2017-04-01

Estimating Short-Term Human Intent for Physical Human-Robot Co-Manipulation

Eric Christopher Townsend
Brigham Young University

Estimating Short-Term Human Intent for Physical Human-Robot Co-Manipulation

Eric Christopher Townsend

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Marc Killpack, Chair
Tim McLain
Michael Goodrich

Department of Mechanical Engineering

Brigham Young University

ABSTRACT

Estimating Short-Term Human Intent for Physical Human-Robot Co-Manipulation

Eric Christopher Townsend
Department of Mechanical Engineering, BYU
Master of Science

Robots are increasingly becoming safer and more capable. In the past, the main applications for robots have been in manufacturing, where they perform repetitive, highly accurate tasks with physical barriers that separate them from people. They have also been used in space exploration where people are not around. Due to improvements in sensors, algorithms, and design, robots are beginning to be used in other applications like materials handling, healthcare, and agriculture and will one day be ubiquitous. For this to be possible, they will need to be able to function safely in unmodelled and dynamic environments. This is especially true when working in a shared space with people.

We desire for robots to interact with people in a way that is helpful and intuitive. This requires that the robots both act predictably and be able to predict short-term human intent. We create a model for predicting short-term human intent in a collaborative furniture carrying task that a robot could use to be a more responsive and intuitive teammate.

For robots to perform collaborative manipulation tasks with people naturally and efficiently, understanding and predicting human intent is necessary. We completed an exploratory study recording motion and force for 21 human dyads moving an object in tandem in a variety of tasks to better understand how they move and how their movement can be predicted. Using the previous 0.75 seconds of data, the human intent can be predicted for the next 0.25 seconds. This can then be used with a robot in real applications. We also show that force data is not required to predict human intent. We show how the prediction data works in real-time, demonstrating that past motion alone can be used to predict short-term human intent. We show this with human-human dyads and a human-robot dyad.

Finally, we imagine that soft robots will be common in human-robot interaction. We present work on controlling soft, pneumatically-actuated, inflatable robots. These soft robots have less inertia than traditional robots but a high power density which allows them to operate in proximity to people. They can, however, be difficult to control. We developed a neural net model to use for control of our soft robot.

We have shown that we can predict human intent in a human-robot dyad which is an important goal in physical human-robot interaction and will allow robots to co-manipulate objects with humans in an intelligent way.

Keywords: Physical Human Robot Interaction, Machine Learning, Short-Term Human Intent Estimation

ACKNOWLEDGMENTS

I would like to thank my family, especially my parents, for all the help and support they gave me. None of this would have been possible without them.

I am very grateful for all the time, technical knowledge, advice, and humor that Dr. Marc Killpack has provided to me and the lab. He was so helpful and because of him I had this opportunity. I also owe much of the success of my research to him as he was constantly available and helping.

I am grateful to all the members of the lab for their technical help, collaboration, and for being friends. We did a lot of work together and had a good time. It would not have been possible to complete the human study without the help of Erich Mielke and Nick Walton.

Dr. Michael Goodrich gave advice many times on the experimental study and the data. Dr. David Wingate's expertise on neural networks made this work successful. He helped me improve the performance of the neural net and troubleshoot my code. Dr. Steven Charles gave detailed and excellent advice on running a study with human subjects.

I would like to acknowledge that this work was funded by the Army Research Laboratory Robotics Collaborative Technology Alliance.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1. INTRODUCTION

While robots have long been used in manufacturing, they are increasingly gaining the capability to work in unstructured and dynamic environments. In the future, this will include applications related to logistics, health care, agriculture, disaster response, and others. However, robots will also be required to successfully interact more naturally with human teammates. Specifically, in this thesis, we aim to create a model for predicting short-term human intent in a collaborative object carrying task that a robot could use to be a more responsive and intuitive teammate.

In this chapter, we outline the contributions of this thesis. Next, we discuss physical human-robot interaction in which we are interested. We then describe the purpose and form of neural networks that are used extensively in this thesis to model the intent of humans and the dynamics of soft robots. Our robot platform used in this research for predicting short-term human intent is described next. Finally, we discuss previous related research that exists in the literature.

1.1 Contributions

The contributions of this thesis include the following:

- An exploratory study of humans co-manipulating an object with 21 human pairs, or dyads.
- Determination that past motion of a human dyad is all that is necessary to predict short-term human intent for at least a limited time horizon, where we define *intent* as how the team will move in the next 0.25 seconds.
- Development of a neural network to predict short-term human intent based on past motion of the human dyad.
- Validation of the short-term human intent prediction in a human-robot dyad.

- Control of soft robots that are difficult to model and could eventually be used in human-robot interaction.

1.2 Physical Human Robot Interaction

In this thesis, we define *co-manipulation* as physical human-robot interaction for the purpose of collaboratively manipulating an object. As part of the work presented in this thesis, we ran an exploratory study with 42 humans (21 human dyads) in order to understand human-human co-manipulation. Each dyad moved a long board representing a table as we measured their movement and forces on the board as in Figure 1.1. Many previous studies have been done on human movement in which one or two people move in tandem. Many of these are done in haptic simulations or with limited degrees of freedom in order to isolate specific behaviors. These studies have given significant insight on things like minimum-jerk motion, negotiation of roles, and task-specific movements. We expect, however, that due to the nature of this past work which has mostly examined a limited number of degrees of freedom, there are limitations to how those results can be interpreted for general purpose six dimensional co-manipulation tasks. We therefore assert it is also necessary to perform studies with natural and realistic human movement without limiting degrees of freedom. This should validate some of what has been learned in other studies and allow further insight and direction for human-robot co-manipulation controller development.

There are different definitions of human intent that have been used by researchers in the past. These include many examples, some of which are the way a human is going to turn a car or when they want to switch lanes in order to create assistive vehicles or to warn against collisions [1–4]. Another example is predicting limb movement for prosthetic applications [5]. A common example is predicting where a human will move so that a robot can avoid it. Examples of this include predicting human arm trajectories in a collaborative assembly task [6], predicting the final target of the human arm in reaching tasks in real-time, [7], and predicting the space a human will occupy based on previously learned human motions for a specific task [10]. All of these allow a robot to work in the space that the human is not predicted to occupy. Some final examples include using hand gestures to understand what a human is doing or going to do, or wants the robot to do [8], using head and feet pose to predict when a human will turn [9], and in our case, predicting

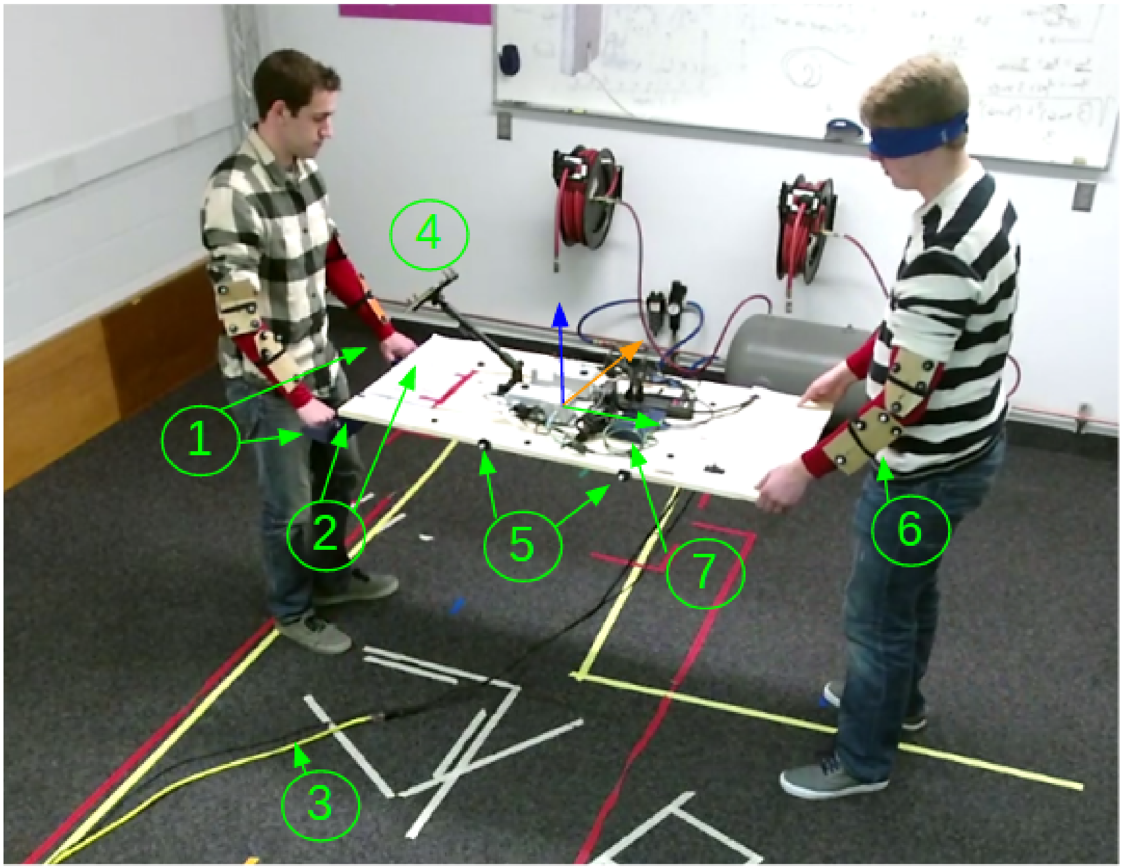

Figure 1.1: A leader and a blindfolded follower performing a table carrying task. The numbers describe the following: 1. Handles 2. Force/torque sensors 3. Power and ethernet cables 4. IPad wih instructions 5. Motion capture IR reflective markers 6. Motion capture sleeves 7. Router, force sensor boxes, and power strip. Green is the longitudinal axis, blue is the vertical axis, and orange is the horizontal axis.

the desired trajectory of a co-manipulated object. The purpose of all of these is so that robots can assist humans or work in the same space as them by predicting their movement.

We specifically focus on predicting intent over a short time horizon on the order of a quarter second. This is about the time that it takes a human to react to stimuli. We hypothesize that a prediction over a short time horizon, while not useful for path planning to a final goal or determining the long term goal of a human partner, is useful for a comfortable and intuitive physical co-manipulation interaction between a human and robot.

There are many sources of information that a robot could use to predict human intent, including motion, force, posture, external context, and verbal communication among others. In our study, we chose to focus on motion and force. We hypothesized that these variables were the most fundamental and easiest to interpret for a robot in order to predict what the person intends to do over a short time horizon. In another paper [11], we discuss insights gained from the force information in this study and how force data can predict initiation of future motion, while in this thesis we use past motion data only to predict where the co-manipulated object will move next over a short time horizon. The interaction between force and motion in a co-manipulation task is complicated but important. In this preliminary work, we have attempted to isolate and explore each physical quantity before we later expect to integrate insights and models from both motion and force for future robot controller development. In large part, this was our approach because our initial prediction of short-term human intent did not perform as well when force inputs were included.

1.3 Neural Networks

Neural networks are a method used in machine learning to approximate any function based on empirical data. Inspired by the function of the brain, they are trained on data and learn to classify or predict. They are especially useful when it is difficult to pick the form of an analytical and possibly nonlinear function to classify or predict the data. Modern neural nets can classify objects in images, recognize speech, or predict economic trends.

A neural net consists of a large number of nodes that mimic the neurons in the brain. These are typically arranged in layers, each containing several nodes as seen in Figure 1.2 which is one of the neural nets that we use in this thesis. The first layer is the input layer and the last layer is the output layer. The rest are called hidden layers. Each layer receives inputs from the nodes in the layer before it in addition to a bias node. Each input is multiplied by a weight and then they are summed. This value is put through an activation function which generates outputs for the next layer. For example, the output of the node in Figure 1.3 would be:

$$Output_1 = activationfunction(\sum_{l=1}^{inputs}(input_l * weight_{l,1}) + bias * weight_{0,1}) \tag{1.1}$$

Modern neural nets can contain millions of nodes. There are many different activation functions but typically they are nonlinear which is required in order for the whole neural net to be

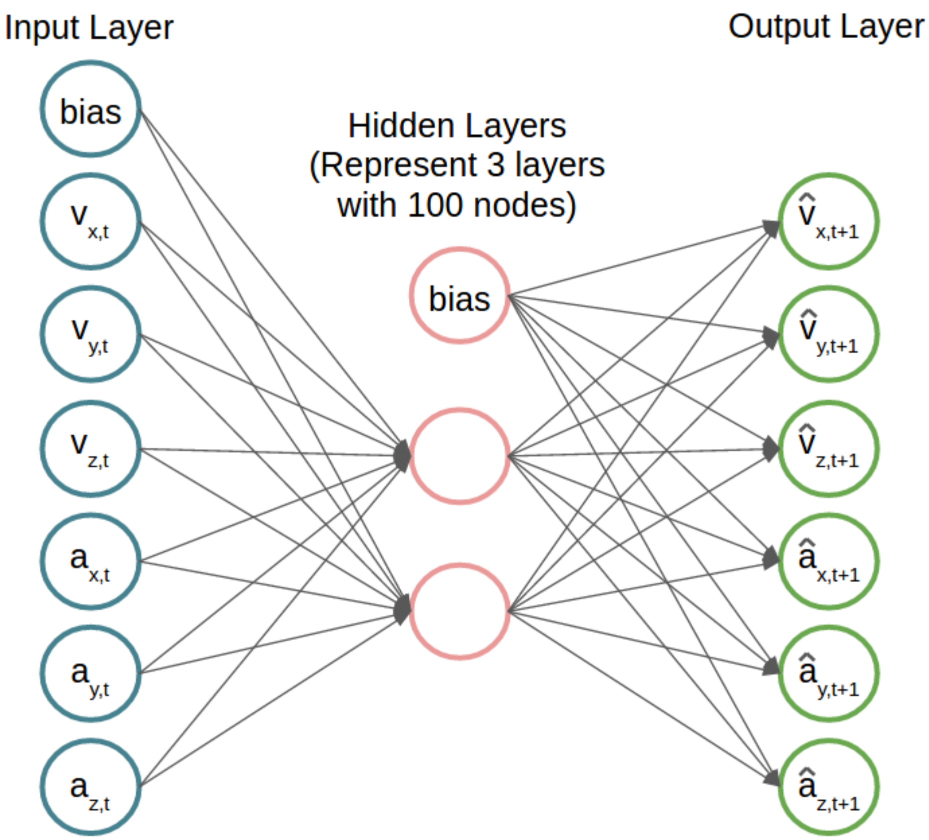

Figure 1.2: Symbolic representation of a neural network. The hat represents a prediction, and $t+1$ indicates that it is one time step into the future.

able to learn a nonlinear function. The first layer takes in the raw inputs and the last layer outputs the prediction or classification.

Neural nets are typically trained using a method called back-propagation. This method inputs a set of recorded data with known inputs and outputs into the neural net and then changes the weights at each node using gradients. There is a cost function, such as the mean squared error between the measured outputs and the predicted outputs, that measures how accurately the neural net predicts or classifies the measured outputs. Changing the weights through back-propagation is continued iteratively until the cost function stops converging or the neural net begins to overfit the data.

Overfitting is prevented by splitting the entire data set into validation and training sets. The neural net is only optimized using the training data set, but we also find the error when validation data set is input to the neural net. The neural net begins to overfit the data when the cost of the

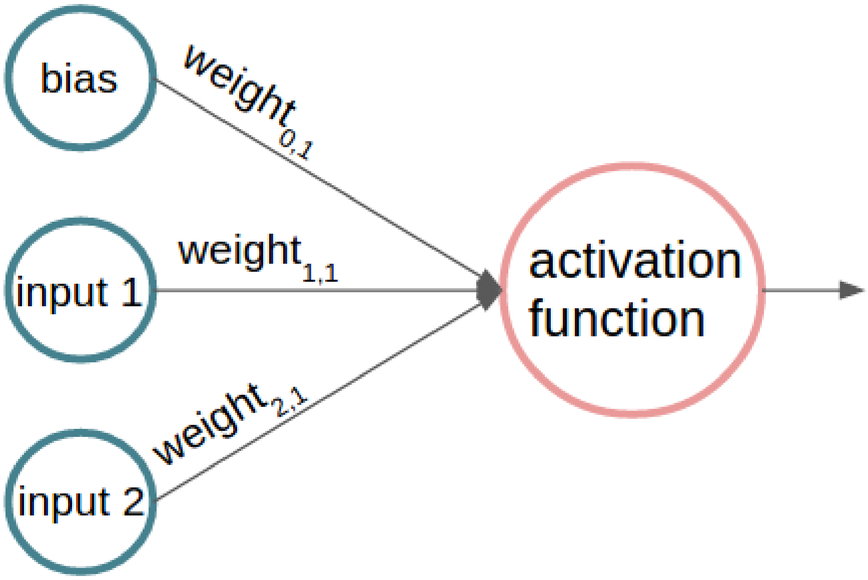

Figure 1.3: Symbolic representation of one node in a neural network.

training data is lower than the cost of the validation data. This becomes a problem when the cost of the validation data increases while the cost of the training data is still decreasing. This is when training is typically stopped.

Neural networks are mostly known for their use in classification, but they can also be used for time series prediction, where past data is used to predict future data, which is how they are used in this thesis. There are many statistical and machine learning methods for time series predictions. Neural nets were the logical first approach given their ability to approximate any function, their simplicity, and and the fact that we get the analytical gradient with no additional needed computation from the neural net due to the training through back-propagation.

We use neural networks in this thesis in order to model the intent of two humans performing a task in collaboration. We also use them to model the dynamic response of a nonlinear, soft, inflatable robots.

1.4 Robot Platform and Controller

Our robot platform for this research was a Rethink Robotics Baxter robot mounted on an AMP-1 holonomic base from HStar Technologies as seen in Figure 1.4. The Baxter robot is a rigid

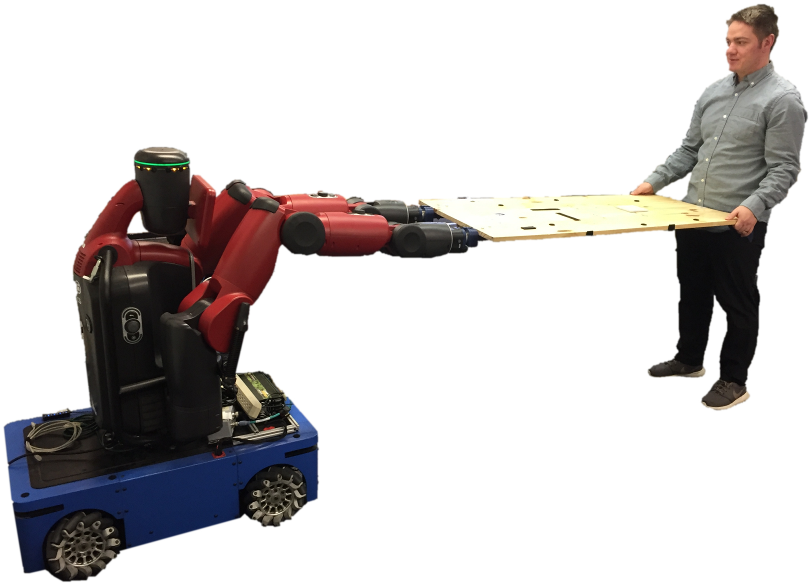

Figure 1.4: Rethink Robotics Baxter robot mounted on HStar Technologies AMP-1 holonomic base carrying the table with a person.

robot that has Series Elastic Actuators [12] in each joint. These are springs that give the robot the ability to sense torque and a natural compliance which makes them somewhat safer for use with humans.

We chose to use a holonomic base with mecanum wheels instead of something like a bipedal robot in order to validate the short-term human intent prediction at speeds similar to two humans moving an object in every day life. This was important to ensure that our estimation and control algorithms for co-manipulation work in real world applications as limiting speed, as is often necessary with bipedal robots, may affect the dynamics of the interaction.

1.5 Related Work

Research has been done in human-robot interaction for physical co-manipulation for several decades because there are many advantages to humans and robots working together. In general, the advantage of human-robot collaboration is that humans provide intelligence and dexterity while robots may provide strength and stability [13]. Robots working with a human also make it possible for less humans to be in dangerous situations. Here we discuss some of the most important developments in human-robot interaction for physical co-manipulation.

1.5.1 Impedance Control

Many different control methods have been used in the last 20 years for human-robot cooperation. Impedance control was first described by Hogan in 1984 as a way for a robot to interact with its environment [14]. One of the first controllers for cooperative manipulation of an object by robots and humans was an impedance controller that could be used with any number of robots and humans holding the object [15]. This became the basis for much of the research that followed.

1.5.2 Human-Human Interaction

To improve human-robot interaction, human-human interaction has also been studied to understand how people communicate haptically. Ikeura et al. showed that when two people move an object, but only one knows the task, the applied force is highly correlated to the velocity and therefore can be modeled by a damping element in an impedance model. Of importance was the result that the spring and mass terms were much less important than the damping terms [16]. Ikeura later showed that a constant damping term does not work well for changing velocities as it slows the human-robot pair significantly compared to the human-human pair. Variable damping that depends on velocity was proposed to allow both fast motion and accurate positioning when using an impedance model [17, 18]. Tsumugiwa et al. continued on this idea by changing the impedance according to the estimated stiffness of the human arm [19]. Duchaine et al. continued this idea by changing the impedance according to the derivative of the force applied by the human [20]. All of these tasks were done in one or two dimensions with no rotation. Other researchers showed how humans adapt to unknown or unstable dynamics and variable impedance controllers have been made based on this principle as well [21–23].

Damping controllers tend to have high internal forces. One of the proposed measures used to determine the effectiveness of a human-robot interaction controller is the interaction force between the human and the robot. As this force is usually counted as wasted energy, the hypothesis was that humans naturally try to minimize it. According to Ito et al., a controller for human-robot cooperation should minimize this interaction force [24]. Motion is also not tracked as accurately because the leader essentially drags the robot into place instead of being assisted by the robot. Variable damping is an attempt to solve this problem, but the user will always feel like the robot is being dragged instead of being an assistant. Because of this, a damping controller is not clearly superior to other methods.

