Contributions of the Human Operator to Supernumerary Robotic Limbs

by

Jacob William Guggenheim

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2020

Author ..... Signature redacted .....
Department of Mechanical Engineering
Oct 11, 2019

Certified by ..... Signature redacted .....
H. Harry Asada
Ford Professor of Engineering
Thesis Supervisor

Accepted by ..... Signature redacted .....
Nicolas G. Hadjiconstantinou
Chairman, Department Committee on Graduate Theses

Contributions of the Human Operator to Supernumerary Robotic Limbs

by

Jacob William Guggenheim

Submitted to the Department of Mechanical Engineering
on Oct 11, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

An expanding literature base has applied Supernumerary Robotic Limbs (Superlimbs) to fields as diverse as heavy industry, robotic surgery, and assistive technology. While the list of applications has grown, and the designs have become more diverse, the research community has focused almost exclusively on the robotic system's role in augmenting the human's capabilities. This represents only one side of the issue; little research has explored the role of the human operator. This thesis represents the first in-depth exploration of the human's contributions to the Superlimb-human system.

We began by examining the control strategy of Superlimbs by asking whether fully manual control of the Superlimbs was viable when the human operator was asked to perform simultaneous and independent tasks with both their robotic and natural limbs. Although we found that the human operator was able to control all four limbs–two robotic, two natural–simultaneously, we found that the human operator performed worse with their natural limbs when controlling all four limbs as compared to when the human operator was only controlling their natural limbs. Thus, when designing Superlimbs for a taskset that requires the human and the robot to perform simultaneous independent tasks, this study points to the need for reducing the number of Superlimb degrees of freedom (DOFs) the human must manually control either through design or control.

In order to achieve this reduction, we next exploited the high redundancy and flexibility of the human body. First, we proposed a methodology for reduced-actuator Superlimbs by exploiting the human operators' ability to manipulate the base of the Superlimb. Based upon this methodology, we realized a lightweight Superlimb that could assist a human operator by opening a door when the human operator's hands are busy. Second, we proposed a novel control input methodology for communicating a rich variety of commands to the Superlimbs while both hands are busy. Based upon this methodology, and in combination with an intermittent control structure, we controlled the reduced-actuator Superlimb described above with action primitives to assist a human operator by opening a door when the human operator was holding a large box.

Finally, as the Superlimb's state changes, that change is reflected as a change in the forces and torques felt by the human operator at the base of the Superlimb. We found that this inherent haptic feedback allowed the operator to both perform closed-loop manually control of the force output of a Superlimb and to supervise the autonomous actions of a Superlimb.

In sum, this thesis explores how Superlimbs can be designed to exploit the benefits while limiting the challenges of being attached to a human operator.

Thesis Supervisor: H. Harry Asada
Title: Ford Professor of Engineering

Acknowledgments

As this acknowledgments section will not be sufficient, I'd like to offer heartfelt thank you to everyone who throughout my time at MIT for making it a thoroughly enjoyable, challenging, and growth-filled time in my life. I hope and intend to make my appreciation felt through my actions rather than the words here. Given that, I'll do my best with words for a few specific thank yous.

I suppose we'll go in (roughly) chronological order here. So first things first, I'd like to thank my family. Aaron, thank you for always being available and always offering help. Mom, thank you for giving me perspective and allowing me to see the important things in life. I truly believe graduate school has been more relaxed and enjoyable because of your influence. Thank you Dad for always listening and offering sound advice. You've helped me immensely in navigating graduate school successfully. Joanie, thank you for tag teaming with my dad in supporting me throughout graduate school. You've been instrumental in helping me approach graduate school with a level head. Jeremy, your sincere curiosity in my work has helped me to see it with a fresh set of eyes. Ina and Peter, thank you setting such great examples (and high bars) for me to follow. Amy, Ed, and Maddie, I've uniformly looked forward to your visits and calls. Your excitement is irresistibly refreshing and encouraging. Ann, thank you for helping me navigate some pretty complex issues. Sheina, Nate, Joel, and Howard, I've loved seeing you all throughout my time here and am looking forward to finally being able to take August off to come hang out at Shelter Island. Finally, thank you Grandma Janet and Grandpa Bernie. You both have offered so much love and support throughout my life.

Isaac and Manu, I cannot imagine better roommates or friends. While I fully expected living with you both to be fun, there is a world of difference between acceptable and extraordinary and you both helped create the extraordinary. Thank you.

Professor Asada, your mentorship and guidance throughout my time at MIT has been invaluable. You've been vital to my growth as a researcher, always pushing me

to think deeper and more critically. Thank you.

