NORTHWESTERN UNIVERSITY

Understanding the Haptic Interactions of Working Together

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Mechanical Engineering

By

Kyle Brandon Reed

EVANSTON, ILLINOIS

June 2007

ABSTRACT

Understanding the Haptic Interactions of Working Together

Kyle Brandon Reed

An understanding of how two people anticipate, adapt, and react to each other's forces and motions could aid in designing machines to work cooperatively with humans and further explain how a single human interacts with the world. Tasks, such as lifting and moving a bulky object, teaching manual skills, dancing, and handing off a baton or a drinking glass, involve haptic interaction, which is a communication channel distinct from spoken language and gestures, but much less studied.

Throughout this thesis, I will discuss my experiments on physical interaction and negotiation between two agents working on a target acquisition task. First, I will evaluate human-human physical interaction. I will show that dyads are faster than individuals, despite applying larger forces. By analyzing the interaction forces through a jointly controlled object, I will reveal a distinctly different completion strategy for dyads that is not available for individuals. Second, I will look at human-robot interaction. By simulating human-human interaction, a robot can surreptitiously replace one of the human partners leading to improved disturbance rejection.

Acknowledgements

I thank my family for encouraging me throughout my entire life. To my parents for not getting mad when I failed a French final exam and understanding that I am better in other areas. To Kerry for encouraging my fun-loving nature. To Kevin for always keeping my mind sharp with a good debate on just about any topic.

I thank everyone in LIMS for the various help and random discussions as well as the many fun times. Thanks to Eric, Songho, Jay, Mark, David, Brian, Aaron, Vik, Tom, Gabriel, James, Laura, Jason, Matthieu, and Peng.

I greatly appreciate the support from the many friends in the Graduate Christian Fellowship at NU. As Kristen told Hannah who told me: Trust God to give you the answers - He knows this stuff better than you or your committee!

I especially thank my advisor, Michael Peshkin, for helping me gain a further understanding of engineering, publishing papers, and teaching. I thank my committee members, Mitra Hartmann and Jim Patton, for their help and ideas throughout this research. Thank you to Kevin Lynch, Ed Colgate, Marcia Grabowecky, Malcolm MacIver, and Peter Vishton for the advice given to me throughout this project.

I thank the National Science Foundation for granting me an NSF Graduate Fellowship. I would not have been able to begin this research without it. I also thank the NSF for continued funding of this project through the Human and Social Dynamics grant.

A very special thanks to my wife, Gigi, for being so wonderful! For the many meals she cooked (and left) for me, for her continued encouragement to work harder and graduate faster, and for the many prayers.

Northwestern University Motto, which I try to uphold: “Whatever is true, whatever is honorable, whatever is just, whatever is pure, whatever is lovely, whatever is gracious, if there is any excellence, if there is anything worthy of praise, think about these things.”
- Philippians 4:8

Table of Contents

List of Tables

List of Figures

Executive Summary

A basic form of human interaction is the physical cooperation necessary to perform a manual task with others. Physical cooperation represents a communication channel distinct from facial expression, gesture, and spoken language, yet one that has been much less studied. Shared physical tasks require participants to adapt, anticipate, and react to each other's forces and motions and are often performed in the absence of any explicit verbal communication. An understanding of how two people physically cooperate, compromise, and guide one another is a fundamentally important aspect of human motor control and it may allow improved cooperation and physical communication between a person and a machine. The forces and motions involved in haptic communication may be coupled directly limb-to-limb, or via a mutually grasped object, which is the focus of this thesis.

As an initial foray into human-robot-human interaction, I devised the simplest experiments to reveal novel effects that arise in dyadic motion control. I designed and built a one-axis motion configuration in which two subjects participate symmetrically and communicate solely through physical interaction. The crank has three possible inputs to control the position: two handles that subjects manipulate and one motor that can apply a torque directly to the crank.

When a person completes a task individually, they necessarily apply all the required force. Adding a partner allows new possible completion strategies where each member can apply less than the total required force. A robotic partner can similarly aid in task completion. If the robotic partner is correctly designed to follow human-human communication standards, a robotic partner can surreptitiously replace one of the human partners.

My experiments reveal a dyadic completion strategy surprisingly different from that of an individual. When two people work together, each member actually applies significantly more force than an individual does. Some of this force is in opposition to their partner and some is in cooperation. The opposition force is internal to the dyad and results in no crank motion. I call this *dyadic-contraction* because it resembles co-contraction in an individual. Dyadic-contraction occurs when each subject applies an equal and opposite force to the their handle. I found that many dyads exerted a significant dyadic-contraction force, in some cases over 40 Newtons. I believe this force is predominantly for stabilizing the interaction of the dyad and also as a means for the members to communicate, since other channels are absent.

Communicating solely through the forces and motions of the crank, the members of a dyad were able to negotiate a distinctly different strategy than an individual. Both members used their past knowledge and collective force to cooperate throughout different task phases. I call this cooperation *specialization*, which consists of dyads temporally dividing the task such that one member takes control to accelerate while the other member takes control to decelerate the crank.

I implemented the human-human interactions discussed above into a robotic partner by simulating dyadic-contraction and specialization. Surprisingly, the robot was unable to elicit specialization in the subjects, but it was able to make subjects believe they were interacting with a person as opposed to a robot. I found that the human-robot pair was better at disturbance rejection than two humans working together.

Despite the opposition forces described above, the completion time of dyads was significantly faster than individuals. When a human worked with a robot simulating a specialized partner, the performance was similar to a single person completing the task. Subjects working with the robotic partner performed faster if they perceived the partner to be a human rather than a robot.

The culmination of my results shows that humans are able to cooperate on tasks that take a significant amount of time, such as specializing their forces based on task phase and improved task performance, but humans are less able to cooperate on rejecting quick force perturbations. A person working with a robotic partner programmed to simulate human-human interaction showed the opposite. Even though subjects believed the robot was a human partner, the robotic partner did not elicit the human-human completion strategy in a human. However, simulating dyadic-contraction did improve the disturbance rejection characteristics of a human-robot pair.

CHAPTER 1

Background

Newborn babies tend to flail their limbs in seemingly random directions. Slowly, over days and weeks, the motions become more controllable and consistent. Within months after they are born, infants are able to reach for and pick up objects. With a little more time, toddlers are able to stand up while balancing on two little feet. Throughout this process and the rest of their lives, humans develop internal models for the way their limbs move and interact with their environment.

With repeated interactions, humans build up a large base of experience for interacting with objects in the world [**50**]. Humans learn to anticipate the weight of an object based upon it's size [**32**] and produce a pre-calculated force that will lift the object along a predetermined path [**45**]. Many internal models are constantly updated as humans interact with the environment around them. The learned trajectories of human hands tend to be roughly straight [**99**] and smooth [**24**], which reduces the total jerk on joints. Smooth motion, quick adaptation, and past knowledge allow humans to manipulate known and unknown objects with many internal degrees of freedom.

Humans are also able to adapt to perturbations from the environment, either from unexpected deviations in a trajectory or from an external source. One method of overcoming external forces is to co-contract opposing muscles on the same joint. Co-contraction is a common strategy when interacting with unstable force fields [**25**][**90**].

Many models are able to predict motion and performance, including animating human athletes [**43**], relating performance to the task [**23**], and recognizing human intent [**54**][**16**], but there are many areas that are not yet fully understood. In particular, it is unclear how the forces and motions of two people combine. Do the same models of performance and motion apply to dyads and individuals alike? Do they respond to perturbation forces similar to an individual or do they develop a new strategy?

There are many ways to classify human-human interaction. Two people can interact by speaking, changing facial expression or body posture, shaking hands or hugging, and written word. Some types of human-human interaction have been studied quite extensively. Since this is a large area of study, I want to start by narrowing it down to interactions where two people are present at the same time, so written communication will not be discussed.

Some interactions can occur completely by sight. When performing a task near another person, such as swinging a leg, two people will tend to synchronize their actions [**85**]. Simply by watching another person, mirror neurons in the brain are developing a representation of how to perform that action [**66**][**81**]. Thus, watching a professional sports player (or maybe an animated sports player) can be a good way to learn the motions of a sport.

Similarly, humans are adept at portraying and identifying feelings based upon facial expressions and body posture. In fact, humans will use different gestures depending on the group of people they are speaking to [**64**]. Visually and audibly, humans can work together on a task, but physical interaction, such as jointly moving a bed, is less understood.

When groups of people cooperate on a physical task, there are various ways to combine the redundant controls [**46**][**52**] of the attached limbs. In some airplane designs, the command given to the airplane control surfaces is actually the result of the average position from the two pilots' flight controls [**94**]. Averaging is a simple strategy, but not necessarily the best combination since the two pilots do not necessarily interact; each pilot can perform independently. A better solution likely consists of a strategy that can exploit the redundant abilities of the group. It has been suggested that groups should be able to perform better than individuals since each person in a group can deal with fewer actions [**49**].

Much like a flight control, devices that mediate the interaction between two people often inhibit physical communication. When two people are working together, it is important to include force feedback, which is a method of adding physical communication between the users of a device. Physical communication allows a giver to know when a receiver has control of a drinking glass and can let go [**83**]. Complications can arise in physical communication since a person perceives self generated forces and received forces differently [**84**]; each person may believe the other person is more commanding than they really are. Additionally, a haptic device may not reproduce the forces perfectly, which makes interaction through a device particularly difficult. Force reproduction can become a significant problem when working over great distances, such as in teleoperation. Much of the research in teleoperation deals specifically with issues of how to recreate forces and less on the interaction between the two remote agents.

The interaction between two people directly connected is important and often omitted. Groups of people have been working together throughout history and there have been many scientific endeavors to understand the interactions between two people. Surprisingly, only a few of these studies examine both the force and motion of two people cooperating on a physical task. Sebanz et al. [**88**] state that the ability of two humans to work together "is crucial for our success as individuals and as a species."

Studying human-human interaction can provide a basic understanding of how humans can cooperate and work together physically, which can also explain how individuals work alone [**88**]. Understanding human-human interaction could lead to more effective human-robot interaction. A system designed to adapt to humans could be more intuitive than forcing people to learn the constraints of the system. Humans and robots can interact in a variety of ways, but only a limited number of designs start with human-human interaction.

Throughout the rest of this chapter, I will elaborate on this discussion about how humans interact with the world, other people, and robots by discussing specific experiments. The discussion consists of three parts. The first is human motion control where I will focus on how a single person moves, particularly the path, performance, and adaptation. The second is human-human interaction where I will focus on the forces and motions involved in physical interaction between two people. The third is human-robot interaction where I will focus on the different implementations of robotic assistants along with the degree of autonomous behavior required.

1.1. Human Motion Control

How do humans move? Human motion involves a complex set of actions including the brain, neurons, and muscles. The brain receives sensory information about where the limbs are and recalculates trajectories based upon changing environmental conditions.

Although there is a tremendous amount happening at the neural level, this discussion will focus on the resulting motions and performance.

Humans tend to have smooth motions. Flash and Hogan [**24**] proposed that people will move to minimize the time integral of the square of jerk in voluntary arm movements. This has been verified and used in many studies to analyze human motion, to predict motion, and to generate trajectories for robot arms. Physically this makes sense; even if the acceleration is high, the jerk can be low (i.e. smooth motion), which is gentler on joints and thus causes less wear and tear.

Smooth motion is also desirable for picking up objects. Grabbing a glass of water with a jerky movement would likely lead to a wet floor. Several studies have examined how humans learn to smoothly pick up unknown objects of varying size and density. Gordon et al. [**32**] and Johansson [**45**] show that subjects can quickly adapt to a change in density by lifting an object one or two times. Johansson examines subjects picking up common objects with a different weight than expected (like an empty can when they are expecting a full can). Their setup consists of regular, fully instrumented, objects that can have weight surreptitiously added. The weight can change prior to each subject lifting an object. During Johansson's experiment, if the object was lighter than expected, the subject lifted it quicker than anticipated and, with about a 100 ms delay, subjects corrected the movement. This correction resulted in a distinct change in the path that is not necessarily smooth, which indicates it was not a planned motion. On the next trial, they picked up the object with a smooth motion and subsequent trials with the known object had little deviation.

Johansson also found that the grip force is tightly coupled to the lifting force. At the moment a subject began to apply an upward force, they also applied a gripping force able to transfer the upward force to the object. With objects that tapered up or down, the force was increased or decreased as needed to ensure that the object did not slip.

Gordon et al. similarly tested subjects with known objects including a book, a soda can, a candle holder, etc. as well as novel objects of unusually high or low density. They found that humans anticipate the weight of an object based upon past knowledge of familiar objects and will scale motor commands appropriately. Both Gordon et al. and Johansson's studies show that humans interact with passive objects by revising models developed from past knowledge.

Scheidt et al. [**84**] similarly showed that people only use the past few trials to generate a motor command. Unlike Gordon et al. and Johansson's passive object experiments, Scheidt et al. used perturbations, which were active and unpredictable. Scheidt et al.'s experiment used a two joint robotic manipulator where subjects were instructed to move from one location to a target in half a second. During the movement, the robot produced random perturbations from trial to trial. Scheidt et al. found that only information from the single past trial was necessary to predict performance. The findings that humans can quickly adapt to changes in their environment suggest that two people should be able to adapt to each other quickly. If each member of a pair were to simply treat the other as an unpredictable force, adaptation should occur based only on the previous few trials. Since both subjects are mutually adapting, the forces should reach an equilibrium fairly quickly. Mutual adaptation will be discussed further in section 5.2.

Quickly adapting to an object requires an existing model describing the environment. Throughout life, the brain develops models that are updated to control how it interacts with the world. Brashers-Krug et al. [**10**] examined how humans learn new motor skills and how the process can be interrupted. In a process called consolidation, learning motor skills continue to develop after practice has ended, yet they show that learning a new motor task can be interrupted if it is immediately followed by learning a different motor task. In their experiment, subjects learned one motor task and either learned another immediately (group I) or waited at least four hours before learning the second motor task (group W). The next day, group I actually performed worse on the first task than they did on the first day whereas group W improved. This failure to benefit from the earlier training in group I suggests that the second motor task disrupted learning the first task, a phenomenon called retrograde interference. They concluded that continuously learning both tasks makes the two tasks appear to be one task and they do not fully adapt to either one. This result poses an interesting conundrum for an experiment where subjects will learn two similar tasks in order to compare results. Although it could be argued that an individual task and a dyad task are two separate tasks, I do not expect this to be an issue since my experiments last less than 30 minutes so there is no time for consolidation between tasks and the order was randomized. If experiments were to occur over multiple days, this issue would need to be addressed.

Performance is often measured and compared, as in the above learning study, since it is an important metric of human ability. Athletes are constantly trying to tune their bodies to achieve the maximum performance possible. Over 50 years ago, Paul Fitts found that performance and target size are directly related. Fitts' law [**23**] states that the performance time (t) for a person to move a distance (D) to a target of size (S) varies linearly with the index of difficulty (ID). The ID is defined as the logarithm of $\frac{D}{S}$. Fitts' law is stated as $t = a + b * \log_2(\frac{D}{S})$, where a and b are constants relating to the task and the person. Put simply, Fitts' law says it will take longer to move to a small target than to a large target for a given distance; for a given target size, it will take longer to move a large distance than a small distance.

Over its long history, Fitts' law has proven to be remarkably robust. It has been observed in tasks of varying numbers of degrees of freedom and complexity. I will briefly explain a few of these areas. Accot and Zhai [**1**] found that Fitts' law applies to following a path of nonuniform size including curved paths. Looser [**55**] extended Fitts' law to include 3D computer games and bimanual input. He found that computer games are a valid substitute for Fitts' law tasks that provide significantly better motivation than traditional target acquisition tasks, which can be boring. Hinckley et al. [**42**] analyzed the speed of a scroll wheel on a mouse using Fitts' law. They found that scrolling time can universally be improved by using an acceleration algorithm. Zhai et al. [**106**] analyzed the magnification of the Dock in Mac OS X. They found that the magnification did not change the time required to select an icon, even though the icon was easier to view. Accot and Zhai [**2**] used Fitts' law to discuss how a person can access the Macintosh menu bar five to ten times faster than the Windows menu bar. The Windows menu is on the second line whereas the Mac menu is at the top, thus the menu bar essentially has infinite height (or size S). Gribble et al. [**37**] show that Fitts' law can even relate the amount of co-contraction to target size. There are many other studies, but I will limit my discussion here.

The coefficients of Fitts' law vary from task to task (and to some extent from person to person), but the linearity in $\log(\frac{D}{S})$ is observed over a wide range of ID, and over a wide variety of tasks, which makes it well suited for studying how two people cooperate in a physical task. Surprisingly, there have only been four experiments that have used Fitts' law to study more than one person at a time: Mottet et al. [**60**], Sallnas and Zhai [**83**], Reed et al. [**75**], and Gentry et al. [**28**]. The former two studies will be discussed in section 1.2 and the later two studies will be discussed in sections 4.1 and 4.2.

Although there are many models describing how individuals move and interact with objects, Sebanz et al. [**88**] state, in a review of joint action, that it may not be possible to fully understand how the mind works by studying people working in isolation. Understanding how two people work together may clarify how an individual works alone.

1.2. Human-Human Interaction

Human-human interaction is a well studied field. There have been studies examining facial expressions between people [**36**][**5**] and even recreating facial expressions in robots [**53**][**8**] as well as studying gesture based interactions [**12**][**64**]. However, much of the focus has been on interaction at a distance. Sebanz et al. [**88**] provide an overview of joint action, which surprisingly, contains few studies focusing on physical interaction. In this section, I will compare the findings from several studies that have examined two people interacting via force and/or motion.

In the previous section, I discussed how a single person will reach for and move a passive object. This relatively simple task becomes more complicated when two people are involved, such as shaking hands or exchanging an object. In a study designed to understand a coordinated action between two people, Sallnas and Zhai [**83**] examined two people exchanging an object in a haptic environment. They used two Phantom devices, rather than direct physical contact, to study performance and the effects of force feedback during the handoff of a virtual object from one person to another. The subjects' task was to exchange various sized objects without dropping them.

According to Sallnas and Zhai, force feedback significantly decreased the error rate for exchanging objects between two people. Without force feedback, 10.7% of the objects were dropped and only 5.2% were dropped when they had force feedback. The drop rate with force feedback seems particularly high if they are trying to compare this to the real world. They give an example of a coordinated action in real life where a flight attendant gives a passenger a cup of coffee, but their percentages suggest that the flight attendant would drop 1 out of 20 cups (probably would not be a flight attendant very long if that continued). The force feedback case does not seem to mimic the real life situation very accurately, but the difference does highlight the importance of a haptic interaction when coordinating physical actions among two people.

Sallnas and Zhai also found that it took about 2.8 ms to move the object both with and without force feedback. Without force feedback, it is difficult to determine who has control of the object, so it surprises me that the time was similar. I would have expected a slightly longer time when they could not feel who has control of the object. They also found that performance time did correlate to the size of the object as expected based upon Fitts' law.

The experiment Sallnas and Zhai conducted consisted of the subjects exchanging an object within certain spatial limits, yet, in everyday life, these spatial limits do not exist.

In many daily activities, the target would be moving, not stationary, much like Mottet et al.'s [**60**] experiment. Mottet et al. used two manipulandum and two rows of LEDs to perform a Fitts' law study with moving targets. Each subject could see their own position and the position of their target where each subject's target was the position of the other person's manipulandum. As each person moved toward the target, the target also moved closer to them, much like two people preparing to shake hands. Mottet et al. showed that this type of task does obey Fitts' law, meaning that for small targets, it takes longer to reach each other than it does for larger targets. They showed that dyads have a stable and tunable interaction by analyzing the variance of the end-points of the motion. The variance was high for low ID's, which indicates that the individuals were operating independently of the other person. For high ID's, the variance was low, which indicates that the individuals were coupling their motion in order to work together. Although this study did not involve physical contact, the tunable nature of the dyadic interaction suggests that humans do have an ability to tune their actions to another person during a physical task in order to cooperate.

Many tasks require two people to cooperate in order to successfully control one object. Shared control occurs in learning to operate some kinds of equipment, such as in the teaching and learning of helicopter control and piloting airplanes. Wegner and Zeaman [**101**] discussed some dual control tasks in everyday life, including the two-handled saw, the see-saw, and the balancing of a tandem bicycle. Shaw et al. [**91**] adds couples dancing and teaching how to swing a golf club. Other everyday activities exhibiting dual control are maneuvering and positioning a board or other object and placing a bed sheet on a bed. In some of these examples, the participants establish a working relationship in which they must compromise (e.g. positioning the bed sheet symmetrically) or in which they must divide authority according to task phase.

Wegner and Zeaman were interested in comparing groups of two and four people working together to individuals working alone. They used a "pursuit rotor" task in which a subject tried to cause a metal stylus to stay in contact with a metal path embedded in a rotating turntable. The task was essentially two-dimensional. A subject manually controlled the stylus via a handle. They added multiple handles to the stylus so that individuals or teams of two and four could be tested similarly.

Wegner and Zeaman found that dyads stayed on the path significantly better than individuals, and quads significantly better than both dyads and individuals. They suggested and methodically tested several possible explanations, including social facilitation (i.e. having an audience), guidance, leadership, group problem solving, and skill learning/transfer. They were unable to suggest a satisfying explanation, possibly since they did not measure the forces exerted by the participants. Interestingly, Wegner and Zeaman reported in a footnote that some of their subjects said that the stylus felt mechanically more stable in team conditions. The investigators tried providing similar stability through mechanical means, but found only a negative effect upon path tracking ability.

Similar to path tracking, the variability of the path tracking increased as the number of subjects increased in Wegner and Zeaman's experiments. Variability was measured in the later trials to reduce the effect of learning. Based on the results from Mottet et al. showing that low ID leads to high variability, I speculate that Wegner and Zeaman's task was a relatively easy task indicating that the subjects may have been working independently. This independence may account for a group's increased ability to track a path. If the subjects were working independently, the joint output of all group members would be a combination of the individual member's independent movements. Independent combinations mean that each member's commands would be diluted, particularly the extraneous movements. Any non-ideal motion would be diluted and would make the group less prone to errors and is thus a possible explanation of the increased ability to track a path.

In Wegner and Zeaman's experiments, they were focusing on the motions and path tracking ability of groups and did not analyze the forces involved. Shergill et al. [**92**] performed an experiment examining the forces between two people without motions. In this experiment, two subjects each put their finger in a lever attached to a force transducer. Each subject was told to push with the same force they felt, but the subjects were unaware of the instructions given to the other subject. Alternately, as instructed, each subject applied the force. Shergill et al. found that in every case, the forces escalated from trial to trial. They explained that the subjects are reporting the true perception of the force and the increasing force is due to neural processing.

In a second set of experiments, Shergill et al. applied a brief constant force to the tip of a participant's finger and asked them to generate the same force from another finger on the other hand. The subjects consistently generated too much force. Shergill et al. suggest that self-generated forces are perceived as weaker than externally generated forces, which has implications for two people working on a task together. Each partner will always feel that they are contributing less to the overall task than the other member, even though that is not and cannot be the case. When working cooperatively with a partner on a task, each subject may want to contribute equally, which could lead to an escalation in performance. This is only speculation since Shergill et al. only discussed escalation in terms of force.

In some applications, forces can relay vital information. If the perception of forces is reduced, as in Shergill el al.'s study, or the transfer of forces is hindered, communication can be significantly diminished. Fly-By-Wire (FBW), a design for airplane control, eliminates the direct mechanical connection (and thus the forces) between the pilot and the plane's control surfaces and also between the two pilots. Depending on the configuration and design of the FBW system, the flight sticks allow little or no haptic interaction between pilots. Summers et al. [**94**] conducted a series of experiments on pilots using a Flight Simulator at NASA Ames Research Center. They examined four different cases, ranked by the pilots in order from most preferred to least preferred: coupled, uncoupled with a disconnect switch, uncoupled with priority logic (essentially the largest input wins), and uncoupled (average of inputs). The pilots significantly preferred the coupled (haptic) FBW more than the uncoupled (non-haptic) FBW. They also found a significant performance decrement when using uncoupled FBW. These results are similar to the object exchange results from Sallnas and Zhai.

The pilots in Summers et al.'s study least preferred uncoupled FBW without force feedback. Interestingly, Glynn and Henning showed that this strategy, average of the commands without haptic interaction, resulted in faster and more accurate task execution than one person alone. In their experiment, the acceleration of an inertial mass is controlled by averaging the force applied by each partner to a joystick. Glynn and Henning setup a maze for subjects to maneuver. There was no force feedback, only visual feedback, so the subjects were not physically interacting. The resulting force on the simulated mass was the average of the dual human input system without haptic interaction. They found that, on average, teams complete the maze in 159.7 seconds while individuals completed the maze in 174.2 seconds, a statistically significant increase of 14.5 seconds for teams. They also found that collisions were reduced from 484 for individuals to 334 for teams. It makes sense that the collisions are less since any faulty motion would be diluted by the other member's action. Statistically, the chance of both subjects producing the same collision-producing motion would be less than that of one person acting alone. The improved performance may also be explained by a diluting effect. Since each subjects' action is diluted, they may not see the effect of their motion, and thus may attempt to push harder than they would alone. This could lead to each subject pushing slightly harder and thus completing the task faster. Since there is no physical interaction between the subjects, neither of them would feel this effect.

Even though performance is better for two people with an average of their commands, it does not necessarily mean that pilots, or other shared control tasks, should use this control mechanism. Field and Harris [**22**] conducted a survey among 157 commercial aircraft pilots, some who have flown planes with a direct mechanical connection and some who have flown using non-haptic FBW. Field and Harris found that communication was lost when using FBW, which could adversely affect the pilots' awareness of current situations. Many pilots in the study stated that being able to feel the motions and forces of the other pilot was important and useful. Many pilots also stated that they wanted to be able to feel what the autopilot was doing so they could determine if the plane was flying correctly. In a direct mechanical connection, the pilots can feel what the other person is doing as well as a response from the plane itself. Shergill et al. explained how an individual can separate their forces from an external force, but I have not found any studies that explain how the pilots are able to separate the force in a flight stick of the other pilot from the force feedback of the plane, which could be an interesting problem to explore.

Glynn et al. [**29**] extended the work of Glynn and Henning's maze study by adding force feedback, so the subjects were interacting through a haptic display, which adds a spring in-between their interaction. They compared four conditions including both force and position control each with and without force feedback. During force control, the added force feedback increased errors while performance time was unchanged. Glynn et al.'s only explanation for larger path deviation was that the feedback interacts with the dynamics of a second order system in complex ways. Just like I discussed about the pilots in Field and Harris's study, it is possible that each member has difficulty separating the force feedback of the device from the other member's forces. Glynn and Henning report that, during position control, the added feedback decreased errors and improved performance. The force feedback is possibly more beneficial during position control since the position does not overlap with the force. With two people interacting physically, they are able to communicate both on position and force, thus there should be no detrimental interaction like there is in the force control experiment.

Although couples dancing has been mentioned as an example of a physical interaction, it had not been studied until Gentry [**27**] examined how this interaction works between dancers. She describes dancing as a finite state machine where dancers move between a limited number of poses and interact through force and motion. Dancers coordinate their actions through various elements, some internal to the dancers and some external. The rhythm, or beat, of the music is an external event that synchronizes the motion of each dancer's movements and is the strongest element coordinating their actions. The poses are a predetermined position of the bodies of the dancers. Some dances, such as the Waltz, have only one pose, whereas Lindy Hop has multiple. The leader moves from position to position while indicating what to do next. The transitions are coordinated based on previous knowledge of a small set of moves.

The physical connection in dancing is maintained though the follower's right hand holding onto the leader's left hand. This physical connection allows the leader to send messages to the follower as well as for both partners to exchange energy. A good follower will keep her hand in the same position relative to her body, which allows the leader to communicate. Most of the communication is based on haptic cues even though the dancers can see each other. Gentry performed experiments on experienced couples dancing blindfolded and found that they were capable of performing quite well with only haptic communication[1]. The experiments in Gentry's thesis are directly relevant to several of my results and, thus, will be discussed further in the relevant sections of 4.1.3, 4.4, and 4.5.

In lower limb rehabilitation, the physical therapist must apply large forces as he guides the motion of the patient's leg through a walking motion. Several previous efforts have tried to create a robot that itself performs aspects of physical therapy, such as the Lokomat [**62**]. Here the robot could take over the task, which would reduce the labor intensive part, but physical therapists value hands-on interaction. It is important to allow the communication of force and motion to the patient, and for the physical therapist to

[1]In talking with Sommer Gentry, she noted that the blindfolded dancers could dance well, but had to have some help in avoiding walls.

feel the patient's response. In the experiments discussed in this thesis, I will explain the important aspects of how two people physically cooperate and compromise to complete a task.

1.3. Human-Robot Interaction

Early robots were not designed to interact with humans. In fact, many factory robots are fenced off to prevent people from interacting with them as they diligently perform their programmed task. In 1979, Robert Williams was the first person killed by a robot [**47**]. Since then, there has been a significant amount of research on how humans and robots can interact and cooperate with different control laws for programming a robot.

There are many different methods to design robots to work with humans. Yanco and Drury [**103**][**104**] discuss several different combinations of how humans and robots can work together. All of their classifications are based on one or many humans giving an order to one or many robots. The robots perform the task and send sensor information back to the human operators. Their taxonomy includes elements relating to the task type, the ratio of humans to robot, and the level of interaction among all the agents. Surprisingly, they do not classify the situation in which there are two humans giving an order to a robot that is acting on those same humans, which is essentially teleoperation or a physical therapist interacting with a patient through a force amplification device.

Not only is the configuration of the humans and robots important, but the mechanism for communication among all the agents is essential. Klingspor et al. [**48**] discuss various methods for communication between humans and robots, including programming commands, listening to natural words, and demonstration. They aim to make it as easy and intuitive as possible for people to work with a robot. To make it easy, human-robot communication in a shared task should follow the implicit human-human communication standards, which are poorly understood. The experiments I explain in chapter 4 will help to decipher human-human physical communication and cooperation. The rest of this section will focus on other studies that have examined human-robot physical interaction.

Rahman et al. [**72**][**71**] started to study humans physically interacting using a 1 DoF placement task. They characterize the humans as either being a master or a follower. A master controls the position of the object. The follower tracks the motion of the master with impedance control. The impedance model of the follower robot, discussed by Rahman et al. [**69**], changed the stiffness and damping throughout the motion. The impedance was high in the beginning and was essentially zero after 0.4 seconds. They determined the impedance mode by analyzing the resistance an arm will apply when it is led through a given path. By using the follower characteristics of a human response found in their studies, Rahman et al. [**70**] implement the same response in a robot. Their aim is to make the robot imitate the response of a human when interacting with another human. Throughout this series of experiments, they do not discuss physical communication between the two humans and they do not take into account how the completion strategy changes from working alone to working with a partner. I believe physical communication and strategy changes are important and will discuss them further in section 4.3.

Lifting and moving a large or awkward object cannot be done individually, so people either get help from another person or, in some cases, a robot. Takubo et al. [**96**] demonstrate a robotic assistance system to empower a single human to work with a single robot to handle unwieldy objects. Their control strategy consists of a dominant human who leads while the robot emulates a virtual non-holonomic constraint. Initially, they kept the device in the horizontal plane allowing only 2-D motion. They experimentally validated the usability of this method by showing that subjects could transport an object to an arbitrary location and orientation with only a small amount of learning by using "skills similar to using a wheelbarrow."

They then extended the method to 3-D space. Interacting with the robot in 3-D space requires a substantial amount of planning in order to place it in an arbitrary location. In order to move the object directly down, the subject must first lower his end of the object and pull the object down along the declined angle to get the robot end at the correct height. At this point, the subject must raise his end and push the object back to the original horizontal location. This same strategy must be completed for each desired direction of motion. Although this control was inspired by human cooperation behavior, they are not trying to simulate a human. Two humans cooperating can perform the task with one person attempting to act as a non-holonomic constraint, but this is not the most natural strategy. There are other cooperative strategies available. Two people can jointly maneuver an object to a desired location in a much more direct path, but this has the complicating factor that both people must know the goal destination. Learning to implement other strategies for how two humans cooperation haptically could help in creating more intuitive communication between a robot and a human.

A recent study has started to examine how a robot could act like a human who lacks knowledge of the end goal. Corteville et al. [**16**] give an example of a blindfolded person who attempts to help a partner complete a task. They discuss how the blindfolded helper would wait for a trigger that indicates when the motion has begun before assisting. Throughout the task, the helper does not know where the leader is going, but the helper will begin to guess based upon the motions of the leader. Using an estimate of the motion, the helper will join in on the motion. Corteville et al. programed a robotic assistant based upon the motion estimation of a blindfolded human helper. The robotic assistant assumes that the path will follow a minimum jerk trajectory. An admittance controller provides assistance once an estimate of the operator's path is obtained. Based on the input from the subject, the controller will provide a scaled force along the predicted path that will aid in reaching the target. This design has potential to directly take human control and apply it to a robotic controller. In its current form, the robot motion is not adaptable to unknown targets though. They assume that the start and end points are known to the controller, which are necessary for calculating the robotic trajectory. They mention that a first step of building versatile robots consists of "building up experience with simple tasks." I believe this experience of learning to implement simple human interaction is important for improved human-robot interaction.