1.5.3 Invariant Models

One way for robots to better assist people is through the development of invariant models of different tasks. By finding a parametrized model for a task, the dynamics can be simplified, but also the robot can predict what a human will do. The most well known example is the minimum-jerk model. Flash and Hogan showed that the human arm tends to follow a minimum-jerk trajectory in many situations. Jerk is the time derivative of acceleration. Minimum-jerk means that humans try to move in a smooth manner. This is especially accurate when force is low and high speed is not an objective [25]. Corteville et al. used the minimum-jerk criterion to train a robot to predict what a human will do in a one-dimensional task [26]. Ikeura used the minimum-jerk model to calculate a damping constant [27] while Maeda et al. used the minimum-jerk model directly to predict what the human is trying to do. Their controller estimates the final time and position that the human is attempting to reach based on the past few steps and uses position control to follow that trajectory along with impedance control to adapt to errors [28].

Kheddar et al. have found that minimum-jerk does not apply to all cooperative movements. One example is when motions are longer and include walking and not just arm movement. In those cases there are phases of constant velocity along with phases of constant acceleration. They used this knowledge to create a controller based on phases of constant velocity that are only changed when a certain force threshold is surpassed [29, 30].

Many models for a controller for cooperative manipulation only have a few changing parameters that include initial and final time and position. They may also include a maximum velocity or other parameter. With just a few parameters, the entire trajectory can be described. One approach to determining these models has been programming by demonstration [31, 32]. These models allow a robot to quickly learn a new task, but they have to be trained for each new task.

1.5.4 Methods that do not Use Force and Motion Exclusively

All the previously mentioned methods only require that the robot have force and position data. Agravante et al. showed that by using cameras for robotic vision, some things can be done that would be difficult otherwise, like keeping a co-manipulated table perfectly flat or keeping a ball from rolling off that table [33]. Choi et al. use EMG signals to predict the movement and force of a monkey's arm which has obvious implications for human-robot interaction although it is overly invasive for many of the tasks that robots would do with humans. [34].

1.5.5 Robot behavior

Past work by Dragan et al. has hypothesized that it is better for a robot to be legible than predictable, or in other words, it is better to move in a way that clearly shows its intent rather than to do what the human thinks it will do. The human should not be surprised by what the robot does. Aside from being predictable, if the robot takes any leadership in the task, such as keeping an object from rotating, it should do so in a way that is understandable to the human [35]. Chipalkatty et al. showed that there is an advantage to making a controller simple enough that a human can learn it quickly. If it is complex and constantly adapting, there is a risk the human will become confused and the cooperation will be inefficient [36].

The work that has been described so far is representative of the current research in the field of physical human-robot interaction for co-manipulation. We propose that most good human-robot interaction controllers for co-manipulation of an object can be judged on certain factors including low interaction forces, quick and accurate motion tracking, safety, simplicity, robot initiative, and predictability. However, little of the past work has been based on real performance data from human-human co-manipulation trials for unconstrained tasks. The purpose of this thesis is to use real human-human data and attempt to develop an estimator that allows us to predict short-term human motion intent for a given set of tasks.

In chapter 2, we discuss the exploratory study that we ran. Chapter 3 discusses the neural network development and accuracy for predicting short-term human intent. Chapter 4 discusses work on modeling the dynamics of soft robots that can be used in human-robot interaction. We conclude in Chapter 5. Appendix A has the documents that were used for exploratory study including consent forms, recruiting materials, and a questionnaire.

CHAPTER 2. EXPLORATORY HUMAN-HUMAN CO-MANIPULATION STUDY

Many studies have been done on human movement in which one or more people move or carry an object together in real life or simulation [15, 37–40]. Many of these are done in haptic simulations or with limited degrees of freedom to isolate specific behaviors. These studies have given significant insight on things like minimum-jerk motion, negotiation of roles, and task-specific movements. All of these insights build a picture of how humans interact in real-world tasks with limited degrees of freedom. We assert it is also necessary to perform studies with natural and realistic human movement without limiting degrees of freedom to validate what has been learned in other studies and to gain more insight. We had 21 pairs, which we refer to as dyads, participate in our study. In this chapter, we describe the specifics of our study. This study was approved by the Brigham Young University Institutional Review Board and all documents used for recruiting and during the study are in Appendix A.

2.1 Experiment Goals

Our purpose in doing this study was to understand how people collaborate while moving a rigid object and without restricting human motion while manipulating the object. We desired to verify that the principles other people have learned in other studies, described in Section 1.5, Related Work, generalize to unbounded human motion and see what new principles we could learn. Our goal was to learn all we could about human-human co-manipulation so that we could program a robot to effectively work as a teammate with a person in object co-manipulation tasks.

We chose a series of tasks that would allow us to understand several different types of common motion. Many of the tasks were designed to isolate certain behaviors, while others were open-ended and incorporated many expected behaviors into one task.

The tasks consisted of two people moving a rigid object on flat terrain as shown in Figure 2.1 from chapter 1, shown here for convenience. This is a task that could be useful in many real-world applications such as rubble removal, materials transport in a warehouse, carrying a stretcher, or moving furniture.

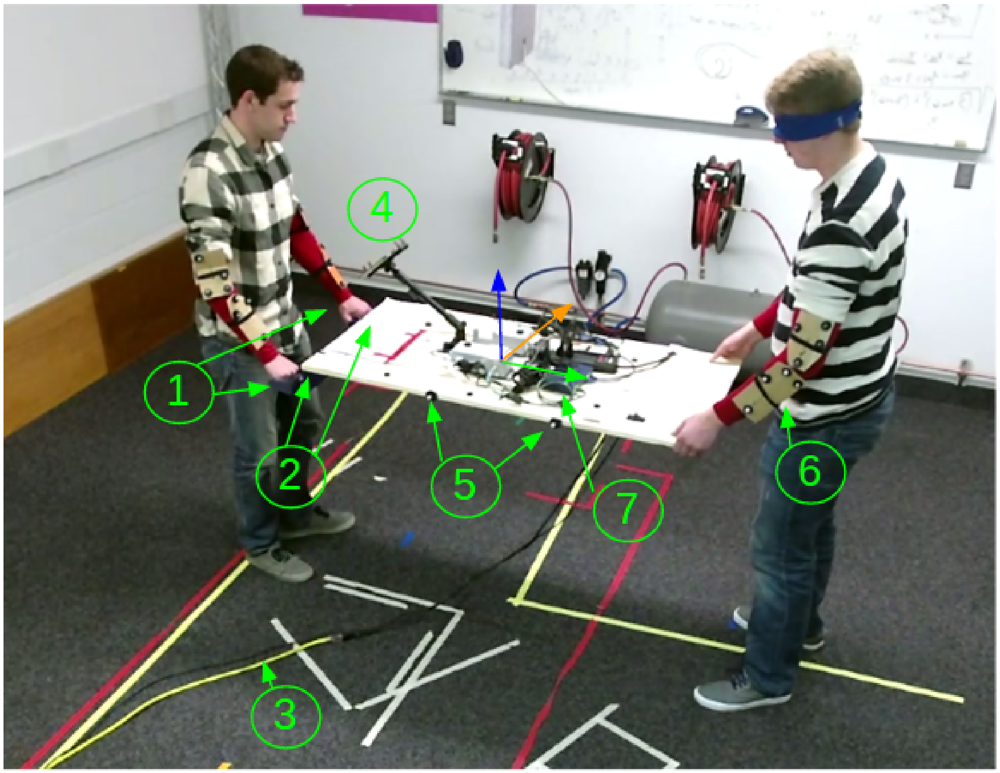

Figure 2.1: A dyad performing a task along with a description of the setup. The leader is on the left holding the handles. The follower is on the right with a blindfold which they wore for half of the tasks. The numbers describe the following: 1. Handles 2. Force/torque sensors 3. Power and ethernet cables 4. IPad wih instructions 5. Motion capture IR reflective markers 6. Motion capture sleeves 7. Router, force sensor boxes, and power strip. Green is the longitudinal axis, blue is the vertical axis, and orange is the horizontal axis.

2.2 Experimental Setup

In describing the experimental procedure, we first describe the room in which the tasks were completed. Then we describe how the table that the dyad carried was set up.

2.2.1 Room Setup

The tasks took place in a space 5 m by 4 m. The floor was set up with colored tape marks that indicated where the subjects were supposed to move for each task as shown on the floor in Figure 2.1.

The room was instrumented with a Motion Analysis motion capture system. The motion capture cameras were on a truss 2.4 meters off the ground that went around the length of the study space. There was also a Microsoft Kinect 2.0 that recorded sound, high-resolution video and the estimated full-body pose of the participants.

A curtain was installed that closed off the study space from the rest of the lab to limit distractions. This curtain restricted the participants' view of everything except two computers for data collection and the space where the subjects moved the table.

2.2.2 Table Setup

A rigid board was used as a table-like object to be transported by the dyad as seen in Figure 2.1. The board was set up to collect data for the task. The table had two handles on one side of the board. Each handle was connected to the board with an ATI Mini45 force/torque sensor in series. These were each connected to a network box that was attached to the table.

The table also carried an iPad that was oriented so that only the leader could see it. This iPad was controlled by the experimenters to show the leader the current task. It showed the leader where they were starting and where they would need to go in the task. The iPad had images with markings that corresponded to the markings on the floor. Each task is shown later in this chapter.

The table had motion capture markers so that it could be tracked. It also had a power strip that powered the force/torque sensors and the iPad. The power strip and Ethernet cables were each connected off the table. One of the researchers was tasked with ensuring that these did not get in the way of the participants and that there was sufficient slack so that it only applied very small forces to the table.

2.3 Experimental Procedure

The following is a description of our experimental procedure. Two subjects, who may or may not have signed up together, came in together. They each signed a video and photo release and informed consent forms as shown in Appendix A. They were then randomly assigned as a leader or follower for the duration of the study. The leader would be the only one to receive instructions for each task. The study was explained to them and the leader was shown how the task instructions would be given on the iPad.

The follower was given a blindfold that was worn for half the tasks as directed by the researchers. When the blindfold was on, the participants were not allowed to discuss the task. When, the blindfold was off, we did not limit their discussion. As many robots will not have vision or the ability to analyze vision like a human can, we performed half the tasks with one person blindfolded. This means the follower depends entirely on force and motion. We also did half the tasks without the blindfold to see what effect that has on motion. Each participant put on compression sleeves with motion capture dots so that their arm configuration could be tracked more accurately than with the Kinect data.

They were then run through one practice task to ensure that they understood everything. This task was shown to them on the iPad in order to ensure that they understood the diagrams. The practice task can be seen in Subsection 2.4.6. We chose a task that would include translation in two axes and rotation in three axes to show them the task complexity. If there were any misunderstandings or questions during the practice task, these were clarified.

The six tasks were randomly ordered and each one was performed blindfolded once and unblindfolded once. The leader was told that he could ask for clarification on each task before they started. In each task, the leader and follower waited for our cue to start. They then lifted the board from off the ground. They completed the tasks as described later in this chapter and then placed the table on the ground. The process was repeated three times for a total of 36 tasks for each dyad. After completing all the tasks, they completed a survey, seen in Appendix A.

2.4 Tasks

In this section we show and describe each task. The pictures were from the screen shots that are shown to the leader on the iPad before each task. The colors around the edge helped orient the leader to the room. The walls in the room each had a colored poster board corresponding to the ones in the diagram. The narrower colored lines corresponded to marks on the floor. Only the ones used for that task are visible in the diagram while all the marks were always visible on the floor.

A start and completion diagram were shown for each task. Some tasks had two nominal finish diagrams. This was because we randomly choose a direction for the participants to go with a random number generator. For example, for the translation task they would be told to translate either to the right or to the left. We had two diagrams that showed one possibility in each, and only one of these diagrams was shown to a participant for any given trial.

The order that the dyad performed the tasks was random. Most of the tasks were designed so that there were at least two possible options for what they would need to do in order to prevent memorization of the tasks. For example, the rotation and translation task started in the same position. The participants would figure out that one particular starting position would always be a rotation or translation task, but they would not know which one it was or which direction they would be going. This was done on purpose since in real life, a teammate may try to anticipate the task but may be wrong.

2.4.1 Translation

In this task, the dyad translated the board laterally 1 meter to the left or right, as indicated by the iPad to the leader (Figure 2.2). Translation is a task that is common in the literature. One difference in our implementation, however, is that the follower did not know that this was a one degree of freedom task beforehand, which allowed us to see if this changed the behavior in any way. They had to differentiate the leader's intent between translation and rotation. In addition the follower had to determine if they were moving left or right.

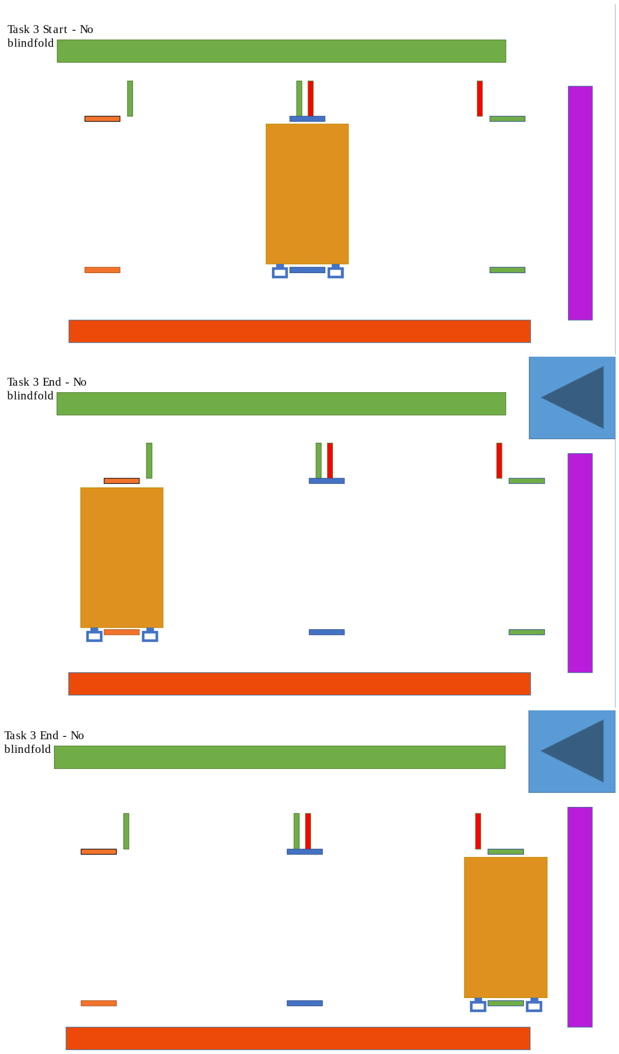
Task 3 Start - No blindfold
Task 3 End - No blindfold
Task 3 End - No blindfold

Figure 2.2: Translation task showing the starting position and possible ending positions.

2.4.2 Rotation

In this task, the dyad rotated 90 degrees around the follower clockwise or counterclockwise, as indicated by the iPad to the leader (Figure 2.3). This allowed us to see how a leader communicated translation versus rotation. This is a very difficult distinction to make and one that has been brought up in literature in the past, but with few solutions.

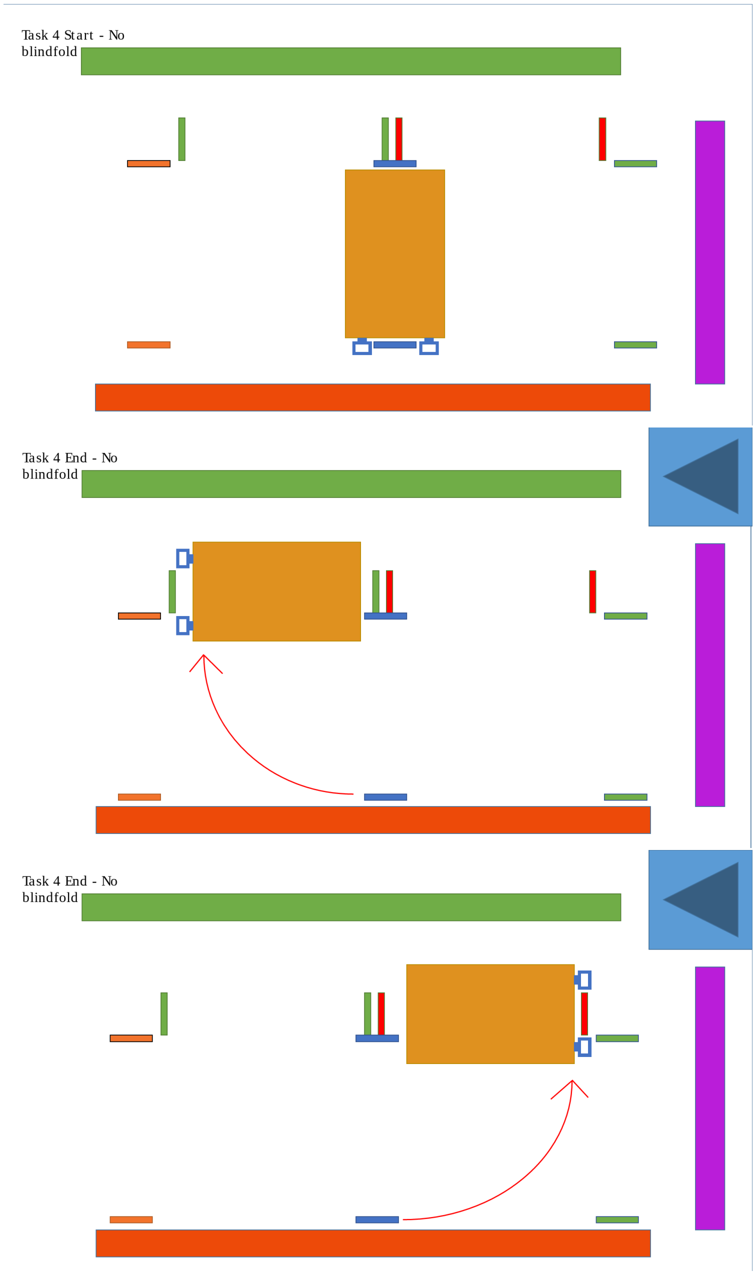
Task 4 Start - No blindfold
Task 4 End - No blindfold
Task 4 End - No blindfold

Figure 2.3: Rotation task showing the starting position and two possible ending positions.

2.4.3 Move table to a specific location avoiding an obstacle

In this task, the table is placed randomly in the room and must be moved to the red square while avoiding a small wooden post that was placed somewhere between the table and the red square (Figure 2.4). The post was a 4 in by 4 in wooden board that stood about 1 meter high. This is the only task that the follower could not predict because the board was started in a random location and the post was in a different location each time. It was different every time it was repeated and the follower had to follow the leaders guidance.

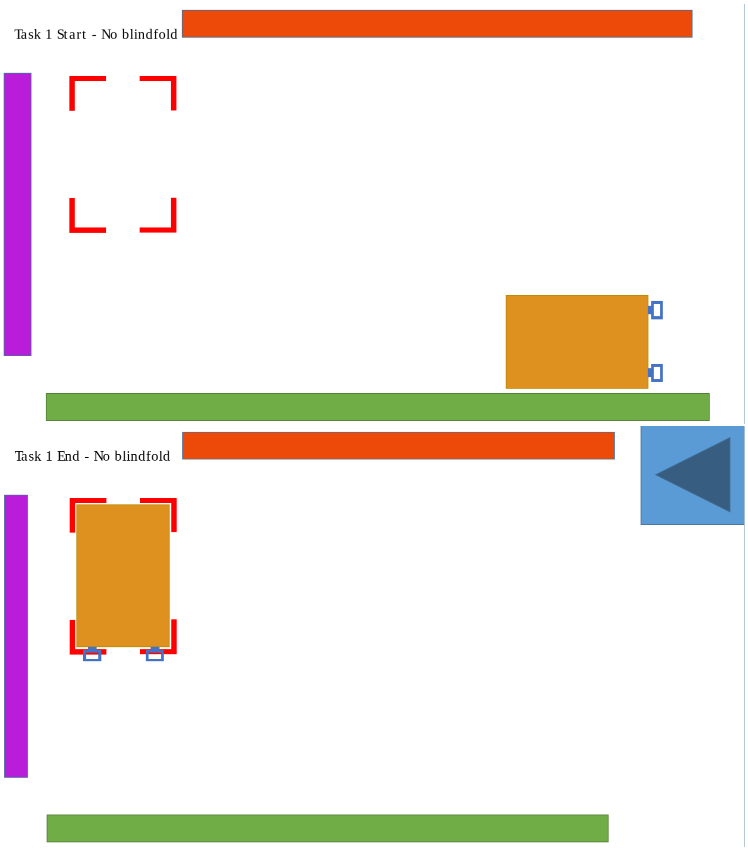

Figure 2.4: Task showing a random initial position and a final position within the red box.

2.4.4 Hallway navigation with leader facing direction of motion

This task required the leader to indicate either a right or left turn to the follower as indicated by the iPad to the leader (Figure 2.5).

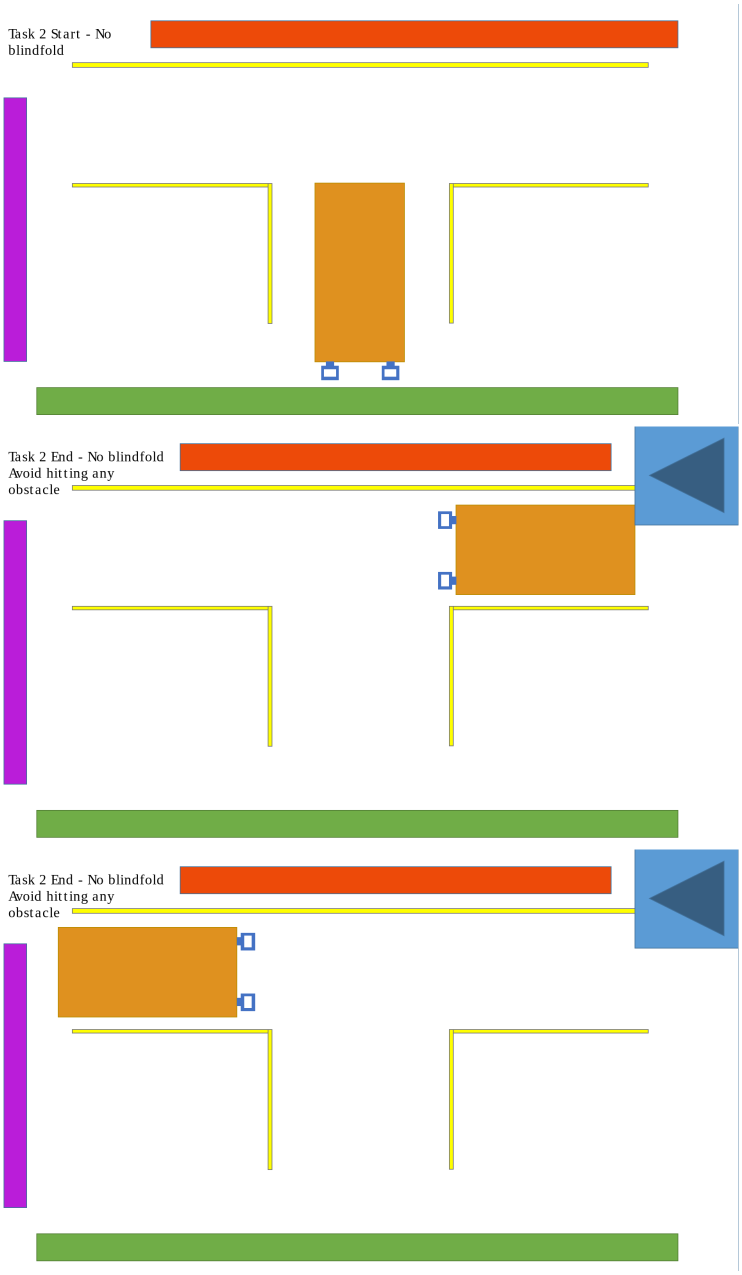
Task 2 Start - No blindfold
Task 2 End - No blindfold Avoid hitting any obstacle
Task 2 End - No blindfold Avoid hitting any obstacle

Figure 2.5: This task shows a hallway that the participants were instructed not to leave. The final two diagrams show two possible endings to the task. The leader pushes the table forward in this task.

