurs within

the first 45 learning trials, with little or no further improvement occurring in the remaining 75 trials. This high degree of similarity suggests that positive curvature occurs predominantly in the final 0.15 m of the movements. This was confirmed by considering the positive and negative curvature in the two portions of the movement: from movement onset until the PF boundary and from the PF boundary until velocity dropped below 0.05 ms^{-1}, at a distance of around 0.23 m. We found that in the early portion of the movement, 731 of the 1080 trials (120 PF trials for each of the 9 subjects) exhibited no positive curvature whatsoever. No trend was seen between trials where positive curvature was present as learning progressed. The mean negative curvature in the early portion of movements exhibited no noticeable trend, replicating the result found when we considered the mean negative curvature over the whole trial, as seen in Figure 15. The trend in mean positive curvature in the late portion of the movement was very similar to that seen in Figure 15. Negative curvature was present in the late period of movements, although it was less prevalent than in the early portion of the movement, with no apparent trend for a decrease as learning progressed.

Student's t-tests (2 tailed, *a priori* α level 0.05) revealed the first trial whose performance score for positive curvature was not significantly different from zero was the 13^{th} trial, i.e., all trials prior to this point had a higher level of positive curvature than the final 4 PF trials. This is very similar to the result for the performance score for maximum deviation that revealed that the 16^{th} trial was the first trial whose score was not significantly different from zero. Following this trial, 93 of the remaining 103 trials (90%) had performance scores not significantly different from zero, whereas 10 trials (10%) had performance scores significantly larger than zero. No trial had a score that was significantly less than zero. This shows that once the mean positive curvature was reduced, it remained, for the most part, at a level similar to that seen in the final 4 trials, highlighting a lower variability in mean positive curvature compared to the maximum deviation measure. Statistical analysis showed that the mean performance score for negative curvature was very consistent within the learning period. Only three widely separated trials had a performance score that was significantly higher than 0, i.e., only

these trials had more negative curvature than the final 4 trials. Early trials did not have greater mean negative curvature than late trials, confirming, as Figure 16B suggests, that no reduction in mean negative curvature occurred during learning. This provides statistical confirmation of the observation that there was very little change in the hand path over the first 0.1 m of each trial during learning. Movements were negatively curved in this region and remained so even after adaptation.

In summary, we have found that positive curvature is predominantly associated with the late portion of the movement, beyond the PF boundary. It is in this late portion of the movement where we saw dramatic adaptation to the PF with learning, both as a reduction in trajectory error and mean positive curvature. Negative curvature showed no trend throughout the learning period remaining similar in magnitude in the early and late learning trials It appears that the negative curvature, predominantly associated with the initial lateral movement that subjects make into the field at the onset of the movement, is not perceived as an important parameter that subjects strive to reduce during learning.

Modification of subjects' CT force profiles during PF learning

Learning in the PF, as characterized by hand path, cannot be taken as a complete characterization of learning. Even though hand paths did not become significantly straighter in later trials does not rule out the possibility that subjects were continuing to modify the way in which they compensated for the PF. Producing the same hand paths with a reduction in the level of cocontraction is the basis of feedback error learning. During early exposure to the PF subjects may employ high levels of cocontraction such that a relatively straight hand path can be produced, even though the motor commands (representing the internal model) issued to compensate for the external field may be inaccurate. Learning may consist of reduction in the level of cocontraction to an optimum level, such that similar hand paths are achieved in a more energy-efficient manner, lowering the overall metabolic cost and limiting the effects of muscular fatigue. Such a learning effect would be visible as changes in the force recorded in CT's, as well as in modification in the feedforward motor command, represented by the EMG. Here we

hypothesized that as subjects continued to be exposed to the PF a more accurate feedforward command to the muscles would be learnt. The force recorded against the channel should more closely resemble the force that would result in a perfectly straight path, in the absence of any cocontraction.

We used the force impulse against the channel wall from 150 ms prior to movement onset until movement reached the PF boundary at 0.1 m as a measure of the feedforward command. This time varied from trial to trial due to small fluctuations in the movement velocity, but averaged 240-270 ms for the 9 individual subjects. Any voluntary response to detection of the channel would probably not have occurred until near the end of this period. As subjects adapted to the PF, the effect of the CT should have become less noticeable and the likelihood of voluntary reactions should have been lower. We saw no distinct drops in the force recorded in CT's over (at least) the first 0.1m, as seen in Figure 17. There was an initial rapid rise and subsequent fall in force, which was consistent with the leftward deviation and negative curvature during the early portion of the movement. It is unlikely that this was a voluntarily reaction to the CT as it occurred at latencies significantly less than 100 ms following movement onset. We assumed that subjects did not react to any perception of the imposition of the CT, at least within the time period relating to the first 0.1 m of movement, and used the force recorded in the CT's as a representation of the feedforward command to move in the PF.

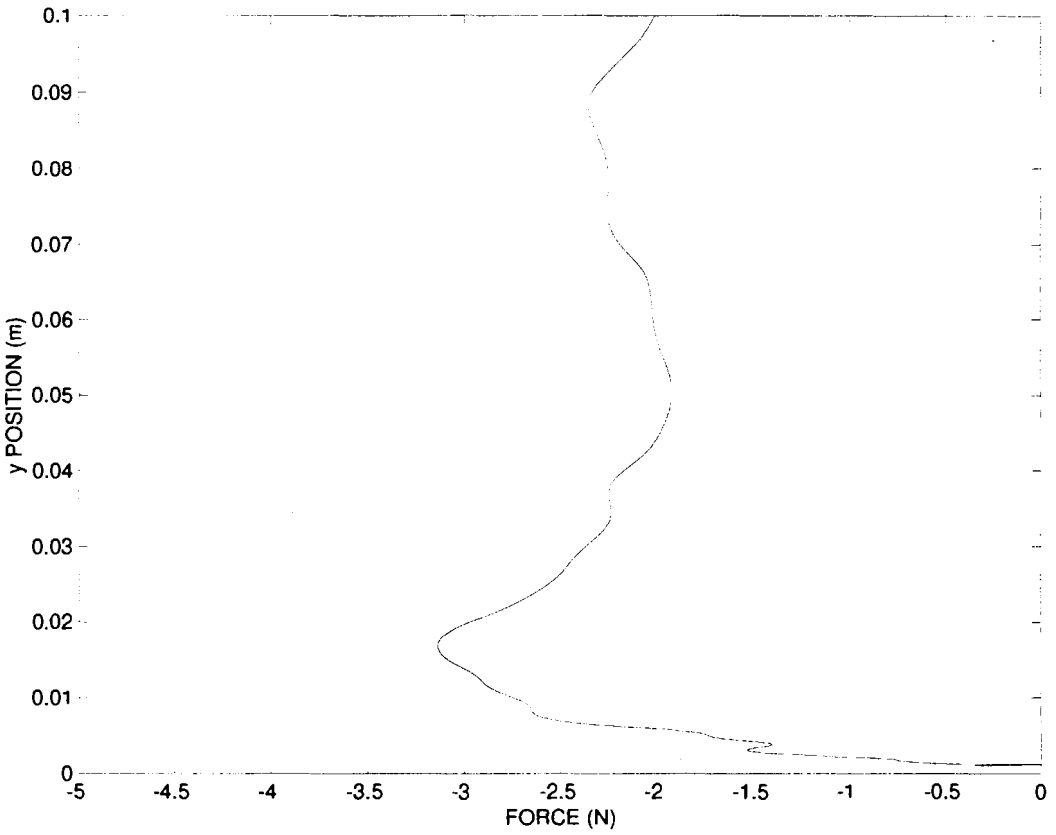

Figure 17. Force profile in the first CT, averaged over all subjects.

A subject's force impulse, I_S, given by the equation

$$I_S = \int_0^T F_{Sx}(t)dt ,$$

where F_{Sx} is the subject's lateral (x) force measured at time t in the CT's and T was the time at which subjects reached 0.1 m amplitude, was used as a measure of the total force that the subject produced over the period in which the PF was active.

From initial inspection of the force impulse in the 30 CT's it was evident that six subjects (MH, US, AN, JJ, DF and RO) had altered their force impulse in CT's as learning progressed, whereas no trend in force impulse was visible for the other three subjects (JH, SM and AR). For this reason, we analysed the data for these two groups separately as

they appeared to have utilized different strategies in compensating for the PF (Figure 18). We used one-way RM ANOVAs to determine if the change in force impulse was significant for both groups of subjects. Four trials, comprising the 1st, 10th, 20th and 30th (final) CT's (CT1, CT10, CT20 and CT30) were included in the analysis to represent various stages of the learning. Using all 30 CT's would have drastically reduced the power of the test, due to lower sphericity (low ε) resulting from the number of degrees of freedom (29 trials). All RM ANOVA's rely on the assumption of sphericity. (ε is a measure of the extent to which the covariance matrix meets the assumption of circularity, i.e., the requirement that all variances between all pair-wise comparisons are equal (Howell, 1997). Increasing the number of degrees of freedom increases the number of pairwise comparison resulting in very low ε and thus very low power.) Note that these 4 CT's follow the 4th, 40th, 80th and 120th exposure to the PF during training.

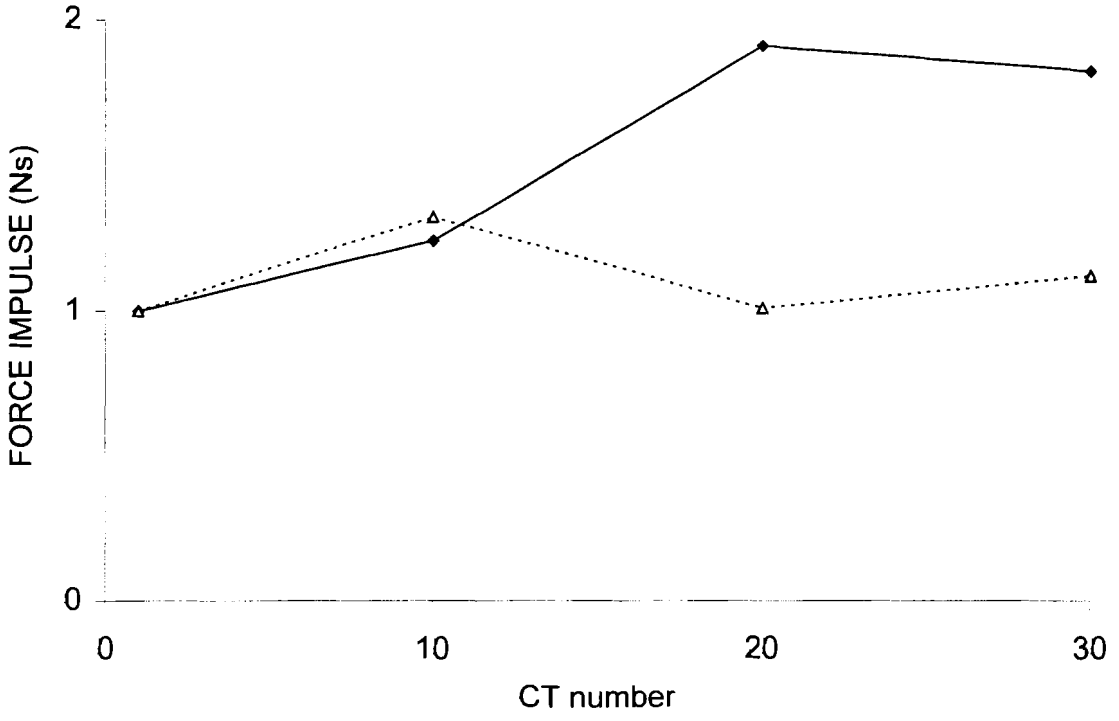

Figure 18. Force impulse in CT's that replaced PF trials during PF learning.
Force impulses are calculated from movement onset until 0.1 m (PF boundary). Data is normalized to the first CT for each subject and is shown for CT's 1, 10, 20 and 30. The dotted line (and Δ) represents the average over the three subjects who appeared not to adapt force impulse during learning. The solid line (and ♦) represents the average over the six subjects who appeared to show an increase force impulse during learning.

We found a significant increase in force impulse for the six subjects who appeared to increase force impulse within PF learning ($F=5.837$ $p=0.01$, Huynh-Feldt sphericity correction) but no change in the force impulse for the three subjects who appeared not to alter the force impulse ($F=5.060$ $p=0.152$, Greenhouse-Geisser sphericity correction). Repeated contrasts revealed that in the six-subject group, force impulses in CT1 and CT10 were both significantly less than the force impulse in CT30 ($p=0.005$, $p=0.034$), but force impulse in CT20 and CT30 were not significantly different ($p=0.734$). In the three-subject group, force impulses in CT1, CT10 and CT20 were not significantly different from the force impulse in CT30 ($p=0.097$, $p=0.079$, $p=0.543$ respectively). This suggests that over the learning period the six-subject group increased force impulse to compensate for the PF, at least until CT20, developing an internal model to compensate for the applied external field. After CT20 it appears that there was no further adaptation in the force produced by the 6 subjects. In contrast, the three-subject group appeared not to change their force impulse at any stage during of the learning that we analyzed.

To determine whether the timing of the peak force against the channel changed significantly as subjects adapted to the PF, force impulses against the channel were calculated in 10 intervals of equal duration, from movement onset until the PF boundary was reached. This would reveal whether subjects altered the timing of their force generation to compensate for the PF, even if the force impulse did not change. Each interval was approximately 20-25 ms in duration. We ranked each interval from 1-10 according to the relative size of the force impulse in that interval compared to the others. Similar to the previous analysis, we conducted a one-way RM ANOVA for each of the two groups of subjects, considering only CT1, 10, 20 and 30. No change in the interval having the highest force impulse was found as learning progressed for either group ($p=0.492$, 0.711, Greenhouse-Geisser corrections, for the 6-subject group and 3-subject group, respectively). Averaged over all nine subjects and the four CT's, the largest force impulse occurred in the 8^{th} interval, i.e., around 160 ms after movement onset (Figure 19). This result shows that subjects did not significantly modify the shape of the force profile during learning.

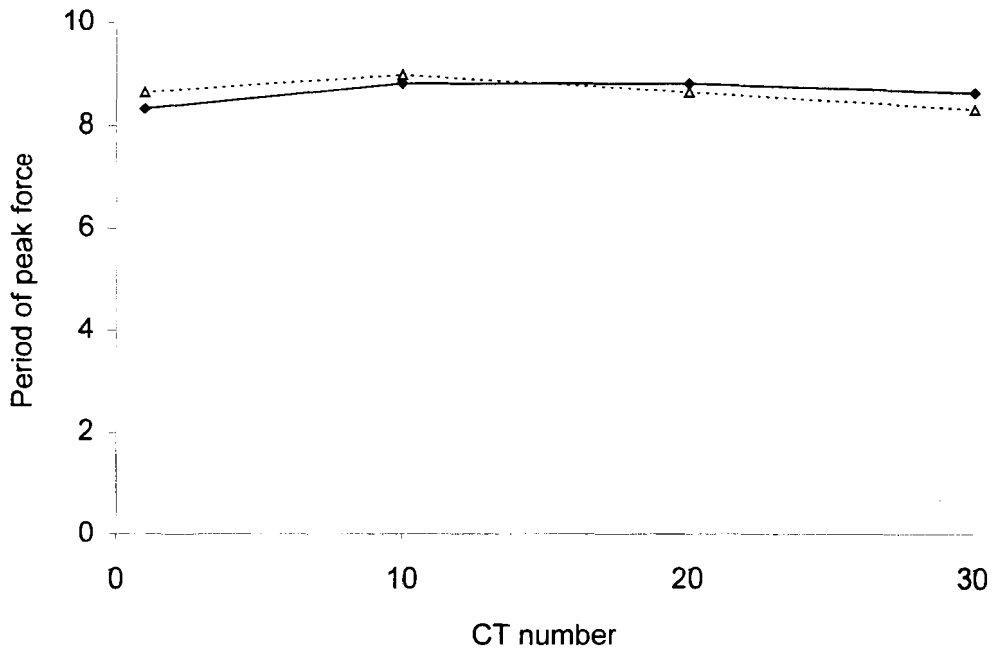

Figure 19. Interval in which the peak force was recorded in CT's 1, 10, 20 and 30.
The dotted line (and Δ) represents the average over the 3-subject group and the solid line (and ♦) represents the average over the 6-subject group. ANOVA revealed no change in the position of the largest force impulse (timing) for either group.

Comparison of CT force profiles to the PF force profile

The previous analysis considered how the total force that subjects produced, in compensation for the PF, changed as they adapted to the PF. We will now consider the actual shape of the force profiles that subjects produce, and compare them to the PF force profile. This will reveal how closely subjects could match the temporal PF force profile i.e., how similar was the force profile that subjects produced, following complete adaptation to the PF, to a profile that would be a perfect compensation to the PF, in the absence of any cocontraction? Previously we showed that there was no significant difference in the force impulse between CT20 and CT30, i.e., trials 100 and 150 in the PF learning, for either group of subjects. For this reason the average force profile in the last 5 CT's (CT's 26-30, i.e., trial numbers 130, 135, 140, 145 and 150) was taken to represent the fully adapted force profile, thereby enabling measurement of the lateral force produced by voluntary muscle activation without having to take into account the

dynamics of the arm or the PFM in the lateral direction. The force over the first 0.1 m represents the force produced as a result of the feedforward motor command, prior to any online corrections. Subjects' force, averaged over the last 5 CT's is compared to the PF force in Figure 20.

Subjects' temporal force profiles in the CT's did not accurately match the PF force. However, some aspects of the CT force profiles were similar for all subjects. At movement onset, subjects produced a force to the left to counteract the force of the PF, which, in the absence of any lateral force produced by subjects or without any cocontraction, would have significantly perturbed hand paths to the right. Subjects increased their force much more quickly than the PF force increased, such that their force was larger than the PF force for the first 20-35 mm of the movement. This overcompensation for the PF is visible in the learnt hand paths, which show initial displacements to the left for all subjects. Following the phase of rapidly increasing force, some subjects reduced their force slightly and then maintained a relatively constant force whereas others continued to increase their force, but more slowly (Figure 20). The peak subject force was generally 4-5 N, 1-2 N below the peak PF force. The point of intersection of the subject force and PF force dictates the point where there is no net lateral force, resulting in zero lateral acceleration. Once subject force dropped below the PF force, the hand began to decelerate in the lateral direction. Shortly thereafter, the hand path began to return towards the centreline. Subject force dropped more slowly than PF force, such that it was once again larger than the PF force around the 0.075 m point. When the PF boundary was reached at 0.1 m subjects were still exerting a force of 2-3 N. Only one subject, JH, produced a force profile that did not exhibit the characteristic fast rise. JH also produced a peak force that was larger than the peak PF force.

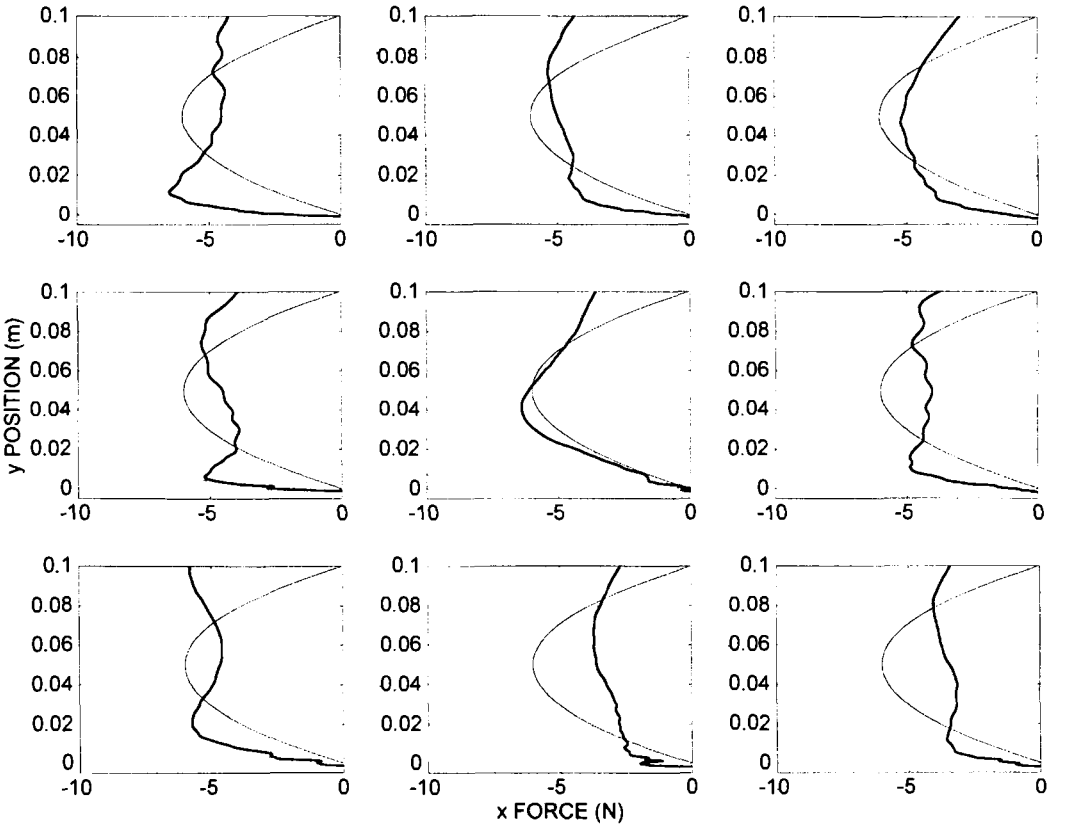

Figure 20. Force profiles for the 9 subjects over the portion of the movement in which the PF was applied. Thin lines represent the PF force, and thick lines represent subjects' actual forces, averaged over the final 5 CT's in the PF learning period. Top left-bottom right subjects MH, US, AN, JJ, JH, SM, DF, RO and AR.

The minimal deviation from a straight line following PF learning, despite the fact that the force recorded in CT's does not exactly match the PF force, suggests that cocontraction or impedance control may play a role in controlling trajectory error.

To investigate the difference between the subject's force and the PF force, over the entire period the PF was active, we computed the PF force impulse. Force impulse for the PF between $t=0$ and $t=T$ is given by the equation

$$I_{PF} = 2400\int_0^T (y(t)^2 - 0.1y(t))dt,$$

where I_{PF} is force impulse, 2400 is the field gain and $y(t)$ is the y position as a function of time.

Partial force impulses were calculated for the PF and for the subject's force between t=0 and all succeeding times until the end of the PF was reached at 0.1 m. If the subject's partial force impulse were less than or greater than the partial PF force impulse, over any interval, this would imply either under-compensation or overcompensation, respectively, for the PF over that time period.

We found that 7 out of the 9 subjects overcompensated for the PF (Figure 21). For these subjects their partial force impulse was larger than the PF partial force impulse over the entire region of the PF. On average subjects' total force impulses (at the PF boundary) were 30% larger than the PF total force impulse, with a range from 15-52%. One subject (RO) under-compensated for the PF by 14%, while one subject (AR) was able to accurately compensate, achieving a total force impulse within 1% of the PF total force impulse (Figure 22). Although AR was able to accurately match the PF total force impulse, this subject's partial force impulse was not accurately matched to the PF partial force impulse. Referring back to Figure 20, one can see that AR initially overcompensated for the PF, producing excess force. AR's force then dropped below the PF force such that over the region of the PF, the total force impulses were very similar.

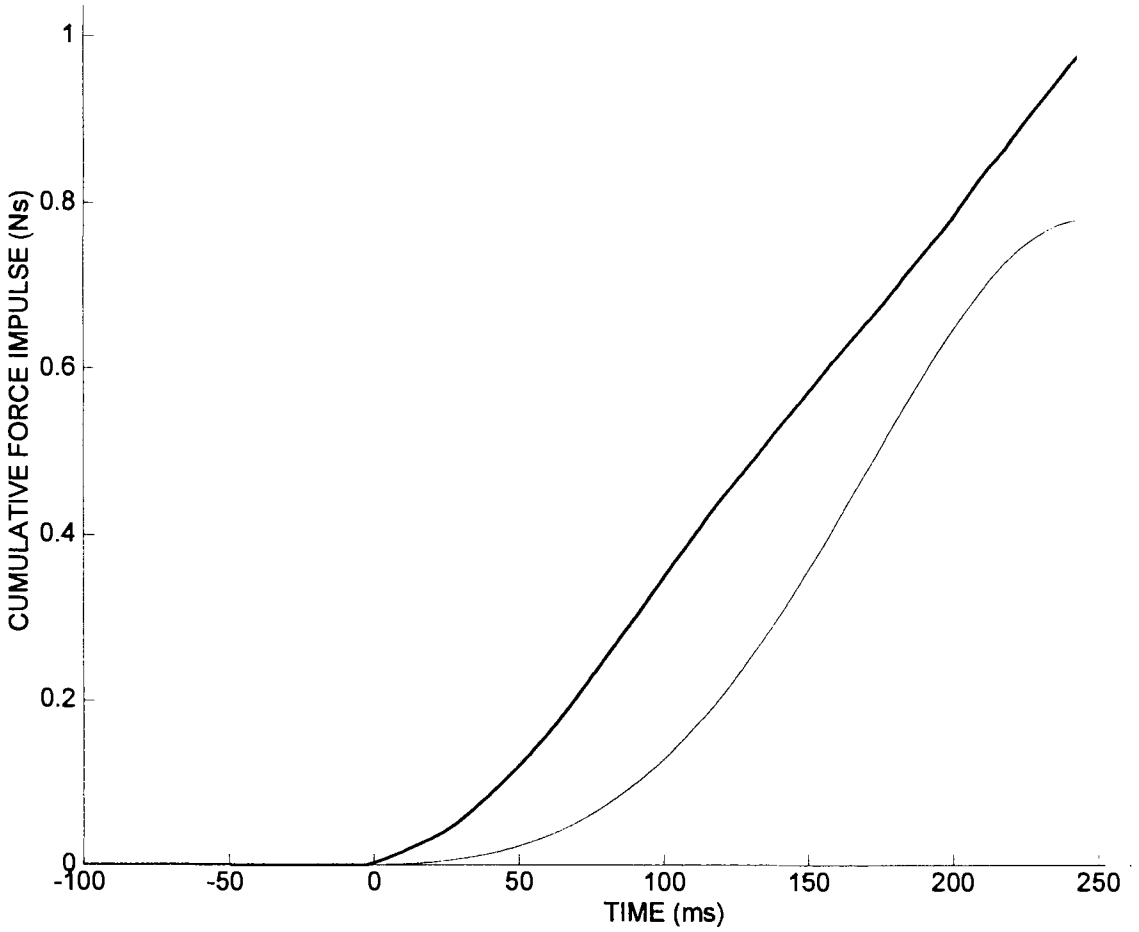

Figure 21. Partial force impulse from movement onset until the PF boundary, following PF learning.
PF force impulse is shown as a thin line and subject's force impulse as a thick line. Data is averaged over the last 5 CT's during learning, for subject SM. Data is representative of 7 out of 9 subjects, who overcompensated for the PF force during the first 0.1 m of movement.

