

The influence of a forward model of arm dynamics on eye behavior in saccadic tracking
of manual reaching tasks

by

Gregory D. Ariff

A thesis submitted to Johns Hopkins University
in conformity with the requirements for the degree of
Master's of Science in Engineering

Baltimore, Maryland

October, 2001

© Gregory D. Ariff 2001

All rights reserved

Abstract

To compensate for the potentially destabilizing effects of sensorimotor delays on human movement, it has been proposed that the brain implements a forward model of dynamics. Such a model would predict future state of the body (or a particular body part) by using efferent copy of motor commands to integrate forward from a delayed sense of state. Simulations of a hypothesized arm controller utilizing a forward model have been shown to generate similar behaviors as humans performing point-to-point reaching movements in novel dynamical environments. These environments are generated by having the subjects grip a handle connected to a robotic manipulandum that is capable of applying perturbing forces to the hand while performing the reaching task. In the experiments described in this thesis, subjects made reaching movements while gripping this same manipulandum and were trained to look at the perceived location of their hands in the absence of visual feedback. The resulting eye motion was predominantly saccadic, having saccade endpoints which predicted arm behavior by several hundred milliseconds. In the presence of movement error (imposed by the robot), both efferent copy of arm motor commands and sensory feedback affected eye motion. The timing and location of saccade endpoints with respect to hand position are shown to be consistent with the output of a forward model of arm dynamics in their manner of predicting future arm state. In the light of previous work in motor learning and eye-hand coordination, these results support the role of forward models for not only control of individual motor systems, but also for the coordination of different motor systems.

Acknowledgements

Throughout my educational experience at Johns Hopkins, I received invaluable help, advice, and encouragement from faculty, students, family, and friends. Within the Laboratory for Computational Motor Control, I would like to thank my mentors: Drs. Reza Shadmehr, Opher Donchin, and Maurice Smith. Their extraordinary knowledge and gifts as scientists were my inspiration for research and were significant contributors to whatever advances I have made as a scientist and thinker. I would also like to thank the numerous wonderful individuals both in the laboratory and in the department of Biomedical Engineering as a whole who not only provided me with advice and encouragement, but also with brains and arms as experimental subjects. Among them, I owe special thanks to Reza Nezafat, Goran Dordevic, Joseph Francis, Vincent Huang, Stephanie Wainscott, Maneesh Dewan, EunJung Hwang, James Hartwell, Paul George, Abhishek Bandyopadhyay, Hillary Huszar, Jeffrey Wallace, Arup Roy, Ramani Balu, Neda Poomipanit, Steven Chase, and, for her love and wisdom, Haiyin Chen. I would like to thank my reviewers, Drs. Eric Young and Kurt Thoroughman, both for taking the time to lend their insights to my work and also for their valuable instruction. Many friends outside of the university were also profoundly important to me in these past few years. James Bubb, William Venner, Jonathan Watts, Oliver Popp, and Denise Rosenthal have my deepest respect and affection. Perhaps the greatest contributing factor to my accomplishments, in graduate school and throughout life, has been the love and support of my parents, Steve and Mickey, and my sister, Megan. I feel more gratitude to them than can be expressed.

Contents

1 Introduction	1
Goal	1
Forward models for motor control	2
Eye-hand coordination	8
Forward models and visual tracking tasks	11
2 Methods	13
Experimental setup and data acquisition	13
Experiments	15
Analysis	22
3 Results	27
Null field: Straight and path-tracing (curved) movements	27
Mixed-fields: Assistive, resistive, CW curl, CCW curl, and null	30
Passive curl movements	35
Force-pulse movements: Null-Kick-(Null/Assistive/Resistive)	38
Force-pulse movements: Varied pulse times	50
Passive kick experiments	58
4 Discussion and Conclusions	62
Restating the question	62
Revisiting the results	63
Difficulties in demonstrating the adaptation of forward models	66
Conclusions	68

Appendix	71
References	73
Scholarly Life	74

List of Tables

2.1	Summary of experiments	19
-----	------------------------	----

List of Figures

1.1	Examples of forward and inverse models of dynamics	6
1.2	Control system for human arm using inverse and forward models	8
2.1	Experimental setup	13
2.2	Example of calibration errors	15
2.3	Target directions relative to the subject	16
2.4	Comparison of CW and CCW curl field trajectories	23
2.5	Example null field trajectories	24
3.1	Hand and eye averages for null field sets	28
3.2	Eye-hand distances for null field sets	29
3.3	Start time of second saccade varies with the hand	32
3.4	Hand and eye averages for null field trials in mixed fields set	33
3.5	Hand and eye averages in assistive and resistive fields	34
3.6	Hand and eye averages in curl fields	36
3.7	Forward model output from simulated curl field movement	36
3.8	Qualitative comparison of eye and forward model for curl fields	37
3.9	Hand and eye averages for passive condition	39
3.10	Eye-hand distance minima for passive and curl field trials	39
3.11	Hand and eye averages and distance minima for No Kick trials	41
3.12	Hand and eye averages and minima for Null-Kick-Null trials	42

3.13	Hand and eye averages for Null-Kick-Assistive and –Resistive trials	44
3.14	Eye-hand distances for Null-Kick-Assistive and –Resistive trials	45
3.15	Saccade perpendicular displacement for Kick trials in three fields	46
3.16	Determination of saccade prediction time for Null-Kick-Null trials	49
3.17	Saccadic prediction of perpendicular error for Kick trials in three fields	49
3.18	Simulation of hand and forward model behavior for three fields	51
3.19	Qualitative comparison of eye and forward model for kick trials	52
3.20	Eye and hand averages for variable kick-time trials	54
3.21	Average of first two saccades for variable kick-time trials	55
3.22	Eye-hand distance minima for variable kick-time trials	56
3.23	Saccadic prediction of perpendicular error for variable kick-time trials	57
3.24	Eye and hand averages for passive kick trials	59
3.25	Average of first two saccades for passive kick trials	60
3.26	Eye-hand distance minima for passive kick trials	61

Chapter 1. Introduction

Goal

In recent years, a growing body of evidence in the field of human motor control has pointed to the existence of forward models. Forward models have been proposed to explain how humans and animals might compensate for characteristics of motor control systems such as temporal delays and self-generated sensory effects. Recently, simulation work in our lab has shown that a control system for the arm which utilizes a forward model of arm dynamics can explain particular features of human arm behavior in point-to-point reaching movements that other models cannot. However, direct methods for examining the forward model output remain elusive.

Another related body of research has been the study of the spatial and temporal relationship of eye and hand motion during tasks such as reaching, grabbing, tracing, and pointing. Studies have shown that both humans and monkeys exhibit significantly better performance at visually tracking self-generated movements than externally generated movements. It is clear that the control system of the eye has some “knowledge” of the hand’s motion, and it has been suggested that the shared information contains the efferent copy of the motor command signals to the arm, as well as sensory information from the arm.

The aim of this thesis is to investigate whether eye motion during visual tracking of arm motion in humans is influenced by the output of a forward model of arm dynamics. Such output would be an estimate of current or future state and has the power

to explain many of the behaviors observed in previous visual tracking experiments. By examining the behavior of the eyes relative to the arm in visually tracked reaching motions with induced errors, we sought more conclusive evidence for the role of a forward model of arm dynamics in eye motion.

Forward models for motor control

As a fundamental part of everyday life, people interact with their environment to achieve various aims. This can be as simple as pushing a button or as complicated as riding a bicycle with only the back wheel on the ground. In many cases, the physical properties of the objects being manipulated are not known before the interaction and must be learned. Through repeated attempts, people learn how to achieve desired results with the objects. Current theories of motor control state that humans learn “internal models” of physical systems. Internal models are the brain’s representation of physical systems, including the human body and objects in the external world, which translate intended motion into motor commands and, hence, actual movements.

Our current understanding of how the brain controls voluntary muscle movement largely comes from the field of robotics. In controlling a robot, the task is to achieve a desired trajectory by sending the appropriate control signal, while also providing error correction through feedback. In the case of the human being, the control signal is the motor command sent from the brain to the muscles, and the feedback is largely from sensory information relayed to the spinal cord and to the brain. The control signal is generated using an *inverse* model, which transforms a desired trajectory into the control signal. However, inherent delays in sensory feedback in biological systems pose

particular problems for feedback control, potentially causing the system to become unstable. Robot controllers often solve this problem with the use of a *forward* model, or a model of plant dynamics which allows for estimation of current or future state based on delayed feedback and an up-to-date control signal. The existence of inverse and forward models in biological systems, including the controller for the human arm, has been proposed. In order to describe the role of a forward model in control of the human arm, we will first provide a brief description of some simple internal models, and work towards the complex controller required for the arm.

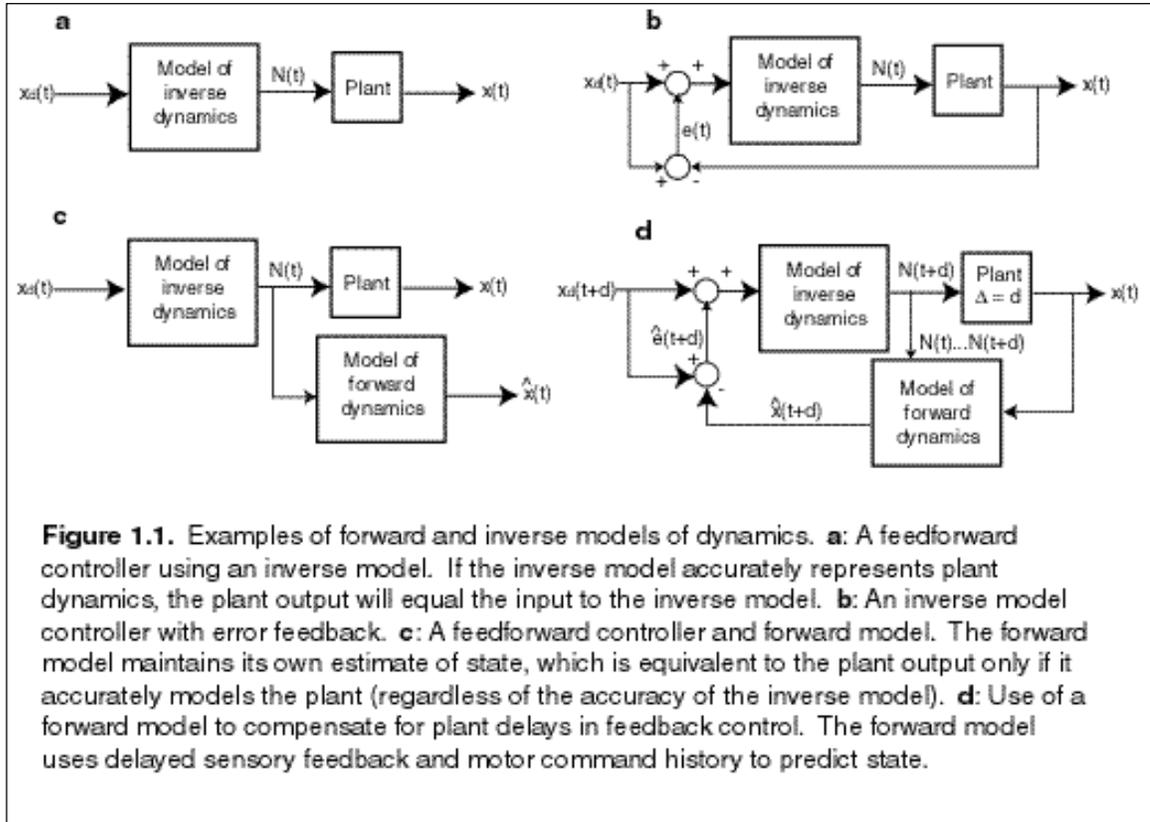
Two types of internal models have been proposed: inverse models and forward models. Both of these models are thought to serve important roles in motor control. An inverse model is an estimation of the inverse dynamics of a system and is used to generate the necessary motor commands for an action. Whereas the physical system (or “plant”) receives a motor command and generates a motion, the inverse model takes as input the desired trajectory and outputs a motor command. If, for example, a person wishes to lift a teacup off a table, the inverse model must take into account the dynamics of the human arm as well as those of the cup and its contents. A very simple feedforward system using only an inverse model and plant is shown in Figure 1.1a. In this system, a desired trajectory is the input to the system. The inverse model computes appropriate motor commands to achieve the desired trajectory, which are then sent to the plant. The plant generates a trajectory. If the inverse model is an accurate representation of the plant dynamics, the actual trajectory will be exactly the same as the desired trajectory. However, if the inverse model is inaccurate, the actual trajectory will not follow the desired trajectory. In the example of the teacup, if there is more or less tea in the cup

than anticipated, the lifting motion will be different than the planned motion. To compensate for deviations of the actual trajectory from the desired trajectory, error feedback can be used to alter the input to the inverse model, as in Figure 1.1b. In this way, inaccuracies of the inverse model can be compensated for as the error is sensed.

In their 1996 review article on forward models, Miall and Wolpert define two types of forward models. The first is a model, which would “aim to mimic or represent the normal behavior of the motor system in response to outgoing motor commands”. The second would “encapsulate knowledge of the physical properties of the environment, and predict the behaviour of the external world”. To achieve a particular task, such as picking up a teacup or making short reaching movements, we can combine these ideas to define a forward model as a model which aims to mimic the behavior of the motor system *and* its interaction with the external world. Moreover, we will limit our discussion to forward models *of dynamics*: explicit estimates of the dynamics of a system which take as input motor commands and sensory information and which output estimates of behavior (or state). The forward model is thus, in some sense, simply a model of the plant. Again, using our example, if a forward model is used to estimate the behavior of the arm holding the cup of tea (as in Figure 1.1c), the accuracy of its estimate is dependent on how well the forward model models the actual plant. If the forward model under- or overestimates the amount of tea in the cup, its estimate of behavior will also be inaccurate.

The forward model in Figure 1.1c, however, is not serving in any way to control the system. As described earlier, a forward model is particularly useful to compensate for delays in sensory feedback. If the controller treats delayed sensory feedback as current, instabilities may occur. However, an appropriate forward model should be able to

estimate current or future state of the plant by integrating the motor command forward from the delayed sensory feedback. This state estimate can then be compared with the desired behavior to estimate current or future error of the plant. A simple modification of our example system to use a forward model which corrects for plant delays is shown in Figure 1.1d. In this system, the estimated error is used to change the input to the inverse model. If the inverse model is accurate (producing the appropriate motor commands and trajectory), the forward model is not needed. But as soon as any error in control occurs, the forward model serves an important function. Using our example, imagine that the inverse model underestimates the amount of tea in the cup, while the forward model accurately describes it. The inverse model provides motor commands to the biceps which are too weak to lift the cup at the desired velocity. However, the forward model, receiving the motor command, begins predicting the slower velocity even before the arrival of error feedback. This prediction is compared to the desired trajectory, generating an error estimate, which is then used to increase the motor command to the biceps, increasing the velocity of the cup. In fact, the forward model need not be precisely accurate to be useful. An estimated state with some error can be more useful for control than an accurate measure of state with considerable delay. The simulations used in this thesis are based on forward models with inaccurate (non-learned) dynamics.



Now that a brief explanation of forward and inverse models has been provided, we describe a proposed model for control of the human arm. Bhushan and Shadmehr (1999) have proposed a controller for the human arm which employs a short-delay spinal reflex as well as a cortical loop with longer delay (see Figure 1.2). In this cortical loop, a forward model of arm dynamics compensates for the considerable system delays and provides an estimate of future arm state for comparison with the desired trajectory, producing an estimated error.

A simplified equation for expressing the actual forward dynamics of the arm-robot system is

$$\ddot{q} = I_s^{-1}[S(q, \dot{q}, \dot{q}^*(t)) + R]$$

where q is the vector of arm joint angles, \dot{q}^* is the desired trajectory of the arm, t is time, I_s is the inertia matrix of the arm, S is the torques on the arm caused by the dynamics of

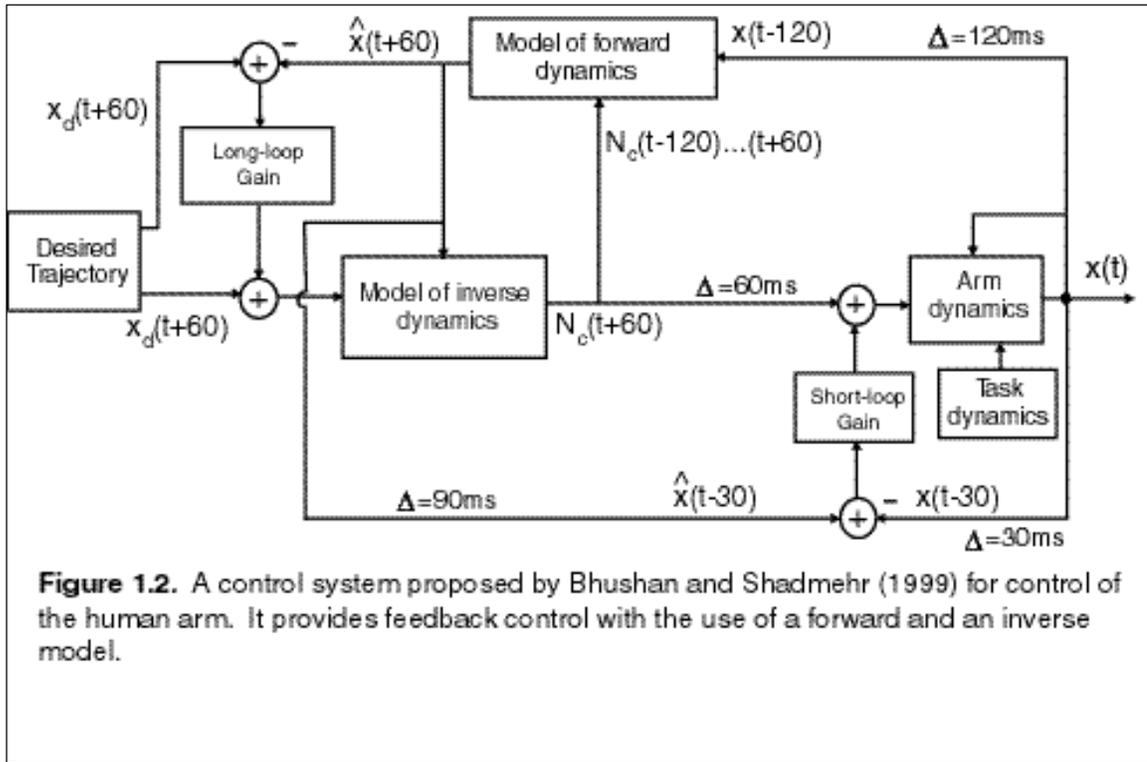
the arm including muscle activity, and R is the torques on the arm joints caused by the dynamics of the robot. R can be a function of q , any of its derivatives, t , or any other variable. The forward model which perfectly models the dynamics of the human arm, but not necessarily those of the robot, estimates the forward dynamics as

$$\hat{q} = I_s^{-1} [S(\hat{q}, \dot{\hat{q}}, \ddot{q}^*(t)) + \hat{R}]$$

where \hat{q} is the forward model's estimate of state, and \hat{R} is the forward model's estimate of robot dynamics. The assumption that S is the same for both equations implies that the forward model estimates the controller's reaction to error (the effect of spinal feedback). In fact, for the state estimate to be reasonable, this is necessary. While an accurate forward model would need to implement a calculation closely resembling the equation above, the neural implementation of such a computation is unknown.

The estimated state and error are inputted to the inverse model to provide feedback control. Simulations using this control strategy were compared to arm behavior in human subjects making point-to-point reaching movements while gripping a robotic manipulandum capable of applying perturbations. To produce trajectories which exactly follow desired trajectories, both the inverse and forward models must be accurate models of arm dynamics. However, simulations suggest that with an accurate forward model and inaccurate inverse model, arm trajectories which closely follow desired trajectories can be produced. These simulations were also shown to mimic sudden changes in arm direction, termed "near path-discontinuities", which appear in arm movements by human subjects when exposed to dynamics dramatically different than those previously learned. While these results support the existence of a forward model of dynamics for the control of the arm, more direct methods to test the forward model are desirable.

The simulation of Bhushan and Shadmehr will be referred to throughout the thesis. For a more detailed explanation of this model and its role in compensating for sensorimotor delays, please refer to their 1999 papers in *Biological Cybernetics* and *Advances in Neural Information Processing Systems*.



Eye-hand coordination

Numerous studies have examined the simultaneous behavior of the arms and eyes of subjects performing tasks which require looking, reaching, pointing, and/or grabbing. These studies have examined saccadic and smooth pursuit eye motion, eye-hand latency, eye-hand gain, and combinations of these and other factors. The most obvious outcome of all these studies is that arm motion affects the activity of the eyes. Due to the enormous size of this body of work, this section will focus on those results which are

most useful for understanding and interpreting the original research to be presented later in this thesis.

Whether saccading to a static visual target or pursuing a moving one, the behavior of the eyes is modified if movement of the subject's arm is performed simultaneously. In tasks where subjects are asked to saccade to and point at a visual target, saccadic response times are shorter than when a saccade is performed in isolation (Lunenberger, et al., 2000). Furthermore, if a subject is asked to point to a visual target and presented with a second saccade target after initiation of the arm movement, the eyes will continue to fixate the first target until completion of the pointing task (Neggers and Bekkering, 2000). Studies such as these suggest that, during manual action, eye behavior is tailored to improve manual performance.

In studies where subjects are asked to visually pursue a moving target, performance is drastically improved if the target is under manual control by the subject or correlates with a simultaneous manual task. Vercher, et al. (1996) reported that subjects had a 150ms latency in initiating pursuit of an externally controlled visual target, 130ms latency for a target correlated with passive movement of the arm, and a 5ms lead in pursuing self-moved targets. Maximum pursuit velocity of a visual target is more than doubled if the target is under manual control (Gauthier, et al., 1988).

By varying the temporal relationship between simultaneous visual and manual tracking tasks, the accuracies of the visual task and the manual task can be affected. If the visual target is a delayed representation of hand motion, pursuit performance decreases as the delay increases (Vercher and Gauthier, 1992). Miall, et al. (2001) used a paradigm in which the appropriate manual behavior to perform a particular task was

correlated with the movement of a visual target and showed that the performance in the manual task varied with the temporal offset between the two. Another study involving visual pursuit of a manually controlled target found that when a frictional perturbation was applied to the hand motion, control subjects suffered reduced smooth pursuit gain while deafferented subjects did not (Scarchilli, et al., 1999).