2.4.5 Hallway navigation with leader facing opposite direction of motion

In the previous task the leader pushed on the board while in this task he pulled (Figure 2.6). From this, we could see what difference there was between pushing and pulling to indicate intent.

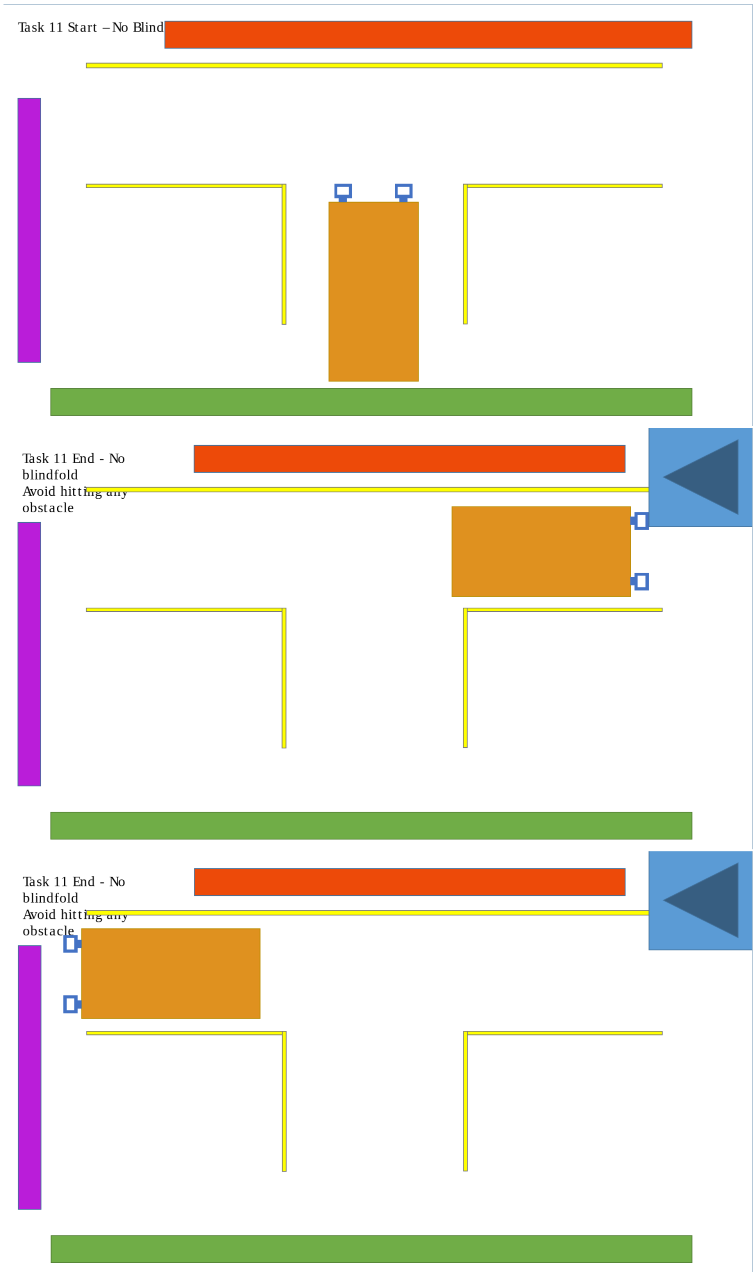
Task 11 Start – No Blind
Task 11 End - No blindfold
Avoid hitting any obstacle
Task 11 End - No blindfold
Avoid hitting any obstacle

Figure 2.6: This task shows a hallway that the participants were instructed not to leave like the previous task. The final two diagrams show two possible endings to the task. The leader pulls the table in this task.

2.4.6 Combined translation and rotation for six degree of freedom motion

In this task, the dyad followed a set circuit that included two obstacles they had to avoid (Figure 2.7). Screenshots of the task are shown in Figure 2.8. This task included rotation around the longitudinal axis (shown in Figure 2.1), which is a task that has not been shown in the literature for two people co-manipulating an object. This task forced the follower to accept that 2D translation and single-axis rotation were not the only possibilities. This was important because leaders often indicated rotation by twisting the board around its longitudinal axis instead of the vertical axis (The axes are described in Figure 2.1). This task forced the follower to abandon the assumption that twisting the board out of the horizontal plane was not needed for the task as this is an unrealistic assumption for many complicated co-manipulation tasks.

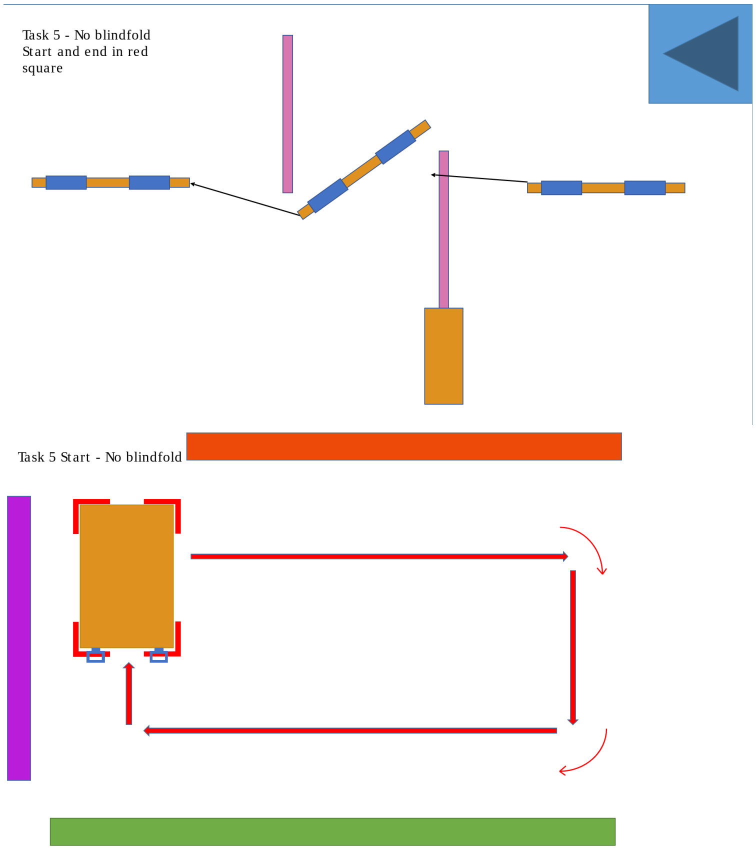

Figure 2.7: The first diagram demonstrates how the dyad must lift the table over one object and lower it under another. The second diagram shows the path they must follow. The objects in the first diagram are along that path.

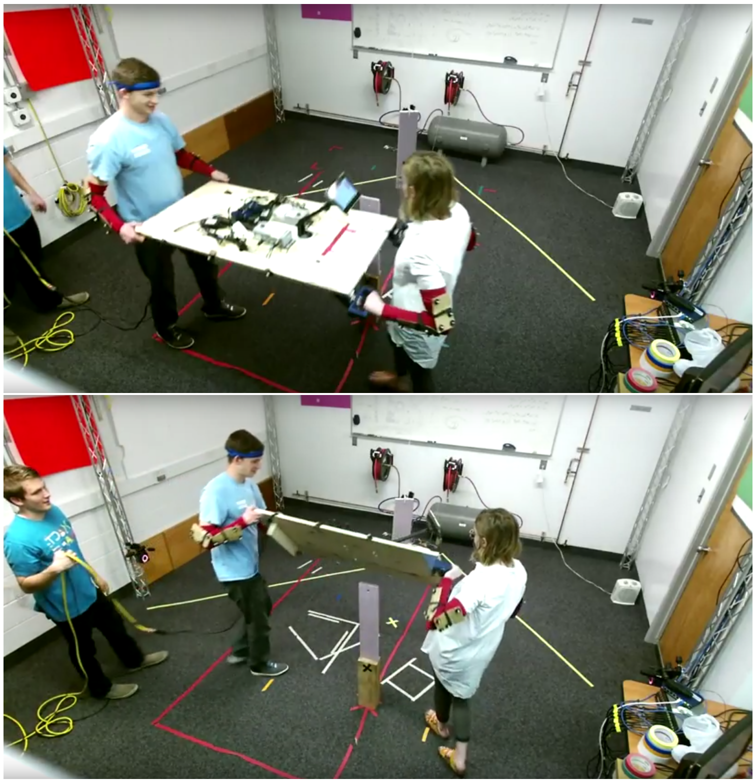

Figure 2.8: These are screenshots from the task that show how the dyad moved the table through the obstacles.

2.5 Recruiting and Participants

Subjects were recruited mostly from Brigham Young University campus. This was done with class announcements, social media, and direct contact with students on campus. Those who were recruited in advance signed up for an available time on a website. Each participant was given $10 for an hour of time. There were a total of 21 dyads that participated. There were 26 men and 16 women ages 18 to 38.

2.6 Data

Motion capture data and force/torque data were each collected at 200 Hz. The force/torque sensors were 36 cm apart and each one collected three orthogonal forces and three orthogonal torques. The Microsoft Kinect data was collected at 30 Hz. All data was recorded through the Robot Operating System (ROS), which is an open-source system for robot software development.

There were moments when a sensor failed and there is missing data. This was mostly when the motion capture system failed to see the table due to occlusion. This was taken into account in all analysis that was performed on the data.

2.7 Conclusion

We designed our study with the goal of understanding how humans communicate as they collaborate to move an object. This is a complex issue that has been studied in the past. In other sections of this thesis, we describe how we have used data from this study to understand human motion. The collection of this data, however, is itself a contribution. The way we use the data in this thesis is only a first application and analysis of the data. Our lab is continuing to analyze the data we have gathered to better learn how humans collaborate when co-manipulating an object.

There are a few changes we could make to the study in future work to improve it. The participants knew that all the movements depended on the colored markings on the floor shown in Figure 2.1 and in each diagram in Section 2.4. Because of this, the followers often tried to predict what they would do next. If the tasks did not depend on the markings on the floor, the tasks would be more random. This could be done by projecting colored lines on the floor for each individual task instead of using permanent marks. Another thing we would try in the future is to decouple vision and hearing in the follower. We allowed the dyad to talk when not blindfolded, but prohibited it when the follower was blindfolded. We would, in another study, prevent all talking during the trial. In future studies, we would also expect that varying the object size, weight, and rigidity may be important to generalize our findings.

CHAPTER 3. NEURAL NETWORK PREDICTION OF SHORT-TERM HUMAN INTENT

This chapter describes the development of a neural network to model short-term human intent. The purpose of the model is to predict where a co-manipulated object will move based on past information so that a robot can better assist a human teammate. While there are many methods of time series prediction, we chose to use a neural network because it can approximate any function and is relatively simple to formulate. In this chapter, we describe the formulation of our neural net, the training process, and one other simple intention prediction method to which we compare our neural net. We then describe the results of the neural net prediction with a human-human dyad and then with a human-robot dyad.

3.1 Neural Network

To develop a nonlinear estimator of human intention, we formulated a neural network (explained in Section 1.3) using the Google Tensorflow API, an open-source machine-learning software library [41].

3.1.1 Topology

We attempted different neural net formulations. We tried to use a convolutional neural net as described in Section 3.1.2 to be able to interpret how forces are used in physical co-manipulation. We used a recurrent neural net which has temporal behavior which seemed appropriate for our application. In addition to being notoriously difficult to train, however, we did not achieve results as good as with a traditional neural net. There may be promise in these other approaches, but in the end we achieved satisfactory results with a traditional neural net. For the neural net to achieve good performance, we had to do significant pre-processing of the data and train it in a specific way to make it stable for iterated prediction. All of that is described in this chapter. Our final estimator

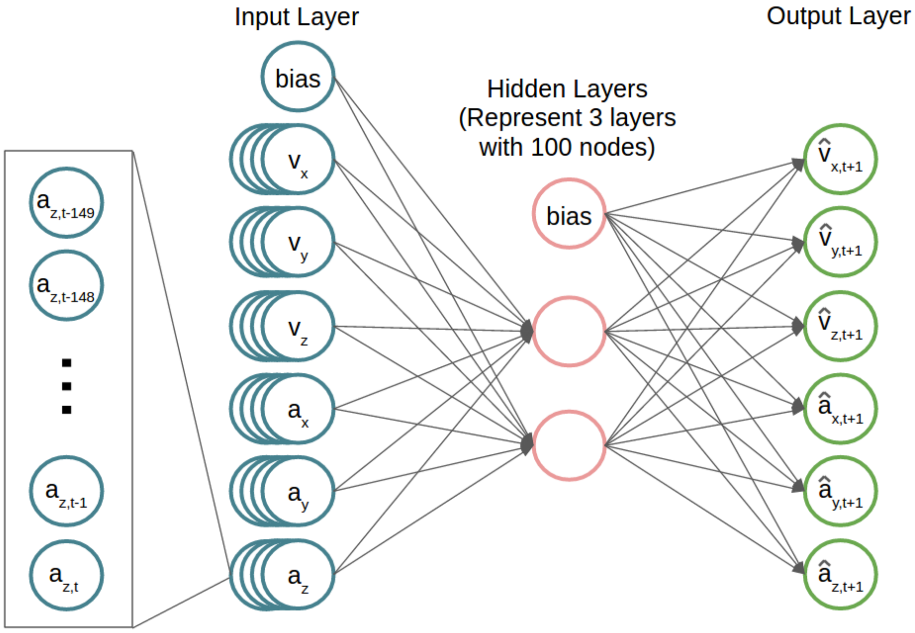

Figure 3.1: Symbolic representation of the neural network. a represents acceleration and v represents velocity. The subscripts, x, y, and z, represent the orthogonal axes. $t+1$ represents one time step into the future.

consisted of two neural networks, one for translation and one for rotation, each with three hidden layers each with 100 nodes as seen in Figure 3.1. Rectified linear units (ReLUs) were used for the activation of each node as these are the current state of the art for neural networks:

$$ReLU = \max(0, x) \tag{3.1}$$

The final layer had no activation function because the output of the neural net needs to be able to predict any arbitrary value, while most activation functions have limited output ranges. The process of choosing a neural network structure and other parameters was not exhaustive and it is possible a that better structure and parameters could be obtained. In the end, a simple structure using only motion data as input worked well. Better structures, including methods other than neural networks, may exist but our purpose in this thesis was to show that short-term human intent estimation was feasible. It was shown by Chipalkatty et al. that more complex predictions of future movement can actually decrease performance if they do not agree with what the human is trying to do. Chipalkatty et al. found that it was more important that the human understand what the robot will do next [36]. In the case of our co-manipulation task, a bad prediction could cause the robot to move in an unexpected way which could be frustrating or even physically dangerous for the human participating in the dyad. This means that it is important that the prediction be not only accurate, but also reliable. The inputs to the neural network were 150 past steps of velocity ($\dot{x}$, $\dot{y}$, $\dot{z}$), acceleration ($\ddot{x}$, $\ddot{y}$, $\ddot{z}$), angular velocity ($\dot{\phi}$, $\dot{\theta}$, $\dot{\psi}$), and angular acceleration ($\ddot{\phi}$, $\ddot{\theta}$, $\ddot{\psi}$) of the board in the longitudinal, horizontal, and vertical direction, $\{x_{t-149}, x_{t-148} \ldots, x_{t-1}, x_t\}$.

$$x_t = \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ \ddot{x} \\ \ddot{y} \\ \ddot{z} \\ \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \\ \ddot{\phi} \\ \ddot{\theta} \\ \ddot{\psi} \end{bmatrix} \tag{3.2}$$

where x, y, and z are the longitudinal, horizontal, and vertical direction. ϕ, θ, and ψ are the body-fixed roll, pitch, and yaw. Roll is about the longitudinal axis. Pitch is about the horizontal axis. Yaw is about the vertical Axis. The axes are defined in Figure 2.1. The outputs were the predicted velocity, acceleration, angular velocity, and angular acceleration in the longitudinal, horizontal, and vertical direction for 1 time step into the future, $\hat{x}_{t+1}$, where $\hat{x}$ indicates a predicted value. The prediction can be iterated to predict 50 steps into the future for velocity and 45 steps into the future for angular velocity as shown in Table 3.1 as described in the next paragraph.

Our neural net formulation used what Engel et al. describe as iterated prediction shown in Algorithm 1 and in Figure 3.2 which they used to predict time series data from a laser in a chaotic state [42]. The neural network itself only ever predicts one time step into the future. Then, the

Table 3.1: Neural network inputs and outputs.

Data required for prediction	150 steps or .75 seconds	velocity, acceleration, angular velocity, and angular acceleration
Prediction length into the future for position	50 steps or .25 seconds	velocity and acceleration
Prediction length for rotation	45 steps or .225 seconds	angular velocity and angular acceleration

prediction, $\hat{x}_{t+1}$, is appended to the input to give $\{x_{t-149}, x_{t-148} ..., x_{t-1}, x_t\}$. The first step of the input is dropped to obtain a new input of past motions for the neural net, $\{x_{t-148}, x_{t-147} ..., x_t, \hat{x}_{t+1}\}$. The new data is input into the neural net which outputs a prediction one step forward, but two total steps into the future, $\hat{x}_{t+2}$. This is then appended to the input. The process is repeated 50 times to obtain a prediction of 50 steps, $\{\hat{x}_{t+1}, \hat{x}_{t+2} ..., \hat{x}_{t+49}, \hat{x}_{t+50}\}$. Because the outputs of each prediction step become the inputs for the next, the inputs and outputs have to be the same variables.

Algorithm 1 Iterated prediction

```
Use last 150 steps of consecutive data
for 50 steps do
    Predict next step
    Append prediction to end of data vector
    Drop the first set of values
```

3.1.2 Effect of Force Data on Prediction

Originally, we expected to train the neural net on the force and motion data from the exploratory study. As we trained the neural net, we discovered that force data was difficult to use and gave inferior prediction of short-term human intent. There are two proposed reasons for the low quality prediction: 1) the relationship between the forces applied to the table by each person and the future motion of the table is complex. This is in part because each human dyad learned to communicate by force in their own way. A dyad will learn their own form of communication and negotiate unique roles as shown in [39]. 2) It is possible that people do not use force for low-level control. Instead, we may act as impedance devices, and force is used to change the set-point in our impedance control, not continuously, but in moments when the leader desires to change the

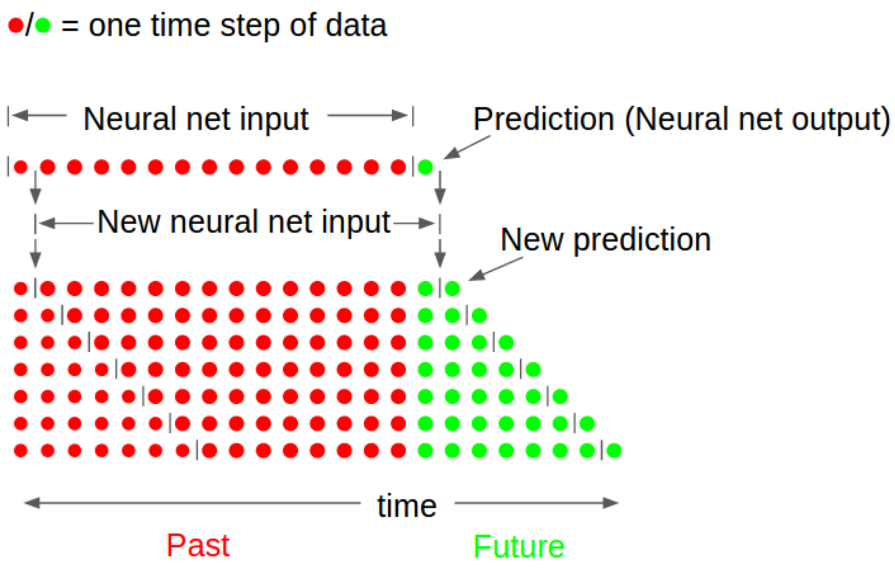

Figure 3.2: Iterated prediction allows us to predict more than one step into the future.

trajectory. We attempted to use a convolutional neural net to be able to detect moments when the leader uses force to communicate, but were unsuccessful. We are currently doing research into the effect of forces to test these hypotheses and use them for better human-robot co-manipulation.

We discovered that more accurate predictions could be made without force data. While overall performance and initiation of motion may be improved with force information, the motion data alone was sufficient for an accurate short term prediction. A more complicated neural net structure would have possibly been able to ignore the force data or even learn the complicated dynamics that govern the relationship between force and motion in human co-manipulation. For our initial implementation, we achieved good results by leaving out the force data. This was advantageous as it meant that the estimation of short-term human intent could be used on robots without force sensors in the wrist. We also determined that the neural net can predict better if two separate neural nets are created, one for translation and one for rotation.

3.1.3 Training

The weights of each node were initialized using Xavier initialization which is the state of the art for initializing the weights for optimal convergence of the neural net when using ReLU activation functions [43].