To my labmates and friends at MIT, thank you. Federico, Faye, Hyeonyu, Min-Cheol, thank you for setting a good example and creating a great lab culture. Meeting you all was instrumental in me choosing the lab. Vincent, thank you for your guidance and expertise during my first year or so. Dgonz, thank you for always being generous with your time. Nick, David, Nithin, Rushina, thank you for carrying me through classes and for remaining great friends. Rachel, you've been the heart of this lab for as long as you've been in it and we'll all be the worse for it when you leave. Abbas, Nick, Filippos, you three have striven for and created constant progress in our lab. Jerry, thank you for always being the guinea pig, though perhaps it would've been better if you didn't always do so well during the tests and set too high of a bar for everyone else...just a thought. Katie and Zoe, your joy and enthusiasm was infectious. Your choices have reshaped the lab for the better. Phillip, Ryan, Tongxi, John, you all joining and excelling has validated the culture our lab has created. It's been an absolute pleasure hanging out with you all. Yuta, Kosuke, Seiichi, Haruhiko, you all have been among the best engineers I've ever met. Thank you for showing me how to approach graduate school with professionalism. Yvette, thank you for helping the lab run smoothly. Thank you to the GAME tennis, soccer, ultimate, volleyball, and softball teams. It's been a pleasure competing (and winning!) with all of you.

To my committee members, Professor Hogan and Dr. Jones, thank you for pushing me to better present my work and offering insightful advice to improve it. I'll carry the things I've learned forward into my future work.

To my new roommates, Molly, Evan, and Maeve. While it's hard to compete with Isaac and Manu, you three have done an incredible job. Thank you for stepping in and providing as welcoming and supportive an apartment as one could hope.

And lastly, thank you Yi. Of all the people for whom I cannot adequately express my gratitude in words, you solidly occupy the front of that list. While I could speak to your great qualities, this is about me, so I'll say you make me want to be better so that I can live up to the high standard you expect and deserve. I hope my actions will continue to make evident my gratitude.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The first paper on Supernumerary Robotic Limbs, known as SRLs or Superlimbs for short, was written in 2012. It introduced Superlimbs as a new type of wearable co-robot that can perform tasks in close coordination with the human operator and demonstrated the Superlimb helping a human perform a drilling task after learning through observation [27].

Llorens et al. argued that Superlimbs could extend the range of available skills and manipulation possibilities beyond what is currently possible with other wearable robotic systems [27]. This is because Superlimbs differ fundamentally from other wearable robotic systems like exoskeletons and prosthetics. Active prosthetics function in series with the existing human body to compensate for lost limb function. Exoskeletons function in parallel with the existing human body to increase joint torques and muscle strength with external actuators or support a load with an elastic structure that stores energy and/or redistributes the load across the body [7] [16]. Importantly, neither of these two wearable robotic systems add additional degrees of freedom to the human body.

Rather than mimicking the structure and kinematics of the human body, Superlimbs are a type of wearable robot that extend the human users' capabilities with extra robotic limbs attached to the human. Because they neither operate in series nor in parallel with the human's natural limbs, Superlimbs have the potential to broaden the functions a single human user can achieve [36] [56].