Steele and Gillespie [**93**] and Griffiths and Gillespie [**38**] have demonstrated a haptic steering wheel that communicates bi-directionally with a driver. That is, the driver is able to exert control over the direction while simultaneously extracting information from the steering wheel. They designed the steering wheel to take advantage of the tactile and kinesthetic sensory capabilities of humans. Typically, driving heavily relies upon vision, so offloading some of the communication to another sense would be beneficial. This steering wheel provides information about the intentions of the controller. The controller aims to stay as close as possible to the reference path regardless of obstacles that may be in the path. It essentially gives the driver an error signal from the reference path. This approach has the same complicating factor discussed above that requires the controller to know the reference path.

Regardless of how the controller knows the path, using the haptic steering wheel reduced the visual demand and reduced the path error. Path error was the sum of the total number of obstacles hit and the average deviation from the reference path. Visual demand was measured by recording the number of times subjects viewed the road. By default, the road was not viewable. The subjects had to press a button to view the road, so visual demand is the sum of the total 'glances' at the road. Steel and Gillespie report that visual demand was reduced by 42% and path error was 50% less when using the haptic steering wheel compared to a passive one. Similarly, Griffiths and Gillespie report that visual demand was reduced by 29%, path error was reduced by 30%, and reaction time was improved by 18 ms.

Many studies have focused on bringing physical information from a remote task site to the operator via force feedback. Barnes and Counsell [**6**] explore a task similar to the haptic steering wheel, but with a haptic joystick for teleoperation where the human and robot are separate. The robot can receive commands from a human to perform tasks while working in an environment that is dangerous to humans. They performed experiments with four different arrangements: 1) teleoperation without haptic feedback, 2) teleoperation with haptic feedback, 3) telerobotics without haptic feedback, and 4) telerobotics with haptic feedback. All experiments were performed with a joystick and monitors. With haptic feedback, the force from the robot was transmitted back to the user via the handle. In the telerobotic case, the robot is autonomous for collision avoidance, but requires the operator to issue other motion commands.

Barnes and Counsell show that the fastest mode was telerobotics without haptic feedback, which is similar to the uncoupled with priority logic case in Summers et al.'s study. The only difference is that Barnes and Counsell's partner is a robot, whereas Summer et al.'s partner was another person. The increased performance of telerobotics without haptic feedback is surprising since this mode offers the least amount of control and information for the operator. The robot may take over to avoid an obstacle at any moment without the operator knowing it due to the lack of force feedback. Adding haptic feedback comes with a slight penalty in performance, but grants the operator greater knowledge about the system, which the authors state is very important. Although the telerobotics modes were not the fastest, they were the safest since they had zero collisions.

Teleoperation is tangentially related to human-human physical interaction, so I will not dwell long on it. There are many aspects of teleoperation, ranging from dealing with internet latencies [**19**][**41**] to reducing mental workload [**18**]. Many of these topics are not related to physical interaction, but there have been several teleoperation studies examining two subjects cooperatively working on one object in a shared virtual environment. Basdogan et al. [**7**] performed experiments where two people jointly manipulated objects in a virtual world. The interaction consisted of visual and haptic feedback via the virtual environment. Basdogan et al. found that haptic interaction gave the users a better "sense of togetherness." Hubbold [**44**] examined two subjects carrying a stretcher in a virtual world. The interaction was both visual and haptic via the virtual environment. Hubbold found that haptic interaction helped give the subjects a better perception of the other member and a "sense of sharing." Many of the details of these studies delve deep into other issues, so I will conclude this discussion of teleoperation.

Figure 1.1 summarizes configurations of several of the above studies which investigated the performance of two people cooperatively manipulating a single object. Two people physically cooperating yield the greatest benefit from working with a partner.

1.4. Joint Biological Systems

Thus far, I have discussed previous research on how humans move and how humans interact with other humans and robots. In looking into how people interact physically, I came across an interesting group of biological systems that are able to effectively carry large loads in groups, the species called *Pheidologeton diversus*. They are otherwise called the Asiatic Ant. In a paper outlining how ants transport food in groups, Moffet [**58**] wrote:

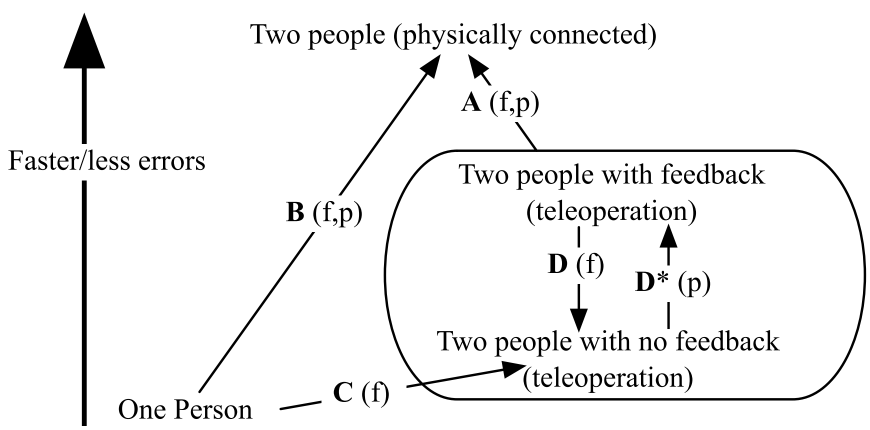

Figure 1.1. Summary of studies comparing the performance of different forms of human-human interactions. All studies involved a real or simulated second order inertial system, except for D*, which was a zero-order system. Comparisons containing an 'f' indicate that their partner's force was conveyed to their partner while a 'p' indicates position was conveyed in the feedback. {**A** - Summers et al. [**94**] and Field et al. [**22**]; **B** - Reed et al. [**76**] and Wegner and Zeaman [**101**]; **C** - Glynn et al. [**30**]; **D** - Glynn et al. [**29**] and Sallnas and Zhai [**83**]}

"Group transport (the carrying or dragging of a burden by two or more individuals) is better developed in ants than in any other animal group." Since their ability is so well carried out, I thought it was worth briefly discussing.

Moffet [**59**] found that the Asiatic Ant is very effective in group transport. The most notable finding from his study is that each ant working in a group carries more weight per ant than an ant carrying an object alone. One ant alone can carry five times its body weight. Yet, in one example, 100 ants worked together to carry a worm that weighed 5000 times the body weight of an ant; each ant carried 50 times its own body weight. Moffet described this increasing load capacity of each individual ant by the equation $W = 26.64 \times N^{2.044}$ where W is the object weight and N is the number of carriers. The ants can carry exponentially more weight with increasing ants until about 11 ants are working together. Moffet speculated that the space around the perimeter of an object limits the effectiveness of ants that can actually work together, thus group effectiveness grew slower for groups larger than 11 ants.

Moffet also found that transport velocity was fairly constant for groups of 2-10 ants, which were twice as fast as individual ants. Groups of 11 or more slowed down to less than half the speed of smaller groups. His efficiency metric of velocity $\times$ object weight showed that the ants were increasingly efficient as more ants worked together upto groups of 11 ants, at which point effectiveness increased at a slower rate, presumably since there is limited space around an object.

Not all ant species use group transport - many break the food down into smaller pieces and carry them individually. In fact, in other ant species if a second worker ant grabs an object, the first ant's progress is halted until the second worker releases the object. It is only some species of ants that have developed this ability to cooperate to carry large work loads.

Although ants cooperating to carry large loads is not directly related to studying how humans interact or how a robot can cooperate with a human, there have been several studies that used ants and other insects as inspiration for developing swarms of smaller robots to work together [**51**][**9**]. These studies have had some success, but I will not go into the specifics since this thesis is primarily concerned with a human working with another human or a robot.

CHAPTER 2

Experimental Apparatus

As discussed in chapter 1, there has been a lack of research examining both the forces and motions involved when two people cooperate on physical tasks. In order to study how two people work together, I designed, built, and programmed a device to allow haptic interactions while preventing other forms of communication such as visual and verbal, which have been extensively studied. My device can also manipulate the user's forces to further probe their interaction and test whether I can create a useful robotic partner based upon human-human interactions.

I chose a 1 DoF device to simulate target acquisition tasks as the simplest configuration which could reveal physical dyadic interaction. Many of the earlier motion studies have examined reaching in one person, much of which was done in 1 DoF. A similar setup would allow me to use the already established frameworks for analyzing human motion. One particularly well established description is Fitts' law, which will be the first experimental results discussed in chapter 4.

This chapter will detail the specifics of the device that allowed me to attain the human-human interaction results in chapter 4 and the human-robot results in chapter 5. This is an abbreviated description of my device. A more elaborate description of my device can be found in my Master's thesis [**74**].

2.1. Mechanical Overview of Device

My device, shown in figures 2.1 and 2.2, consists of two handles that can spin freely about a center shaft. Two users can each hold one of the two handles. The force from each user can be measured by strain gauges mounted along the handles. An encoder, mounted to the motor under the table, measures the position of the handles. The velocity is found by differentiating the position signal. An accelerometer, mounted on the bottom of one of the handles, measures the acceleration of the crank. The motor can apply a force based upon any of these parameters.

The handles are connected via a rigid link that can spin freely at the center where a direct drive motor[1] is attached under the table as show in figure 2.2. A cross sectional view of the device (figure 2.3) reveals the linkage shaft between the motor and the handles. A combination of set screws and keyways allow the linkage shaft and motor to spin together

[1]Moog G413 Series brushless motor.

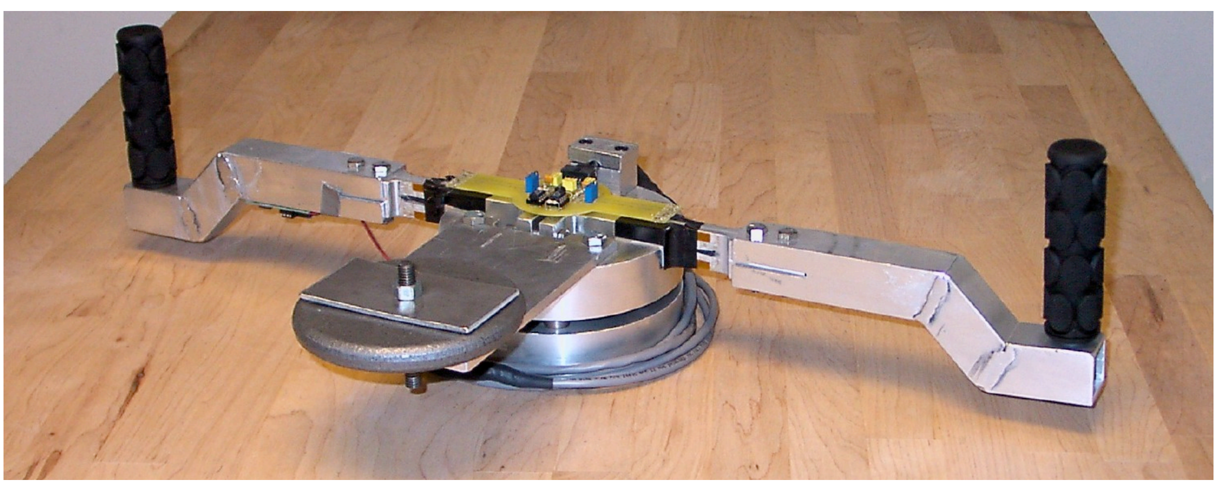

Figure 2.1. Top view of the mechanical device used for experiments. Two users grab the handles, one on either side. This device can measure the force applied to each handle, the angular acceleration of the crank, and the angular position of the crank.

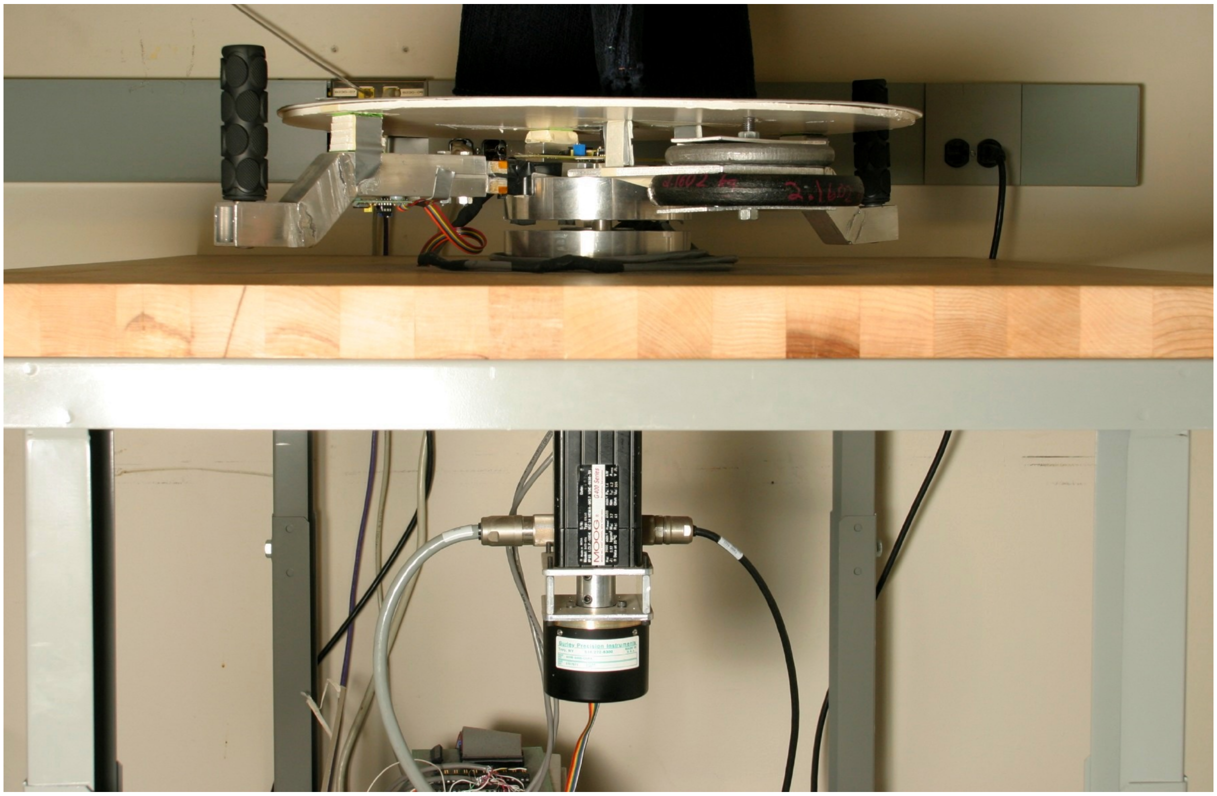

Figure 2.2. View of the device from the side, both above and below the table. The direct drive motor is attached under the table and out of the view of the users.

while transmitting torque. The motor can provide 3.7 Nm of continuous torque and 13.0 Nm of peak torque.

The linkage shaft is made of steel since it is supporting the entire moment and torsion of the handles including any force applied by the users. All other machined parts are aluminum. The shaft is held by two angular contact bearings[2] arranged so the lines of contact intersect outside the ball bearing envelope.

[2]Single row angular contact ball bearings - ABEC1 for 15mm shaft diameter. 35mm outer diameter (item 6680K13) from McMaster.

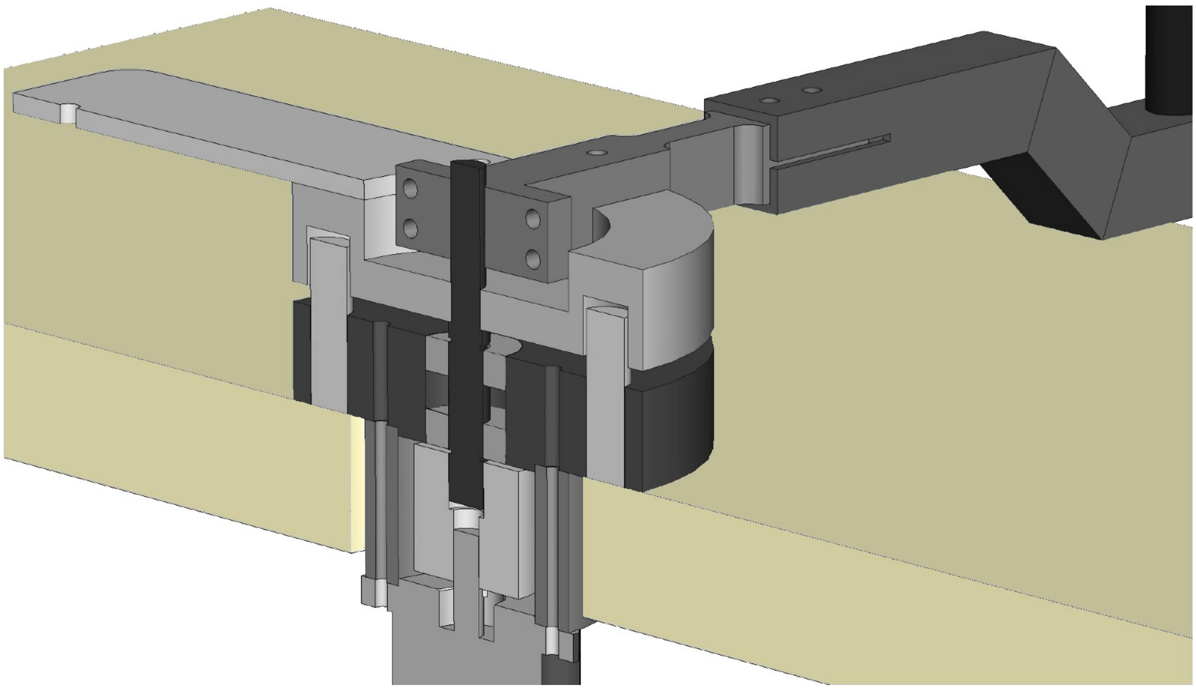

Figure 2.3. Internal view of the center of the device. This view shows the internal connections between the motor, handles, and table. The bearings support all forces and moments from the handles.

Steel shafts prevent the handles from spinning more than 110° to ensure safety. Each shaft contains a rubber pad to cushion the impact at the edge of the workspace. The groove in the bottom of the top plate, shown in figure 2.4, shows the extreme positions of the top plate. Even if the motor were to spin wildly out of control, the handles are not able to spin in a complete circle. This method to prevent spinning is also necessary to prevent the power and sensor cables from getting tangled around the circular plates since the strain gauges and accelerometer are located on the spinning handle bars.

On either side of the top circular plate, the protruding bar of the handles narrows for 3 cm. This narrowing helps to concentrate the bending of the bar without drastically increasing the maximum deflection at the handles and gives a better reading from the mounted strain gauges.

In order to be able to change the moment of inertia of the handles, I added a weight plate. It sticks out 12 cm. The weights are held in place by a thin aluminum plate pressed down to the weight plate with a bolt.

Each handle is 31.75 cm from the center of rotation. Both handles consist of freely rotating bicycle handle grips connected to a rigid aluminum shaft. The handle lengths are adjustable, but I never changed the length.

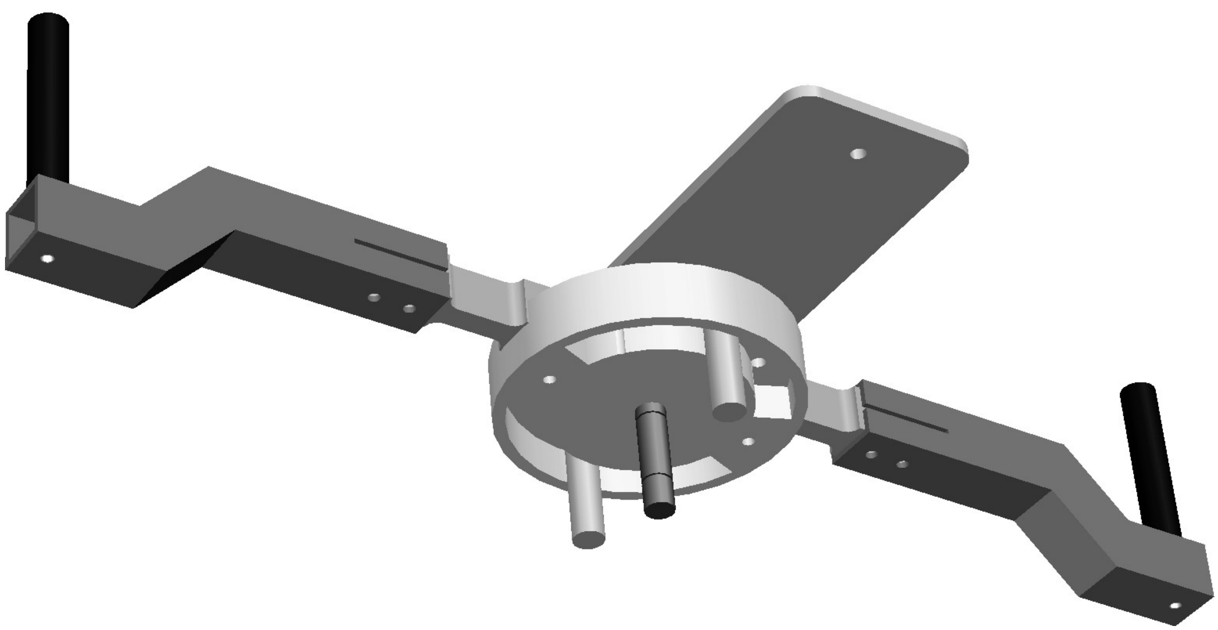

Figure 2.4. View of the crank handle bars from underneath. This reveals the groove in the bottom of the top plate that prevents the crank from spinning in a complete circle. The removable stopper shafts are mounted to the bottom plate (not shown) to limit the range of motion.

2.2. Data Acquisition

I control this device using a PC running QNX with a ServoToGo card for input and output. The computer samples the sensors and outputs at 1 KHz based on hardware interrupts. I record data to the hard drive for later analysis. The graphics are updated at 60 Hz. The computer screen output is mirrored on a CRT monitor and on the projector from above. I designed the software to be run from a second computer[3] via a terminal window to allow the entire QNX screen and projector to be used for displaying graphics. The remote login user can monitor the progress of the experiments and sensors as well as send commands to the program.

I mounted four foil strain gauges[4] on each handle to provide force information about each users input force. I arranged them in a full Wheatstone bridge configuration with two strain gauges on each side of the bar. This configuration only measures bending in the direction tangent to rotation while other directions are cancelled by the Wheatstone bridge. I use a 0-20 Ω potentiometer on one branch of the bridge to zero the bridge. I glued the strain gauges to the bar using a heat cured epoxy.

An instrumentation amplifier compares the outputs from the Wheatstone bridge with a gain of 500. The output is then filtered using an RC filter with a cutoff of 338 Hz. This setup has a resolution of 0.16 Newtons and a theoretical maximum reading of 586 Newtons for each handle, although I have only calibrated it and used it with a force less than 200 Newtons.

[3]Using telnet from Mac OS X inside of X11 works very well. It allows all of the Curses terminal displays to work with the exception of color.
[4]Omega item number SG-7/350-LY13. 1.4 cm by 0.8 cm.

I mounted an encoder on the bottom of the motor (visible in figure 2.2) to determine the position and velocity of the handles. This is a Gurley Precision Instruments encoder with 120,000 counts per revolution.

This encoder provides very clean position data and, after differentiating, provides a clean velocity data. Although this is a very high quality encoder with a high resolution, the signal cannot be double differentiated to achieve a clean acceleration signal. The acceleration signal contains too much noise to be useable and after filtering, the signal contains too much time delay to be used in realtime feedback. It does provide a nice acceleration profile for offline analysis.

Since the encoder cannot be double differentiated in realtime to attain a good acceleration signal, I decided to add an accelerometer to one of the handles. Only one is required since the handles rigidly move together. I used a Analog Devices[5] accelerometer. I filter the output from the accelerometer using an RC filter with a cutoff of 234 Hz. This setup has a resolution of 0.0013 g's and a maximum reading of 2 g's.[6] Calibration of the accelerometer involved rotating it 90° in each direction to measure the acceleration of gravity. To double check, it was then measured against the offline double differentiated and highly filtered encoder signal. The signals matched well.

2.3. Canceling and Amplifying One User's Force

When the motor applies a force, both users can feel it. Due to the rigid configuration of the crank, it is not possible to selectively apply a force to only one user. I can however command a force based upon the measured force from one or both of the handles. I can

[5] Analog Devices IC accelerometer LP 14-cerpack, item number: ADXL202JQC-1-ND.

[6] The accelerometer specifications only claim up to 2 g's and state that it may be non-linear outside of this range.

nearly cancel the effects of one user's force by applying a force in the opposite direction, such as in equation 2.1:

$$\tau_m = -xF_B * d \tag{2.1}$$

Where τ_m, F_A, and F_B are defined in figure 2.5 and d is the distance to the handles (31.75 cm). x is the proportion of cancelation. x can range from 0 (no canceling) to 1 (complete canceling). This makes the canceled user feel very weak and they have to apply much larger forces to have any effect on crank motion.

Alternatively, I can amplify the effect of one of the participants by applying a torque in the same direction he is pushing, such as in equation 2.2:

$$\tau_m = yF_A * d \tag{2.2}$$

Where y can range from 0 (no amplifying) to ∞ (very large amplification), although practically I do not set y to be larger than about 7.

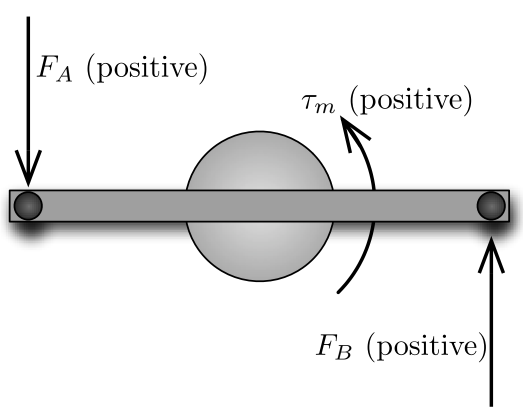

Figure 2.5. Definition of forces and torques on the crank. τ_m is the motor torque calibrated to match an equivalent force at the handle. F_A is the force at one of the handles and F_B is the force at the other handle.

Combining equations 2.1 and 2.2, and setting $x = y = n$ allows a variable amount of control to be assigned to each user ranging from equal control to almost canceling one user's input force[7].

$$\tau_m = n(F_A - F_B) * d \tag{2.3}$$

The effect of one user's force can be almost entirely canceled by using equation 2.3 with $n = \pm 0.9$. For stability reasons, one user's force cannot be completely counteracted (i.e. $n = 1$). Using this method, the motor can assist one user to make them apply $1/20^{\text{th}}$ the force of the opposite user (e.g. if user A is pushing with 2 Newtons, then user B has to push with 40 Newtons to keep the handle from moving). Depending on the value of n, the following effects will be felt.

$$n \begin{cases} \leq -1 \implies & \text{unstable} \\ = -0.9 & \downarrow \\ \vdots & \text{decreasing advantage to B} \\ = -0.1 & \downarrow \\ = 0 \implies & \text{no advantage} \\ = 0.1 & \downarrow \\ \vdots & \text{increasing advantage to A} \\ = 0.9 & \downarrow \\ \geq 1 \implies & \text{unstable} \end{cases} \tag{2.4}$$

[7]Changing who has control is the most common demo, particularly for younger groups. When one person can use only one finger to resist the force from an entire arm, the effects are immediately noticed. Appendix A describes my demonstrations in further detail.

2.4. Simulated Inertia

Similar to canceling a users force, I can use the accelerometer to make the device feel heavier or lighter by applying a force proportional, but negative, to the acceleration. When the simulated inertia is set to make the device appear lighter and no user is holding the handle, it becomes unstable. When a user holds it, the inherent impedance in the human musculoskeletal system stabilizes the system.

2.5. Filtering the Motion

I can use the user forces and crank velocity to filter out certain frequencies of motion. By filtering one of these, the motor will apply the opposite force back to the users, which will either make the crank feel damped (if using velocity) or heavier (if using forces) at certain frequencies.

2.5.1. Cancel Steady State Forces

As I will further discuss in section 4.5, two subjects in a dyad can apply an equal and opposite force against each other with no resulting motion. Due to the rigid connection between the two handles, there is no way to completely eliminate these opposing forces from both users simultaneously. However, using the motor, it is possible to eliminate the steady state force from one of the two users by applying a force opposite to that of one user.

In order to cancel the steady state force from user A, user A's force will be measured and low pass filtered. The opposite of this slow changing force will be applied, which will nearly eliminate the effect of any slow changing force this user is applying. User B will not feel any steady state force from user A.

Alternatively, canceling the steady state forces could be position dependent. I do not want to interfere with the two subjects moving the handle and am only interested in canceling the steady state force while the handles are in the target and the subjects are moving relatively slowly. By slowly canceling the steady state forces as the subjects approach the target, the overall movement will not be impeded, yet one of the subjects will not feel any force from the other subject while in the target.

2.5.2. Filter High Frequency Oscillations

I have implemented two methods to reduce the oscillations that will be discussed in section 4.4.

2.5.2.1. High Pass Filter the Force. One method is to high pass filter the force from one user and apply the opposite force. This has a slight noticeable delay. When a user pushes on the handle, the handle does not respond as expected. Because of the high pass filter delay, the handle will respond quickly and then the delayed force will counteract the motion which feels very unnatural, almost like the device is adding energy to the user. The large high pass filter delay makes it hard to display an output force based upon the input force from the user.

2.5.2.2. Frequency Dependent Damper. The other method, which feels more natural, is to create a frequency dependent damper. A standard damper is $\tau_m = -b * V$, where b is a constant and V is the velocity of the crank. Equation 2.5 shows a frequency dependent damper which is instead based upon the frequency that the crank is moving.

$$\tau_m = -b * \hat{V} \tag{2.5}$$

Where $\hat{V}$ is the high pass filtered velocity.

Equation 2.6 shows how this would be programmed with a recursive filter.

$$\tau_m(t) = -b * (a_1 V(t) + a_2 V(t-1) - c_1 \tau_m(t-1)) \tag{2.6}$$

Where a_1 and a_2 are the filter coefficients in the numerator and c_1 is the filter coefficient in the denominator. This is a high pass filter, so high frequencies are damped. The net result is a low pass frequency dependent damper.

Since the motor output is no longer a function of input force, the delay is not as noticeable and achieves the desired effect. Similarly, a low pass frequency dependent damper of the same design allows gross motion, but attenuates high frequency oscillations such as the oscillations from a subject inside the target at the end of the trial.

2.5.3. Choice of Filter

The frequency dependent damper is using a first order high pass Butterworth filter with a cutoff of 2 Hz. This works well to reduce high frequency oscillations while leaving low frequency steady state motions mostly untouched.

Using a second order Butterworth filter does not work nearly as well. In response to a step function, it overshoots. A similar response occurs in a Chebyshev filter. Recursive filters higher than first order tend to have a similar response. Figure 2.6 shows how a first, second, and fourth order high pass filter respond to a step response. The overshoot

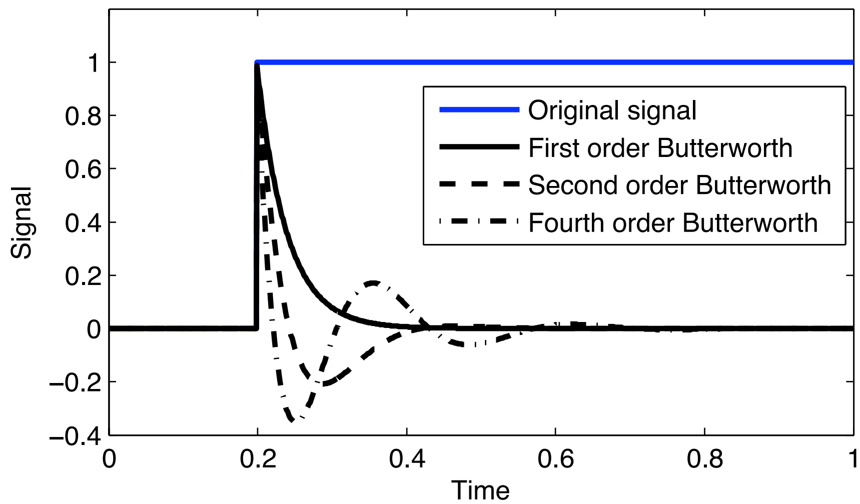

Figure 2.6. The step response for a first, second, and fourth order high pass Butterworth filter. The Butterworth filter is a linear recursive filter. Notice that second and higher orders have an overshoot, which feels very poor when used to filter the Velocity.

causes significant problems. When a user applies a sudden force, the filtered response results in a resistance force, but the delayed overshoot pushes in the opposite direction. This makes the crank feel active instead of passively resisting. A filter that overshoots does not work for this application.