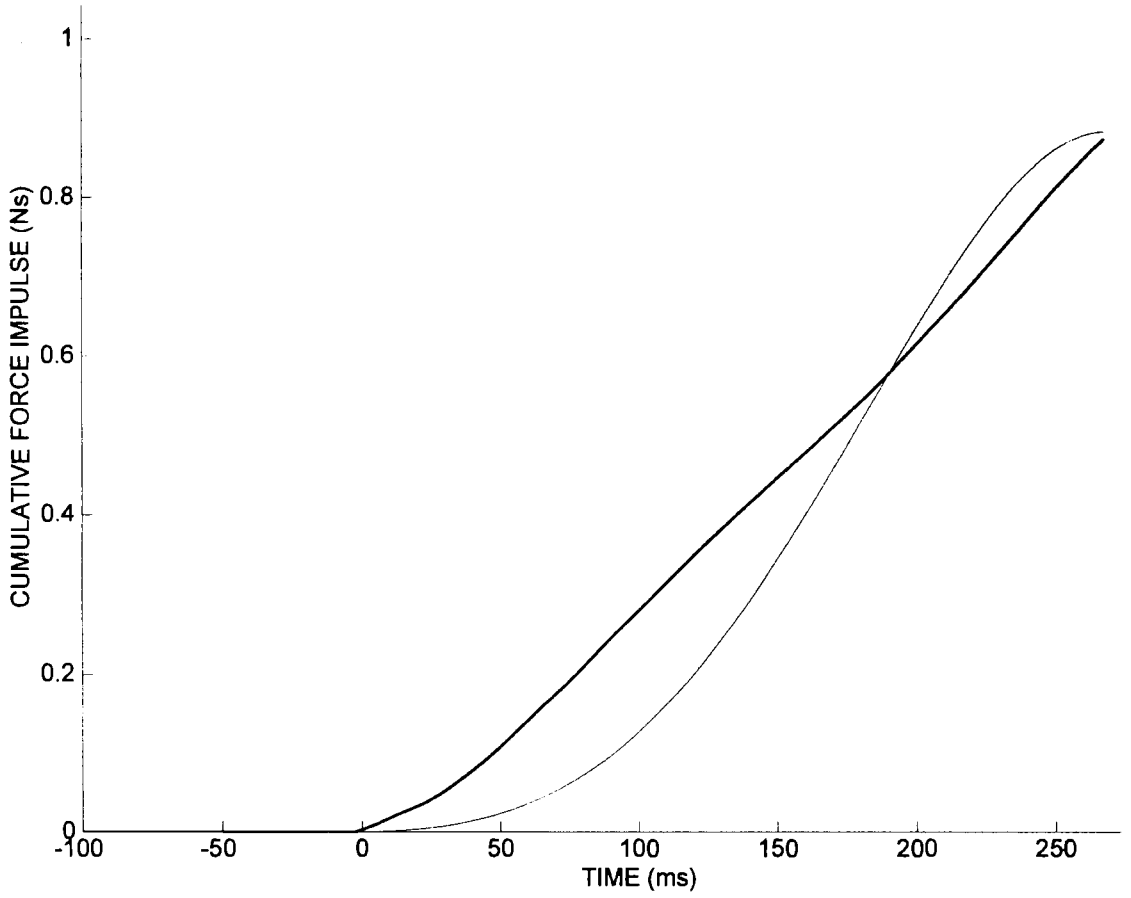

Figure 22. Partial force impulse from movement onset until the PF boundary, following PF learning.
PF force impulse is shown as a thin line and subject's force impulse as a thick line. Data is averaged over the last 5 CT's during learning for subject AR, who accurately compensated for the total PF force impulse during the first 0.1 m of movement.

Joint torque predictions in PF learning trials

Subjects' adaptation to the PF can also be investigated in terms of changes in the shoulder and elbow joint torques in compensating for the PF. Analysis of joint torques may provide clues as to why subjects were not able to completely adapt to the PF so that even following learning they produced bowed hand paths. Using anthropometric data for subjects' upper and lower arm segment lengths we converted the end point position of the hand (**r**), in Cartesian co-ordinates, into elbow and shoulder joint angular positions (**q**) using inverse kinematics. We then determined time-varying elbow and shoulder joint torques using inverse dynamics from end point force, as measured on the PFM handle.

We calculated the total torque, i.e., the torque that subjects produced at the elbow and shoulder using a two-segment planer arm model (Hollerbach and Flash, 1982) to move the inertial mass of the arm-PFM system but with additional components to account for the interaction torques between the subject and the externally applied force field. (Franklin et al, 2003). Here the Jacobian matrix, J, was used to transform endpoint forces, F, to joint torques, τ.

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

$$\tau_e = \ddot{\theta}_e Y + \ddot{\theta}_s(X\cos\theta_e + Y) + \dot{\theta}_s^2 X\sin\theta_e - l_2\sin(\theta_s+\theta_e)F_x + l_2\cos(\theta_s+\theta_e)F_y$$

where

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

I is the moment of inertia about the centre of mass of the segment (c_m), which is represented as a distance from the centre of rotation of the limb segment. l and m represent the segment lengths and masses. Subscript 1 and 2 refer to the upper arm and forearm, respectively, while subscripts s and e relate to the shoulder and elbow and c to the thermoplastic cuff that we used to prevent motion of the wrist. Inertia and mass of limb segments were estimated from the segment lengths and scaled anthropometrical data (Winter 1990).

Using kinematic data from a previous study in which subjects learnt to move in a null field, i.e., subjects moved without an additional applied force, we predicted the joint

torques that would be required by each subject to move on the same, approximately straight, NF path in the PF. This analysis was completed for the 7 subjects for whom NF data was available. The validity of this technique was verified by checking that there were no significant differences between the movement kinematics in the NF and PF, i.e., we confirmed that the movement durations and the peak values of the reach velocity were similar in the two fields. The predicted joint torques were compared to subjects' joint torques, as calculated from actual trials. We aimed to determine how accurately subjects could match the required joint torque profiles, and suggest possible reasons for any disparities. Quantitatively, predicted joint torque in the PF varied slightly across subjects, due to anthropometric variations and slight variations in the paths subjects learnt in the NF, but the torque profile shapes were very similar across all subjects.

Differences in the joint torques (Figure 23) reflect the addition of the PF between 0 and 0.1 m along the y-axis. To compensate for the PF and produce a trajectory similar to that in the NF, i.e., move at the same speed and along an identical hand path to that in the NF, substantially higher shoulder flexor torque was required, compared to that needed in the NF. The peak predicted value in the PF, of approximately 3 Nm, is 3-4 times greater (depending on the subject) than the peak shoulder flexor torque recorded in the NF, and occurred slightly later in the movement. Extensor torque was required at the elbow for both the NF and PF. In the PF, the peak predicted elbow extensor torque was actually slightly *lower* than in the NF. This being said, the main characteristic of the PF predicted elbow torque is that the rate of change of torque following peak extension torque, back towards zero torque, was faster than in the NF.

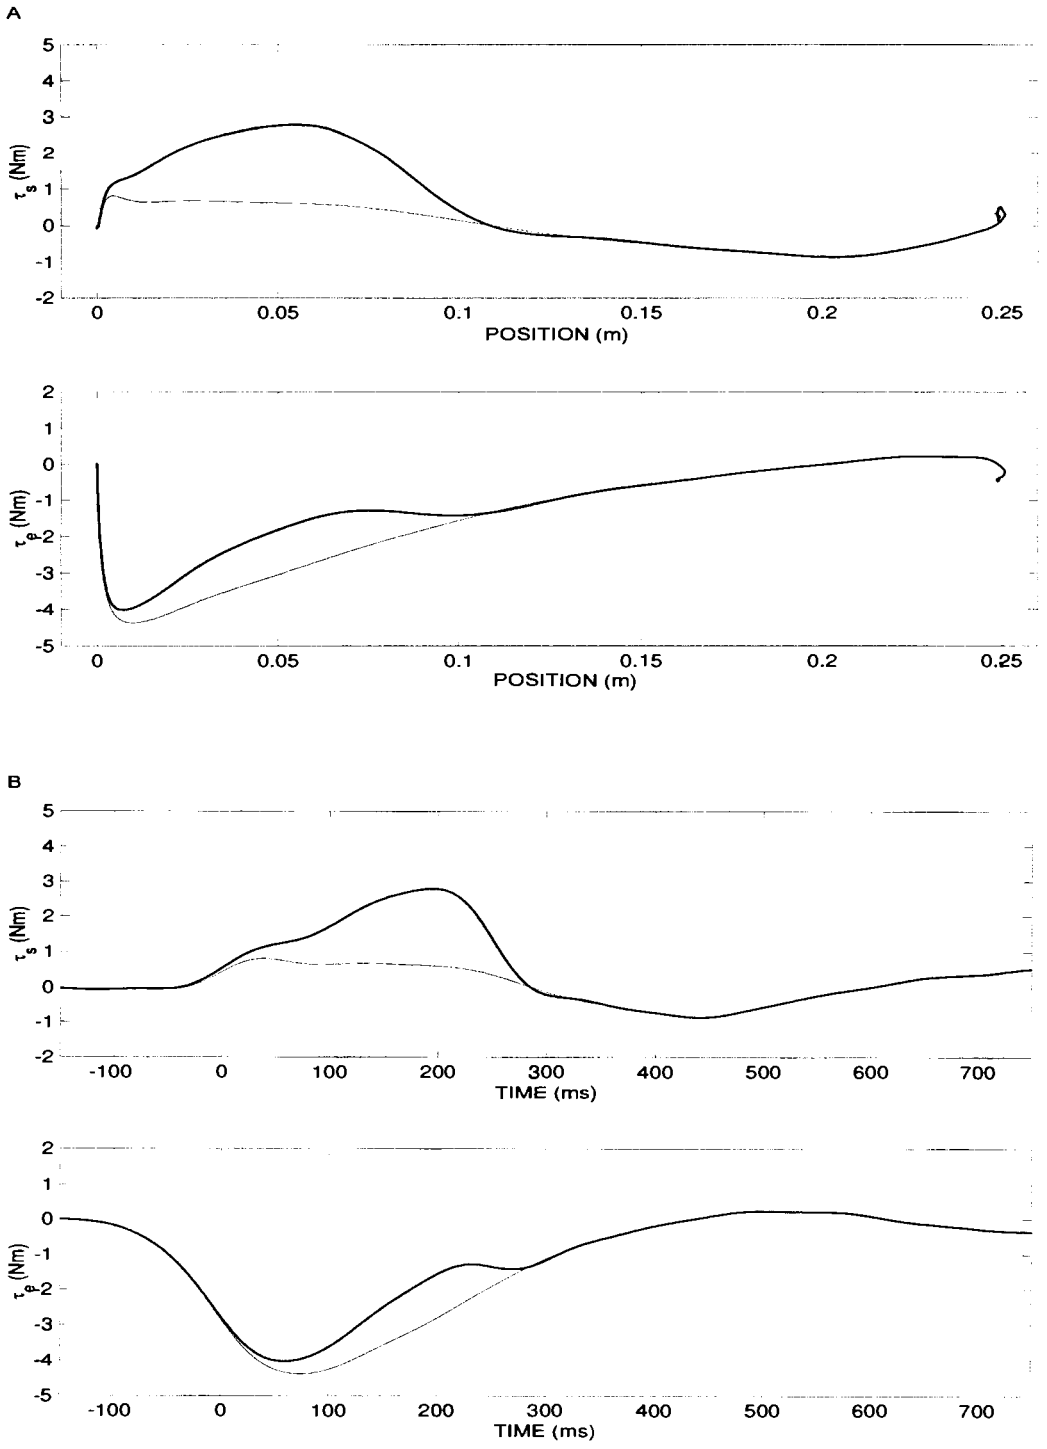

Figure 23. Predicted shoulder and elbow joint torque as a function of A) position and B) time.
Flexion torque is positive, and extension is negative. The thick traces show the torques required to move in the PF, on a path identical to that in the NF. The thin traces show the joint torques calculated for NF movements. Data is shown for subject US, and is representative of all subjects.

To analyze learning we considered how the joint torques produced by subjects changed as they learnt to move in the PF, with particular interest in how closely subjects could match the predicted torque profiles. We compared averages of groups of trials at several stages of learning with the first 4 trials. We used the averages of 4 consecutive trials to reduce the variability. All trials were compared to the values of predicted torque as calculated from the NF data (Figures 24 and 25).

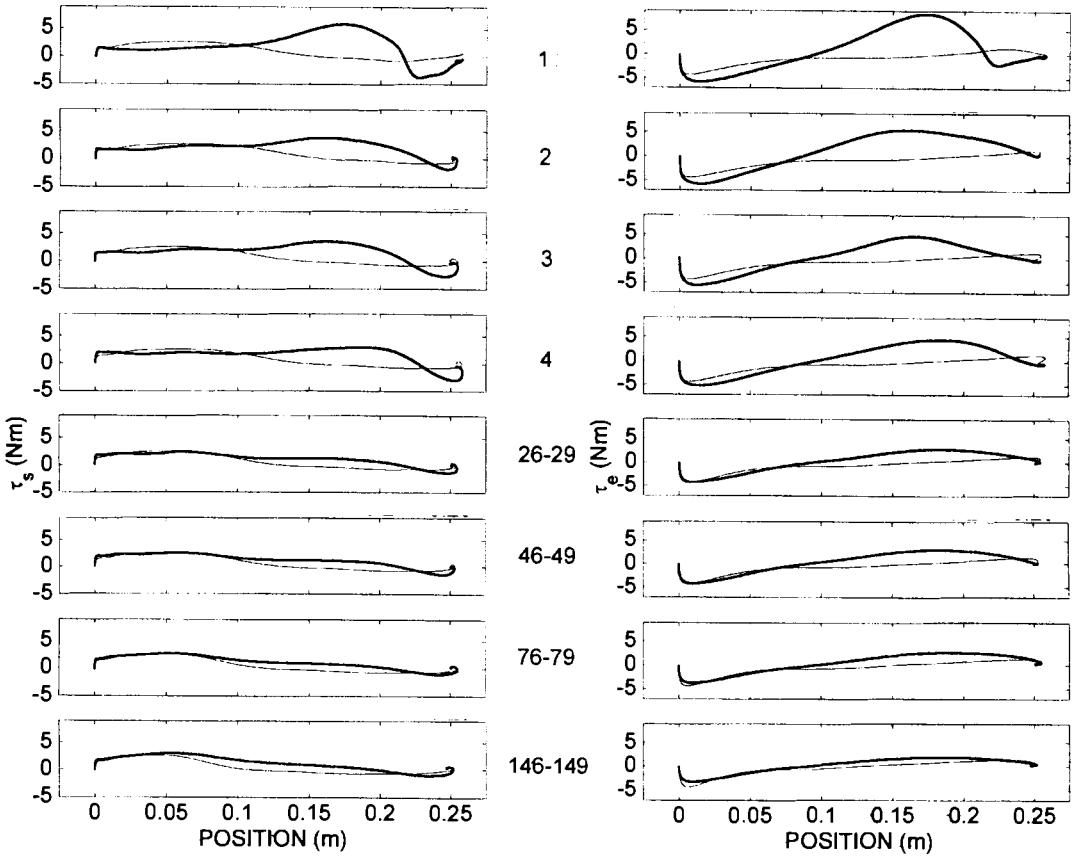

Figure 24. Shoulder (τ_s, left) and elbow (τ_e, right) torque as a function of position during PF learning. Flexion torque is positive, and extension is negative. The thin trace shows the joint torque required to move, in the PF, on a path similar to the NF. The thick traces show the torque produced by subject MH in PF learning trials. Data is representative of all subjects. Trial numbers are shown between the traces. As learning progresses the subject matches the required joint torques more accurately.

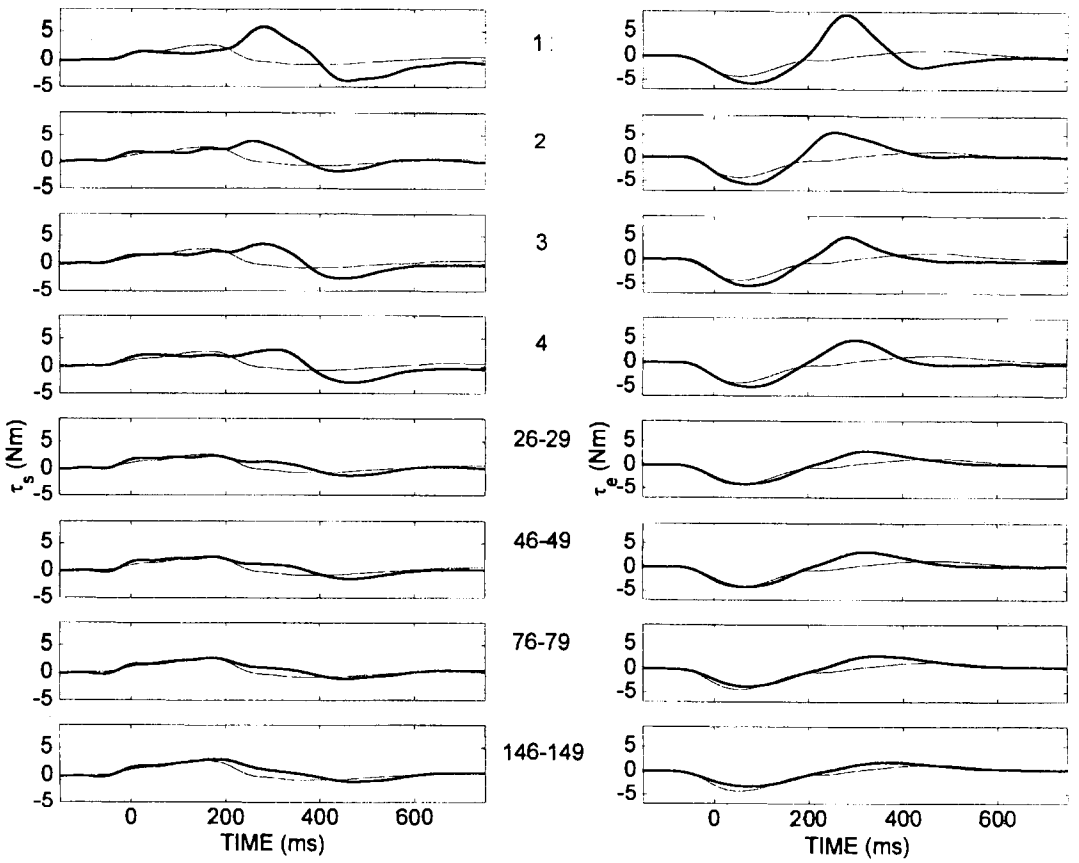

Figure 25. Shoulder (τ_s, left) and elbow (τ_e, right) torque as a function of time during PF learning. Flexion torque is positive, and extension is negative. The thin traces show the joint torque required to move on a path, in the PF, similar to the NF. The thick traces show the torque produced by subject MH in PF learning trials. Data is representative of all subjects. Trial numbers are shown between the traces.

Figure 23 implies that a lack of shoulder flexor torque, based on the premise that subjects would issue motor commands similar to those used in the NF, would be responsible for the initial rightward displacement observed in hand paths on the first PF trial. This disparity is seen in the upper left panel of Figures 24 and 25.

In the position range 0.1-0.15 m (~240-450 ms) shoulder and elbow flexor torques are much higher than the predicted values in the early trials (Figures 24 and 25). It is unlikely that much of this torque increase is a corrective shoulder and elbow torque due to a stretch reflex response because, as we previously reported the velocity of joint angular

perturbations are too low to induce a sizeable reflex response. It is more likely a voluntary, online error correction, in an attempt to correct initial deviation to the right. This voluntary response causes overcompensation for the PF, as seen in the torque traces, resulting in deviation to the left of the straight line joining the start and target positions (see hand path learning kinematics). Following the large shoulder flexor torque in these early trials, causing deviation to the left, an increase in shoulder extension torque was used to locate the target. A similar pattern is seen in elbow torque profiles, where online overcompensation occurs (too much elbow flexion torque), followed by slightly greater than predicted extensor torque to reach the target.

Figure 26 shows that, averaged over the 7 subjects for whom we had the NF data needed to predict the PF torques, a lack of shoulder torque was responsible for causing the rightward displacements in the first PF trials. Following complete adaptation (lower panels, Figure 26), we note that subjects initially overcompensated for the PF, producing a small excess of shoulder flexor torque, explaining the characteristic initial movements to the left, following adaptation. The peak elbow extensor torque was well matched to the predicted value, although later in the trial neither shoulder and elbow torque were well matched to PF values. This may be because subjects did not follow the NF path in the early portion of the movement, requiring amended torque profiles later in the trial to locate the target.

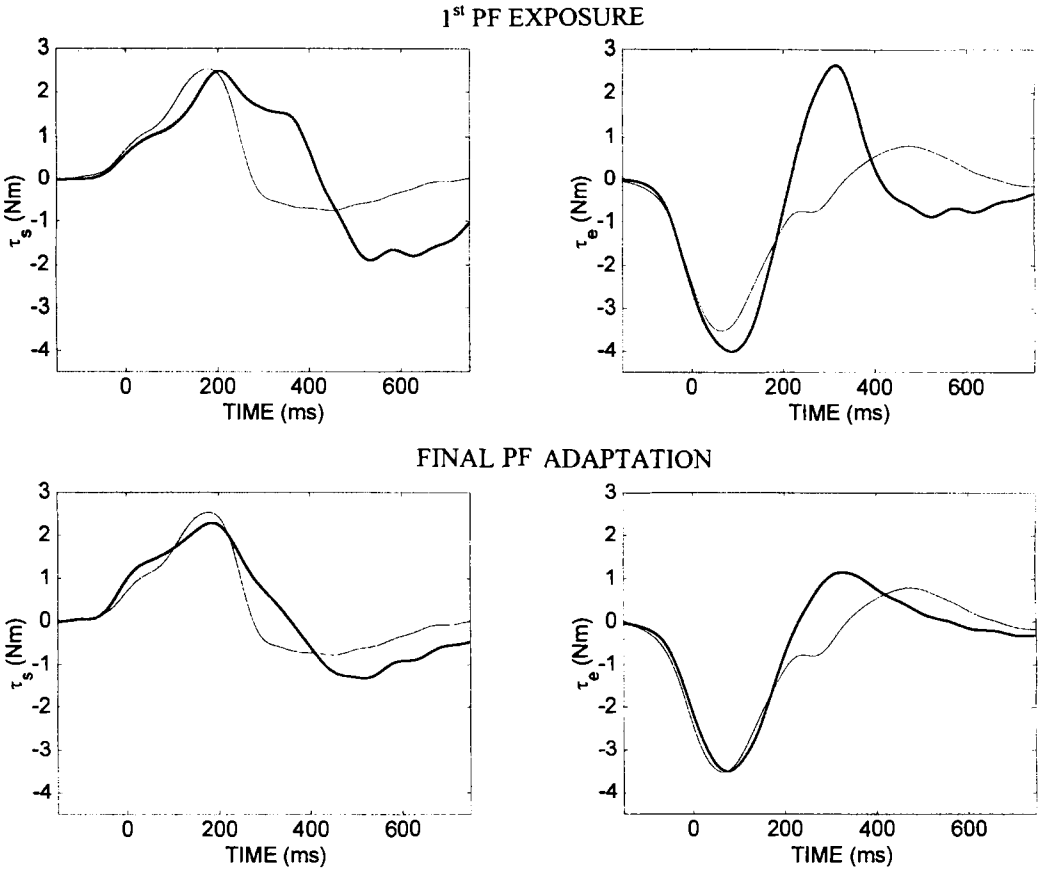

Figure 26. Shoulder (τ_s, left) and elbow (τ_e, right) torque as a function of time during PF learning. Flexion torque is positive, and extension is negative. The thin traces show the joint torque required to move on a path, in the PF, similar to the NF. The thick traces show the torque, averaged over the 7 subjects for whom we had NF data, for the first trial in the PF (top) and following complete adaptation to the PF, i.e., averaged over trials 146-149 (bottom). The PF boundary occurred at ~240 ms, after which there was no external force applied.

Subtle differences in the shoulder torque during the earliest phase of the movement indicate that the feedforward command to the muscles did not accurately reproduce the torque required to compensate for the PF. Larger differences are evident later on. The main difference in the shoulder torque profile occurred following peak torque production, at around 0.05 m, where subjects reduced the shoulder flexor torque too slowly. This resulted in surplus shoulder flexor torque, relative to the predicted values, beginning around 0.07 m. In contrast, subjects matched peak elbow extension torque well, but then

extension torque was reduced too quickly. Elbow torque continued to change, culminating in excess elbow flexor torque, relative to the predicted torque values, also beginning around 0.07 m.

The small excesses in shoulder and elbow flexion torque, relative to the predicted values, were compensated during the latter part of the movement, such that subjects were still able to locate the target accurately. Over the latter portion of the movement, from 0.10 m to 0.25 m, subjects presumably learned to amend their joint torque such that they compensated for the initial differences between the predicted and measured joint torques.