To train our net, we preprocessed the data to improve the results. The velocity, acceleration, angular velocity, and angular acceleration data were scaled to have zero mean and standard deviation of one over the entire set of data. This was then inverted on the output to show the results in their proper units. This same scaling can be used on new data even though the mean and standard deviation will be different. All trials with any missing or corrupted data were entirely thrown out. 79 percent of the trials had no missing or bad data. Most of the bad data was because of missing poses from motion capture where too many motion capture markers were occluded from the cameras. The entire set of data for every dyad and task, excluding trials with missing data, consisted of 1.6 million steps. Each dyad was assigned randomly to training and validation sets. 75 percent of the data were assigned to the training set and the other 25 percent to the validation set.

Algorithm 2 Neural Net Training for Iterated Prediction

```
while Neural Net continues to converge in a timely manner do
    Time = 0
    while ThresholdCounter $\leq 5$ and Time $\geq 1000$ seconds do
        Randomly choose 32 batches from the training set
        Evaluate batches through the net
        Update neural net weights through back-propagation
        Randomly choose 32 batches from the validation set
        Run batches through the net
        if MSE is less than threshold then
            ThresholdCounter $+= 1$
        else
            ThresholdCounter $= 0$
    end
    Calculate predictions for every possible set of 150 data points
    A new data set is created with prediction appended and the oldest data point dropped
    This data set is added to the previous data sets
end
The neural net is converged
```

The neural net had to be trained in a specific way in order to make the iterated prediction $\hat{x}_{t+1}$ stable beyond the first step as shown in Algorithm 2. This process is adapted for neural nets from [42] where they used a similar process with a kernel recursive least squares algorithm. For our training, batches of data were created that randomly pulled in 32 random sets of 150 steps of data from the entire training set. These numbers were chosen somewhat arbitrarily after some experimentation. Another 32 sets of 150 steps were created from the validation set. The neural net was trained on new training batches until the cost function reached a value less than a threshold that we chose for five consecutive validation batches. We used the mean squared error (MSE) for the cost function.

$$MSE = \frac{\sum_{n=1}^{q} \sum_{n=1}^{p} (\hat{x}_{n,t+1} - x_{n,t+1})^2}{p * q} \tag{3.3}$$

where $\hat{x}_{n,t+1}$ is the predicted value from the neural net and $x_{n,t+1}$ is the actual measured value from real data. The set in the batch is represented by n. The size of the batch is p which is 32 in this case, and q is the number of variables in x. $t+1$ represents a value one step into the future. Once the MSE is below the threshold for 5 consecutive batches and the neural net has been training for at least 1000 seconds, a prediction is calculated for every possible set of 150 steps of data. A new data set is created that used 149 steps of real data and the prediction appended to the end, $\{x_{t-148}, x_{t-147}..., x_{t-1}, x_t, \hat{x}_{t+1}\}$. This new set is as large as the original set because a prediction is made for every possible set of 150 steps. The neural network was then trained on a combined data set that included the original data and the new data set that included the prediction. Once this training was complete, a new data set was created with 148 steps of real data and two predictions after it, $\{x_{t-147}, x_{t-146}..., x_{t-1}, x_t, \hat{x}_{t+1}, \hat{x}_{t+2}\}$. The same neural net was then trained again. This is continued until the neural net no longer converges in a reasonable time. By training the neural network on data that includes predictions, the stability of the prediction was improved.

The length of the prediction was limited by our computational resources as we train the neural network. The translation neural net was trained to predict for 50 steps, or .25 seconds into the future. We stopped training the neural net once one step took more than half a day to train. The neural net can be used to predict further into the future, but is only reliable for 50 steps. The rotation neural net predicted 45 steps, or .225 seconds into the future. We do not know if this is because a fundamental limit on the ability to predict after that amount of time exists or not. We imagine there would be such a limit as the dyad has time to make decisions on where to move, but determining how long it would be for different teams is for future work. It could also be due to limitations in the implementation of our current neural network formulation.

Table 3.2: The mean squared error of each polynomial predictor between the predicted and actual value 0.25 seconds from the beginning of the prediction.

Order	MSE
2	0.58
4	0.22
6	0.08
8	0.08
10	0.26
12	0.65
14	4.20

After obtaining these results, we cross-validated the neural net. Cross-validation is a technique for ensuring that the neural net is generalizable and not overfit to the training data. As mentioned above, the data was split into training and testing steps before training. To cross-validate, the entire data set was repartitioned randomly. We did this multiple ways. Some of the partitions kept every task from each dyad in the same set. Other times, the tasks were all split randomly. The neural net was retrained on each training set and then validated on the validation set. Each time, the neural net performed about the same on each validation set. This ensured that the neural net was general and that it worked just as well on new data from old dyads and new data from new dyads. Since the neural nets all gave similar results, all the results here are on the last neural net we created.

3.2 Polynomial Fit Predictor for Short-Term Human Intent

We also developed a polynomial fit estimator to compare to our neural network. A polynomial regression does not have to be trained on data and can be implemented much more simply then a neural net. We compared the neural net to this method because if the polynomial fit is comparable, then it should be used instead of the neural net. Also, it was important to know that the neural net was not just fitting something like a polynomial to the data, but that it was learning a dynamic, time-dependent model. Like the neural net, it also took in the previous 150 steps, but fit an 8th order polynomial to that data. The 8th order was chosen because it had the best performance of different polynomial orders as shown in Table 3.2. The polynomial was then extrapolated forward 50 steps just like the neural net prediction.

3.3 Robot Controller

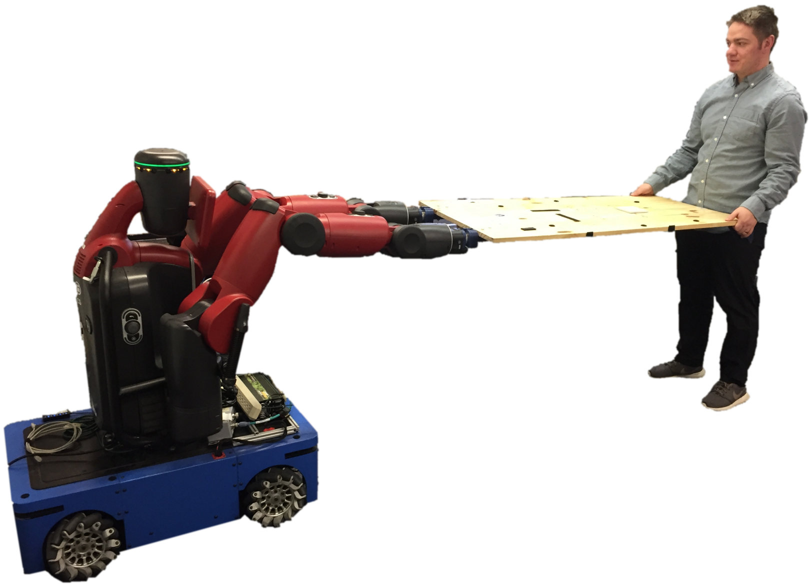

Figure 3.3: Rethink Robotics Baxter robot mounted on HStar Technologies AMP-1 holonomic base carrying the table with a person.

The purpose of predicting short-term human intent was to have the robot use the prediction in a closed-loop feedback scheme. While we validated the prediction on human-human data, this did not ensure that adding a robot in the loop with a simple controller would not affect the prediction. This was because it was unclear if the predictor would be affected by a different type of dyad. We implemented a simple controller on the robot (described in Section 1.4) that would allow it to quickly follow the human's motion while they both carry the table. The arms of the robot had low impedance controllers implemented that allowed the human to move the table easily. The base then tracked the end effector of the arms. This allowed the human to move the table easily as far as desired as the robot followed along supporting the other side of the table. This controller does not use the short-term human intent prediction for control. That is future work. Instead, we calculate the predicted intent of the human-robot dyad using the neural net prediction and compare this to the actual data. For control, the Baxter robot arms were rigidly attached to the table. Each arm was running a low impedance controller with commanded joint angles specified by the position of the arm relative to the table before the task began. For this initial implementation, when the arms were displaced, the mobile base displaced by the same amount to put the arms in their original pose relative to the base. This was implemented for translation but not rotation. We then recorded motion data as the human and robot carried the table in order to test our short-term human intent estimator.

3.4 Results and Discussion

In this section, we show that neural net predicts short-term human intent well on real data of both human-human and human-robot dyads. We also show that the performance of the neural net is better than the polynomial fit estimator. We show this in individual tasks as well as the average performance over many tasks. We then show that the neural net is fairly robust to noise while the polynomial fit is not.

3.4.1 Neural Net Performance

Figures 3.4 through 3.6 show the neural network predictions of velocity in the longitudinal, horizontal, and vertical directions for a single hallway navigation task as described in Section 2.4.4. The predictions are only shown for 50 steps at a time with a gap of another 50 steps between each prediction for clarity. The neural net was trained for data that is collected at 200 Hz. The neural net is capable of predicting 50 steps into the future at 2500 Hz, which means a prediction can be made every time a new data point is collected. The actual measured velocity is shown for the whole task in blue. The predicted velocity is shown in red starting every second and each one continues for 50 time steps or .25 seconds. Figures 3.7 through 3.9 shows the same thing for angular velocity for a rotation task as described in Section 2.4.2. As seen, the predictions are very accurate for that time scale. Here we only show velocity, but the acceleration data must also be predicted because each velocity prediction depends on the prediction of acceleration for the time step before it. Without acceleration data, the neural net performance degrades.

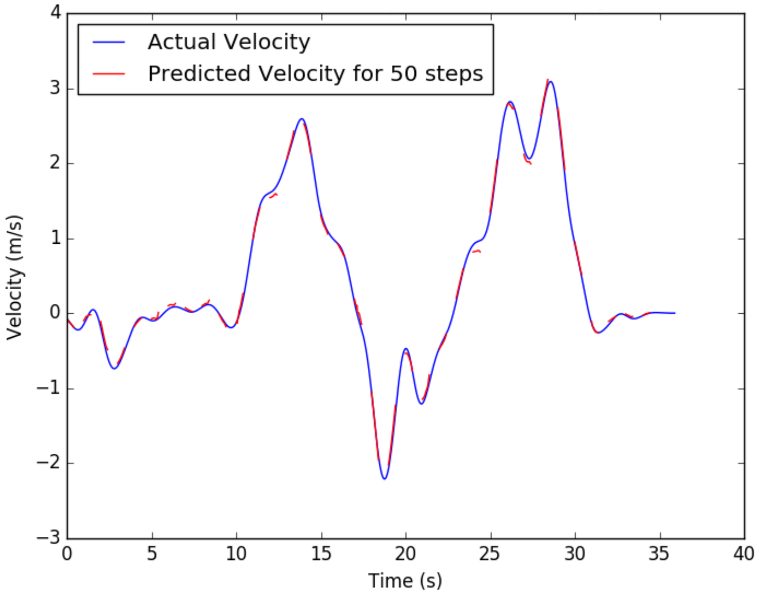

Figure 3.4: Comparison of the velocity prediction to actual future data in the longitudinal direction while a human dyad moves the table. The directions are body-fixed to the table frame. Each red line is a separate 50 step prediction using the 150 steps before it.

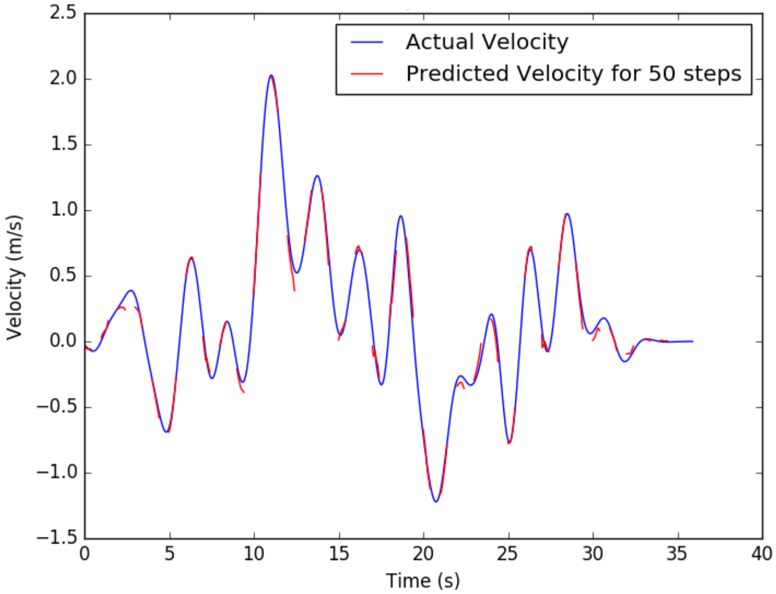

Figure 3.5: Horizontal direction prediction versus actual velocity.

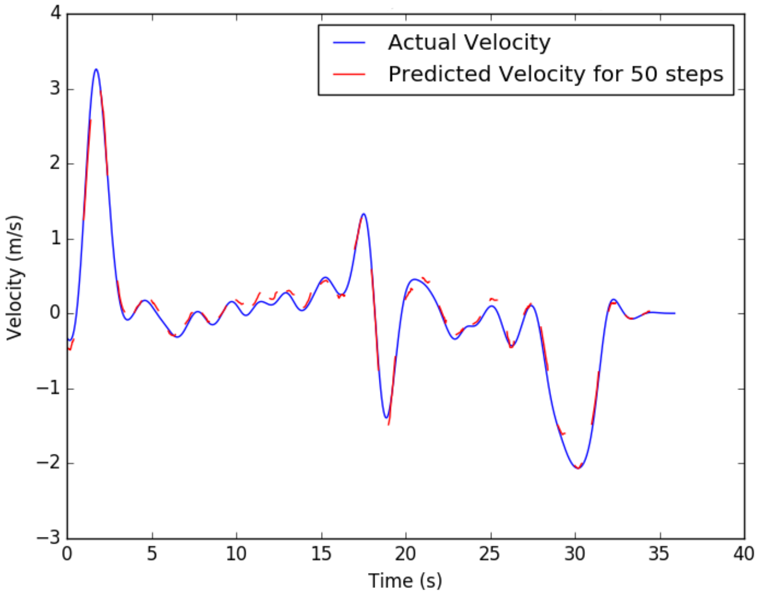

Figure 3.6: Vertical direction prediction versus actual velocity.

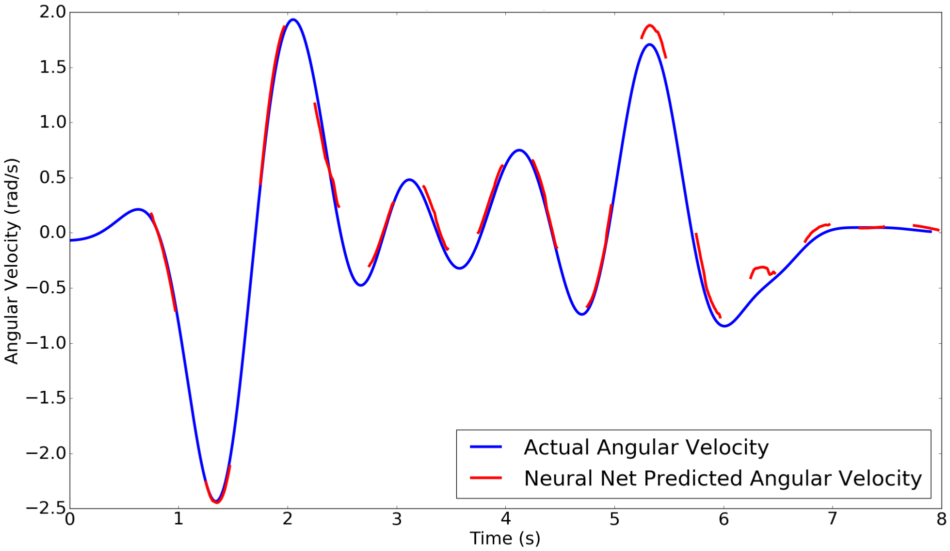

Figure 3.7: Comparison of roll angular velocity prediction to actual future data while a human dyad moves the table. These are in the table frame. Each red line is a separate 50 step prediction using the 150 steps before it.

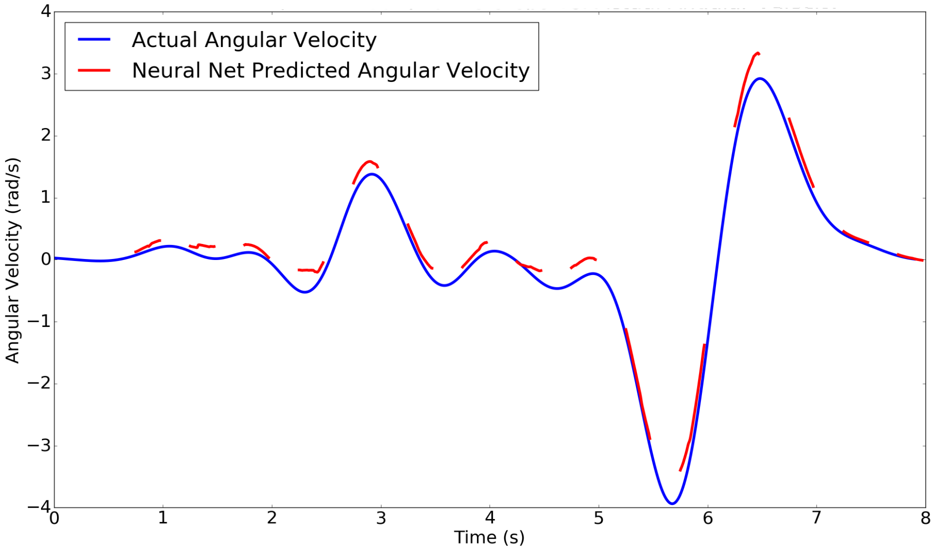

Figure 3.8: Pitch angular velocity prediction versus actual angular velocity.

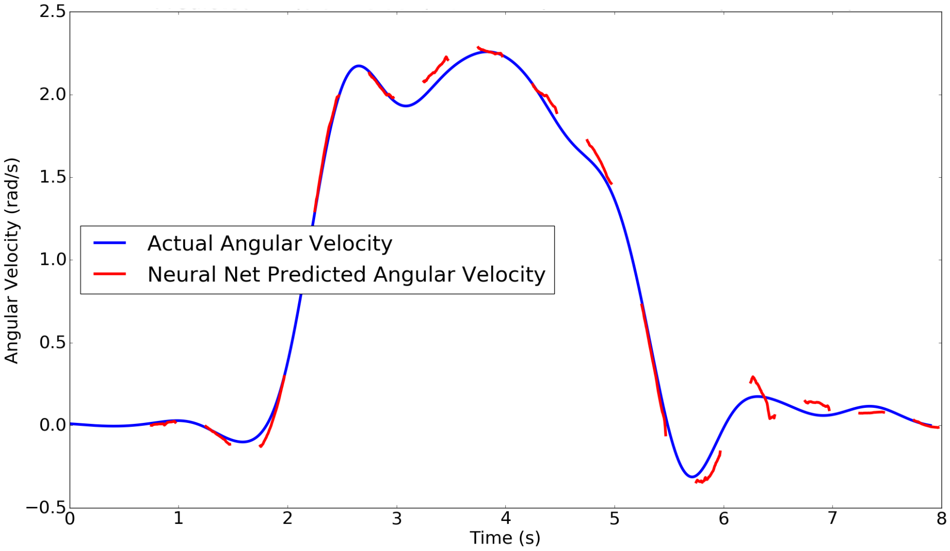

Figure 3.9: Yaw angular velocity prediction versus actual angular velocity.

3.4.2 Comparison to Polynomial Fit Predictor

The neural net prediction is compared to the polynomial fit prediction for a single randomly selected task in Figure 3.10. The polynomial fit was accurate in many cases especially when velocity changed slowly, but was more prone to large errors. Figure 3.11 shows how added white noise affects the polynomial fit more than the neural network. The amplitude of the noise was consistent with the empirically measured noise of the end effector position on our robot which was 0.05 m/s which gives a signal-to-noise ratio of about 60. This is a major issue as filtering the function to be smooth adds lag to the signal which takes away time from the prediction so that it cannot estimate as far into the future. Figure 3.12 shows how the prediction degrades over time for each predictor. The polynomial predictor degrades smoothly but quickly. The neural net is accurate for 50 steps and then degrades quickly in some cases.

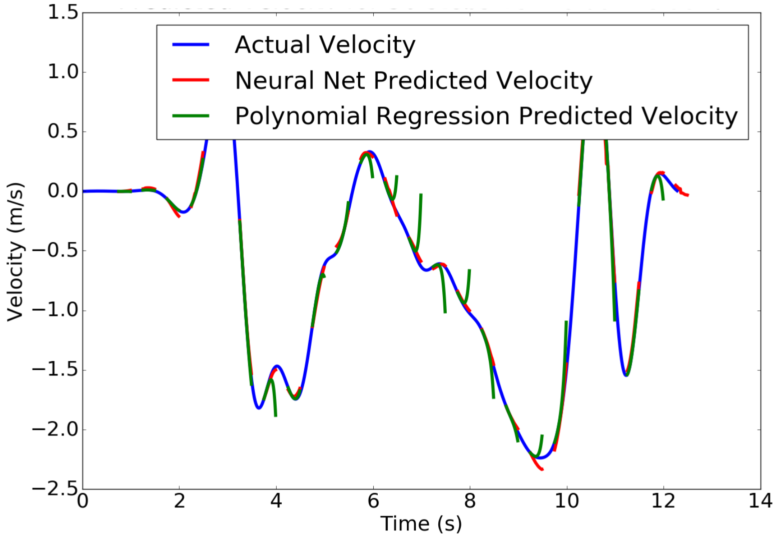

Figure 3.10: Predictions using the neural net and a polynomial estimator for velocity in one task. While both are accurate in most cases, in several cases the polynomial prediction is far from the actual data.

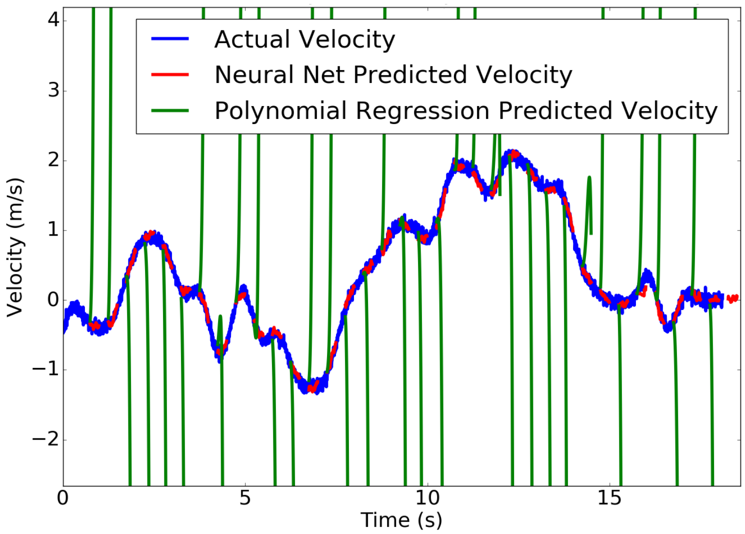

Figure 3.11: Predictions using the neural net and a polynomial estimator for velocity in one task where white noise has been added. The noise causes the polynomial prediction to go completely unstable while the neural network prediction is fairly robust.