2.6. Calculating the Impedance of a User's Arm

In some of my experiments, I wanted to determine the impedance of a user's arm at various locations in the workspace. The first method I tried was to apply a 100 ms 10 N force and measure the resultant position, velocity, and acceleration. I then fit the state measurements to a second order model of the arm. This is the method used to determine arm impedances in several studies [**61**][**31**]. I was able to measure the inertia and the stiffness reasonably well when I added a known inertia and attached a spring. I also added an eddy effect damper using strong magnets hovering near an aluminum bar, yet the method described above was unable to differentiate between having damping and not having damping. In fact, when damping was added the inertia and spring values would also change. The results from this experiment are shown in Table 2.1. These results made me doubt the ability to reliably measure the impedance.

Since measuring each of the individual components of impedance is not reliable, I decided to apply the same 100 ms 10 N force and measure the resulting position, which is a combined measure of the entire impedance. This had much better results and was used in the experiments discussed in section 5.6. The impedance will be referred to in terms of the displacement, which does not have a direct physiological meaning. In order to give a little more intuition about what each displacement means, I modeled the system as a second order system and will display some values that relate to how the inertia, damping, and stiffness changes the displacement. Since the displacements are small and the motions are nearly straight lines along the circumference o the workspace, I did not model the system as a rotational system. Thus, the system is represented by equation 2.7 over the 100 ms.

Table 2.1. Results from calculating inertia, damping, and stiffness of known quantities by applying a short force impulse. Inertia and stiffness are accurately measured. Damping is poorly measured and correlated to inertia and stiffness.

device setup:	Inertia ($kg * m^2$)	Damping ($\frac{N*m*s}{degrees}$)	Stiffness ($\frac{N*m}{degree}$)
crank alone	0.12	0.01	0.03
crank + 0.06 $kg * m^2$	0.18	0.04	0.00
crank + 0.16 $kg * m^2$	0.28	0.06	0.00
crank and spring	0.08	0.00	0.94
crank and damper	0.12	0.03	0.21
crank + 0.16 $kg * m^2$ and damper	0.30	0.07	0.00

$$\ddot{x} = \frac{-b * \dot{x} - k * x + 10}{m} \tag{2.7}$$

Where m is the mass of the arm in kilograms, b is the damping in $\frac{Newton*second}{meter}$, and k is the stiffness in $\frac{Newton}{meter}$. The initial conditions are zero velocity and zero position ($\dot{x}(0) = 0$ and $x(0) = 0$). Table 2.2 shows several resultant displacements based upon several values of the impedance. The displacements do not necessarily specify the exact impedance, but help to give an understanding of what effect the relative change in mass, damping, and stiffness have on the final displacement.

Table 2.2. An analytic solution to the second order model of the arm after a 100 ms 10 N force. This table is here to provide an intuition about what effect the impedance has on the final displacement.

mass (kg)	damping ($\frac{N*m*s}{degree}$)	Stiffness ($\frac{N*m}{degree}$)	Displacement (cm)
0.7	0.0	200	5.60
0.7	0.0	1000	1.80
0.7	10.0	200	3.76
0.7	10.0	1000	1.46
1.7	0.0	200	2.66
1.7	0.0	1000	1.75
1.7	10.0	200	2.22
1.7	10.0	1000	1.49

CHAPTER 3

Experimental Setup

3.1. General Experimental Procedure

I have performed several iterations of experiments. Early experiments used target sizes based upon Fitts' law [**23**], while the later experiments used a constant target size. The general experimental procedure for all experiments is explained in this chapter. When the results of an experiment are discussed in chapters 4 and 5, any experimental changes will be explained.

3.1.1. Participants

Unless otherwise noted, all participants are from Northwestern University's Psychology pool and participated after giving informed consent. There are four main sets of participants that I will be discussing. The human-human studies (chapter 4) consisted of 58 subjects[1] (34 female; 55 right-handed) with five different target sizes (conducted in Spring, 2004) and a second set of 30 subjects (20 female; 28 right-handed) with one target size (conducted in Winter, 2005). The human-robot studies (chapter 5) consisted of 22 subjects (15 female; 14 right-handed) with a robotic partner (conducted in Fall, 2005) and 22 subjects (12 female; 14 right-handed) for the perturbation studies (conducted in Fall, 2006).

[1]One of the 58 subjects was uncooperative and kept letting go of the handle, so the data from this dyad was not analyzed.

Figure 3.1. Two-handled crank setup used for experiments. A projector displays the targets from above. A curtain hangs between the two subjects to prevent non-physical communication.

3.1.2. Stimuli and Apparatus

Two randomly selected participants standing on opposite sides of the table (figure 3.1) held the rigidly connected and freely spinning handles. A black line, mounted to the crank, displayed the handle's angular position. A projector (not shown) mounted above the curtain displayed identical targets to each participant. The setup is shown in figure 3.2.

Participants were instructed to move the handle into the target as quickly as possible and to hold it there until a new target appeared. The new targets appeared at a randomly selected time between 700 and 1700 ms to prevent the subjects from adapting to a predictable delay. The target changed color when the black line, representing crank position, was inside the target. The projector also displayed a performance measure to encourage participants to perform to the best of their abilities.

Each target subtended 6° of the 50.5 cm diameter disk (2.6 cm at the perimeter of the disk) with a distance between consecutive targets of 70° ± 10° (30.8 ± 4.4 cm). The targets for the Fitts' law experiments varied between one of five sizes: 2° (0.9 cm around the circle at the handle), 4° (1.8 cm), 7° (3.1 cm), 13° (5.7 cm), and 25° (11.0 cm) with 70° ± 10° between targets. The Index of Difficulty, as defined by Fitts' law [**23**], covers a large range to account for different behaviors at varying sizes.

In the human-human experiments, five-sixths of the trials required a reversal of handle rotation from the previous trial; in one-sixth of the trials, handle motion in the same direction was required ("catch-trials"). The catch trials were designed to prevent the subjects from predicting the direction of the next target and were thus not included in data analysis. No two consecutive trials were catch-trials, so I also discarded trials immediately after a catch trial since the direction to move could easily be guessed.

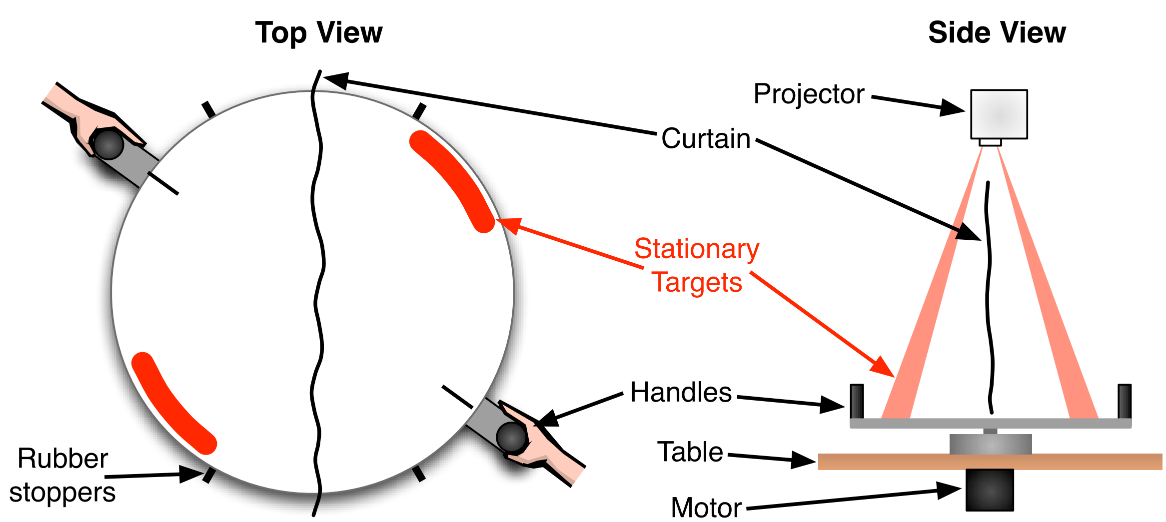

Figure 3.2. Diagram showing the setup of the device used for experiments.

3.1.3. Procedure

An experimental run started with an individual, or a dyad, performing a block of trials (target acquisitions). Each block of trials was anywhere from 24 to 154 trials, depending on the experiment. Half of the participants completed one block of trials individually, and then one block as a dyad (A, B, AB). The other half performed as the dyad first (AB, A, B). The sequence was performed twice so possible combinations are:

Dyad: A and B together (AB)		(A) A individually
A individually (A)		(B) B individually
B individually (B)		(AB) Dyad: A and B together
Dyad: A and B together (AB)	OR	(A) A individually
A individually (A)		(B) B individually
B individually (B)		(AB) Dyad: A and B together

The experimental apparatus was identical when the participants were working as individuals and dyads, except that the rotational inertia of the crank ($0.113\ kg \times m^2$) was doubled in the dyad condition. The entire experiment took less than 30 minutes. Force and common motion were recorded at a sampling interval of 1 msec. The completion time is the time at which the cursor entered and stayed within the target.

The term dyad can have different meanings depending on the context. When talking about human-human experiments, a dyad refers to both subjects holding the handles and completing a task jointly. When talking about human-robot experiments, a dyad refers to a subject working with a robot. The humans apply force from the handles whereas the robot applies a torque directly from the motor mounted underneath the table.

3.2. Graphics

There are several graphics that are displayed on the top plate during the experiments. Figure 3.3 shows the graphics that are displayed throughout the experiment. In addition to the targets discussed above, there are two messages that can be displayed: "Grab the handle and get ready to go" or "Release the handle and keep clear of it." These messages are displayed so I do not have to explicitly tell the the subjects when to grab the handle and when to let go. This assures that I do not indicate to them when they are working individually or together. This is particularly important during the human-robot

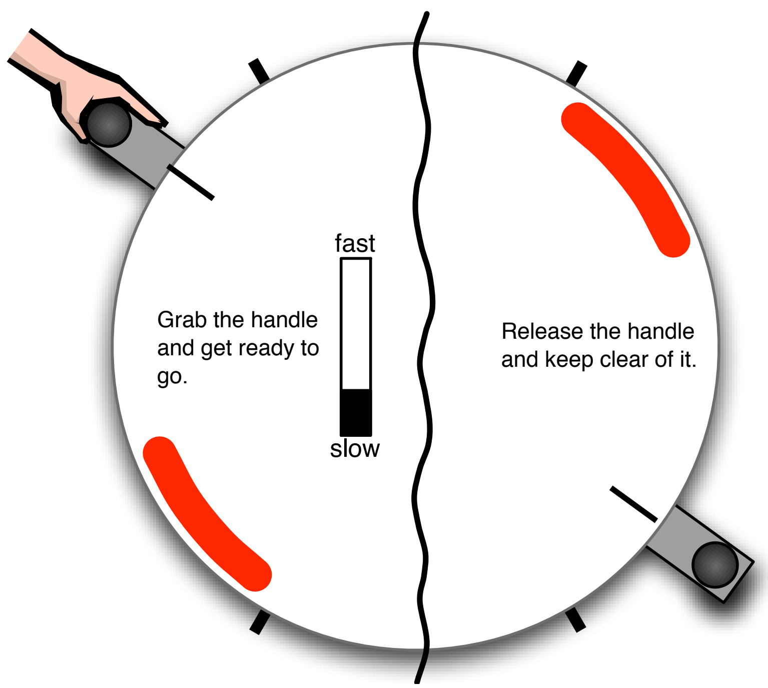

Figure 3.3. Display of the graphics shown on the top plate of the crank. The projector mounted above the crank and curtain displayed the targets, a message, and a performance measure. The targets change color when the crank is moved inside. The messages change at the end of each block of trials to inform the subjects to either grab the handles or let go of it. The performance measure indicates how fast the subjects are performing and attempts to keep them engaged and actively trying to complete the task.

studies. Each of these messages is constant throughout one block of trials. There is also a performance measure that indicates whether the subject is performing “fast” or “slow.” This ensures that the subjects will maintain motivation and continue trying throughout the experiment. The performance measure is only shown when the subject is performing the task. Early experiments were tainted with lazy participants.

One particular issue with the QNX graphics engine is a time lag. From the moment the code issues a draw command, it takes about 30 ms until the image is drawn on a CRT screen. I tested this using a photodiode connected to ServoToGo. I issued a draw command and measured how long until the photodiode measured a change on the screen. Using the LCD projector adds about 15 ms more beyond the CRT monitor. I measured this by oscillating the handle back and forth at a known frequency while drawing the current location as a dot. I adjusted the frequency until the handle and dot were 180° out of phase, which is easily seen with a naked eye. I then calculated the delay.

This delay is not noticeable in most situations such as UTLA [**100**][**102**], but it is noticeable when the screen is projecting down onto the device that is moving, such as mine. Originally, the projector was displaying the location of the handles onto the device, but this time lag was noticeable and distracting, so I replaced it with the black line near the handles as discussed above. This time lag can be compensated for when I have complete control of the motion, such as when the targets are displayed. Compensating for realtime events is possible by predicting where the position will be 45 ms in the future based on current position, velocity and acceleration, but the reaching tasks are not necessarily predictable and I did not want the results to be skewed due to a poor position tracker.

CHAPTER 4

Human-Human Interaction

In this chapter, I will discuss several results from my human-human experiments. This chapter is primarily focused on understanding how two people can cooperate on physical tasks.

4.1. Fitts' law Applies to Two People

As an initial foray into human-human interaction, I first performed an experiment based on the prototypical Fitts' law task [**23**] (described in section 1.1). The goal for my experiment was to determine whether Fitts' law describes the performance of two people engaged in a physically shared task with a common goal. Additionally, does Fitts' law hold true for two people with slightly different tasks in which reaching the goal requires coming to a physical compromise with each other?

4.1.1. Procedural Changes

The procedure is similar to the generalized experimental procedure described in chapter 3 with a few exceptions. Instead of displaying a target to each subject on the crank, each subject views a monitor, which shows a cursor and a target as shown in figure 4.1. The cursor moves horizontally according to the angle of the two-handled crank with the subjects holding the handles. Also, there was no curtain in this experiment so the subjects could see their partner as they moved the crank and the crank inertia was not doubled in the dyad case.

This experiment consisted of four separate blocks:

Random order:
- Subject A performing alone (24 targets)
- Subject B performing alone (24 targets)
- Dyads working together with common target (24 targets)
- Dyads working together with different, but overlapping, targets (60 targets)

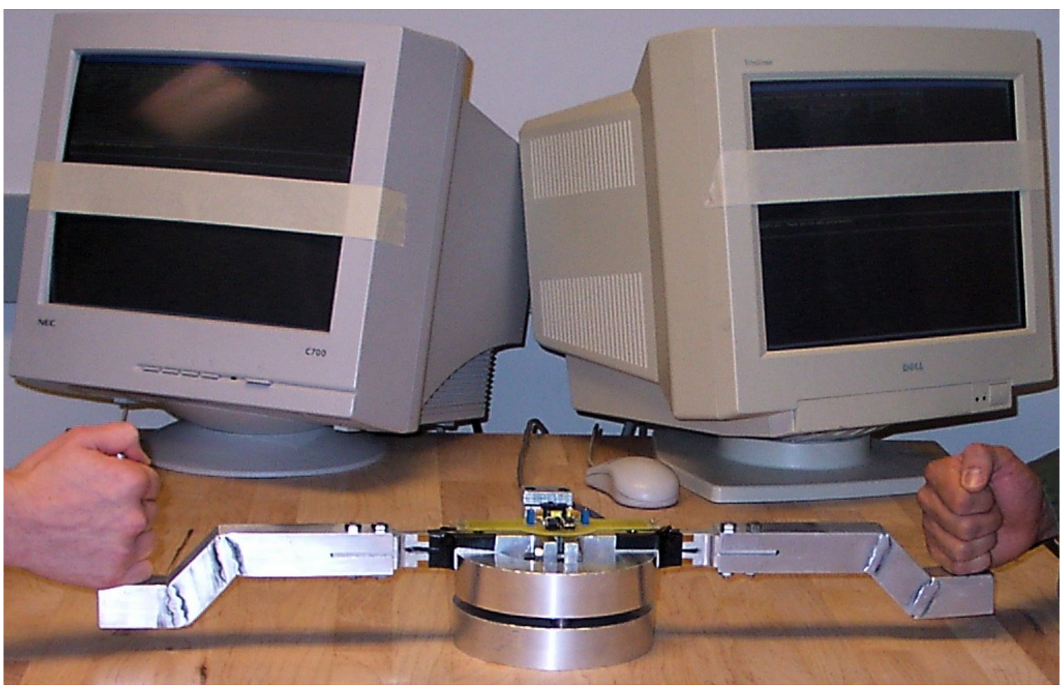

Figure 4.1. Initial device setup without curtain and without top plate. The targets are displayed on the computer screen along with the current position of the crank.

In all blocks, the targets varied between one of four sizes: 2.4° (1.3 cm around the circle at the handle), 5.6° (3.1 cm), 10.4° (5.7 cm), and 16.8° (9.2 cm) to cover a broad range of the index of difficulty. When the cursor is successfully brought and held inside the target, the current target disappears and a new target appears. Target sizes appear in random order.

In the common target part of the experiment, two subjects jointly control the position of the crank, coupled physically by the rigid connection. Although they view separate monitors, the monitors show the same cursor and target, so that the subjects are physically working together with the same goal. Figure 4.2a shows an example of this arrangement. Since the inertia and friction of the device are small, there is little to get in the way of the haptic communication between the subjects. This experiment addresses physical communication with no conflict in goal, and in which no physical compromise is required.

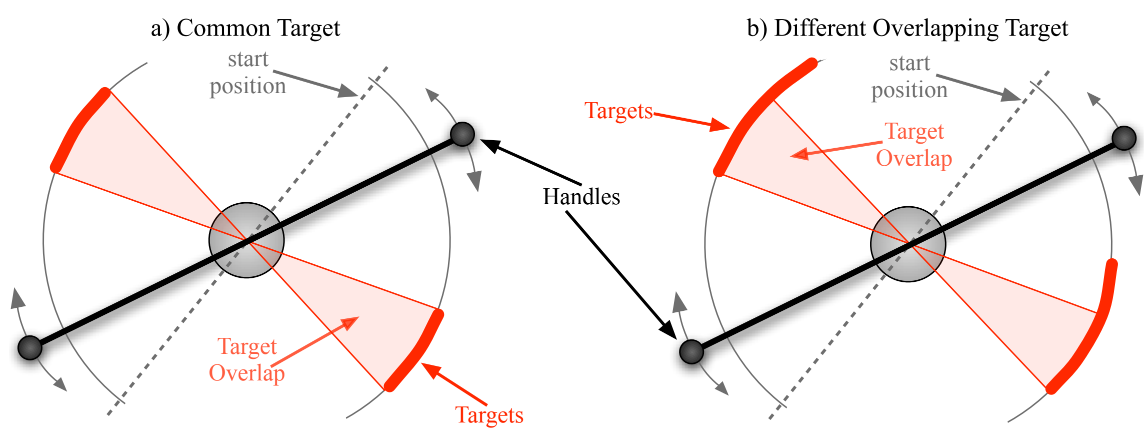

Figure 4.2. Comparison of the two parts of the Fitts' law experiment. a) The common target experiment (shown on the left) requires no negotiation between the subjects since their targets are exactly the same. b) The different overlapping target experiment requires the subjects to compromise by moving into a subset of their target that also satisfies their partners target. The subjects only see their own target, not their partner's.

In the different target part of the experiment, two subjects control the cursor together, but the target seen by subject A is not the same as the target seen by subject B. Each subject can only see their own task. In this case, the two subjects must physically communicate and compromise to move the cursor to a position where both subjects are satisfied. Figure 4.2 shows a comparison of the common and different target parts of the experiment.

4.1.2. Participants

Eight students (all men; 1 left-handed), age 22-28, from the Laboratory for Intelligent Mechanical Systems participated. Subjects were randomly paired to make four dyads.

4.1.3. Results and Discussion

For the eight individual subjects performing alone, the completion times show a high correlation to Fitts' law ($R^2 = 0.96$). The performance data for the four dyads also obeyed Fitts' law ($R^2 = 0.90$). This shows that two people working together do continue to obey Fitts' law.

When the two subjects have slightly different goals, I must clarify what 'target size' means. It seems reasonable that the target size should now be interpreted as the overlap distance. Using this definition, I do not find a Fitts' law type behavior for this task ($R^2 = 0.44$). Since, historically, a very wide class of tasks have shown Fitts' law behavior, even for definitions of ID that are quite diverse and task-specific, my result is surprising. This result suggests that the perception of the participants' interaction (either cooperative or competitive) can influence how they perform and interact with each other.

Bruning et al. [**11**] found that competition led to faster results than cooperation when subjects can see each other. They explain that competition requires no pooling of information whereas cooperation requires the overhead of negotiating. Using different target locations for each subject puts them somewhere in-between cooperation and competition. Each subject tends toward the center of their target, but in order to successfully complete the task, they both must be in the target, which means they cannot be in the center.

Inspired by my published results [**75**], Gentry et al. [**28**] performed a similar experiment. Gentry et al.'s experiment used a cyclical (not discrete time) target acquisition task and found similar performance results as mine. They used five subjects and tested all ten possible combinations for the dyads. Even though subjects were a member of multiple dyads, there was not a preferred completion pattern associated with any individual.

4.2. Dyad Performance Increase

4.2.1. Participants and Setup

These results are from the 30 subjects, setup, and procedure described in chapter 3.

4.2.2. Results and Discussion

It is commonly accepted that people prefer to work alone on tasks that require accuracy, finding a partner to be an impediment. Many of the participants in this study reported this same opinion of difficulty when working with a partner; few reported cooperation. This perception of poor cooperation likely stems from the increased force exerted by each member of the dyad. In fact, each dyad member averaged 2.1 times larger forces than under individual conditions with most of this force being expended in opposition to the other member.

Despite the increased forces and lack of perceived cooperation, dyads completed the task faster than individuals. The completion time for dyads decreased by an average of 54.5 msec compared to the average completion time of the two constituent individuals working alone (paired-samples, $t(15) = 5.95$, $p < 0.0001$, $d = 0.81$). The average completion time for individuals was 680 ms. Results are shown in figure 4.3. Increased

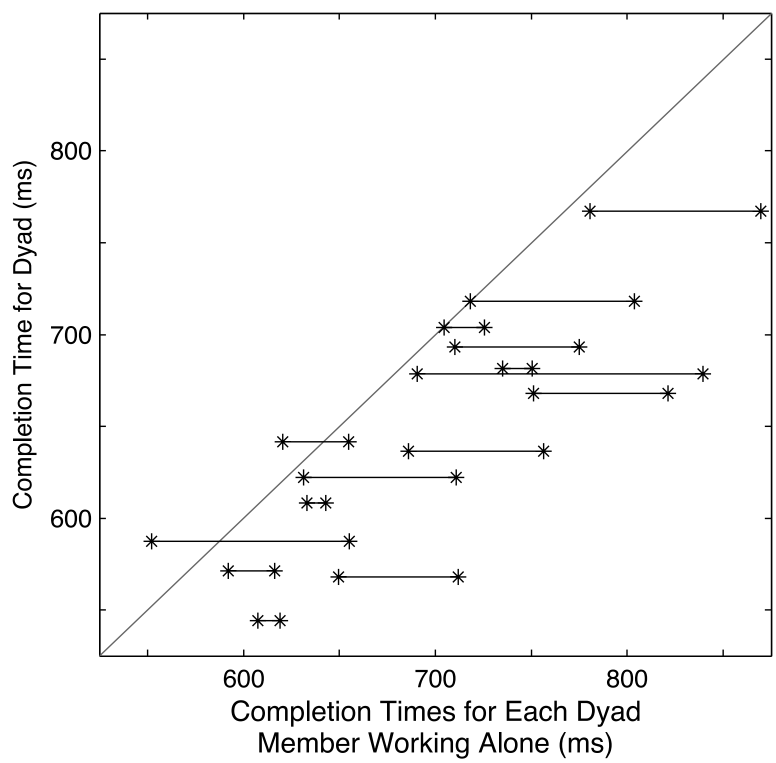

Figure 4.3. This graph shows completion times for each individual performing alone (x-coordinates of the two "∗" symbols) and for the same two individuals as part of a dyad (y-coordinate of the "∗" symbols). A horizontal line connects the two members comprising each dyad. Hence, symbols to the left of the diagonal represent individuals who outperformed their dyad (only two of thirty). Dyads performed faster than individuals, despite expending significant force in opposition to each other.

force associated with the dyad condition might result from a faster participant pulling along a slower one. However, dyads averaged 24.8 msec faster than the faster of the constituent individuals working alone ($t(15) = 2.76$, $p < 0.0154$, $d = 0.39$). Only two of 30 subjects were faster than their dyad. Improved dyad performance was established quickly when a dyad began to work together (within 20 trials), for both the A, B, AB sequence and the AB, A, B sequence.

The explanation of faster performance cannot be sharing the load because I doubled the cranks rotational inertia in the dyad condition. Another theory for why dyads are faster is social facilitation, which will be addressed in section 5.4. These performance results are published in Psychological Science [**76**].

4.3. Specialization

4.3.1. Participants and Setup

These results are from the 30 subjects, setup, and procedure described in chapter 3.

4.3.2. Results

To analyze the forces present during a dyadic target acquisition task, it is convenient to transform the forces from member A (F_A) and B (F_B) into a "net force" and a "difference force." The net force ($F_{net} = F_A + F_B$) is the sum of the members forces and is the task-relevant force that accelerates the crank. The difference force ($F_{diff} = F_A - F_B$) is a measure of the disagreement of the members, and has no physical effect on crank motion. The difference force is a measure of the force expended that does not go into accelerating the crank. As will be discussed in section 4.5, the difference force is a possible channel of communication between the participants.

Figure 4.4 shows the net, difference, and each member's force for a single trial by two different dyads. The vertical line near 600 msec in each plot indicates the completion time. As expected, individuals performing alone produce force trajectories similar to those predicted by a minimum jerk trajectory [**24**]. Interestingly, the net force profiles for dyads also produce a minimum jerk trajectory even though the constituent members of a dyad produce very different force profiles.

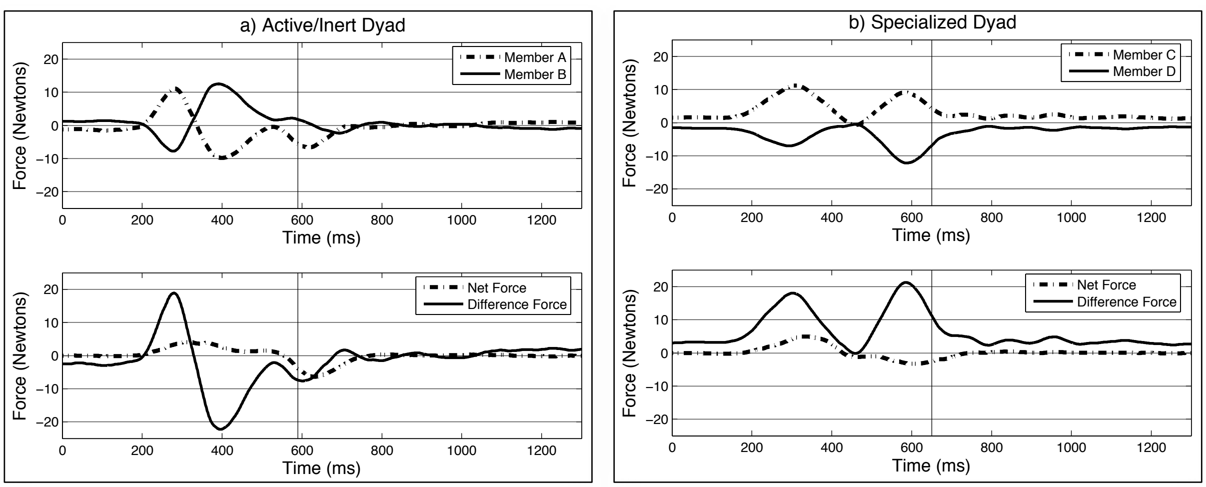

Figure 4.4. Forces in a single trial for a) an "active/inert" dyad and b) a "specialized" dyad. The top graphs show the forces produced by each member of the dyad when working together. Forces are recorded as positive when applied toward the target. The bottom graphs show the forces transformed into the net and difference forces: the net force (the sum) contributes to crank motion while the difference force is a measure of the force expended in opposition to one another. In the specialized dyad, the difference force is always the same sign: member C accelerates the crank forward while D is pulled along, and then C continues to push forward (fails to reverse) while D brakes.

An active/inert dyad is shown in figure 4.4a. Early in the trial member A provides a force toward the target while member B's passive inertia creates a counterproductive force. Late in the trial member A provides all the deceleration force as well, while member B is again passive. Member A is active and member B is inert.

Figure 4.4b shows a very different pattern, which I denote "specialized." Member C pushes toward the target early in the trial to accelerate the crank while member D either passively or actively resists the motion. Member D then pushes away from the target late in the trial to decelerate the crank while member C continues pushing toward the target. Member C primarily contributes during the acceleration phase and member D primarily contributes during the deceleration phase. This specialization is clearly revealed by inspecting the difference force, which remains always the same sign. Only the characteristic shape of the difference force matters – it could be mirrored around the x-axis if the subjects were standing on opposite sides of the crank.

The acceleration phase begins when the net force deviates from the initial resting period and ends when the net force switches from positive to negative. The deceleration phase begins when the net force switches from positive to negative (end of the acceleration phase) and ends when the crank is both in the target and the velocity is less than a given threshold. In figure 4.4b, the acceleration phase lasts from 190 ms until 430 ms. The deceleration phase lasts from 430 ms until 740 ms.

To find the contributions, I integrated the force of each member over each phase, and divided by the integrated net force for that phase as shown in equations 4.1.

$$C_{member,phase} = \frac{\int F_{member,phase}}{\int F_{net,phase}} \qquad (4.1)$$

Where *phase* is either acceleration or deceleration, *member* is either A or B, C is the contribution of each *member* during each *phase*, F is the force of each *member* during each *phase* and F_{net} is the net force.

The result provides four fractional contributions of each member of the dyad during the acceleration and deceleration phases. The two members' contributions for each phase necessarily sum to one. A negative contribution indicates that the member was working against the motion of the crank (accelerating during a deceleration phase, or decelerating during an acceleration phase, even if only due to passive inertia). A contribution greater

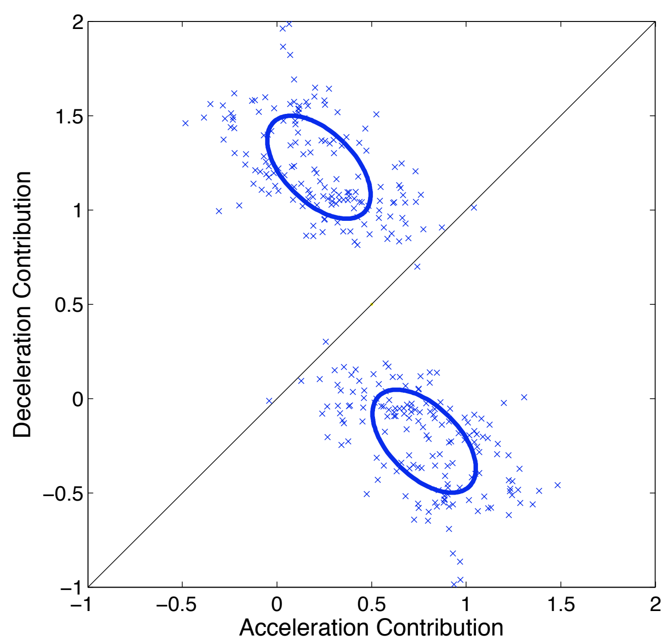

Figure 4.5. Specialization based on acceleration and deceleration phases of the movement. A dyad's degree of specialization can be characterized by plotting each members contribution to the acceleration and deceleration phases. The other member's contribution is necessarily opposite and appears as a dot mirrored about the center of the box $(0.5, 0.5)$. The distance from the $x = y$ line is the degree of specialization.

than one indicates that this member had to compensate for the negative contribution of the other dyad member.

Figure 4.5 shows two clusters of dots, one for each member of a specialized dyad, with each dot representing one trial. For every point in one of the clusters, there is a point mirrored around $(0.5, 0.5)$, which represents the other member of the dyad. As an example, a dot at $(0, 1)$ represents a trial in which this member contributed 100% of the acceleration and 0% of the deceleration. In this case, the other member necessarily contributed 0% of the acceleration and 100% of the deceleration. The average forces for this dyad, including the specialization, can be seen in the appendix in figure B.11.