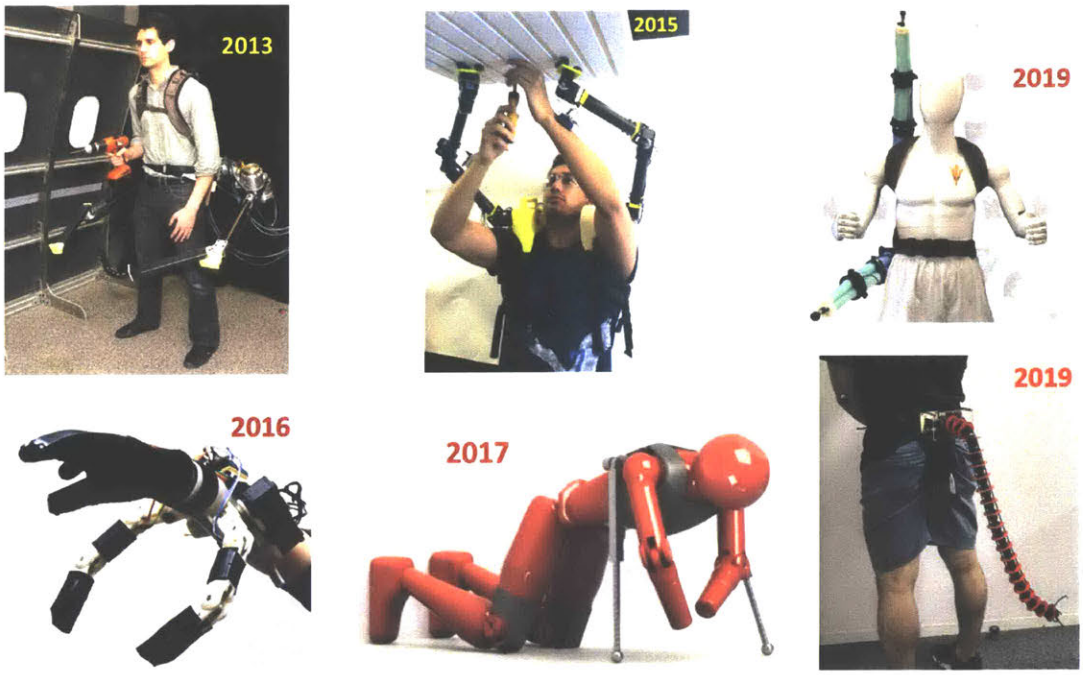

Figure 1-1: Examples of Superlimbs: Over the years, the Superlimb research community has explored different designs and placements. Images for the figure taken from [37] [8] [56] [25] [31] [33]

Since the initial publication in 2012, significant progress has been made both on the design and control front as shown in Fig. 1-1. In 2013, Parietti et al. argued that control performance of Superlimbs could be hindered by unpredictable motions of the human operator. Thus, they used dynamic analysis and state estimation to reject the disturbances caused by the human [37]. In 2014 and 2015, Superlimbs moved from being exclusively waist mounted to being mounted on other parts of the body, including the shoulder [8] [50]. New potential applications of Superlimbs were also identified, including having the Superlimb brace the body of a heavy industry worker and having the Superlimb function as an assistant in domestic tasks [34] [38] [51]. In 2016 through 2018, new control methodologies were applied to Superlimbs, including using impedance control for body support, using myoelectric signals from Auricularis muscles for the control of an extra robotic thumb, and using neural signals from a non-invasive electroencephalogram (EEG) in combination with some on-board autonomy to control an Superlimb [30] [25] [42]. Finally, 2018 and 2019 have seen the

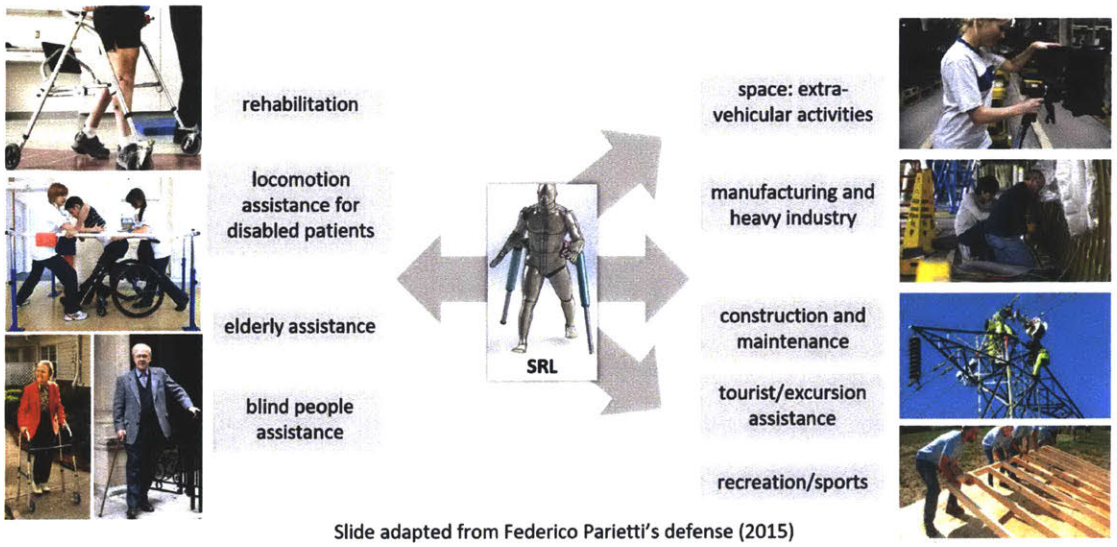

Figure 1-2: Sampling of proposed applications for Superlimbs. Figure adapted from Federico Parietti's defense, 2015.

introduction of soft Superlimbs, low cost 3D printed Superlimbs, and snake shaped Superlimbs [12] [11] [3].