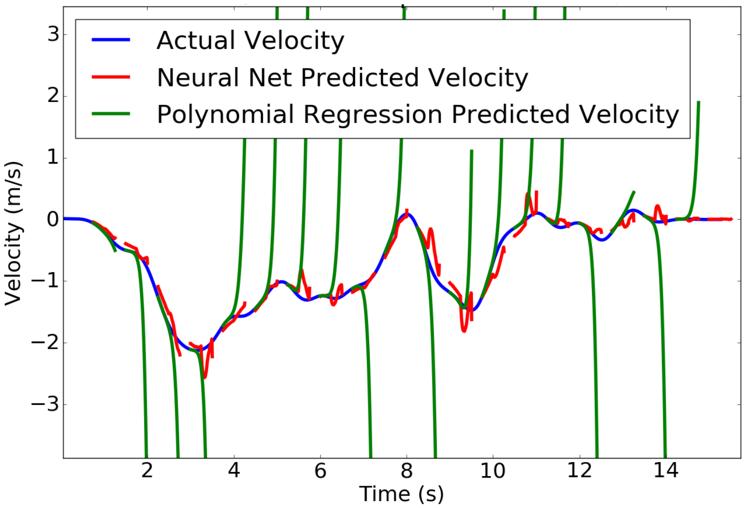

Figure 3.12: Comparison of the velocity prediction to actual future data in the x direction while a human dyad moves the table. Each red line is a separate 100 step prediction using the 150 steps before it. The green line is the polynomial prediction. This shows how the quality of the neural net prediction decreases significantly after 50 steps as expected.

3.4.3 MSE of each Predictor

Figure 3.13 shows the average MSE (Equation 3.3) of the prediction from the start of the prediction to 0.5 seconds, or 100 time steps, using predictions from every time step from 20 randomly chosen tasks from the training set. This shows that the estimator is general enough to work on many tasks. We only expect good performance for the first 50 time steps. Figure 3.14 shows the same thing using the validation set. The similarity of these shows that the neural net did not overfit the data. As seen, the polynomial prediction was very good for short time horizons but degrades quickly. The MSE of the polynomial prediction reaches about 1500 (m/s)^2 at 0.5 seconds. The neural net is very accurate for the 0.25 seconds that it was trained for, after which it quickly degrades. Interestingly, the first 50 steps are all predicted with the same accuracy. This was due to the way the neural net was trained. Figure 3.15 shows the MSE with white noise added. Note the change in scale of the y-axis compared to the other MSE graphs. The polynomial prediction degraded very quickly. Figure 3.16 shows the mean squared error along with the 10th percentile and 90th percentile for translation prediction. Figure 3.17 shows this for rotation data. Figure 3.18 shows the translation data again but with noise added. Notice that the scale on this figure is different than the others and that the noise causes the polynomial prediction to be unstable while the neural net is robust to the noise.

3.4.4 Estimation with a Robot in the Loop

We ran the estimator while a human-robot dyad moved the table for a single trial. This was not guaranteed to work because the dyad could not be expected to interact like a human-human dyad. Figures 3.19 through 3.21 show that the neural net predicted the future motion very well even though the dynamics were different.

It is clear that the neural net estimator was able to predict short-term human intent over a short time horizon. There may be additional methods that can improve on this prediction. However, the results in this thesis show that an accurate and reliable prediction can be made and on a time scale that should allow very responsive human-robot interaction behavior.

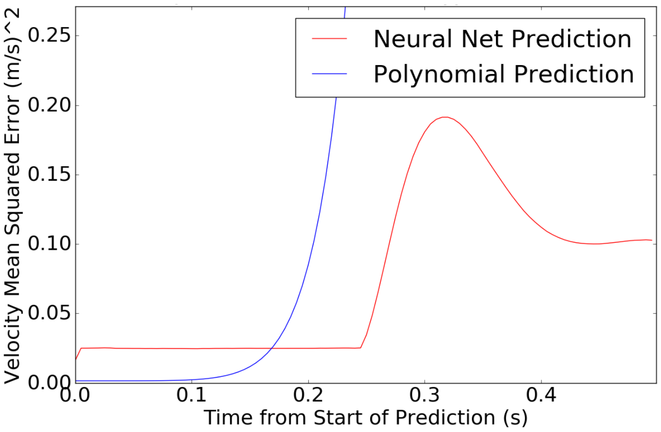

Figure 3.13: Mean squared error of the neural network prediction and the polynomial prediction for 0.5 seconds using the training data set. The neural net is specifically trained for the first 0.25 seconds which are flat, after which the performance of the neural network significantly degrades. The polynomial prediction MSE reaches 1500 (m/s)^2 by 0.5 seconds. Interestingly, the neural net prediction increases and then decreases after the first 50 steps. It is difficult to know why neural nets behave in certain ways and this warrants further investigation.

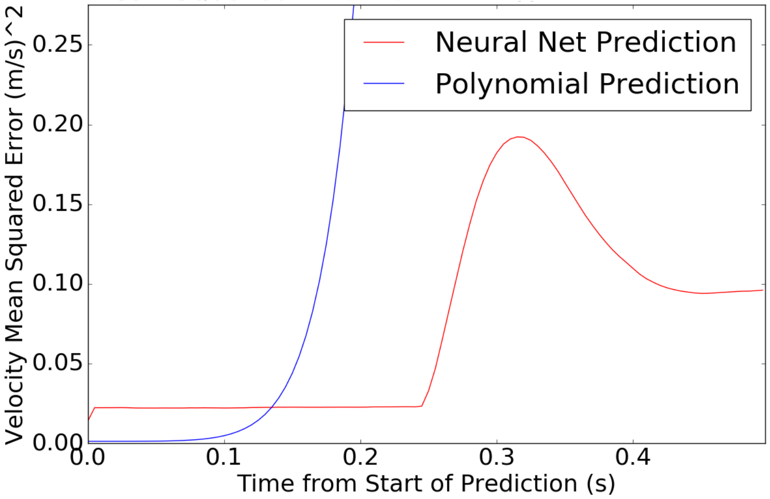

Figure 3.14: Mean squared error of the neural network prediction and the polynomial prediction for 0.5 seconds using the validation data set. There is very little difference between the prediction on the validation and training data sets. The polynomial prediction MSE reaches 1500 (m/s)^2 by 0.5 seconds.

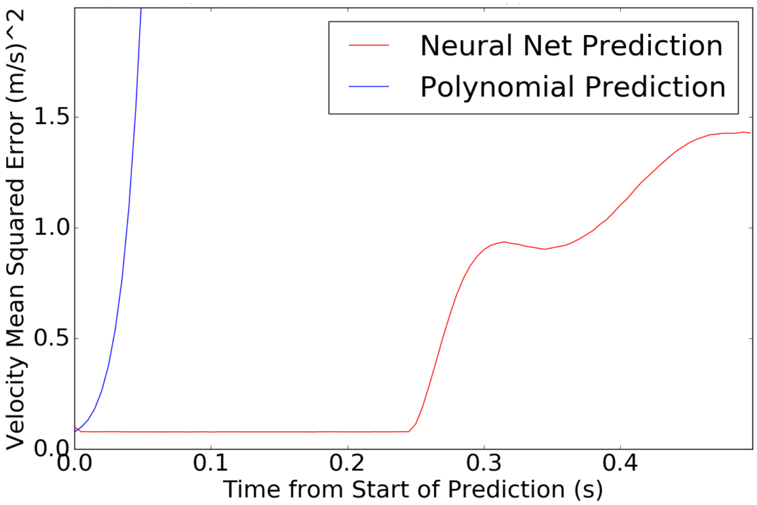

Figure 3.15: Mean squared error of the neural network prediction and the polynomial prediction for 0.5 seconds using data with added noise similar to the noise from the encoders on our robot. While both predictions are degraded, the polynomial prediction becomes useless. An 8th order polynomial does not extrapolate well when there is noise. Notice that the scale of this graph is different than the other MSE graphs. The polynomial prediction MSE reaches 1,000,000 $(m/s)^2$ by 0.5 seconds.

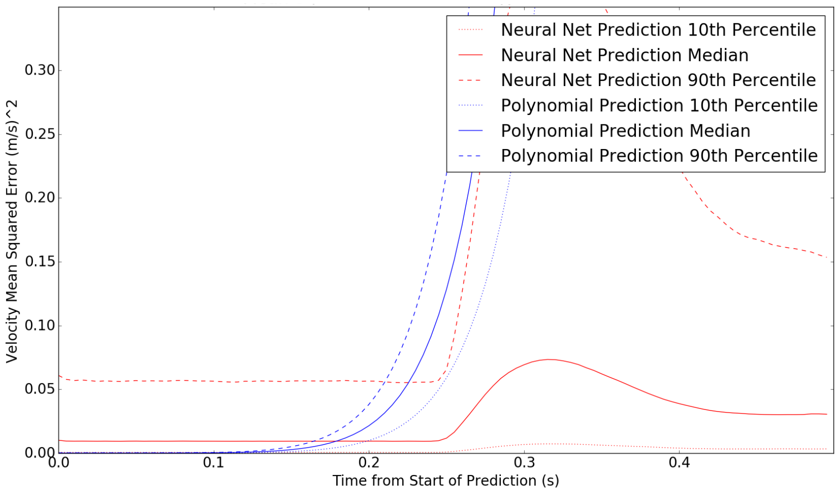

Figure 3.16: 10th percentile, median, and 90th percentile squared error plots for the neural net prediction in red and the polynomial prediction in blue.

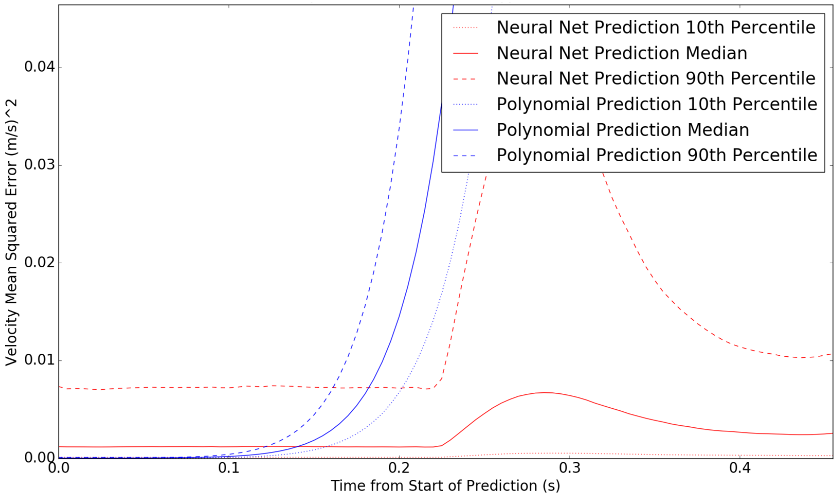

Figure 3.17: 10th percentile, median, and 90th percentile squared error plots for the neural net prediction in red and the polynomial prediction in blue for rotation.

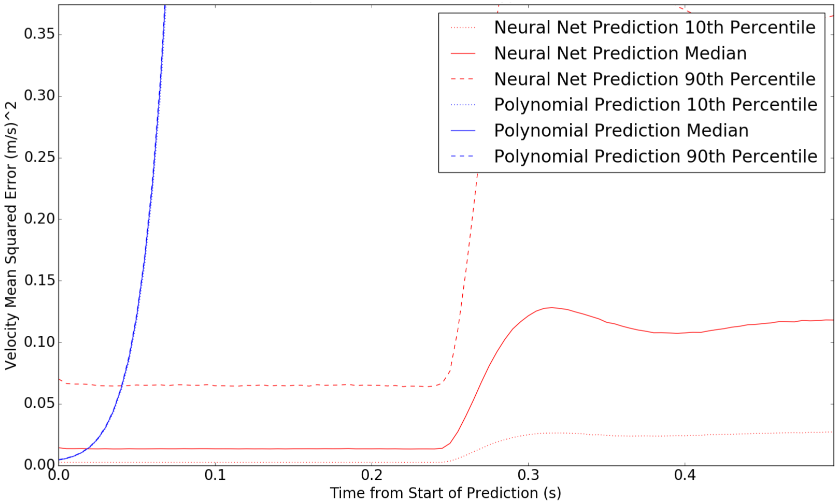

Figure 3.18: 10th percentile, median, and 90th percentile squared error plots for the neural net prediction in red and the polynomial prediction in blue when noise is added to the data. The polynomial prediction lines all overlap. Notice that the values for the neural net prediction are comparable to the ones in Figure 3.16.

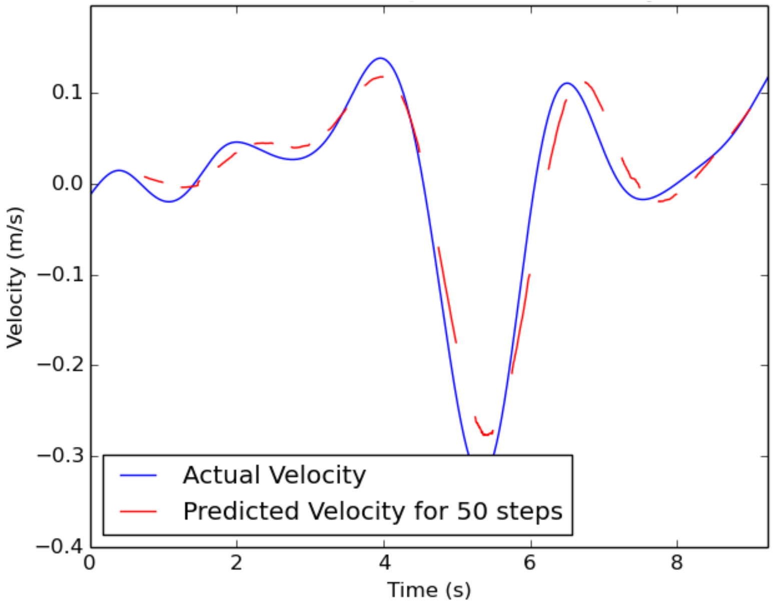

Figure 3.19: Comparison of velocity prediction to actual future data in the longitudinal direction while a human-robot dyad moves the table. Each red line is a separate 50 step prediction using the 150 steps of real data before it.

3.5 Conclusion

We have shown that short-term human intent can be estimated accurately from previous motion of the object that is being co-manipulated and that extrapolating instead with simple polynomial fits does not work sufficiently well for control. We achieve 0.02 average mean squared error over the first 0.25 seconds over about 20 tasks. We have also shown that with a mobile robot using a very simple controller in the loop, the prediction method is still valid. This work lays the foundation for future controller development which could use the human intention as a direct input or which could mediate the short-term human intention and modify it to improve overall performance through haptic or other means of communication. Ideally, the human and robot would share the leadership of the task. Advantages for this capability to share leadership include the following:

- The robot could be a more efficient leader.
- The robot could better keep away from its own joint limits.

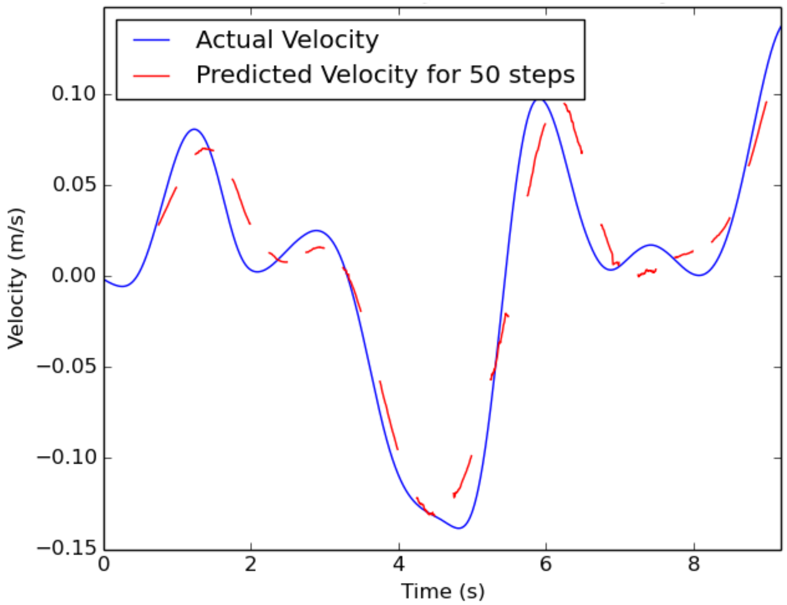

Figure 3.20: Horizontal direction comparison.

- The robot could keep the human from violating constraints of the task (e.g. the object can not rotate more than 15 degrees if the object is a person on a stretcher).
- The robot could have knowledge of the environment not known by the human operator [44].

As described in Section 1.5, without any prediction of human intent, the best controllers that have been shown are variable impedance controllers which cause the robot to be a burden to be dragged instead of an assistant. If the robot can predict even short-term human intent, then it can stay caught up with the human.

Interestingly, although even a short prediction into the future should allow the robot to work better with the human, we have been able to predict motions over a time horizon that is comparable to human reaction time. This seems to imply that if used in a closed-loop control scheme, this estimator would be enough to help a robot perform more like a teammate and less like a tool. We expect that this research along with future controller development will allow robots to work more intuitively and effectively with humans in collaborative object manipulation tasks.

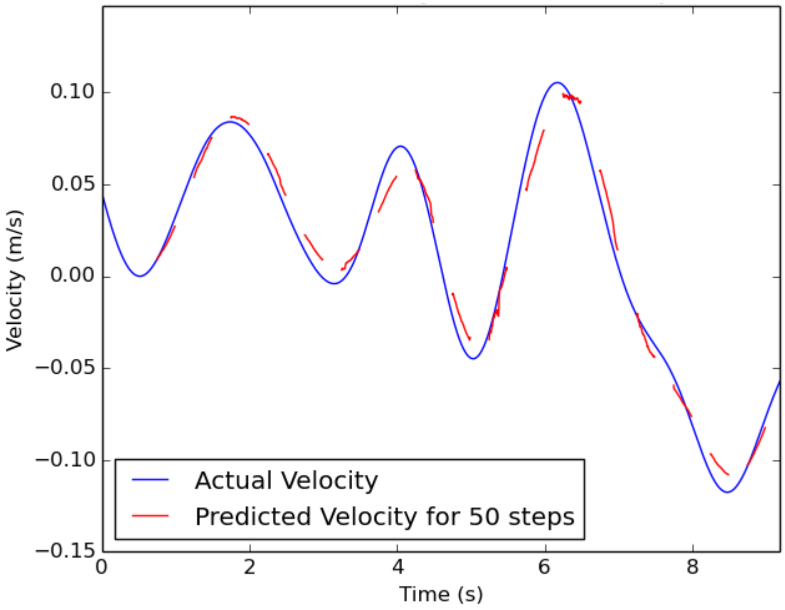

Figure 3.21: Vertical direction comparison.

CHAPTER 4. NEURAL NETWORK MODEL OF SOFT ROBOTS

The advent of soft robots is one of the many things that may make it more feasible for robots to work closely with people. Some soft robots can still apply high forces but they have low inertia, which makes collisions less dangerous. These are the types of robot that we imagine being common working in close proximity with humans.

This chapter is related to the rest of the thesis in two main ways. The first is that we are using the same tool that we used to model human intention to model these nonlinear pneumatic systems. We show that similar to its ability to predict behavior for a very nonlinear and difficult to model human, neural networks can model the dynamics of these soft robots as well. Secondly, although the algorithms that we have proposed for modeling human intent have already been implemented on a more traditional robot, this is not the end goal. We expect that co-manipulation may be easier and more effective between a human and a soft robot. This work is a first step towards being able to better control and model soft robots for this purpose.

In addition to being good at modeling the hysteresis and nonlinearity of soft robots, neural nets have another major advantage when used for control. The neural network is trained using gradients from every input to every output, so the gradients are easily available after the neural net has been trained. There is no need to linearize the output of the neural net and this saves significant computation time.

In this work, we create a controller for the robot in Figure 4.1 which we call a “grub”. Before the work reported here, there was approximately a year of developing analytical models and tuning them on this robot to obtain good control [45]. However, our work with neural nets took two weeks of development. This shows that it is possible to develop empirical models rapidly using machine learning.

In collaboration with Charles Best and Morgan Gillespie, we developed a neural network to model the dynamics of a soft, pneumatically-actuated robot for closed-loop, model-based control.

4.1 Background

4.1.1 Robot Description

The grub was developed and built by Pneubotics, an affiliate of OtherLab. Besides electronics such as IMUs and pressure sensors, the grub is entirely made of ballistic nylon fabric with internal bladders to prevent air leakage. The structure of the grub comes from an inflatable main bladder which is pressurized to between 6.9-13.8 kPa (1-2 PSI) gauge.

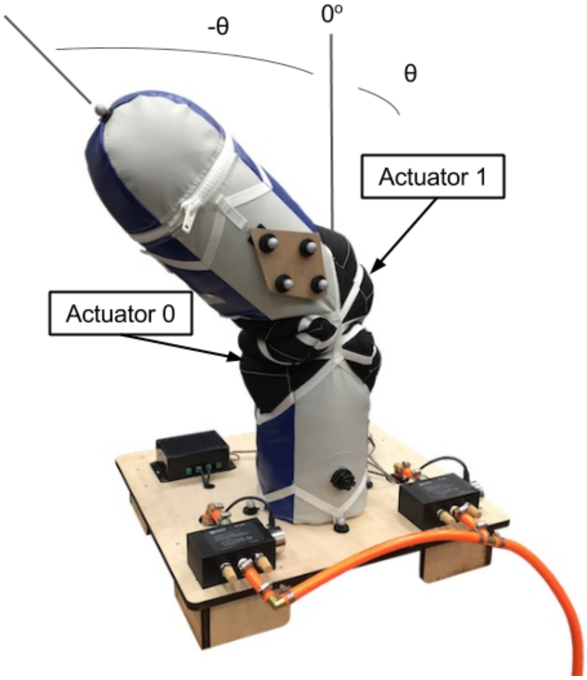

Figure 4.1: This is a single degree of freedom soft robot platform that we call a grub.