These studies and numerous others demonstrate several interesting characteristics about the natural coordination between eye and arm movements. First, eye motion and timing is altered when arm motion is simultaneously performed, and vice versa. Second, the greater the correlation (in a spatio-temporal sense) between eye and hand tasks, the better the performance of both. Third, the eyes utilize information about the arm motion. It is clear that this information includes both arm sensory signals and some signal related to the planned motion of the arm. This second signal is generally thought to be a copy of the arm motor command, termed *efferent copy*.

Forward models and visual tracking tasks

Models describing eye-hand coordination have been proposed, but remain, in general, vaguely defined (for example, see Gauthier, et al., 1988). For eye motion to track hand motion with zero or negative temporal delay, the eye controller must have access to an estimate of current or future hand state. As we have described, this estimate is the natural output of a forward model of arm dynamics. We hypothesize that a forward model of arm dynamics influences the motion of the eye during visual tracking.

This hypothesis is not simple to test. If arm movement is made without error, the forward model output is indistinguishable from the desired trajectory. In this case, the eye can track the hand by performing a coordinate transformation of the desired trajectory from hand coordinates to eye-head coordinates. On the other hand, the fact that the eye responds to sensory feedback from the arm is not, in itself, supporting evidence for the hypothesis. In fact, it is possible that a weighted combination of desired trajectory and delayed sensory feedback could be used to guide eye motion during tracking. Unless this information is combined through a model of plant dynamics, it cannot be called a *forward model* in the sense that we intend.

In the set of experiments described in this thesis, subjects performed point-to-point reaching movements in the absence of visual feedback while attempting to visually track the perceived location of their hands. Due to the lack of visual feedback, subjects generated eye trajectories which were predominantly saccadic with only brief periods of smooth pursuit. As will be shown, saccades tended to lead hand motion by more than one hundred milliseconds. This eye-hand prediction is significantly greater than the 60ms prediction of the forward model in Figure B2. However, if the forward model has

access to cued motor commands yet to be sent, it can predict with arbitrary lead. It is also possible that separate forward models of arm dynamics might exist: one for controlling the arm and another for controlling the eye, each with different prediction times. This possibility will not be explored further in this thesis.

Chapter 2. Methods

Experimental setup and data acquisition

Subjects made 10cm reaching movements while gripping a handle attached to the end of a two-degree-of-freedom robotic manipulandum capable of moving in the horizontal plane. This manipulandum is similar to a human arm, having an “upper” and “lower” segment, and “shoulder” and “elbow” joints. Two motors at the base of the arm were used to independently apply torque to the two joints of the manipulandum. Optical encoders were used to acquire position and velocity of each motor. A PC at a rate of 100Hz performed torque motor control and data collection. See Figure 2.1, below.

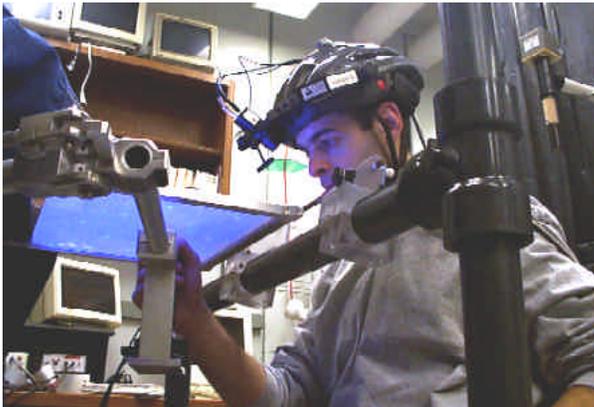


Figure 2.1. Experimental setup with the subject gripping the robotic manipulandum. A drape (not shown) surrounded the horizontal screen to prevent the subjects’ vision of their own arms. A bite bar was used for head fixation, while a helmet-mounted camera recorded eye position. A LCD projector (outside the field of view, above the subject’s head) projected visual stimuli onto the screen.

Visual stimuli were projected from a Philips ProScreen 4600 LCD projector, suspended above the head of the subject, onto a horizontal screen suspended directly above the robot workspace. A cursor, indicating robot handle position, was backprojected onto the screen by a red LED mounted to the top of the handle. The same PC controlled the robot, projector, and LED.

Gaze position was measured using the SMI iView system with Polhemus head tracker. This system consists of a video camera mounted to an adjustable bicycle helmet. The subject viewed the world through a pane of glass which also reflected an image of the eye to the camera. In addition, the camera imaged the reflection of a near infrared light source aimed at the eye. The iView software used the position of the pupil and the near IR source to determine the rotational position of the eye. The Polhemus transmitter was mounted to the top of the bicycle helmet and communicated to a fixed receiver mounted to a stand behind the subject to detect head position and orientation in space. After calibrating the system by having the subject look at a series of points in the robot workspace, the iView software computed gaze position in the workspace which is recorded at 100Hz by a PC synchronized to the robot's PC.

To minimize head movement and increase the length of time before recalibration became necessary, subjects used a bite bar during calibration and movement sets. In addition to the iView system calibration, a second calibration was used to determine the transformation from iView gaze coordinates to robot coordinates. This was performed as needed during the experiments by asking the subjects to visually track a moving target which visited a series of points of known location within the workspace. Over repeated trials, the existing calibration became increasingly inaccurate and the system was recalibrated periodically. Figure 2.2 shows the calibration error in fixating the movement origin for a single subject over 45 trials in the Assistive-Kick-Assistive/Null paradigm (to be described later). Generally, loss of calibration over repeated trials is thought to be due to slipping of the helmet or the subjects becoming tired, causing the eye to close slightly

so that eyelashes or eyelids blocked the cameras view of the pupil. Periodically readjusting the helmet and camera and recalibrating the system corrected these problems.

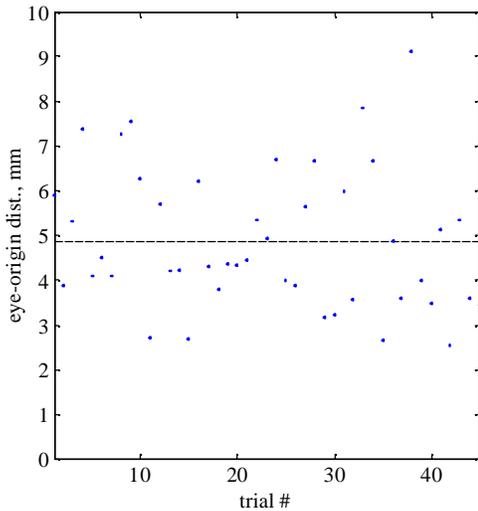


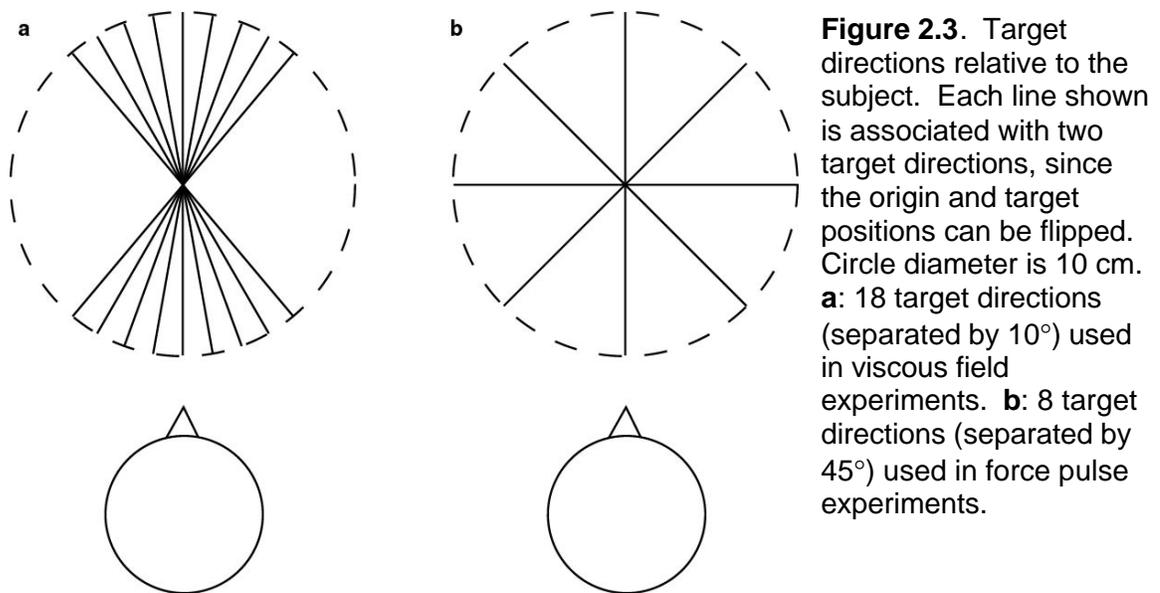
Figure 2.2. Calibration errors for visual fixation of movement origins for 45 consecutive movements in a single subject performing a force pulse task. Dashed line indicates mean error. With proper helmet adjustment and initial calibration, accuracy could be maintained over many trials, as shown here.

Experiments

Subjects sat, gripping a handle on the end of a two-degree-of-freedom robotic manipulandum capable of moving in the horizontal plane. A horizontal screen surrounded by a drape and suspended just above the plane of motion blocked the subject's view of the robot motion. Visual stimuli and cues were projected down onto the screen from a LCD projector above the subject's head. A red LED attached to the robot's handle was used to back-project (from below) a cursor, indicating the handle position in the workspace.

Before each movement, the robot moved the subject's hand to the start point (origin) of the movement, indicated by a marker on the screen. The red handle cursor was also illuminated at this time. After the subject fixated the origin for 0.5 seconds, the origin marker and robot cursor were extinguished, and the target cursor was illuminated.

After the subject fixated the target cursor for 0.5 seconds, it was extinguished and the origin and handle cursor were illuminated again along with a backdrop of random dots which covered the workspace. With the subject looking at the origin, the origin was extinguished, cueing the subject to begin the movement to the target. Except in the passive movements, subjects were instructed to move the robot handle to the remembered target location while looking at the perceived location of the hand with their eyes. All origins and targets lie on opposite sides of a 10cm circle (not visible to the subject), with its center in front of the centerline of the subject's body, as shown in Figure 2.3.



A successful movement was indicated to the subject by explosion of the target and a pleasing sound. Successful movements were movements which acquired the target within set windows in movement space (a final distance from the target which was adjusted as a function of subject performance) and time (mean of 1s) or in space alone. If a time criteria was used, a change in target color indicated whether the movement was too

fast or too slow. Time was counted from movement onset, when the hand speed exceeded 30mm/s. “Success” of a trial was not a criterion for later analyses, but was used as an incentive for subjects to have consistent desired trajectories.

In all experiments, subjects visually tracked their hands without any visual feedback (nvfb) of hand position during the movement. Subjects were instructed to follow the perceived locations of their hands with their eyes throughout the movement. The ability to satisfactorily track the hand took repeated practice. All subjects were trained in the null field to eventually track their hands without visual feedback. Generally, subjects were taken through a process of visual tracking with full visual feedback, then intermittent visual feedback (a flashing cursor at 5 Hertz), and, finally, no visual feedback. In the nvfb condition, subjects made eye movements which were predominantly saccadic. For nvfb movements, the handle cursor was extinguished when the handle exceeded a velocity threshold of 30mm/s, indicating movement initiation, and was illuminated only after the movement ended. Subjects perceived that the cursor was extinguished at the instant of movement onset.

Subjects were both male and female with ages ranging between 22 and 37. Some subjects were naïve, while others had participated in previous studies using the manipulandum. Most subjects participated in several of the experiments described here. Because experiments were designed to avoid any motor learning, the naïveté or experience of subjects was assumed to be irrelevant to the results.

Individual trials were included in the analysis if they met two criteria. First, the calibration of gaze point in the workspace had to be acceptable. This was determined by measuring the position of the gaze point during the periods where subjects looked at the

origin and the target before the movement. The measured gaze point was required to be within 7.5mm of each marker for at least 0.5 seconds. Secondly, movements with false starts were not accepted. These were cases when the subject began the movement prematurely (before the origin cursor and visual feedback were removed) and was required to return to the origin and begin the movement again. Roughly 1/3 of all movements were excluded from analysis based on these two criteria, mainly due to poor calibration of the eye.

In what follows, two groups of experiments will be described: viscous field experiments and force-pulse experiments, summarized in Table 2.1, below. In descriptions of both types of experiments, any reference to a force field means a velocity-based (viscous) field in which forces are applied to the robot handle as a function of the velocity of the handle, following the equation $\underline{F} = B\underline{\dot{x}}$, where \underline{F} and $\underline{\dot{x}}$ are 2x1 vectors and B is a 2x2 matrix. The “null” field is the condition in which the robot is passive (B is the zero matrix). In the force-pulse experiments, the terms “pulse” and “kick” are used interchangeably and refer to the unexpected and brief application of forces to the robot. All references to “eye position” mean the point of gaze within the workspace plane. Hand and eye positions are always described in Cartesian space where the target direction is the direction of the movement origin to the target, and the perpendicular direction is 90 degrees clockwise from the target direction. The words “parallel direction” will be used synonymously with “target direction”.

Table 2.1. Summary of experiments

Viscous Field Experiments		subjects
Straight Null Field		GM, HJ, PH, PG, AR, SW
Curved Null Field (path tracing)	curved left; curved right	GA, GD, RS
Mixed fields experiments	assistive; resistive; CW curl; CCW curl; null	GA, RS, JW, MD, AB, PG, GM, HH, LH
Passive movements	curved left; curved right	RS, AR, SW
Force-Pulse Experiments		subjects
Null field, no-kick sets		AR, MD, PG, RS, SW
Null-Kick-(Null/Assist)	Null-Kick-Null; Null-Kick-Assist; No-Kick	PG, GA, RS, AR, MD, SW
Null-Kick-(Null/Resist)	Null-Kick-Null; Null-Kick-Resist; No-Kick	GA, RS, SW, AR, PG
Varied pulse times	No-Kick; Null-Kick-Null at 200, 250, 300, 350ms	GA, RS, AR, SW
Passive movements	No-Kick; Kicks at 200, 250, 300, 350ms	GA, RS, AR, MD, SW

I. *Viscous field experiments*. These movements took place in 1 of 18 directions separated by 10 degree intervals, centered on the subject's body centerline (as shown in Figure 2.3a). No force pulses were applied during these experiments.

- a. *Null field sets*. Subjects made a series of movements in which no external forces were applied by the robot during the movement. Two types of null field sets were performed: straight and path-tracing. *Straight sets*: Subjects generally moved in a straight path from origin to target. 6 subjects participated, generating a total of 100 trials for analysis (after removal of trials with poor calibration and/or false starts). *Path tracing sets*: Subjects were shown a curved path to trace (approximating the paths effected by a viscous

curl field) which appeared and disappeared along with the target marker preceding movement onset. The displayed path alternately deviated to the left and right of the straight-line movement on consecutive trials. Subjects were asked to trace the remembered path, ending at the target. 4 subjects generated 300 trials.

- b. *Mixed fields experiment.* Subjects made a series of movements in which, from one trial to the next, any one of five velocity fields was pseudorandomly applied: assistive (pushing in direction of motion, $[10\ 0; 0\ 10]$), resistive (against motion, $[-25\ 0; 0\ -25]$), CW curl (perpendicular to motion $[0\ 25; -25\ 0]$), CCW curl ($[0\ -25; 25\ 0]$), or null. These fields were applied with probabilities of 25%, 25%, 12.5%, 12.5%, 25%, respectively. 9 subjects generated 479 trials.
- c. *Passive movements.* Subjects were instructed to hold the handle and allow the robot to move the hand. The robot then followed a pre-programmed trajectory starting near the origin and ending near the target. (Motor torques for passive movements were generated from position, velocity, and acceleration data of actual perturbed arm trajectories, “replayed” using an inverse model of robot dynamics and error feedback.) Visual feedback was displayed as in all other experiments, and the trajectories were similar to those in the “active” experiments, with movements beginning at the origin and ending near the target. The paths simulated either null or curl field movement (in either direction) and were given in a pseudorandom order, with equal probability of being straight, right-curved, or left-curved. 3 subjects generated 54 trials.

II. *Force-Pulse Experiments.* Targets were given in 8 directions separated by 45 degree intervals, encompassing the full 360 degrees of the circle in a pseudorandom order. Force pulses (also termed *kicks*) consisted of a 25 Newton force superimposed on an assistive field of $[10\ 0; 0\ 10]$ and lasted 50ms. The 25 Newtons were applied perpendicular (either -90° or 90°) to the target direction. (One subject, PG, did not receive the superimposed assistive field on kicks during the Null-Kick-(Null/Assist) experiment.)

- a. *Null field, no-kick sets.* Subjects made a series of movements in which no external forces were applied by the robot during the movement. These sets (consecutive null field, no kick trials) were performed either immediately preceding or following varied pulse-time experiments or passive pulse experiments. 5 subjects generated 166 trials.
- b. *Null-Kick-(Null/Assist).* Subjects made movements, starting in the null field, 2/3 of which received a force-pulse perpendicular to the direction of movement 200ms into the movement, and 1/3 of which did not in a pseudorandom manner. If a force-pulse was given, there was a 50% chance of returning to the null field after the kick and a 50% chance of having an assistive ($[10\ 0; 0\ 10]$) field imposed. If no pulse was applied, the dynamics remained null field. 6 subjects generated 722 trials.
- c. *Null-Kick-(Null/Resist).* Same as (b), except that subjects received force pulses in 50% of the trials and had a 50:50 chance of being exposed to either a null field or a resistive ($[-20\ 0; 0\ -20]$) field following a kick. 5 subjects generated 657 trials.

- d. *Varied pulse times.* All movements were made in the null field either with or without a pulse. Each movement had equal probability of having a pulse at 200ms, 250ms, 300ms, 350ms, or not at all. 4 subjects generated 456 trials.
- e. *Passive pulse.* Pulse timing was varied as in (d), except that subjects were instructed to hold the robot lightly as the robot moved through a trajectory. These trajectories simulated human movements with and without force pulses. (The robot generated trajectories with a stiffness field, attracting the handle to the target, and kicks were created by a superimposed 25 Newton pulse.) 5 subjects generated 531 trials.

Analysis

In all experiments, subjects generated predominantly saccadic eye trajectories. Due to high variance in both saccade timing and placement, we found it most informative, in general, to look at averages over large numbers of similar movements. To increase the number of comparable movements within each experiment, both eye and hand trajectory data were rotated in x-y space according to the target direction so that, in the new coordinate system, all movements had overlapping origins and targets. Furthermore, for movements which inherently had a negative displacement perpendicular to the target direction (due to curl fields, curved path tracing, or force pulses) were “mirrored” to become positively displaced along the perpendicular axis. This is possible because fields in opposing directions generate similar (but mirrored) trajectories, as shown in Figure 2.4. Using these two coordinate transformations, it was possible to combine more data

for analysis under a given trial type, regardless of the direction of imposed perpendicular displacement.

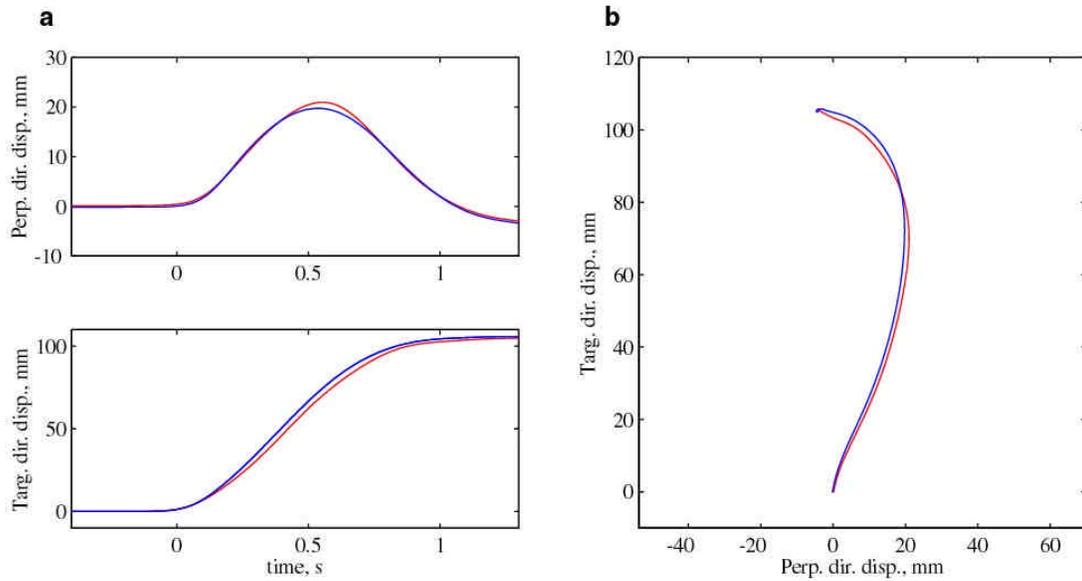


Figure 2.4. Comparison of average trajectories in clockwise (CW) and counterclockwise (CCW) viscous force fields from mixed fields experiment. Blue = CW. Red = CCW. Trajectories are shown in transformed coordinates. **a:** Parallel and perpendicular components of trajectories versus time. **b:** Movement paths in x-y space.

One method used to average eye movements was to simply take the mean gaze position (in x-y space) at each data sample. While individual eye trajectories tended to be saccadic, making them “step” in time, the average trajectory was a much smoother signal. A computer algorithm was also used to detect saccade startpoints and endpoints in eye trajectories from individual trials (see Appendix for a description). With these saccades detected, it was then possible to take the mean startpoint and endpoint for all saccades falling within a certain bin of time, as well as to take the mean time and positions for the first one or two saccades across multiple movements. The graphical representations of the data use a combination of these techniques in comparing eye and hand trajectories. Figure 2.5a shows an example null field trajectory. Figure 2.5b shows the average null

field trajectory and all saccades from null field movements in the viscous field experiments. As can be seen, saccade positions are highly variable. By averaging the saccade start and endpoints within bins of time, it is possible to estimate typical behavior.