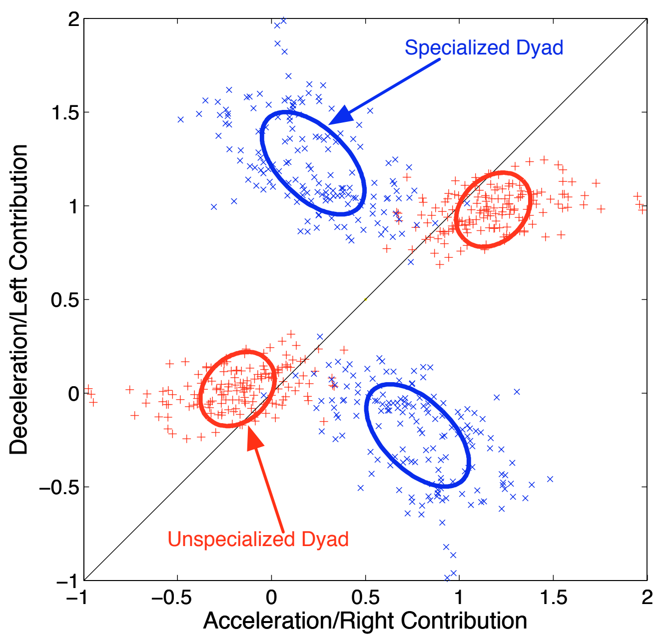

Figure 4.6. Specialization by acceleration/right and deceleration/left phases of the movement. The graph shows superimposed datasets for two different dyads, one of which exhibits significant specialization. The specialized dyad is A/D while the unspecialized dyad is L/R.

Dots near the $x = y$ line represent a member's equal contribution during the acceleration and deceleration phases. The perpendicular distance of a dot from the $x = y$ line is a measure of the degree of specialization. The ellipses show the standard deviation along the $x = y$ and $x = 1 - y$ lines. Using these measures, a non-specialized dyad is distinctly different than the specialized dyad. An example is shown in figure 4.6, which shows four clusters of dots. Two of the clusters are the same specialized dyad from figure 4.5 and two of the clusters are from an active/inert dyad. The average forces for the active/inert dyad can be seen in the appendix in figure B.10.

To assess the significance of the observed acceleration/deceleration (A/D) specialization, I compared A/D and left/right (L/R) specialization. L/R specialization is another

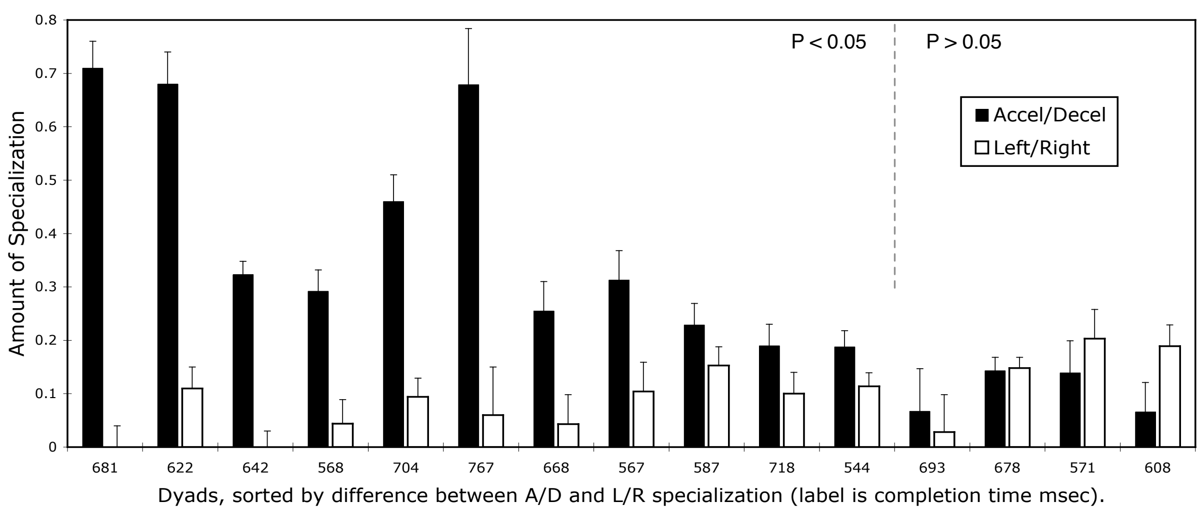

Figure 4.7. Each bar shows a dyad's average degree of specialization over all trials, and a 95% confidence interval. A value of 0.707 (the distance from the $x = y$ line to the point $(1, 0)$, meaning one member's contribution to acceleration is 100% and to deceleration is 0%) can be considered "fully specialized." 11 of the 15 dyads show significantly more acceleration-deceleration (A/D) specialization (black bars) than left-right (L/R) specialization (white bars) ($p < 0.05$). Many, but not all, dyads show a qualitatively high degree of A/D specialization. The x-axis label shows the average completion time for the given dyad.

reasonable division of the task, but one that was generally not adopted by these dyads. If dyads had adopted L/R specialization, this would be seen as one member pushing only to the left while the other member would only push to the right. Figure 4.7 compares A/D specialization to L/R specialization. Eleven of the fifteen dyads show significantly more A/D specialization than L/R specialization ($p < 0.05$).

These specialization results have been published in [**77**].

4.3.3. Discussion

When a single person becomes part of a dyad, many new solutions to completing the task develop due to redundant limb motion [**46**][**52**]. There is no longer a one-to-one correspondence between dynamics and kinematics. For example, in a dyadic task, one subject can choose not to perform at all, to help with only the braking phase, or to use only elbow flexor muscles. Knoblich and Jordon [**49**] suggest that when groups work together, they might perform better because each person has fewer actions to deal with. Many dyads specialize, which is one way they can solve the redundancy problem. They get out of each others way and only deal with a few actions. Additionally, humans can precue an action so that when some other event happens they perform a certain prepared action [**82**]. In the case of two people working together, the deceleration specialist possibly waits for some cue, such as reaching a certain location or velocity and subsequently begins to decelerate the crank while the accelerator focuses on other aspects of the task.

Dyads learned to specialize their force production temporally to produce a net force profile similar to a single person, but faster. The subjects arrived at this division of labor solely through a haptic communication channel since no other communication was allowed. Feygin et al. [**21**] found that haptic guidance helped subjects learn the temporal aspects of tasks better than visual guidance did.

A single person executing a target aquisition task in my experiment would be expected to use the triphasic burst pattern of muscle activity, in which an agonist muscle burst starts the movement, and is followed by an antagonist muscle burst to brake the movement and another agonist burst to help hold the limb at the final position [**39**][**40**][**35**]. These bursts represent careful planning based on prior knowledge, rather than feedback received during the task [**34**][**33**]. Moreover, these patterns represent optimal movements that best accomplish the task within rather limiting physiological constraints such as the rates at which muscles can be turned on and off [**33**][**73**] and the limited torque generating capacity in different areas of the workspace [**68**]. Consequently, the rate of force onset and offset becomes less critical if one person can be ramping up while the other is ramping down, which is presumably what occurs during specialization.

A similar specialization strategy has been observed in a single individual performing a bimanual task. Reinkensmeyer et al. [**80**] show that an individual holding a pencil between two fingers on different hands will accelerate with one hand and decelerate with the other hand, which might be taken as a bimanual model for my observed two person A/D specialization. For a single individual the inward force allows the pencil to accelerate and decelerate while being rigidly held and not dropped. The tight neural coupling between two arms in an individual allows a person to coordinate the actions of each arm [**95**]. In my dyadic task there is, of course, no neural coupling between the individuals, so the developed strategy must have occurred by haptic communication instead since the only communication between participants was physical.

4.4. Difference Force Oscillations

The forces exerted by humans tend to exhibit oscillations at around 8-20 Hz due to isometric contractions and force tremor [**20**][**56**]. Many of my experiments showed that two subjects tended to exhibit about five times more oscillations around 6-10 Hz than a single subject did. Initially, I thought these oscillations were the result of some interaction between the two subjects, but this is not necessarily the case. These same oscillations could be reproduced by simply attaching a mass to the other handle while one person worked alone.

Regardless of whether these additional dyadic oscillations are from a joint action on the other handle, each subject will feel more 8 Hz oscillations when working with a partner than when working alone. These oscillations may be beneficial for the two subjects by allowing each member to synchronize their timing to the other member. Instead of just seeing the target appear and moving like they do alone, these oscillations may serve to build a rhythm between the two people. Gentry et al. [**28**] showed that dyads are more capable of harmonic motion than individuals. I believe this rhythm may allow the subjects to perform better.

It may be possible to synchronize two people by introducing additional oscillations. Subjects could either be synchronized to the task (such as $1 \ldots 2 \ldots 3 \ldots$go) or to each other (oscillating at a random time before the task starts). Synchronizing both subjects to the task is easier and very likely to occur. This is similar to two people swinging an object back and forth as they prepare to throw it. Synchronizing to each other, and not necessarily the task, is harder, but likely possible. Several studies have shown that dyads will couple their rhythmic motions visually [**85**][**88**]. Although none of these studies have tested this in a physical environment, I believe it could be possible.

An experiment that could test whether oscillations aid in communication consists of disrupting the oscillations. Subjects would perform a similar reaching task where their high frequency oscillations are damped (as discussed in section 2.5.2) when the dyad is inside the target and moving slower than some speed. This will remove the oscillations in the target, but continue to allow them to move the crank for large motions. If their performance increases, then the oscillations are just a consequence of the brain/motor correction speed. If their performance gets worse, then the oscillations are beneficial and are possibly helping them to synchronize their movements.

4.5. Dyadic Contraction

At the end of each trial, the members bring the crank to rest as they wait for the next target to appear. If a participant is working alone, then there is no way to apply any force to the handle without also causing the crank to accelerate. If there are two members, each member can apply a force even though the crank is not moving, as long as the forces are equal and opposite. This generates a difference force with no net force, hence no acceleration.

I found that dyad members exert a significant force in opposition to one another (averaging 4 Newtons). Figure 4.8 shows the distribution of the average difference force for each pair at the end of each trial. This force could serve to increase stiffness for the dyad in the same way that muscle co-contraction increases arm stiffness for an individual by co-contracting both the agonist and antagonist muscles to account for perturbations [**63**][**57**].

Co-contraction is a common strategy when interacting with unstable force fields [**25**][**90**]. Since the difference force at the end of the trials resembles co-contraction in individuals, I call the dyadic version "dyadic-contraction." Dyadic-contraction is a strategy similar to those used in parallel robotics and in human bimanual control [**65**][**14**].

Gribble [**37**] has shown that the extent of co-contraction in individuals correlates inversely with target size in a Fitts' law type task. Similarly, my results [**79**] for dyads showed a correlation of opposition force with target size.

Dyadic-contraction could also serve as a simple message between partners that they are in fact working together. Without applying any force, there is no way to know what is happening or who is on the other side of the curtain. By applying a small force, each person feels a resistance that helps them to determine what is opposite them. In order to

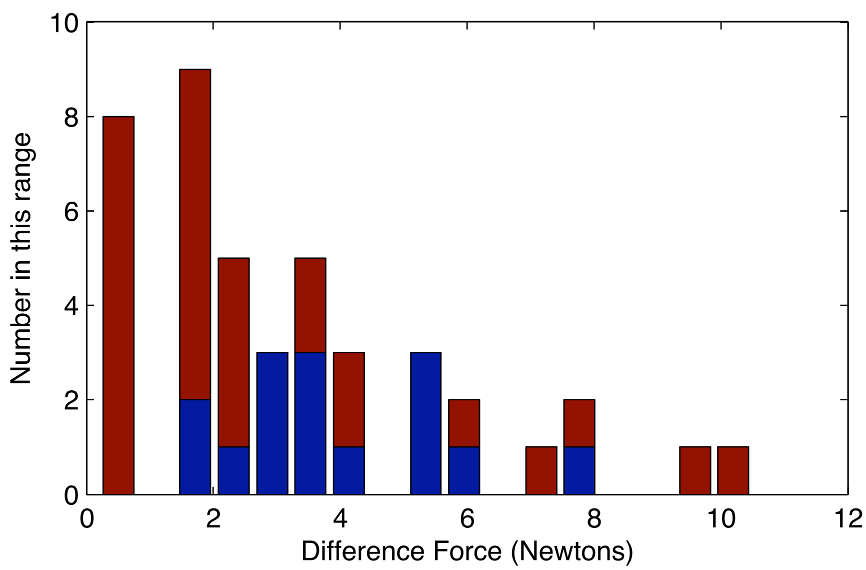

Figure 4.8. The average difference force at the end of the trial for 43 dyads. The 15 in blue on bottom had a single target size and the 28 in red on top [**79**] had multiple target sizes. In order to complete the task, no force is required at this point, yet many dyads exert a significant amount of extraneous force.

explore a surface, a person may attempt to maintain an optimum haptic force [**15**][**13**]. In order to learn about their partner, each member may be trying to keep the optimum force.

A possible experiment to test whether dyadic-contraction is a consequence of an optimal haptic force involves adding a slowly increasing external force while subjects complete a dyadic reaching task. If the subjects are attempting to maintain an optimal force, both members will want to push near the force level that they can receive the most information. In order for both of them to push near the same limit, they will presumably start pushing in the same direction against the motor when the motor exceeds twice the optimal haptic force. If they are not looking for an optimal haptic force, the subjects may divide the task in a number of other ways.

4.6. Discussion

The data from the HH studies is very rich, which inevitably leads to many questions about correlations between measurements. This section will detail some of these correlations. My hope is that this clarifies what effects are related to other effects. The correlations will include:

- Reaction: This is the difference in the reaction time of each individual working alone. The reaction time difference is calculated by subtracting the reaction time of subject B from that of subject A. The absolute value of their difference is shown on each graph except where the other axis is related to a specific member of the dyad. Each person's reaction time is on the order of 200 ms.

- Starter: In each dyad trial, one member begins to apply force toward the target before the other member. The 'starter' measure is a percentage of time that member A starts first. The starter percentage from member A and member B necessarily sum to 100%, so I arbitrarily plot one member to give an idea of who starts first. This is measured by finding who first pushes beyond the dyadic-contraction force toward the target. Since this is an arbitrary assignment for each dyad, some correlations do not make sense and are omitted.
- Specialization: This is the degree of specialization as described in section 4.3 and figures 4.4 and 4.7. A larger amount indicates more specialization (i.e. dividing the task among the dyad members) with a maximum value of 0.707 indicating fully specialized.
- Performance: This is the difference between the average individual completion times of the two subjects and the dyad composed of those two subjects. I use this measure rather than absolute completion time since I am focusing on the effects of the two subjects combining their effort. The performance times were discussed further in section 4.2.
- DC Force: This is the dyadic-contraction force that was discussed in section 4.5. At the end of each trial, the subjects pushed with an average of this force, even though no force was required to stay inside the target. Only the absolute number matters, except where the other axis is related to a specific member of the dyad. The positive and negative values of this relate to which direction the subjects were pushing.

Table 4.1 outlines how the measurements are correlated. Although most of these measurements do not correlate, I wanted to display these correlations to help answer any questions that may exist. I will discuss the particularly interesting graphs in more detail.

There are a couple of these graphs that are marked *omitted*. This is an indication that the correlation does not make any sense. An example would be specialization compared to starter. Specialization is a joint metric that compares how the two subjects performed, whereas starter is a metric about the percentage of trials that an arbitrary member started. I am arbitrarily picking to plot the percentage of member A, so if the subjects had switched locations, the value for starter would be 1−starter and specialization would remain the same value. In all other graphs, I can arbitrarily flip member A and B and the interpretation of the graph will not change. In the case of an *omitted* graph, if the members had switched places, the graph would change, which means that the graph is nonsensical. All other graphs would remain the same and are thus informative comparisons.

Table 4.1. This table shows the correlations for all of the measurements. It also shows where each graph is so you can easily find the graph you are looking for.

	Reaction	Starter	Specialization	Performance	DC Force
Reaction	—	$\boldsymbol{R^2 = 0.25}$ Fig. 4.10	$R^2 < 0.10$ Fig. 4.11	$R^2 < 0.10$ Fig. 4.12	$R^2 < 0.10$ Fig. 4.13
Starter	$\boldsymbol{R^2 = 0.25}$ Fig. 4.10	—	omitted	omitted	$\boldsymbol{R^2 = 0.73}$ Fig. 4.14
Specialization	$R^2 < 0.10$ Fig. 4.11	omitted	—	$R^2 < 0.10$ Fig. 4.15	$R^2 < 0.10$ Fig. 4.16
Performance	$R^2 < 0.10$ Fig. 4.12	omitted	$R^2 < 0.10$ Fig. 4.15	—	$R^2 < 0.10$ Fig. 4.17
DC Force	$R^2 < 0.10$ Fig. 4.13	$\boldsymbol{R^2 = 0.73}$ Fig. 4.14	$R^2 < 0.10$ Fig. 4.16	$R^2 < 0.10$ Fig. 4.17	—

4.6.1. Highlighted Comparisons

Figure 4.10 compares the difference in reaction time to the percentage of dyad trials that member A reacted first. When each individual is performing the task, there is a certain delay before they move due to their reaction time. I am using the difference between the two individual reaction times since the difference between the two subjects is what each member will feel when working as a dyad. It would make sense that the subject that reacts faster working alone would continue to react faster as a dyad, but that is not necessarily the case. The correlation ($R^2 = 0.25$) is quite small. There are many instances where the subject with the faster individual reaction time does not start as a dyad. This could be an issue of social loafing, but this explanation does not make sense. If the cause was social loafing, I would expect there to be a similar correlation between the difference in reaction time and the difference force. Subjects that do not want to exert any more effort than is required will not continue to push when it is unnecessary or start any earlier than is required. Figure 4.13 shows the lack of correlation between the difference in reaction time compared to the difference force. Some subjects that start slower have a large difference force.

Figure 4.15 compares the performance to the amount of specialization. Being faster as a dyad does not necessitate that the dyad specialized. There are other ways in which a dyad can speed up relative to an individual. Also, specializing does not imply that the dyad will speed up. There are possibly other reasons for specializing such as reducing the mental effort, getting out of each other's way, or reducing the overall work required.

Figure 4.9 compares the reaction times during the individual trials to who actually started the motion during the dyad trials. For this comparison, I am only including the dyads and constituent members that actually did show specialization since other dyads did not divide the task in such a manner. On the graph, each member of a dyad is connected with a line to indicate who was working together. There is not a statistically significant difference in reaction time between the accelerators and decelerators (paired-samples, $t(11) = -0.6904$, $p = 0.506$), which is surprising. I expected that the person with the faster reaction time would start first and apply a significant amount of the force for the acceleration phase. But, this was not the case. The subjects' strategy was distinctly different than the individual strategies.

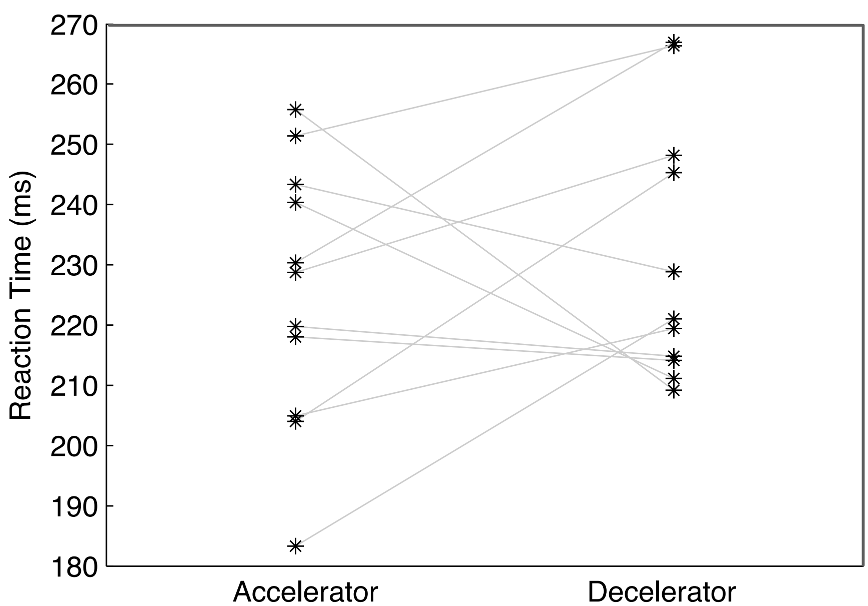

Figure 4.9. Comparing the reaction times of the individuals to who became the accelerator and decelerator. Only the 11 dyads that showed significant specialization are shown. Each member of the dyad is shown as a "$*$" and a line connects the accelerator to the decelerator. In four of the eleven dyads, the accelerator is the member with the slower reaction time, which implies that reaction time alone cannot explain why one person becomes the accelerator.

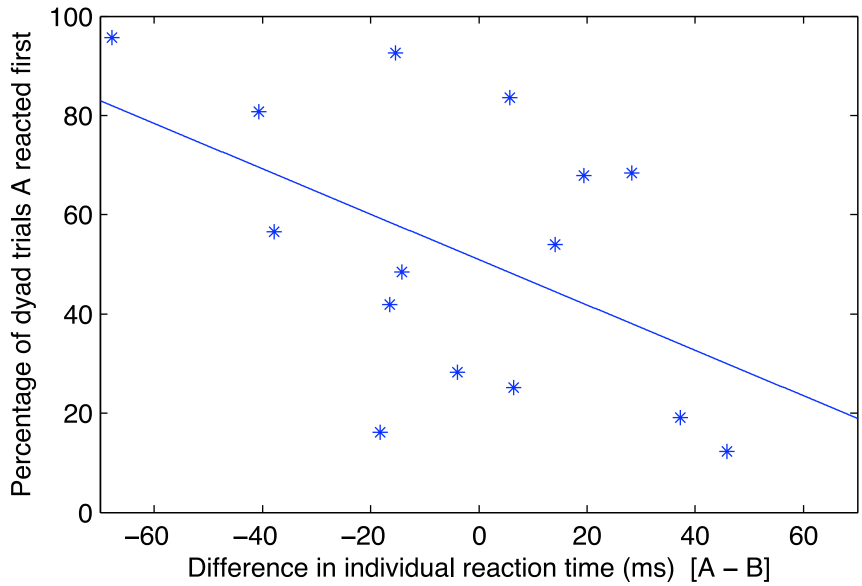

Figure 4.10. Comparing the difference in reaction time (negative means A is faster; positive means B is faster) to the number of times member A started first when working as a dyad. There is a very small correlation here, $R^2 = 0.25$. This is surprisingly low since the member with the faster reaction time would be expected to start nearly all the time, particularly when the difference is large.

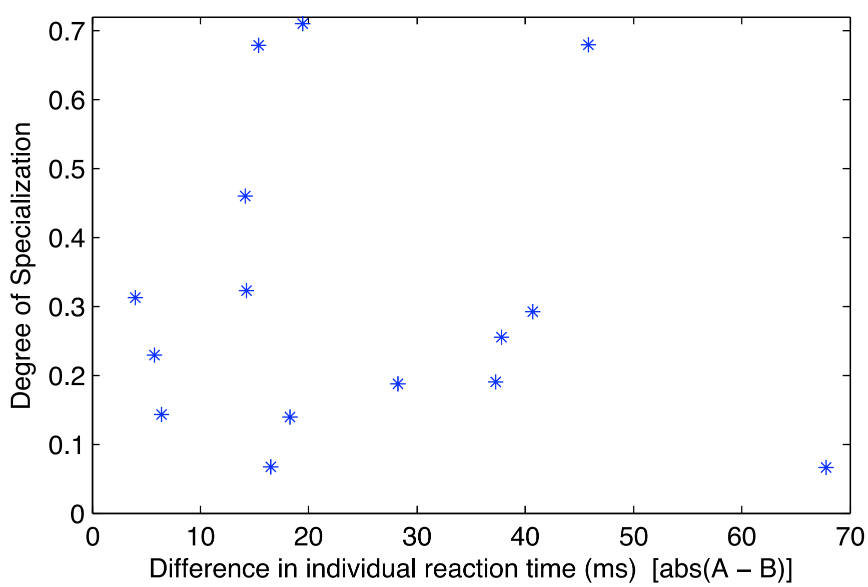

Figure 4.11. There is no correlation between the difference in reaction time compared to the degree of specialization.

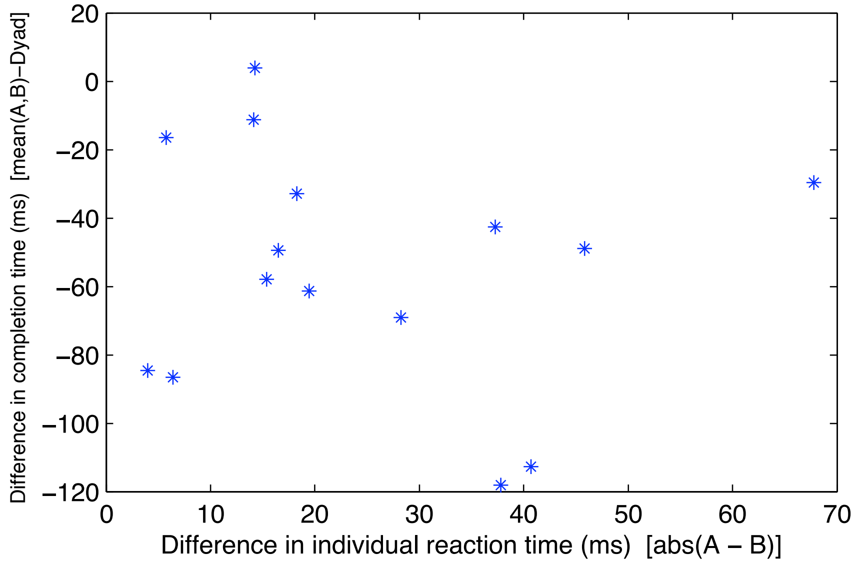

Figure 4.12. There is no correlation between the performance and the difference in reaction times. This would have made sense if this was simply one of the subjects starting motion before the other subject typically started, but it is not that simple.

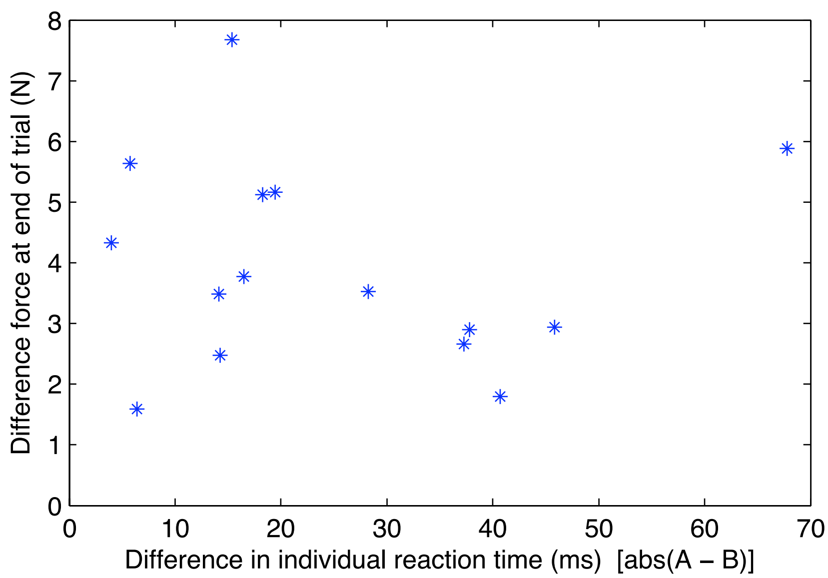

Figure 4.13. There is not a correlation between the difference in reaction time and the dyadic-contraction force. The force does not seem to be the result of the two subjects trying to start at different times, which would have suggested competition as opposed to a collaboration. Competition generally leads to faster performance than cooperation [**11**].

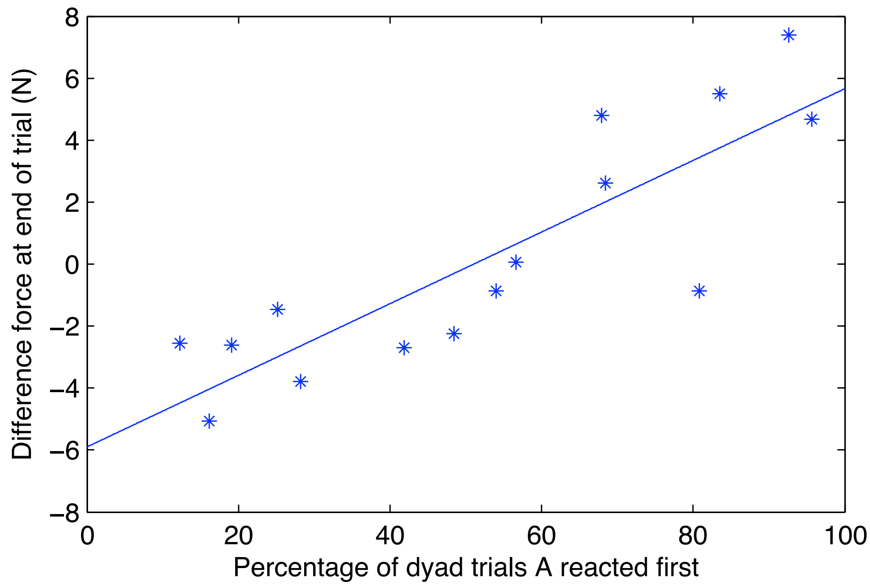

Figure 4.14. There is a correlation between one subject starting first and the difference force ($R^2 = .73$). The subject that is already pushing toward the target due to dyadic-contraction tends to push harder first. This could be due to the different ramp up/down characteristics of muscles [**3**].

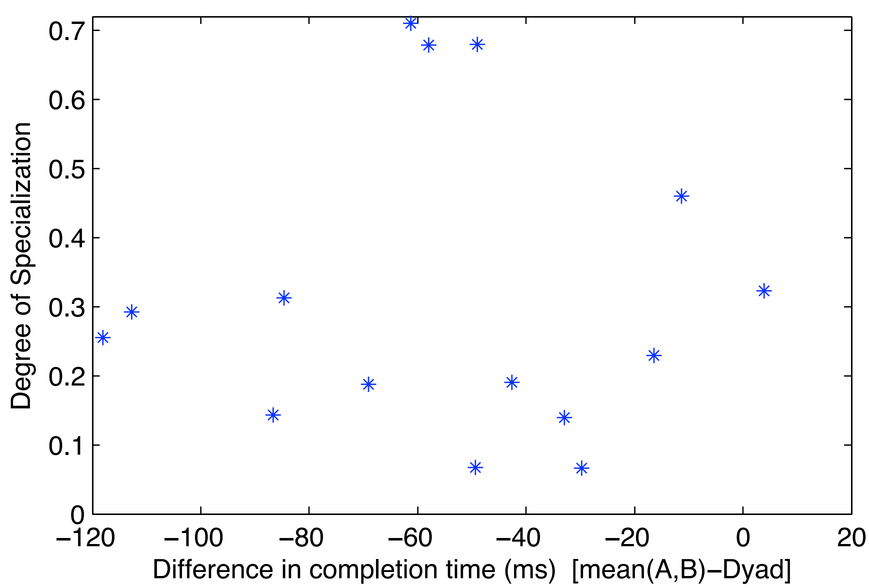

Figure 4.15. There is not a correlation between specialization and performance. This means that it is not necessary to specialize in order to perform faster as a dyad. Regardless of the lack of any correlation, each of these results are interesting by themselves.

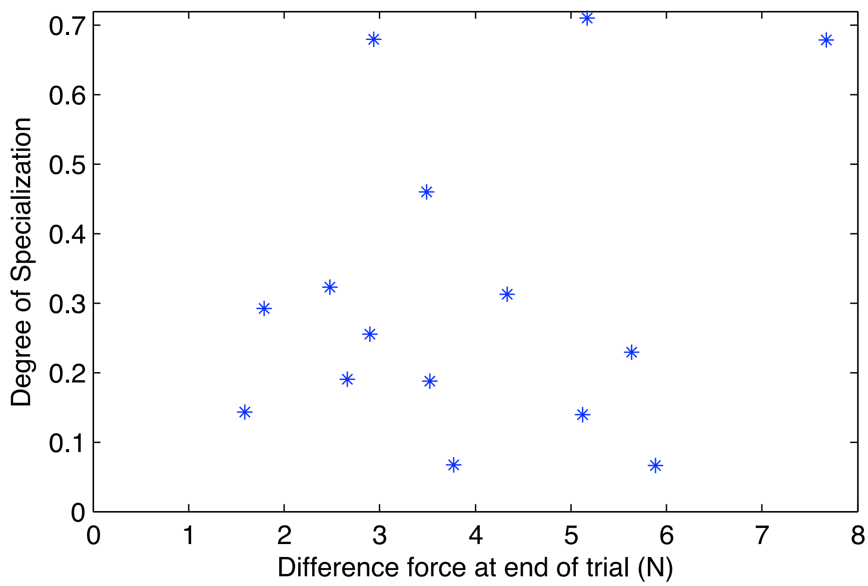

Figure 4.16. There is not a correlation between the difference force and specialization.