The analysis of joint torques have highlighted that there are disparities between the measured joint torques and the joint torques required to move on a path similar to that in the NF, but in the presence of the PF. Specifically, on the first PF trial we found that shortly after movement onset subject's shoulder flexor torque was lower than the shoulder flexor torque required to move on the NF path. This net deficit of shoulder flexor torque resulted in hand paths that were displaced to the right. Following the practice period, subjects adapted their shoulder flexor torque such that it was initially larger than the predicted shoulder flexor torque. This led to the characteristic bowing of hand paths to the left, before a reversal of sign in the net shoulder flexor torque, i.e., subjects' torque dropping below the predicted torque, resulted in hand paths returning towards the centreline. In the region where the PF was inactive, there was an initial surplus of shoulder flexor torque, followed by a surplus of shoulder extension torque (see Figure 26, lower left panel) that explains why hand paths bowed to the right before returning towards the centreline. In the region where the PF was active there was very little difference between the actual and predicted elbow torque profiles, suggesting that the differences in shoulder torque was responsible for the bowing of the hand path. In the region where the PF was inactive, there was initially a surplus of elbow extensor torque, followed by a deficit, relative to the predicted values, which would have also contributed to the rightward, then leftward bowing of hand paths.

Modification of the EMG during learning

To fully appreciate the nature of the motor command that the CNS issues, we must consider a direct measure of that command, rather than simply focusing on the outcome of that command in terms of performance (hand path error) or force output. Subjects may indeed be progressively adapting to the PF by reducing the amount of muscular cocontraction without any visible change in either hand paths of force profiles. Such an adaptation is an important aspect of learning which must not be overlooked. We used surface EMG recordings of 6 muscles as a measure of the motor activity on each trial in the PF. In a similar manner to the kinematic and force analysis, we considered how the motor commands represented by our EMG recordings varied, as subjects adapted to the PF. To reduce the amount of variability in the dependent measure, which limits the practicality of using raw EMG signals without first averaging over many trials, we calculated root mean square (rms) EMG over selected time intervals in all PF trials. This enabled us to reduce variability without having to average over trials, which would have limited the quality of information that could be obtained during the learning period, where motor commands may have been continually amended. Two intervals were considered during each trial. Early rmsEMG was defined as rms EMG in a 300 ms interval beginning 150 ms prior to movement onset, i.e., until 150 ms after movement began (Figure 27). Given that it took subjects approximately 240 ms to reach the PF field boundary at 0.1 m, and considering electromechanical delay, EMG in this early interval represents the feedforward motor command, issued by subjects to initiate movement and compensate for the PF, as well as any stretch reflex activity due to perturbations caused by the PF. Late rmsEMG was defined as rms EMG in the interval beginning 150 ms following movement onset until data collection ceased, 1250 ms following movement onset (Figure 28). This late interval covered both the latter stages of the movement including all online responses to the field as well as subsequent stabilization at the target. All PF trials except for the initial trial that, as we explained previously, can be designated the stimulus to initiate learning, were included in this analysis. All 30 CT's were excluded, primarily because they may have had a suppressed reflex component, based on the analysis of their kinematics compared to PF trials, and were unlikely to have included

any online corrective responses. Furthermore, subjects may have detected the channels and amended their motor command late in the movement. This left 119 trials in each regression analysis.

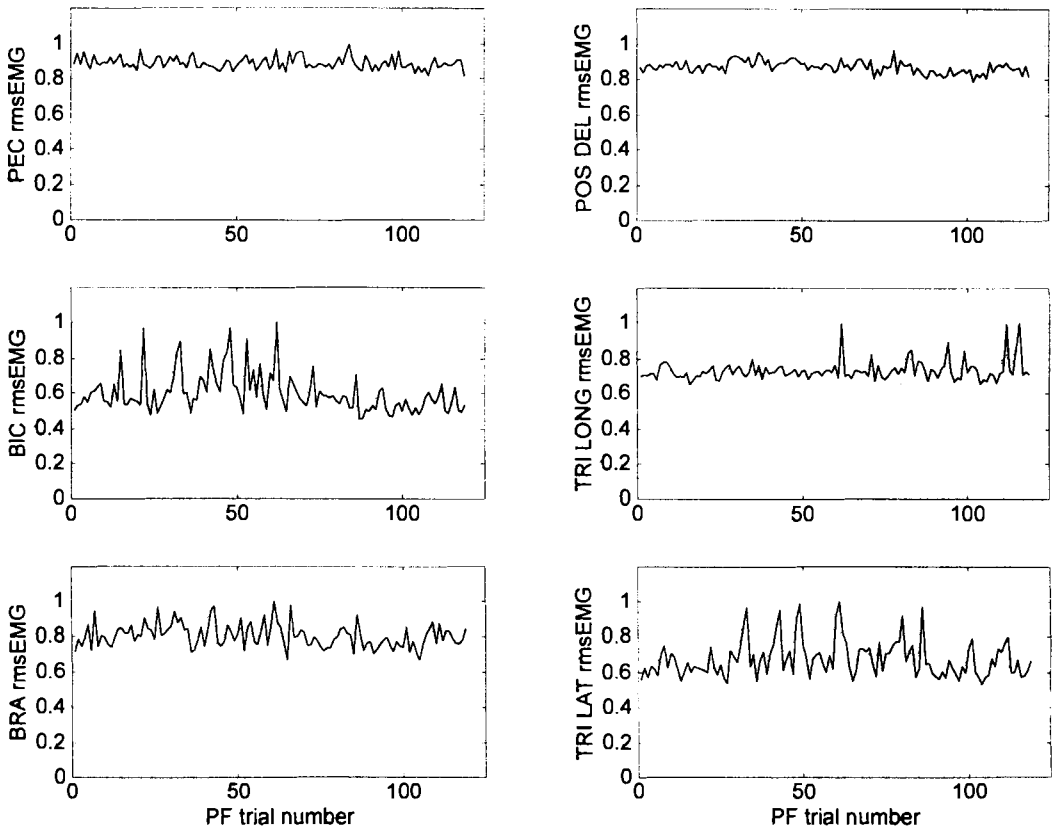

Figure 27. rmsEMG in the early interval of the movements.
Data is shown for two shoulder muscles (pectoralis major, PEC, and posterior deltoid, POS DEL), two biarticular muscles (biceps brachii, BIC, and long head of the triceps, TRI LONG) and two elbow muscles (brachioradialis, BRA, and lateral head of the triceps, TRI LAT) All PF trials for subject DF are shown. Data are normalized to the maximum rms value obtained for each muscle, and trends are representative of all subjects.

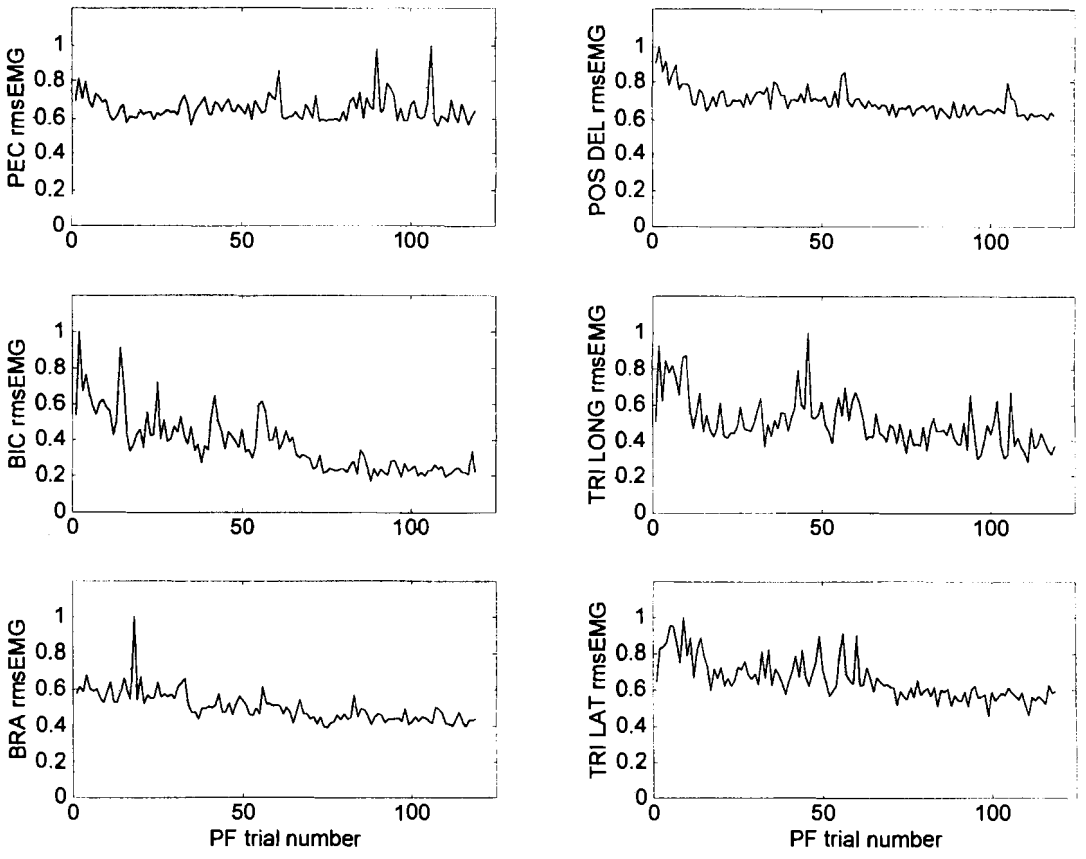

Figure 28. rmsEMG in the late interval of the movements.
Data is shown for two shoulder muscles (pectoralis major, PEC, and posterior deltoid, POS DEL), two biarticular muscles (biceps brachii, BIC, and long head of the triceps, TRI LONG) and two elbow muscles (brachioradialis, BRA, and lateral head of the triceps, TRI LAT) All PF trials for subject DF are shown. Data are normalized to the maximum rms value obtained for each muscle, and are representative of all subjects.

A linear regression was used to determine if the early and late rmsEMG changed systematically as subjects adapted to the PF, i.e., was the gradient of the least-squares linear fit significantly different from zero? Contrasting results were found for the early and late periods of the trial, confirming the impression from Figures 27 and 28. We found very weak correlations between the early rmsEMG and trial number. For the 6 muscles for the 9 subjects, i.e., 54 linear regressions, R-squared values were very low, generally less than 0.05, with only 2 values larger than 0.20. However, 31 of the 54 regressions exhibited gradients significantly less than zero ($p<0.05$), suggesting a reduction in EMG

over the early period, although 4 regressions exhibited positive gradients ($p<0.05$) whereas 29 gradients were not significantly different to zero. Taken as a whole, the low R-squared values and inconsistency of the gradients suggests no overall systematic trend for a reduction in the early motor activity. In contrast, we found much more consistent results for the late rmsEMG. Of the 54 regressions 49 had gradients that were significantly less than zero ($p<0.02$ for all 49 tests and generally $p<0.0001$) with R-squared values that were generally greater than 0.20 and as high as 0.68. These gradients tended to be at least an order of magnitude greater than the corresponding gradients for the early rmsEMG. Of the remaining regressions, 4 had negative gradients but were not significantly different from zero at the *a priori* significance level. Only 1 of the 54 regressions exhibited a positive gradient (increase in rmsEMG), although this was not significantly different from zero.

The result from the rmsEMG analysis is evident in the temporal EMG traces shown in Figure 29. There was little change in the early EMG (−150 ms to 150 ms) between trials at the beginning and end of the learning regimen. Large decreases in the late EMG (150 ms onward) are apparent at the end of learning, compared to the beginning of learning.

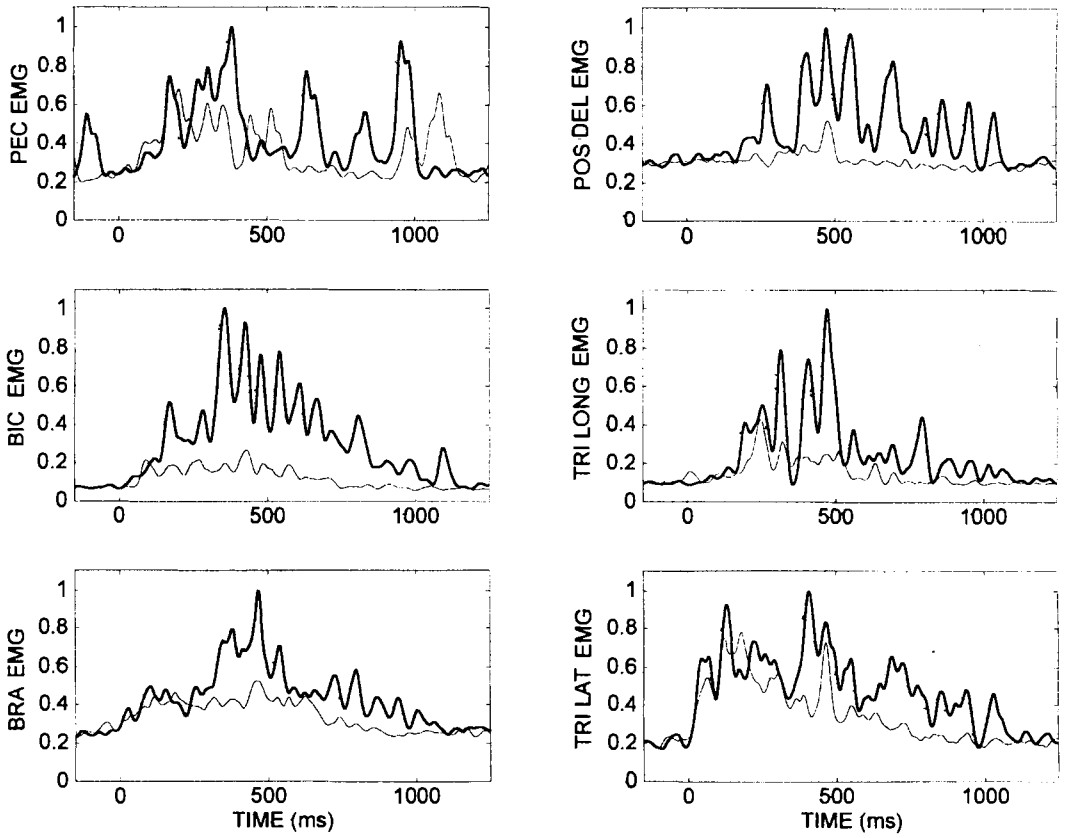

Figure 29. EMG activity in the 6 muscles for subject DF.
Thick lines represent average activity in the first four trials of learning while thin lines show average activity in the final 4 PF trials at the end of learning. Large decreases in the late EMG reveal that it is in this period that cocontraction drops significantly during learning.

Since the PF was only present during the first 0.1 m of the movement, muscular activity to compensate for the PF should only have occurred before reaching the force field boundary. No reduction in early rmsEMG was seen in any muscle, suggesting that at the onset of movement cocontraction may play a part in compensating for the PF, even following learning. In the final 0.15 m of the movement the PFM exerted no lateral force. There was little need for cocontraction during this period, especially following learning when subjects' feedforward model should have become more accurate. Consistent with this, a clear reduction in muscular activity in the late portion of the movement is visible both in Figure 28 and 29.

Beyond the PF boundary the only muscular activity required was that needed to extend the arm to the target and subsequently decelerate it to achieve a stationary final position. This required shoulder flexion torque and elbow extension torque. Linear regressions revealed that the activity of all muscles in the late portion of the movement decreased, although we noted that the magnitude of the gradients of the pectoralis major (shoulder flexor) and long head of the triceps (biarticular extensor) were less than the other four muscles, averaged across all subjects. The fact that we see a less obvious reduction in these muscles is consistent with the fact that a certain amount of activity in these muscles is required to move to the target. The large reduction in the biceps muscle, signified by the largest negative gradient in the regression (averaged across all 9 subjects), suggests that this biarticular muscle helps in stabilizing the arm through cocontraction, in the early trials, which is greatly reduced as learning progresses.

In summary, we have shown that subjects adapt quite well to the PF. We have analysed several kinematic measures and shown that significant improvement in performance occurred within the learning period. Following learning subjects produced hand paths in the PF with a characteristic bowing to the left, before returning to the start-target line near the PF boundary, at 0.1m, while the latter part of the movement was relatively straight. The paths were noticeably different from those in the NF. We found that subjects were able to adapt their lateral force production in CT's, modifying the internal model to more accurately compensate for the PF. Following learning, subjects produced excess shoulder flexor torque in the earliest portion of the movement, relative to that required to compensate for the PF, resulting in leftward deviation. Later, a deficit in shoulder flexor torque resulted in movement back towards the start-target line. Elbow extensor torque was initially well matched to that required to compensate for the PF. Later disparities for both shoulder and elbow torque may reflect amended feedforward commands to account for the early differences between PF and actual shoulder torque.

We found adaptation in the level of muscle cocontraction used by subjects in the latter part of the movement enabling them to move in a more energy-efficient manner without

any associated decrement in performance. The next section of results will focus on how subjects adapted to the perturbed strength fields, and how the PF was re-learnt between these ΔPF sets.

Results- ΔPF perturbation sets

Following the learning of the PF, subjects were exposed to 315 trials in which the strength of the PF field was unexpectedly altered (see Methods). Here we considered how subjects learned both of the perturbed fields, as well as how quickly subjects relearned the PF, following sets of perturbed strength trials. In both cases, kinematic and electromyographic data were analyzed, using similar techniques to those used in the PF learning analysis.

Did the ΔPF sets affect the PF hand paths?

The 8 PF trials prior to the introduction of any ΔPF trials ("learnt-PF") were compared to 8 PF trials late in the experiment, following exposure to the ΔPF fields ("late-PF") to determine whether any significant change in the PF hand paths occurred throughout the experiment (Figure 30). We hypothesized that because exposure to the low and high strength perturbations throughout the experiment did not alter the overall mean field strength, motor commands and, hence, hand paths would remain unchanged. If there was a significant change in PF hand paths, we would have to consider how to approach the analysis of the ΔPF trials carefully- a comparison of ΔPF hand paths to PF hand paths would be of little benefit if PF hand paths themselves evolved after subjects were exposed to ΔPF sets.

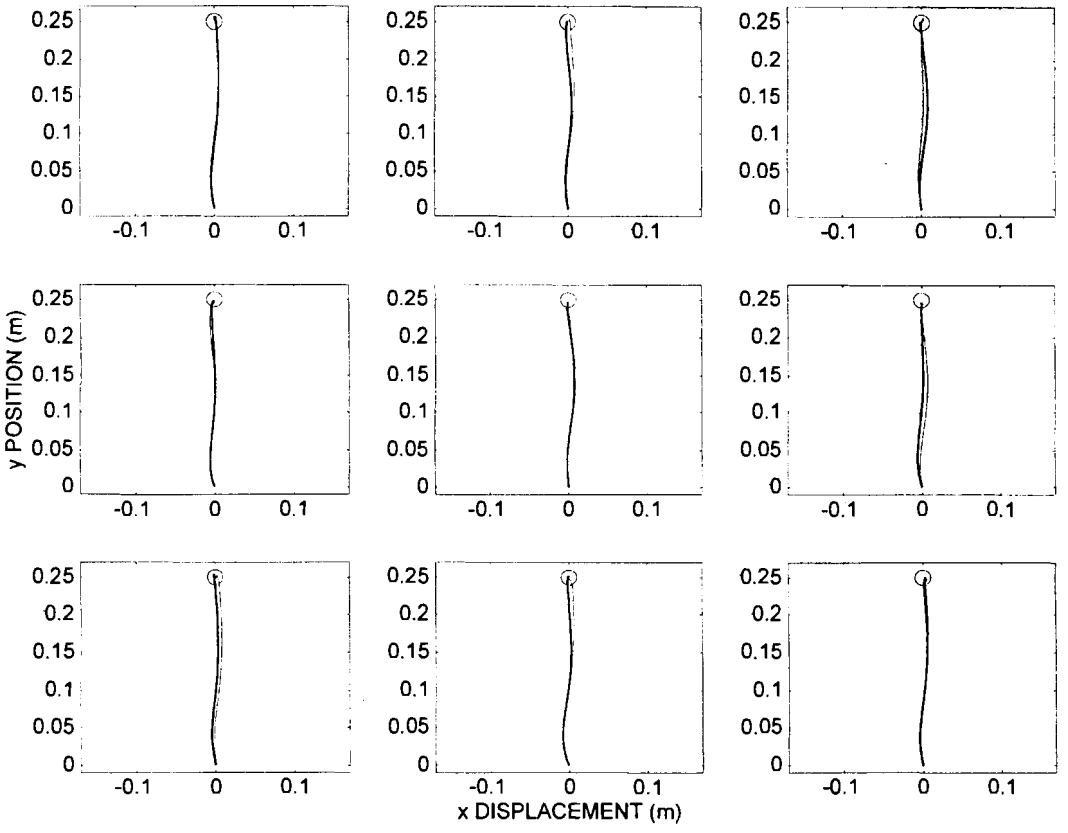

Figure 30. Hand paths for the 9 subjects in the learnt-PF and late-PF trials.
Hand paths are shown for trials immediately prior to (learnt-PF, thick lines), and following (late-PF, thin lines), all ΔPF trials. Traces are averaged over 8 trials in each condition. Top left-bottom right subjects MH, US, AN, JJ, JH, SM, DF, RO and AR.

Absolute and signed area between the hand paths and the straight line joining the start position and the target, termed absolute and signed hand path error, respectively, were calculated for learnt-PF and late-PF trials. Two-way repeated measures ANOVA revealed a non-significant field main-effect for both of the dependent variables ($p=0.632$ absolute hand path error; $p=0.583$ signed path error). The trial number main-effect and the interaction between the field and trial number were also highly non-significant. We conclude that subjects' performance, in terms of the *straightness* of the trajectory, measured by the absolute hand path error, and signed path error, was not altered due to the random application of sets of perturbed strength trials. This result enabled us to analyze ΔPF trials within the experimental set with respect to the learnt-PF path, without

the need to compensate for change in the learnt-PF path, i.e., the reference trajectory remained unchanged throughout the experiment.

Although we found no significant change between hand paths in the PF prior to and following imposition of all ΔPF sets, there may have been a change in the associated muscle activity. One possible scenario may have been that the uncertainty caused by the random changes of field strength to ΔPF_{low} and ΔPF_{high} had caused an increase in the level of cocontraction. To determine whether there were any changes in EMG, we conducted a two-way RM ANOVA on the rmsEMG values in the PF-learnt and PF-late trials. We considered both the feedforward EMG (early rmsEMG), from 150 ms prior to movement onset until 150 ms after onset, and EMG later in the trial (late rmsEMG), from 150 ms until the end of data collection. We found no significant differences in any of the six muscles, in either of the defined periods between the learnt-PF and late-PF trials (timing main effect: $p>0.05$ for all tests, with $p>0.1$ for 11 out of 12 tests). The trial main effect and the interaction were also non-significant ($p>0.1$), suggesting that there was no trend in the rmsEMG within the 8 trials of each set. Since we found hand paths did not change between the learnt-PF and late-PF trials, this rmsEMG analysis indicates that both the internal model and level of cocontraction were similar, prior to and following imposition of the ΔPF sets. We conclude that subjects' performance in the PF was similar, prior to and following the ΔPF sets, both in terms of the hand paths and the motor commands issued to produce the movements. It appears that subjects were able to retain the well-learnt motor commands required to move in the PF, despite the interference caused by the ΔPF sets.

Relearning of the PF field following perturbation trials to the lower and higher ΔPF

As explained in the methods section, a randomly chosen number of PF trials, between 6 to 8, occurred following each ΔPF set, such that subjects could readapt to the previously learned PF, prior to exposure to the subsequent ΔPF set. It was important that subjects be able to re-establish baseline performance in the PF between exposures to the ΔPF. To determine how quickly subjects were able to readapt to the PF field following exposure to

ΔPF sets we compared PF trials following ΔPF trials, i.e., post-perturbation PF trials, to learnt-PF trials. Performance in the post-perturbation trials was analysed separately for trials that followed ΔPF_{high} and ΔPF_{low}. Each post-perturbation PF trial was compared to the final 4 PF trials prior to the commencement of perturbed strength force fields. In the same manner as we analyzed adaptation to the PF, a performance-related score was assigned, based on how straight the post–perturbation PF trials were relative to the 4 learnt-PF trials. A score of 2 was assigned if the maximum deviation from a straight line joining the start and target positions in the post-perturbation trial was larger than that of all 4 learnt-PF trials. A score of 1 was assigned if the maximum deviation in the post-perturbation trial was larger than that of 3 of the learnt-PF trials, with a score of 0 indicating that the post-perturbation trial was straighter than 2 and less straight than 2 learnt-PF trials. Scores of –1 and –2 related to the post-perturbation trial being straighter than 3 or all 4 learnt-PF trials, respectively. The score from each of the post-perturbation trials was averaged across all subjects for the 12 trials corresponding to the same position in the 1^{st}-6^{th} PF trials, the 8 trials in the 7^{th} position and the 4 trials in the 8^{th} position (Figure 31).

Student's t-tests (2 tailed, *a priori* α=0.05) were carried out to determine whether the score for each trial position was significantly different from zero. The first 2 PF trials following an increase in the strength of the force field were found to be less straight than the learnt-PF trials ($p<0.0001$, $p=0.016$, respectively). Performance on the remaining 6 trials was not significantly different from performance on the learnt-PF trials. Following a decrease in the strength of the force field, the first 3 PF trials were significantly less straight than the learnt-PF trials ($p<0.0001$, $p<0.001$ and $p<0.001$, respectively). The next 5 trials were not statistically different from the 4 learnt-PF trials, i.e., the score was not significantly different from zero. This result indicates that subjects were able to readapt to the PF within 2 to 3 trials, despite the interference associated with adaptation to the perturbed strength fields. This allowed us to proceed with the analysis based on the substantiated premise that subjects were performing at baseline levels in the PF prior to each ΔPF set.