At the joint, there are two antagonistic actuators which can be filled to pressures between 0-172.4 kPa (0-25 PSI) gauge. For this research we use pressures between 0-103.4 kPa (0-15 PSI) gauge to reduce pinching effects in the main chamber at higher pressures. The source pressure is provided by an air compressor regulated to 105.5 kPa (15.3 PSI) gauge. Pressures are reported in absolute with an atmospheric pressure of 101.4 kPa (14.7 PSI).

Enfield LS-v25 five port spool proportional flow valves are used to control the variable flow of air from the pressure source to the actuation bladders or from the actuation bladders to atmosphere. This platform utilizes only one output port of the Enfield valves, effectively using the valves as three port spool valves. As seen in Figure 4.2, each actuation bladder has an individual valve for control of air flow, while both bladders share the same pressure source.

We use the Robot Operation System (ROS) to access pressure sensor data and IMU data as well as to send valve and pressure commands. Our controller code is operating in non-realtime on an Ubuntu workstation. Data for the pressure sensors can be read at approximately 1000 Hz. Pressure for each bladder is controlled by an underlying PID controller also operating at approximately 1000 Hz. Commanded pressures are published over ROS, and the valves are actuated to achieve commanded pressure values.

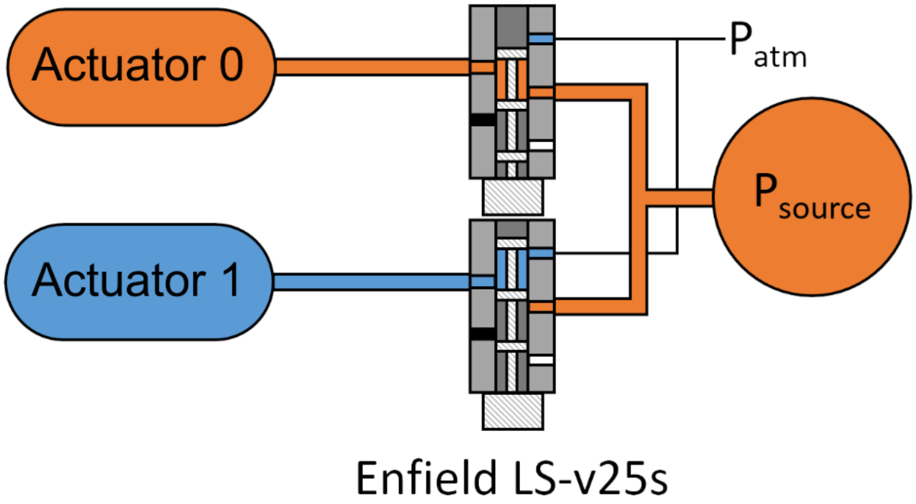

Figure 4.2: Representative figure of valve and actuator bladder configuration.

4.1.2 Model Predictive Control

In this work we use model predictive control (MPC) to control the robot. MPC is a method of control that uses a dynamic system model to calculate a discrete optimal trajectory across a finite horizon. The trajectory is calculated at every time step, but only the first control step is used in the system. The optimizer uses the system dynamics, the current state, and any constraints applied to calculate the trajectory.

Originally used in chemical engineering processes that were slower, MPC has been used recently in robotics as increased computing speeds have made that viable. We use it because it allows us to easily add state and input constraints and use our neural net model.

4.1.3 Previous Control Work

In previous work [45], a model was developed from first principles to describe the dynamics of an inflatable joint with pneumatic actuation. These models included the joint angle, angular velocity, and pressures in each actuator as states. The model was then discretized to predict future states from arbitrary inputs which was then used with MPC for effective angle tracking.

A feature of MPC is that we can replace the system model with a different model without changing the control scheme. Figure 4.3 shows the diagram for the MPC controller used in this work. The neural net in the diagram directly replaces the model developed previously.

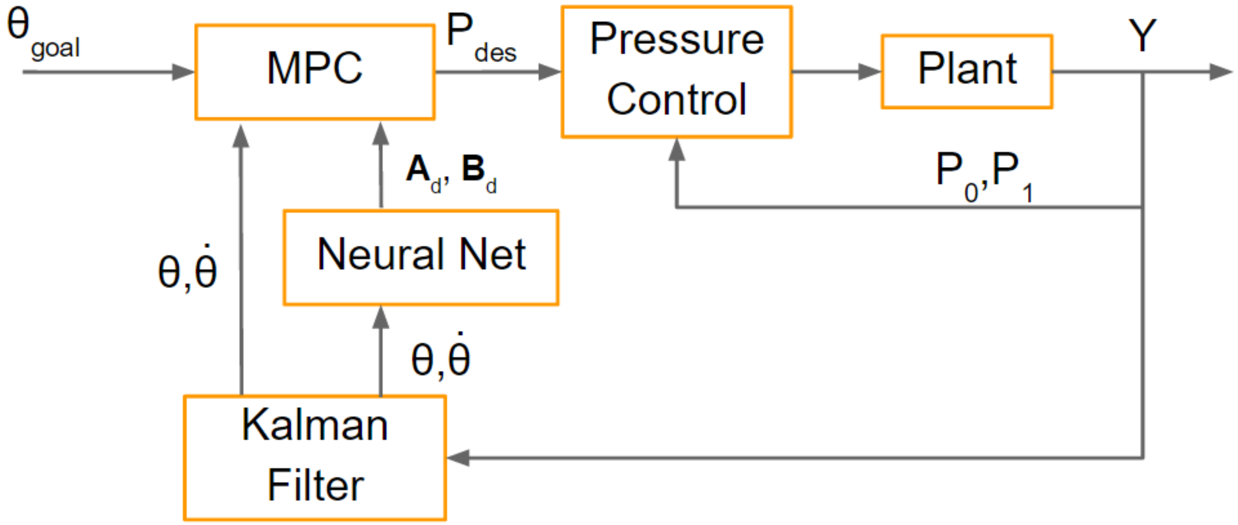

Figure 4.3: MPC flow chart showing the neural net linearization and the low level PID pressure controller

4.2 Neural Net Model of Soft Robot

The neural net was developed using Google's Tensorflow API from data collected on the grub and the data can be seen in Figure 4.4. Data were collected for 1000 seconds at 30 Hz for a series of random steps in both position and stiffness using the method developed in [46] for control. Variations in the joint stiffness were necessary so that the data included joint angles at different pressures. The inputs and outputs each were scaled by dividing by the infinity norm of that variable for the entire dataset. This decreased optimization time and improved convergence.

The neural net consisted of three hidden layers each with 200 nodes. A sigmoid activation was used on each hidden layer. The output layer had no activation which is necessary for time series prediction neural nets. This allowed the neural net to output an arbitrary continuous value. From the data, 90 percent was used for training and 10 percent for validation. The net was trained on batches of 100 units at a time until the entire training set had been used to train 1500 times for a total of 20 minutes. The length of training time was based on when the cost function on the validation data stopped decreasing. The cost function was the mean squared error of the neural net prediction versus the actual data for each of the four normalized states.

$$CostFunction = \frac{\sum\limits_{l=1}^{100} \sum\limits_{m=1}^{4} (\hat{x}_{m,l} - x_{m,l})^2 / x_{m,\infty}}{100} \tag{4.1}$$

Here, l is the index of the 100 batches and m is the index for the four states. $x_{m,\infty}$ is the norm of each state over the entire data set, $\hat{x}_{m,l}$ is the predicted value of each state, and $x_{m,l}$ is the actual value of each state.

The inputs to the neural net were the joint angle, joint velocity, pressures in both actuation chambers, and commanded pressures in both actuation chambers. The commanded pressures were the control inputs. While the outputs were the states (joint angle, joint velocity, and pressures) at the next time step. The neural net prediction of future states can be represented by the equation

$$\mathbf{x}[k+1] = f(\mathbf{x}[k], \mathbf{u}[k]), \tag{4.2}$$

where $f(\cdot)$ is the nonlinear discrete function representing the neural net, $\mathbf{x} = \left[\dot{\theta}, \theta, P_0, P_1\right]^T$, $\mathbf{u} = \left[P_{0,des}, P_{1,des}\right]^T$ and k is the time step index. $\dot{\theta}$ and θ are the angle and angular velocity of the joint, respectively. P_0 and P_1 are the pressures in each of the bladders. Because the data was collected at 30 Hz, the state prediction using Equation 4.2 automatically has a time step of 0.033 seconds. A more general neural net would include data with different time steps and the time step size would be included as an input to the neural net. This would allow the neural net to be used in controllers at different rates.

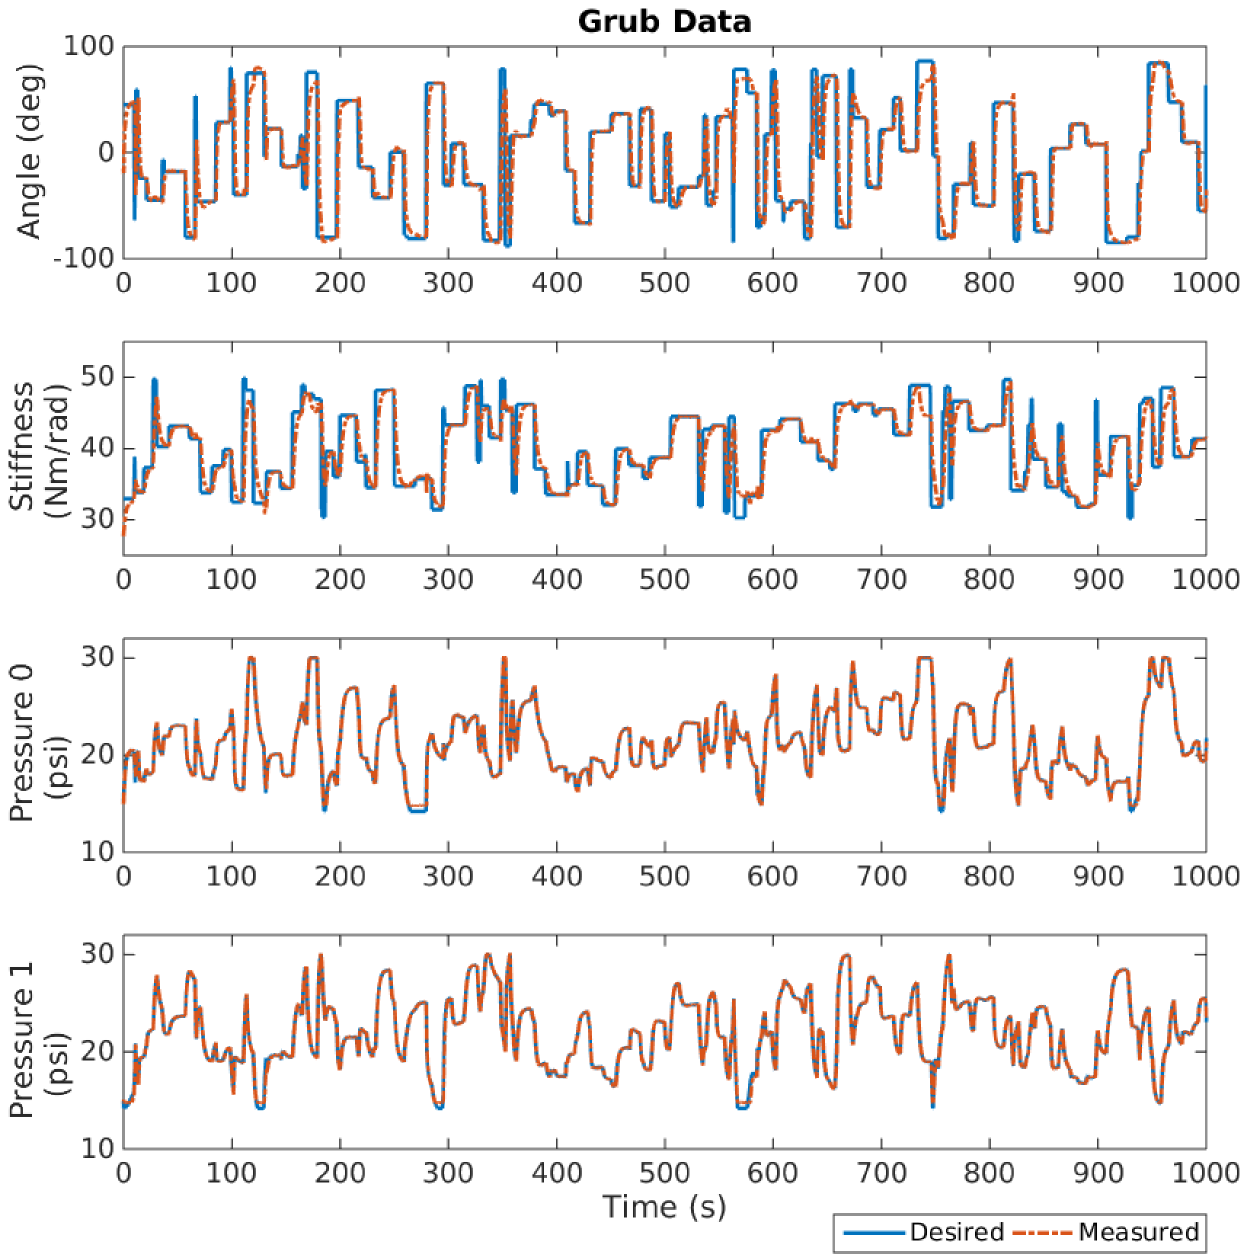

Figure 4.4: Data collected from the grub used to train the neural net.

An open-loop prediction from the neural net compared to measured data using only the the initial conditions and desired pressures can be seen in Figure 4.5. There is no feedback in these graphs as they depend only on the initial conditions and the pressure inputs. The trend of the neural net response matches the measured data for all four states making the neural net an ideal model for model predictive control. Using the neural net directly for state predictions is a feasible option for control (especially given that we have the analytical gradients of the net with respect to the states and inputs to use for optimization) but the nonlinear nature of the neural net can, depending on the optimization routine, increase solve times significantly making MPC intractable. Faster optimization techniques or a linearization of the neural net would make MPC tractable.

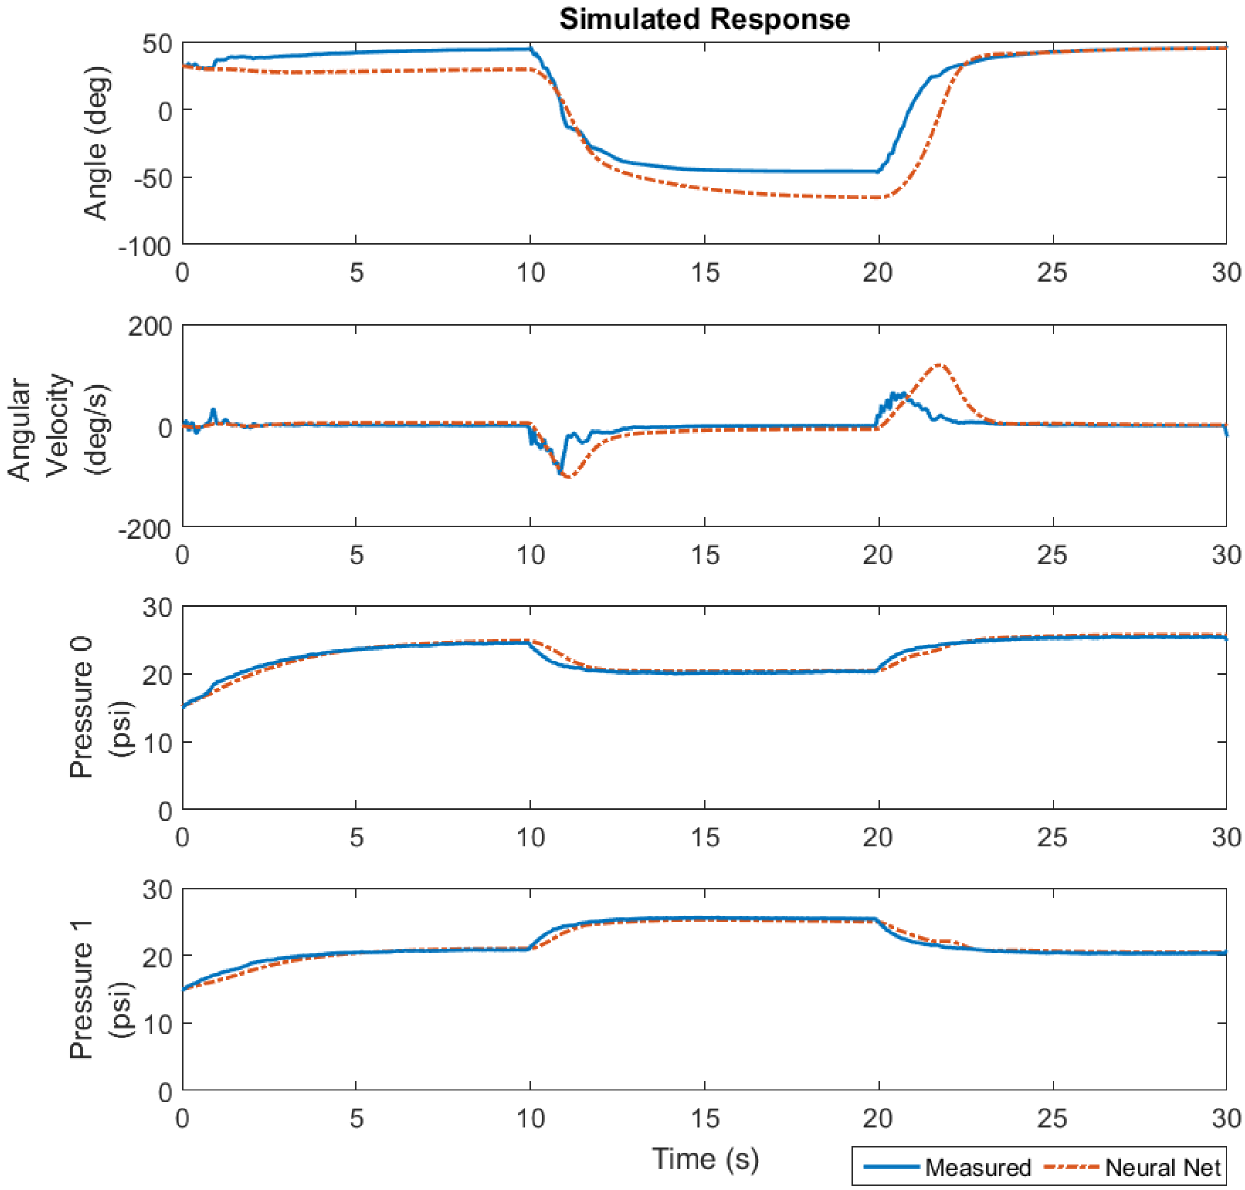

Figure 4.5: Neural net predictions compared to actual data.

Instead of using the neural net to predict states directly, we use the analytical gradient directly from the neural net. The neural net uses the gradients already in the optimization so the gradients are free. Using these, a linearized discrete state space model was formulated. The state matrix for the system becomes

$$\mathbf{A}_d = \begin{bmatrix} \frac{\partial f_1}{\partial \dot{\theta}} & \frac{\partial f_1}{\partial \theta} & \frac{\partial f_1}{\partial P_0} & \frac{\partial f_1}{\partial P_1} \\ \frac{\partial f_2}{\partial \dot{\theta}} & \frac{\partial f_2}{\partial \theta} & \frac{\partial f_2}{\partial P_0} & \frac{\partial f_2}{\partial P_1} \\ \frac{\partial f_3}{\partial \dot{\theta}} & \frac{\partial f_3}{\partial \theta} & \frac{\partial f_3}{\partial P_0} & \frac{\partial f_3}{\partial P_1} \\ \frac{\partial f_4}{\partial \dot{\theta}} & \frac{\partial f_4}{\partial \theta} & \frac{\partial f_4}{\partial P_0} & \frac{\partial f_4}{\partial P_1} \end{bmatrix}, \tag{4.3}$$

and the input matrix becomes

$$\mathbf{B}_d = \begin{bmatrix} \frac{\partial f_1}{\partial P_{0,des}} & \frac{\partial f_1}{\partial P_{1,des}} \\ \frac{\partial f_2}{\partial P_{0,des}} & \frac{\partial f_2}{\partial P_{1,des}} \\ \frac{\partial f_3}{\partial P_{0,des}} & \frac{\partial f_3}{\partial P_{1,des}} \\ \frac{\partial f_4}{\partial P_{0,des}} & \frac{\partial f_4}{\partial P_{1,des}} \end{bmatrix}. \tag{4.4}$$

Each entry in the state and input matrices come from the gradients within the neural net that relate the states and inputs. Using the gradients as a model for the system, model-based control can be applied. The gradients come from the neural net model which is already discrete. The specific discrete form of the system model becomes

$$\mathbf{x}[k+1] = \mathbf{A}_d \mathbf{x}[k] + \mathbf{B}_d \mathbf{u}[k]. \tag{4.5}$$

4.3 Control with Neural Net Model

The model predictive controller which uses the neural net as a model and is described in this section was created by Charles Best and Morgan Gillespie [45]. It is described here to compare the controller performance when using the neural net as compared to the analytical model created by them previously. The discretized matrices $\mathbf{A}_d$ and $\mathbf{B}_d$, the current states $\dot{\theta}[k]$, $\theta[k]$, $P_0[k]$, and, $P_1[k]$, the inputs $P_{0,des}[k]$ and $P_{1,des}[k]$, the final goal angle θ_{goal}, and the model constraints and weights are used by a MPC solver at 30 Hz. A flow chart for the control process can be seen in Figure 4.3.

The output from the model predictive controller is desired pressures that are maintained in each actuator by a PID pressure controller running at 1000 Hz. A Kalman filter is used to estimate the joint angle and angular velocity from the IMU accelerometer and gyro measurements.