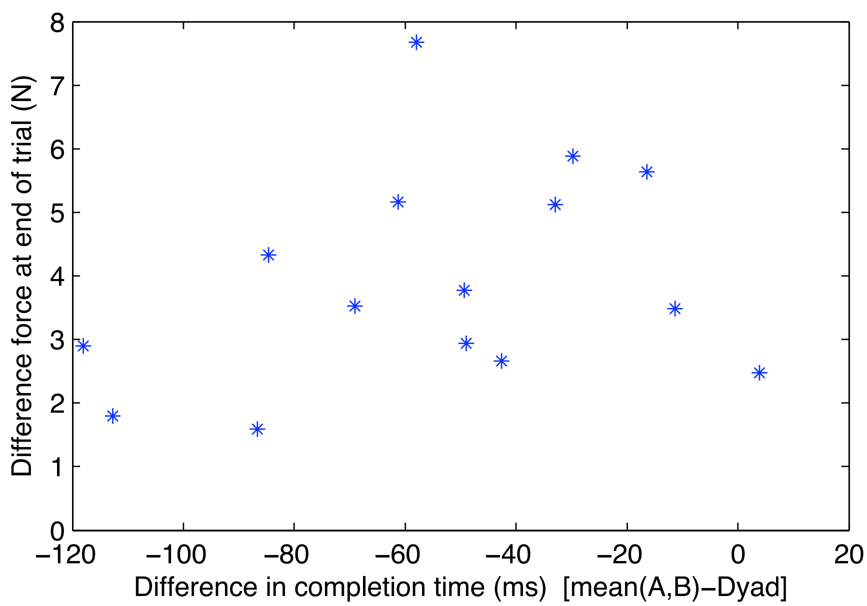

Figure 4.17. Initially, I thought that dyadic-contraction would correlate to performance. I expected the subjects that pushed harder against each other would also have a tendency to perform faster, but this is not the case.

CHAPTER 5

Human-Robot Interaction Results

In this chapter, I will discuss several results from my human-robot experiments. This chapter is primarily focused on understanding how a human can work with a robot programmed to simulate a human-human behavior on a physical task.

5.1. Simulated Person

Since two subjects can perform a strategy similar to a single person performing a bimanual task (discussed in section 4.3.3), would a simulated partner that displays an acceleration/deceleration specialized trajectory be enough to elicit the same response? If this model were to elicit a similar response in a human, then this model provides a complete and natural strategy for human-robot interaction in reaching tasks. If this did not elicit a similar response, the model is likely incomplete and other human-human interaction strategies may need to be added, such as dyadic-contraction.

Additionally, I wanted to examine whether a human could determine the difference between a human and this simulated partner, much like a Turing Test [**98**]. Alan Turing originally proposed a test for evaluating a computer's ability to produce human like conversation, but it has since been used more broadly for other human like attributes.

5.1.1. Hypothesis

My human-human results suggest that humans have an ability to cooperate and effectively work together by combining their forces in a specialized manner (section 4.3). I designed the simulated person experiment described below to replicate the interaction of two people on a simple task by replacing one person with a simulated partner. The simulated partner mimics one partner's role in the specialization I found in human-human physical interaction. I expected subjects to work with the robotic partner in a similar way as they did with a human partner.

It is possible that subjects who know they are working with a robot will behave differently than if they believe they are working with a human. As a control, some experiments involved a confederate. The confederate stood in as if they were the subject's partner, but the confederate did not work with the subject, the robot did. I hypothesized that a perceived human partner combined with a robotic physical partner would elicit the same response from a subject as actually working with a human partner.

5.1.2. Participants

Twenty two students (7 men; 1 left-handed), age 18-24, from Northwestern University's Psychology participant pool participated after giving informed consent.

In eleven of the experiments, a confederate also participated. The confederate pretended to lack knowledge of the experiment. Although the confederate never actually worked with a subject, the subjects would assume they were working with a person. The confederate stood across from the subject during the experiment, but did not physically interact with the subject. The confederate did perform some trials of the task alone to ensure the subjects believed the confederate was participating. In the other eleven experiments, the subjects knowingly worked without a human partner.

5.1.3. Stimuli and Apparatus

I used the same experimental apparatus as the human-human experiments shown in figure 3.1, except one of the participants was simulated using the motor mounted under the table. The motor simulated a person by displaying an average torque trajectory equivalent to the participant's own force during individual trials. The motor force[1] was only displayed during the acceleration phase to elicit specialization (i.e. no deceleration force as shown in figure 5.2).

5.1.4. Robotic Partner

The robotic partner is composed of two parts: an active force production and a simulated inertia. The first part mimics the behavior of a specialized partner who has taken on the role of accelerating the crank. The acceleration part is a modified version of the subject's own force trajectory that was recorded and averaged during the individual trials. When working alone, the subject necessarily did both the acceleration and deceleration parts of the task. I captured the acceleration part of the subject's own force profile, and used it later as the robot's force profile. This acceleration force trajectory was multiplied by 2.1, which is the typical amount an individual increased his/her force by when he/she becomes part of a dyad (section 4.2). This modified force trajectory becomes a typical force trajectory that could be found in a specialized member of a dyad. I used a recorded

[1]Motors apply torques, but I will refer to the equivalent force at the handle that the motor torque generates since the interesting parameter is the force at the handle.

version of the subject's forces to account for variations in forces and completion times among subjects, so that any differences can be attributed to the subject working with a partner and not because the robotic partner is faster or slower than the subject.

The force trajectory for the robotic partner ($RP(t)$) is summarized in equation 5.1.

$$RP(t) = 2.1 \times \sum_{i=1}^{100} \frac{\underline{f_i}(t)}{100} \tag{5.1}$$

where $f_i(t)$ is a vector containing the forces for individual trials, $i = 1$ to 100, t is the time since the target was shown ($t = 1$ to 1300 ms), and $\underline{f_i}(t)$ only allows positive values of the force as defined in equation 5.2.

$$\underline{f_i}(t) = \begin{cases} f_i(t) & : \quad f_i(t) > 0 \\ 0 & : \quad f_i(t) \leq 0 \end{cases} \tag{5.2}$$

The second part of the robotic partner consists of a simulated inertia. In order to simulate the mass of an arm holding the handle, an inertia similar to that of a human arm was added. Four different confederates assisted throughout the experiments and the robot simulated an average of their arm inertias. The inertia (I) of each of their arms was calculated by grabbing the crank with the same grip subjects used in the experiments. The motor applied a torque (τ) and I measured the angular acceleration (α) over a frequency from 1 to 35 Hz. The inertia was then found from the equation $I = \frac{\tau}{\alpha}$, which is similar to the method discussed in section 2.6. The calculated average inertia used was 0.24 $kg \times m^2$. The simulated force was increased since it would have to accelerate it's own simulated arm as well as the crank. The control loop for the robotic partner ran at 1 kHz.

5.1.5. Procedure

In half the trials, I employed a confederate so the subjects would think they were working with a person to test whether a subject would cooperate differently if they knew they were working with a person or a robot. In this group (human-robot-confederate), an experimental run started with the individual (I) or the confederate (C) performing a block of trials individually. Then the other person completed a block of trials individually. Next, the individual worked with the robotic partner (I_m), which the subjects believed to be the other person (i.e. the confederate). This sequence was performed twice, so six subjects performed (I,C,I_m,I,C,I_m) and five subjects (C,I,I_m,C,I,I_m). Presentation order made no significant difference.

A confederate was not present in the other half of the experiments. In this group (human-robot), an experimental run started with the individual performing a block of trials alone followed by a block of trials in which the individual worked with the robotic partner. The subjects knew there was not a human on the other side of the curtain. This sequence was performed twice, so eleven subjects performed (I,I_m,I,I_m).

The experimental apparatus was identical when the subjects were working as individuals and when subjects were working with the robotic partner, except that the small rotational inertia of the crank ($0.113\ kg \times m^2$) was doubled when they were working with the robotic partner. I doubled the inertia since there is twice the available force; both the subject and the robotic partner are applying forces. Also, since my human-human studies in chapter 4 doubled the inertia in the dyad case, doubling the inertia in the human-robot experiments allows the results to be directly compared. Each experiment took less than 30 minutes, for a total of 600 trials.

This experiment duplicates the setup and procedure from the human-human experiments described in chapter 4, except for the catch trials. The catch trials in the human-human experiment required the subjects to move in the same direction on two consecutive trials. In the human-robot experiment, this would not work since the robotic partner had only learned a trajectory for aiming at a target on the opposite side of the workspace. Collecting enough data to adequately and believably replicate this motion would have taken too long since fatigue becomes an issue with long experiments.

5.2. Robot Believability

The eleven subjects without a confederate knew they were working with a non-human agent. Ten of the eleven subjects with a confederate present said they thought they were working with a person and were surprised at the end of the experiment when I revealed that they worked with a robot and never worked with the other person. One of them had some doubts, but was not sure either way and this subject's results were not significantly different than the other ten subjects who worked with a confederate. At a conscious level, the subjects working with a confederate present believed they were working with a human partner. This indicates that the robotic partner was able to cognitively pass a Turing test.

5.3. Performance

The catch trials in the human-human experiment were different than the catch trials in the human-robot experiment, but the difference in completion times between one person and a dyad (or a person and a person working with the robotic partner) is comparable between experiments. Any difference of having or not having catch trials is eliminated because I am comparing the change in completion time within each experiment.

When two people work together, the dyad performs on average 54.5 ms faster than the average of the subjects working individually. When a person works with a simulated partner and a confederate is present, the human-robot pair is on average 5.8 ms faster than the subject is when working alone. There is a significant difference of 48.8 ms between the human-human and human-robot completion times ($t(11, 15) = 3.02, p = 0.006$). When there is not a confederate present, the subject working with the motor is 24.8 ms slower than working alone, but not a statistically significant amount ($t(11) = 1.52, p = 0.14$). The completion time results are summarized in figure 5.1.

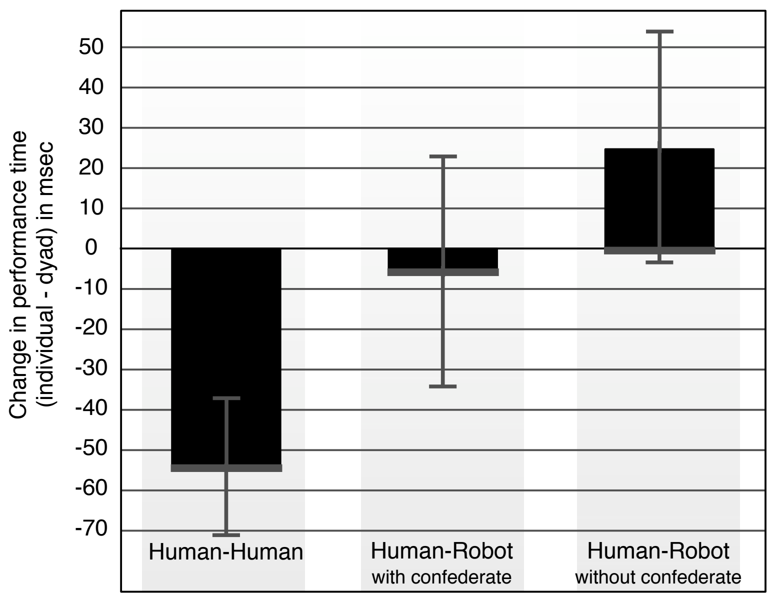

Figure 5.1. Comparing the improvement time between the three experimental groups. A negative value means the two people (or human-robot pair) worked faster together than the individuals. Two subjects working together are significantly faster than an individual working with a motor simulating a partner. When working with the robotic partner, subjects with a confederate present are faster than subjects without a confederate present, but not significantly.

When the subjects knew they were not working with a person, the robotic partner hindered the performance relative to working alone. When the subjects believed their partner was human, performance was similar to when the subjects worked alone, which suggests that the origin of forces in physical collaboration can affect the way in which a person will interact with them.

5.4. Social Facilitation

Since my experiments have two people working together, it is possible that social facilitation [**105**] is the cause of the faster performance in the dyadic condition. Social facilitation research has a long history [**97**] with many studies showing that simply having a person in a room observing a subject will lead to better performance on the given task. This is typically explained by the mere presence of others elevating drive levels. Mere presence tends to improve performance on simple, or well mastered, tasks and inhibits the performance on complex or, poorly mastered, tasks [**87**].

Sebanz et al.'s study [**89**] examined how two people in close proximity work when they are both working on different, but related, tasks. Each of the two people influenced the other's actions, sometimes beneficially and sometimes not, depending on the task. Schmidt and Turvey [**86**] and Amazeen et al. [**4**] studied how two people can coordinate a swinging pendulum only by looking at another person swinging a pendulum. Knoblich and Jordon [**49**] suggest that group coordination may be beneficial since each person in the group has fewer actions to deal with. Wegner and Zeaman [**101**], when studying groups and singles on a pursuit rotor task, said that social facilitation may have been responsible for some of the increased performance of groups, but it could not account for all of the effects they observed.

One might think my results for the improved performance of two people working together is due to each participant mutually watching their partner, which causes both participants to perform better. Social facilitation could possibly also explain why the human-robot-confederate group was faster than the human-robot group. Social facilitation effects could possibly account for part of the performance benefit, but the experimenter was watching the subjects in all cases, so there was always someone visually present, which is the sole requisite for improved performance in most social facilitation studies. Also, social facilitation has only been studied in terms of visual interaction, not physical interaction. During my experiments, the subjects could not see each other due to the curtain[2] and, thus, could only feel each other through the device. The only aspect of the task that is changing is whether they are holding the handle or not. It is possible this physical change could elicit the same effect, but in all the social facilitation studies I found, the two subjects are physically disconnected from each other. They are predominantly communicating through vision, which is the basis of most of the social facilitation literature.

My experiments examine physical communication, a different channel of conveying information to a partner, which also shows increased performance. This performance increase is likely due to the motor control systems of both people working cooperatively and not for the reasons explained in the social facilitation literature. My results suggest

[2]The first set of experiments I performed (section 4.1) did not have a curtain and similarly showed a performance increment, but these results are confounded by the lack of additional inertia for dyads and a different setup.

that social facilitation could be extended to include a similar "haptic presence" effect, which has not been studied in the literature. Even if haptic presence had a similar effect to social facilitation, it would not be able to fully explain all of the improved performance for dyads. Dyads improved by 54.5 ms, whereas the difference of performance increase between the human-robot and human-robot-confederate groups was only 24.8 ms, half the time.

5.5. Specialization

When working with the robotic partner, the subjects are given an easy and natural way to specialize (section 4.3). Each subject is completely responsible for all the force during deceleration, whereas the subject is free to choose their force during the acceleration phase. The motor applies enough force to accelerate both the crank and the subject's arm. Figure 5.2 shows the average force of all subjects working alone compared to the same

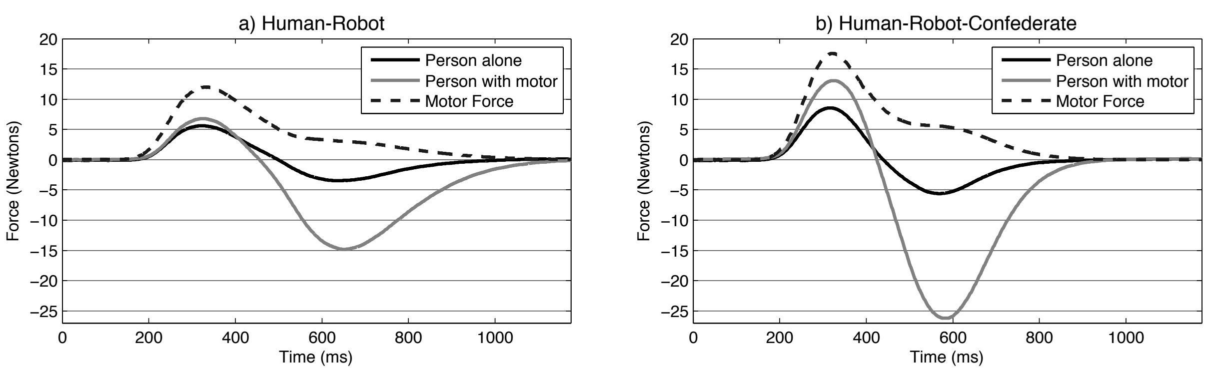

Figure 5.2. The average force profiles for individuals working with the motor. All subjects in (a) worked alone and knew they were not working with another human. All subjects in (b) assumed they were working with a confederate, but were working with the robotic partner as in (a). Both graphs show the average profile for the subjects when they worked alone, worked with the motor, and the force from the motor.

subjects working with the robotic partner. The two force profiles are averaged among all subjects where: a) the confederate is not present and b) the confederate was present. The difference between the two groups' overall completion time (and hence the magnitude of the forces) is partially due to social facilitation whereas the change from individual to dyad is the result of the perceived partner. The effects of social facilitation do not change during the experiment, so the difference in completion time from individual to dyad are comparable between groups. I want to study the effect of partner perception, which is why I am comparing the change from individual to dyad.

Even when the subjects consciously believed they were working with a human partner, the subjects did not develop specialization. In terms of a Turing test, this is an interesting split between consciously believing and physically acting different. Their forces did not show a significant change even though the subjects consciously believed the robot was human. The human-robot force profiles shown in figure 5.2 convey a very different strategy than two specialized people working together (figure 4.4b). The human-robot force profiles show a strategy that is remarkably similar to an individual performing the task alone.

With a subject's own amplified force applied as a feedforward force, the subjects are given an easy way to specialize. There is no motor force during deceleration, so the subject is required to apply all the force, whereas they are free to choose their force during the acceleration phase. Comparing the applied force from a subject working alone to a subject working with a motor shows that they do not change their feedforward force during the acceleration phase. They tend to accelerate similarly and actually apply slightly more peak force, but switch to deceleration earlier than they do when working alone. It seems that the subjects are pushing with a preprogrammed feedforward force similar to their previous trials alone, but begin to correct it and slow down the crank earlier.

The positive peak force was actually larger in the with-robot case than in the alone case. The larger peak force caused some subjects to apply a negative force earlier in the with-robot case. The robotic partner was designed to accelerate the crank, but the subjects also accelerated the crank, which resulted in too much acceleration force. Thus, the subjects working with the robot had to slow down the crank earlier. The subjects actually worked against the robotic partner and applied a larger overall force than was necessary. As a result, the subjects working with the robotic partner pushed harder during the deceleration phase than in the alone case and the dyad case.

When working alone, each subject knows what the result of their action will be, so each person can accurately predict the outcome of their action. When working with a human partner, the outcome is less predictable since a partner's action is unknown. Sebanz et al. [**89**] show that a person will internally represent the actions of a person nearby when working on a complementary action. Presumably, haptic interactions also allow two people working together to depend on their partner to complete the complementary action of specialization. I expected that a person would also learn to depend on the robotic partner to complete the complementary action of specialization, but this was not the case. Scheidt et al. [**84**] show that people can adapt to unpredictable forces within one trial. The robotic partner's forces are more predictable than a human's forces. Thus, it is very surprising that the subjects did not learn to work with a predictable robotic partner in the same way as they did when working with an unpredictable human partner.

The slight variations of a human partner are possibly beneficial to developing an advantageous cooperation. Human-human interaction requires a mutual adaptation where both subjects are constantly adapting to what their partner did on the previous trial. But, the other partner is also adapting, so they both are working with the knowledge from the past trial. This allows the two members to jointly explore different strategies, unlike adapting to the unchanging robotic partner.

In my human-robot cooperation experiment, the human is required to adapt to the robot. Corteville et al. [**16**] discussed a similar arrangement where their robot is adapting to the user's forces, but as the robot assists, the user's force changes. The human and robot are mutually adapting to each other. I expect that a revised robotic partner that uses both specialization, a learning algorithm, and slight variations from trial to trial would produce better results than specialization or a learning admittance control alone. These results have been published in [**78**].

5.6. Force Perturbations

The experiments up to this point have been primarily dealing with interaction throughout the entire trial where each trial can last a little over a second. This relatively long interaction allows two people to specialize their forces and to develop dyadic-contraction. In this section, I will examine a near instantaneous cooperation. I will discuss how effectively two subjects (or a person and a simulated partner) can reject a quick force pulse over the smaller timeframe of 100 ms.

5.6.1. Hypotheses

My first hypothesis is that dyads should be able to asymmetrically divide the response to a disturbance force. In section 4.3, I showed that two people working together can divide a task based on task phase, a temporal specialization. Since subjects can divide a task over a period of one second, I suspected that they may be able to asymmetrically divide a similar task, such as rejecting forces, over a smaller time period. It has been shown that subjects can adjust the impedance of their arm to fit the task at hand [**25**][**26**][**17**], but there have been no studies showing how two people can collaboratively complete such a task.

My second hypothesis examines why dyadic-contraction develops. As I discussed in section 4.5, I have two suppositions about why dyadic-contraction exists. One is that it serves as a mechanism for the two participants to feel the presence of the other. By pushing on the crank, each subject can feel the resistance of the other member, much like a firm handshake. My second is that this force serves to stabilize the subjects, much like co-contraction does in a single person. The experiment in this section will focus on the stabilizing effects of dyadic-contraction.

5.6.2. Participants

Twenty two (10 male; 8 left-handed) participants, age 18-24, from Northwestern University's Psychology participant pool participated after giving informed consent.

5.6.3. Apparatus and Stimuli

I used the same experimental apparatus as previous experiments, shown in figure 3.1. The procedure is similar to the experiments used in chapter 4 except the motor is applying additional forces.

I used the motor to apply a perturbation force for 100 ms after the subjects are inside the target for 700 ms. I measured the resulting displacement at the moment the perturbation force ended. This displacement relates to the impedance of the individual or dyad. This method is further discussed in section 2.6. The pulse is 10 N for individuals and 20 N for dyads, which accounts for twice the available impedance from two people, similar to doubling the inertia. One in six trials contained a perturbation.

The perturbation force varied in terms the side of the workspace and the direction. I will refer to the dominant and non-dominant sides of the workspace which correspond to the subjects' dominant hand. The dominant hand is the one they prefer to use and the one they used throughout this experiment. I will refer to the direction of the force perturbation as CW (clockwise) and CCW (counterclockwise) by looking at the device from above.

Additionally, during individual trials, I simulated the dyadic-contraction force. The force acts similar to a spring such that it pulls to the center of the workspace with a force proportionate to the distance from the center, but the force reaches a constant maximum force near the targets. The maximum force simulates a constant dyadic-contraction force near the targets while smoothly transitioning from side to side. The maximum constant force near the handles is one of three forces, either 0, 5, or 10 Newtons. The maximum constant force is the same for 48 continuous trials and then changes. The individual blocks

(154 trials) consist of the following sequence:

10 trials		Warm up phase (no force or perturbations)
48 trials each with 8 perturbations (random order)	{	No dyadic-contraction force
		5 N dyadic-contraction force
		10 N dyadic-contraction force

I chose to apply a force toward the center to keep the system, at worst, marginally stable. Another approach is to apply the spring force away from the center of the workspace, but this is inherently unstable and could have adverse affects if a subject were to let go of the handle. The dyadic-contraction force would be similar in either case since dyads necessarily have to apply both directions of force when they work together, so I chose the safer option.

5.6.4. Procedure

The trials consisted of the following sequence presented in random order:

154 trials each (random order)	{	A alone with simulated dyadic-contraction force
		B alone with simulated dyadic-contraction force
		A and B together with no dyadic-contraction force

Repeat same sequence once

Each block started with a warm up phase consisting of ten trials with no force perturbations. Each block of 154 trials contains 24 perturbations so on average $^1/_6$ of the trials contain a perturbation. In the individual case, there are two perturbations in each direction on each side (four combinations) with each of the three weight groups for a total of twelve possible combinations. There are two duplicates per block and four duplicates per 30 minute experiment. In the dyad case, there are two directions on each side and no weight groups so there are four possible combinations. There are six duplicates per block and twelve duplicates per 30 minute experiment.

5.6.5. Results

This data contains several interrelated factors, so I used an ANOVA. The dependent variable is the displacement and the interesting factors are the perturbation type and the simulated dyadic-contraction force in the individuals. The ANOVA tables are shown in tables 5.1, 5.2, and 5.3. The individual analysis includes an extra factor for the simulated dyadic-contraction force. Although the dyads had a dyadic-contraction force, the members jointly choose the force whereas I specified the dyadic-contraction force in the individual trials.

All of the ANOVA tables contain factors called ‘subject’ and ‘trial’. The ‘subject’ factor compares whether there was a statistically significant difference between the subjects. In all cases, there was at least one subject that was statistically significantly different. This is to be expected since different people perform differently. Similarly, the ‘trials’ factor determines if there is a learning effect, which was not present in any of these analyses.

The ‘type’ factor represents the four combinations of the location in the workspace and the direction of the perturbation force. The configuration of the arm and workspace can be seen more easily in figures 5.4 and 5.5. The four types are:

- Non-dominant CW: The crank is on the opposite side of the workspace from the subject's dominant arm (i.e. left side for a right handed subject) and the force perturbation is in a clockwise direction.
- Non-dominant CCW: The crank is on the opposite side of the workspace from the subject's dominant arm and the force perturbation is in a counterclockwise direction.
- Dominant CW: The crank is on the same side of the workspace from the subject's dominant arm (i.e. right side for a right handed subject) and the force perturbation is in a clockwise direction.
- Dominant CCW: The crank is on the same side of the workspace from the subject's dominant arm and the force perturbation is in a counterclockwise direction.

Due to the multiple types, there are two ways in which a dyad could be formed. One is by two same-handed individuals, which could consist of either two right handed or two left handed subjects. The other is by two different-handed individuals, which consists of one left handed and one right handed individual. In all previous experiments, this difference has not been statistically significant, but due to the asymmetric disturbance rejection characteristics of the arm configuration, it is relevant in this experiment. Since there is a significant difference, I will analyze the two dyad and the individual configurations separately.

Table 5.1. ANOVA Table for the individual trials. There is a statistically significant difference in the 'force' and 'type' factors.

Factor	$\sum$ of Squares	D.o.F.	Mean Square	F	Prob<F
Subject	200.95	21	9.569	19.54	0
Trial	21.57	47	0.459	0.94	0.5947
Force	25.83	2	12.913	26.37	0
Type	321.34	3	107.113	218.76	0
Error	480.83	982	0.49		
Total	1071.35	1055			

Table 5.2. ANOVA Table for the dyads with the same handedness. There is a statistically significant difference in the 'type' factor.

Factor	$\sum$ of Squares	D.o.F.	Mean Square	F	Prob<F
Subject	62.319	4	15.5797	52.63	0
Trial	8.831	47	0.1879	0.63	0.9662
Type	152.777	3	50.9256	172.05	0
Error	54.76	185	0.296		
Total	320.747	239			

Table 5.3. ANOVA Table for the dyads with different handedness. There is not a statistically significant difference in any of the interesting factors.

Factor	$\sum$ of Squares	D.o.F.	Mean Square	F	Prob<F
Subject	6.8191	3	2.27305	8.86	0
Trial	9.6372	47	0.20505	0.8	0.8107
Type	0.0279	1	0.02793	0.11	0.7419
Error	35.9234	140	0.2566		
Total	52.3992	191			

5.6.5.1. Individuals. Table 5.1 shows that there is a statistically significant difference for the force factor. Figure 5.3 shows that the displacement decreases as the dyadic-force increases. All three force levels (0, 5, and 10 Newtons) show a significant difference. Applying an external force can decrease the displacement due to a perturbation force, thus the impedance of the arm has increased. This decrease in displacement implies that dyadic-contraction would benefit dyads rejecting a disturbance, but, as I will discuss in section 5.6.5.2, this result does not scale to two humans.

Table 5.1 also shows that there is a statistically significant difference for the type factor. Figure 5.4 shows the configuration and displacements of each type of perturbation. There is a statistically significant difference between the dominant side cw and ccw and the non-dominant side. The non-dominant side cw and ccw are not statistically different.

The difference between the dominant and non-dominant sides is largely due to the configuration of the arm. In the non-dominant position, the arm is straight so the force is conveyed to the upper arm and the torso. On the dominant side, the elbow is bent so less mass is able to resist the motion and a larger disturbance occurs.

The difference in the directions on the dominant side is more interesting. The different impedances show that the subjects are rejecting the disturbance asymmetrically. The

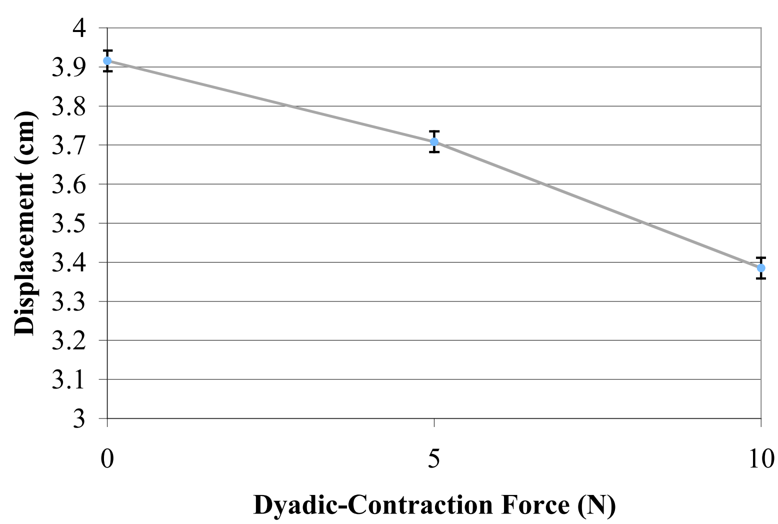

Figure 5.3. The displacement as a result of applying a perturbation force as well as the dyadic-contraction force. As the dyadic-contraction force is increases, the displacement decreases indicating that applying a constant external force can help a subject reject disturbance forces.

displacements do not change the dyadic-contraction force from the motor, so the only difference is how the muscles respond.

The larger displacements occur when the perturbation is in the direction that the subject is pushing (i.e. away from the center). The smaller displacements occur when the perturbation is in the direction opposing the direction the subject is pushing. This shows that the impedance of the arm is increasing with larger forces. These results match well with the stiffness ellipses published in several studies [**67**][**31**].

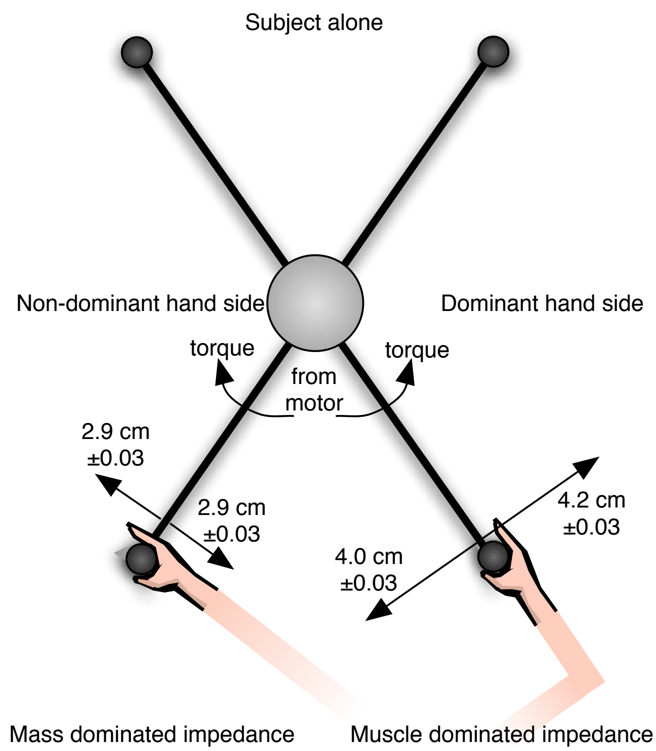

Figure 5.4. Configuration of the individual trials. The motor applies a torque that simulates the dyadic-contraction force. The motor also applies a perturbation force that results in a displacement (shown with the confidence intervals). The displacement difference between the dominant and non-dominant sides is largely due to the arm configuration. The displacement difference on the dominant side is largely due to different muscle impedance in each direction.

5.6.5.2. Dyads. Table 5.2 shows the ANOVA results for only the 5 dyads that were composed of same handed individuals. Table 5.3 shows the ANOVA results for only the 6 dyads that were composed of different handed individuals (i.e. a left and right handed). Figure 5.5 shows the configuration of same and different handed dyads along with their displacements for each type.