Fig. 1-2 shows a sampling of proposed applications for Superlimbs. This figure highlights the scale of problems Superlimbs might be able to address. Parietti's Superlimbs for balance augmentation targeted helping the elderly limit the risk of falling by increasing the support polygon during walking; the 80+ years old population's risk of death due to falls is 90 times higher than that for the below 60 year old population [38]. Wu's Supernumerary Robotic fingers aimed to help the one in six Americans who suffer from some form of limb mobility and dexterity problem by providing a robotic system that allowed the operator to perform two-handed tasks–like taking the cap off a bottle–with a single hand[55]. Kurek's Mantisbot sought to address a portion of the 240,000 workplace injuries in the manufacturing and agriculture sectors by providing a Superlimb capable of supporting the person during near ground work [25]. Clearly, it is worth pursuing the improvement of Superlimbs.

Despite this significant progress and promise, one commonality among all this work is that it focused on how to improve the robotic system, whether that be through design or control. However, a fundamental property of Superlimbs is that they are

mounted to a person. This begs the question; what are the contributions of the human to the Superlimb-human system and should we care about these contributions?

It is important to note the dearth of research on this topic within the Superlimb community is not indicative of the importance of the topic but rather represents a common evolution of the thinking within a particular field. For example, the exoskeleton research community tracked a similar path. It has been argued that some of the surprising failures within the exoskeleton research community–for example exoskeletal aids that fail to reduce metabolic energy consumption–can be attributed to the failure to account for the constraints of the human control system. For example, in 2003 van den Bogert proposed a design using a single spring intended to reduce muscle forces associated with body weight support, ankle push-off, and leg swing. However, when implemented in 2011 by van Dijk et al., the exotendon design increased metabolic rate despite reducing muscle forces. It is hypothesized this was due to, at least in part, difficulties human users may have had in controlling it [13]. As a result, more recent work in the exoskeleton field has focused on addressing the human element of the human-exoskeleton field both through mathematical modeling and empirical studies. This has led to both passive and active devices that have successfully decreased the metabolic cost of walking [13] [29].

This recent trend in exoskeleton research, and the resultant performance improvements, are mirrored in the prosthetics community [20] and the collaborative robotics community [4]. Given this, one could argue that in order to achieve the promise of Superlimbs, the Superlimb research community must similarly focus on the contributions of the human.

There do exist a few works that explicitly address the human. In one study, users were given control of either two virtual hands that mirrored the movements of their natural hands or three virtual hands, two of which mirrored the movement of the natural hands and the third which mirrored the movement of their foot. When presented with three virtual falling objects, users preferred to use three virtual hands to catch the objects as compared to two virtual hands. However, no significant difference was found in success of catching the falling objects or in average effort

during the tasks when comparing three virtual hands and two virtual hands [1]. While this study does represent an important step toward recognizing the importance of the human in the human-Superlimb system, it still treats the human as something that is being acted upon by the robotic system rather than treating the human as a partner within the system.

Building upon these initial works, and drawing inspiration where possible from neighboring fields, this thesis focuses on the contributions of the human to the human-Superlimb system to inform the control architecture of, design of, communication with, and feedback from Superlimbs. In particular, the key contributions of this thesis are:

- Quantifying the Feasibility of Fully Manual Control for Concurrent Tasks
- Exploiting the Human Operator's High Degrees of Freedom to Design and Control a Reduced-Actuator Supernumerary Robotic Limb
- Quantifying the Contributions of Inherent Haptic Feedback from Supernumerary Robotic Limbs

Chapter 2 discusses an experiment designed at understanding which control architectures are best suited for Superlimbs through human subject testing. In it, we find that fully manual control leads to worse performance with the human operator's natural limbs when asked to perform a task with two natural and two robotic limbs as opposed to with just their natural limbs. This leads to a discussion on the appropriate control architectures given a taskset. Chapter 3 provide two methodologies that leverage the high redundancy and degrees of freedom in the human to address some of the key challenges in the design and communication with Superlimbs. The first methodology exploits the ability of the human operator to directly position the Superlimb through movements of the human body to design a lightweight, reduced-actuator Superlimb. The second methodology exploits a type of redundancy of force at the fingertips in order to send command inputs to the Superlimb despite the human hand's being busy performing another task. Both methdologies in Chapter 3

and Chapter 4 are used to design and control a lightweight, reduced-actuator Superlimb that can assist a human operator by opening a door when one's hands are full. Chapter 4 details two studies on the inherent haptic feedback from Superlimbs. As the Superlimb's state changes, that state change is often reflected as change in forces and torques at the base of the Superlimb that can be felt by the human operator. The first study shows that this feedback, which we have termed inherent haptic feedback, can be leveraged to allow the human operator to perform fully manual closed loop control over the force ouptput of the Superlimb. The second study shows that the human operator can use the inherent haptic feedback to supervise the autonomous actions of a Superlimb. Finally, Chapter 5 provides a conclusion of the work done for this thesis along with recommendations for future directions of work for the Superlimb community based upon some incomplete work with unanticipated outcomes done during the completion of this thesis. This work is discussed with an eye toward future directions.