The solver that we used for MPC was generated using CVXGEN (see [47]), a web-based tool for developing convex optimization solvers, with a horizon of five time steps. The cost function minimized across the horizon T is

$$minimize \sum_{k=1}^{T} \left(\|\theta_{goal} - \theta[k]\|^2_{Q_{pos}} + \|\dot{\theta}[k]\|^2_{Q_{vel}} + \|\Delta P_{0,des}[k]\|^2_{R_1} + \|\Delta P_{1,des}[k]\|^2_{R_2} \right), \quad (4.6)$$

subject to the system model as constraints, as defined in equation 4.5, as well as the following additional constraints:

$$|\theta| \leq \theta_{max}, \quad (4.7)$$

$$\Delta P_{i,des}[k] = P_{i,des}[k] - P_{i,des}[k-1] \qquad i = 0,1, \quad (4.8)$$

$$P_{i,des}[k] \leq P_{max} \qquad i = 0,1, \quad (4.9)$$

$$P_{min} \leq P_{i,des}[k] \qquad i = 0,1, \quad (4.10)$$

where Q_{pos}, Q_{vel}, R_1, and R_2 are scalar weights manually tuned for performance, θ_{max} is the joint limit, and P_{min} and P_{max} are atmospheric pressure and regulated source pressure respectively. In Equation 4.8, when $k = 1$, $P_{i,des}[0]$ is the control input applied at the previous time step.

An integrator was applied to the desired pressure outputs of MPC to reduce the steady state error on the angle signal. The integrator either added or subtracted from both the desired pressures to increase the pressure differential between the two actuation chambers depending on the sign of the error. To prevent integrator wind up, the integrator error only accrued when the error was less than 12° and the angular velocity was less than 2 degrees per second. The integrator was necessary in our case to obtain low steady state error. If the neural net were trained on more data, it is likely that the steady state error would be less of an issue. However, the integrator was easily applied to the low level pressure control to obtain the desired angle.

4.4 Results and Discussion

A series of 40° step angle commands ranging from -80° to 80°, changing in increments of 10 seconds was commanded. The results with an integrator can be seen in Figure 4.6. As seen from the step response, the controller is able to track the desired angle and the solver finds solutions on average at around 5000 Hz which is definitely sufficient for control at 30 Hz. The neural net was trained on data that was collected at 30 Hz. Had it been trained on data collected at a faster rate, the controller could be run at a higher rate. One of the most interesting results is that looking at the desired and actual pressure shows that the model is making full use of the multi-input system by driving the pressures in opposite directions in order to get the joint to move more quickly. This is something that we expected to see when we explicitly modeled torque as a function of the two actuation pressures in past work. However, in this case, the behavior of the system was learned automatically by the neural net and exploited by the model predictive controller. The controller does perform worse near the edge of joint limits around 80 degrees most likely because there was insufficient data from the training set in this region for the neural net. Another reason for poor performance in this region is because of the change in dynamics as the joint reaches the joint limits. At the joint limits, fabric dynamics and interactions between the actuators become increasingly significant. Including more data from this region within the neural net could alter how the net predicts future states to match the dynamics of the joint limits for the whole range of motion.

We next present a comparison between the results using this empirical learned model for control and results from Morgan Gillespie and Charles Best's past work (see [45]) that took significant modeling effort *and* significant effort to gather actual torque data from the real system. Figure 4.7 shows the empirically determined mapping from pressure in the bladders to an equilibrium joint angle, θ_e. This equilibrium angle was used in a torsional spring model to model joint torque. K_s here is an empirically determined spring constant.

$$\tau = K_s(\theta_e - \theta) \tag{4.11}$$

The control response using the neural net gradients compared to the control response using a torque model derived from first principles shows similar performance and can be seen in Figure

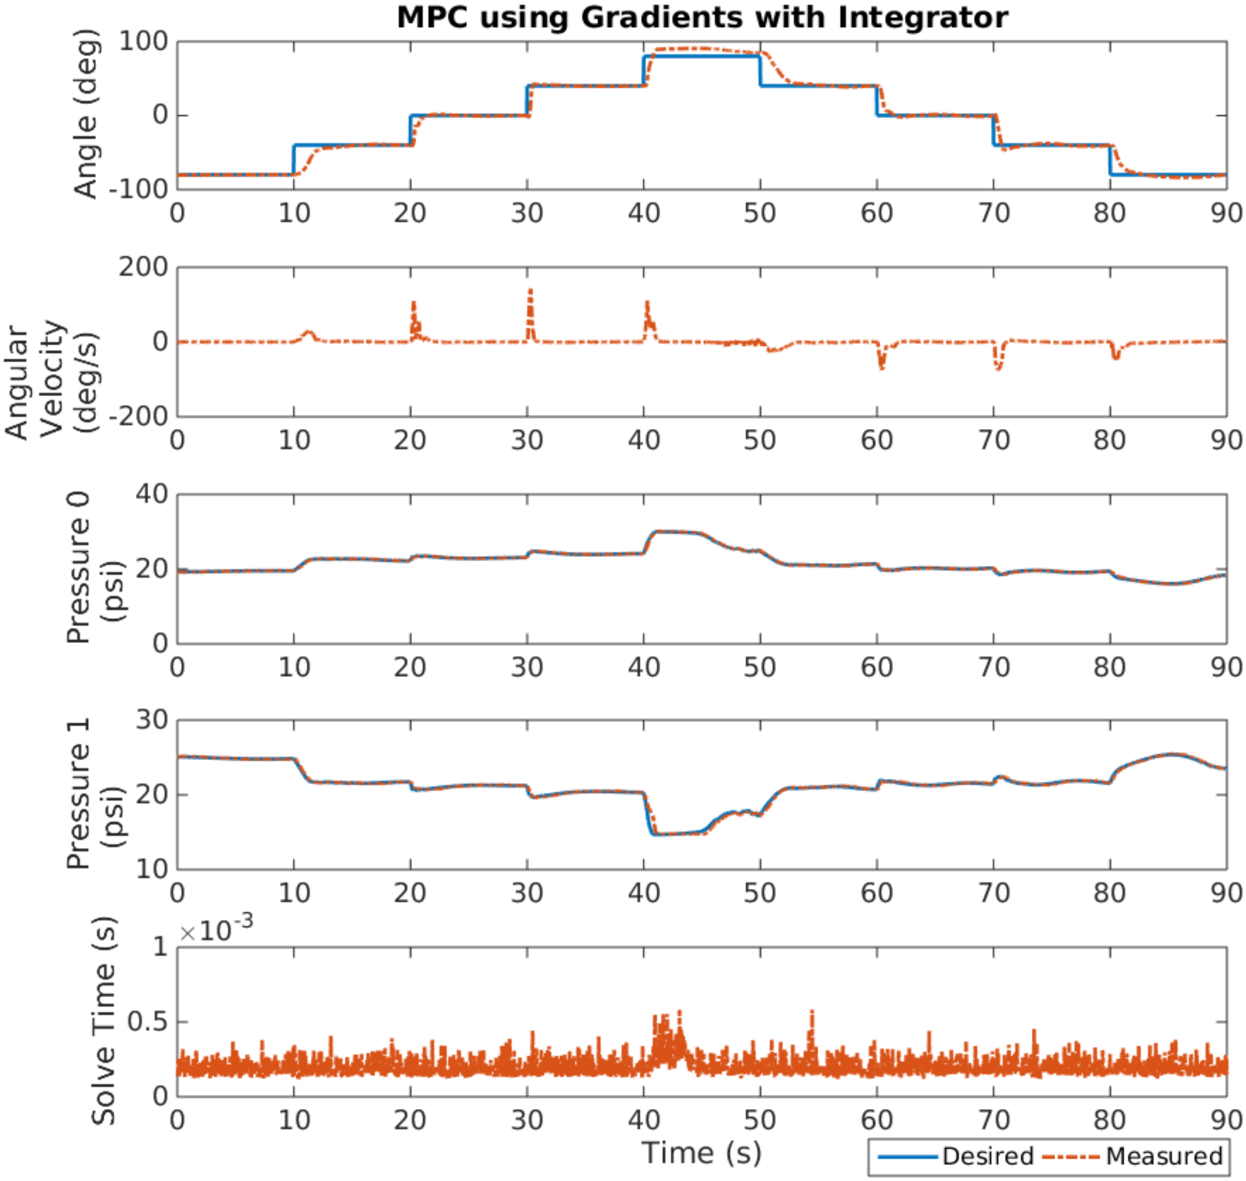

Figure 4.6: MPC with the gradients from the neural net at 30 Hz using an integrator

4.8. There is more overshoot from control using the neural net gradients and the controller using the torque model has a faster rise time. The torque model however is not robust to change in the system over time and will require the development of a new model or at least system identification for new parameters. The entire pressure to angle map in Figure 4.7 would have to be remeasured. Updating the neural net model only requires a minimal amount of new data from the system for control. This can be much more efficient than the development of a new model for each change in the system and may even be done online.

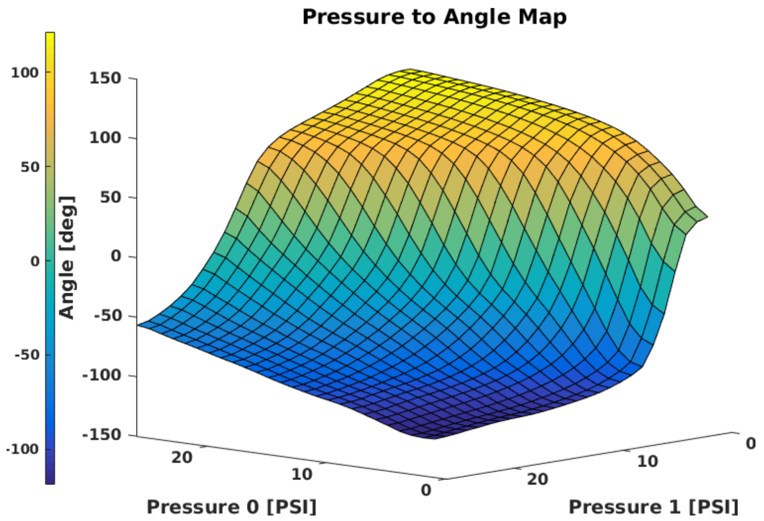

Figure 4.7: The angle of the soft robot joint based on the pressures in the bladders.

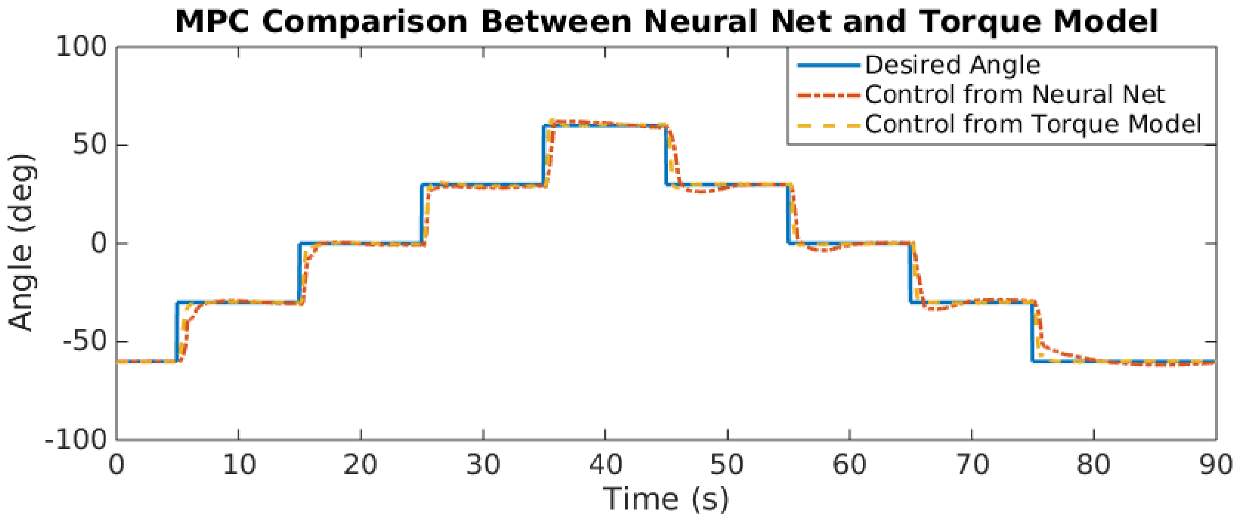

Figure 4.8: This shows how MPC using the neural net gradients compares to MPC using first principles models in past work.

4.5 Conclusion

We have shown that using a neural network with no initial knowledge about a complicated nonlinear dynamical system except for assumed state variables and inputs, we can develop a high-performing model-based controller. Additionally, we have shown that the neural network model predictive controller performs comparably to a model predictive controller with a model that took significantly more time to develop and that will be harder to adapt each time we have to deflate or repair part of the inflatable soft robot. Whereas for this controller, we can simply collect data online for a few minutes before having a new model that we can control with. Controlling more joints would require a more complicated neural net. Each joint adds more inputs and significantly more data would be needed to learn how the joints interact. It is important to collect enough data for every situation. If the robot is in a mode that has not been learned, it could be unpredictable or unstable. Neural network control makes controlling soft robots much simpler so that they can more easily be developed to be used in human-robot interaction.

CHAPTER 5. CONCLUSION AND FUTURE WORK

Here we discuss future directions that follow from this work and the contributions made to the field of physical human-robot interaction for co-manipulation by this work.

5.1 Future Work

We have only begun to use the data from the human study to understand how humans collaborate. We can predict human intent for .25 seconds. However, there is much to be learned. For example, we know that the force data gives important information to the follower on where and how fast to move, but we are still trying to understand that relationship. Perhaps each dyad learns a new system of communication. It is also possible that the humans act as impedance devices and forces are used to change the set-point of the impedance.

We have shown that velocity and angular velocity can both be predicted, but we have not yet shown how the robot will react while predicting both velocity and angular velocity at the same time. It is possible that people prefer to only rotate or translate at one time. In that case, the force information may be used to decide if translation or rotation is preferred [11]. From there, the controller would only use the prediction from either the translation or rotation neural net.

Neural nets are the method that we chose for prediction. However, there are many other methods for time series prediction like recursive least squares or support vector machines that may have better results. It would be useful to compare these to see which performs best depending on the type of task, object, data, or even robot.

It would be valuable to compare our work directly to other work that has been done in human-robot interaction. Variable impedance controllers developed by Ikeura in [18] would be a good choice.

There are several things that could be done to improve the data collected in the user study. Some of these will be discovered as we do more analysis on the results that we have. Some of the things that would be useful to study that we have already considered are:

- Changing the dimensions of the object.
- Changing the weight and weight distribution of the object.
- Using non-rigid objects.
- Add trials where the dyad can talk while the follower is blindfolded and where the dyad cannot talk when there is no blindfold.

We also desire to know what additional sensors would help us gain insight into how humans move, like cameras to perform robotic vision and pose estimation, EEG, or EMG. Specifically, knowing which insights can be gained from each sensor would allow us to know what is lost by not using that sensor and which sensors are redundant.

Finally, in our work, we assumed that the human would be the leader and the robot would be the follower. While this is an easier problem to deal with, the work we have done would also be useful in designing a controller for a robot that could lead a human during certain parts of a task. There is much work to be done in determining who should lead and who should follow in different situations and how the robot can safely and intelligently lead.

5.2 Contributions

This thesis presents the following contributions which are described in more detail below.

1. We did an exploratory study of humans co-manipulating an object with 21 dyads that will continue to be used to understand human-robot co-manipulation.
2. We developed several new insights into how humans interact haptically while co-manipulating an object.
3. We developed a neural net based model that predicts short-term human intent for 0.25 seconds into the future based entirely on past movement.

4. We have shown that the prediction can be used with a robot in the loop.
5. We have used neural nets to control soft robots that are difficult to model but we expect will be useful in human-robot interaction.

5.2.1 Exploratory Study

We did an exploratory study of human collaboration. Many similar studies have been done in the past, but ours had some distinct differences. We found no other studies that allowed the human pair to move and rotate in three dimensions except in simulation. Our study also included 21 dyads which is significantly more than the majority of studies that we observed. This allows us to see differences in different dyads and ensure that our results are somewhat general. We had the dyads perform six different tasks three times unblindfolded and three times with one person blindfolded. While it is valuable to limit degrees of freedom in some cases in order to isolate behavior, it also important to allow the dyad to work without restrictions on movement to see if past trends found in prior research or still apply and if there are new models that also apply.

5.2.2 Human Co-Manipulation Insights

Several insights came from our analysis of the exploratory study. It became clear the position data alone was sufficient to predict future positions. While we expect that force does make a difference, much can be estimated from past movement alone. From what we have learned so far, we believe that human participants act like impedance devices that change their impedance set point by predicting where they should go. Many people have been trying to minimize interaction forces between the two members of the dyad. These are the forces that are placed on the table that do not contribute to its movement. Despite past work that predicted that humans minimize this force, however, our data shows that these forces are usually present for most tasks [11]. There may be several reasons for this. By applying a force to the table, we may be more sensitive to its movement which improves communication. Some dyads may apply a constant force and then stop applying it when it is time to stop moving. Whatever the cause of these forces, they do not seem to have a negative effect until they exceed a certain threshold. In other words, conserving energy may not be the human's number one goal, because small forces do not tire us. There are a number of other goals that may be more important that include speed, stability, or effective communication.

5.2.3 Short-Term Human Intent Prediction

We have shown that we can accurately predict short-term human intent. We can predict the velocity with .02 $(m/s)^2$ mean squared error over the first 0.25 seconds. Interestingly, this is time horizon is close to the average reaction time for a human. Position data for the last 1.5 seconds is all that is required. This prediction was then run in real time with a human-robot dyad. The results were still quite good in spite of the fact that the robot is not acting in the same way that a human would. We expect that as the robot relies on the prediction for control, it will act more effectively and the prediction will be as good or better.

5.2.4 Soft Robot Control

We have shown that the same tools used to generate nonlinear models of human intention could be used to model soft robots for improved control and eventual HRI. The neural net was only trained on a long set of random inputs with the associated output state of the robot. The control was fairly good even though the robot is highly nonlinear with significant hysteresis. This is important because it means model-based control of new soft robots can be developed in a short time period. Combining this with fast design times and fast build times means that soft robots can be iterated on the order of days. These will be used extensively in human-robot interaction because they are safe around humans so this is an important part of human-robot interaction.

5.3 Conclusion

This thesis showed that short-term human intent can be predicted using only the previous velocities of the dyad, whether it be human-human for six degrees of freedom or human-robot which we have shown for three degrees of freedom. This prediction can be used in controllers that will be developed in the future that allow the robot to cooperate with the person in a useful way for physical co-manipulation tasks.

REFERENCES

[1] Doshi, A., and Trivedi, M. M., 2009. "On the roles of eye gaze and head dynamics in predicting driver's intent to change lanes." *IEEE Transactions on Intelligent Transportation Systems,* **10**(3), pp. 453–462. 2

[2] Doshi, A., Morris, B. T., and Trivedi, M. M., 2011. "On-road prediction of driver's intent with multimodal sensory cues." *IEEE Pervasive Computing,* **10**(3), pp. 22–34. 2

[3] Cheng, S. Y., and Trivedi, M. M., 2006. "Turn-intent analysis using body pose for intelligent driver assistance." *IEEE Pervasive Computing,* **5**(4), pp. 28–37. 2

[4] Jain, A., Koppula, H. S., Raghavan, B., Soh, S., and Saxena, A., 2016. "Car that knows before you do: Anticipating maneuvers via learning temporal driving models." *Proceedings of the IEEE International Conference on Computer Vision,* **11-18-Dece**, pp. 3182–3190. 2

[5] Ojakangas, C. L., Shaikhouni, A., Friehs, G. M., Caplan, A. H., Serruya, M. D., Saleh, M., Morris, D. S., and Donoghue, J. P., 2006. "Decoding movement intent from human premotor cortex neurons for neural prosthetic applications.." *Journal of clinical neurophysiology : official publication of the American Electroencephalographic Society,* **23**(6), pp. 577–84. 2

[6] Mainprice, J., Hayne, R., and Berenson, D., 2015. "Predicting Human Reaching Motion in Collaborative Tasks Using Inverse Optimal Control and Iterative Re-planning." *International Conference on Robotics and Automation*, pp. 885 – 892. 2

[7] Perez-D'Arpino, C., and Shah, J. A., 2015. "Fast Target Prediction of Human Reaching Motion for Cooperative Human-Robot Manipulation Tasks using Time Series Classification." *International Conference on Robotics and Automation*, pp. 6175–6182. 2

[8] Nehaniv, C. L., Dautenhahn, K., Kubacki, J., Haegele, M., Parlitz, C., and Alami, R., 2005. "A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction." *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication,* **2005**, pp. 371–377. 2

[9] Unhelkar, V. V., Stirling, L., and Shah, J. A., 2015. "Human-Robot Co-Navigation using Anticipatory Indicators of Human Walking Motion." *International Conference on Robotics and Automation*, pp. 6183 – 6190. 2

[10] Mainprice, J., and Berenson, D., 2013. "Human-robot collaborative manipulation planning using early prediction of human motion." *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 299–306. 2

[11] Mielke, E. A., Townsend, E. C., and Killpack, M. D., 2017. "Analysis of Rigid Extended Object Co-Manipulation by Human Dyads: Lateral Movement Characterization." *ArXiv e-prints*. 4, 60, 62

[12] Pratt, G. A., and Williamson, M. M., 1995. Series elastic actuators. 7

[13] Kazerooni, H., 1988. "Human Machine Interaction via the Transfer of Power and Information Signals.". 7

[14] Hogan, N., 1984. "Impedance Control: An Approach to Manipulation." *IEEE American Control Conference*, pp. 304–313. 8

[15] Kosuge, K., Yoshida, H., and Fukuda, T., 1993. "Dynamic control for robot-human collaboration." *Proceedings of 1993 2nd IEEE International Workshop on Robot and Human Communication*, pp. 398–401. 8, 12

[16] Ikeura, R., Monden, H., and Inooka, H., 1994. "Cooperative motion control of a robot and a human." *Proceedings of 1994 3rd IEEE International Workshop on Robot and Human Communication*, pp. 2–3. 8

[17] Ikeura, R., and Inooka, H., 1995. "Variable impedance control of a robot for cooperation with a human." *Proceedings - IEEE International Conference on Robotics and Automation,* **3**, pp. 3097–3102. 8

[18] Ikeura, R., Morita, a., and Mizutani, K., 1997. Variable damping characteristics in carrying an object by two humans. 8, 60

[19] Tsumugiwa, T., Yokogawa, R., and Hara, K., 2002. "Variable impedance control based on estimation of human arm stiffness for human-robot cooperative calligraphic task." *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292),* **1**(May), pp. 644–650. 8

[20] Duchaine, V., and Gosselin, C., 2009. "Safe, Stable and Intuitive Control for Physical Human-Robot Interaction." *2009 IEEE International Conference on Robotics and Automation*, pp. 3383–3388. 8

[21] Yang, C., Ganesh, G., Haddadin, S., Parusel, S., Albu-Schäeffer, A., and Burdet, E., 2011. "Human-like adaptation of force and impedance in stable and unstable interactions." *IEEE Transactions on Robotics,* **27**(5), pp. 918–930. 8