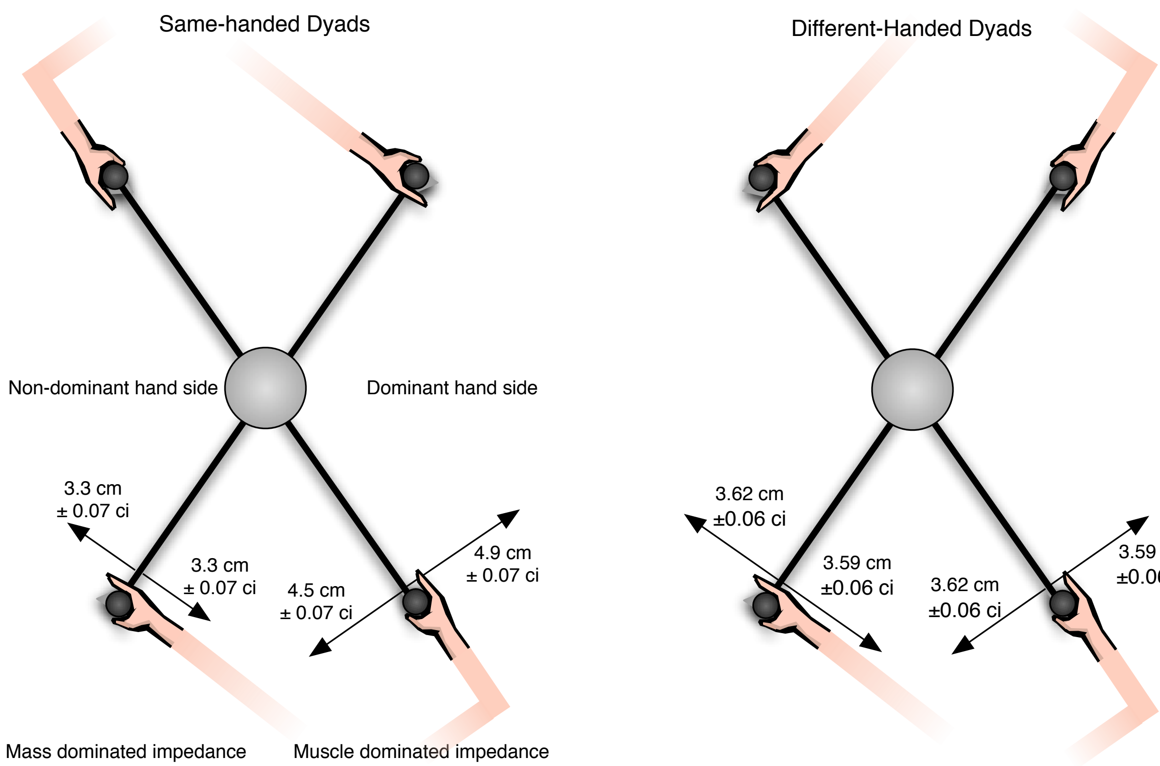

Figure 5.5. Configuration of the dyad trials. The motor applies a perturbation force that results in the displacement (shown with confidence intervals). There are two arrangements of dyads: same-handed or different-handed. The displacement difference between the dominant and non-dominant sides of the same-handed dyads is largely due to the arm configuration. The displacement difference on the dominant side is largely due to different muscle impedance in each direction. The different-handed dyads use a combination of the high and low impedance on each side, which results in a constant displacement for each side and location.

Different-handed dyads do not have a statistically significant difference between side or direction of force perturbation. This makes sense since the configuration of the arms is opposite, so the highest impedance of one member will be paired with the lowest impedance from the other member. Due to the configuration of their arms, the dyad does asymmetrically divide the force during different motions.

Table 5.2 shows that there is a statistically significant difference for the type factor in same-handed dyads. In fact, same-handed dyads show the same pattern of displacements as individuals do. There is a statistically significant difference between the dominant side cw and ccw and the non-dominant side. The non-dominant side cw and ccw are not statistically different. Same-handed dyads do not divide the restoring force between them. There is no specialization like there was during the task movement.

The only difference between the individual and same-handed dyads is that the dyads can choose the amount of dyadic-contraction force. Figure 5.6 shows a histogram of all

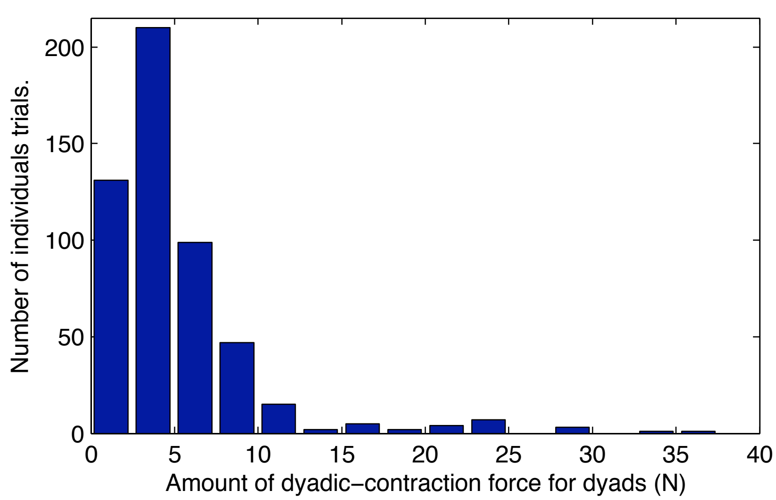

Figure 5.6. A histogram of the amount of dyadic-contraction between each dyad during the perturbation trials.

the dyadic-contraction forces from all the dyads. The forces are not drastically different than the 0, 5, and 10 Newton forces used in the individual trials.

5.6.5.3. Discussion. Figure 5.7 shows a comparison of the individuals to both kinds of dyads. Although same-handed dyads and individuals exhibit the same pattern of increased displacement, the dyads always have significantly larger displacements. In each

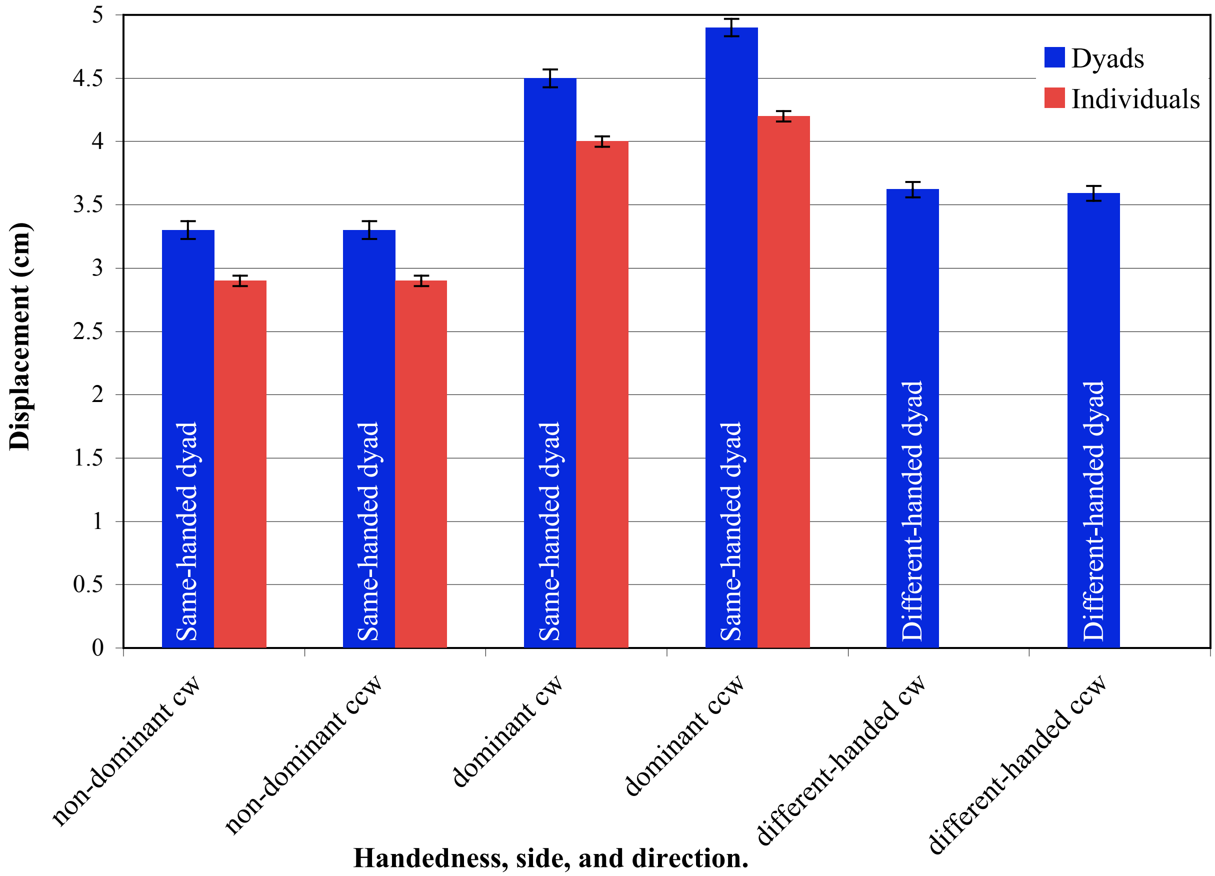

Figure 5.7. Bar graph showing the average displacement of all the individual and dyad perturbation trials. The same-handed dyads shows a similar pattern as individuals, but the dyads had a larger displacement in every case. Unlike specialization, the same-handed dyads were not able to improve their performance by joining forces. The different-handed dyads were able to use the best of each subject's ability for each type to create a constantly good joint effort.

of the four types, the same-handed dyads had a larger displacement than individuals. The two subjects do not simply add the impedance of their arms, something less efficient is occurring. The impedance of each arm should add as if they are in parallel, but they do not. I expected that the impedance would at least double, which would have resulted in the same displacement since I doubled both the perturbation force and the inertia of the crank. In the case of different-handed dyads, the impedance does appear to be in the middle of a dominant and non-dominant individual, which is closer to expected if the impedances were adding in parallel.

5.7. Discussion

Two humans working together are able to perform faster than their constituent members and specialize over the entire task phase, but they are worse at rejecting forces. When a human works with a robot programmed to mimic certain parts of the human-human interaction, they are slower than the individual and do not specialize, but this pair is better at rejecting forces. This set of experiments suggests that humans are able to communicate physically at a relatively slow rate, such as over an entire motion and not on quick tasks such as rejecting perturbations. A robotic partner is the opposite since the human-robot pair did not specialize, but was better at rejecting perturbations.

References

[1] J. Accot and S. Zhai. Beyond fitts' law: Models for trajectory-based hci tasks. *Proceedings of the CHI 97*, pages 295–302, 1997.

[2] Johnny Accot and Shumin Zhai. More than dotting the i's - foundations for crossing-based interfaces. *In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, 4(1):73–80, 2002.

[3] Kenzo Akazawa, Ted E. Milner, and Richard B. Stein. Modulation of reflex emg and stiffness in response to stretch of human finger muscle. *Journal of Neurophysiology*, 49(1):16–27, 1983.

[4] Polemnia G. Amazeen, R. C. Schmidt, and M. T. Turvey. Frequency detuning of the phase entrainment dynamics of visually coupled rhythmic movements. *Biological Cybernetics*, 72:511–518, 1995.

[5] Michael Argyle and Janet Dean. Eye-contact, distance and affiliation. *Sociometry*, 28(3):289–304, 1965.

[6] D. P. Barnes and M. S. Counsell. Haptic communication for remote mobile manipulator robot operations. *Proceedings of the 8th Topical Meeting on Robotics and Remote Systems, American Nuclear Society, Pittsburgh, PA*, 1999.

[7] Cagatay Basdogan, Chih-Hao Ho, and Mandayam A. Srinivasan. An experimental study on the role of touch in shared virtual environments. *ACM Transactions on Computer-Human Interaction,*, 7(4):443–460, 2000.

[8] Ansgar Bittermann, Kolja Kuhnlenz, and Martin Buss. On the evaluation of emotion expressing robots. In *International Conference of Robotics and Automation (ICRA)*, Rome, 2007.

[9] E. Bonabeau, M. Dorigo, and G. Theraulaz. Inspiration for optimization from social insect behaviour. *Nature*, 406:39–42, 2000.

[10] Thomas Brashers-Krug, Reza Shadmehr, and Emilio Bizzi. Consolidation in human motor memory. *Nature*, 382:252–255, 1996.

[11] J. L. Bruning, D. K. Sommer, and B. R. Jones. The motivational effects of cooperation and competition in the means-independent situation. *Journal of Social Psychology*, 68:269–274, 1966.

[12] T.L. Chartrand and J. Bargh. The chameleon effect: the perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76(6):893–910, 1999.

[13] Vikram S. Chib, James L. Patton, Kevin M. Lynch, and Ferdinando A. Mussa-Ivaldi. Haptic identification of surfaces as fields of force. *Journal of Neurophysiology*, 95:1068–1077, 2006.

[14] V.S. Chib, J.L. Patton, K.M. Lynch, and F.A. Mussa-Ivaldi, . Haptic discrimination of perturbing fields and object boundaries. *HAPTICS 2004*, 0-7695-2112-6/04, 2004.

[15] Seungmoon Choi, Laron Walker, Hong Z. Tan, Scott Crittenden, and Ron Reifenberger. Force constancy and its effect on haptic perception of virtual surfaces. *ACM Transactions on Applied Perception*, 2(2):89–105, 2005.

[16] B. Corteville, E. Aertbelien, H. Bruyninckx, J. De Schutter, and H. Van Brussel. Human-inspired robot assistant for fast point-to-point movements. *International Conference of Robotics and Automation (ICRA), Rome*, 2007.

[17] Mohammad Darainy, Nicole Malfait, Paul L. Gribble, Farzad Towhidkhah, and David J. Ostry. Learning to control arm stiffness under static conditions. *journal of Neurophysiology*, 92:3344–3350, 2004.

[18] B. P. DeJong, E. L. Faulring, J. E. Colgate, M. A. Peshkin, H. Kang, Y. S. Park, and T. F. Ewing. Lessons learned from a novel teleoperation testbed. *Industrial Robot*, 33(3):187–193, 2006.

[19] I. Elhajj, H. Hummert, N. Xi, and Y. Liu. Real-time haptic feedback in internet-based telerobotic operation. *IEEE International Conference on Robotics and Automation, San Francisco*, 2000.

[20] C. M. B. Evans and S. N. Baker. Task-dependent intermanual coupling of 8-hz discontinuities during slow finger movements. *European Journal of Neuroscience*, 18:453–456, 2003.

[21] David Feygin, Madeleine Keehner, and Frank Tendick. Haptic guidance: Experimental evaluation of a haptic training method for a perceptual motor skill. *Proceedings of the 10th Symposium On Haptic Interfaces For Virtual Environments and Teleoperator Systems (HAPTICS)*, 2002.

[22] Edmund Field and Don Harris. A comparative survey of the utility of cross-cockpit linkages and autoflight systems' backfeed to the control inceptors of commercial aircraft. *Ergonomics*, 41(10):1462–1477, 1998.

[23] P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47:381–391, 1954.

[24] T. Flash and N. Hogan. The coordination of arm movements: An experimentally confirmed model. *Journal of Neuroscience*, 5(7):1688–1703, 1985.

[25] David W. Franklin, Etienne Burdet, Rieko Osu, Mitsuo Kawato, and Theodore E. Milner. Functional significance of stiffness in adaptation of multijoint arm movements to stable and unstable dynamics. *Experimental Brain Research*, 151:145–157, 2003.

[26] David W. Franklin and Theodore E. Milner. Adaptive control of stiffness to stabilize hand position with large loads. *Experimental Brain Research*, 152:211–220, 2003.

[27] Sommer Gentry. *Dancing cheek to cheek: Haptic communication between partner dancers and swing as a finite state machine*. Ph.D thesis, Massachusetts Institute of Technology, 2005.

[28] Sommer Gentry, Eric Feron, and Roderick Murray-Smith. Human-human haptic collaboration in cyclical fitts' tasks. *International Conference on Intelligent Robots and Systems*, 2005.

[29] S.J. Glynn, R. Fekieta, and R.A. Henning. Use of force-feedback joysticks to promote teamwork in virtual teleoperation. *Virtual Teleoperation Proc. Of the Human Factors and Ergonomics Society 45th Annual Meeting*, 2001.

[30] S.J. Glynn and R.A. Henning. Can teams outperform individuals in a simulated dynamic control task? *Proceedings Of The Human Factors And Ergonomics Society Annual Meeting (ISSN: 10711813)*, 6:141–144, 2000.

[31] Hiroaki Gomi and Rieko Osu. Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *Journal of Neuroscience*, 18(21):8965–8978, 1998.

[32] Andrew M. Gordon, Goran Westling, Kelly J. Cole, and Roland S. Johansson. Memory representations underlying motor commands used during manipulation of common and novel objects. *Journal of Neurophysiology*, 69(6):1789–1796, 1993.

[33] G. L. Gottlieb, C. H. Chen, and D. M. Corcos. Nonlinear control of movement distance at the human elbow. *Experimental Brain Research*, 112:289–297, 1996.

[34] G. L. Gottlieb, D. M. Corcos, and G. C. Agarwal. Organizing principles for single-joint movements. i. a speed insensitive strategy. *Journal of Neurophysiology*, 62:342–357, 1989.

[35] G. L. Gottlieb, D. M. Corcos, G. C. Agarwal, and M. L. Latash. Principles underlying single joint movement strategies. *J.M. Winters, S.L.Y. Woo, "Multiple Muscle Systems," Springer-Verlag, New York*, pages 236–250, 1990.

[36] Ewan C. Grant. Human facial expression. *Man*, 4(4):525–536, 1969.

[37] P. L. Gribble, L. I. Mullin, N. Cothros, and A. Mattar. Role of cocontraction in arm movement accuracy. *Journal of Neurophysiology*, 89:2396 2405, 2003.

[38] Paul Griffiths and R. Brent Gillespie. Shared control between human and machine: Haptic display of automation during manual control of vehicle heading. *Haptics 2004*, pages 358–366, 2004.

[39] M. Hallett, B. Shahani, and R. Young. Emg analysis of stereotyped voluntary movements in man. *Journal of Neurosurgy Psychiatry*, 38:1154 – 1162, 1975.

[40] B. Hannaford and L. Stark. Roles of the elements of the triphasic control signal. *Expimental Neurology*, 90:619 – 634, 1985.

[41] J. P. Hespanha, M. McLaughlin, G. S. Sukhatme, M. Akbarian, R. Garg, and W. Zhu. Haptic collaboration over the internet. *Fifth Phantom Users Group Workshop, Aspen, CO*, 2000.

[42] K. Hinckley, E. Cutrell, S. Bathiche, and T. Muss. Quantitative analysis of scrolling techniques. *Conference on Human Factors in Computer Systems*, pages 65–72, 2002.

[43] Jessica K. Hodgin, Waybe L. Wooten, David C. Brogan, and James F. O'Brien. Animating human athletes. *In Proceedings of the ACM SIGCHI Conference on Computer Graphics and Interactive Techniques*, pages 71–78, 1995.

[44] Roger J. Hubbold. Collaborative stretcher carrying: a case study. *Eighth Eurographics Workshop on Virtual Environments*, pages 7–12, 2002.

[45] Roland S. Johansson. Sensory input and control of grip. *Sensory guidance of movement: Novartis Foundation Symposium*, 218:45–59, 1998.

[46] A. Karniel, R. Meir, and G. F. Inbar. Exploiting the virtue of redundancy. *International Joint Conference on Neural Networks, Washington, DC*, 1999.

[47] Tim Kiska. Death on the job jury awards $10 million to heirs of man killed by robot at auto plant., August 11, 1983, The Philadelphia Inquirer.

[48] V. Klingspor, J. Demiris, and M. Kaiser. Human-robot communication and machine learning. *Applied Artificial Intelligence Journal*, 11:719–746, 1997.

[49] Gunther Knoblich and Jerome Scott Jordon. Action coordination in groups and individuals: Learning anticipatory control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(5):1006–1016, 2003.

[50] J. W. Krakauer, M. F. Ghilardi, and G. Ghez. Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neuroscience*, 2:1026–1031, 1999.

[51] C. Ronald Kube and Eric Bonabeua. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30:85–101, 2000.

[52] F. Lacquaniti and C. Maioli. Distributed control of limb position and force. *Requin GESaJ (ed) Tutorials in Motor Behavior II*, pages 31–54, 1992.

[53] Seung-Ik Lee, Gunn-Yong Park, and Joong-Bae Kim. Natural emotion expression of a robot based on reinforcer intensity and contingency. In *International Conference of Robotics and Automations (ICRA)*, Rome, 2007.

[54] Ming Li and Allison Okamura. Recognition of operator motions for real-time assistance using virtual fixtures. *Haptic Interfaces for Virtual Environment and Teleoperator Systems Symposium*, pages 125–131, 2003.

[55] J. C. A. Looser. *3D games as motivation in Fitts' law experiments*. Masters thesis, University of Canterbury, 2002.

[56] J. H. McAuley, J. C. Rothwell, and C.D. Marsden. Frequency peaks of tremor, muscle vibration and electromyographic activity at 10hz, 20hz and 40hz during human finger muscle contraction may reflect rhythmicities of central neural firing. *Experimental Brain Research*, 114:525–541, 1997.

[57] Theodore E. Milner. Adaptation to destabilizing dynamics by means of muscle cocontraction. *Experimental Brain Research*, 143:406–416, 2002.

[58] Mark W. Moffett. Cooperative food transport by an asiatic ant. *National Geographic Research*, 4:386–394, 1988.

[59] Mark W. Moffett. Ant foraging. *National Geographic Research*, 8:220–231, 1992.

[60] D. Mottet, Y. Guiard, T. Ferrand, and R. J. Bootsma. Two-handed performance of a rhythmical fitts task by individuals and dyads. *Journal of Experimental Psychology: Human Perception and Performance*, 27(6):1275–1286, 2001.

[61] F. A. Mussa-Ivaldi, N. Hogan, and E. Bizzi. Neural, mechanical, and geometric factors subserving arm posture in humans. *Journal of Neuroscience*, 5(10):2732–2743, 1985.

[62] Rehabilitation Institute of Chicago. Gait restoration/lokomat project. http://www.smpp.northwestern.edu/mars/Project2.htm, 5/18/2006.

[63] Rieko Osu, David W. Franklin, Hiroko Kato, Hiroaki Gomi, Kazuhisa Domen, Toshinori Yoshioka, and Mitsuo Kawato. Short- and long-term changes in joint co-contraction associated with motor learning as revealed from surface emg. *Journal of Neurophysiol*, 88:991–1004, 2001.

[64] A. Ozyurek. Do speakers design their cospeech gestures for their addressees? the effects of addressee location on representational gestures. *Journal of Memory and Language*, 46:688–704, 2002.

[65] J. L. Patton and P. Elkins. Training with a bimanual-grasp beneficially influences single limb performance. *Society for Neuroscience, Orlando, FL, USA*, 2002.

[66] G. di Pellegrino, L. Fadiga, L. Fogassi, V. Gallese, and G. Rizzolatti. Motor facilitation during action observation: a magnetic stimulations study. *Journal of Neurophysiology*, 73:2608–2611, 1992.

[67] Eric J. Perreault, Robert F. Kirsch, and Patrick E. Crago. Effects of voluntary force generation on the elastic components of endpoint stiffness. *Experimental Brain Research*, 141:312–323, 2001.

[68] J. Prodoehl, G. Gottlieb, and D. Corcos. The neural control of single degree-of-freedom elbow movements: Effect of starting joint position. *Experimental Brain Research*, 153:7–15, 2003.

[69] M. Rahman, R. Ikeura, and K. Mitzutani. Impedance characteristics of human arm for cooperative robot. *International Conference on Control, Automation and Systems*, pages 1455–1460, 2002.

[70] M. Rahman, R. Ikeura, and K. Mizutani. Control characteristics of two humans in cooperative task and its application to robot control. *Proceedings of 2000 International Conference on Industrial Electronics Control and Instrumentation*, pages 1773–1778, 2000.

[71] M. Rahman, R. Ikeura, and K. Mizutani. Analysis of cooperation characteristics of two humans in moving an object. *Proceedings of the International Conference on Mechatronics and Information Technoloy*, 19:454–458, 2001.

[72] M. Rahman, R. Ikeura, and K. Mizutani. Cooperation characteristics of two humans in moving an object. *Machine Intelligence and Robotic Control*, 4(2):43–48, 2002.

[73] C. F. Ramos, S. S. Hacisalihzade, and L. W. Stark. Behavior space of a stretch reflex model and its implications for the neural control of voluntary movement. *Medical and Biological Engineering and Computing*, 28:15–23, 1990.

[74] Kyle Reed. *Specialization in Dyadic Shared Manual Tasks*. Masters thesis, Northwestern Univerisity, 2004.

[75] Kyle Reed, Michael Peshkin, J. Edward Colgate, and James Patton. Initial studies in human-robot-human interaction: Fitts' law for two people. *International Conference of Robotics and Automation (ICRA), New Orleans*, 2004.

[76] Kyle Reed, Michael Peshkin, Mitra J. Hartmann, Marcia Grabowecky, James Patton, and Peter M. Vishton. Haptically linked dyads: are two motor control systems better than one? *Psychological Science*, 17(5):365–366, 2006.

[77] Kyle Reed, Michael Peshkin, Mitra J. Hartmann, James Patton, Peter M. Vishton, and Marcia Grabowecky. Haptic cooperation between people, and between people and machines. *International Conference on Intelligent Robots and Systems (IROS)*, 2006.

[78] Kyle B. Reed, James Patton, and Michael Peshkin. Replicating human-human physical interaction. *International Conference of Robotics and Automation (ICRA), Rome*, 2007.

[79] Kyle B. Reed, Michael Peshkin, Mitra J. Hartmann, J. Edward Colgate, and James Patton. Kinesthetic interaction. In *International Conference on Rehabilitation Robotics (ICORR)*, pages 569–574, Chicago, IL, 2005. IEEE.

[80] David J. Reinkensmeyer, Peter S. Lum, and Steven L. Lehman. Human control of a simple two-hand grasp. *Bilogical Cybernetics*, 67(6):553–564, 1992.

[81] G. Rizzolatti, L. Fogassi, and V. Gallese. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2(9):661–670, 2001.

[82] D. A. Rosenbaum. Human movement initiation: Specification of arm, direction, and extent. *Journal of Experimental Psychology: General*, 109:444–474, 1980.

[83] E. Sallnas and S. Zhai. Collaboration meets Fitts' law: Passing virtual objects with and without haptic force feedback. *Proceedings of INTERACT, IFIP Conference on Human-Computer Interaction*, pages 97–104, 2003.

[84] Robert A. Scheidt, Jonathan B. Dingwell, and Ferdinando A. Mussa-Ivaldi. Learning to move amid uncertainty. *Journal of Neurophysiol*, 86, 2001.

[85] R. C. Schmidt, Claudia Carello, and M. T. Turvey. Phase transitions and critical fluctuations in the visual coordination of rhythmic movements between people. *Journal of Experimental Psychology: Human Perception and Performance*, 16(2):227–247, 1990.

[86] R. C. Schmidt and M. T. Turvy. Phase-entrainment dynamics of visually couple rhythmic movements. *Biological Cybernetics*, 70:669–376, 1994.

[87] Bernd H. Schmitt, Thomas Gilovich, Natan Goore, and Lisa Joseph. Mere presence and social facilitation: One more time. *Journal of Experimental Social Psychology*, 22:242–248, 1986.

[88] Natalie Sebanz, Harold Bekkering, and Gunther Knoblich. Joint action: bodies and minds moving together. *TRENDS in Cognitive Sciences*, 10(2):70–76, 2006.

[89] Natalie Sebanz, Gunther Knoblich, and Wolfgang Prinz. Representing others' actions: just like one's own? *Cognition*, 88:11–21, 2003.

[90] Reza Shadmehr and Ferdinando A. Mussa-lvaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5):3208–3224, 1994.

[91] R.E. Shaw, E. Kadar, M. Sim, and D.W. Repperger. The intentional spring: A strategy for modeling systems that learn to perform intentional acts. *Journal of Motivational Behavior*, 24(1):3–28, 1992.

[92] S. S. Shergill, P. M. Bays, C. D. Frith, and D. M. Wolpert. Two eyes for an eye: The neuroscience of force escalation. *Science Magazine*, 301:187, 2003.

[93] Micah Steele and Brent Gillespie. Shared control between human and machine. *Human Factors and Ergonomics Society 45th Annual Meeting, Minneapolis, MN*, 2001.

[94] L.G. Summers, J.H. Shannon, T.R. White, and R.J. Shiner. Fly-by-wire sidestick controller evaluation. *SAE Technical Paper 871761, Aerospace Technology Conference and Exposition, Long Beach, CA, Society of Automotive Engineers*, 1987.

[95] Stephan P. Swinnen and Nicole Wenderoth. Two hands, one brain: cognitive neuroscience of bimanual skill. *TRENDS in Cognitive Sciences*, 8(1):18–25, 2004.

[96] T. Takubo, H. Arai, Y. Hayashibara, and K. Tanie. Human-robot cooperative manipulation using a virtual nonholonomic constraint. *The International Journal of Robotics Research*, 21:541–553, 2002.

[97] Norman Triplett. The dynamogenic factors in pacemaking and competition. *American Journal of Psychology*, (9):507–533, 1898.

[98] Alan Turing. Computing machinery and intelligence. *Mind*, LIX(236):433–460, 1950.

[99] Y. Uno, M. Kawato, and R. Suzuki. Formation and control of optimal tra jectory in human multijoint arm movement. minimum torque-change model. *Biological Cybernetics*, 61(2):89–101, 1989.

[100] M.Y. Yambay Valiente. *Design of a Unicycle Cobot*. Masters thesis, Northwestern University, 2001.

[101] N. Wegner and D. Zeaman. Team and individual performance on a motor learning task. *Journal of General Psychology*, 55:127–142, 1956.

[102] Tom Worsnopp. *Design of a Unicycle Cobot Controller*. Masters thesis, Northwestern Univerisity, 2003.

[103] H. A. Yanco and J. L. Drury. A taxonomy for human-robot interaction. *Proceedings of the AAAI Fall Symposium on Human-Robot Interaction, AAAI Technical Report FS-02-03, Falmouth, MA*, pages 111–119, 2002.

[104] Holly A. Yanco and Jill Drury. Classifying human-robot interaction: An updated taxonomy. *IEEE International Conference on Systems, Man and Cybernetics*, pages 2841–6, 2004.

[105] Robert B. Zajonc. Social facilitation. *Science Magazine*, 149:269–274, 1965.

[106] S. Zhai, S. Conversy, M. Beaudouin-Lafon, and Y. Guiard. Human on-line response to target expansion. *Proceedings of CHI 2003, ACM Conference on Human Factors in Computing Systems, Fort Lauderdale, Florida*, pages 177–184, 2003.

Vita

Kyle Brandon Reed

EDUCATION

Northwestern University, Evanston, Illinois

Ph.D. in Mechanical Engineering **June 2007**

M.S. in Mechanical Engineering **December 2004**

University of Tennessee, Knoxville, Tennessee

B.S. in Mechanical Engineering **May 2001**

- Minors in Material Science and Engineering Communications

Los Alamos, New Mexico

Graduated high school **May 1997**

Born and raised in Los Alamos, NM **August 20, 1978**

JOURNAL PUBLICATIONS

K. Reed, M. Peshkin, M.J. Hartmann, M. Grabowecky, J. Patton, and P.M. Vishton. "Haptically Linked Dyads: Are Two Motor-Control Systems Better Than One?" *Psychological Science*, 17(5), May, 2006.

CONFERENCE PUBLICATIONS

K. Reed, M. Peshkin, M.J. Hartmann, and J. Patton. "Replicating Human-Human Physical Interaction," *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, Rome, April 2007.

K. Reed, M.J. Hartmann, J. Patton, P.M. Vishton, M. Grabowecky, and M. Peshkin . "Haptic Cooperation Between People, and Between People and Machines," *Proc. of IEEE International Conference on Intelligent Robots and Systems (IROS)*, Beijing, October 2006.

K. Reed, M. Peshkin, M.J. Hartmann, J.E. Colgate, and J. Patton. "Kinesthetic Interaction," *Proc. of the 9th International Conference on Rehabilitation Robotics (ICORR)*, Chicago, June, 2005.

K. Reed, M. Peshkin, J.E. Colgate, and J. Patton. "Initial Studies in Human-Robot-Human Interaction: Fitts' Law for Two People," *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, New Orleans, April 2004.

APPENDIX A

Demonstrations

Throughout my time in LIMS, I helped to organize and give many lab demonstrations. This included tours for Cub Scouts, Bring Your Daughters to Work Day, Career Day for Girls, and students from surrounding schools. A part of these tours was demonstrating my device. As you can see in figure A.1, many of the kids greatly enjoyed interacting with my 'robot'.

The most popular demo was arm wrestling. This generally consisted of me asking everyone if they had an older brother or sister. I then asked if they ever wanted to be stronger than their older sibling. Most of the time, this would intrigue them and they would get curious about how they could beat them at arm wresting. So, I would have one of the kids hold the handle with one finger and the other with their entire hand. I

Figure A.1. Kids happily interacting with my device during a demonstration.

Figure A.2. Arm wrestling: one finger against an arm.

would amplify their force as described in section 2.3 so one kid would appear to have an incredibly strong finger while the other kid was straining. It gets the point across very quickly. Figure A.2 shows two people arm wresting. If you look close, the two people on the left are only holding on with a single finger.

Once I have them all intrigued and excited about enhanced strength, I begin to explain how this device works and why it is useful. If you have not figured out why this device is useful, I suggest you go back and reread my thesis. Figure A.3 shows me explaining how the sensors, motor, and computer all work together.

Figure A.3. Explaining how my device works during a demonstration.

APPENDIX B

Force Profiles

This Appendix includes data for the 15 dyads used to publish the Psychological Science paper [**76**]. The figures show the acceleration, velocity, and position of the crank along with the force from each subject as well as the net force. I have plotted the ensemble average (solid line) along with the ensemble standard deviation (dotted lines). These are simple averages and standard deviations, so only 68% of the data is necessarily represented. They are sorted in order of increasing average completion time.

Specialization is clearly shown in figures B.11, B.13, and B.15. It can be seen to a lesser extent in figures B.2, B.7, and B.8.

Dyadic-contraction is clearly shown in figures B.2, B.4, B.5, B.6, B.11, and B.15.