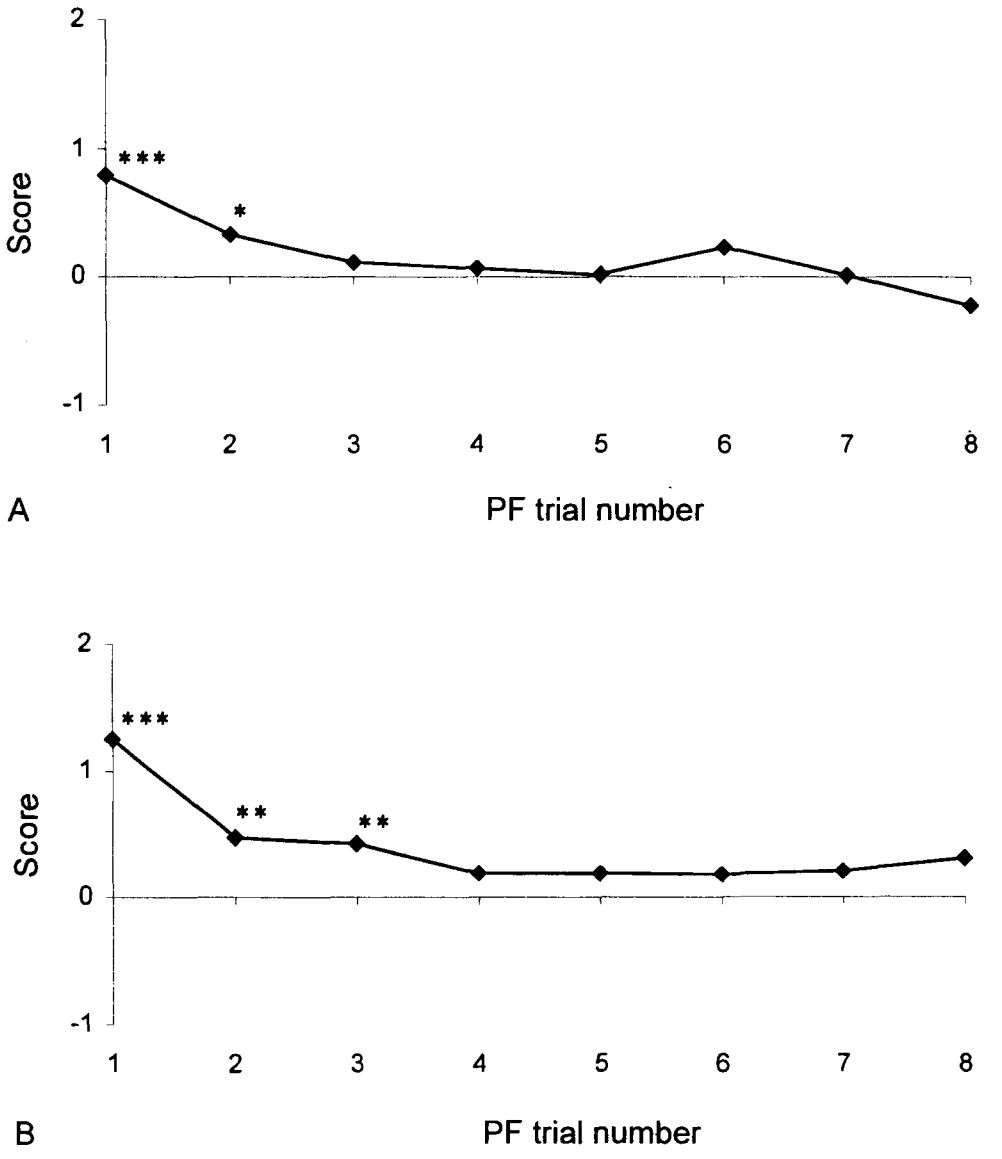

Figure 31. Performance scores for maximum deviation

A) Scores for the 8 post-perturbation trials following ΔPF_{high} trials and B) following ΔPF_{low} trials. Data are averaged over the 6-12 trials in each position and over the 9 subjects. *, ** and *** indicate scores significantly different than zero ($p<0.002$, $p<0.001$ and $p<0.0001$, respectively).

The effect of CT's within ΔPF trials

As described in the Methods, 25% of ΔPF trials in positions 2-5 were replaced with CT's for the purpose of isolating feedforward commands from feedback commands. The CT's were intended to eliminate reflex responses from the EMG by minimizing trajectory error so that only feedforward motor commands would effectively contribute to the EMG during the first 200 ms of the movement. The very nature of CT's, i.e., restricting subject's to a straight line between the start and target positions, is such that no error feedback is available on these trials, i.e., subjects are unable to judge how accurate the movement resulting from their issued motor command would have been in the ΔPF, relative to a straight line (or their learnt path). In the following ΔPF trial, subjects would have been unable to update their motor commands based on error feedback from the previous trial. We hypothesized, as with the PF learning, that the absence of error feedback in the CT would adversely affect performance in ΔPF trials that followed CT's, relative to ΔPF trials that were preceded by another ΔPF trial. We assumed that the effect of the CT would not vary with the number of exposures to the ΔPF, as we did not find any consolidation of learning between ΔPF sets, i.e., subjects were no more adept at moving in the ΔPF in the later sets than the first set, in terms of hand path straightness (see Hand paths in the ΔPF trials). To enable us to average over all ΔPF trials in the same position within the ΔPF set (a process which would reduce the variability of our measures and lead to more powerful conclusions), we first determined if performance in all ΔPF trials in the same position was similar. If ΔPF trials that followed CT's exhibited significantly worse performance in all positions within the ΔPF set, then including them in such an average would be unwise. To determine the effect of a CT on the subsequent ΔPF trial we compared ΔPF trials in positions 3, 4 and 5 that were preceded by a CT with those in the same position that were preceded by a ΔPF trial (Figure 32). No CT's occurred in position 1 of the ΔPF, hence, trials in position 2 were not included in this analysis.

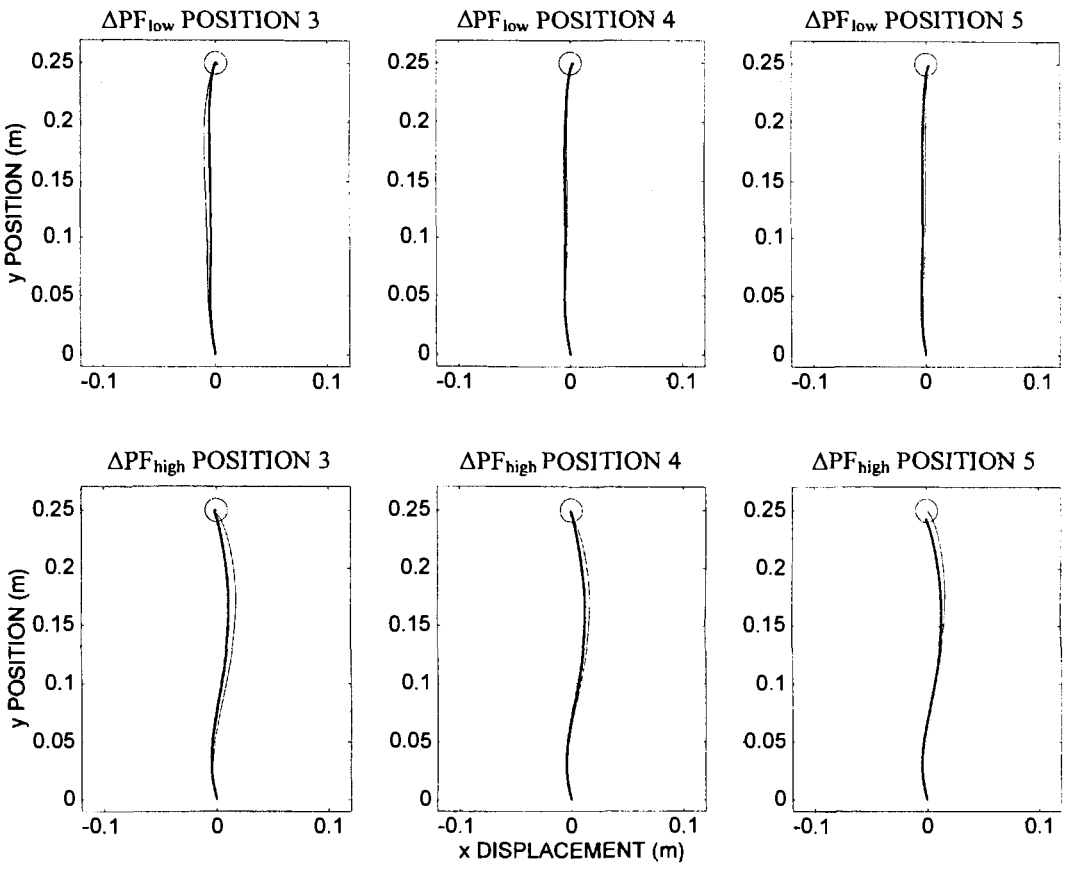

Figure 32. Average hand paths for the 9 subjects in the ΔPF_{low} and ΔPF_{high}.
Thick traces represent trials preceded by another ΔPF trial, while thin traces show trials that were preceded by a CT. Absolute hand path error was statistically larger for ΔPF_{high} in positions 3 and 4 that followed CT's, compared to those that followed other ΔPF trials.

Absolute area between the actual hand path and the learnt path in the PF was used as a measure of absolute hand path error. We used Student's paired t-tests (two-tailed *a priori* α=0.05) to compare absolute hand path error in the 3 ΔPF trials preceded by a CT and 6 ΔPF trials preceded by another ΔPF trial, over the 9 subjects. A two way ANOVA was not possible due to the different number of trials in each condition. Analyses for ΔPF trials in positions 3, 4 and 5 of the increased and decreased strength force fields were carried out individually, yielding six distinct t-tests.

Averaged over the 9 subjects, larger absolute hand path error was found in trials that followed a CT compared to those that followed another ΔPF trial in positions 3 and 4

with the higher force field strength ($p<0.05$ in both cases). Although larger absolute hand path error was also found in trials in the 5th position that followed a CT this was not significant ($p>0.15$). With the lower force field strength all three tests revealed non-significant differences in the absolute hand path error between the trials that were preceded by a CT and those preceded by a ΔPF trial ($p> 0.15$, 0.40 and 0.80 for positions 3, 4 and 5, respectively). Overall, this result suggests that performance in ΔPF trials was affected by whether the previous trial was a CT or another ΔPF trial in some cases, but not all cases. Performance was affected more in the higher strength field than in the lower strength field.

One approach that we considered was to analyze each individual's performance separately. In this manner, we would have been able to determine if the CT's affected certain subjects' performance differently. Here a non-parametric test would have been used, due to the fact there were only 3 and 6 trials in the two groups i.e., data could not be presumed to be normally distributed. Although this could have been done, our very low sample sizes would have resulted in critical U values almost as large as the total degrees of freedom, i.e., the test would have been extremely conservative. Given the variability in subjects' paths, a significant result would have been almost impossible to achieve. To circumvent this problem we calculated the difference between the absolute area between each of the 3 ΔPF trials that followed a CT trial and each of the 6 ΔPF trials that followed another ΔPF trial. This yielded 18 differences in absolute area for each subject, for each of the perturbed field strengths and 3 trial positions. A Kolmogorov-Smirnov test was performed to ensure that each set of 18 values was consistent with being normally distributed, thereby permitting us to use a parametric test. We tested whether the mean difference for each subject, trial position and field strength combination was significantly different from zero. A z-score was computed based on the mean and standard error of the differences in absolute area for each combination of subject, field strength and trial position. In the ΔPF_{high}, 5 out of 9 subjects had worse performance following CT's in positions 3 and 4, i.e., the mean difference in absolute area between the ΔPF trials that followed CT's and those that followed ΔPF trials was significantly greater than zero

($p<0.05$). In position 5, performance was worse for only 2 of the subjects. In the ΔPF_{low}, 4 subjects' performance was worse following the CT's in positions 3 and 4. The performance of only 1 subject was statistically worse in position 5.

The effect the CT's have on the subsequent ΔPF trials is not conclusive, and the effect we hypothesized is certainly not true in all cases It appears that performance is affected by the imposition of CT's, although this is not consistent across subjects, field strength or trial position. In subsequent analysis of the ΔPF learning we did not consider any trials immediately following a CT. When analyzing trials in positions 3, 4 and 5 of the ΔPF, we only considered 6 trials, i.e., although there were 12 sets of trials in each ΔPF field, we excluded the 3 ΔPF trials that were replaced with a CT and the 3 ΔPF trials that followed a CT. We considered all 12 trials in positions 1, as no CT's were present in position 1, and 9 trials in position 2, as there were 3 CT's in position 2, but position 2 was never a post-CT.

Hand paths in the ΔPF trials

Prior to the analysis of ΔPF sets, we wanted to ensure that there was no consolidation of learning between each consecutive exposure to the ΔPF_{low} or ΔPF_{high}. To enable us to average over consecutive exposures to the perturbed strength field, we needed to ensure that performance on these trials was independent of the number of exposures to the field, i.e., we needed to ensure that subjects did not improve their performance in the ΔPF trials on each subsequent exposure. If we found that performance did improve with repeated exposures, averaging over trials would be of little benefit, and an alternative analysis would have to be considered. We did this by testing whether performance in the ΔPF was a function of the number of times subjects were exposed to ΔPF sets. Using absolute hand path error as our measure of performance, a RM ANOVA was used to determine if performance was affected by the number of exposures to the ΔPF sets. We found that in both the lower and higher strength perturbed fields the exposure number main effect i.e., the number of times the subjects had been exposed to the ΔPF (1-12), and the interaction between the number of exposures and the trial position in the ΔPF set i.e., position 1-5 in

the ΔPF set were non-significant ($p>0.14$ in all cases, Greenhouse-Geisser correction). This confirms that learning was not consolidated across sets of ΔPF trials, allowing us to average across trials in the same position within the ΔPF sets.

On the first ΔPF trial of each set, subjects were perturbed to the left, in the ΔPF_{low}, and to the right in the ΔPF_{high}, relative to the preceding PF trial. The displacement on these trials was approximately symmetric, relative to the preceding PF trial, such that both trials had a similar perturbing effect, albeit in opposite directions. On the second ΔPF trial, trajectories were substantially straighter than on the first ΔPF trial, but displacements did not appear to be further reduced in the 3rd, 4th and 5th trials, such that the 5th ΔPF trials generally remained displaced relative to the learnt-PF path. Figures 33 and 34 show the learnt-PF trajectories together with the 1st and 5th trials in the low and high strength perturbed fields. We have not shown the 2nd, 3rd and 4th ΔPF trials in these figures for clarity, as these trajectories were very similar to those of the 5th ΔPF trials. The 5th ΔPF_{low} trials were displaced to the left of the start-target line almost along their entirety, whereas the learnt-PF trials showed the initial bowing to the left, before rightward displacement occurred later in the trial. The 5th ΔPF_{high} trials were also initially displaced to the left before larger displacements to the right occurred. The position at which the maximum rightward displacement occurred in the 5th ΔPF_{high} trials was around 0.15 m, unchanged from the learnt PF path. The maximum displacement to the left, relative to a straight line, on the 5th ΔPF_{low} trials was more variable in terms of location across the subjects, but was smaller than the rightward displacements on the 5th ΔPF_{high} trials for the majority of subjects.

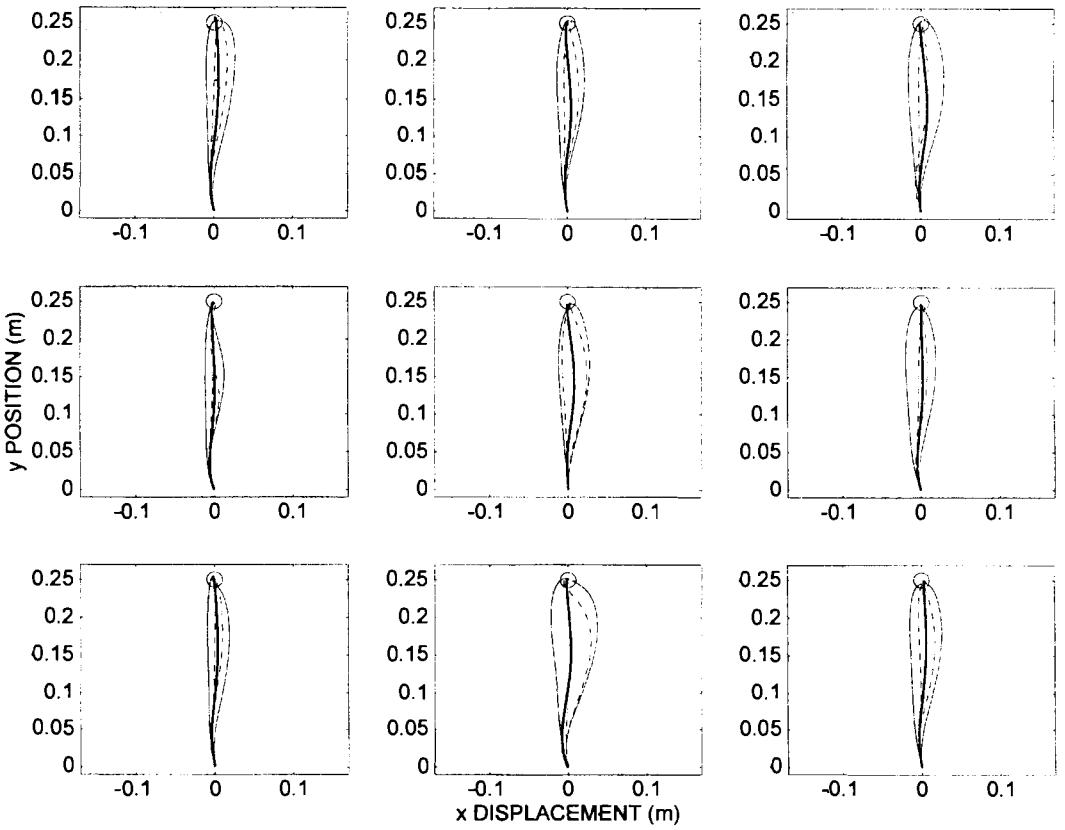

Figure 33. Hand paths in the perturbed strength fields.
The 1st ΔPF_{low} and ΔPF_{high} are shown as thin traces displaced, to the left and right, respectively, of the thick trace representing the learnt-PF path. The 5th ΔPF_{low} and ΔPF_{high} trials (dashed traces) are generally still displaced relative to the learnt-PF, but reveal significant learning in the perturbed strength fields. Each plot shows a single subject and represents the average over the 8 learnt-PF, 6 ΔPF_{low} or 6 ΔPF_{high} trials. Top left-bottom right: subjects MH, US, AN, JJ, JH, SM, DF, RO and AR.

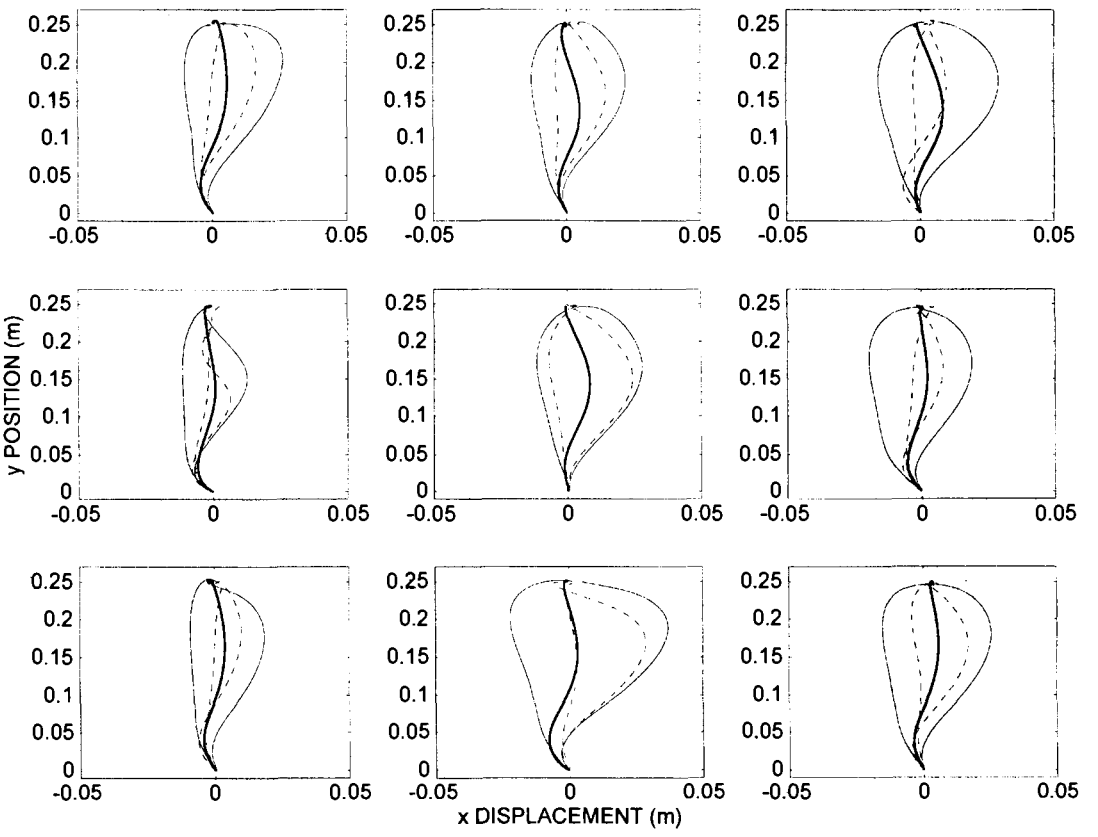

Figure 34. The same data as shown in Figure 33, but on an expanded scale.
The larger scale enables a clearer impression of the amount of perturbation on the 1st ΔPF trials and subsequent straightening achieved by the 5th ΔPF trials.

To determine whether there was any evidence of progressive straightening of hand paths between the 2nd and 5th ΔPF trials, a two-way RM ANOVA was used to compare the absolute hand path error in the ΔPF trials in all 5 positions for both the low and high strength fields. This analysis was of interest because we had found that subjects progressively decreased hand path error in the PF, as measured by incremental reduction in maximum deviation and absolute hand path error, over 15-20 trials (see Figure 13). Therefore, it was somewhat surprising that subjects did not appear to improve performance following the 2nd ΔPF trial. We wanted to determine statistically if there was any reduction in trajectory error following the 2nd ΔPF trial. To enable an ANOVA to be undertaken and make pairwise contrasts if a significant effect was found, we selected 6 of

the 12 PF trials in the 1st position and 6 of the 9 trials in the 2nd position in the ΔPF, to compare to the 6 ΔPF trials in the 3rd, 4th and 5th positions, such that we had the same number of trials in each position (Figure 35).

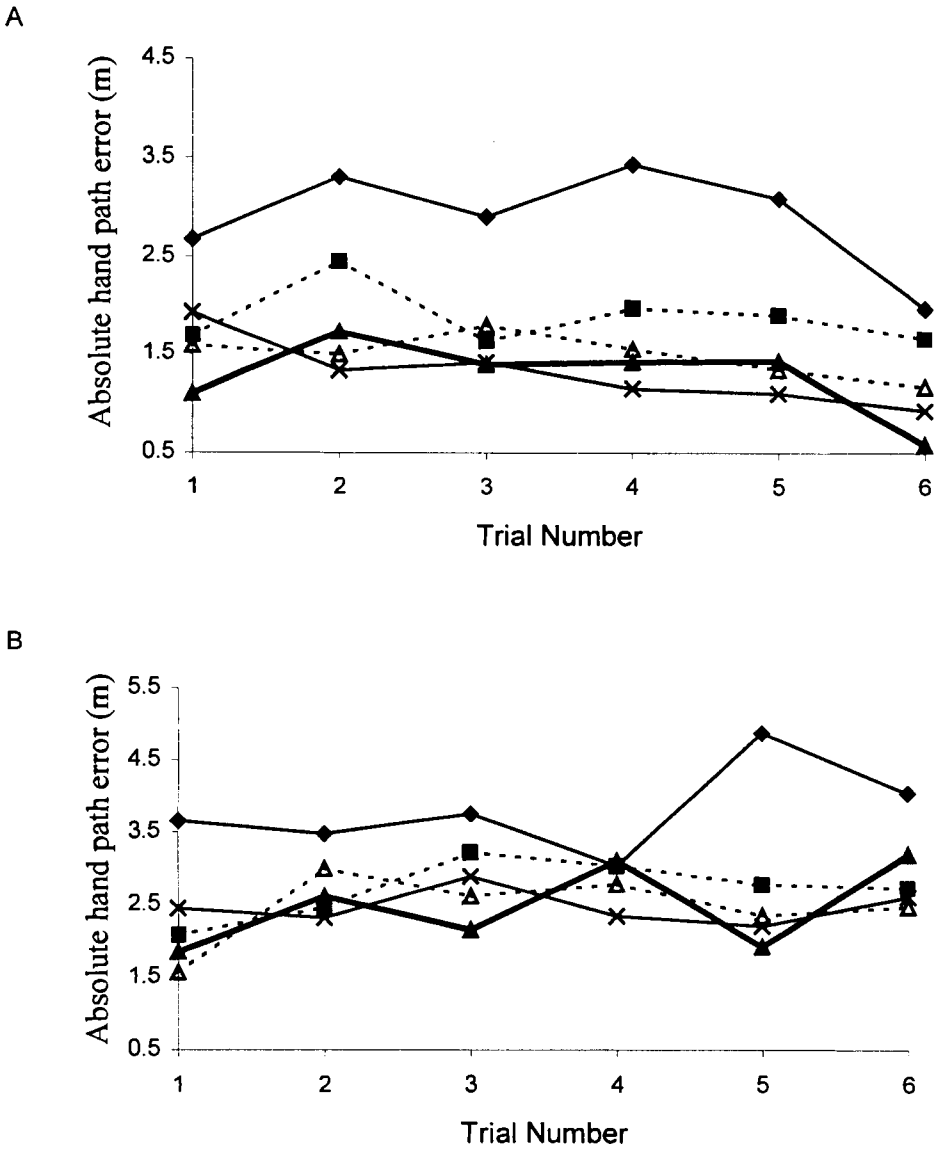

Figure 35. Absolute hand path error in the 5 lower, and higher, strength perturbed trials
A) The 5 ΔPF_{low} trials and B) the 5 ΔPF_{high} trials. Data are averaged over the 9 subjects. 1st trial: ◆, thin line. 2nd trial: ■, dotted line. 3rd trial: Δ, dotted line. 4th trial: X, thin line. 5th trial: ▲, thick line.