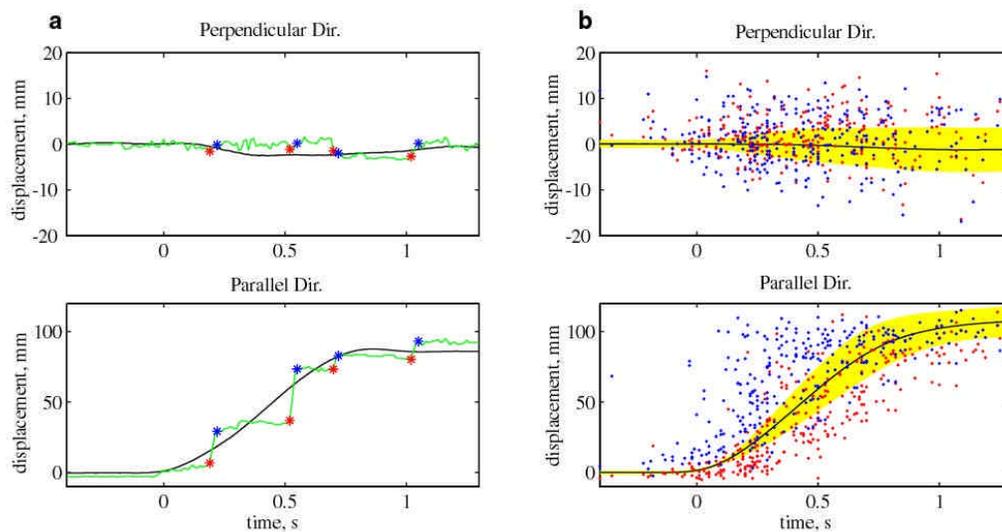


Figure 2.5. Null field trajectories. Black = hand trajectory. Green = eye trajectory. Red = saccade startpoint. Blue = saccade endpoint. **a:** Example trajectory. **b:** Average hand trajectory shown with saccade start and endpoints from all null field trials in viscous field experiments. Standard deviation about the mean for hand trajectories is shown in yellow.

The most common analyses applied to the data (as presented in Chapter 3) will be described here. Instead of using example figures in this section, the text will refer to figures in Chapter 3 where the actual results are discussed. It may be useful to refer to the following paragraphs while reading Chapter 3 for a better understanding of the analytical techniques applied.

Saccade averaging.

To calculate average saccade start and endpoints, saccades were grouped into 30ms wide bins with centers at 10ms intervals. Within each bin, all saccade start and endpoints were averaged. For an example of this type of data representation, see Figure 3.1a. Accompanying these plots is a histogram showing the estimated likelihood of a saccade occurring within each 10ms interval, also filtered by averaging over the 30ms wide bins. (This likelihood is calculated by dividing the number of saccades within each bin by 3 times the number of trials contributing to the figure. For a probability density, in probability/ms, divide this result by 10.)

Average saccade times, startpoints, and endpoints for the first and second saccades of trials were also calculated. Saccades are counted starting 400ms before movements onset (after origin and target fixations have ended). Only movements with at least two saccades are included in the averaging. Figure 3.1c shows the average first two saccades in straight null field movements, as well as the average eye trajectory.

Three-dimensional saccade prediction plots.

If the eyes are believed to be an indicator of past, present, or future hand position, it is of interest to know how saccades at any instant in the movement relate to hand position throughout the movement. For this reason, the average saccade endpoint for each time (in 30ms wide bins, as before) is compared to the average hand positions within a range of times. This type of analysis is presented in a three-dimensional figure (Figure 3.2a, for example), with the time from movement initiation on the y-axis and each x-axis value representing the time relative to the saccade. The third dimension, color, represents the

distance between the saccade endpoint position at the y -axis value and the hand position x milliseconds from that time, where x is a value on the x -axis. For example, if the color at ($x = 120\text{ms}$, $y = 550\text{ms}$) represents 0mm , this is interpreted as meaning that at 550ms the average saccade ended at the location that the hand would visit 120ms later (at 670ms). By finding the minimum eye-hand distance for each y -value (as in Figure 3.2c), the “prediction time” of the eye can be estimated. This prediction time will be seen to change as the movement progresses and can be approximated by a constant value up to some time along the y -axis, followed by a quadratic relating prediction time to movement time. Nonlinear fitting is used to optimize the constant-to-quadratic transition time and the parameters in each segment of fit.

Hand-versus-eye perpendicular position plots.

Once the “prediction time” of the average saccades for the average hand has been established, we might want to ask how well the saccade placement in a given movement matches the hand position for the same movement. This will prove to be particularly interesting for the perpendicular displacement of force pulse trials. A range of times is chosen and the perpendicular components of each saccade within this range are plotted (on the x -axis) against the hand position (on the y -axis) for the same movement at the time for which the saccade is estimated to predict. Figure 3.17 shows such a plot, along with a linear fit. The slope and quality of fit indicate how well the eye is predicting hand position for individual movements.

Chapter 3. Results

Null field: Straight and path-tracing (curved) movements

Null field trials are from sets of unperturbed movements, when subjects' movements are likely to closely approximate their desired trajectories. In these movements, eye trajectories approximate hand trajectories, generally leading them in time. As shown in Figure 3.1a through d, the earliest saccades originate around the time of movement onset and end close to a future position of the hand. Path tracing (curved) trials were, on average, slower than straight trials, and are not completed within the time shown in these plots.

Typically, multiple saccades are made for any one movement (3.5 ± 1.1 , mean \pm sd for straight trials), with each saccade beginning near the hand, and ending along the future hand path. This can be seen clearly in Figure 3.1c and d, where the average time and positions of the first two saccades are shown. Shown in Figures 3.1a and c, the hand eventually surpasses the eye position, resulting in an eye endpoint closer to the movement origin than the hand's endpoint. As we will see, this occurs for nearly all movement types.

Figure 3.2 is a representation of the distance between averaged saccade endpoints and hand positions at different moments. Color represents the distance between the endpoint of the average saccade at y ms and the hand x ms relative to time y . (See the Analysis section in Chapter 2 for a more detailed description.) The minima (for each saccade time) early in the movement in Figure 3.2d (from 0ms to about 400ms along the y -axis) suggest that the eye is attempting to predict, on average, roughly 240ms into the

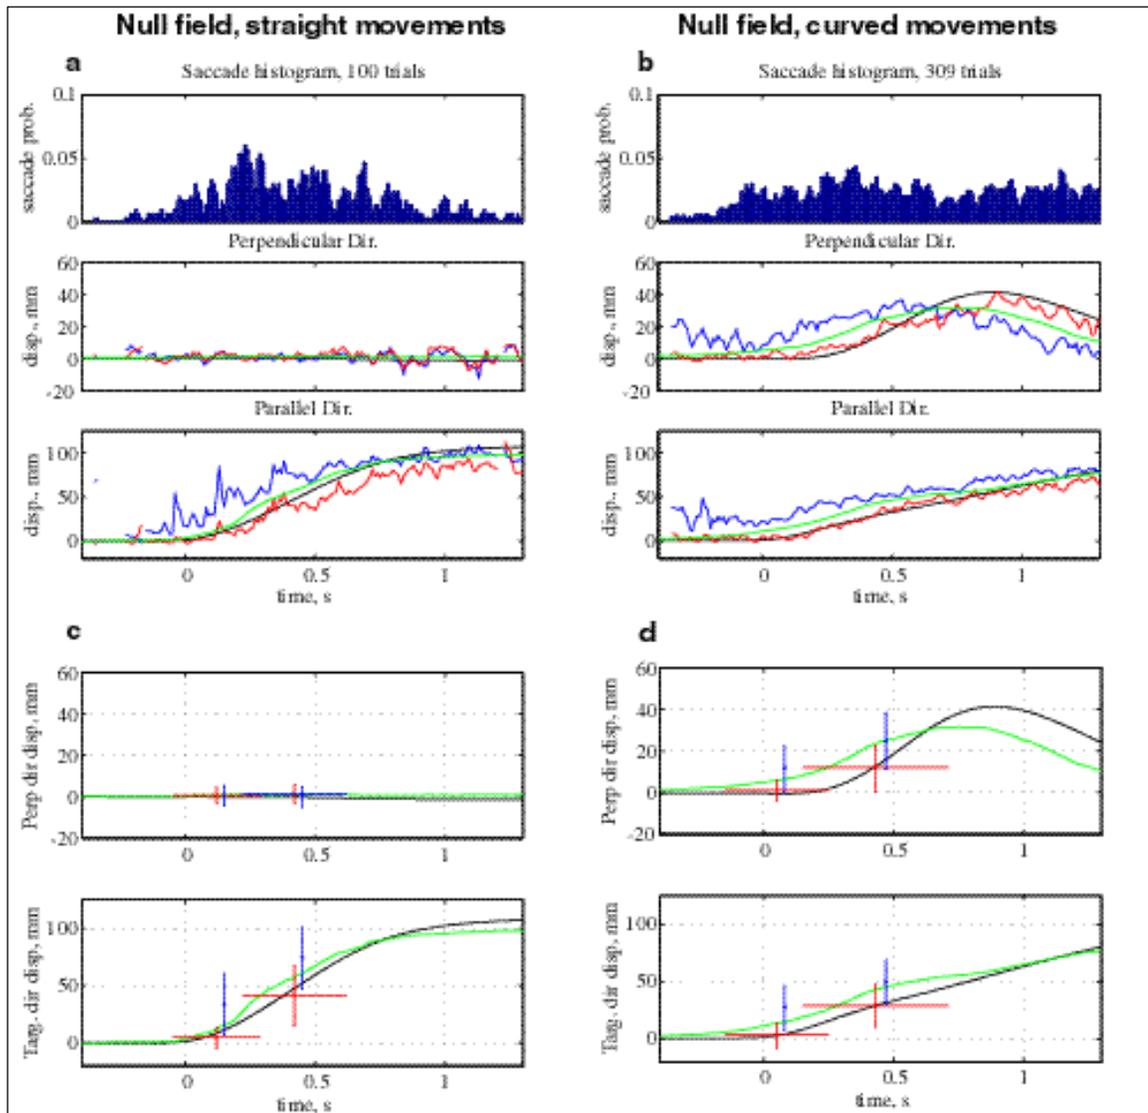


Figure 3.1. Hand and eye averages for movements in null field sets. red = saccade start; blue = saccade endpoint; green = average eye trajectory; black = average hand trajectory. **a, b:** Saccade histograms, along with average saccades in 30ms bins at 10ms intervals. **c, d:** Averages of first two saccades shown with standard deviation bars in time and position. Standard deviation in time is shown only for saccade start points, since the same variance exists for endpoints.

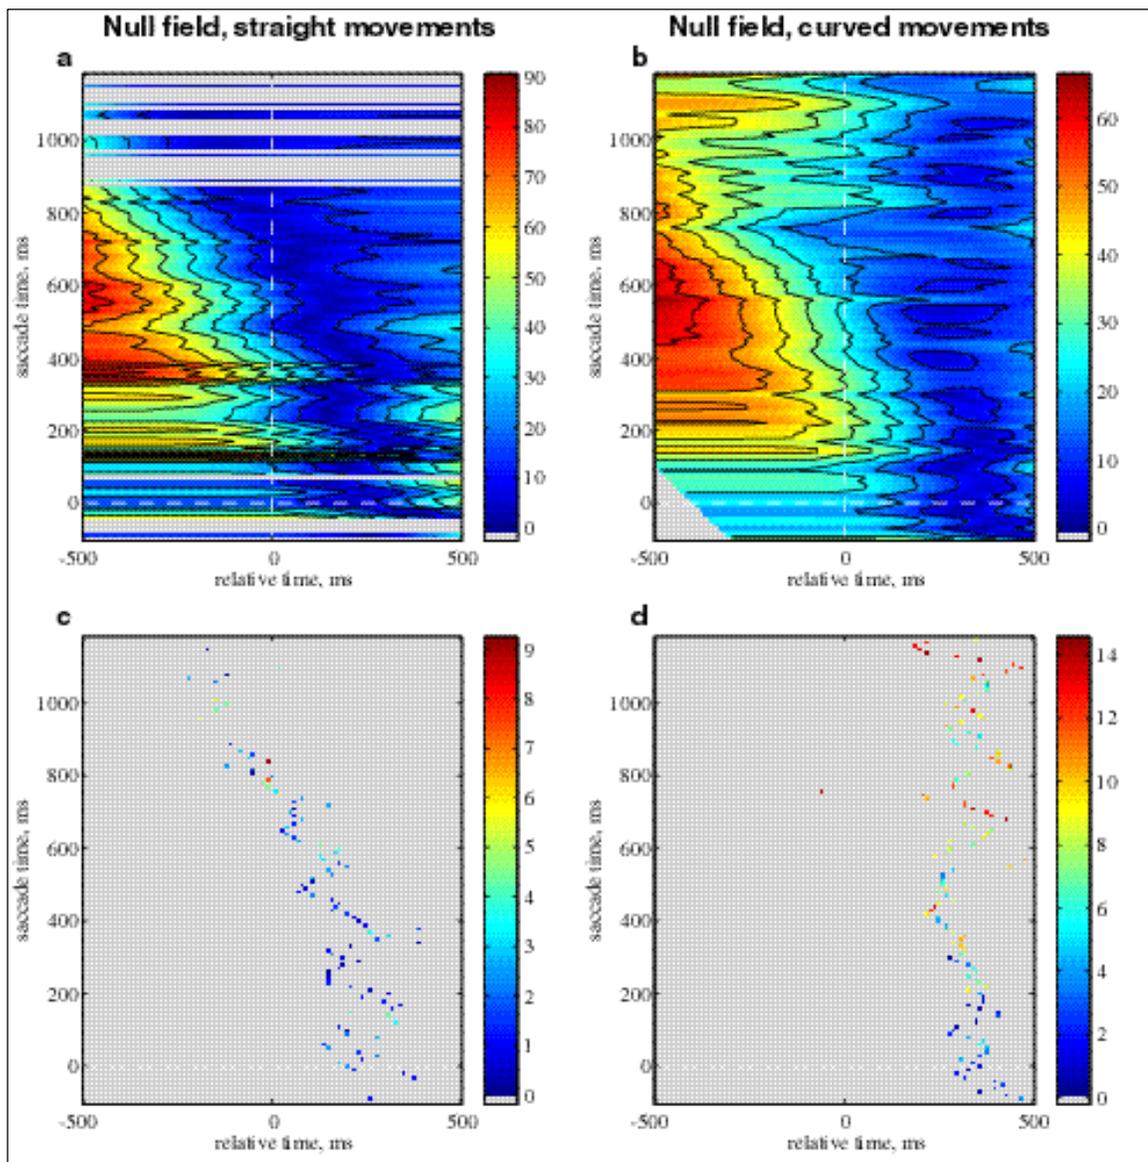


Figure 3.2. **a,b:** Eye-hand distance for average of saccade endpoints and mean hand trajectories. Y-axis represents time of saccade. X-axis represents time along hand trajectory relative to saccade. Color represents distance between saccade endpoints and hand position, in mm. Gray regions represent lack of data due to insufficient saccades in bin (less than 2) or no data collected. Saccades collected in 30ms wide bins, separated by 10ms intervals. **c,d:** Minimums of each row in a and b; exclusion of other points.

future. For the curved movements, this prediction time is closer to 330ms. The growing underestimation of hand position in the target direction described before is represented here by the decreasing prediction times as the straight movements progress. The decreased prediction time occurs later in the slower, curved movement so that it is not captured in this figure.

Mixed-fields: Assistive, resistive, CW curl, CCW curl, and null

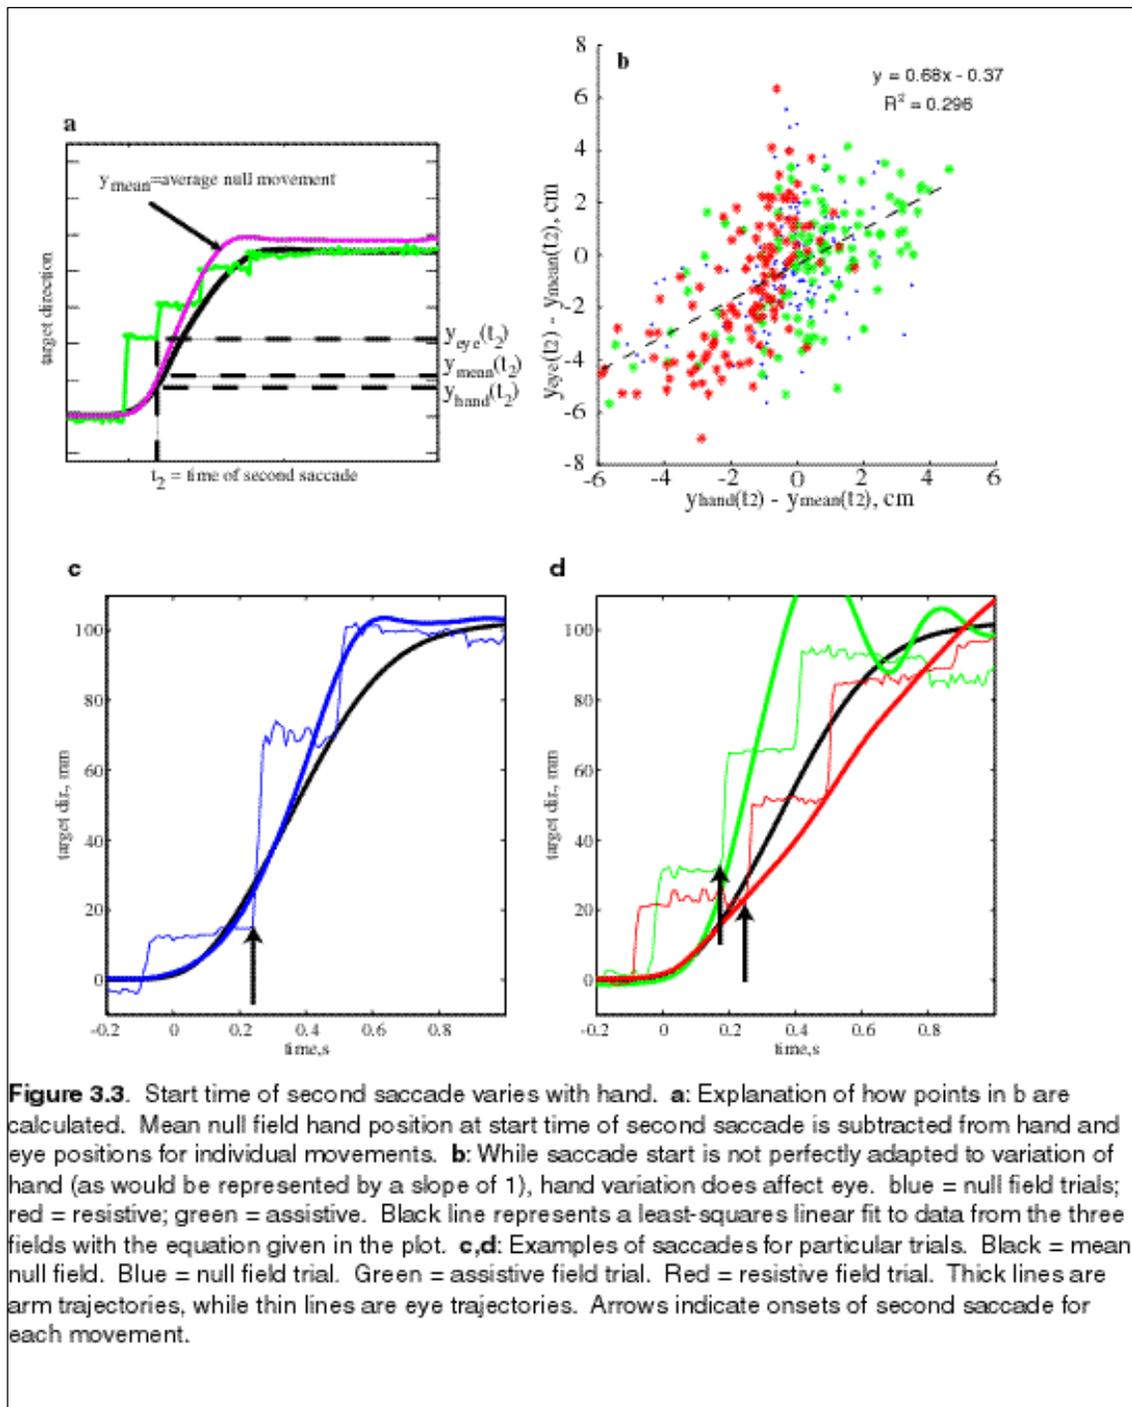
The data from null field sets demonstrates that the average eye movement closely predicts the average hand trajectory in the null field. By applying perturbations to the movements, we can examine how the eye movement is affected by error feedback from the hand. In these trials, subjects received perturbations in five velocity fields presented in a pseudorandom order. Due to the randomness of field presentation, subjects were unable to anticipate the field and made large movement errors.

As was discussed above, the averages of saccade start points are close to the mean hand trajectory for unperturbed movements (Figure 3.1), but are individual eye trajectories related to the hand trajectory for the same trials? Fields which are assistive (in the direction of movement) and resistive (against the direction of movement) will cause the hand to move faster or slower than expected. Figure 3.3a demonstrates a method for comparing the timing of the second saccade in a given movement to the hand behavior in that same movement. This is performed by subtracting the mean hand trajectory for null fields (the assumed desired trajectory) at the time of the second saccade from the eye position at that time and comparing this value to the hand position for that movement with the same mean trajectory subtracted. Figure 3.3b shows the result for

these calculations, demonstrating that deviations of the hand from the mean trajectory coincide with deviations of saccade start points for the same movements. The slope of the fit is positive but less than unity, indicating that the eyes do respond to hand error but are also influenced by the expected hand movement. The small y-intercept value supports the idea that when the hand trajectory is close to the mean trajectory, the eyes saccade shortly after the hand passes the gaze point. Figures 3.3c and d show examples of null, assistive, and resistive trials where the startpoint of the second saccade approximate hand position. Figures 3.4 and 3.5 show the hand and eye averages for null, assistive, and resistive fields for this experiment. The average startpoints of the second saccades for assistive and resistive fields show their inability to completely compensate for movement error by lagging or leading the hand more than in the null field trials.