Table B.1. All of the average times for each dyad and constituent members along with the 95% confidence intervals.

Average dyad CT	Member A's average CT	Member B's average CT
544.3 ± 74.3	619.0 ± 118.1	607.6 ± 99.7
568.0 ± 61.3	649.6 ± 146.6	711.9 ± 125.2
571.2 ± 74.8	592.0 ± 74.1	616.3 ± 115.5
587.3 ± 99.4	655.1 ± 132.8	552.2 ± 85.6
608.4 ± 77.9	643.1 ± 129.9	633.2 ± 81.8
622.1 ± 76.1	711.0 ± 119.4	631.2 ± 54.7
636.7 ± 86.9	756.4 ± 115.4	686.1 ± 117.5
641.6 ± 90.4	654.9 ± 50.9	620.4 ± 53.9
668.1 ± 121.0	821.2 ± 217.8	751.2 ± 140.0
678.5 ± 66.5	839.6 ± 141.1	690.5 ± 110.6
681.4 ± 58.6	750.4 ± 97.7	735.0 ± 82.3
693.2 ± 126.2	775.0 ± 96.0	710.1 ± 124.5
703.9 ± 55.2	704.6 ± 56.2	725.7 ± 70.0
718.2 ± 79.4	803.6 ± 114.7	718.2 ± 62.0
767.1 ± 129.9	780.7 ± 163.5	869.4 ± 159.3