In both the low and high strength perturbed fields we found a significant position main effect, i.e., the absolute hand path error varied across trials in positions 1-5 in both ΔPF_{low}

and ΔPF_{high} ($p=0.002$, $p=0.001$ for low and high strength fields, respectively; Greenhouse-Geisser correction). The exposure number main effect and interaction between the exposure and position were non-significant in both fields ($p>0.14$ in all cases, Greenhouse-Geisser correction), confirming that subjects' performance was not affected by the number of times that they were exposed to each perturbed strength field within the experiment, i.e., learning was not consolidated across sets of trials in the same field. Simple pairwise contrasts, i.e., comparisons of performance in the 1st, 2nd, 3rd and 4th ΔPF trials to the performance in 5th ΔPF trial in a pairwise manner, revealed that the absolute hand path error on the 1st ΔPF trial for both the lower and higher strength force fields was, indeed, significantly larger than that of the 5th ΔPF trial ($p=0.002$ in both cases) The 2nd trial in the ΔPF_{low} was also found to be less straight than the 5th ΔPF_{low} trial ($p=0.025$), although the magnitude of the difference was very small compared to the magnitude of the difference between the 1st ΔPF_{low} and the 5th ΔPF_{low} trials (Figure 35 A). All other comparisons to the 5th ΔPF trial remained non-significant ($p>0.12$ in all cases) in both fields. Overall, this result confirms that performance did improve in the ΔPF field, although most of that improvement occurred between the 1st and 2nd ΔPF trials, with little further improvement in the next 3 trials. Possible reasons for the minimal amount of learning, in terms of reducing the absolute hand path error, following the 2nd ΔPF trials, are considered in the Discussion.

A two-way RM ANOVA was used to compare the absolute hand path error in the learnt-PF trials with that of the 5th ΔPF_{high} and the 5th ΔPF_{low} trials. In this way, we could gauge how well subjects had adapted to the perturbed strength fields within 5 trials, compared to the complete adaptation to the PF. We compared the 6 PF trials prior to any perturbations, i.e., learnt-PF trials, with the 6 ΔPF trials in the 5th position for both the lower and higher strength force fields (Figure 36). We found a significant force field main-effect ($p=0.003$), i.e., the absolute hand path error differed between the learnt-PF, 5th ΔPF_{high} and 5th ΔPF_{low} trials. The trial main-effect was non-significant ($p=0.788$), indicating that subjects' performance in the ΔPF field was not a function of the number of times they were exposed to it, i.e., learning was not consolidated across ΔPF sets. The

interaction between the force field strength and trial number was also non-significant (p=0.322). Pairwise comparisons were used to determine where significant differences between trial types occurred. We found no difference between the learnt-PF and ΔPF_{low} trials (p= 0.454), but significant differences between both learnt-PF and ΔPF_{high} trials (p=0.038) and ΔPF_{high} and ΔPF_{low} trials (p=0.012).

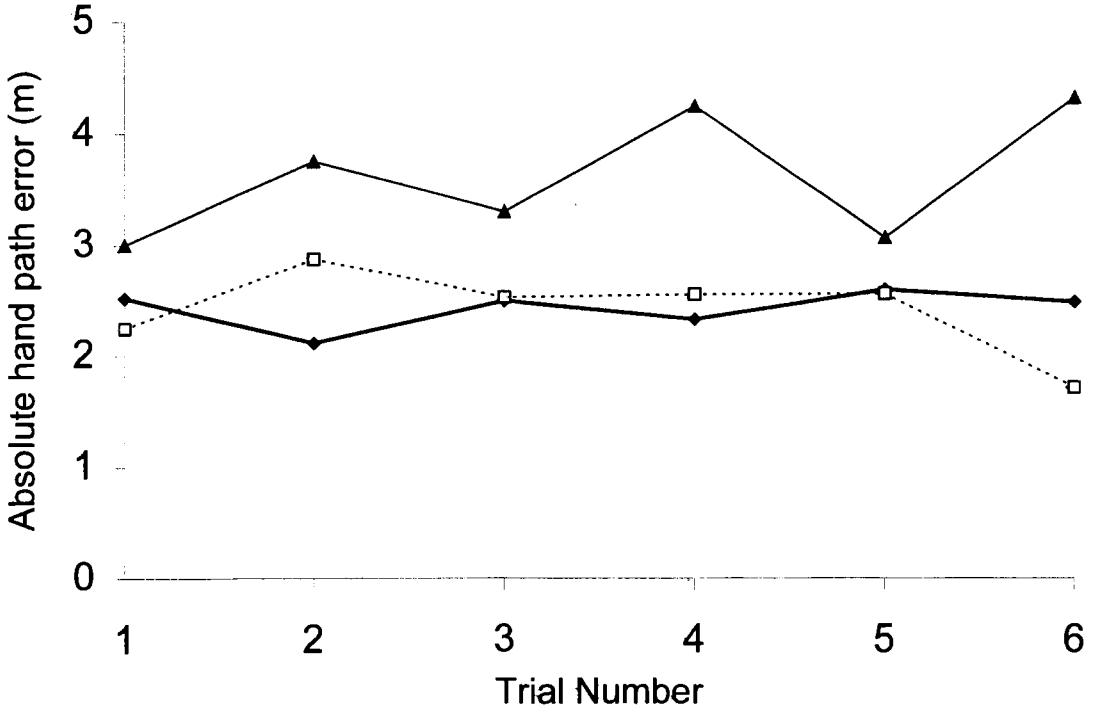

Figure 36. Absolute hand path error in learnt-PF and adapted ΔPF trials.
Absolute area in the 6 learnt-PF trials prior to any perturbations (◆, thick line), the 6 ΔPF_{low} trials in the 5th position (□, dotted line) and the 6 ΔPF_{high} trials in the 5th position (▲, thin line). Values are averaged across the 9 subjects.

After 5 trials in the ΔPF_{low} subjects produced hand paths that were as straight, relative to a straight-line path, as those following complete adaptation to the PF. Given that the instruction to subjects was to move as straight as possible, it appears that performance was as good in the 5th ΔPF_{low} as in the learnt-PF (following the 150 trial adaptation period). In contrast, the 5th trials in the ΔPF_{high} were still significantly less straight, relative to the straight line, than both learnt-PF and ΔPF_{low} trials.

Force on CT's within the ΔPF_{low} and ΔPF_{high} sets

As we previously noted, reducing trajectory error is only one aspect of learning. Learning to make similarly straight movements with less cocontraction, while forming a more accurate internal model of the task dynamics is an aspect of learning that cannot be determined simply by analysing hand paths. Analysis of force, both in terms of force impulse and temporal force profiles, would enable us to determine if subjects adapted the internal models that they had learnt in the PF, to more accurately compensate for the perturbed dynamics of the ΔPF fields. Even though minimal hand path straightening occurred following the second ΔPF trial, learning may have still continued by modifying the internal model while reducing muscular cocontraction. This can be investigated through analysis of force and EMG (see EMG analysis of PF learning). To determine whether subjects continued to learn in the ΔPF, even though hand paths did not straighten noticeably after the 2nd ΔPF trial, we considered the force profile in CT's that randomly replaced 25% of ΔPF trials in positions 2-5. As in the PF learning period, force was analysed by calculating the force impulse from 150 ms prior to movement onset until the PF boundary was reached at 0.1m, about 240 ms following movement onset The force profile up to this point should have been relatively unaffected by voluntary reactions even if subjects detected the channel. The force impulse provided a measure of the feedforward command issued on the premise that the ΔPF was expected. If an internal dynamics model was being continuously modified, we hypothesized that the force impulse in the CT's that replaced ΔPF_{low} trials in positions 2 to 5 would progressively decrease relative to the force recorded in CT's at the end of the PF learning period. In CT's that replaced the ΔPF_{high} trials in positions 2 to 5, we hypothesized that internal model learning would be characterized by progressive increases in the force to compensate for the higher strength force field. We also hypothesized that even though learning might not be evident in straightening of hand paths, the level of muscular cocontraction associated with the initial exposure to a new force field would drop, accompanied by changes in muscle activation to more specifically compensate for the external force.

Over the 9 subjects, the change in force impulse with exposure to the ΔPF_{low} and ΔPF_{high} was relatively symmetric (Figure 37). Subjects increased (in ΔPF_{high}) or decreased (in ΔPF_{low}) the force impulse on CT2, relative to the force impulse at the end of PF learning. The force impulse was further modified on CT3 and CT4, but a reversal in the adaptation was seen on CT5, i.e., the force impulse decreased in the ΔPF_{high} field and increased in the ΔPF_{low} field.

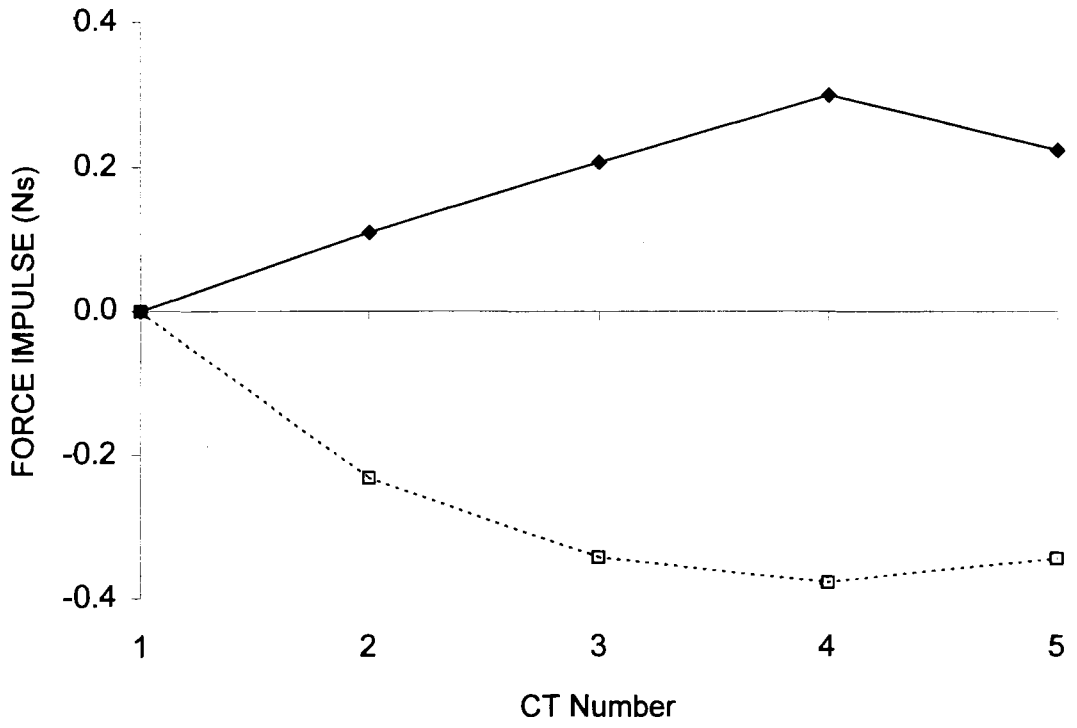

Figure 37. Force impulse recorded in CT's that replaced ΔPF trials in the 2nd-5th position
Force impulse in CT's that replaced ΔPF_{high} trials (◆, solid line) and ΔPF_{low} trials (□, dotted line) are expressed as differences relative to the force impulse recorded in CT's at the end of the PF learning, shown as CT 1. Data are averaged over the 3 CT's in each position and over the 9 subjects.

A one-way repeated measures ANOVA was conducted to determine if, over the 9 subjects, the adaptations were significant. In both the low and high strength fields there was a significant trial main-effect ($p=0.006$, $p=0.016$ for ΔPF_{high} and ΔPF_{low}, respectively), reflecting the fact that subjects progressively adapted to the perturbed fields.

Repeated contrasts revealed a significant difference between the force impulse on the CT's representing the learnt-PF and on CT2, for both the higher and lower strength force fields ($p=0.02$, $p=0.018$, respectively). Therefore, subjects had begun to adapt to the perturbed strength field, altering their IDM to produce higher or lower forces to compensate for the perturbed field strength following the feedback from a single trial in the ΔPF. Despite the trend evident in Figure 37, no significant difference was found between the force impulse on CT2 and CT3, CT3 and CT4 or CT4 and CT5 for either perturbation in force field strength ($p>0.05$). This is probably due to the variability in the force impulse across the 9 subjects.

We found that all three subjects who *did not* adapt their force significantly over the 4 CT's that we considered within the PF learning (see Figure 18) *did* appear to adapt their force in the both the high and low strength ΔPFs. Figure 38 shows force impulse data for subject, AR, who showed no modification of force impulse between CT's 1, 10, 20 and 30 in the PF learning. In contrast, AR, like all 9 subjects, appeared to adapt to the higher and lower strength force fields by modifying force impulse in the CT's that replaced ΔPF trials in the 2^{nd}-5^{th} position.

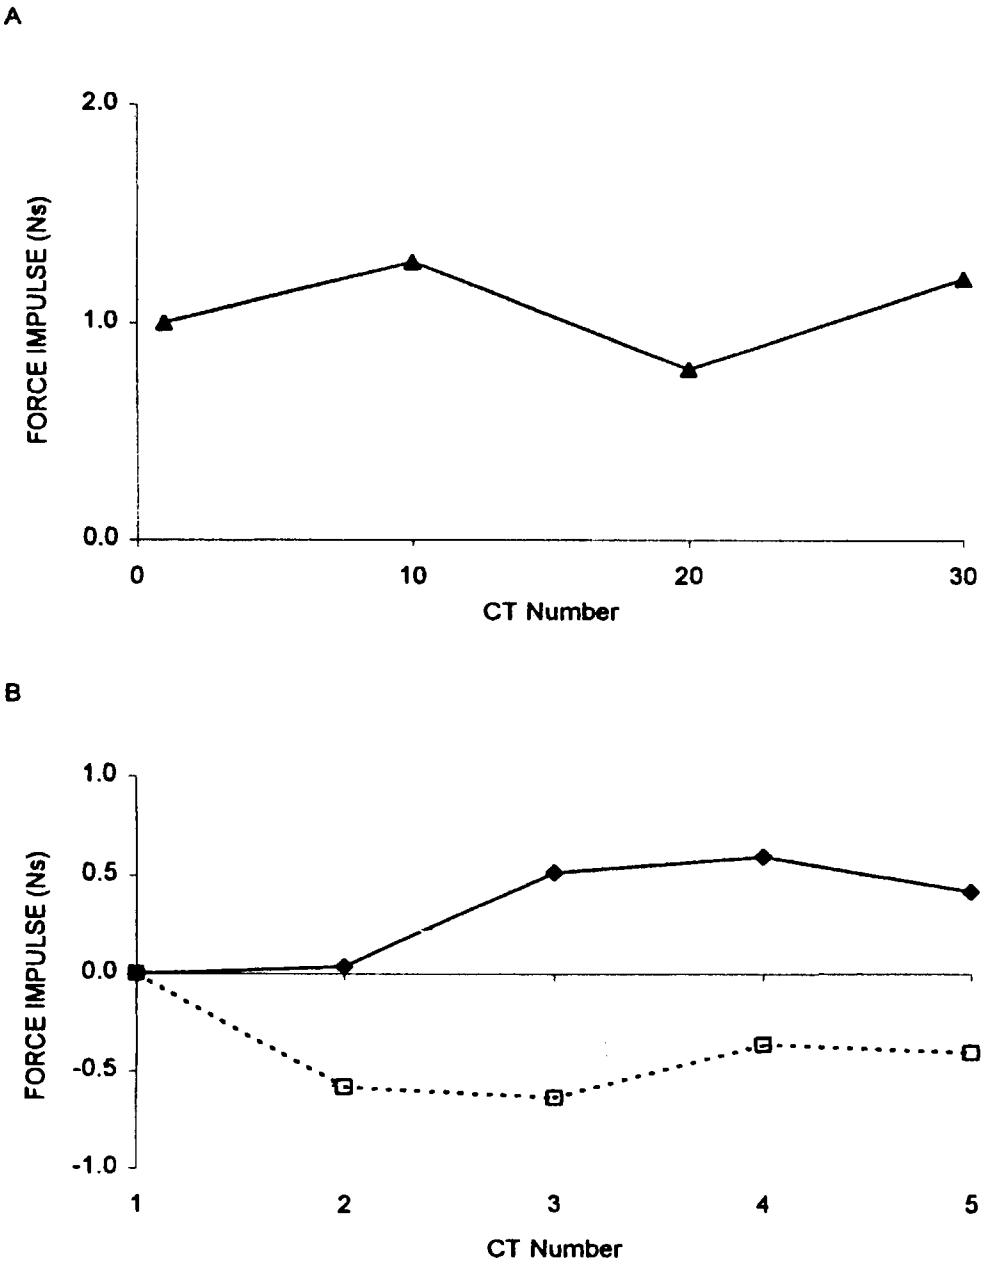

Figure 38. Force impulse recorded for subject AR in CT's that replaced PF and ΔPF trials

Force impulse in CT's that A) replaced PF trials within the PF learning period and B) replaced ΔPF trials in the 2nd-5th position for ΔPF_{high} (◆, solid line) and ΔPF_{low} (□, dotted line). 4 CT's (CT 1, 10, 20 and 30) were used to represent learning in the PF. Values in B) are expressed as differences relative to the force impulse recorded in CT's at the end of the PF learning, shown as CT 1, and are averaged over the 3 CT's in each position. Data is representative of the 3 subjects who were categorized as not adapting force in the PF learning.

Temporal force profiles in ΔPF trials

As well as considering how the force impulse changed as subjects were exposed to more ΔPF trials, we also considered the temporal profile of that force following the 5th ΔPF trial. This would give an indication of the type of the characteristics of the internal model being formed during adaptation by revealing how closely the force profile subjects produced resembled the force profile required to compensate for the ΔPF. Given that subjects were unable to produce straight hand paths after extensive exposure to the PF we decided that it was not appropriate to compare the force profiles produced in the ΔPF to a profile representing perfect compensation to the ΔPF through internal model formation. (Perfect compensation force is defined as the forces required to compensate exactly for the external force field, resulting in a perfectly straight hand path). Instead, we compared the force profile which subjects produced in the ΔPF to the force profile required to move on a path, identical to the learnt-PF path, but now in the perturbed strength field. Our hypothesis was that subjects would aim to recreate the PF hand paths in the ΔPF, rather than aim to achieve a straight line path. To determine whether this hypothesis had merit we first checked that the variability of the learnt-PF movements was not greater than that of the learnt-NF movements. We assumed that subjects' trajectories were stable from trial to trial in the NF. Showing that subjects' trajectories were no more variable at the end of learning in the PF than learnt-NF trajectories would indicate that they were equally stable. The relative variability in the two fields (for the 7 subjects for whom we had NF data) was determined within the last 5 learning trials in each field. The absolute area between the hand path of each of these 5 trials and their mean hand path was calculated for each force field. The standard deviation was used as a measure of the variability of the hand paths. A Student's two-tailed paired t-test (*a priori* α=0.05) revealed that there was no significant difference in the variability between the two fields (p=0.47), supporting the merit of the hypothesis.

The force required in the ΔPF to move along the learnt-PF path was determined by taking the force profile in the final 5 CT's in PF learning and adding or subtracting the change in magnitude of the parabolic force due to the change in field strength. This comparison was

carried out for the first 0.1 m of the reach, where the PF was active (Figures 39 and 40). In this manner, we considered primarily force produced by feedforward commands, without including voluntary reactions to detection of the channel. Like the analysis undertaken in the PF, this analysis enabled us to examine how closely the hand force matched the predicted force at specific locations, rather than simply comparing force impulses.

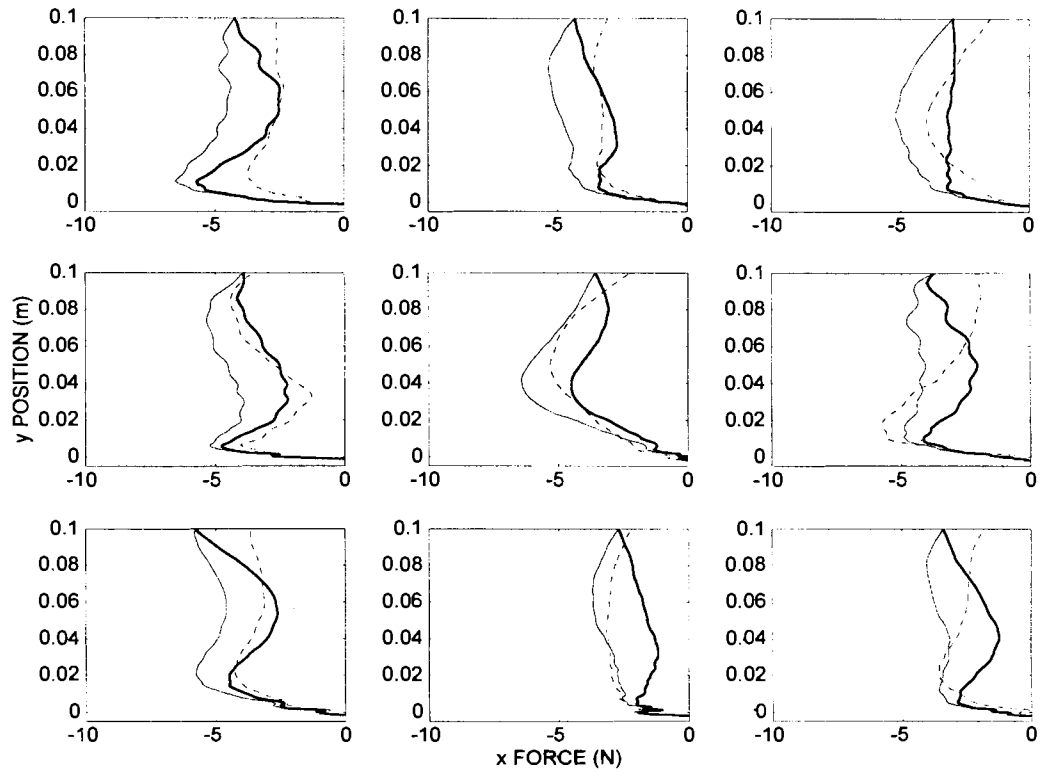

Figure 39. Force profiles in the ΔPF_{low}.
The thin solid line represents the force in CT's at the end of PF learning. The thick solid line represents the predicted force required to move along a path similar to the learnt-PF path in the ΔPF_{low}. The dashed line represents the force recorded in the CT's that replaced ΔPF_{low} trials in position 5. Data is averaged over the 3 CT's in the ΔPF field and over the last 5 CT's during PF learning. The 9 individual subjects are shown in the 9 panels (Top right- bottom right: MH, US, AN, JJ, JJ, SM, DF, RO, AR).

Adaptation to the ΔPF_{low} was not consistent among subjects (Figure 39). The initial rapid increase in force, seen in the PF, was less obvious in the ΔPF. Forces generally increased

to levels above the force required to move on a path similar to the PF path, resulting in trajectories that were displaced to the left of the learnt-PF paths. Although force profile shapes differed considerably among subjects, for 5 of the subjects (MH, US, AN, JJ and DF) the ΔPF_{low} force profile resembled a scaled down version of the PF force profile (compare the dashed and thin lines in Figure 39). There were only very small differences between force profiles for CT2 (not shown) and CT5. It appears that most of the adaptation to the change in field strength occurred after a single exposure to the perturbed field.

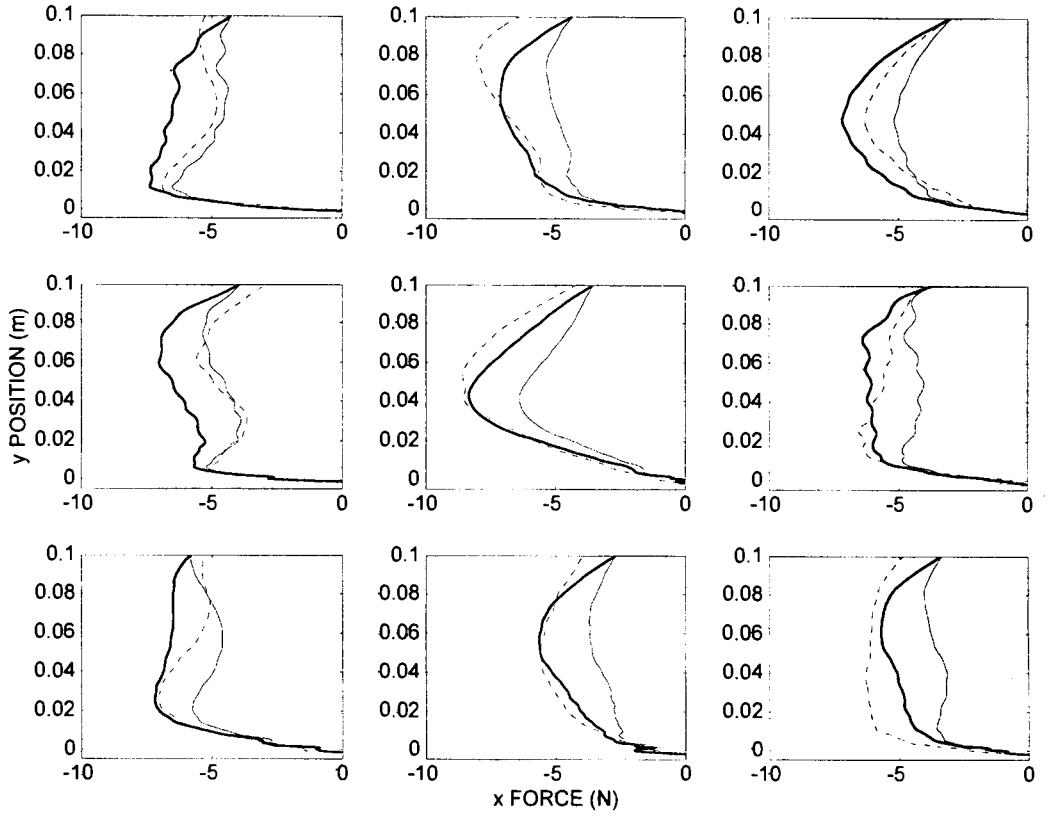

Figure 40. Force profiles in the ΔPF_{high}.
The thin solid line represents the force in CT's at the end of PF learning. The thick solid line represents the predicted force required to move along a path similar to the learnt–PF path in the ΔPF_{high}. The dashed line represents the force recorded in the CT's that replaced ΔPF_{high} trials in position 5. Data is averaged over the 3 CT's in the ΔPF field and over the last 5 CT's during PF learning. The 9 individual subjects are shown in the 9 panels (Top right- bottom right: MH, US, AN, JJ, JJ, SM, DF, RO, AR).