If the intended movement is straight, any perpendicular displacement is due to error. Because curl fields introduce significant perpendicular displacement to the direction of intended motion, they offer an excellent opportunity to observe saccade placement as a function of error. In these fields, the first saccade occurs close to movement onset, at a time when no perturbation is expected or sensed, and therefore has almost no perpendicular component. The second saccade, however, begins to reflect the error with a perpendicular component in the direction of error but still underestimates the magnitude of the perpendicular error (see Figure 3.6).

Eye behavior in the curl field has an interesting characteristic which is consistent with the idea that a forward model of the arm drives eye behavior. A simulation based on the control strategy described in Chapter 1 (Figure 1.2) was used to generate Figure 3.7 and 3.8b, which show the simulated hand position and forward model output for a



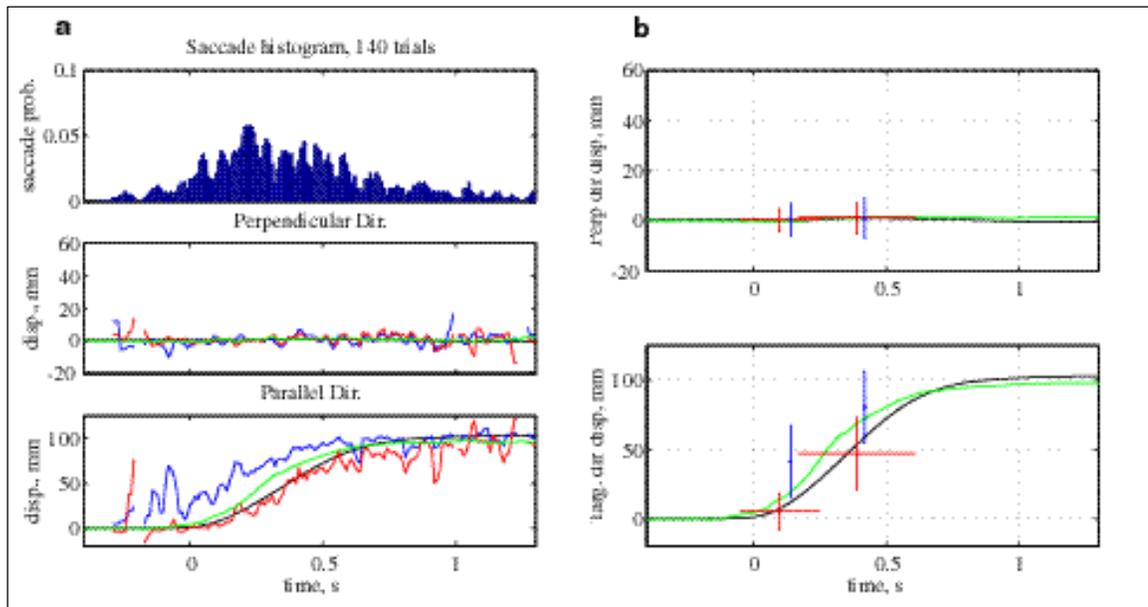


Figure 3.4. Saccade and hand averages for null field trials in mixed fields experiment. Saccades are distributed across the movement time and predict future hand state.

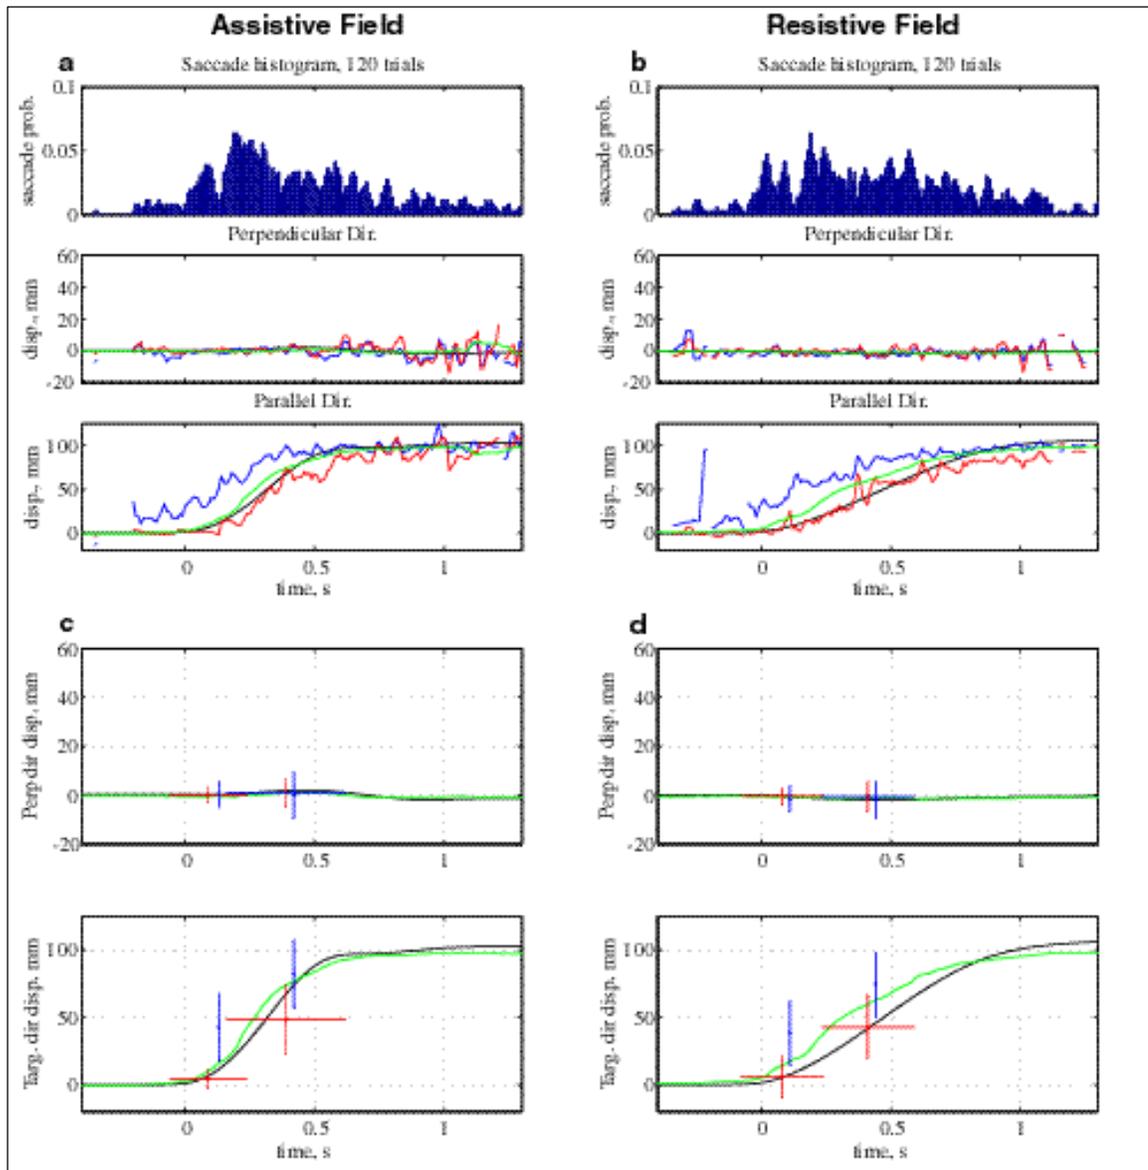


Figure 3.5. Hand and eye averages for movements in assistive and resistive fields. **a, c:** Saccade histograms and averages for assistive field trials. **b, d:** Saccade histograms and averages for resistive field trials.

movement in a non-learned curl field movement. In this simulation, the forward model's estimate of hand position underestimates the perpendicular component of position at all times. This behavior is also true of the mean eye trajectory in Figure 3.6. We can understand this if we imagine that the eye movement is a sampled representation of the continuous output of the forward model. The perpendicular component of each consecutive sample would be increasingly positive (away from zero) until about the time of peak perpendicular displacement, and then in the negative direction (toward zero) after the peak. This behavior is evident for the averaged saccades in Figure 3.6a. In the model, this behavior occurs because once error feedback is available, the forward model output is influenced by the predicted error correction. However, because the curl field dynamics are still "unknown" to the forward model, the prediction remains inaccurate. Figure 3.8 is a comparison in x-y space of actual movements and the simulation. For the human data, lines connect the start and endpoints of saccades with the hand path at the same time. Lines are used to connect the simulated forward model output to the simulated hand position at particular times. The lines demonstrate the similarity of the eye-hand relationship and the forward model-hand relationship in simulations.

Passive curl movements

Subjects were instructed to relax and allow the robot to move their arms, presumably resulting in arm motor commands which were decreased or less relevant to arm motion. As a result, eye movements should have been driven more strongly by error feedback than desired trajectory or arm motor command.

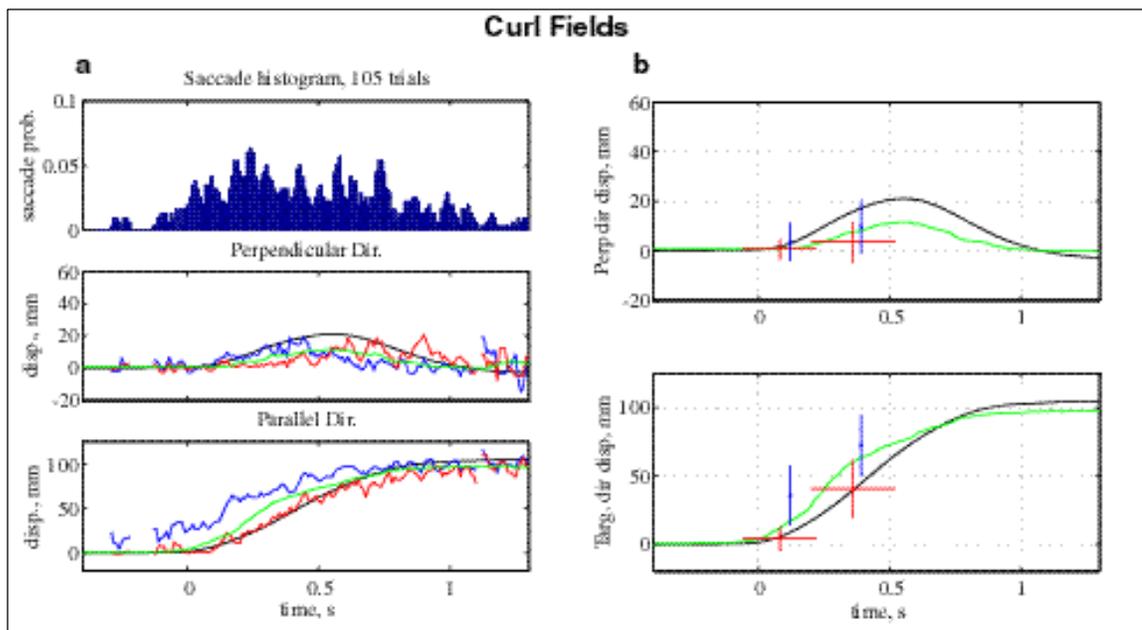


Figure 3.6. Hand and eye averages for movements in clockwise (CW) and counter-clockwise (CCW) curl fields. Values of perpendicular displacement for CCW fields are multiplied by -1 so that all movements have measured perpendicular displacement in positive direction.

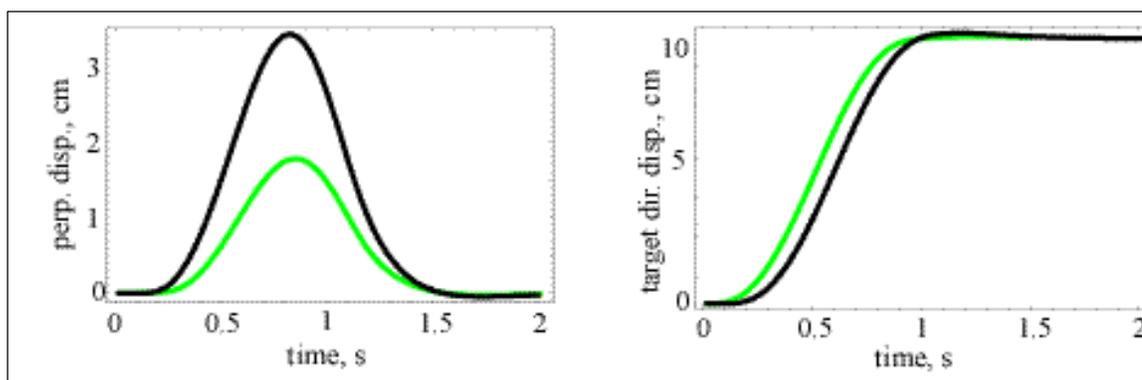


Figure 3.7. Output of the forward model from a simulation of a curl field movement based on the control strategy described in the Background section. green = forward model output; black = hand position. Compare to the mean eye and hand in Figure 3.6b.

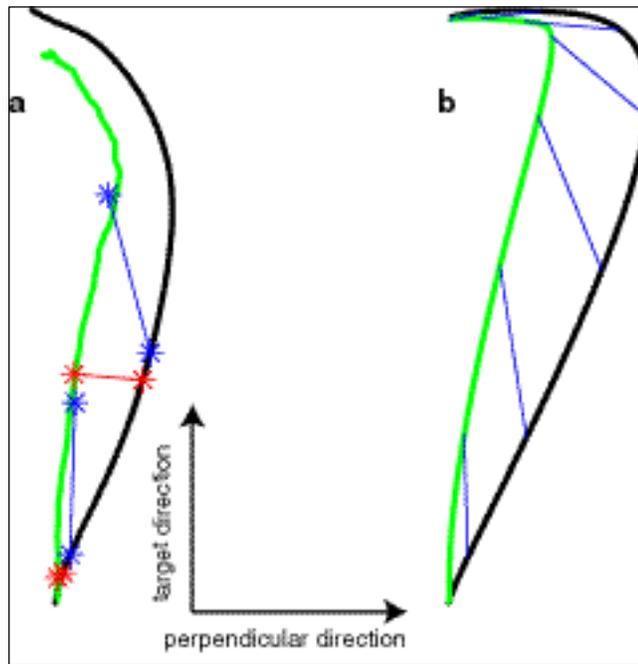


Figure 3.8. Qualitative comparison of eye behavior and simulated forward model output for curl field trials. **a:** Average hand path (black) and average eye path (green). Saccade startpoints and endpoints are shown with red and blue asterisks, respectively, for average of first two saccades. Lines connect saccades to hand position at the same time. **b:** Simulated hand path (black) with simulated forward model output (green). Lines connect forward model output to hand position at specific times.

Figure 3.9 shows mean eye and hand behavior for curved passive movements. As might be expected, few saccades occur near the onset of movement since onset cannot be predicted. There is a high occurrence of saccades roughly 150 milliseconds after movement onset. Both saccade averages and trajectory averages show eye movements which tend to lag the hand movement, although some prediction does appear to occur, as evidenced by the saccade endpoints which lead hand trajectory for some portions in Figure 3.9a. In Figure 3.10, the predictive ability of saccade endpoints in the passive case is compared to that in the curl field trials, using eye-hand distance minimums. While saccade endpoints are nearly always closest to future hand position in the curl field trials (Figure 3.10b), the minima are more scattered across time in passive movements and tend to be larger, on average, than those in the curl fields. These results indicate that while the eye may attempt to predict future hand position in the passive movements, the lack of a relevant motor command prohibits the eye from making consistent, accurate predictions.

Force-pulse movements: Null-Kick-(Null/Assistive/Resistive)

The purpose of the force-pulse experiments was to generate a condition in which, after the introduction of some error, the internal model of system dynamics (inverse and forward models) would be in agreement with the post-pulse dynamics. Before performing the force-pulse trials, subjects made a series of non-kicked movements in the null field, assuring that the null field was the “learned” field before the force-pulse movements.

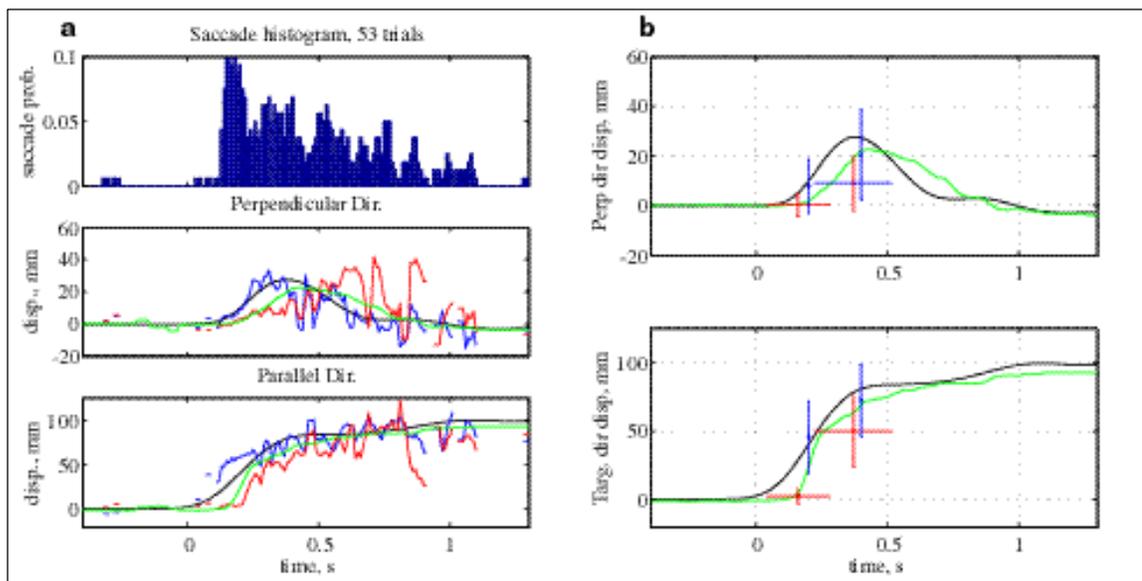


Figure 3.9. Hand and eye movements for trials under passive condition. Perpendicular displacement occurred in either direction pseudorandomly. Saccade endpoints tend to lag the hand in both the target and perpendicular directions.

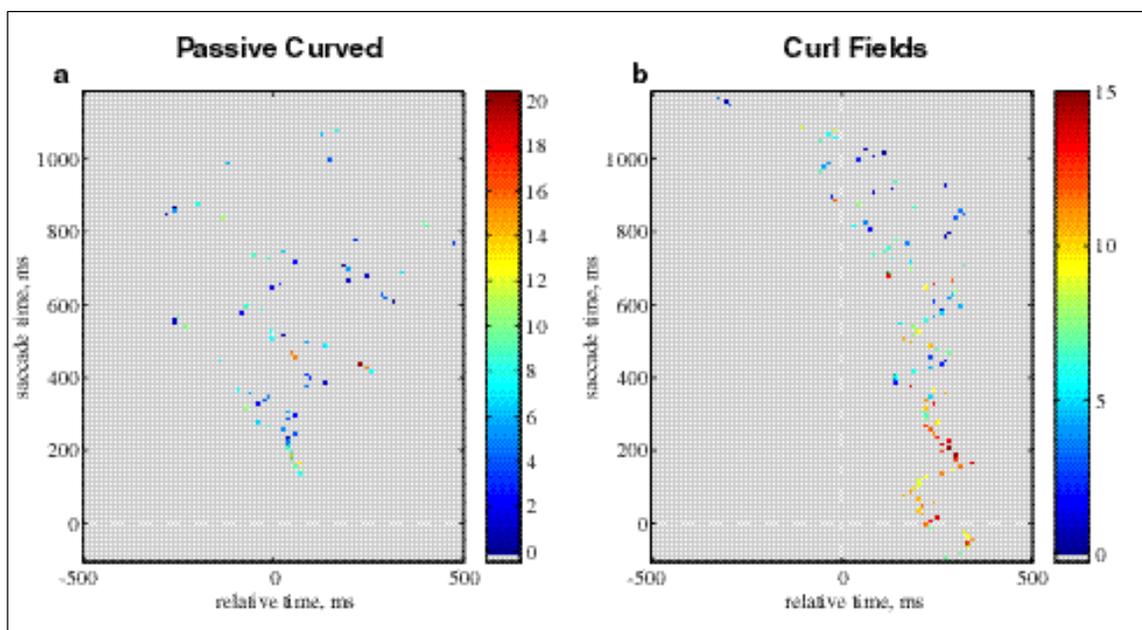


Figure 3.10. Eye-hand distance minima, in mm, for **a:** passive movements and **b:** curl field movements. While a consistently positive prediction time is indicated by the minima for curl fields, the minima are scattered across times for the passive trials.

For purposes of analysis, the No Kick and Null-Kick-Null trials from the Null-Kick-(Null/Assistive) and Null-Kick-(Null/Resistive) experiments have been averaged together. Figure 3.11 and 3.12 show the average eye and hand behaviors for No Kick and Null-Kick-Null trials, respectively. Because the dynamics in No Kick trials and null field trials from the viscous field experiment are identical, the behavior is also similar with saccades that consistently predict future hand position. For pulsed trials, the force-pulse is applied at 200ms and lasts 50ms, causing a sudden change in perpendicular displacement of the hand at 200ms. The eyes reflect this change sometime after the hand, and the saccade histogram shows a peak roughly 150ms following the kick (Figure 3.12a). As in curl field trials, saccade endpoints initially lag the perpendicular displacement and later lead as the hand returns towards the target. The mean eye trajectory, however, does not reflect this lead as clearly as in the curl field trials.

The eye-hand distance plots in Figure 3.12c and d show that the eye accurately predicts future hand position for the first 50 to 100ms of the movement and then becomes inaccurate. This is explained by the fact that the eye is unable to predict the kick onset at 200ms. Following the kick, the eye is inaccurate until shortly before 400ms, and then accurately predicts again for the rest of the movement, with a tendency towards negative prediction times as in No Kick and null field trials. The roughly 300ms gap in accurate prediction (from 100ms to 400ms) is due to the additive effects of the eye-hand prediction time and the sensorimotor delay for eye response to hand perturbation.