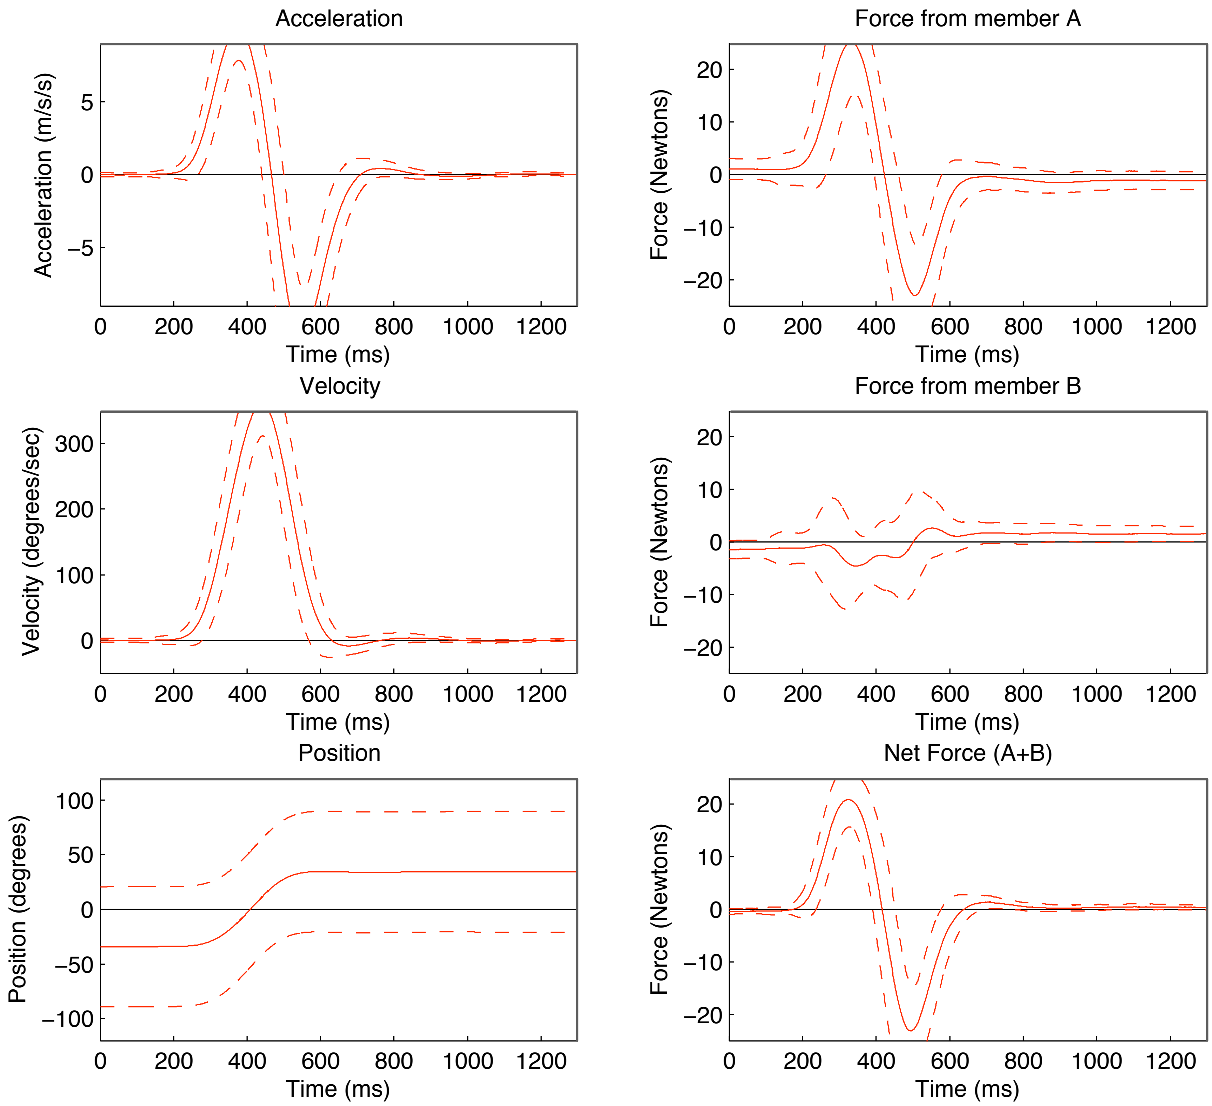

Figure B.1. Average Completion time: 544 ms

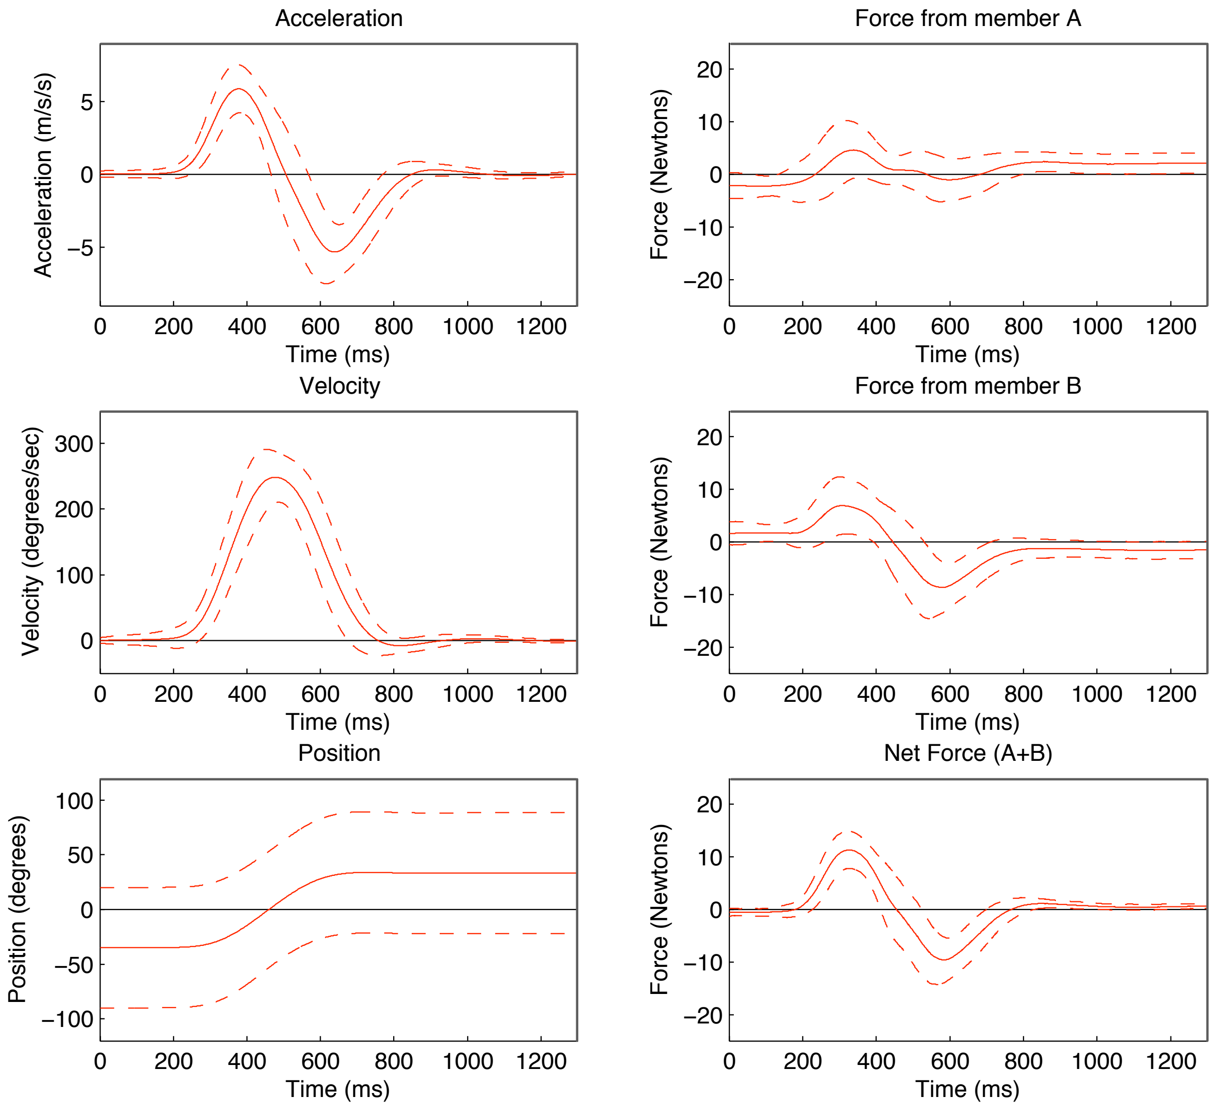

Figure B.2. Average Completion time: 567 ms

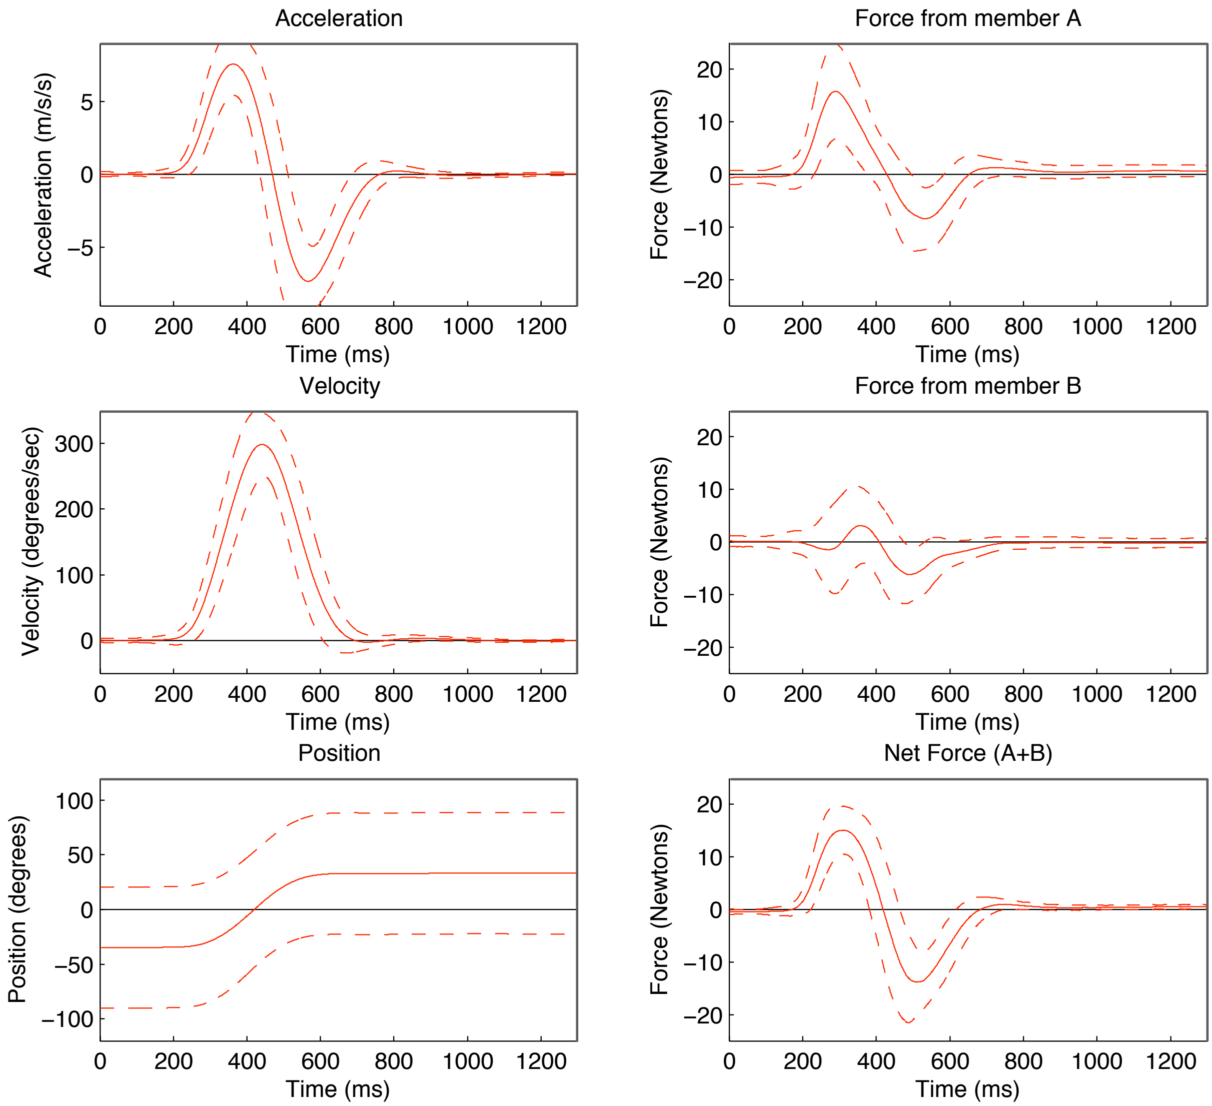

Figure B.3. Average Completion time: 568 ms

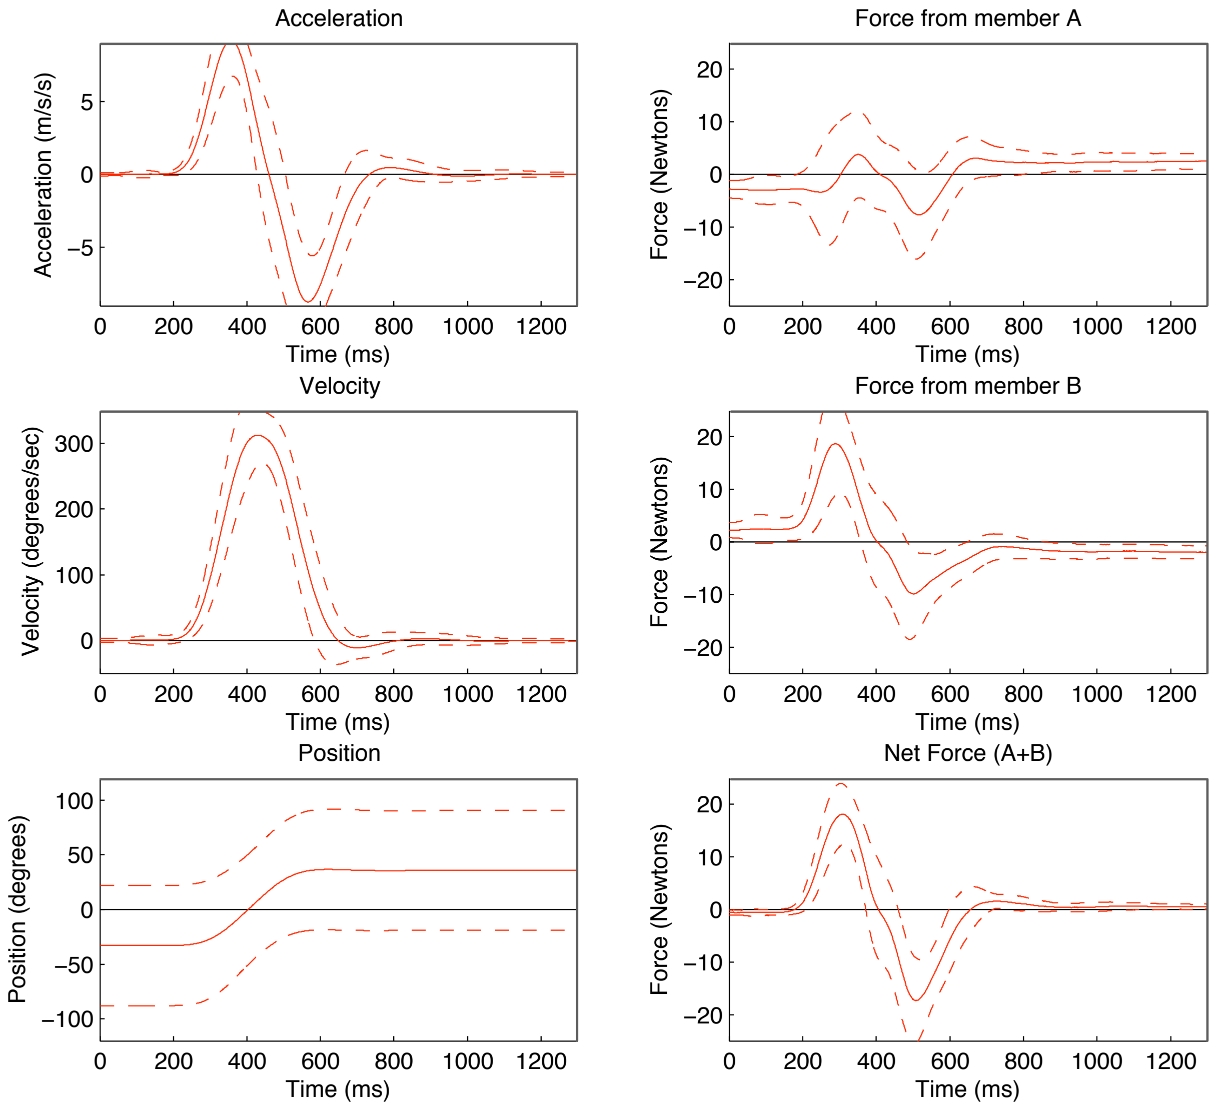

Figure B.4. Average Completion time: 571 ms

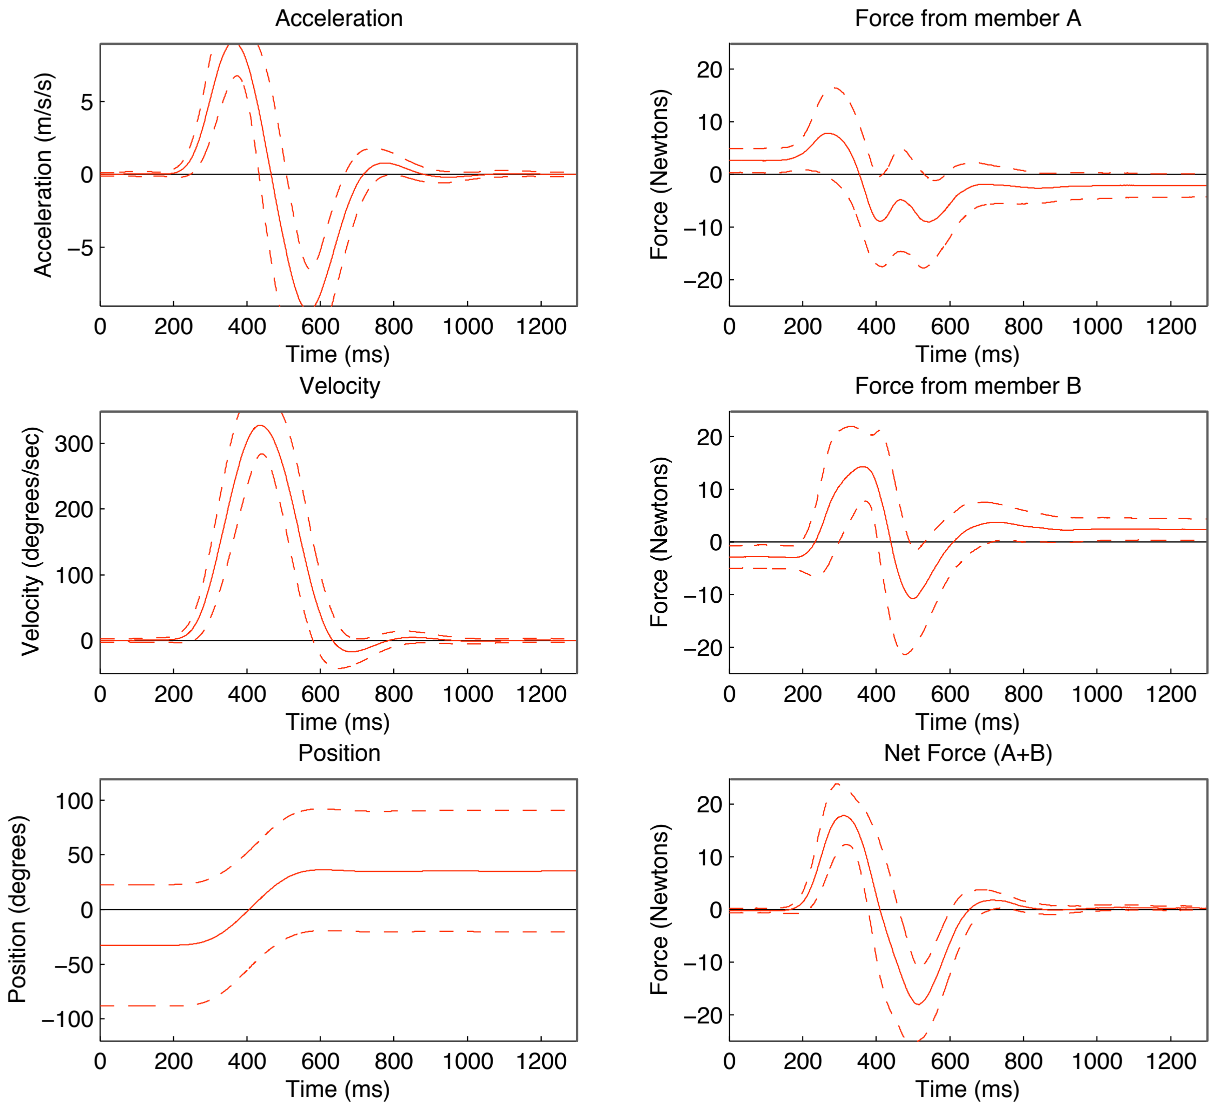

Figure B.5. Average Completion time: 587 ms

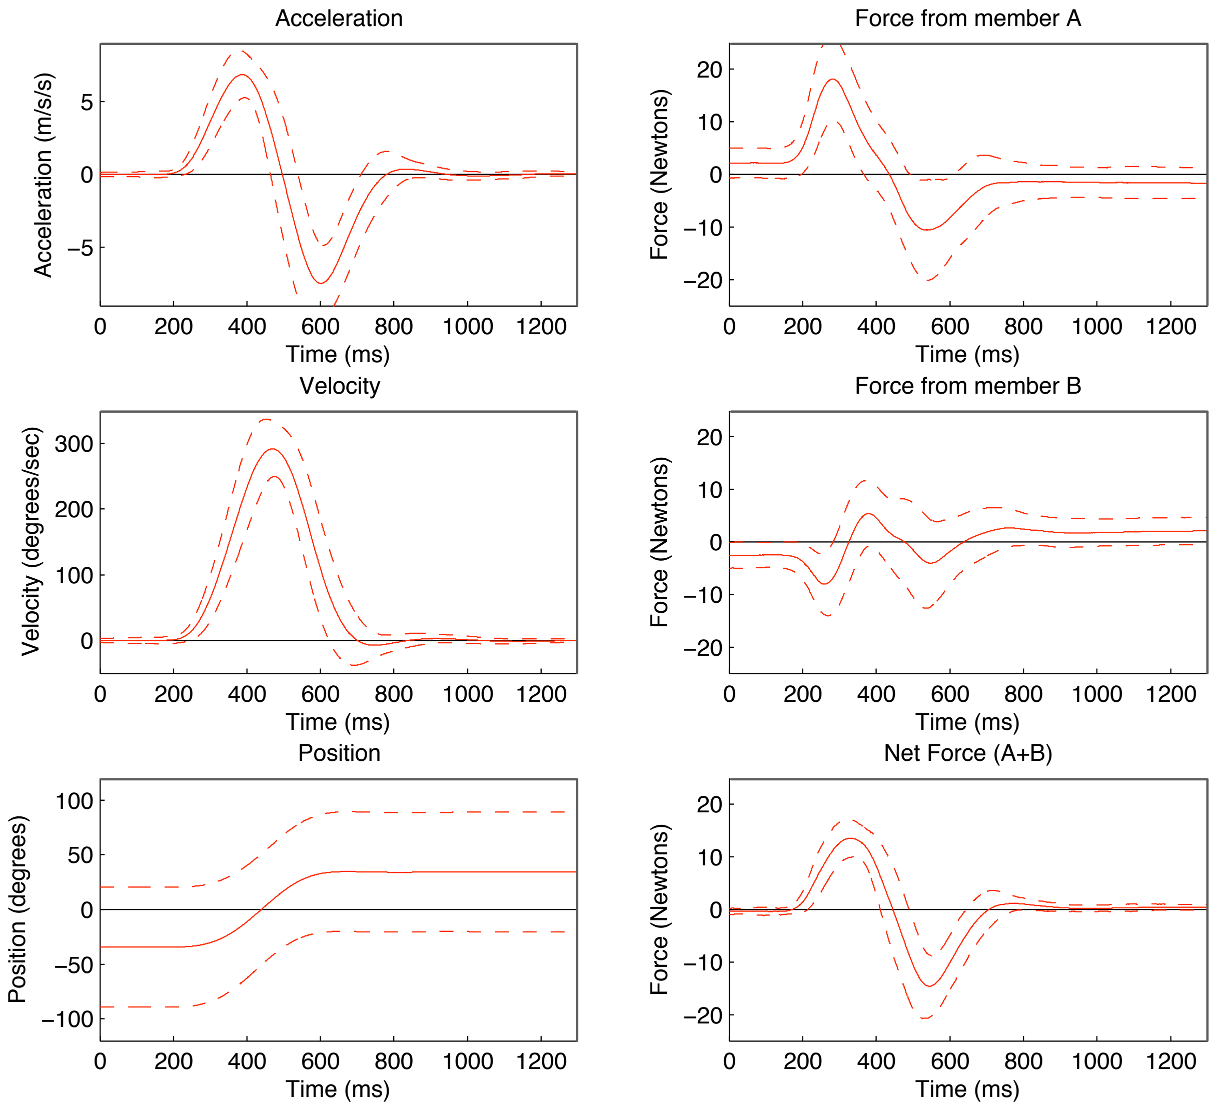

Figure B.6. Average Completion time: 608 ms

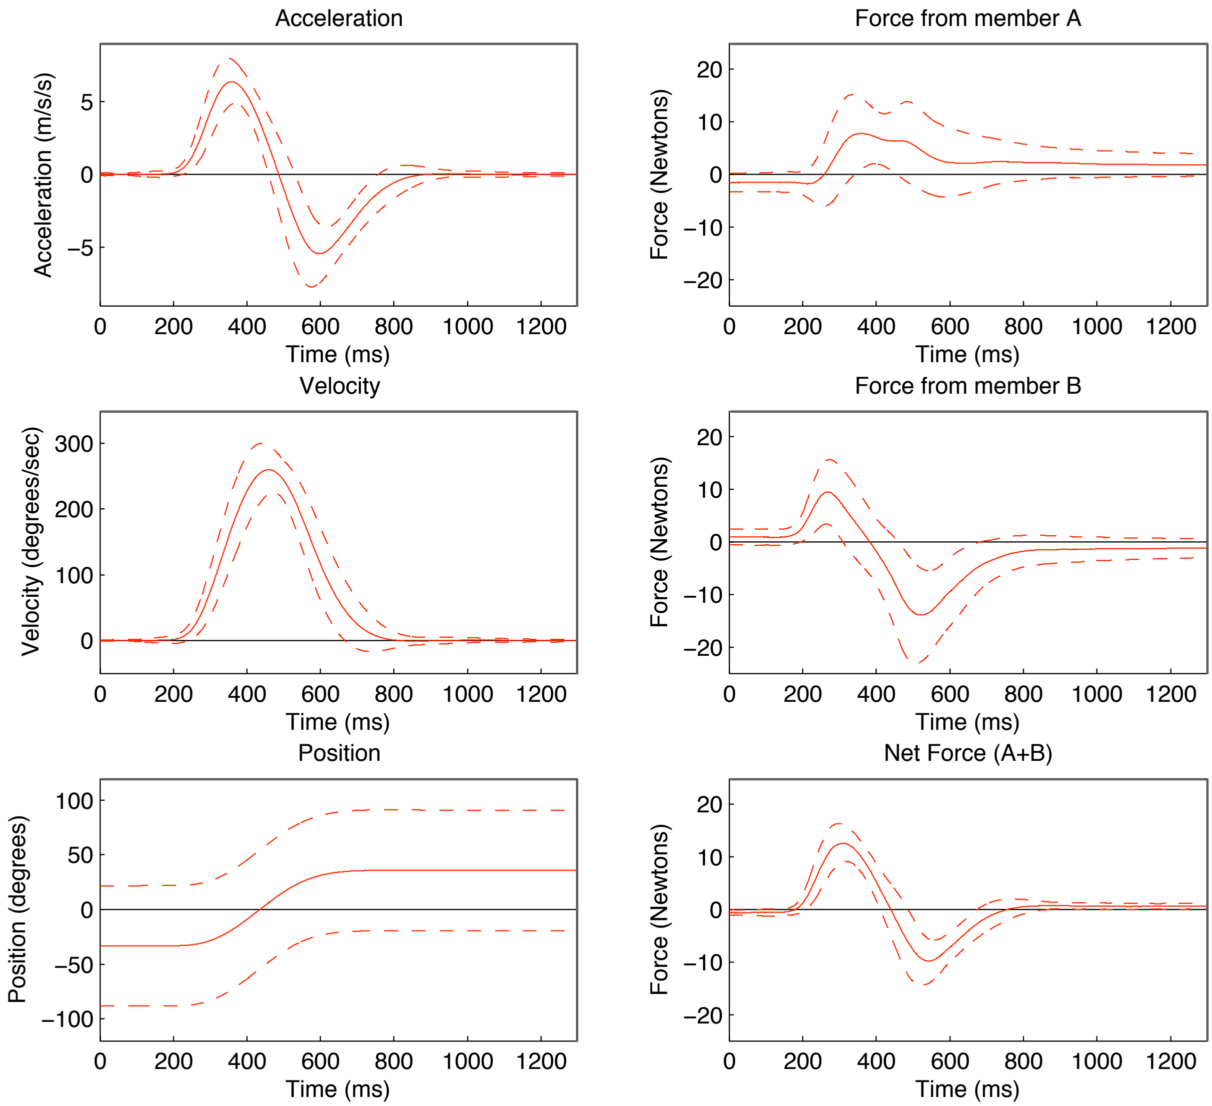

Figure B.7. Average Completion time: 622 ms

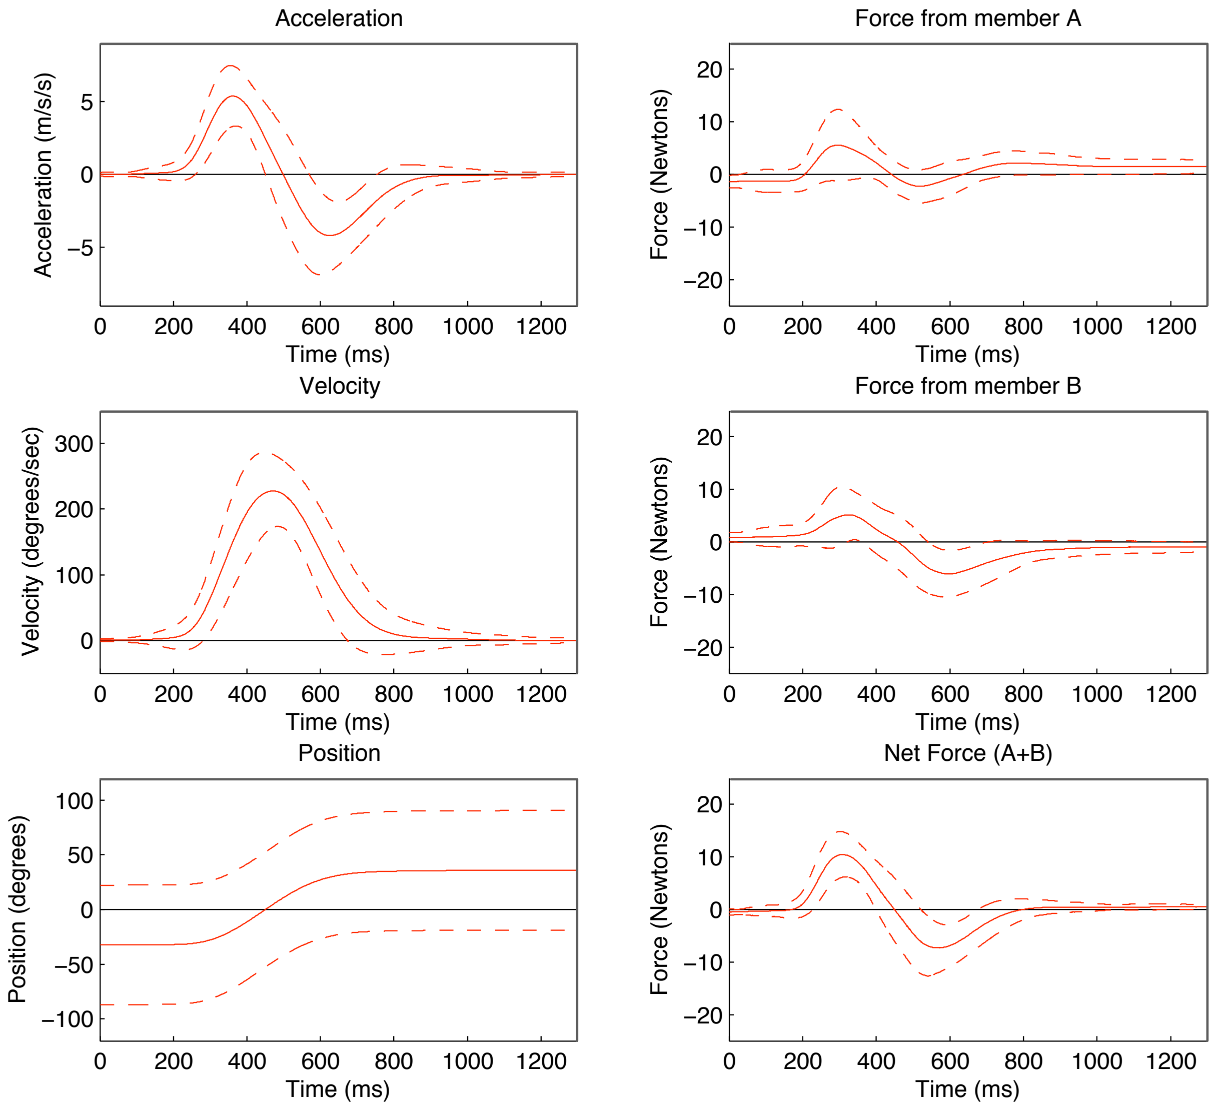

Figure B.8. Average Completion time: 642 ms

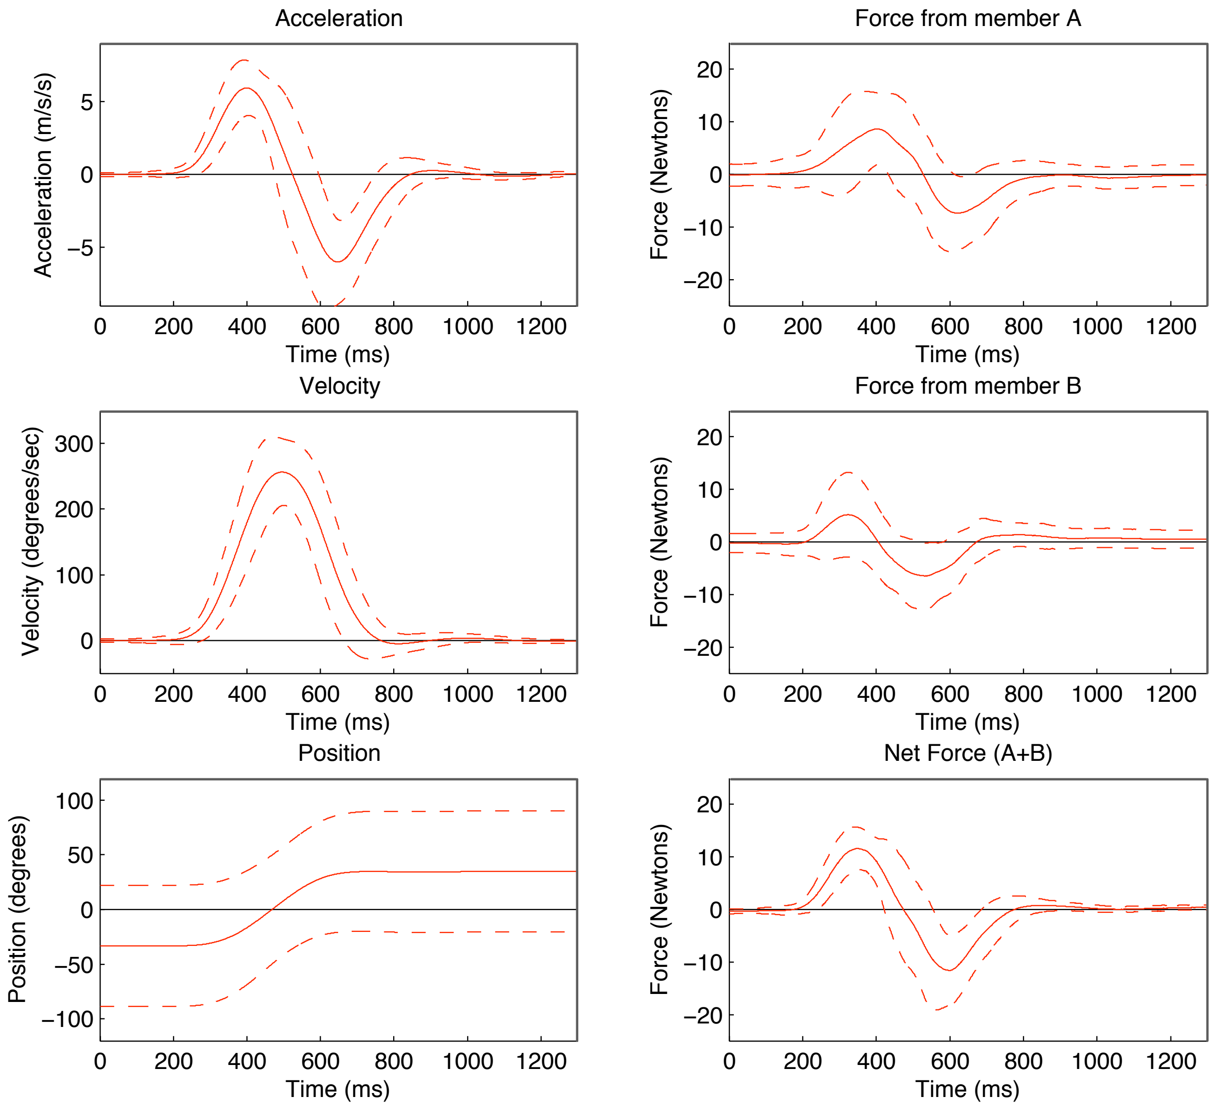

Figure B.9. Average Completion time: 668 ms

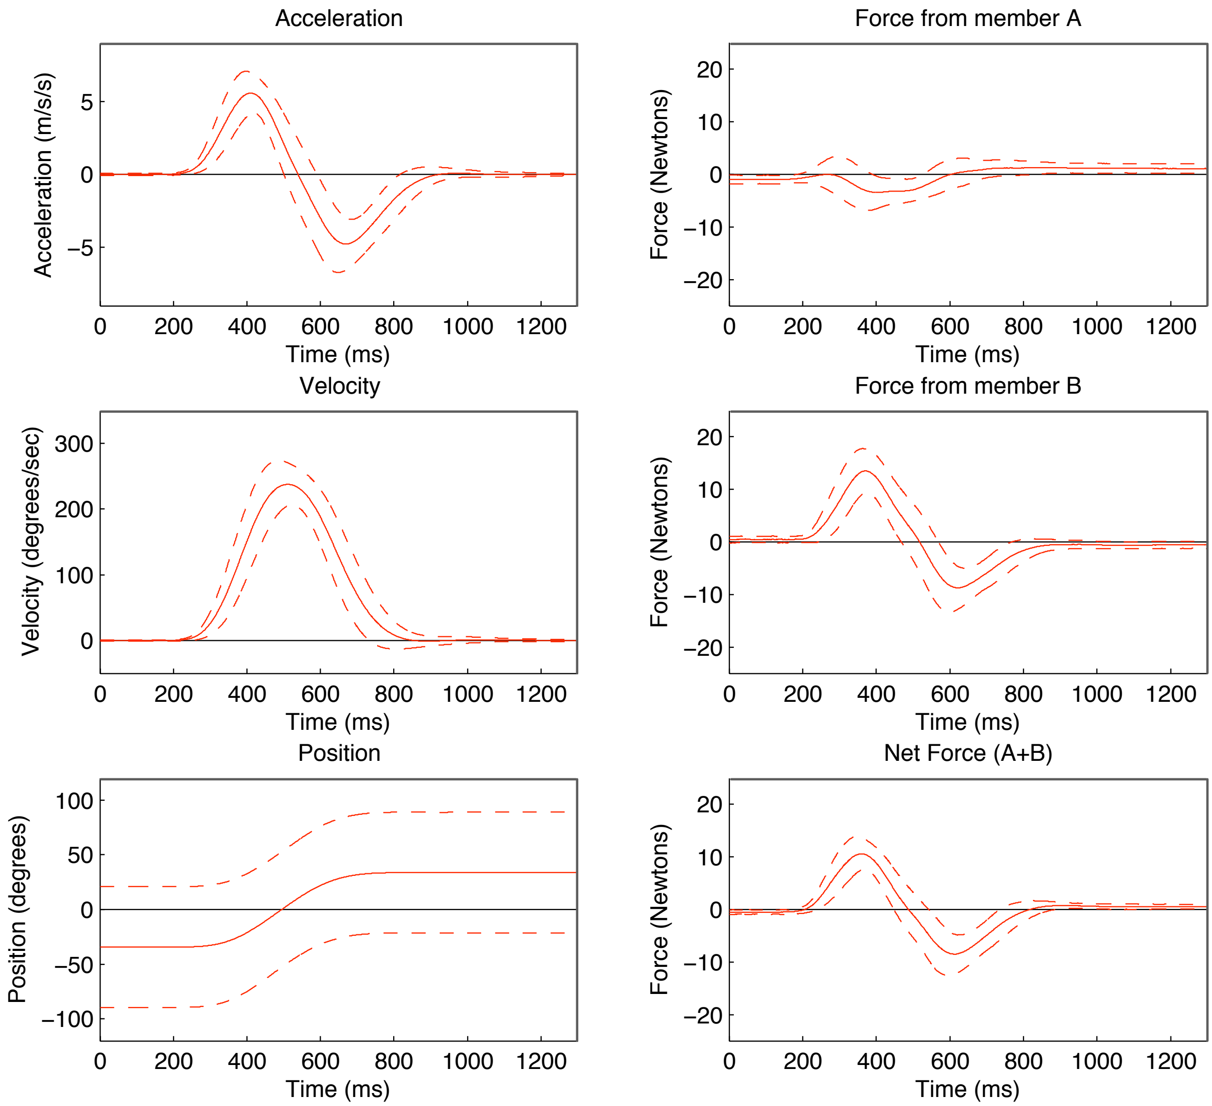

Figure B.10. Average Completion time: 678 ms

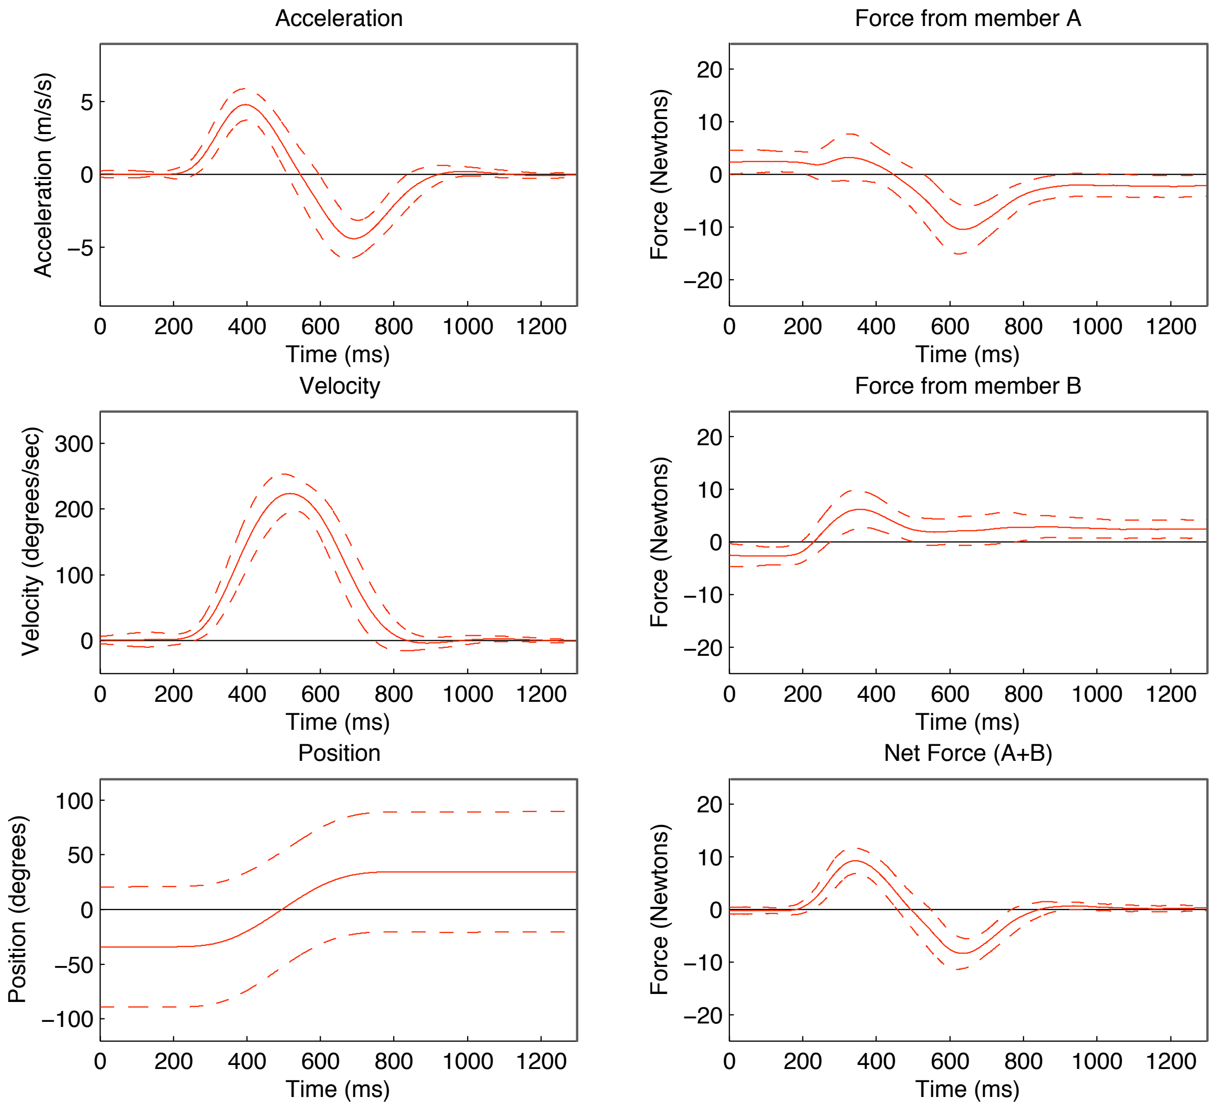

Figure B.11. Average Completion time: 681 ms

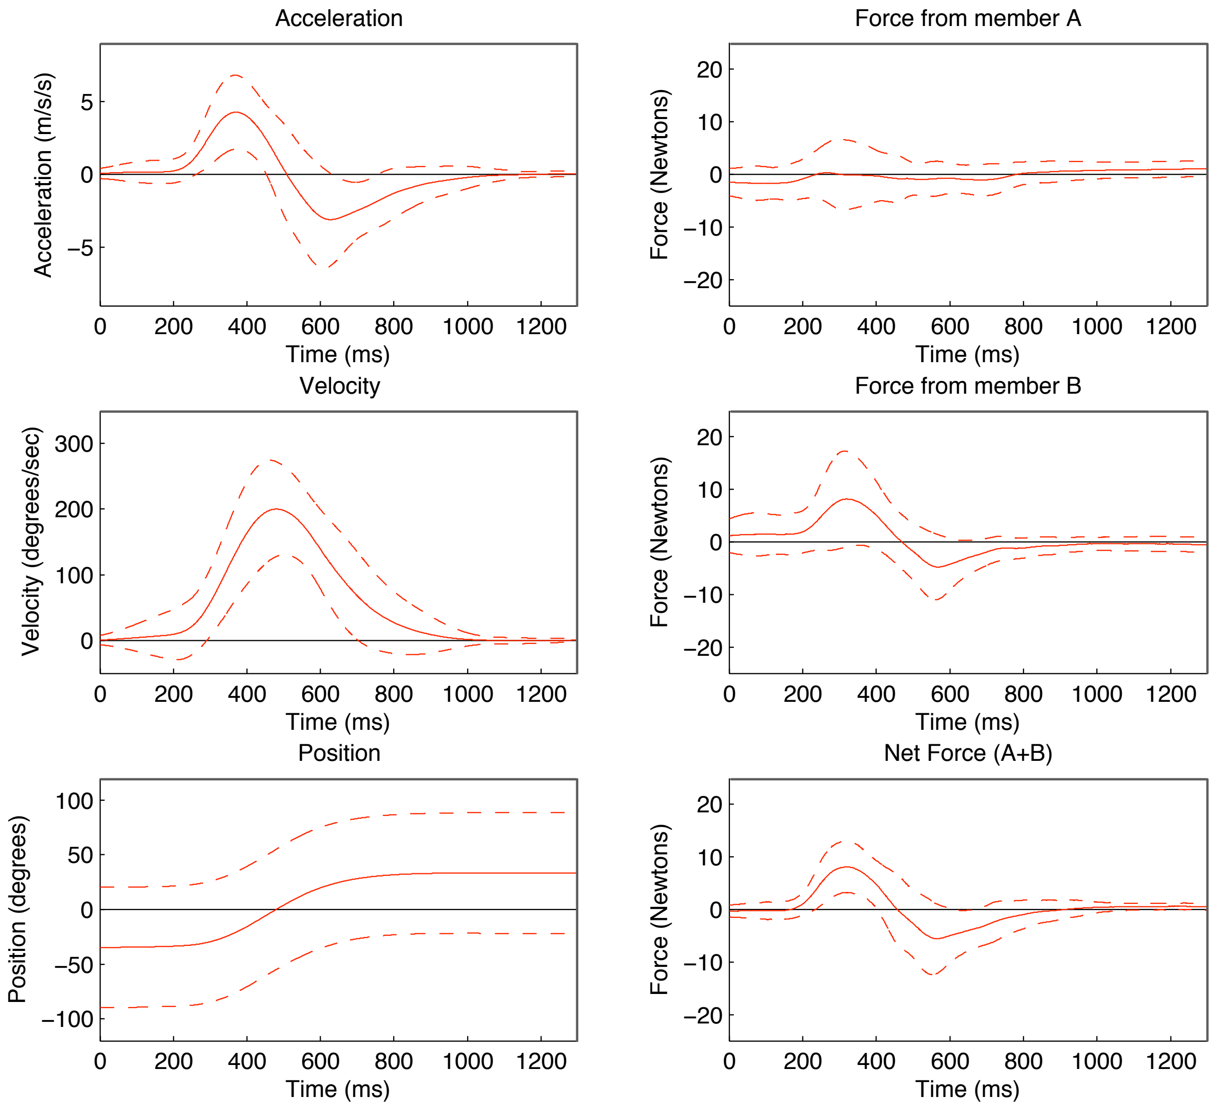

Figure B.12. Average Completion time: 693 ms

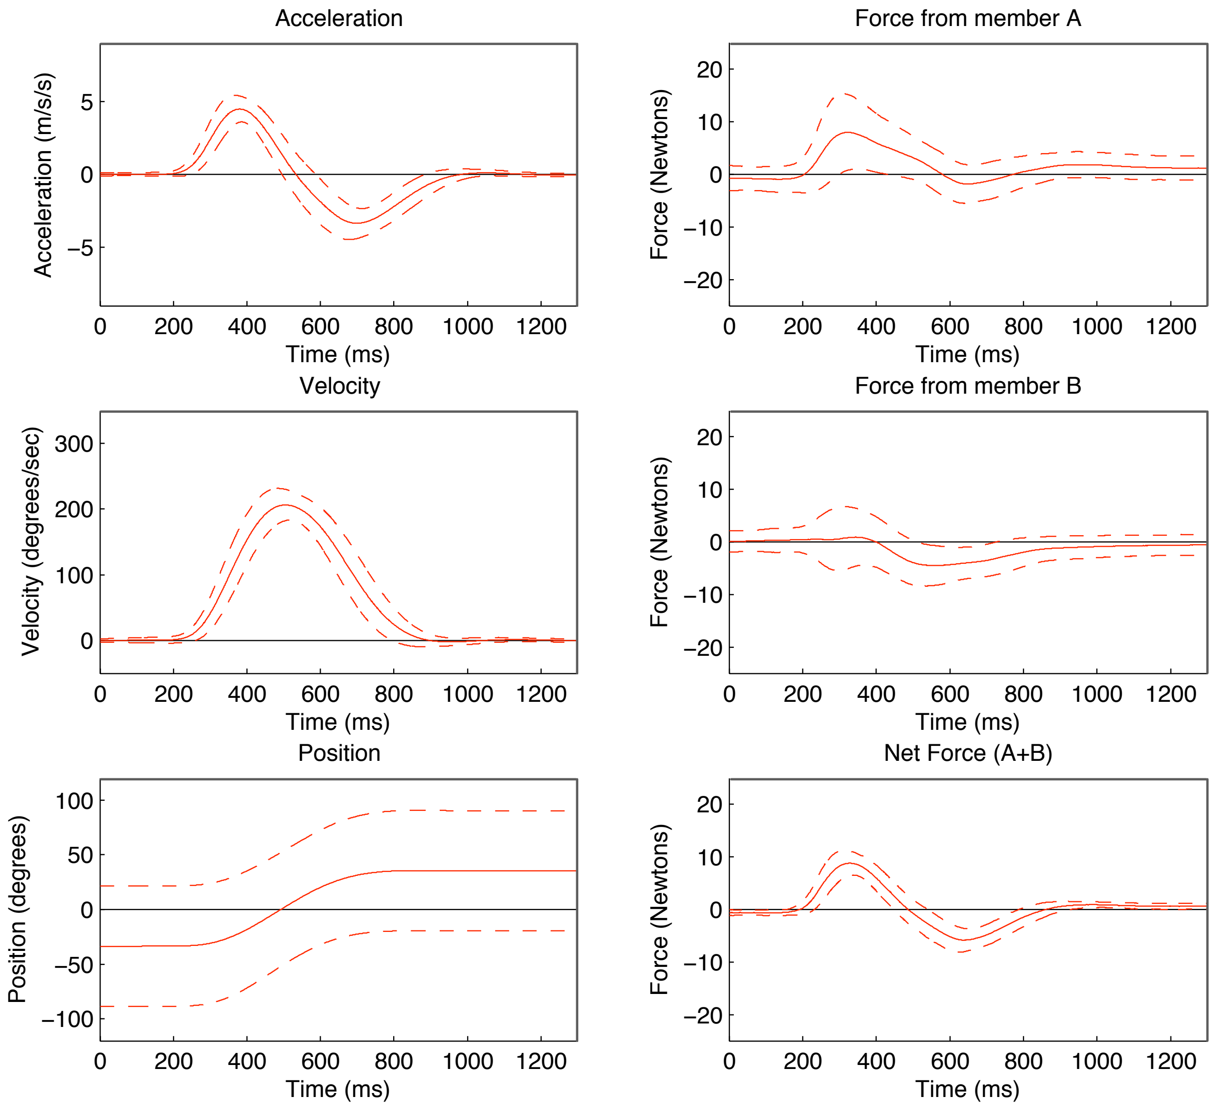

Figure B.13. Average Completion time: 704 ms

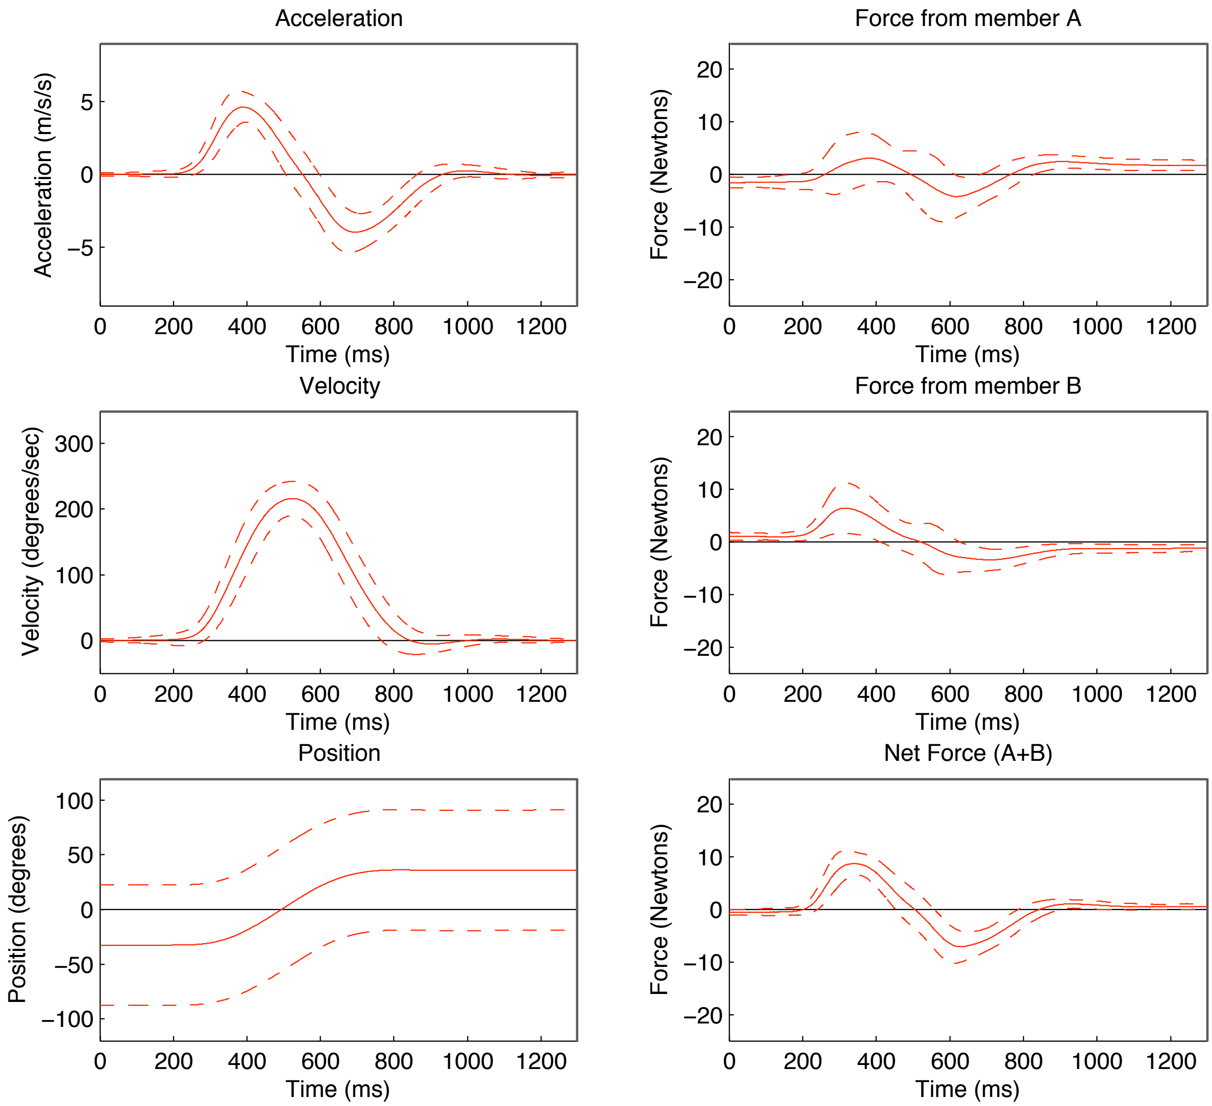

Figure B.14. Average Completion time: 718 ms

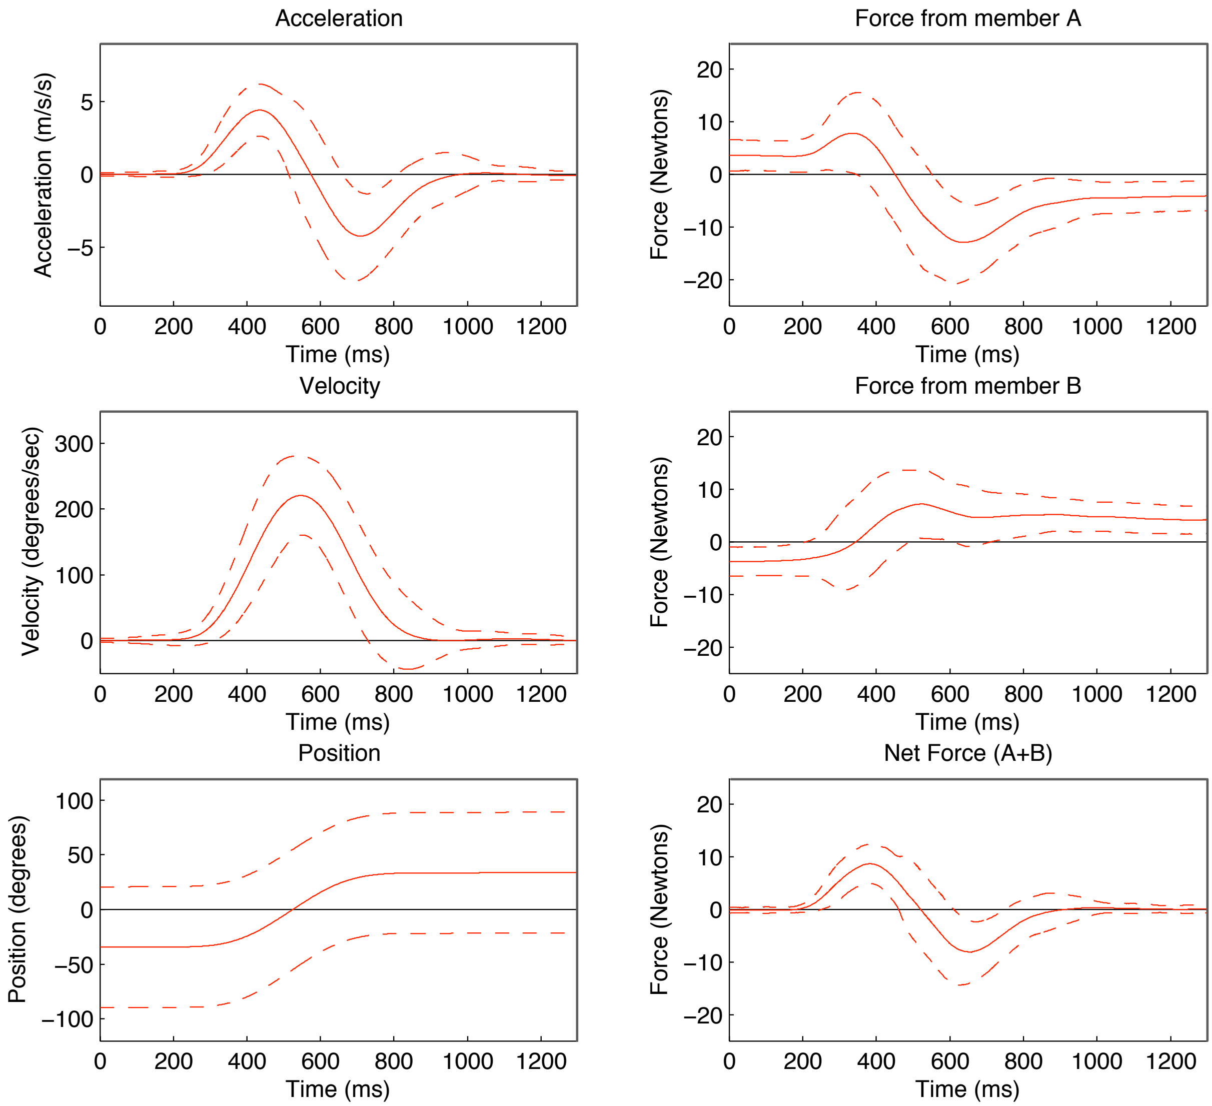

Figure B.15. Average Completion time: 767 ms