The force profiles in the ΔPF_{high} resembled a scaled up version of the PF force profile (compare the thin and dashed lines in Figure 40) for 8 of the 9 subjects. Again, we found very little change in the force profile between CT2 and CT5. Generally subjects produced less force than that required to move along their learnt-PF path, resulting in ΔPF_{high} paths that were displaced to the right of the PF path.

Joint torque analysis in the ΔPF

Similarly to the PF analysis, we analyzed joint torques to determine the manner in which subjects adapted to the perturbed strength fields. Joint torques required to produce hand paths in the ΔPF similar to those in the PF, were computed for both of the perturbed strength force fields (Figure 41). We note that in the ΔPF, the joint torques required to compensate for the force field are scaled, relative to the torques required to compensate for the PF, by a factor equal to the change in field strength, i.e., in the ΔPF_{low} 50% less torque is required compared to the PF, whereas in the ΔPF_{high} 50% more torque is required. Predicted peak shoulder flexor torque is approximately 2.2 and 3.7 Nm for ΔPF_{low} and ΔPF_{high}, respectively In contrast, the peak elbow extension torque did not vary significantly between the PF, ΔPF_{low} and ΔPF_{high}, indicating that it was required principally to accelerate the arm at the onset of movement where the magnitude of the force field was low. However, the required rate of change of torque following peak extension torque, where the force field reached it greatest magnitude, was slower (for the ΔPF_{low}), or faster (for the ΔPF_{high}), than that recorded in the PF.

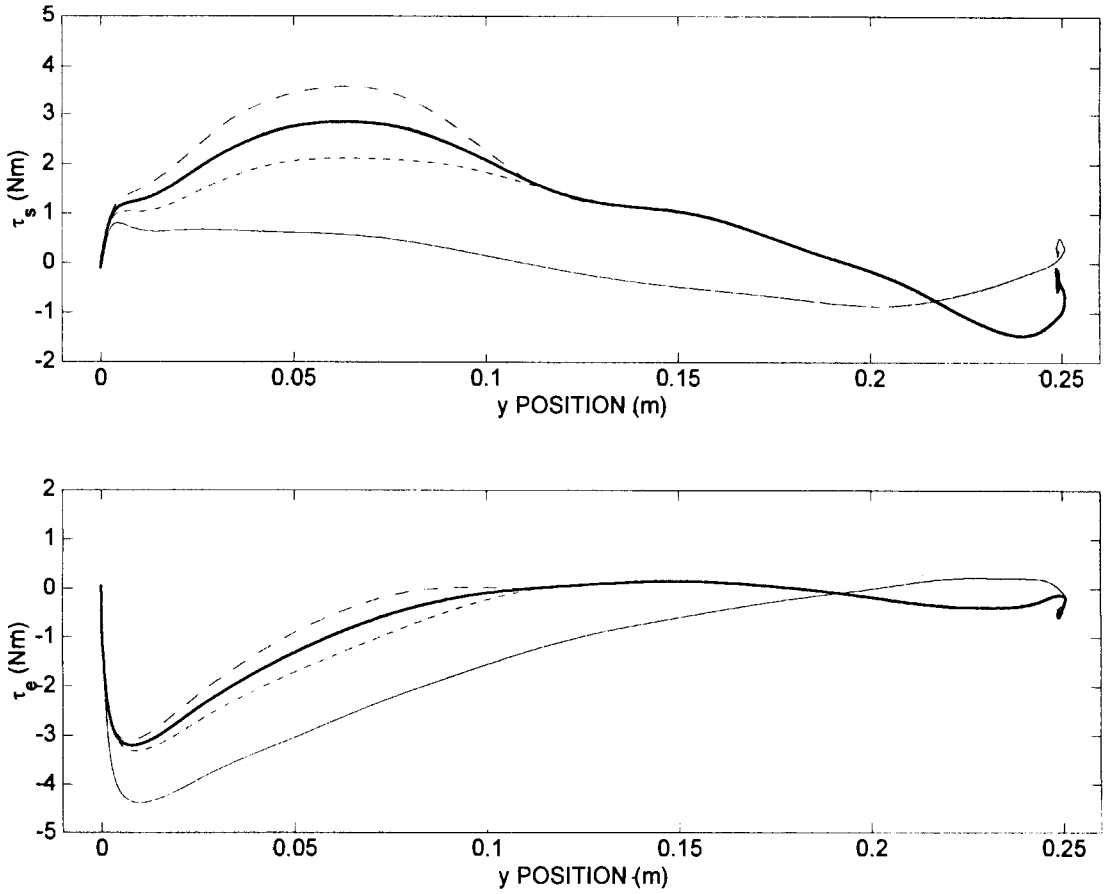

Figure 41. Shoulder (τ_s) and elbow (τ_e) joint torque calculations for the ΔPF_{low} and ΔPF_{high}.
Solid lines show the actual torques measured in the NF (thin line) and learnt-PF (thick line). Joint torque required to move on a path similar to the learnt-PF path in the ΔPF_{low} and ΔPF_{high} are shown as the dotted and dashed lines, respectively. Data are shown for US, and are representative of all subjects.

To determine the extent to which subjects were able to match the predicted joint torque profiles following exposure to ΔPF trials, we compared the computed and recorded differences in torque, with respect to the learnt-PF torque, as a function of position for the first 0.1 m of the movement, representing the region in which the PF was applied (Figure 42). As previously described, subjects moved the 0.1 m in approximately 240 ms. Thus, torques over this region represent predominantly feedforward commands. To reduce the level of variability in the analysis we combined all ΔPF trials in position 3, 4 and 5 to represent the motor activity following learning in the ΔPF. The validity of this procedure is based on there being no change in the absolute hand path error or force impulse

between ΔPF trials in positions 3, 4 and 5. ΔPF trials that followed CT's were not included in the analysis.

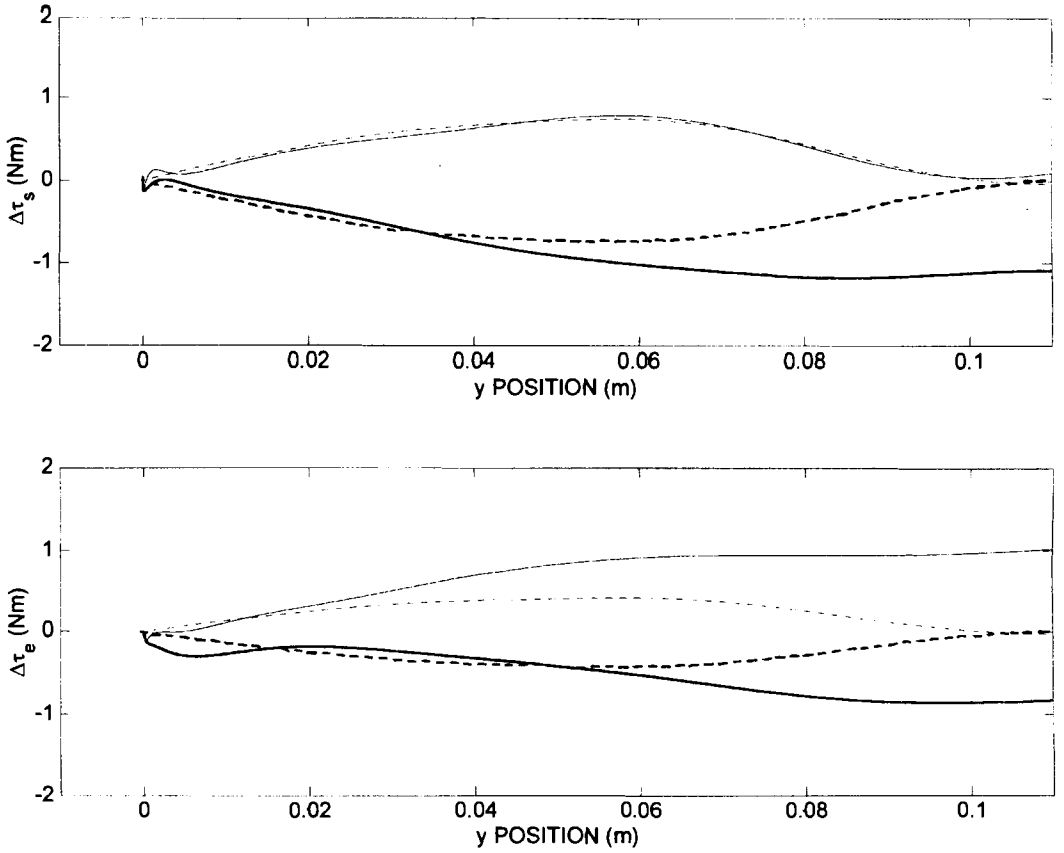

Figure 42. Differences in shoulder ($\Delta\tau_s$) and elbow ($\Delta\tau_e$) joint torques in the ΔPF fields, with respect to the actual PF torque, as a function of position.
Computed values are shown for ΔPF_{low} and ΔPF_{high} as thick and thin dotted traces, respectively. Actual torques are shown for ΔPF_{low} and ΔPF_{high} as thick and thin solid lines, respectively. Data are shown for the region where the PF field was active, for subject RO, and are representative of all subjects.

Following learning in the ΔPF_{low}, subjects were generally good at matching the required shoulder flexion torque for the initial part of the movement, although subjects began to reduce shoulder flexion torque too quickly, resulting in a deficit of flexion torque, compared to the computed torque profiles. Peak elbow extension torque was well matched, but was often followed by torque reductions which were too slow, resulting in net elbow extension, relative to the computed values.

In the ΔPF_{high}, subjects matched the computed shoulder torque very well (compare the thin dashed line and thin solid line in the top panel of Figure 42) throughout the period while the PF was active. Elbow extensor torque was also generally well matched in the ΔPF_{high} for the initial part of the movement. Following peak elbow extension, elbow extension torque was reduced too quickly compared to the computed values and often reversed to elbow flexion torque at amplitudes greater than 0.06 m.

Although it is useful to compare the recorded joint torques to computed values, this technique is limited in assuming that adaptation to the perturbed strength force fields will occur through updating an internal model. Any increase in stiffness due to co-activation of antagonistic muscle pairs will limit the value of such comparisons. We note that, compared to the PF, the ΔPF_{low} and ΔPF_{high} require less, and more shoulder flexion torque, respectively. A lower or higher rate of change of elbow torque is also predicted, in the ΔPF_{low} and ΔPF_{high}, respectively (Figure 41). By analyzing muscle activity it might be possible to determine whether adaptation to the perturbed strength force fields was the result of alteration of muscle activity that would result in the computed torque changes, i.e., whether subjects adapted their internal models, or whether there was a generalized increase in muscle activation, i.e., subjects used cocontraction in adapting to the perturbed strength force fields.

EMG analysis of ΔPF learning

Kinematic analysis of hand paths had revealed that little adaptation of ΔPF hand paths had occurred following the 2nd ΔPF trial. This was further highlighted by temporal force profiles that did not exhibit significant change in shape of amplitude following the 2nd ΔPF trial. A final aspect of learning that we must consider is how the motor commands may have changed as subjects were exposed to the 5 consecutive ΔPF trials.

Force analysis suggests that the significant straightening of hand paths between the initial ΔPF exposure and the 2nd ΔPF occurred through adaptation of the PF internal model, as we noted an increase (in ΔPF_{high}) or decrease (in ΔPF_{low}) in the force recorded on the

CT's, although progressive changes in the force were not found between the 2^{nd}-5^{th} ΔPF trials. Furthermore, hand paths did not become progressively straighter as subjects were exposed to the 5 ΔPF trials. A reduction in the level of cocontraction, noticeable as a reduction in EMG of both muscles of an antagonistic pair without any associated worsening of performance in terms of hand path improvement, may have occurred. This is an important aspect of motor learning that would enable the task to be undertaken more efficiently, with less metabolic cost.

To determine how subjects modified their feedforward motor command during exposure to the perturbed strength force fields, both in the period of ΔPF learning where we noted improvements in hand path (from the initial exposure to the 2^{nd} ΔPF trial), and in subsequent trials where no improvement was noted (from the 2^{nd}-5^{th} ΔPF trials), we considered rmsEMG in CT's that randomly replaced ΔPF trials. In this manner, we were able to investigate the feedforward motor command in isolation from any stretch reflex activity caused by the perturbation relative to the learnt path. Like the EMG analysis for the PF learning, we calculated the rmsEMG, but in this case we only examined early rmsEMG. EMG later in the movement may have been affected by voluntary responses if subjects detected the channel.

A two-way RM ANOVA was used to detect any differences in rmsEMG as subjects adapted to the ΔPF. We considered the final three PF trials prior to exposure to any ΔPF trials to represent the feedforward EMG in the PF. We hypothesized that since subjects had adapted completely to the PF, the EMG of the PF trials themselves would constitute an adequate representation of the feedforward command. By considering three consecutive PF trials, rather than the final 3 CT's within the PF (which were the 140^{th}, 145^{th} and 150^{th} trials in the PF learning regimen and hence were non-consecutive), we felt that a more consistent measure of the EMG would be obtained. In fact, comparison of EMG traces of the final 3 CT's and final 3 PF trials during learning were not noticeably different, providing evidence that stretch reflex responses or online corrections were essentially absent in the PF trials. We compared the 3 PF trials to the 3 CT's that replaced

ΔPF trials in positions 2, 3, 4 and 5 of ΔPF sets to investigate any progressive changes in the feedforward command with exposure to the ΔPF. Separate ANOVA's were used for each of the 6 muscles and field strengths.

In the ΔPF_{low}, we found a non-significant trial position main effect for 5 of the 6 muscles, i.e., rmsEMG did not change as subjects adapted to the ΔPF. Activity of the posterior deltoid (shoulder extensor) was the only muscle to show a drop in activity ($p=0.004$). Although pectoralis major rmsEMG did not show significance at the *a priori* significance level, the p value of 0.074 suggests that was greater than 90% probability that there was less activity in this shoulder flexor as subjects adapted to the ΔPF_{low}. Given that these two aforementioned muscles work as an antagonistic pair, this result suggests that there may have been a reduction in the cocontraction at the shoulder as subjects adapted to the ΔPF_{low}. The trial number main effect was non-significant for all 6 muscles ($p>0.2$ in all cases), verifying that performance was not affected by the number of times subjects were exposed to the ΔPF. The interaction was also highly non-significant for all muscles ($p>0.4$ in all cases). In the ΔPF_{high}, we found a non-significant trial position main effect, trial number main effect, and interaction, for all 6 muscles ($p>0.1$ for all comparisons)

We previously showed that subjects reach the PF boundary at around 240-270 ms after movement onset. Therefore, changes in EMG to compensate for the change in field strength may continue beyond 150 ms. To determine if changes in rmsEMG reached statistical significance when we considered a longer time interval, we analysed the rmsEMG from 150 ms prior to movement onset until 240 ms following onset. Although this cannot be assumed to be solely feedforward EMG, it provides a more accurate representation of the adaptation, since we consider the whole period in which the force fields are active. The analysis was undertaken on the same trials as in the previous analysis. Again, in the lower strength force field we found a similar drop in the posterior deltoid rmsEMG ($p=0.04$), while pectoralis major again approached significance for a reduction in activity ($p=0.071$), suggesting a possible decrease in the cocontraction at the shoulder. In the higher strength force field all main effects and interactions remained non-

significant ($p>0.1$), at a level similar to that in the analysis undertaken up to 150 ms. We noted a trend for a decrease in posterior deltoid ($p=0.11$), suggesting that subjects may have been increasing shoulder flexion torque in the period while the ΔPF_{high} was active, as would be predicted from the torque analysis.

In summary, this particular analysis of rms EMG has provided few clues as to how the motor command may have changed as subjects adapted to the ΔPF. The analysis did reveal that there was a tendency for a decrease in cocontraction at the shoulder as subjects adapted to the ΔPF_{low}, although we did not find a corresponding decrease in cocontraction for the ΔPF_{high}. The non-significant trial number main effect in the ANOVA suggests that there was no systematic change in the motor command as subjects were exposed to the ΔPF set more times. This result justifies averaging the EMG over the repeated exposures in subsequent analyses, to produce a less variable, more reliable estimate of motor activity.

Change in feedforward command between PF and the second ΔPF trials

To investigate possible changes in EMG that occur both early, i.e., up to 150 ms after movement onset, and later, i.e., from 150 ms after onset until the end of data collection, between trials, we considered PF trials and ΔPF trials, rather than CT's. PF trials were used, as opposed to CT's as we wanted to analyze EMG throughout the trial. EMG late in CT's may be prone to online compensation, as subjects may have sensed the presence of the channel. In contrast, the learnt-PF trials should consist of only feedforward motor commands for the duration of the trial. Based on the theory of error feedback learning, we would expect a change in the feedforward command on the 2nd ΔPF trial as a result of updating the command based on feedback from the 1st ΔPF trial. We investigated this by comparing the rmsEMG in the 2nd ΔPF trials and the corresponding pre-perturbation PF trial, for perturbations to lower and higher strength force fields using a two-way RM ANOVA. We considered both the early rmsEMG, as a measure of the feedforward command, and late rmsEMG, as previously described, in the 9 pairs of trials (the 9 ΔPF trials in position 2 that were not CT's, and the 9 preceding PF trials) for all subjects and

both field strengths. Since we found little adaptation to the ΔPF between ΔPF trials 2-5, we focused on the 2nd ΔPF trials and did not compare further ΔPF trials. Surprisingly, we found no significant differences in either early or late rmsEMG values for all 6 muscles, between pre-perturbation PF trials and 2nd ΔPF_{low} trials ($p>0.1$ in all cases). We did note that there appeared to be a trend for a decrease in early rmsEMG for both the pectoralis major and biceps muscles. This suggests that subjects may have realised the requirement for less shoulder flexion in the ΔPF_{low}, compared to the PF, although this did not reach significance at the *a priori* level. Similarly, no significant differences in late rmsEMG were found between the PF and 2nd ΔPF_{low} trials ($p>0.1$). When we considered the higher strength perturbations, we found a significant increase in early rmsEMG for the pectoralis major muscle in 2nd ΔPF_{low} trials, compared to the corresponding PF trials ($p=0.019$), suggesting subjects adapted to the field by increasing net shoulder flexion torque. No significant differences were found in early rmsEMG for the other 5 muscles ($p>0.1$). Significant increases were found in late rmsEMG for pectoralis major, anterior deltoid, biceps, triceps long head and brachioradialis ($p=0.011$, 0.015, 0.03, 0.009 and 0.0006 respectively) in the 2nd ΔPF_{high} trials compared to the corresponding preceding PF trials. Only the activity of the lateral head of the triceps remained non-significant, although we did note a trend for increased activity in this muscle as well.

The preceding EMG analysis was inconclusive in determining whether there was a significant change in muscular activity between the PF and ΔPF trials. A possible reason for this is the limited number of ΔPF trials in each position over which we could average (also see Experimental Limitations for a power analysis). To reduce the level of variability in the analysis we combined all ΔPF trials in position 3, 4 and 5 to represent the motor activity following learning in the ΔPF, similar to the procedure used to analyze torque. EMG was normalized for each subject to the learnt-PF activity to allow for better inter-subject analysis.

Early rmsEMG was calculated over an interval of 300 ms, beginning 150 ms prior to movement onset, representing the feedforward command. Late rmsEMG was calculated

in the interval from 150 ms following onset until the end of data collection. Students t-tests (2 tailed, *a priori* α=0.05) were undertaken to compare rmsEMG in learnt-PF and ΔPF_{low} and learnt-PF and ΔPF_{high} trials.

Significantly less early rmsEMG was found in ΔPF_{low} trials, compared to learnt-PF trials, in the posterior deltoid and brachioradialis muscles ($p<0.01$, $p<0.05$, for posterior deltoid and brachioradialis, respectively) (Figure 43). Given that subjects produce shoulder flexor torque and elbow extensor torque while moving in the ΔPF_{low} field, reducing the activity of posterior deltoid would increase net shoulder flexor torque while reducing the activity of brachioradialis would increase net elbow extensor torque. The increase in net elbow extensor torque resulting from a decrease in brachioradialis is consistent with the torque changes seen, with respect to the learnt-PF, in Figure 42. In contrast, the net increase in shoulder flexor torque that would occur due to a decrease in posterior deltoid activity is not apparent in Figure 42. To the contrary, Figure 42 shows subjects *reduce* their shoulder flexor torque compared to that in the learnt-PF, suggesting that reduced posterior deltoid activity can only represent a reduction in co-contraction. The activity of pectoralis major, or other shoulder flexors, must also be reduced, by a greater amount, such that the net decrease in shoulder flexor torque seen in Figure 42 occurs. Indeed, the average change in magnitude of the rmsEMG for pectoralis major was larger than the change in magnitude of rmsEMG for posterior deltoid, although a larger variability preventing a statistically significant p value. All other muscles showed no significant change in activity ($p>0.1$), although we note that there was also a trend for a decrease in the activity of the biceps.

Torque analysis had revealed that subjects reduced their shoulder flexion torque in the ΔPF_{low}, relative to the learnt-PF torque (Figure 42), although we do not find conclusive evidence to support this when analyzing muscle activity, i.e., we do not see any significant reduction in the activity of pectoralis major or biceps. It appears that the rmsEMG was not sensitive enough to reveal a significant reduction in muscle activity, which probably occurred in these muscles. Late rmsEMG in the ΔPF_{low} was also lower

than in learnt-PF trials for the posterior deltoid ($p<0.01$) and brachioradialis ($p<0.05$). We note that individual subjects did not all show the same responses, with some subjects lowering activity in biceps and pectoralis major more than others.

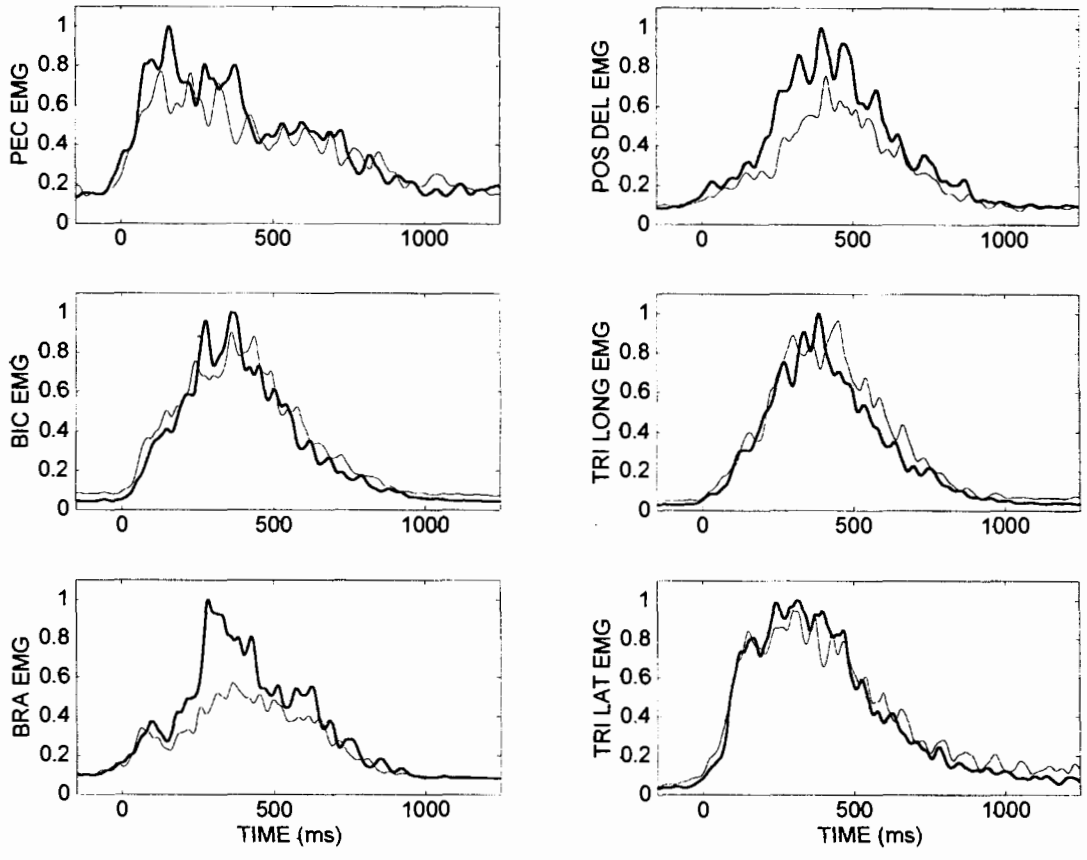

Figure 43. Comparison of EMG in the learnt-PF (thick lines) and ΔPF_{low} (thin lines) trials.
Data are averaged over 18 trials for subject JH. Decreases in rmsEMG (averaged across all subjects) were found in the ΔPF_{low} trials compared to learnt-PF trials for posterior deltoid and brachioradialis for both the early and late portions of the trial.