Figures 3.13 and 3.14a through d show the average eye and hand behaviors for Null-Kick-Assistive and Null-Kick-Resistive trials. The assistive field following the

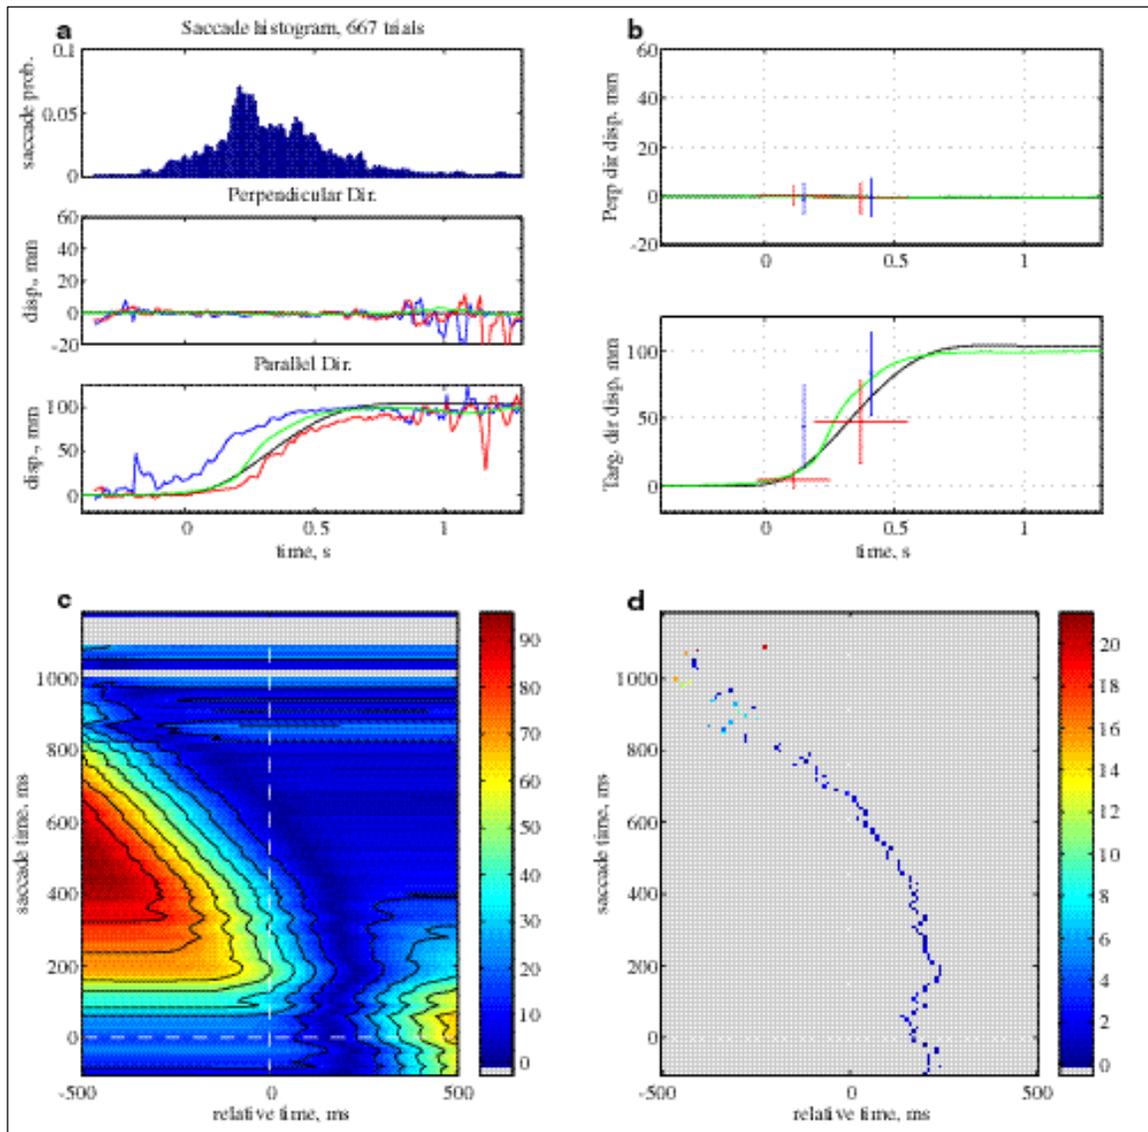


Figure 3.11. Hand and eye averages for No Kick movements in force-pulse experiments. **a, b:** Eye and hand averages. **c, d:** Eye-hand distances and minima, in mm.

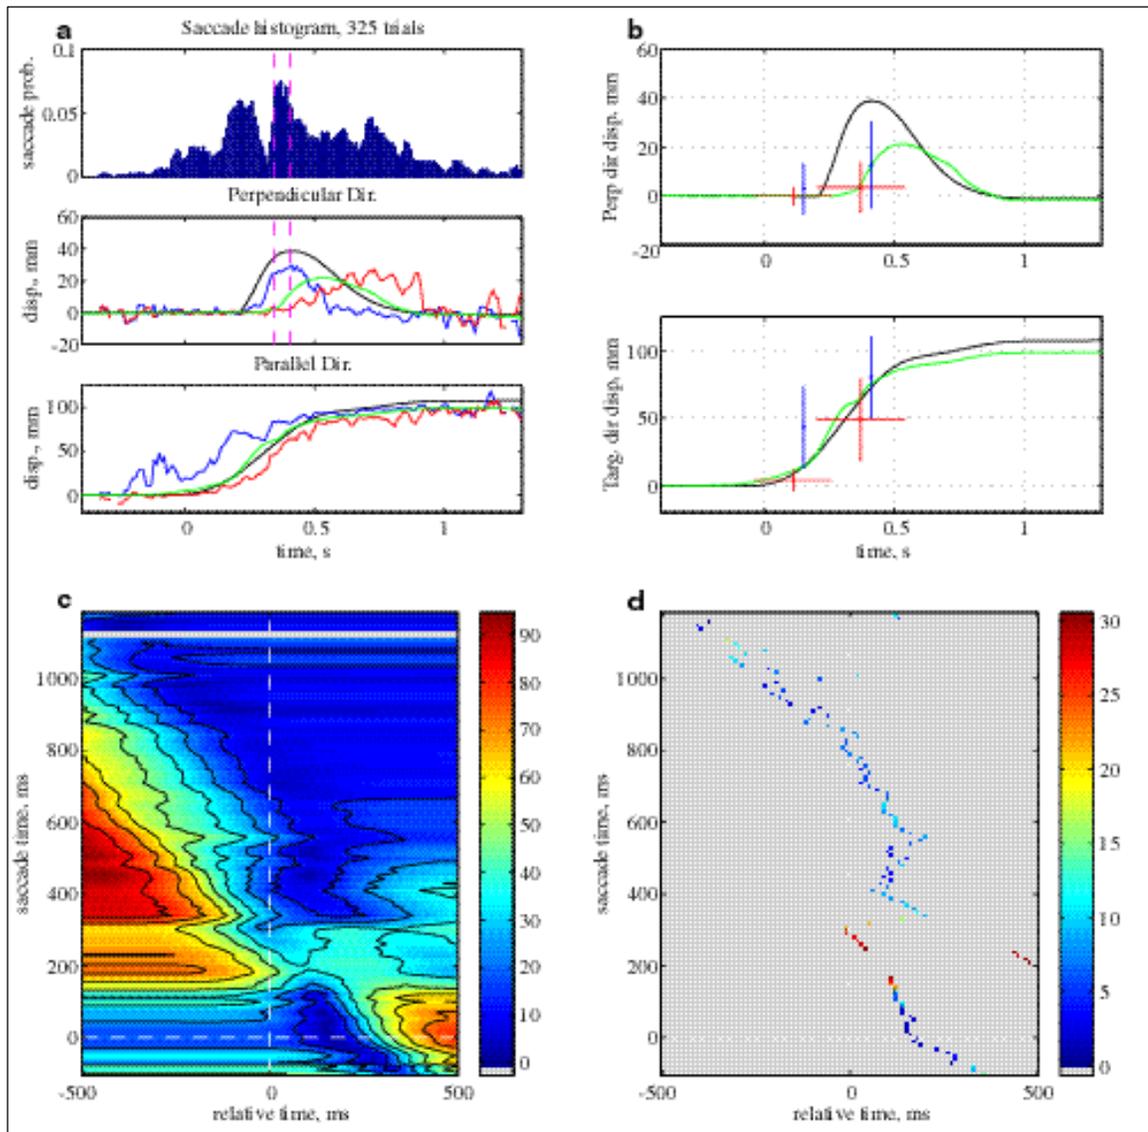


Figure 3.12. Hand and eye averages for Null-Kick-Null movements in force-pulse trials. **a, b:** Eye and hand averages. **c, d:** Eye-hand distances and minima, in mm.

pulse has the effect of amplifying the perpendicular displacement with respect to the Null-Kick-Null trials, while the resistive field has the effect of damping it (Figure 3.13). The eyes reflect the effects of the fields less strongly than the hand. Figure 3.14 shows the eye-hand distance for the Null-Kick-Assistive and Null-Kick-Resistive trials. In Null-Kick-Assistive, the eyes begin to accurately predict hand behavior much later in the movement with respect to the time of the pulse. In Null-Kick-Resistive, the eyes appear to predict more accurately than in the –Assistive trials.

In Figure 3.15, which shows the average saccade endpoints for No Kick, Null-Kick-Null, Null-Kick-Assist, and Null-Kick-Resist trials together, we can see that the earliest saccades with reasonably large perpendicular displacement (falling within the bin indicated by vertical dashed lines) are nearly the same across fields. The saccade displacements quickly change following this, with an increase in -Assistive and a decrease in -Resistive with respect to the Null-Kick-Null trials.

Do the eyes in the Null-Kick-Assistive and Null-Kick-Resistive predict hand behavior in the Null-Kick-Null trials? If the eyes “expect” a null field following the kick, their behavior should, initially after the pulse, better predict the Null-Kick-Null arm behavior than the –Assistive or –Resistive behaviors. One method to attempt to answer this question is to subtract the average saccade endpoints from the Null-Kick-Assistive or –Resistive trials from the mean hand trajectory for Null-Kick-Null trials. The minimums of these distances are shown in Figure 3.14e and f. Eye behavior in Null-Kick-Assistive appears to predict better for the Null-Kick-Null hand than the Null-Kick-Assistive hand, as expected. This does not appear to be true for the Null-Kick-Resistive eye in Figure 3.14f. This is possibly due to the similarity of arm trajectories in Null-Kick-Null and

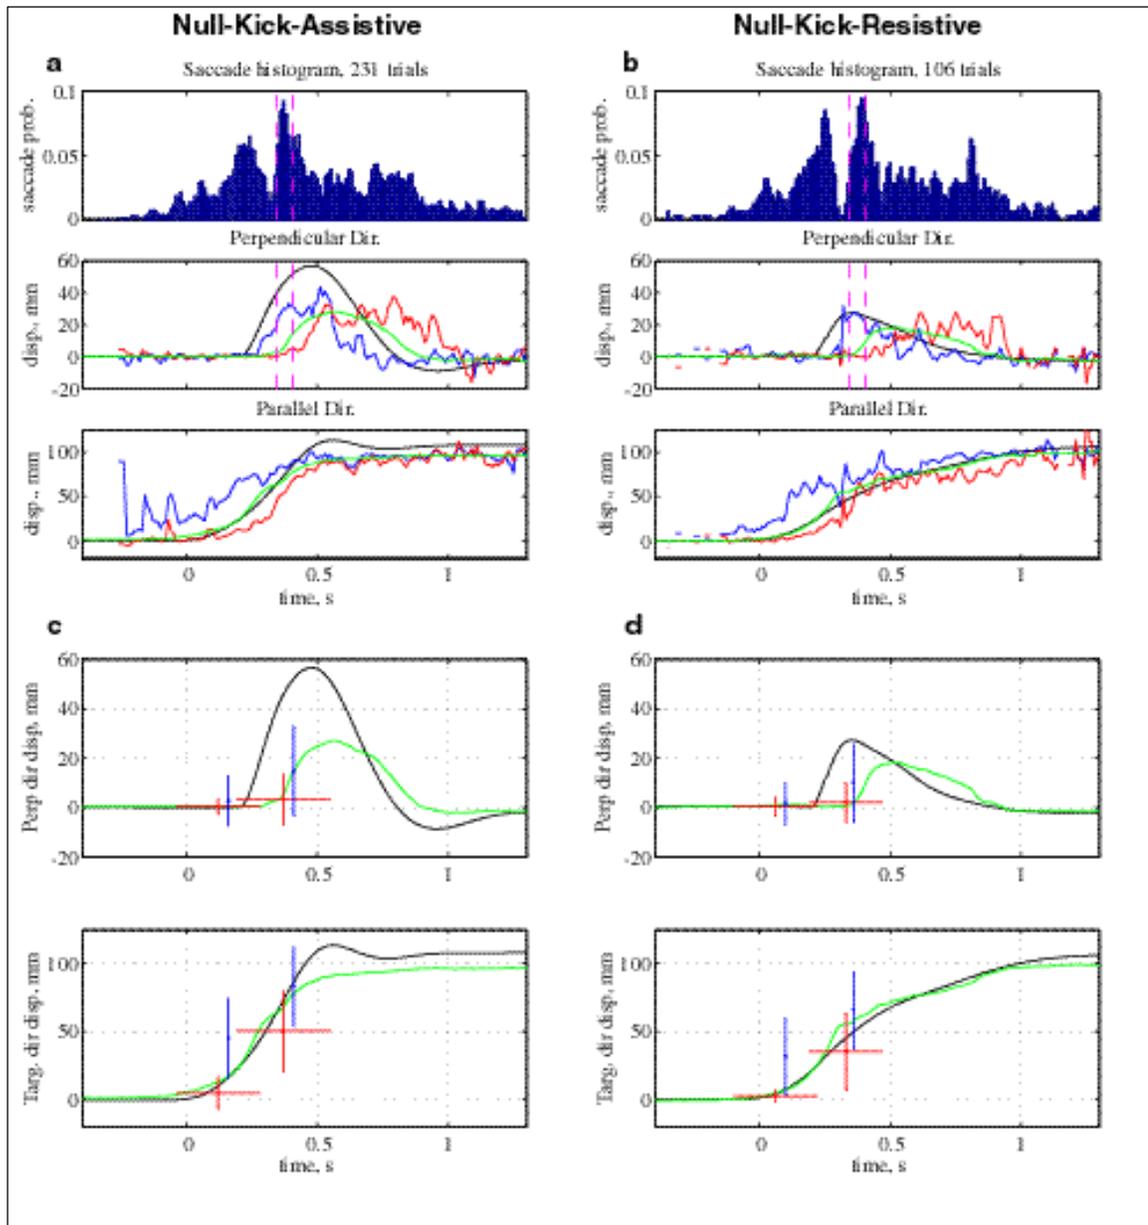


Figure 3.13. Hand and eye average for force-pulse experiments. **a,c:** Null-Kick-Assistive trials. **b,d:** Null-Kick-Resistive trials.

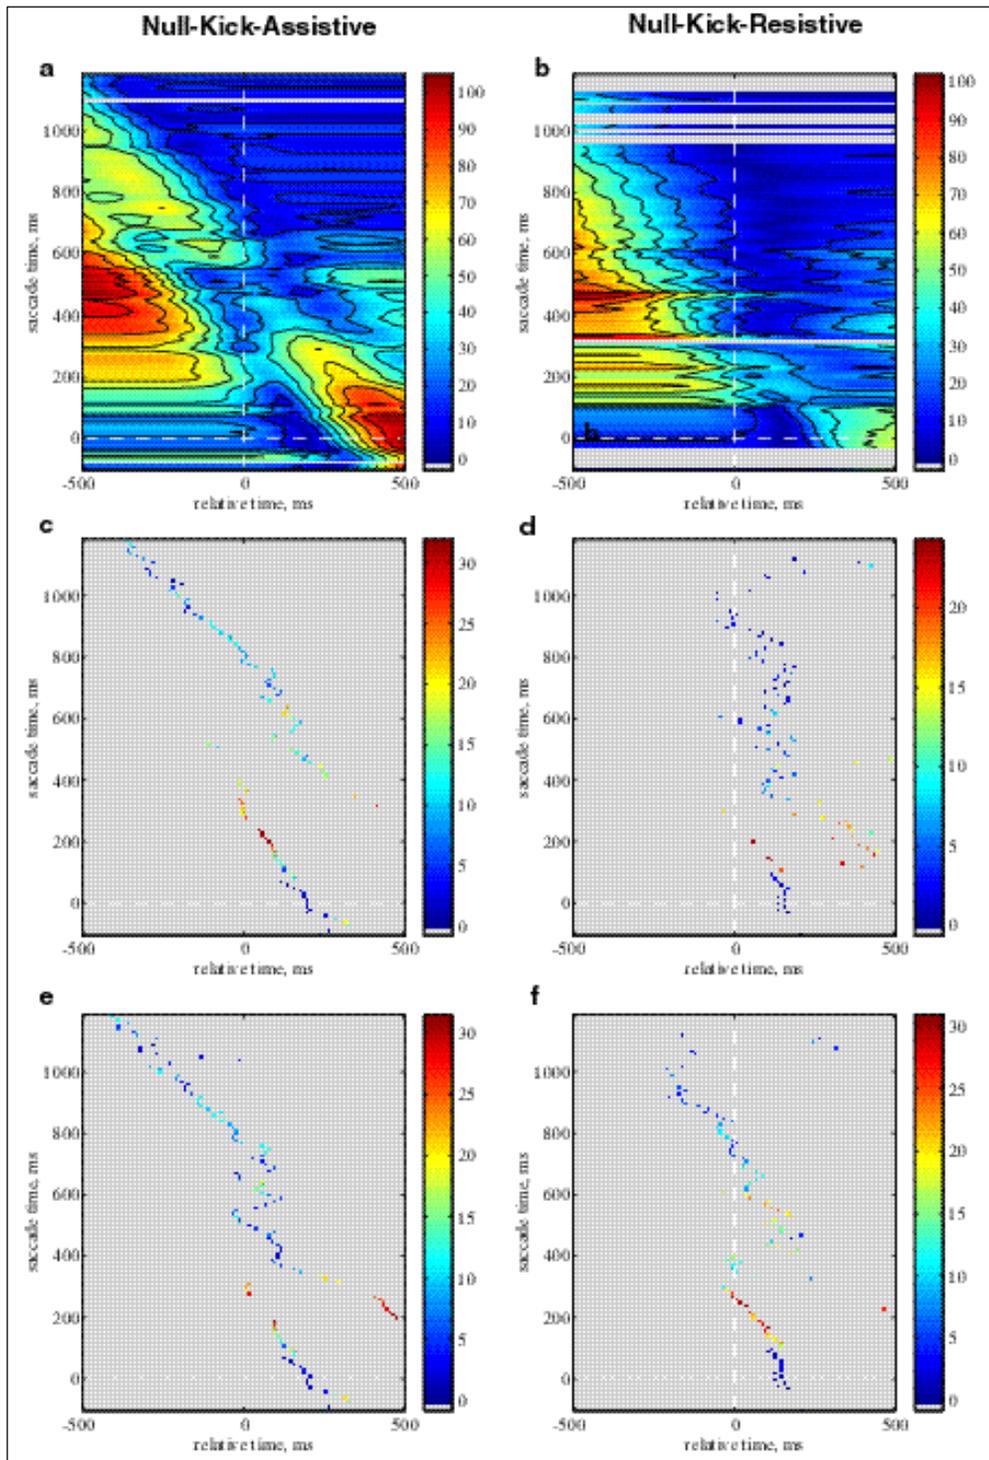


Figure 3.14. a,b,c,d: Eye-hand distances and minimums, in mm, for Null-Kick-Assistive and Null-Kick-Resistive trials. e: Average saccade endpoints from Null-Kick-Assistive trials and mean hand trajectories from Null-Kick-Null trials. f: Average saccade endpoints from Null-Kick-Resistive trials and mean hand trajectories from Null-Kick-Null trials, in mm.

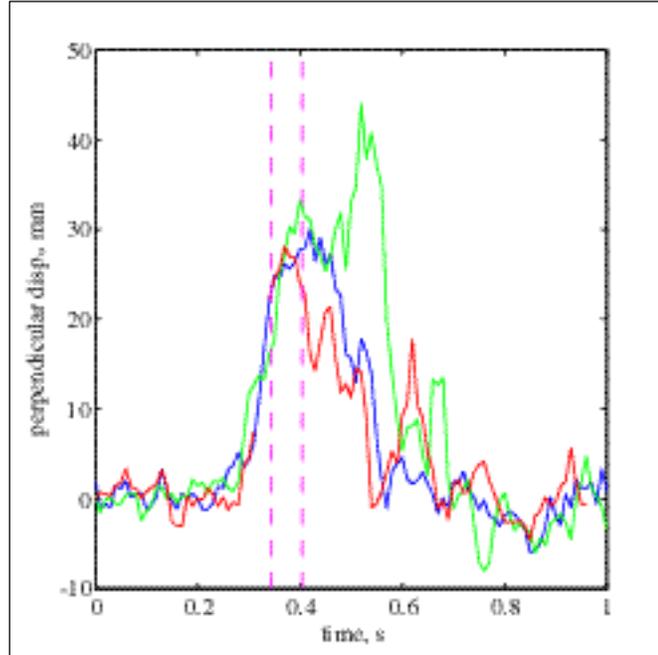


Figure 3.15. Perpendicular displacement for saccade endpoint averages in Null-Kick-Null (blue), Null-Kick-Assistive (green), Null-Kick-Resistive (red) trials. Dashed lines indicate times 0.35s and 0.4s, containing the peak in the saccade histogram, which will be used in subsequent analyses.

–Resistive trials. As we are about to show, other measures do indicate inaccuracies of the eye in Null-Kick-Resistive trials.

Figures 3.12 through 3.14 demonstrate the relationship of the average eye and average hand for different force fields. However, even for a single type of trial, considerable movement variation exists due to the effects of the force-pulse on the arm in different configurations and due to numerous other parameters (such as intended motion, arm stiffness, movement velocity). By examining saccades within a thin slice of times, it is possible to look at the relationship of eye and arm movements for individual trials. Figures 3.12a, 3.13a, and 3.13b show a 50ms slice of time encompassing the peak of the saccade histogram following the force-pulse, from 350ms to 400ms. To make a comparison of saccade placement at this time to the hand position, we must estimate the prediction time discussed earlier. No Kick trials indicate that, on average, the prediction time for saccade endpoints lies between 150 and 250ms.