[22] Kadiallah, A., Franklin, D. W., and Burdet, E., 2012. "Generalization in Adaptation to Stable and Unstable Dynamics." *PLoS ONE,* **7**(10). 8

[23] Ganesh, G., Albu-Schäffer, A., Haruno, M., Kawato, M., and Burdet, E., 2010. "Biomimetic motor behavior for simultaneous adaptation of force, impedance and trajectory in interaction tasks." *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 2705–2711. 8

[24] Ito, S., Yuasa, H., Ito, M., and Hosoe, S., 1999. "On an adaptation in distributed system based on a gradient dynamics." *IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.99CH37028),* **1**(1), pp. 200–205. 9

[25] Flash, T., and Hogan, N., 1985. “The coordination of arm movements: an experimentally confirmed mathematical model..” *The Journal of neuroscience : the official journal of the Society for Neuroscience,* **5**(7), pp. 1688–1703. 9

[26] Corteville, B., Aertbelien, E., Bruyninckx, H., De Schutter, J., and Van Brussel, H., 2007. “Human-inspired robot assistant for fast point-to-point movements.” *Proceedings - IEEE International Conference on Robotics and Automation*(April), pp. 3639–3644. 9

[27] Ikeura, R., and Mizutani, K., 1998. “Control of robot cooperating with human motion.” In *Proceedings of 1998 IEEE International Workshop on Robotics and Human Communication*, pp. 525–529. 9

[28] Maeda, Y., Hara, T., and Arai, T., 2001. “Human-robot cooperative manipulation with motion estimation.” *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180),* **4**, pp. 2240–2245. 9

[29] Miossec, S., and Kheddar, A., 2008. “Human motion in cooperative tasks: Moving object case study.” *2008 IEEE International Conference on Robotics and Biomimetics, ROBIO 2008*(i), pp. 1509–1514. 9

[30] Bussy, A., Gergondet, P., Kheddar, A., Keith, F., and Crosnier, A., 2012. “Proactive behavior of a humanoid robot in a haptic transportation task with a human partner.” *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*(2), pp. 962–967. 9

[31] Evrard, P., Gribovskaya, E., Calinon, S., Billard, A., and Kheddar, A., 2009. “Teaching physical collaborative tasks: Object-lifting case study with a humanoid.” *9th IEEE-RAS International Conference on Humanoid Robots, HUMANOIDS09*, pp. 399–404. 9

[32] Rozo, L., Bruno, D., Calinon, S., and Caldwell, D. G., 2016. “Learning Optimal Controllers in Human-robot Cooperative Transportation Tasks with Position and Force Constraints.” *IEEE/RSJ Internation Conference on Intelligent Robots and Systems*, pp. 1024–1030. 9

[33] Agravante, D. J., Cherubini, A., Bussy, A., Gergondet, P., and Kheddar, A., 2014. “Collaborative Human-Humanoid Carrying Using Vision and Haptic Sensing.” *International Conference on Robotics and Automation*. 10

[34] Choi, K., Hirose, H., Iijima, T., and Koike, Y., 2005. “Prediction of four degrees of freedom arm movement using EMG signal.” *Engineering in Medicine and Biology 27th Annual Conference*(2005), pp. 5820–5823. 10

[35] Dragan, A. D., Bauman, S., Forlizzi, J., and Srinivasa, S. S. “Effects of Robot Motion on Human-Robot Collaboration.” pp. 51–58. 10

[36] Chipalkatty, R., Member, S., Droge, G., and Member, S. “Less Is More : Mixed Initiative Model Predictive Control With Human Inputs.” pp. 1–9. 10, 27

[37] Kheddar, A., 2011. “Human-robot haptic joint actions is an equal control-sharing approach possible?.” *4th International Conference on Human System Interaction, HSI 2011*, pp. 268–273. 12

[38] Groten, R. K., 2011. “Haptic Human-Robot Collaboration : How to Learn from Human Dyads.”. 12

[39] Reed, K. B., Peshkin, M., Hartmann, M. J., Patton, J., Vishton, P. M., and Grabowecky, M., 2006. “Haptic cooperation between people, and between people and machines.” *IEEE International Conference on Intelligent Robots and Systems*, pp. 2109–2114. 12, 28

[40] Feth, D., Groten, R., Peer, A., Hirche, S., and Buss, M., 2009. “Performance related energy exchange in haptic human-human interaction in a shared virtual object manipulation task.” *Proceedings - 3rd Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, World Haptics 2009*, pp. 338–343. 12

[41] GoogleResearch, 2015. “TensorFlow: Large-scale machine learning on heterogeneous systems.”. 25

[42] Engel, Y., Mannor, S., and Meir, R., 2004. “The kernel recursive least squares algorithm.” *IEEE Transactions on Signal Processing,* **52**(8), pp. 2275–2285. 27, 30

[43] Glorot, X., and Bengio, Y., 2010. “Understanding the difficulty of training deep feedforward neural networks.” *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS),* **9**, pp. 249–256. 29

[44] Evrard, P., and Kheddar, A., 2009. “Homotopy switching model for dyad haptic interaction in physical collaborative tasks.” *Proceedings - 3rd Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, World Haptics 2009*, pp. 45–50. 45

[45] Best, C. M., Gillespie, M. T., Hyatt, P., Killpack, M. D., Rupert, L., and Sherrod, V., 2016. “Model Predictive Control for Pneumatically Actuated Soft Robots.” *IEEE Robotics & Automation Magazine,* **3**(9), p. 31. 47, 50, 54, 56

[46] Best, C. M., 2016. “Position and Stiffness Control of Inflatable Robotic Links Using Rotary Pneumatic Actuation.”. 51

[47] Mattingley, J., and Boyd, S., 2012. “CVXGEN: a code generator for embedded convex optimization.” *Optimization and Engineering,* **13**(1), pp. 1–27. 55

APPENDIX A. IRB HUMAN STUDY MATERIALS

A.1 Recruiting Materials

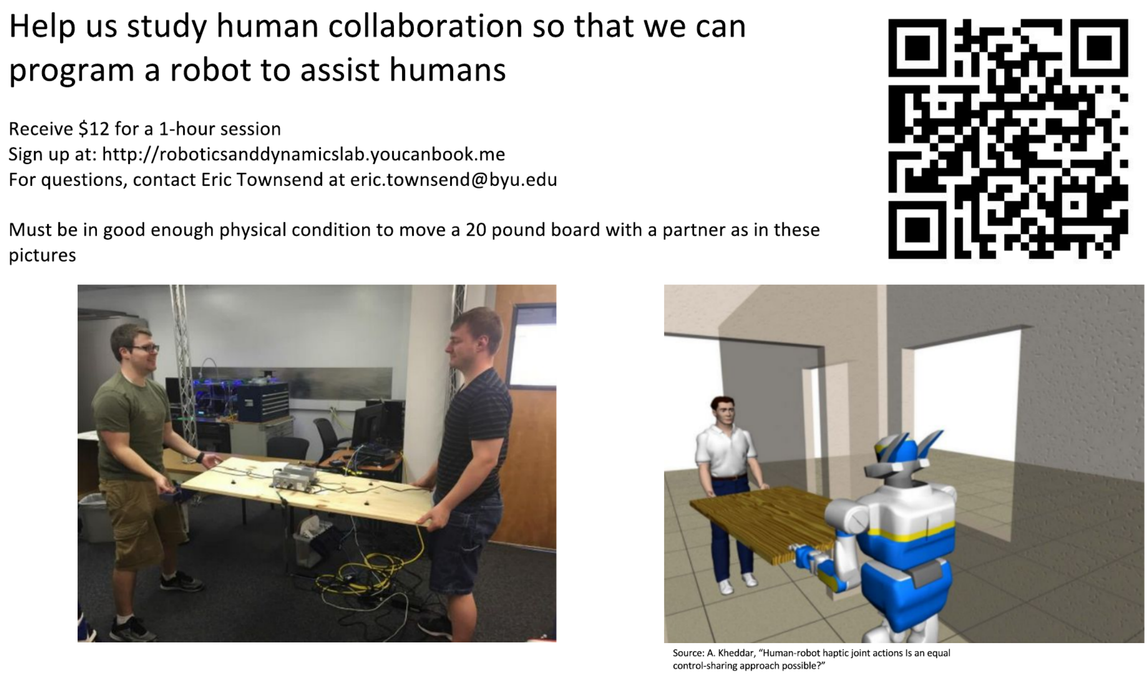

Help us study human collaboration so that we can program a robot to assist humans

Receive $12 for a 1-hour session
Sign up at: http://roboticsanddynamicslab.youcanbook.me
For questions, contact Eric Townsend at eric.townsend@byu.edu

Must be in good enough physical condition to move a 20 pound board with a partner as in these pictures

Source: A. Kheddar, "Human-robot haptic joint actions Is an equal control-sharing approach possible?"

BYU Institutional Review Board 10-14-2015 Approved 10-13-2016 Expires

Help us study human collaboration so that we can program a robot to assist humans

Receive $12 for a 1-hour session

Sign up at: http://roboticsanddynamicslab.youcanbook.me

For questions, contact Eric Townsend at eric.townsend@byu.edu

Must be in good enough physical condition to move a 20 pound board with a partner as in these pictures

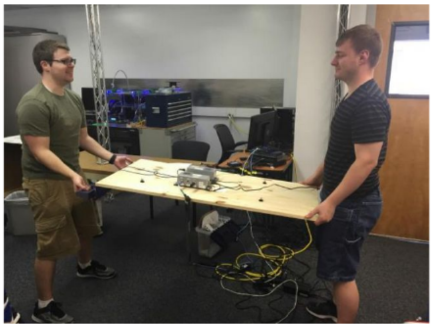

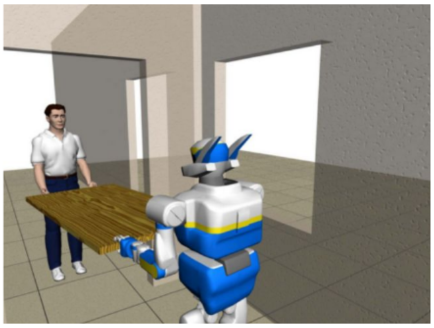

Source: A. Kheddar, "Human-robot haptic joint actions Is an equal control-sharing approach possible?"

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study
Sign up at http://roboticsanddynamicslab.youcanbook.me
For questions, email eric.townsend@byu.edu

Human-Robot study

A.2 Informed Consent and Releases

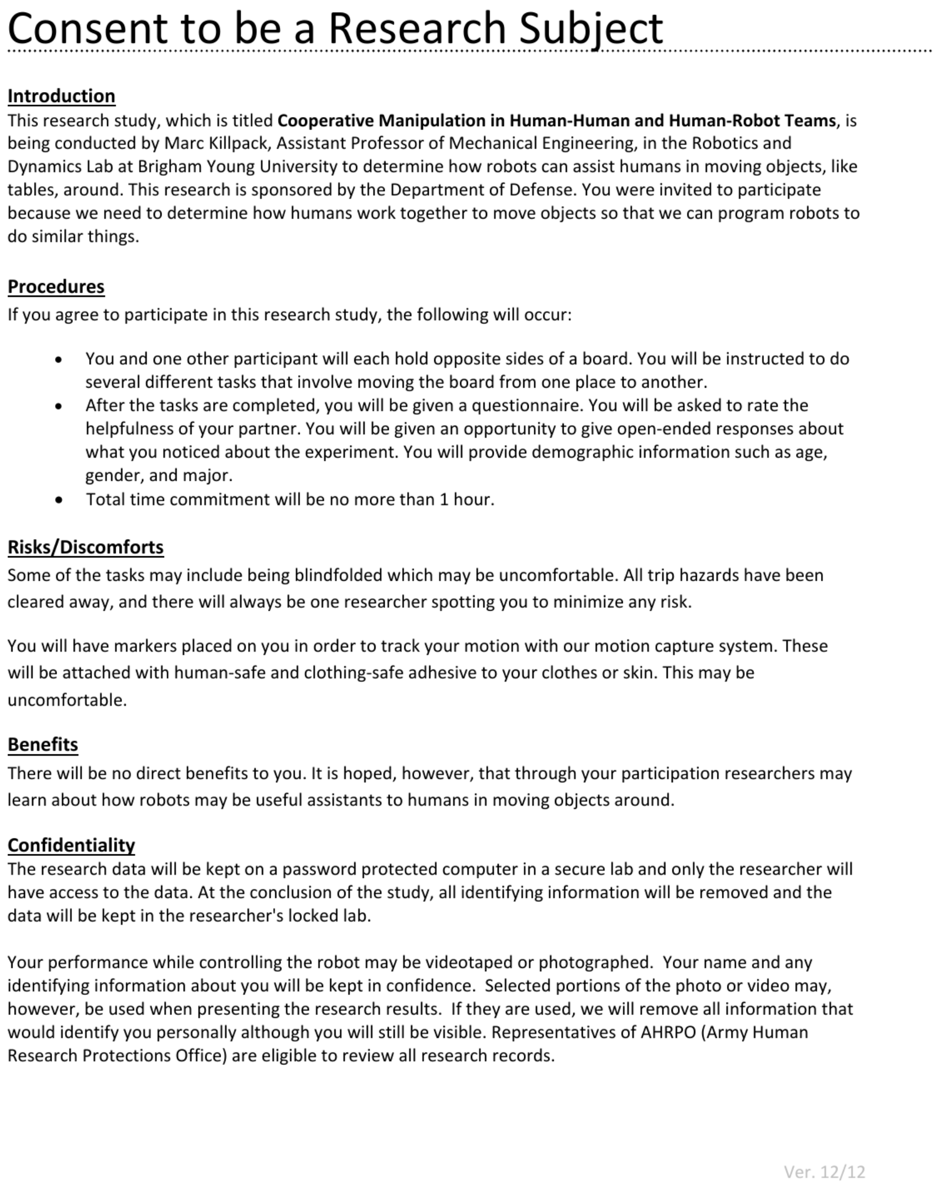

Consent to be a Research Subject

Introduction

This research study, which is titled **Cooperative Manipulation in Human-Human and Human-Robot Teams**, is being conducted by Marc Killpack, Assistant Professor of Mechanical Engineering, in the Robotics and Dynamics Lab at Brigham Young University to determine how robots can assist humans in moving objects, like tables, around. This research is sponsored by the Department of Defense. You were invited to participate because we need to determine how humans work together to move objects so that we can program robots to do similar things.

Procedures

If you agree to participate in this research study, the following will occur:

- You and one other participant will each hold opposite sides of a board. You will be instructed to do several different tasks that involve moving the board from one place to another.
- After the tasks are completed, you will be given a questionnaire. You will be asked to rate the helpfulness of your partner. You will be given an opportunity to give open-ended responses about what you noticed about the experiment. You will provide demographic information such as age, gender, and major.
- Total time commitment will be no more than 1 hour.

Risks/Discomforts

Some of the tasks may include being blindfolded which may be uncomfortable. All trip hazards have been cleared away, and there will always be one researcher spotting you to minimize any risk.

You will have markers placed on you in order to track your motion with our motion capture system. These will be attached with human-safe and clothing-safe adhesive to your clothes or skin. This may be uncomfortable.

Benefits

There will be no direct benefits to you. It is hoped, however, that through your participation researchers may learn about how robots may be useful assistants to humans in moving objects around.

Confidentiality

The research data will be kept on a password protected computer in a secure lab and only the researcher will have access to the data. At the conclusion of the study, all identifying information will be removed and the data will be kept in the researcher's locked lab.

Your performance while controlling the robot may be videotaped or photographed. Your name and any identifying information about you will be kept in confidence. Selected portions of the photo or video may, however, be used when presenting the research results. If they are used, we will remove all information that would identify you personally although you will still be visible. Representatives of AHRPO (Army Human Research Protections Office) are eligible to review all research records.

Ver. 12/12

Compensation

You will receive $12 for your participation as soon as your appointment is over; compensation will not be prorated.

Participation

Participation in this research study is voluntary. You have the right to withdraw at any time or refuse to participate entirely without jeopardy to your class status, grade, or standing with the university.

Questions about the Research

If you have questions regarding this study, you may contact Eric Townsend at eric.townsend@byu.edu or (360)202-8989 for further information. To contact the principal investigator, Marc Killpack, email marc_killpack@byu.edu or call 801-422-6342.

Questions about Your Rights as Research Participants

If you have questions regarding your rights as a research participant contact IRB Administrator at (801) 422-1461; A-285 ASB, Brigham Young University, Provo, UT 84602; irb@byu.edu.

Statement of Consent

I have read, understood, and received a copy of the above consent and desire of my own free will to participate in this study.

Name (Printed): ______________________ Signature: _________________________ Date: ____________

Witness (Printed): ____________________ Signature: _________________________ Date: ____________

Photographic Release Form

As part of this project, I will be taking photographs of you during your participation in the research. Please indicate what uses of these photographs you are willing to permit, by initialing next to the uses you agree to and signing at the end. This choice is completely up to you. I will only use the photographs in the ways that you agree to. In any use of the photographs, you will not be identified by name.

_______ Photographs can be reviewed by the research team.

_______ Photographs can be used for project illustrations.

_______ Photographs can be used for classroom presentations.

_______ Photographs can be used for academic conference presentations.

_______ Photographs can be used for fundraising presentations/proposals.

_______ Photographs can be used for newspaper or magazine publication

_______ Photographs can be posted to a website.

I have read the above descriptions and give my express written consent for the use of the photographs as indicated by my initials above.

Name (Printed): ______________________ Signature __________________________ Date: _____________

Ver. 12/12

Video Release Form

As part of this project, I will be making video recordings of you during your participation in the research. Please indicate what uses of this video you are willing to permit, by initialing next to the uses you agree to and signing at the end. This choice is completely up to you. I will only use the video in the ways that you agree to. In any use of the video, you will not be identified by name.

_______ Video can be studied by the research team for use in the research project.

_______ Video can be used for scientific publications.

_______ Video can be shown at scientific conferences or meetings.

_______ Video can be shown in classrooms to (elementary/middle/high school/college) students.

_______ Video can be shown in public presentations to non-scientific groups.

_______ Video can be used on television or the audio portion can be used on radio.

_______ Video can be posted to a website (i.e. YouTube)

I have read the above descriptions and give my express written consent for the use of the video as indicated by my initials above.

Name (Printed): _______________________ Signature ___________________________ Date: ______________

Ver. 12/12

A.3 Questionnaire

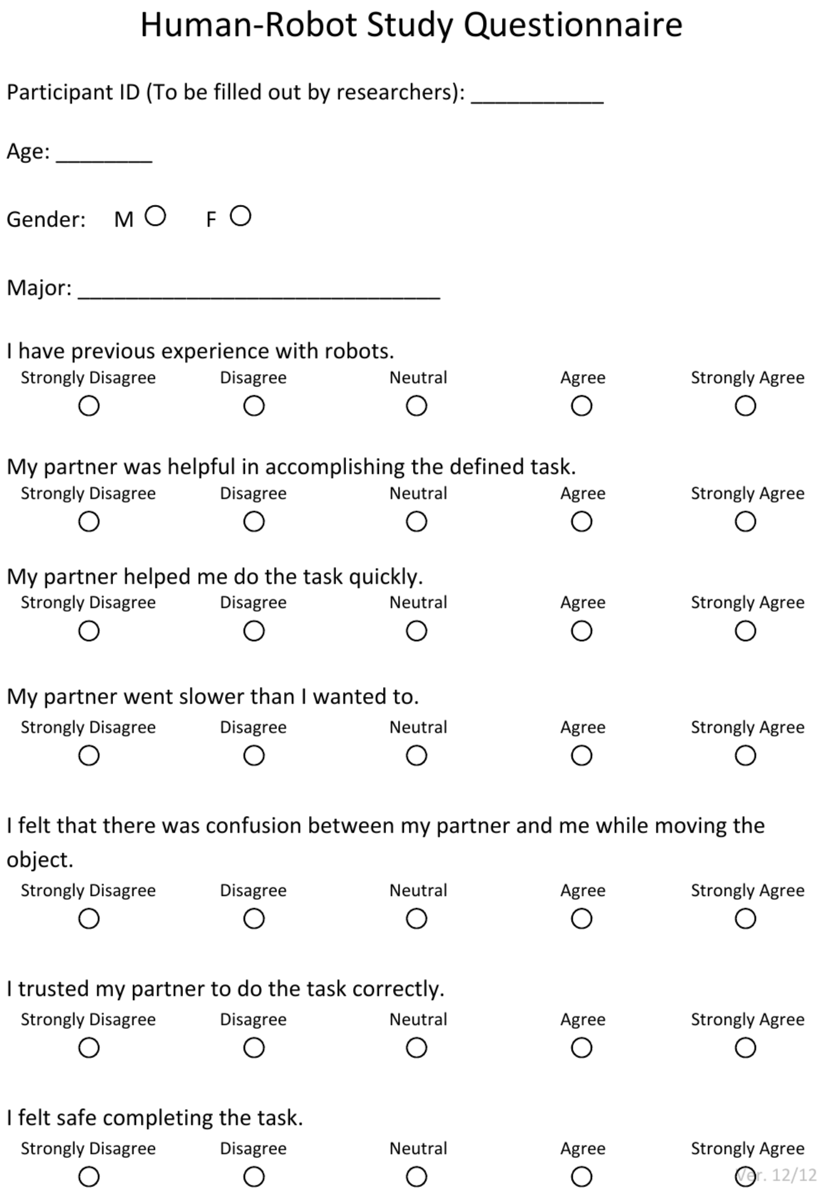

Human-Robot Study Questionnaire

Participant ID (To be filled out by researchers): ___________

Age: ________

Gender: M ○ F ○

Major: ______________________________

I have previous experience with robots.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

My partner was helpful in accomplishing the defined task.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

My partner helped me do the task quickly.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

My partner went slower than I wanted to.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

I felt that there was confusion between my partner and me while moving the object.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

I trusted my partner to do the task correctly.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

I felt safe completing the task.

Strongly Disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly Agree ○

Ver. 12/12

I trusted my partner to move at appropriate speeds.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

My partner did not push or pull too hard.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

My partner moved in a predictable way.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

If your partner was a human assigned as a follower:
I felt like I could complete the task as effectively when they were blindfolded.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

My partner helped me do the task better than I could by myself.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

My partner equally shared the task.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

I consider myself to be assertive.

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
○	○	○	○	○

Please leave any comments you have about your experience today.

____

____

____