When comparing the early rmsEMG between learnt-PF and ΔPF_{high} trials, we found a significant increase in the rmsEMG of the biceps ($p<0.02$) (Figure 44). Pectoralis major showed a trend for an increase in activity, although this was not significant ($p>0.1$). All differences in the early rmsEMG of the other muscles remained highly non-significant ($p>0.30$). This would produce an increase in shoulder flexor torque and a reduction in elbow extensor torque. Torque analysis revealed that an increase in shoulder flexion torque and a reduction in elbow extensor torque, relative to that required in the PF, was

needed to compensate for the ΔPF_{high}, and was observed. The increased activity of the biceps muscle would achieve both of these actions. This provides evidence that subjects amended their internal models to more accurately compensate for the novel field. In the late portion of the movements, we found significant increases in the activity of all muscles, except for the posterior deltoid. ($p<0.05$, $p<0.02$, $p<0.001$, $p<0.01$ and $p<0.05$, for the pectoralis major, biceps, long head of the triceps, brachioradialis and lateral head of the triceps, respectively), although posterior deltoid did also show a trend for an increase in activity. This suggests that subjects used generalized cocontraction in the latter part of the movement suggests to increase the stiffness of the arm. It appears that subjects update the internal model to compensate for the ΔPF_{high}, but also employ impedance control during later periods of the movement to reduce the perturbing effect inaccuracies in the feedforward model.

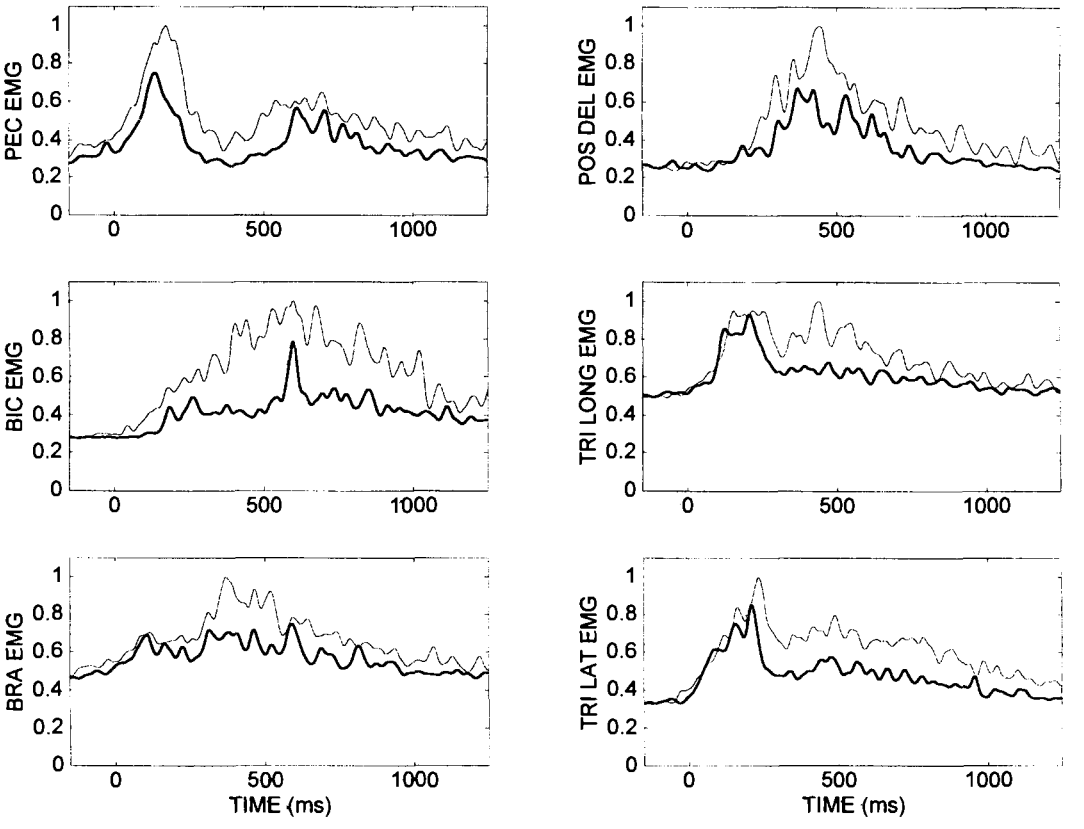

Figure 44. Comparison of EMG in the learnt-PF (thick lines) and ΔPF_{high} (thin lines) trials.
Data are averaged over 18 trials for subject SM. Increases in rmsEMG (averaged across all subjects) were found in ΔPF_{high} trials compared to learnt-PF for the biceps in the early portion of the trial, and for the all muscles except posterior deltoid in the late portion of the trial.

Discussion

We hypothesized that stretch reflex EMG, resulting from trajectory error, would provide an indirect estimate of the afferent error signal available to the brain. However, we were unable to identify any stretch reflex in surface EMG recordings. We found no evidence of stretch reflex activity in the first ΔPF trial following PF trials, and in PF trials following ΔPF sets, i.e., there appeared to be no change in the EMG signals in the first ΔPF trial, relative to the preceding PF EMG, or in the PF trial relative to the preceding ΔPF EMG. We note that variability in the EMG signals may have prevented identification of any stretch reflex activity, although we measured the rate of joint perturbation and found that these "perturbation velocities" were very low, compared to those previously used to evoke sizeable reflex responses (Nakazawa et al, 2001). Although the magnitudes of the trajectory perturbations were sizeable, the rate at which these perturbations occurred was insufficient to evoke reflexes. In the absence of stretch reflex, we suggest that subjects used both proprioceptive and visual feedback to update their feedforward commands on subsequent trials.

PF learning

Analysis revealed that even following extensive learning in the PF, movement trajectories exhibited significant bowing, first to the left and then back towards the centreline. Deviations within the first 0.1 m, i.e., while the PF was active, were of similar magnitude following learning to that in the 2nd or 3rd exposure to the PF. In contrast, trajectory error in the final 0.15 m of movements was significantly reduced throughout the learning regimen. Path straightness in the early portion of the reach may have been less important to subjects than a straight approach to the target. Alternatively, the amount of displacement may not have been great enough for subjects to perceive as deviation from a straight-line path, or path straightness may have been constrained by intrinsic muscle properties, i.e., subjects were unable to produce the required muscle activation patterns that would compensate accurately for the PF.

Complexity of task may limit adaptation

Evidence to suggest that subjects were unable to produce the required muscle activation to compensate for the PF was found in subjects' force profiles, as measured in isolation from PF force in CT's. We found that subjects' force rose quickly, initially overcompensating for the PF, before PF force became larger. At the PF boundary subjects' force remained above zero, again resulting in overcompensation for the field. This suggests it is difficult for subjects to temporally match the PF force profile. The initial overcompensation for the PF force shows subjects were able to generate force quickly enough to compensate for the force field since they increased lateral force more rapidly than the force field. Previous work has shown that subjects could increase force fast enough, and to high enough levels, to allow for adequate compensation for the PF. Ghez and Gordon (1987) investigated how subjects control force in goal-oriented isometric contractions of the elbow muscles to reach a target force. Subjects were asked to increase force at various rates whilst maintaining accuracy. Subjects utilized cocontraction of antagonist muscles when generating forces that increased quickly, to prevent overshoot of the target force, enabling them to limit the maximum force output by counteracting the torque produced by agonist muscles without limiting the rate of increase in the force. Forces could be increased at a maximum rate of approximately 2500 Ns^{-1}, reaching maximum values of around 100 N when cocontraction was employed and 140 N without cocontraction. The rates of increase and force levels which subjects could achieve in Ghez and Gordon's (1987) experiment were an order of magnitude greater than those required in the present study, where maximum rate of change of force was approximately 100-200 Ns^{-1} and maximum forces were 6N for the PF, rising to 8N for the ΔPF_{high}.

We note, however, that there are competing requirements on certain muscles in the present study. In the Ghez and Gordon (1987) study, subjects performed isometric contractions, i.e., there was no movement. In our study, subjects needed to accelerate their arm in the target direction while simultaneously compensating for the force field. It may be difficult to produce muscle activity such that movement occurs at the desired

speed and in the desired direction while simultaneously compensating for the force field. This may be one of the reasons why we found that cocontraction was used in conjunction with subjects forming internal models.

We calculated the required torque profiles to move along an NF path while exposed to the PF (Figure 23). In the case of shoulder torque, we note an inflection in the profile, requiring an increase in the rate of shoulder flexor torque, approximately 100 ms following movement onset. In the case of elbow torque, there is an inflection in the profile slightly after 200 ms, requiring extensor torque to stop decreasing and increase slightly, for a short period of time. We suggest that these torque inflections may be difficult for the neuromuscular system to achieve. Furthermore, shoulder torque must decrease faster than it increased, while elbow torque must decrease more slowly than it increased to accurately compensate for the PF (Figures 23 and 26), i.e., the torque profiles for the two joints are oppositely skewed. Although muscles are known to generate asymmetric force profiles with much shorter (faster) rise times than fall times, different temporal torque patterns can be achieved using an antagonistic muscle pair. Brown and Cooke (1990) trained subjects to produce horizontal forearm movement velocity profiles that were symmetric, positively skewed such that the velocity increased faster than it decreased, or negatively skewed such that velocity decreased faster than it increased. Subjects were able to reproduce profiles of varying skew (symmetry ratio 0.4-2.0) by altering the characteristics of the phasic muscular bursts within an antagonistic pair of muscles. This proves that asymmetries in velocity, and therefore, presumably, in acceleration and torque profiles can be achieved using pairs of antagonistic muscles. We suggest that the difficulty in perfectly compensating for the PF is due to the required shoulder and elbow joint torque profiles being oppositely skewed. The biceps and long head of the triceps are biarticular muscles, responsible for generating torque at both joints. Given that a proportion of the total joint torque at the elbow and shoulder will be due to the activity of these muscles, oppositely skewed profiles would be hard to achieve. Interestingly, actual torque profiles following adaptation to the PF show some evidence of a compromise solution to this problem. Shoulder torque decreases too slowly and

elbow torque too quickly, compared to the computed values (Figure 26, lower two panels).

Force profiles in CT's were affected by arm and PFM inertia

In our initial analysis of channel trials, we estimated how much force the channel constraint introduced by preventing acceleration of the arm. The peak force was generally found to be 1-2.5 N, and occurred shortly after movement onset. Inspection of the force profiles measured in CT's, such as those seen in Figure 20, suggest that the force that subjects produced was higher than the PF force for the initial 0.03-0.04 m of the movement. This led to the conclusion that subjects were overcompensating for the PF in the initial region of the force field. Of course, this is true, as we noted that all subjects initially moved laterally leftward following PF learning. Without a net force to the left, i.e., overcompensation for the PF force, there could not have been any leftward motion. The difference between the internal model force i.e., the force that subjects produced to compensate explicitly for the PF, and the actual PF force, was probably less than suggested in Figure 20. Had we subtracted the force that would otherwise have been used to accelerate the arm, the initial difference between the subject's force and the PF force in Figure 20, would have been reduced. Given that the force needed to accelerate the arm rose to its maximum value shortly after movement onset, before falling and remaining close to zero for the remainder of the period of interest, i.e., until 0.1 m amplitude was reached, we suggest that this force produced an overestimate of the compensatory force which subjects produced during the initial portion of the CT force profiles, but would have had little effect on the estimate of compensatory force over the remaining region of the force field.

Simultaneous use of internal models and cocontraction

The bowing of trajectories in the early part of movements, characterized by lateral movements "into" the field, suggests that subjects were able to anticipate and compensate for the perturbing effects of the PF as early as the 2nd or 3rd trial. This finding implies that subjects produced an initial lateral force large enough to overcome the PF, enabling the

inertial mass of the arm and PFM to be accelerated laterally, and is conclusive evidence that as early as the 2nd or 3rd trial subjects had begun to develop an internal dynamics model to compensate for the external force. We suggest that the initial overcompensation may have aided in reducing subsequent trajectory error later during the movements.

Computed torque profiles revealed that shoulder flexion and elbow extension were required to move in the PF, along a path similar to NF paths. We noted activity in all muscles groups, contributing to both shoulder and elbow flexion and extension, suggesting that subjects employed a certain amount of cocontraction, in conjunction with the internal model, while moving in the PF. Progressive learning of the PF resulted in a decrease in shoulder flexion torque and a decrease in elbow flexion torque, 200-400 ms following movement onset (Figure 25). Linear regression analysis of rmsEMG during learning suggested that the activation of *all* muscle groups was progressively reduced in the "late" period of the movement, i.e., after 150 ms, as subjects adapted to the force field. Decreases in activation of muscle groups that did not contribute to changes in the torque profiles suggest that the level of cocontraction was progressively reduced during learning. We found no associated worsening of performance as the late rmsEMG dropped, suggesting that subjects were aiming to reduce cocontraction to an optimal level, i.e., they learnt to move in a more energy efficient manner without any effect on performance. We suggest that as subjects' internal model became more accurate in compensating for the PF, there was less need to employ high levels of cocontraction to increase the stiffness of the arm during the later portions of the movement, i.e., subjects utilized both internal model formation and impedance control (cocontraction) as a means of adapting to the PF. This result replicates previous studies that suggest internal model formation and impedance control as a dual strategy in compensating for external loads (Takahashi et al., 2001; Franklin et al., 2002a, b; Osu et al., 2002). It appears, as Osu et al. (2002) suggest, that cocontraction is initially high, when the internal model is inaccurate, but is reduced as feedforward commands more accurately compensate for the external dynamics. Further evidence to suggest formation of an internal model, rather than a simple strategy of cocontracting, was apparent in analysis of force impulses

recorded in CT's. We found that subjects produced force in an opposite direction to that of the applied PF field. Although force profiles did not exactly match the PF force, the net lateral force in the period while the PF was active was significantly reduced, thereby reducing the perturbing effect of the field. Evidence suggests that internal model formation, along with cocontraction was used in compensating for the PF.

Although we note a progressive decrease in the level of cocontraction in the "late" periods of movements, we suggest that this may be due to subjects initially cocontracting while the PF is active. Cocontraction while the PF was active would be appropriate in reducing inaccuracies in the feedforward model. We surmise that subjects were slow in reducing levels of cocontraction, such that changes in cocontraction were found even when the PF was no longer active. Alternatively, the PF could be considered as a perturbing pulse that produces a long-lasting displacement, whose effect is still apparent even after the PF itself was no longer active. Such a displacement may be reduced by cocontraction after the initial perturbing effect has been removed, as found in the present study.

We found no change in early rmsEMG in any muscle as learning progressed. As discussed, only shoulder flexion and elbow extension were required in the initial period of the trial, although we noted activity in all muscle groups, including those responsible for actions opposite to those required, i.e., activity was found in posterior deltoid and brachioradialis that provide shoulder extension and elbow flexion torques, respectively. We suggest that no significant change in muscle activity in the early period of movements was found because it represents activity associated with both internal model formation and cocontraction. We suggest that no change was found in the shoulder extensor muscle and elbow flexor muscle due to EMG variability or a relatively high level of residual cocontraction as learning progressed in this early period. We note a major limitation of our study was that we did not have subjects perform NF trials prior to imposition of the PF. Had we done this, we would have been able to determine the activation levels expected in muscles that were not compensating directly for the PF.

Optimization strategy to reduce metabolic cost

It may have been possible for subjects to achieve straighter initial hand paths by excessive muscle cocontraction. However, this can result in muscular fatigue and significant metabolic cost over many trials. Therefore, subjects may have limited co-activation to a level that allowed small deviations from the straight line. There is a trade-off between the level of muscular fatigue and overall movement accuracy. Increasing cocontraction could presumably reduce trajectory error, such that it was negligible, but with a large metabolic cost. After a number of trials subjects would become fatigued, resulting in lower force generating capabilities and an associated worsening of performance. In tolerating a certain amount of trajectory error, subjects can move in a more efficient manner. Reducing endpoint error to an "acceptable" level, while limiting the metabolic cost, may be a more optimal strategy than attempting to move in a perfectly straight line to the target.

Force impulse in CT's

In analyzing subjects' force impulse in CT's throughout PF learning, we found evidence that 6 subjects amended the force impulse as learning progressed. In contrast, 3 subjects appeared not to change their force impulse at any stage of the same learning period. To explain this dichotomous result we propose that the group of 3 subjects used a different strategy in compensating for the PF, namely more cocontraction, such that trajectory error was reduced to a greater extent by the stiffness of the arm than for the other group of subjects, who relied more heavily on an internal dynamics model. Further evidence to support this strategy is apparent in the rms EMG analysis. The average gradient of the linear regression for late rmsEMG, i.e., rms EMG from 150 ms after movement onset until data collection ceased, for the three subject group was lower, in 5 out of 6 muscles, than the respective average gradient for the group of 6 subjects. This did not reach significance at *a priori* α=0.05, although this may be due to the small number of data points, i.e., we were only able to compare 3 gradients to 6 gradients. Although all 9 subjects were found to reduce late rmsEMG as learning progressed, i.e., for the 9 subjects and 6 muscles 49 of the 54 gradients were significantly less than zero, the apparent

difference in gradients suggests that the group of 3 subjects relied more heavily on cocontraction than internal model formation, compared to the 6 subject group, throughout the learning period. In hindsight, it would have been useful to implement several after-effect trials at various stages of the learning in order to assess the relative levels of cocontraction.

Effect of CT's on the subsequent PF trial

We found that subject's performance was adversely affected by the imposition of CT's throughout the learning period. This result is in agreement with that of Scheidt et al. (2001), who reported that subjects altered their feedforward motor commands on the basis of feedback gained in the single previous trial. By nullifying the error feedback normally available in PF trials, using a CT, performance was significantly affected on the subsequent PF trial. Subject's motor commands in the post-CT trials had been amended, resulting in higher hand path error on these trials relative to pre-CT trials. Subjects benefited from the error feedback of a single re-exposure to the PF in the post-CT trial, such that performance immediately improved to that prior to the intervening CT. An earlier study by Scheidt et al. (2000) suggested that imposition of a channel following adaptation to a force field would reduce the rate at which subjects re-adapt to the null field. The lack of kinematic feedback-error in the channel trials resulted in persistence of motor commands more akin to moving in the force field, rather than the null field. This is somewhat different from the present study, as we implemented CT's *during*, and not *following* learning. We found larger trajectory error in trials following CT's, suggesting that subjects amended their feedforward commands to be more appropriate for moving in a lower strength, or null field, rather than persisting with the commands for movement in the PF. A possible reason for this result is that CT's restricted subjects to a straight line. We found that PF paths were not straight, i.e., there was a difference, although not statistically significant, in CT and PF paths, suggesting that subjects *could* benefit from kinematic feedback in the CT's to update their feedforward commands. This is different to the study of Scheidt et al. (2000) where learnt force field paths and NF paths were not significantly different. Throughman and Shadmehr (2000) investigated how subjects were

able to adapt to viscous force fields. Following adaptation they considered whether error feedback from a trial in which the force field was removed, i.e., a catch trial, undertaken in a certain movement direction, had an effect on performance on the subsequent movement in that same direction. They reported that error, measured by perpendicular distance from a straight line, was significant larger in the trial following the catch trial, compared to the error in the trial immediately preceding the catch trial. Our result agrees with this finding, suggesting that the error feedback from the channel trial is not specific to moving in the PF, thus performance in the subsequent PF trial was worse, compared to that in the preceding PF trial.

ΔPF perturbation trials

Ability to adapt to the fields is not determined by physical strength

We do not believe that the ability to straighten hand paths in either the PF or ΔPF fields was limited by the actual physical strength of subjects. The subjects within the study had wide-ranging anthropometric characteristics, yet we found little difference in the magnitude of hand path errors among the subjects in either the learnt-PF, 5^{th} ΔPF_{high} and ΔPF_{low} movements. Furthermore, we found that all subjects moved "into" the field, resulting in deviation to the left of the start-target line, in the ΔPF_{high}, i.e., subjects were able to overcompensate for the PF, at least in the initial part of the trial. Indeed, one of the female subjects, AN, actually had a larger magnitude leftward deviation in the ΔPF_{high} than in the PF (Figure 34, top right panel). This maximum leftward deviation occurred at approximately 0.05 m, where the ΔPF_{high} force was at the maximum strength, adding weight to the argument that strength was not a limiting factor in adapting to any of the fields within the study.

Apparent differences in the level of adaptation to the three fields

We have shown that after 5 trials in the ΔPF_{low}, subjects were able to move along a path that was as straight, relative to a straight line, as the PF path following 150 learning trials. In contrast, the ΔPF_{high} path was significantly less straight than the PF and ΔPF_{low} paths after 5 trials. There are two potential explanations for the apparent differences in the

amount of adaptation to the two perturbed strength fields. The ability to compensate for a given force field may be dictated by neuromuscular factors that constrain the rate of change of joint torque. Since the lower strength force field demands less of the neuromuscular system, one may predict that adaptation to the force field would result in a straighter path, relative to the straight line, than in the PF. In contrast, adaptation to the ΔPF_{high} would result in paths that were less straight than the PF because it required greater changes in the rate of torque production. Alternatively, it may be that the rate at which learning of a novel field occurs is dependent on the relative strength of the fields, although this was not found in a previous study (Lai 2002). In this case the straightness of the 5^{th} ΔPF_{low} trials would be due to an accelerated rate of learning compared to the PF, which in turn was learned faster than the ΔPF_{high}. If this were true, many more trials would have been necessary in the ΔPF_{high} to produce similarly straight paths. In hindsight, it would have been interesting to compare hand paths following complete adaptation in both ΔPF_{high} and ΔPF_{low} fields to the PF paths, thereby testing this idea.

No apparent learning following the 2^{nd} ΔPF trial

We found that subjects produced feedforward adaptive responses to the perturbed strength fields within a single trial, modifying their motor commands to significantly straighten hand paths in the 2^{nd} ΔPF trials, as a result of feedback gained on the 1^{st} ΔPF trials. We noted that very little learning occurred between the 2^{nd} and 5^{th} ΔPF trials, such that the 5^{th} ΔPF trials were no straighter than the 2^{nd} ΔPF trials. This result was somewhat unexpected, given the rapid improvement in learning that was apparent in the PF learning in the first 5 trials. We suggest that the lack of improvement in performance in the ΔPF may be a residual effect of the imposition of CT's within the ΔPF trials. We noted in the PF that the imposition of CT's within the learning period had a negative effect on performance on the subsequent PF trial (Figure 11A), before performance recovered to pre-CT performance on the following trial (Figure 11B). We also found that performance in ΔPF trials that followed CT's was somewhat affected, compared to the corresponding ΔPF trials that had followed another ΔPF trial (Figure 32). Subjects were exposed to 12 ΔPF trials in position 1 of each ΔPF, i.e., CT's replaced no trials in position 1. In

positions 2-5, 3 CT's replaced ΔPF trials, resulting in 9 ΔPF trials in each position. The 9 ΔPF trials in position 2 were all preceded by a ΔPF trial. Of the 9 ΔPF trials in position 3, 3 were preceded by CT's (post-CT trials), while 6 did not follow a CT. In position 4, 3 trials were immediately preceded by CT's (post-CT trials), 3 trials were preceded by a ΔPF trial, which itself had been preceded by a CT (post-CT+1 trials) while 3 trials were preceded by 2 ΔPF trials but no CT trials. ΔPF trials in position 5 were either post-CT trials, post-CT+1 trials or post-CT trials+2. No trials in position 2 were affected by lack of error feedback due to preceding CT's. It is therefore not surprising that we saw obvious improvement in performance between the 1st and 2nd ΔPF trials. We suggest that as in the PF, the CT will cause a decrement in performance, due to the absence of error feedback. Therefore, we may expect performance in the post-CT trials in position 3 to be worse than in position 3 trials that did not follow CT's. Using the same logic, post-CT and postCT+1 trial performance in position 4 may not be as good as those ΔPF trials that followed only other ΔPF trials. All ΔPF trials in position 5 followed a CT (either in position 2, 3 or 4), so performance in all of these trials may not have been as good as otherwise expected. Given the small number of repeated exposures to the perturbed strength fields and the effect of preceding CT's on performance, it is not surprising that we did not find significant improvement in performance following the 2nd ΔPF trial. A larger number of ΔPF sets, some without CT's, would have been of benefit in determining if there was, in fact, progressive learning throughout the 5 ΔPF trials.

EMG during ΔPF trials

We compared EMG in the learnt-PF trials to that in the ΔPF trials in positions 3, 4 and 5. Results were more conclusive in the ΔPF_{high} than in the ΔPF_{low}. In the ΔPF_{high} we found a significant increase in biceps activity for the period while the PF was active. We suggest this shows that subjects adapted to the PF by amending their internal model, increasing shoulder flexor torque and reducing elbow extension torque as required. In the later period of the movement we saw increased in activity of 5 out of 6 muscles compared to the learnt-PF. This indicates that increased cocontraction was used later in the movements to reduce trajectory error as a result of inaccuracies in the newly amended

internal model. Again, combined use of internal models and cocontraction was apparent. Had subjects been allowed to fully adapt to the ΔPF_{high}, we would have expected to see further changes in the EMG activity as the new internal model became more accurate and the associated cocontraction was reduced.

In the ΔPF_{low}, we found a reduction in activity of the posterior deltoid and brachioradialis muscles. Pectoralis major activity was reduced, although this result was not significant, probably due to large variability. The decrease in posterior deltoid activity would have caused an increase in the net shoulder flexor torque. Without a larger decrease in shoulder flexor torque, as a result of a decrease in the activity of pectoralis major or other shoulder flexors, we cannot account for the reduction in shoulder flexion torque, relative to the learnt-PF torque, that was evident in Figure 42. It is possible that our surface EMG electrodes may not have been optimally placed to detect these changes in muscle activity or that the changes occurred in motor units located deep within the muscle. Furthermore, we did not record the EMG of anterior deltoid, which can contribute to shoulder flexion in the horizontal plane. It is quite likely that combined reduction in activity in anterior deltoid, pectoralis major and biceps would have outweighed the effect of the reduction in the posterior deltoid activity to produce a net decrease in shoulder flexor torque. The relative reduction in the activity of each of these muscles may have varied from trial to trial preventing a statistically significant result for any single shoulder flexor muscle. The reduction in posterior deltoid activity would suggest that there was also a reduction in cocontraction relative to muscle activation patterns in learnt-PF, i.e., subjects reduced the stiffness of the arm because they detected the smaller perturbing effect of the lower strength force field. A power analysis to determine the reasons why a number of our EMG analyses were inconclusive can be found in the Experimental Limitations section.