To obtain an improved estimate of the prediction time for saccades at the post-pulse histogram peak, we compare hand positions at various times with respect to the saccade endpoints according to the following equation:

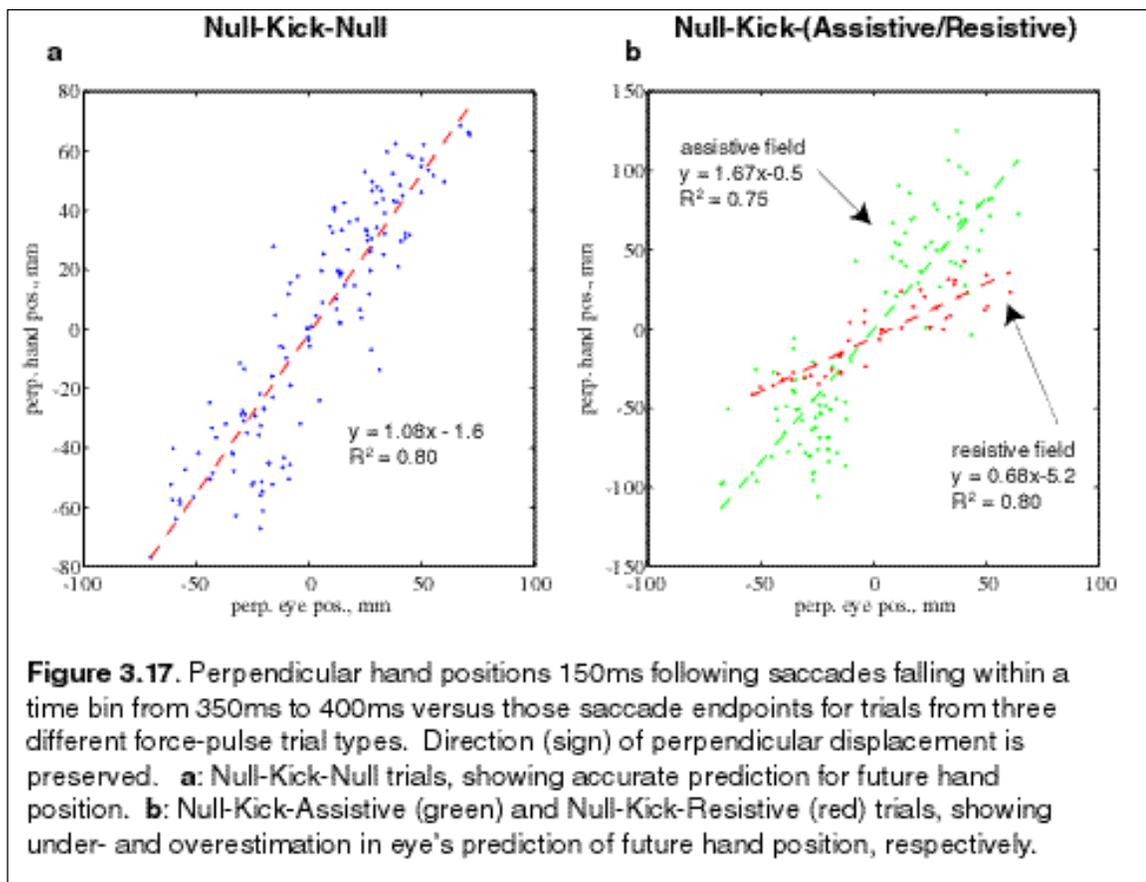
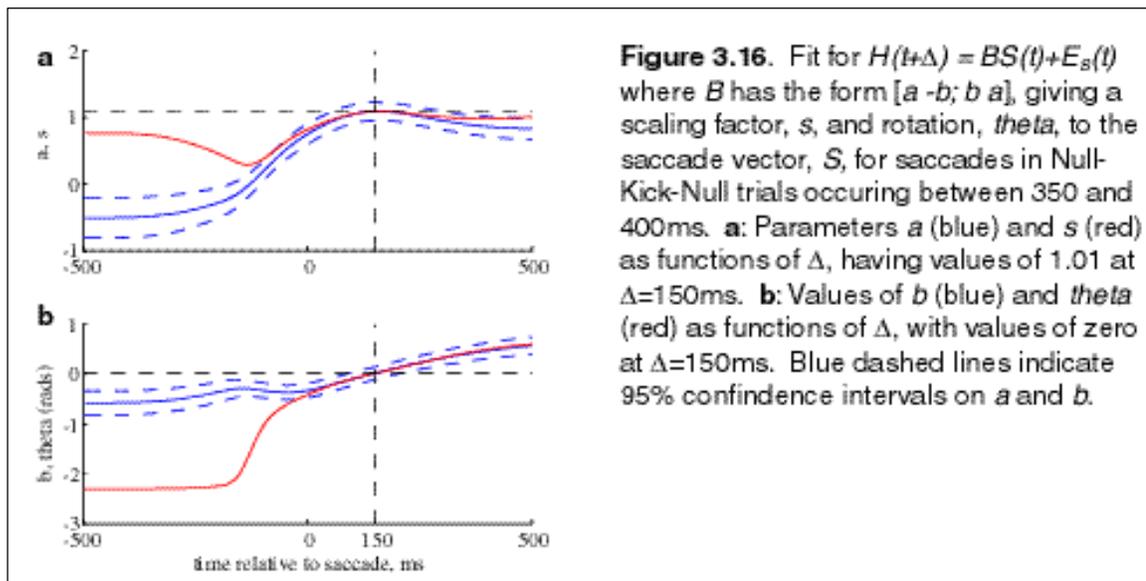
$$H(t+\Delta) = BS(t) + E_s(t),$$

where $E_s(t)$ is the vector representing the startpoint of the saccade, $S(t)$ is a vector representing the direction and magnitude of the saccade, and $H(t+\Delta)$ is the vector of hand position at some time Δ with respect to the saccade. By constraining the matrix B to have the form $\begin{bmatrix} a & -b \\ b & a \end{bmatrix}$ and computing B using a nonlinear least squares fit, we can obtain the scaling and rotation of the saccade vector which best approximates the hand position for a

particular Δ . For the saccade to be made to the location of the hand at some point in time, both the direction and magnitude of the saccade must be appropriate. For such a saccade, a would be one, while b would be zero. This analysis was performed for all saccades within the 350-400ms bin in the Null-Kick-Null trials and for values Δ from -500 ms to 500 ms. Figure 3.16 shows the results of this analysis, plotting the values of a and b , as well as the scaling factor, s (where $s = \sqrt{a^2 + b^2}$), and the rotation angle, $\theta = \tan^{-1}(b/a)$, as functions of Δ . The rotation parameter θ equals zero (meaning b equals zero) at $\Delta=150$ ms, at which time a and s roughly equal 1, indicating that the magnitude and direction of the saccade are both appropriate and that 150ms is, on average, the prediction time for saccades within this bin.

In Figure 3.17, all perpendicular components of saccade endpoints falling within this slice are plotted against the hand's perpendicular displacement for the same movements at 150ms following the saccade, for each field type. For this estimated prediction time, saccade endpoints fit perpendicular displacement best in the Null-Kick-Null condition with a slope of approximately 1 (Figure 3.17a). Saccades under-predict Null-Kick-Assistive hand position and over-predict Null-Kick-Resistive hand position (Figure 3.17b). This is an indication that in each trial the eye behavior predicts what the arm behavior for that trial *should be* if the null field follows the pulse.

Does the output of the forward model in the simulation yield similar results? Figure 3.18 shows the modeled hand trajectory and the forward model output for a force pulse in Null-Kick-Null, Null-Kick-Assistive, and Null-Kick-Resistive paradigms. While the prediction time of the forward model in the simulation is shorter than the measured prediction time for the eye, the perpendicular displacement of the hand is only accurately



represented by the model in the Null-Kick-Null trial, as in the eye behavior. (The simulated inverse and forward models “know” the null field, like the subjects in the experiment.) The forward model underestimates the perpendicular displacement of the hand in the Null-Kick-Assistive trial and overestimates in the Null-Kick-Resistive. Figure 3.19 shows the movement paths for the Null-Kick-Null trials and simulations. As in the comparison made for the curl field, the forward model behavior shows some similarities to that of the eyes for these trials.

The saccade behavior of the pulsed trials is particularly similar to the output from a forward model of dynamics in the sense that the saccade goes to a position that is appropriate for the *known* dynamics, whether or not the known dynamics are the actual dynamics. When the known dynamics are correct (in the Null-Kick-Null trials), the eyes are able to predict future hand position even when that hand position does not lie along the desired trajectory (as indicated by non-zero perpendicular displacements). The difference between these pulsed trials and the viscous curl field experiment is that while the unknown dynamics were present throughout the movement in the curl field, they quickly disappear in the pulsed trials. For this reason, after the introduction of error from the pulse, the forward model has the ability to predict future hand position because its model of dynamics is correct for the remainder of the movement.

Force-pulse movements: Varied pulse times

Force-pulse movements with pulse times of 200ms, 250ms, 300ms, and 400ms, as well as No Kick movements, were pseudorandomly given with equal probability of occurrence. The saccade histograms in 3.20a through d, for each force-pulse time, show that the peak

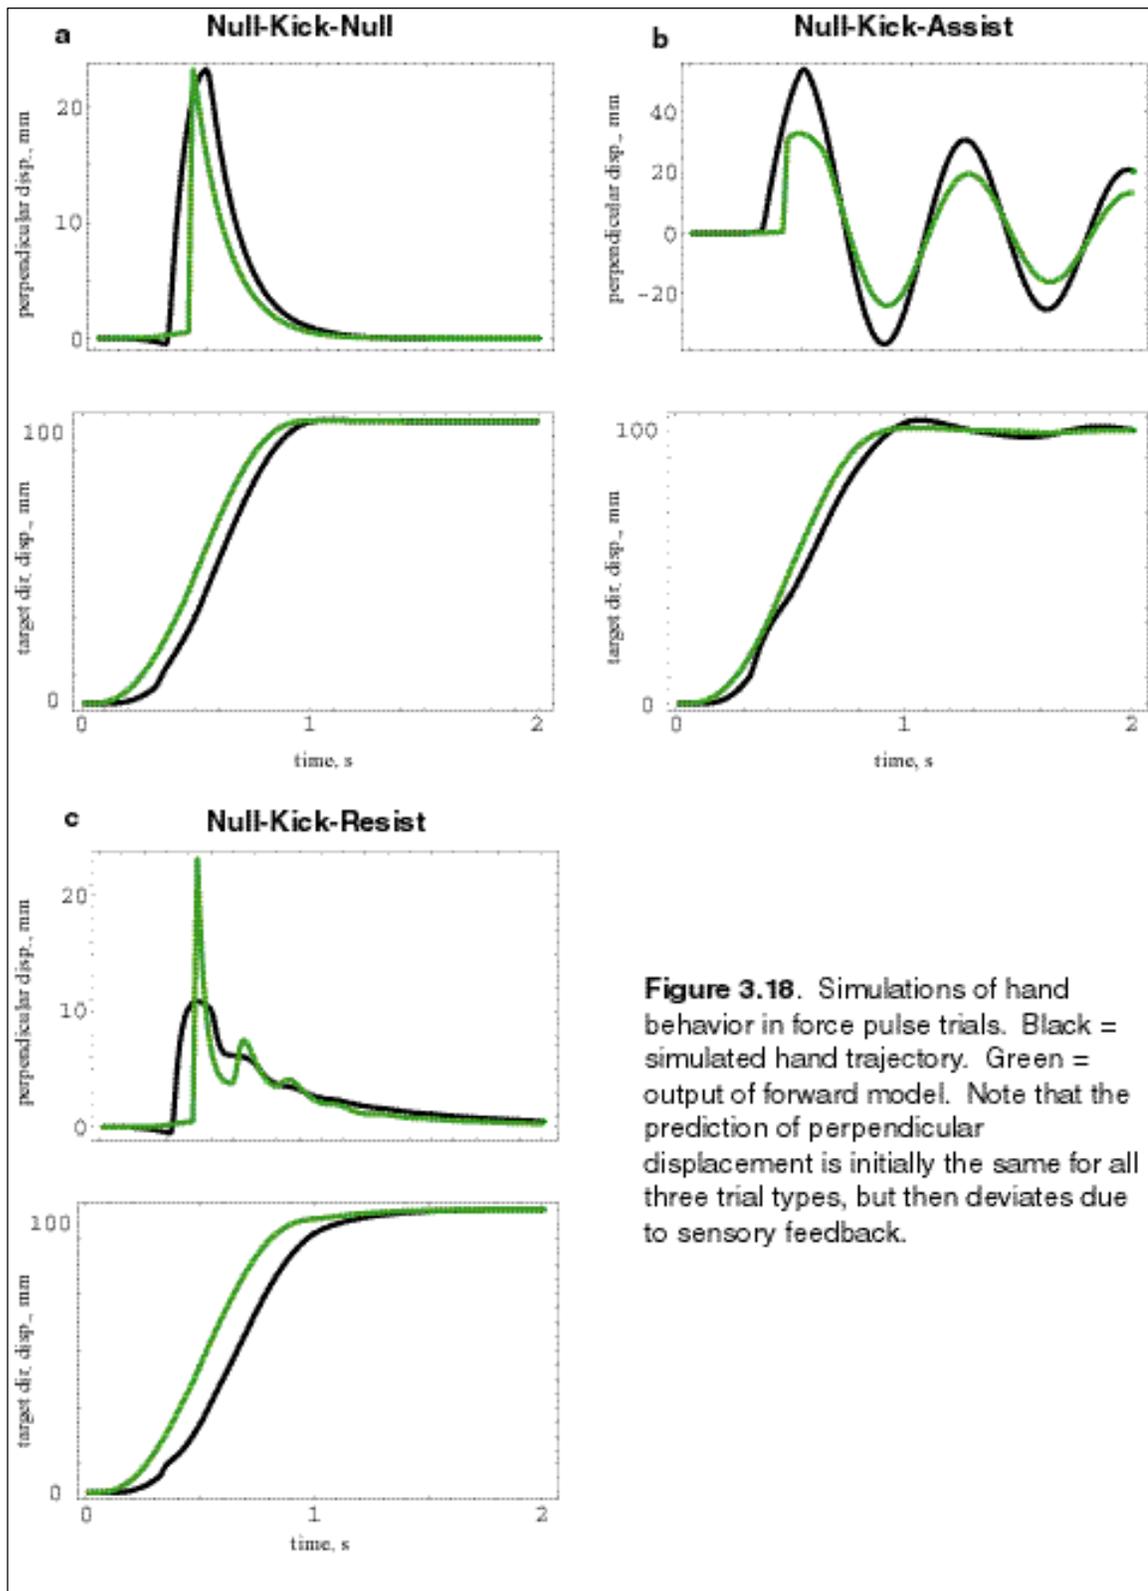


Figure 3.18. Simulations of hand behavior in force pulse trials. Black = simulated hand trajectory. Green = output of forward model. Note that the prediction of perpendicular displacement is initially the same for all three trial types, but then deviates due to sensory feedback.

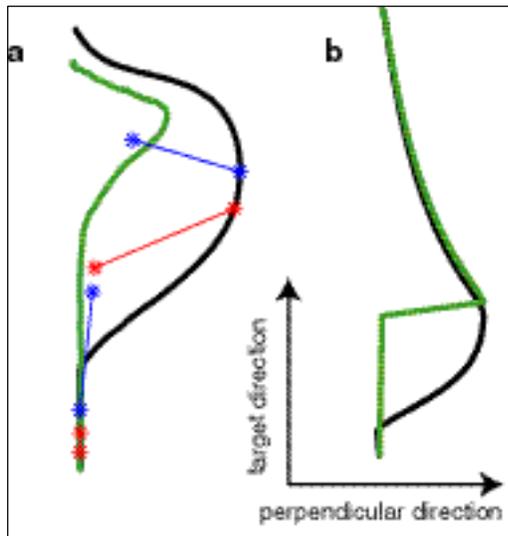


Figure 3.19. Qualitative comparison of eye behavior and simulated forward model output for Null-Kick-Null trials. **a:** Average hand path (black) and average eye path (green). Saccade startpoints and endpoints are shown with red and blue asterisks, respectively, for average of first two saccades. Lines connect saccades to hand position at the same time. **b:** Simulated hand path (black) with simulated forward model output (green).

in saccade frequency following the kick (as seen for 200ms kicks in Figure 3.12) is time-locked to the pulse time, following the pulse by approximately 150ms. The dip in the histogram (possibly indicating saccade suppression) follows the kick by approximately 120ms.

Average behavior of eye and hand for each pulse time are shown in Figures 3.20 and 3.21. For a given trial, because the second saccade can occur over a range of times, there is some probability that the second saccade will follow the pulse with sufficient time to reflect the perpendicular hand displacement. This probability decreases as the time of the pulse increases. This results in the second saccade averages having a decreased perpendicular component with increasing pulse time.

The predictive ability of saccade endpoints is examined in Figure 3.22 where, for the 200ms pulse time, behavior is consistent with that observed in the Null-Kick-Null trials from the Null-Kick-Null/Assistive experiment. With increasing pulse time, saccades in the initial part of the movement are accurate for longer periods, as might be expected. However, the eye-hand fit for an assumed 150ms prediction is not consistent across pulse times (Figure 3.23), increasing from 1.18 for the 200ms pulse time to approximately 1.8 for the 350ms pulse time. In each case, the saccades were taken in 50ms-wide bins starting 150ms after the pulse. This underestimate of perpendicular displacements later in the movement may be related to the growing underestimate of target direction displacement in null field trials described earlier.

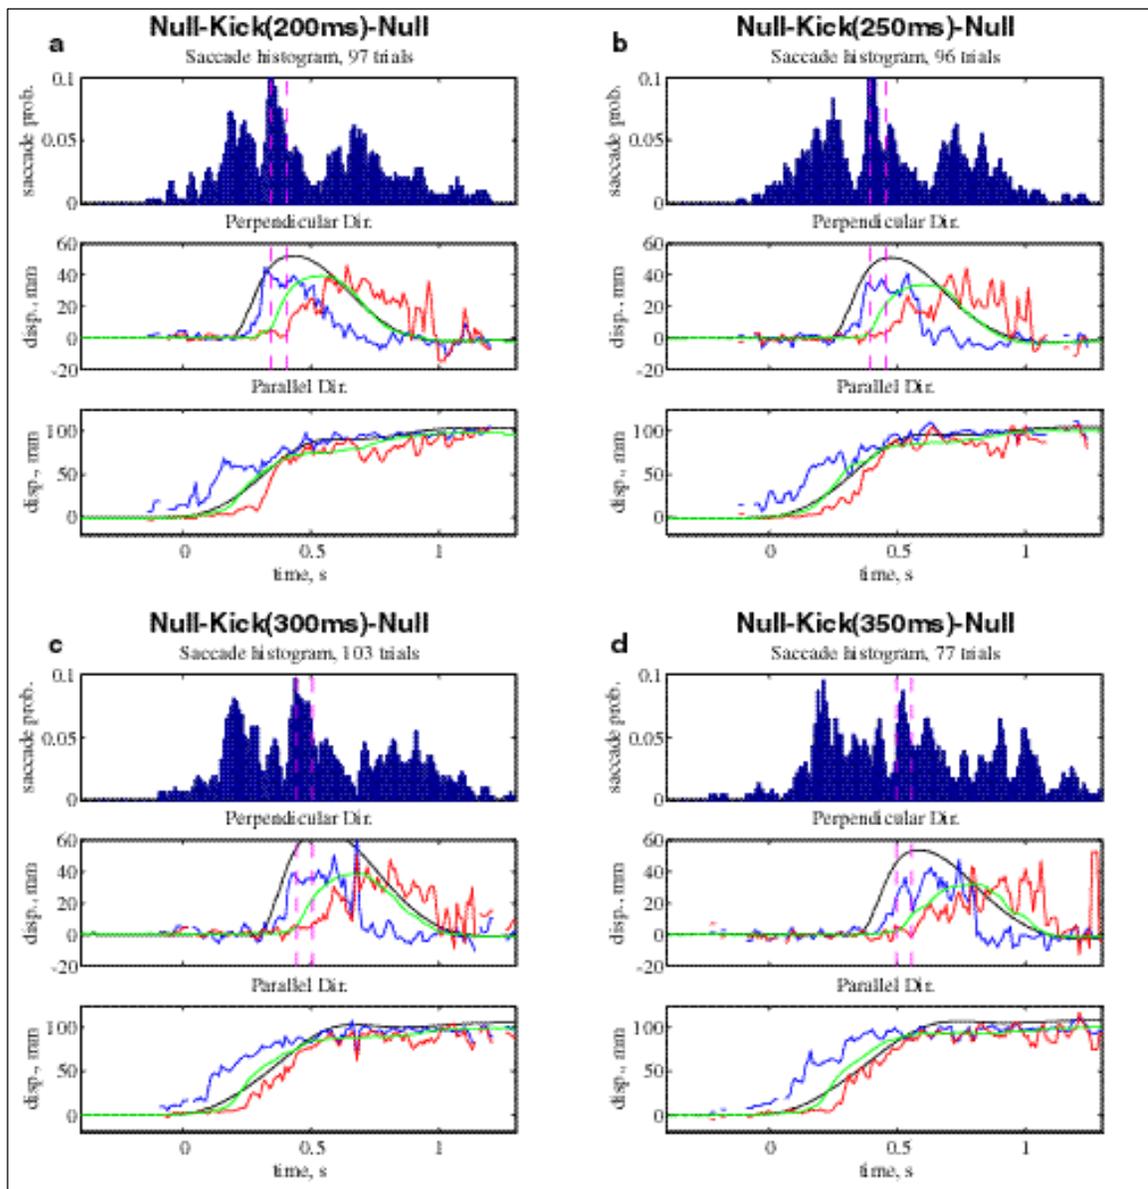


Figure 3.20. Eye and hand averages and saccade histograms for variable kick-time experiments, showing all four kick times. Time bins enclosed by dashed magenta lines contain saccades used for eye vs. hand analysis in Figure 3.24.