Force Impulse in CT's that replaced ΔPF trials

We found that the three subjects who did not appear to adapt their force impulse significantly over the 4 CT's that we considered within the PF learning period *did* adapt their force impulse in the both the high and low strength ΔPFs. This is, in fact, consistent

with their behaviour in the PF, i.e., in both the PF and ΔPF all subjects adapted their internal model in the first 5 trials. After 5 trials in the PF, the group of three subjects appeared to stop improving their internal model, such that no improvement in performance was seen after the first CT. Had subjects been allowed to adapt fully to the ΔPF, for many more than 5 trials, we would have been able to see if, as in the PF, the group of 3 subjects was different to the group of 6, i.e., one group continued to improve their internal model, while the other group showed no further improvement, simply reducing cocontraction.

The effect of CT's within ΔPF sets

We showed that CT's adversely affected performance in the PF learning. We hypothesized that the CT's we have a similar effect in the ΔPF trials, although when considered whether CT's in positions 2-4 of the ΔPF had affected the performance in the following ΔPF trial, the result was inconclusive. We speculate that that the larger variability in performance in the ΔPF, due to the fact learning was at a very early stage, may have led to the non-significant results.

Experimental Limitations

EMG analysis resulted in a number of statistically non-significant results. This was especially prevalent in cases where we compared rms EMG in trials following adaptation to the ΔPF to the rms EMG in trials following full adaptation to the PF. Torque analysis had revealed that subjects amended their joint torques when adapting to the ΔPF fields, compared to the joint torques produced in the learnt-PF. Changes in muscle activation should therefore also be expected, corresponding to the changes in joint torque. Analysis of rms EMG failed to find differences between rms EMG in learnt-PF and ΔPF_{low} trials, although increases in rms EMG in ΔPF_{high} trials, relative to the learnt-PF were found. A large number of the EMG and kinematic comparisons that we made within the study were significant. The repeated measures nature of the experiments enabled us to undertake ANOVA's, which yielded more detailed results than simple t-test comparisons.

Rather than too few subjects, significance of our some of our results probably suffered due to the relatively small number of repeated exposures to the perturbed strength fields over which we were able to average. This is of particular importance when we consider that imposition of channel trials affected performance, limiting the number of trials over which we could average. If, for example, we had only replaced 10% of ΔPF trials in each position with CT's, and increased the number of ΔPF exposures from 12, to 30, in each strength field, we may have obtained more significant results. EMG has high inherent variability because of its stochastic nature and requires averaging over several trials to reduce the variability of the signal sufficiently that reliable inferences can be made. In certain cases, conclusions were more definite because we were able to average over a larger number of similar trials.

Power Analysis

To investigate whether lack of statistical power in our rms EMG analysis could have been responsible for some of our non-significant results, a post-hoc power analysis was undertaken. An *a priori* power analysis to determine the number of subjects or repeated exposures required to yield sufficiently high power given a certain effect size, although good scientific practice, was not possible as we had no way of estimating the variability between EMG signals prior to data collection. Furthermore, a power calculation based on Cohen's effect sizes (Cohen, 1988) would have been of limited benefit in accurately predicting the number of subjects, or repeated exposures to our ΔPF trial conditions, required to give a specified power, as we had little idea of the effect size we would have expected. (The effect size would be dependent on many variables, such as how well the subjects adapted to the force fields, and the variability of the EMG signal.)

Post-hoc power calculations were carried out for the comparisons of rms EMG in the early and late periods of the trial, between learnt-PF and ΔPF_{low}, and between learnt-PF and ΔPF_{high}. Recall that early rmsEMG refers to rms EMG in the 300 ms period beginning 150 ms prior to movement, while late rmsEMG refers to rms EMG from 150

ms following onset until the end of data collection in that trial. We determined that for the comparisons of late rmsEMG, between learnt-PF and ΔPF_{high} trials in the 2nd position of each set, the statistical power ranged from 0.44-0.83 for the 6 muscle comparisons (Howell, 1997). Note that for this comparison, we found statistical differences in the rmsEMG in 5 of the 6 muscles ($p<0.05$). In the early rmsEMG comparison between the learnt-PF ΔPF_{high} trials, where we found a significant difference for only one muscle (pectoralis major), power was much lower, and did not exceed 0.17. Similarly, in the late and early rmsEMG comparisons between the learnt-PF and ΔPF_{low}, where no significant differences where found, power did not exceed 0.21.

Although these power values may be low, the primary reason for the low power appears to have been very small differences between the mean rms EMG values that were being compared. Indeed, the variability (standard deviation) in the rms EMG was very similar in all four comparisons addressed here, while the differences in means were considerably larger (in most cases, at least an order of magnitude) in the comparison between the late rmsEMG in the learnt-PF and ΔPF_{high}, i.e., where significant differences were detected. In cases where we found no significant difference between rms EMG values, the effect size never exceeded 0.16, and was generally less than 0.06, while in cases where we found significant differences in the rms EMG, the effect size ranged from 0.20-0.34. Cohen (1988) defines an effect size of 0.20 as a small effect where the effect is "real but difficult to visually detect". Rather than simply estimating the number of samples required to increase the power of the test, such that one would have higher probability of detecting a statistical difference between the values, we must consider that differences where the effect size is so small are not real, i.e., the null hypothesis that there is no change in the rms EMG may, in fact, be true. With an effect size similar to that recorded in the comparison of late rmsEMG between the learnt-PF and ΔPF_{high} (where a true difference in the rms EMG was apparent, albeit a fairly small effect size of 0.20-0.34), we calculated that with about three times as many trials, the statistical power would be at least 0.85, compared to the range in power that we achieved of 0.44-0.83. This would involve increasing the number of ΔPF exposures from 12, to 30, in each strength field,

and reducing the percentage of trials from 25% to 10% of ΔPF trials in each position that were replaced with CT's. Furthermore, a higher number of trials should reduce the variability of the EMG, increasing the effect size and thus resulting in a further increase in power above that achieved by simply increasing the number or trials. We also note that to increase the power in analyses of CT's, we would require more than 3 CT's in each position. This would further increase the total number of trials required to achieve the desired power in analysis of both ΔPF trials and CT's.

Increasing the number of exposures to ΔPF_{low} and ΔPF_{high} is not all that simple. A three-fold increase in the number of ΔPF exposures, together with the previously suggested complete adaptation to the ΔPF fields, would increase the length of the experimental session from about 90 minutes to well over 240 minutes. Due to the possibility of muscular fatigue, and discomfort that subjects would suffer from sitting in the same position for extended periods while simultaneously holding their arm in the desired position, a number of separate experimental sessions would be needed. This would allow the subject time to rest for periods between sessions of the experiment, preventing fatigue and excessive discomfort. One problem with this protocol would be that subjects might require some trials to allow re-adaptation to the field following each rest period. All sets would have to be conducted within the same day to allow comparison between the EMG in the different sessions.

Electrode placement and contact

We must also consider other reasons for the why we failed to find any significant change in muscle activity in a number of our comparisons, despite changes in the joint torques, that could only have been produced with associated changes in muscle activity. Incorrect electrode placement and poor skin-electrode contact are two potential factors. If the electrodes were not optimally placed, changes in activity within the targeted muscle may not have been fully represented in the EMG data. Furthermore, there is a possibility that inaccurate electrode placement resulted in activity from other non-targeted muscle groups being recorded. Had this been the case, the reliability of the recorded EMG would have

been lowered. Poor electrode contact with the skin would have resulted in a higher inter-electrode resistance, resulting in less detectable changes in the electrical signals produced within the muscle.

Failure to record NF trials and allow full adaptation to the ΔPF fields

One major shortcoming of the experiment was that we did not record any NF trials, prior to the PF learning set. This meant that we were unable to compare EMG data in the PF and ΔPF to NF data. Such a comparison may have yielded interesting results, such as which muscles showed the largest changes in activation. Changes in muscle activation between the NF and the PF would probably have been much larger than the change that we saw between the PF and ΔPF trials. We found only small changes in EMG between the PF and ΔPF trials. We were fortunate to be able to draw on the kinematic data from a previous study, allowing hand path comparisons between the various fields, as well as enabling the predictions of joint torques. However, because this study was not performed on the same day, it was not possible to compare the EMG. As previously mentioned, it would have been of interest to determine the complete adaptation to the ΔPF fields following the experiment, which we actually conducted. Electrodes would not have been removed between these sessions, such that we could compare EMG in the 5 ΔPF trials to a complete adaptation. Comparisons of fully adapted hand paths in the PF, ΔPF_{low} and PF_{high} would have also been of benefit in drawing conclusions from our results.

Perturbations did not elicit noticeable stretch reflex responses

We found that perturbations caused by the CT, both in PF learning and in the ΔPF trials, did not produce joint perturbation velocities large enough to induce significant stretch reflex activity in the stretched muscles. The magnitude of the initial hand path perturbations in the ΔPF trials, relative to the learnt-PF trajectory did not exceed 4 cm. Although this is actually quite large compared to most other studies, had the change in strength between the various fields been larger, a larger perturbation would have been expected. This may have yielded better results when subjects adapted to the ΔPF fields. The small magnitude perturbations may have not provided sufficient motivation for the

subjects to adapt; once they reduced the magnitude of the perturbation following the initial exposure no further significant learning was observed. This may be, as explained above, due to the imposition of CT's, but alternatively may have been due to subjects' lack of motivation to further straighten paths.

Force field strength was lower than expected

Following initial analysis of data, we determined that the strength of the force fields was a factor of 2 smaller than had initially been desired, i.e., the maximum forces in the PF, ΔPF_{low} and ΔPF_{high} were intended to be 12 N, 8 N and 16 N, respectively, rather than 6 N, 4 N and 8 N, which we actually measured. This apparent error in the programming of the motor would have caused smaller trajectory errors when the field strength was perturbed. Subjects may have adapted more had the perturbations been of the originally intended magnitude.

Summary of experiment 2

At the start of the investigation we outlined 4 main questions that we hoped to answer during this experiment. We asked 1) what are the relative contributions of IM formation and impedance control in early adaptation to a novel force field 2) How does the level of cocontraction vary with repeated exposure to the field? 3) How quickly can subjects adapt their IM when the expected environmental force is perturbed? 4) How does the level of cocontraction vary when perturbed field strengths are initially encountered and subsequently learnt?

We were able to address each of these questions with our analyses. Firstly, we found that initially there appears to be a high level of cocontraction when learning to move in a novel field. However, the PF learning revealed that as early as the second or third trial in a novel force field subjects began to develop an internal model, although there was still a high level of associated cocontraction. Secondly, we found that the level of cocontraction fell as subjects adapted to the PF, i.e., as the internal models became more accurate. Thirdly, we found that internal models were adapted to the perturbed strength fields

following feedback from a single trial, i.e., on the second trial in the perturbed field subjects' internal model had been amended to be more suitable for the perturbed external dynamics, compared to the well-learned PF. Fourthly, we found an increase in the cocontraction on the second exposure to the high strength field. This remained at an elevated level for all subsequent ΔPF_{high} trials, relative to that in the PF. Although our experiment did not allow subjects to completely adapt to the perturbed strength fields, we suggest that, as in the PF, the level of cocontraction would also fall as learning progressed and the amended internal models became more accurate.

Thesis Summary

Two distinct, yet complementary experiments were undertaken within this thesis. Experiment 1 provided a direct test of the equilibrium point hypothesis (EPH) of movement control. A torque motor was used to provide negative damping, which assisted subjects' movements to the target. The level of negative damping was randomly reduced on 10% of trials, and we measured the affect of this perturbation, both in terms of movement kinematics and feedforward EMG signals. On these trials there was an undershoot, which increased in a linear fashion, as the magnitude of the reduction in assistance increased. This result is contrary to the EPH, which predicts that, under a load such as the negative damping, muscular reflex responses should drive the limb to the same final position, irrespective of the change in load, given that the feedforward command did not change (something that we did in fact show to be true). The dependence of the final position on the change in the assisting load is entirely consistent with predictions based on the formation of an internal dynamics model. Following extensive learning, subjects would learn to generate the correct amount of torque to reach the target, and stabilize the limb at this final position for specific dynamics. When dynamics were altered by reducing the level of assistance, the net torque, i.e., the combination of the motor torque and the torque subjects produced, was insufficient to drive the limb to the target, resulting in an undershoot of the final desired position.

As experiment 1, and considerable previous research suggests, subjects appear to adapt to novel dynamics by forming an internal model of the required task. We designed a second experiment to more closely investigate the learning process. We considered both long-term learning, through many consecutive exposures to the same external force field (PF), and short-term learning, by repeatedly exposing subjects to a different force field strength (ΔPF) 5 times, but preventing full adaptation by interrupting learning with trials conducted in the original force field. To our knowledge, no previous study has used this technique to closely study the early-learning period. Mechanical channel trials were used to test for evidence of internal model formation.

In the long term learning of the PF force field, we found evidence of internal model formation in the second or third trial. Subjects produced lateral force to counteract the force field. The temporal profile of the force that subjects produced did not perfectly match the PF force, resulting in bowed trajectories that were characteristic in all subjects. Hand path error was progressively decreased over ~ 45 trials, after which variability in performance dominated over any further improvement. We noted that most of the progressive learning occurred over the portion of the movement beyond the region of the PF. As learning progressed, we found significant decreases in the EMG in all 6 shoulder and elbow muscles once subjects had passed beyond the force field. We suggest that subjects modified their internal model with extended practice to more accurately compensate for the PF, such that the cocontraction could be reduced, without any associated degradation in performance, as measured by hand path errors.

Subjects quickly readjusted to the PF following each ΔPF set, suggesting that they retained the motor commands required to compensate for the PF. It took only 2 or 3 PF trials following each ΔPF set for subjects' performance to return to the same level as that following complete adaptation to the PF. On the first ΔPF trial of each set, subjects were displaced in the direction of the change in force. The displacement, relative to the PF path, was similar for both the low and high strength perturbations. On the second trial, trajectories were substantially straighter, suggesting that subjects benefited from error feedback from a single ΔPF trial to significantly reduce hand path error. Hand path error was not reduced significantly more in the 3^{rd}-5^{th} ΔPF trials. We suggest that this may have been due to the channel trials interfering with progressive learning during the 5 ΔPF trial set. Subjects adapted their internal model to more accurately compensate for the higher strength ΔPF, evident as changes in EMG activity corresponding to the required changes in joint torque, while simultaneously employing higher levels of cocontraction, characterized by increases in activity of antagonistic pairs of muscle later in the movement. In the lower strength ΔPF, reduction in the activity of antagonistic pairs of muscles suggest the level of cocontraction decreased, perhaps reflecting the lower perturbing forces associated with this field.

Overall, this study has complemented the existed literature in the field of motor control. Our first experiment provided a less ambiguous test of the EPH, while the second experiment considered how adaptation to a novel force field occurs, both in the initial few trials, and over an extended time period.

There is a great deal more work that can be done related to the experiments reported within this thesis. In experiment 1 we showed that the EPH does not hold for the specific experimental condition that we tested. To understand more precisely the conditions under which the EPH could potentially explain experimental results, and test the generality of our findings, we must consider different movement directions, with varying movement velocities and perturbation types. We have recently begun work on a follow up study in which we hope to run experiments where subjects undertake both flexion and extension wrist movements.

We would also like to test how the magnitude of the reflex gain affects the final position in the perturbed movements. Subjects would be screened, using a test to categorize their stretch reflex activity and subsequently divided into two groups (high and low reflex gains). Both groups would then participate in a similar experimental protocol to that described in experiment 1. The strength of the reflex gain may result in differences in final position on perturbation trials, between the two groups.

There are many follow-up experiments linked to experiment 2. A similar experiment in which subjects were allowed to fully adapt to the ΔPF following the sets of 5 trials would allow greater insight into the rate of learning in force fields of different strengths, as well as determining what the final adaptation to the various strength fields would be. Further interesting experiments may involve exposing subjects to a similar position dependent field, but in the final 0.1 m, as opposed to the first 0.1 m, of the reach. Strategies used by subjects to adapt to the two varying load conditions may differ, due to the timing of the error feedback available to the subjects.

Appendix 1- Ethics Approval

SIMON FRASER UNIVERSITY

OFFICE OF RESEARCH ETHICS

BURNABY, BRITISH COLUMBIA
CANADA V5A 1S6
Telephone: 604-291-3447
FAX: 604-268-6785

July 8, 2003

Mr. Mark Hinder
Graduate Student
School of Kinesiology
Simon Fraser University

Dear Mr. Hinder:

Re: Adaptive control of goal-oriented human arm movement
NSERC - Discovery Grant

I am pleased to inform you that the above referenced Request for Ethical Approval of Research has been approved on behalf of the Research Ethics Board. The approval for this project is for the term of the period of the grant, as defined by the funding agency. If this project does not receive grant support, the term of the approval is twenty-four months from the above date.

Any changes in the procedures affecting interaction with human subjects should be reported to the Research Ethics Board. Significant changes will require the submission of a revised Request for Ethical Approval of Research. This approval is in effect only while you are a registered SFU student.

Your application has been categorized as 'minimal risk" and approved by the Director, Office of Research Ethics, on behalf of the Research Ethics Board in accordance with University policy R20.0, http://www.sfu.ca/policies/research/r20-01.htm. The Board reviews and may amend decisions or subsequent amendments made independently by the Director, Chair or Deputy Chair at its regular monthly meetings

"Minimal risk" occurs when potential subjects can reasonably be expected to regard the probability and magnitude of possible harms incurred by participating in the research to be no greater than those encountered by the subject in those aspects of his or her everyday life that relate to the research.

Please note that it is the responsibility of the researcher, or the responsibility of the Student Supervisor if the researcher is a graduate student or undergraduate student, to maintain written or other forms of documented consent for a period of 1 year after the research has been completed.

Best wishes for success in this research.

Sincerely,

Dr. Hal Weinberg, Director
Office of Research Ethics

c: Dr. Theodore Milner,
Supervisor & Co-investigator

References

Asatryan D.G., Feldman A.G. (1965). Functional tuning of the nervous system with control of movement or maintenance of a steady posture-I. Mechanographic analysis of the work of the joint or execution of a postural task. *Biophysics* 10, 925-935.

Bizzi E., Polit A. (1978). Processes controlling arm movements in monkeys. *Science* **201,** 1235-1237.

Bizzi E., Polit A. (1979). Characteristics of the motor program underlying arm movements in monkeys. *Journal of Neurophysiology* **42**, 183-194.

Bizzi E., Dev P., Morasso P., Polit A. (1978). The effect of load disturbances during centrally initiated movements. *Journal of Neurophysiology* **41,** 542-556.

Bizzi E., Accornero N., Chapple W., Hogan N. (1982). Arm trajectory formation in monkeys. *Experimental Brain Research* **46**, 139-143.

Burdet E, Osu R, Franklin DW, Milner TE, Kawato M (1999). Measuring stiffness during arm movements in various dynamic environments. In: *Proceedings of the 1999 ASME Annual Symposium on Haptic Interfaces and Virtual Environments of Teleoperator Systems*. 421-428

Burdet E, Osu R, Franklin DW, Milner TE, Kawato M (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. Nature **414**,446-449.

Cohen J. (1988). Statistical power analysis for the behavioural sciences (2nd edition). Academic Press, New York.

Colgate JE, Hogan N (1988) Robust control of dynamically interacting systems. *Int J Control* **48**,65-88.

Conditt M.A., Gandolfo F., Mussa-Ivaldi F.A. (1997). The motor system does not learn the dynamics of the arm by rote memorization of past experience. *Journal of Neurophysiology* **78**, 554-560.

Dizio P and Lackner JR (2001) Coriolis-force-induced trajectory and endpoint deviations in the reaching movements of labyrinthine-defective subjects. *Journal of Neurophysiology* **85**, 784-789.

Feldman A.G., Ostry D.J., Levin M.F., Gribble P.L., Mitnitski A.B. (1998). Recent tests of the equilibrium-point hypothesis (lambda model). *Motor Control* **2**, 189-205.

Feldman A.G. (1986). Once more on the Equilibrium-Point Hypothesis (λ Model) for motor control. *Journal of Motor Behavior* **18**, 17-54.

Feldman A.G. (1966a). Functional tuning of the nervous system with control of movement or maintenance of a steady posture-II. Controllable parameters of the muscle. *Biophysics* **11**, 565-578.

Feldman A.G. (1966b). Functional tuning of the nervous system with control of movement or maintenance of a steady posture-III. Mechanographic analysis of execution by man of the simplest of motor tasks. *Biophysics* **11**, 766-775.

Feldman, A. G.; Adamovich, S. V., and Levin, M. F.(1995) The relationship between control, kinematic and electromyographic variables in fast single-joint movements in humans. *Experimental Brain Research* **103**, 440-50.

Franklin D.W. (2000). Biomechanics and neural control of limb position. M.Sc. thesis, Simon Fraser University.

Franklin D.W., Osu R., Burdet E., Kawato M., Milner T.E. (2003a) Adaptation to stable and unstable dynamics as achieved by combined impedance control and inverse dynamics model. *Journal of Neurophysiology* (in press)

Franklin D.W., Burdet E., Osu R., Kawato M., Milner T.E. (2003b) Functional significance of stiffness in adaptation to stable and unstable dynamics. *Experimental Brain Research* **151** 145-157

Gandolfo, F.; Mussa-Ivaldi, F. A., and Bizzi, E. (1996) Motor learning by field approximation. *Proceedings of the National Academy of Science U.S.A* **93**,3843-6.

Gomi H., Kawato M. (1996). Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* **272**, 117-120.

Goodbody SJ and Wolpert DM (1998) Temporal and amplitude generalization in motor learning. *Journal of Neurophysiology* **79**, 1825-1838.

Gribble, P.L., Ostry, D.J., Sanguineti, V., Laboissière, R. (1998). Are Complex Control Signals Required for Human Arm Movement? *Journal of Neurophysiology* **79**, 1409-1424.

Gribble PL and Ostry DJ (2000) Compensation for loads during arm movements using equilibrium-point control. *Experimental Brain Research* **135**, 474-482.

Hollerbach M.J., Flash T. (1982). Dynamic interactions between link segments during planar arm movement. *Biological Cybernetics* **44**, 67-77.

Howell D.C. (1997). Statistical Methods for Psychology (4th edition). Duxbury Press, Wadsworth Publishing Company, Belmont California.

Hogan N. (1985). The mechanics of multi-joint posture and movement control. *Biological Cybernetics* **52**, 315-331

Kawato M (1999) Internal models for motor control and trajectory planning. *Current Opinions in Neurobiology* **9**, 718-727.

Lackner J.R., Dizio P. (1994). Rapid adaptation to coriolis force perturbations of arm trajectory. *Journal of Neurophysiology* **72**, 299-313.

Lai E.J., (2002). Sensitivity and complexity of adaptive models in human motor control. M.Sc. Thesis, University of British Colombia.

Latash, M. L. and Gottlieb, G. L. (1991) Reconstruction of shifting elbow joint compliant characteristics during fast and slow movements. *Neuroscience* **43**, 697-712.

McIntyre J, Mussa-Ivaldi FA, Bizzi E (1996) The control of stable postures in the multijoint arm. *Experimental Brain Research* **110**:248-264.

Milner TE and Cloutier C (1993) Compensation for mechanically unstable loading in voluntary wrist movement. *Experimental Brain Research* **94**, 522-532.

Milner T.E., Cloutier C. (1998). Damping of the wrist joint during voluntary movement. *Experimental Brain Research* **122**, 309-317.

Milner T.E. (2002). Adaptation to destabilizing dynamics by means of muscle co-contraction. *Experimental Brain Research* **143**, 406-413.

Mussi-Ivaldi F.A., Hogan N., Bizzi E. (1985). Neural, mechanical and geometric factors subserving arm posture in humans. *Journal of Neuroscience* **5**, 2732-2743.

Nakazawa, K., Yamamoto, S. -I., Ohtsuki, T., Yano, H., and Fukunaga T. (2001). Neural Control: novel evaluation of stretch reflex sensitivity. *Acta Physiol Scand* **172**, 257-268.

Osu R., Franklin D.W., Kato H., Gomi H., Domen K., Yoshioka T., Kawato M. (2002) Short- and long-term learning in joint co-contraction associated with motor learning as revealed from surface EMG. *Journal of Neurophysiology* **88**, 991-1004.

Popescu F.C., Rymer W.Z. (2000). End points of planar reaching movements are disrupted by small force pulses: An evaluation of the hypothesis of equifinality. *Journal of Neurophysiology* **84**, 2670-2679.

Rothwell J.C., Traub M.M., Day B.L., Obeso J.A., Thomas P.K., Marsden C.D. (1982). Manaul performance in the deafferented man. *Brain* **105,** 515-542.

Sanes J.N. (1986). Kinematics and end-point control of arm movements are modified by unexpected changes in viscous load. *Journal of Neuroscience* **6**, 3120-3127.

Sanes J.N., Evarts E.V. (1983). Effects of perturbations on accuracy of arm movements. *Journal of Neuroscience* **3**, 977-986.

Scheidt R.A., Reinkensmeyer D.J., Conditt M.A., Rymer W.Z., Mussa-Ivaldi F.A. (2000). Persistence of motor adaptation during constrained, multi-joint, arm movements. *Journal of Neurophysiology* **84**, 853-862.

Scheidt R.A., Dingwell J.B., Mussa-Ivaldi F.A. (2001). Learning to move amid uncertainty. *Journal of Neurophysiology* **86**, 971-985.

Shadmehr R., Mussa-Ivaldi F.A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience* **14**, 3208-3224.

Sokal R.R., Rohlf F.J. (1981). Biometry: The principles and practice of statistics in biological research (second edition). W.H. Freeman and Company, San Francisco.

St-Onge, N.; Adamovich, S. V., and Feldman, A. G. (1997) Control processes underlying elbow flexion movements may be independent of kinematic and electromyographic patterns: experimental study and modelling. Neuroscience; **79**, 295-316.

Thoroughman K.A., Shadmehr R. (2000). Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742-747.

Vidasager M. (1978). Nonlinear systems analysis (New Jersey): Prentice-Hill

Weeks, D. L.; Aubert, M. P.; Feldman, A. G., and Levin, M. F (1996). One-trial adaptation of movement to changes in load. *Journal of Neurophysiology* **75**, 60-74.

Winter D.A. (1990) Biomechanics and motor control of human movement. (New York): John Wiley and Sons.