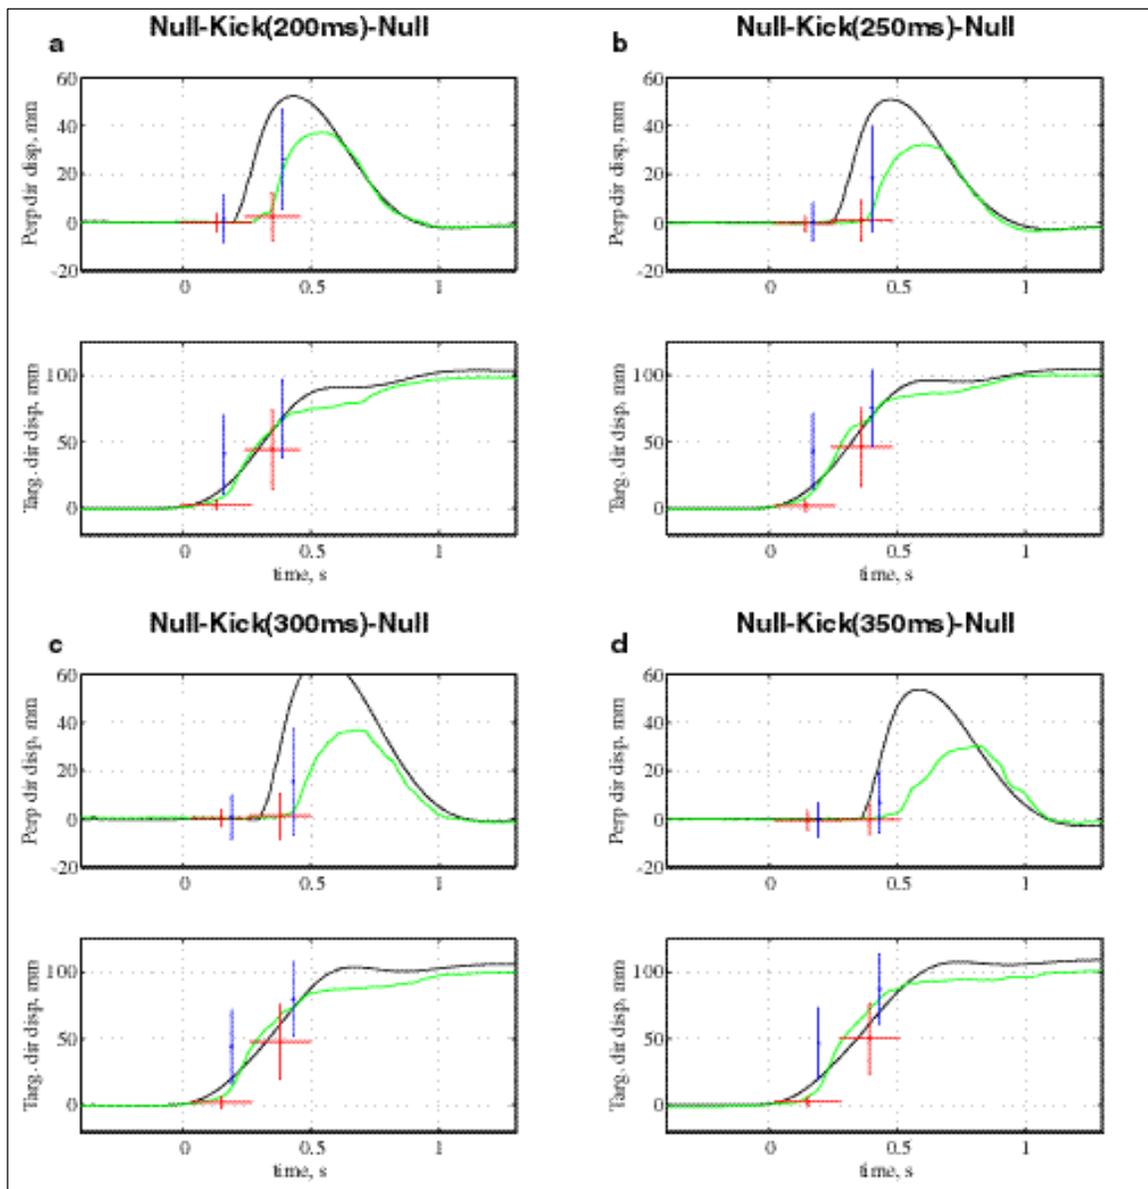


Figure 3.21. Eye and hand average trajectories and averages of first two saccades for variable kick-time experiments, showing all four kick times.

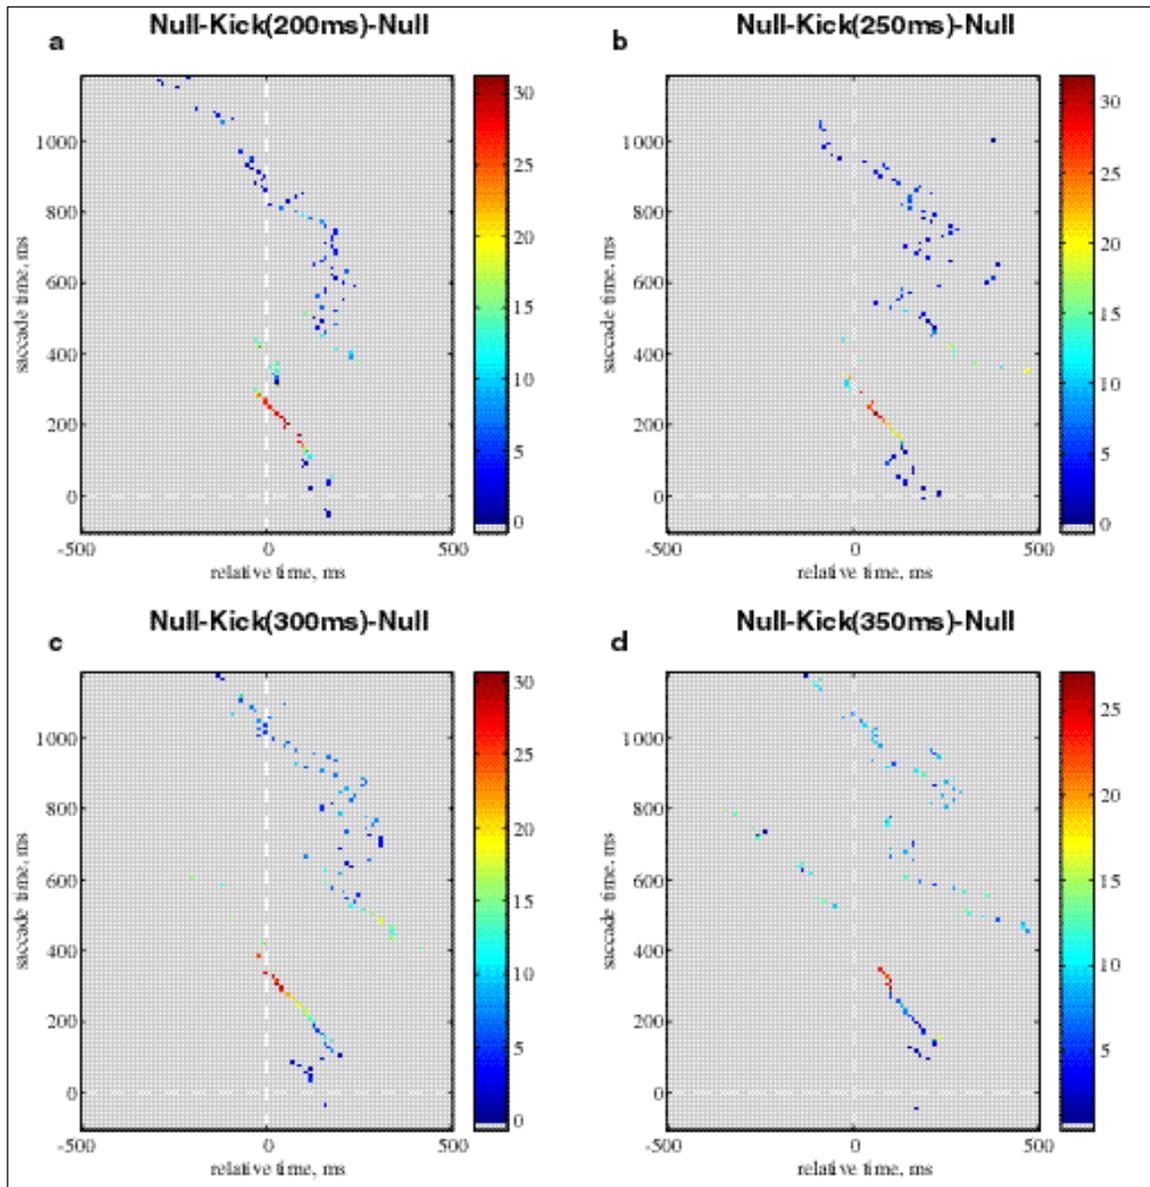


Figure 3.22. Minima of eye-hand distances, in mm, for variable kick-time experiments, showing all four kick times.

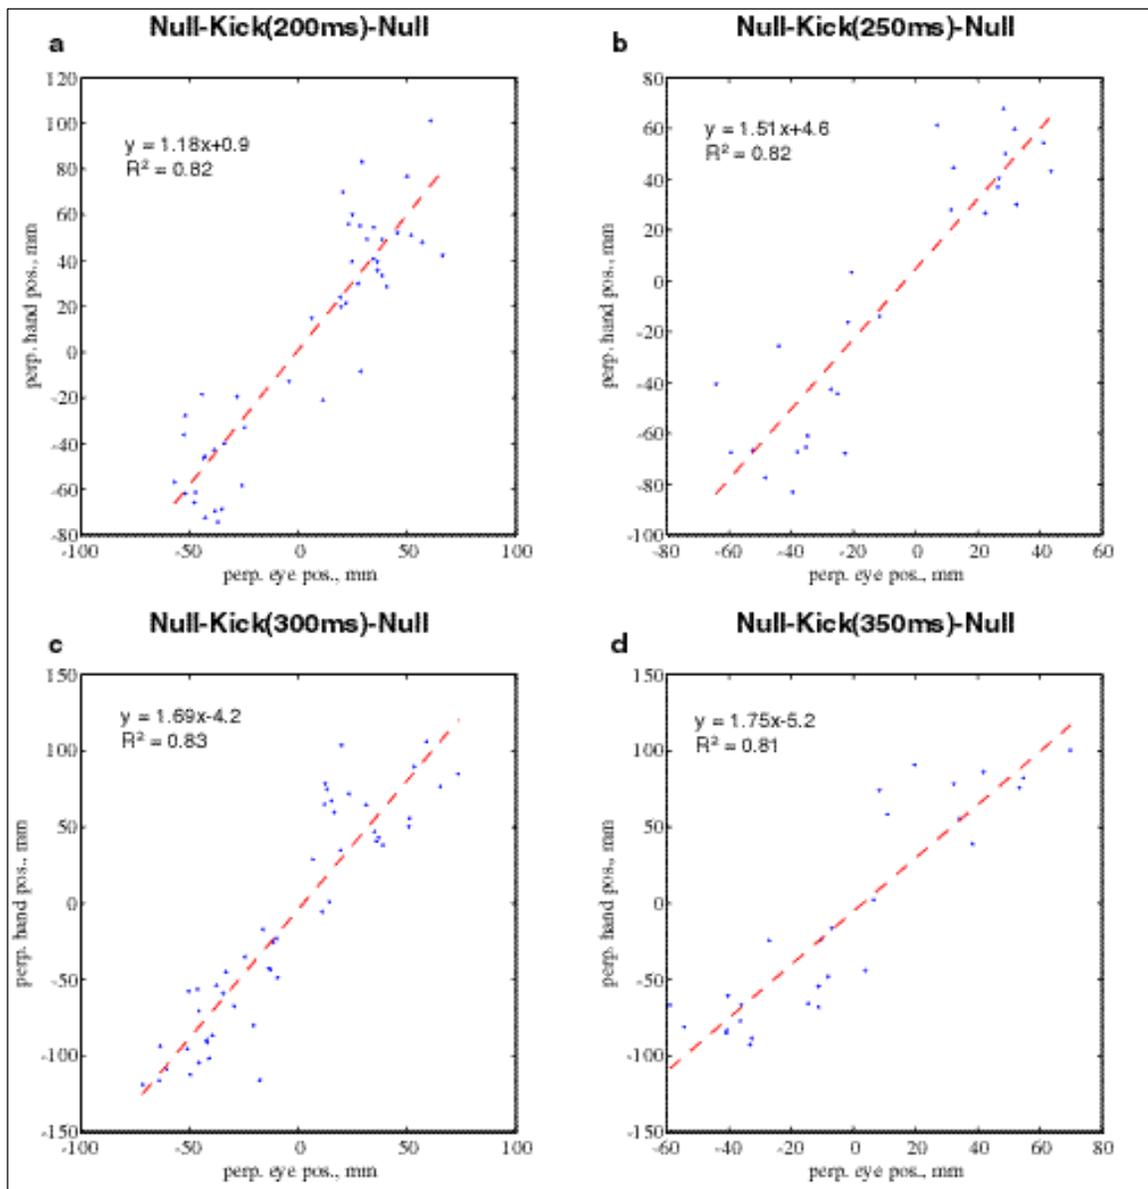


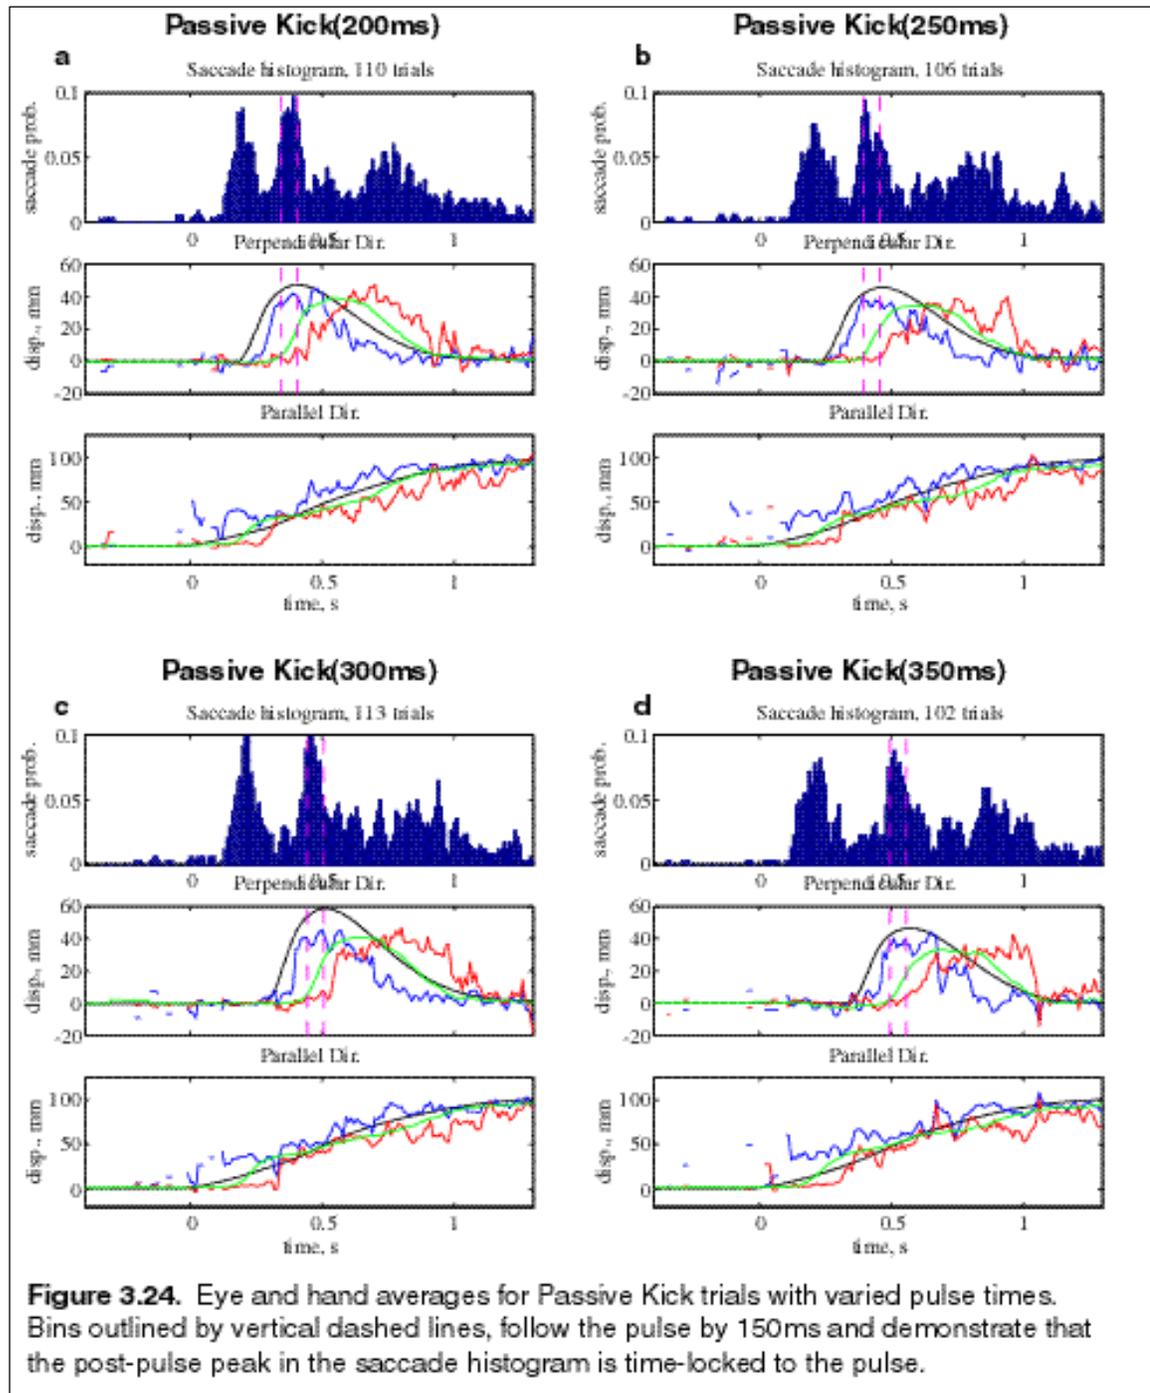
Figure 3.23. Perpendicular hand positions 150ms following saccades falling within a 50ms time bin (shown in Figure 14) versus the perpendicular component of the saccade endpoints within the bin. The bin in for each pulse time begins 150ms following pulse onset. Direction (sign) of perpendicular components are preserved.

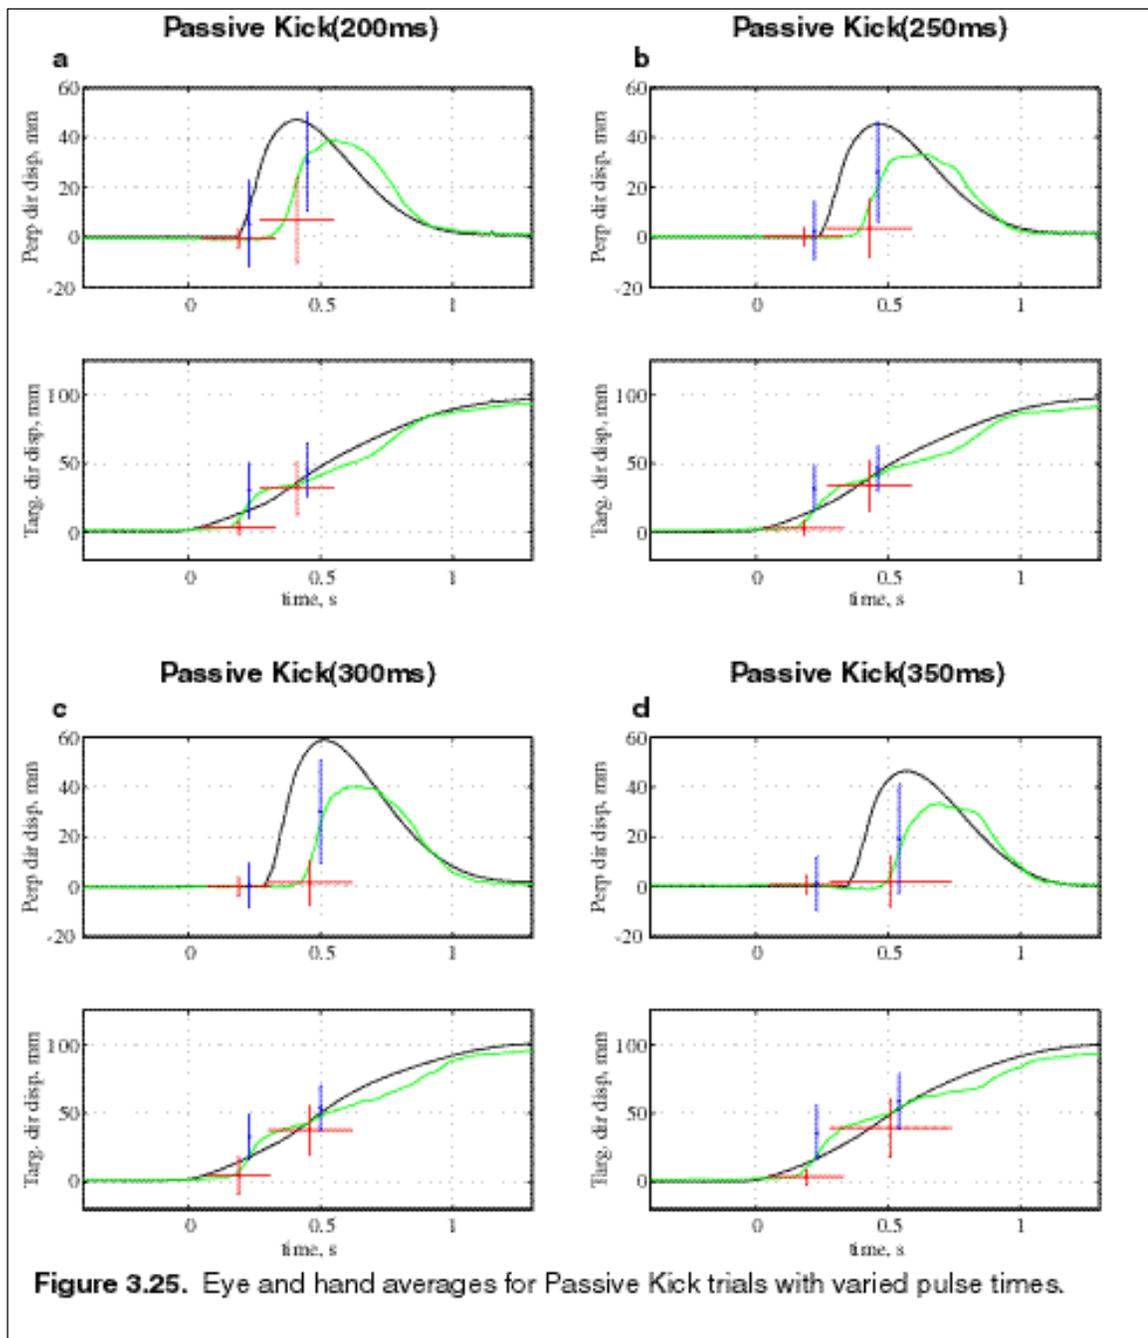
Passive kick experiments

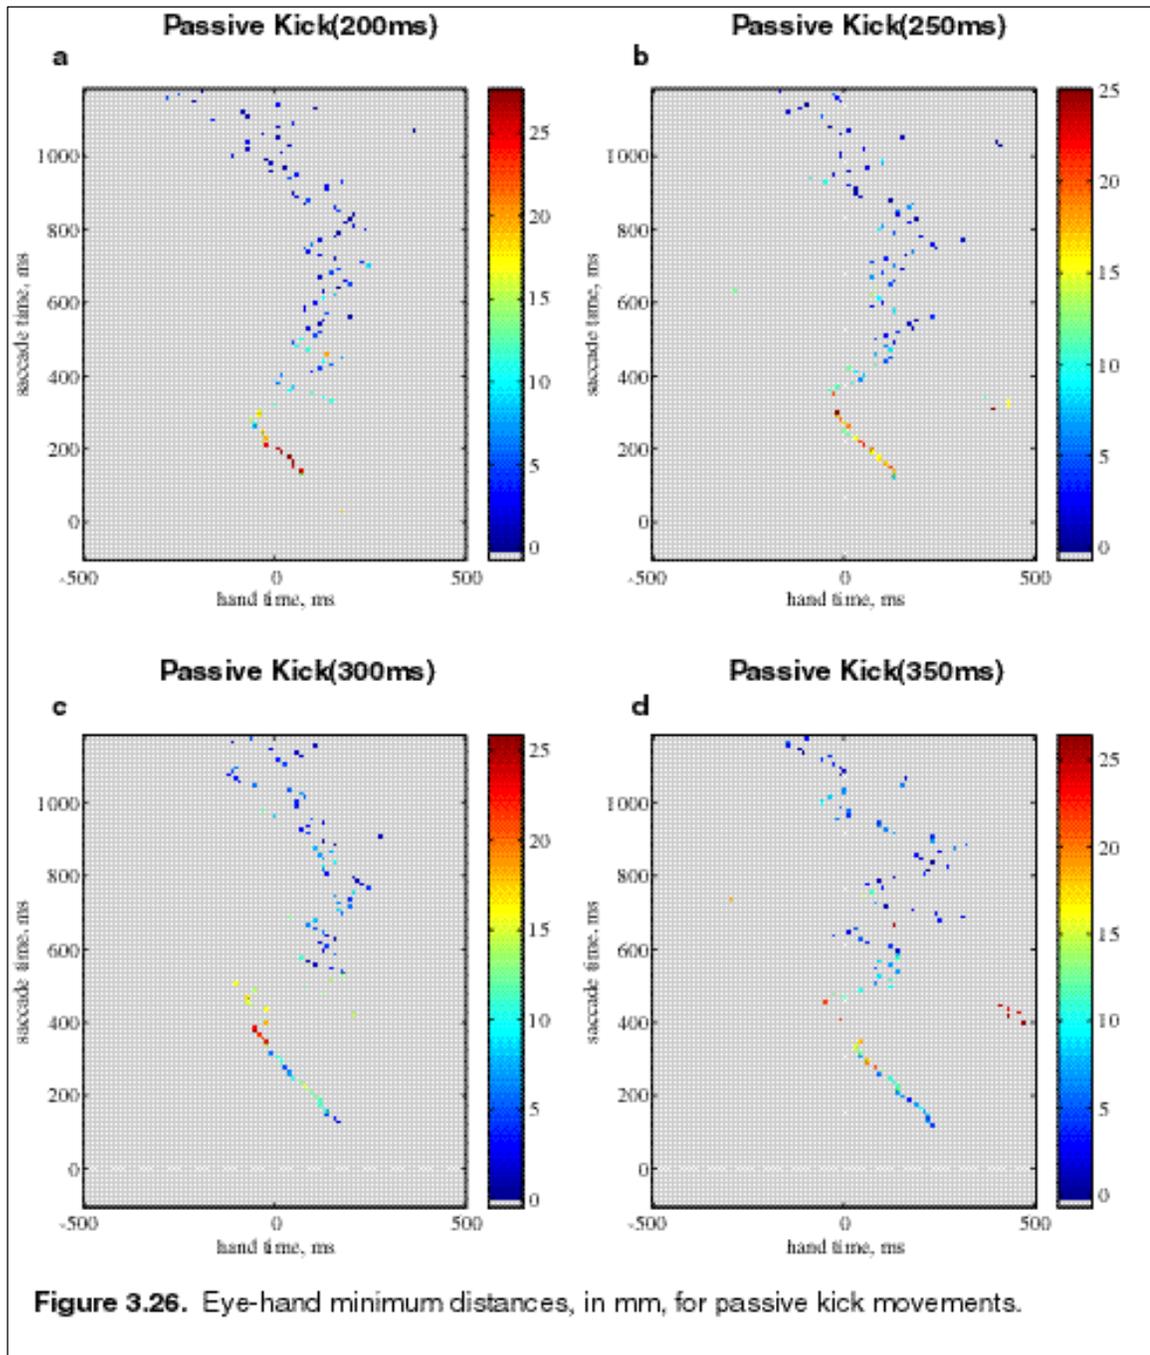
As was done for the viscous field experiments, an experiment was run in which the robot simulated force-pulse trajectories with subjects passively gripping the handle. In several ways, these results closely resemble passive curved movements described earlier. Saccade histograms in Figure 3.24 show that saccades generally lagged movement onset by roughly 150ms. Figures 3.24 and 3.25 reveal that saccade startpoints lagged the hand position more than in self-generated hand movements, and saccade endpoints led the hand in the target direction less consistently. However, in these trials, it appears that the eye is better able to predict future hand position than in the passive curved movements. This is supported by Figure 3.26 which shows a consistent prediction by saccade endpoints following force-pulses at 200ms, 250ms, 300ms, 350ms. Figure 3.27 further confirms this by showing that the eyes performance for each trial in predicting the hand for that same trial 150ms later is even better than for self-generated movements.

These results for passive kicked trials may seem surprising. They can be partially explained, however, by our concept of a forward model. While the subject attempts to be passive, he is most likely unable to sufficiently inhibit the spindle reflex. The spindle reflex will cause the subject to generate forces opposite the direction of error generated by the force pulse. This reflex-generated movement will complement the robot's planned movement since the robot's movement is already designed to simulate human reaction to a kick. Because a forward model of dynamics must be able to predict the results of arm reflex mechanisms, the eyes are able to predict the arm trajectory with a high degree of accuracy. The increasing underestimation for active pulsed trials is not evident in these passive trials, suggesting that the underestimation in the perpendicular direction (and,

perhaps, in the target direction as well) is related to the cortical component of the motor command or the intended goal of reaching the target for active trials.







Chapter 4. Discussion and Conclusions

Restating the question

The forward model is of particular interest in the field of human motor control because it has the power to explain how the brain deals with sensorimotor delays. While simulations allow us to examine the expected output of a forward model (and any other component, for that matter) for controlling the arm, the influences on human behavior are not so easily separated. The arm's response to error is thought to be the result of such properties as inertial dynamics, stiffness, and spinal reflexes in addition to the effects of the inverse and forward models.

The natural coordination of the eye and hand in performing visually-guided manual tasks caused us to ask the following question: Does a forward model of arm behavior play a role in the coordination of eye and hand? Over the course of two years, we set out to design experiments and collect data which would serve to answer this question. Guiding our search was the simulations of the arm using a forward model of dynamics. Our faith in this model was based on its intuitive sense as well as the fact that previous research in our lab had demonstrated its ability to explain characteristics of human movements and motor learning.

We were well aware of the problems that asking such a question create. If we found indisputable influence of a forward model for the arm in eye behavior, the investigation would be considered successful. (We would have to rely, however, on the simulation results of Bhushan and Shadmehr to demonstrate its use by the arm controller, since proving the existence of a forward model does not necessarily demonstrate its

application in control). On the other hand, failing to find the influence of a forward model for the arm on the behavior of the eye did not rule out the possibility that the forward model existed for use by the arm controller alone. Nonetheless, the research was driven by the exciting connection that the eye was an excellent estimator/predictor for arm behavior and that a likely candidate for the source of this estimation would be forward model output (should the forward model exist).

To convincingly demonstrate that the eye controller responded to a forward model of arm dynamics for saccadic tracking tasks, a few criteria should be met: 1) In the presence of error, the eye exhibits behavior that is qualitatively similar to the simulated output of a forward model for arm dynamics; 2) As the subject learns the dynamics of a reaching task, eye behavior adapts in a manner which is in agreement with an adapting forward model. The criteria must be met in situations where the eye behavior cannot be explained by a simplistic method for combining desired trajectory and error.

Revisiting the results

The first of the two criteria stated above was met with considerable success, while the second was not. This section will provide a summary of the results of this thesis as well as a discussion of why demonstrating the learning of a forward model was found to be exceedingly difficult.

We have found that, when subjects attempted to track the perceived location of their arms during point-to-point reaching movements with their eyes in the absence of visual feedback, eye motion was predominantly saccadic and responsive to motor commands as well as sensory feedback. Despite the fact that eye behavior is variable

under seemingly static conditions, convincing results were achieved by having subjects perform numerous movements and taking average behaviors the movements from numerous subjects. Our results may be summarized as follows:

1. Saccades generally lead the motion of the hand by originating near the location of the hand and terminating along the future path of the hand. This prediction time can be as long as several hundred milliseconds.
2. Saccade timing and placement are adjusted in response to error by reflecting the trajectory of hand motion.
3. When the arm is passively moved slowly, saccades may still attempt to predict future hand position but are less accurate, presumably due to the decreased influence of motor command on arm trajectory.
4. As a reaching movement progresses, the eyes increasingly underestimate the distance traveled by the hand. There is some evidence that this effect exists both in the target direction and in the perpendicular direction. It should be noted that Wolpert et al. (1995) also observed a change in the bias of hand location estimate as a function of time for reaching trials. While the bias tended to be positive (overestimate), several important differences exist between that experiment and the ones described in this thesis. Most importantly, only endpoint estimates were made for near-constant velocity movements of different durations.
5. When the arm is exposed to sudden perturbation, saccades are suppressed approximately 120ms following the perturbation and then experience an increased likelihood 30ms later. Saccades occurring within this “burst”

strongly reflect the perturbation by having a significant component in the direction of perturbation.

6. Eye behavior qualitatively resembles a discretized (saccadic) representation of the simulated output of a forward model of arm dynamics. This resemblance is particularly evident in response to perturbation from curl fields and force pulses. As predicted by the simulations, the eyes underestimate perpendicular displacement of the hand in curl fields. Also in accordance with the simulation, eye placement predicts perpendicular components of hand trajectories following a force pulse if the null field then follows the pulse. Eyes underestimate displacement due to assistive fields following pulses and overestimate for resistive fields. Saccades occurring roughly 175 ms after the pulse (during the peak in the saccade histogram) predict, on average, the hand placement 150ms later.
7. For passive movements, eye accuracy is improved if the trajectory causes rapid changes in error, presumably because an insuppressible spinal reflex is triggered whose effects are estimated by the forward model. While the eyes tend to under predict perpendicular displacement of pulses later in active movement trials, the eyes continue to predict accurately for passive trials.

The sixth point above is the crucial evidence in meeting the first criterion, particularly in the force pulse experiments. Just as the forward model simulation suggests, the eyes are able to accurately predict future hand position even in the presence

of error *when the dynamics are known*. If the dynamics are unknown, the eyes behave appropriately for the forces that are expected.

Difficulties in demonstrating the adaptation of forward models

Pilot experiments were designed and performed to examine the adaptation of eye movements as subjects learned force fields. These were intended to be compared to simulations of forward models in which the imposed force fields were only partially modeled. This task proved to be more difficult than expected.

According to internal model theories, the internal model increasingly represents the initially unknown, imposed forces over consecutive trials. This means that both the inverse and forward models would adapt. If adaptation of these models occurred simultaneously and at the same rate, they would represent the same force field in any given movement, even when this representation is incorrect (or partially correct). In this case, the output of the forward model would equal the desired trajectory for the initial, feedforward portion of the movement. This occurs because the inverse model produces motor commands intended to produce the desired trajectory, while the forward interprets the motor commands as producing the desired trajectory. Due to the agreement of inverse and forward models, initial saccades would be expected to be aimed towards the target, regardless of the degree of adaptation which has occurred. In the pilot experiments, the eyes did appear to saccade consistently toward the target, though the data is not shown in this thesis.

If the forward model and inverse model were not identical, initial saccades could be expected to take place in directions other than towards the target. If, for example, the

forward model accurately modeled the force field while the inverse model did not, the initial saccades could be expected to take place in the direction of initial movement error before the error occurred. This would take place because although the inverse model would produce motor commands that were inappropriate for the desired trajectory, the forward model would accurately estimate the error those motor commands would produce. While Bhushan and Shadmehr suggested that, in some cases, the forward model may indeed learn faster than the inverse model, these results are suggestive at best and the conditions under which such a situation would occur is not clear.

Another complicating factor for demonstrating adaptation of the eye movement is the high rate of learning in the early trials when exposed to a novel force field. As we have seen, the eye movements demonstrate their dependence on a forward model in the way that they reflect error. When there is no movement error, little can be said about the role of a forward model. (For control, there is, in fact, no need for a forward model in the absence of error.) If error is quickly reduced, few trials containing sufficient error exist for analysis, and, as we have seen, the inherent variability of saccade behavior necessitates large number of trials for analysis.

As a solution to these difficulties, it was thought that by having subjects perform force pulse trials in a learned force field, an adapted forward model could be demonstrated. A few subjects were adapted to the assistive field by performing 100 non-pulsed movements in this field in different directions. After adaptation, subjects performed randomly pulsed movements in which either the assistive field or the null field followed the pulse. The hypothesis was that eye behavior would accurately predict the assistive field behavior after the kick, but would overestimate the null field behavior,

demonstrating that the subjects had developed a forward model of the assistive field. However, subjects instead accurately predicted the null field behavior and underestimated the effects of the assistive field. In retrospect, we now believe that this inability to predict the effect of a learned field is the result of poor generalization in the velocity space. Subjects trained in the assistive field at speeds that were considerably lower than those imposed by the force pulses. While internal model theories state that forces may be learned as a function of velocity, it is unclear how these models of forces generalize to velocities that have not been visited during training.

Conclusions

While forward models for motor control have been hypothesized to compensate for sensorimotor delays, evidence for their existence is lacking. To date, the most convincing results have come from simulations of arm motion utilizing a forward model. These models produce simulated arm behaviors that exhibit some of the more complex features of human arm behavior in point-to-point reaching movements.

To our knowledge, the saccadic tracking experiments described in this thesis demonstrate the first direct evidence for forward models of limb dynamics. It is not surprising that during ocular tracking (either pursuit or saccadic), eye behavior is influenced by a combination of both efferent copy and sensory signals. However, the forward model hypothesis makes specific claims as to how these two sources of information are integrated. By using the motor commands to integrate arm state forward from the delayed state obtained from sensory feedback, an estimate of current or future state is computed.

Under each of the experimental conditions tested, saccade endpoints were consistent with a sampled forward model predicting several hundred milliseconds into the future. When the forces applied by the robot were predictable (or *learned*, presumably modeled by the forward model), saccades went to a future location of the hand. This behavior occurred even in the presence of trajectory error following a force pulse of short duration. When the forces were not expected (or *nonlearned*) by the subject, deviations of saccades from the hand trajectory were consistent with the error of the forward model output in the simulation. In the passive curved trials, the eyes were unable to consistently predict future hand state due to the lessened import of efferent copy. For passive kicked trials, saccadic prediction was improved, most likely because the forward model predicted the spinal reflex dominating hand motion following the pulse.

While it appears that the forward model of the arm is a significant contributor to eye behavior, our simulations do not predict saccade placement completely. For example, the simulation does not predict the underestimation in active movements in both the target and perpendicular directions. It is possible that these errors are the result of inaccuracies of the forward model's dynamics, errors in sensory feedback, or the influence of information other than the forward model output. Further research in this area may reconcile these differences between the expected forward model output and actual behavior of the eyes.

Despite some inaccuracies in the eye's predictive ability, these results support the existence of a forward model of arm dynamics, whose primary purpose is most likely to permit stable feedforward control of arm motion by compensating for sensorimotor delays. It appears that this model is also accessible by the brain to control eye motion, as

hypothesized. Such use of the forward model in eye-hand coordination may serve to increase the usefulness of visual input in natural reaching situations, by providing visual information about spatial positions in advance of the hand's arrival, even in the presence of error. More generally, forward models of dynamics may play an important role in both controlling individual body systems and in the coordination of different systems by allowing the brain to predict the consequences of internal and external influences on the body as a whole. Further research into the behavior and learning of eye-hand coordination and the coordination of other motor effectors will aid in the understanding of forward models, as well as understanding the ability of the human being to act within and interact with his environment.

Appendix: Detection of saccades from gaze point position data

An algorithm was implemented for determining saccade start and end times from the eye position data. This algorithm used a number of criteria and was designed to distinguish saccades from noise in the position signal. Criteria and parameters were empirically chosen with the goal that the algorithm detected the same events as would be detected by an analysis of the position trace performed by a human researcher identifying saccades.

The algorithm was applied to each dimension of eye position (*perpendicular* and *target*, or *x* and *y*) independently. First, eye position was convolved with a template of an “ideal” saccade shape in time. The absolute value of this convolution was then filtered using a Savitsky-Golay filter. Next, the filtered version was subtracted from the absolute value of the convolution, creating what we will term the “difference signal”. For the time of each peak (local maximum) in the difference signal exceeding a set threshold value, the following steps were applied:

1. The slope of the eye position in time is calculated from the first pre- and post-peak eye positions.
2. Working from the peak time backward in time, the saccade start time is determined to be the first time for which the absolute value of the position slope is less than one fourth the absolute value of the slope at the peak.
3. Working from the peak time forward in time, the saccade end time is determined to be the first time for which the absolute value of the position slope is less than one fourth the absolute value of the slope at peak.

4. The average slope of the position and the duration of the saccade are calculated using the start and end times determined above. If the slope of the saccade is less than a set threshold or the saccade duration is greater than a set threshold, the saccade is discarded.
5. With saccades thus far detected independently from the x and y position data for each trial, these saccades are combined and discrepancies eliminated. If saccade start and stop times are determined to be different from the x and y positions for the same saccade, the actual saccade start time is taken to be the earlier start time and the actual saccade end time is taken to be the later end time.
6. Only saccades of lengths greater than or equal to a set threshold value are included. All others are discarded.

The result of this algorithm was that only saccades of significant size were detected, avoiding the acceptance of noise in the eye position signal and of “micro-saccades”. The parameters chosen worked well for the data collected from the experimental paradigms used. For a different task (for example, one with a higher proportion of smooth pursuit), other parameters may have worked better.

References

- Bhushan N. and Shadmehr R. (1999). Computational nature of human adaptive control during learning of reaching movements in force fields. *Biological Cybernetics*, 81:39-60.
- Bhushan, N. and Shadmehr, R. (1999). Evidence for a forward dynamics model in human adaptive control. *Advances in Neural Information Processing Systems*, 11:3-9.
- Gauthier, G.M., Vercher, J.L., Mussa-Ivaldi, F., and Marchetti, E. (1988). Oculo-manual tracking of visual targets: control learning, coordination control and coordination model. *Experimental Brain Research*, 73:127-137.
- Lunenburger, L., Kutz, D.F., and Hoffmann, K.P. (2000). Influence of arm movements on saccades in humans. *European Journal of Neuroscience*, 12:4107-17.
- Miall, R.C. and Wolpert, D.M. (1996). Forward models for physiological motor control. *Neural Networks*, 9(8):1265-79.
- Miall, R.C., Reckess, G.Z., and Imamizu, H. (2001). The cerebellum coordinates eye and hand tracking movements. *Nature Neuroscience*, 4(6):638-644.
- Neggers, S.F.W. and Bekkering, H. (2000). Ocular gaze is anchored to the target of an ongoing pointing movement. *Journal of Neurophysiology*, 83(2):639-51.
- Vercher, J. and Gauthier, G.M. (1992). Oculo-manual coordination control: Ocular and manual tracking of visual targets with delayed visual feedback of the hand motion. *Experimental Brain Research*, 90:599-609.
- Vercher, J., Gauthier, G.M., Guedon, O., Blouin, J., Cole, J., and Lamarre, Y. (1996). Self-moved target eye tracking in control and deafferented subjects: roles of arm motor command and proprioception in arm-eye coordination. *Journal of Neurophysiology*, 76(2):1133-44.
- Wolpert, D.M., Ghahramani, Z., and Jordan, M.I. (1995). An internal model for sensorimotor integration. *Science*, 269:1880-82.

Scholarly Life

Gregory David Ariff was born in Voorhees, New Jersey to Steven G. and Muriel A. Ariff. After attending Piman High School in Pitman, New Jersey, Gregory obtained a bachelor's degree in Engineering Science and Mechanics from Virginia Polytechnic Institute and State University, where he concentrated in Biomedical Engineering. During his undergraduate years, he spent several summers and semesters working in industry with Baxter Healthcare and DuPont Tyvek, and in academic research at the University of Minnesota and with the Materials Response Group at Virginia Tech. Upon completion of his BS in 1999, Gregory enrolled at Johns Hopkins University for his masters degree in Biomedical Engineering. He completed his degree in October of 2001, having spent two years performing research of eye-hand coordination in reaching tasks under the guidance of Dr. Reza Shadmehr in the Laboratory for Computational Motor Control. Gregory has since accepted a position with Directed Technologies, Inc., an engineering firm in Arlington, Virginia.