Chapter 2

Quantifying the Feasibility of Fully Manual Control for Concurrent Tasks

Superlimbs have the potential to broaden the functions a single human user can achieve by helping with holding, grasping, and manipulating objects as well as supporting and bracing the human body [36], [56]. The Mantis-Bot Superlimbs, for example, can brace the human body while the human is working on or near the floor [25]. The human does not have to use one of the natural arms to support the body; the Superlimbs support the body so that both human hands are available for a task. The Superlimbs on the shoulder can hold an object above the head while the human is fixing the object on the ceiling [8].

Many of these examples of initial work within Superlimbs have focused on tasks that can be solved solely through measurements internal to the human-Superlimb system and using a leader-follower paradigm; measurements of the human, the leader, fully dictate the actions of the Superlimb, the follower. For example, the Supernumerary Robotic (SR) Fingers relied on measurement of the human's natural fingers to dictate the position of the SR Fingers [56]. Similarly, the robot on the shoulder used two IMUs to measure the actions of the human's wrists to dictate the robot's position [8]. However, Superlimbs show theoretical promise for tasks requiring perception external to the human-Superlimb system. Because of their portability, Superlimbs could readily be integrated into household tasks such as adding seasoning during

food preparation or opening a door when one's hands are full [32]. In fact, a recent system named Soft Poly-Limbs (SPL) demonstrated their soft robotic wearable arm operating a card swipe and opening a door for a user but relied on a secondary user to control the SPL [33]. Within manufacturing and assembly, Superlimbs could be used to help workers by operating secondary tools, grabbing parts from disparate areas and bringing them together for assembly, changing the orientation of objects during assembly, or holding cables out of the way during assembly. Broadly, barring a secondary operator to control the Superlimb, these tasks require concurrent, simultaneous and independent, actions of the Superlimbs and the natural limbs and perception of objects external to the human-Superlimb system.

Traditional robotics might solve these problems by increasing the manual control of the robot. Placing the human in the loop leverages human's superior perception/planning abilities to solve challenging external perception problems [47]. However, Superlimbs differ from traditional tele-robotics in that ideally the user is able to operate both their natural body and the Superlimb concurrently. This both limits the choice for control inputs and increases the degrees of freedom (DOFs) the user must control. Direct teleoperation of a high DOF robot with a low-dimension input device can lead to high cognitive and physical load [2]. Raising similar questions for Superlimbs, is there any limitation to exploiting extra wearable robotic limbs? Will users be able to take advantage of the extra DOFs offered by Superlimbs when given direct control over them?

Early related research into concurrent operation of Superlimbs and natural limbs on independent tasks is promising. In one study, users were given control of either two virtual hands that mirrored the movements of their natural hands or three virtual hands, two of which mirrored the movement of the natural hands and the third which mirrored the movement of their foot. When presented with three virtual falling objects, users preferred to use three virtual hands to catch the objects as compared to two virtual hands. However, no significant difference was found in success of catching the falling objects or in average effort during the tasks when comparing three virtual hands and two virtual hands [1].

Looking at the example tasks detailed above, many of the tasks, such as adding seasoning to food or operating secondary tools, are currently done sequentially in the absence of Superlimbs. Thus, our goal is not simply to make an existing task easier by adding Superlimbs as previously studied, but to increase the number of actions that can be accomplished by the user concurrently.

In the current work, we aim to address whether a user can control both SRLs and the natural arms concurrently; and how the user coordinates the tasks. To this end, we present a human study designed to explore how humans perform tasks that require concurrent actions, whether natural or robotic. We use simple, but real physical Superlimbs attached to the subject's body. By designing the study to use a pair of natural limbs and a pair of Superlimbs, we can investigate intra-robotic limb coordination as well as intra-robotic-natural limb coordination. The study occurred over multiple days to see if these results were stationary.

2.1 Experimental Setup and Control Signal

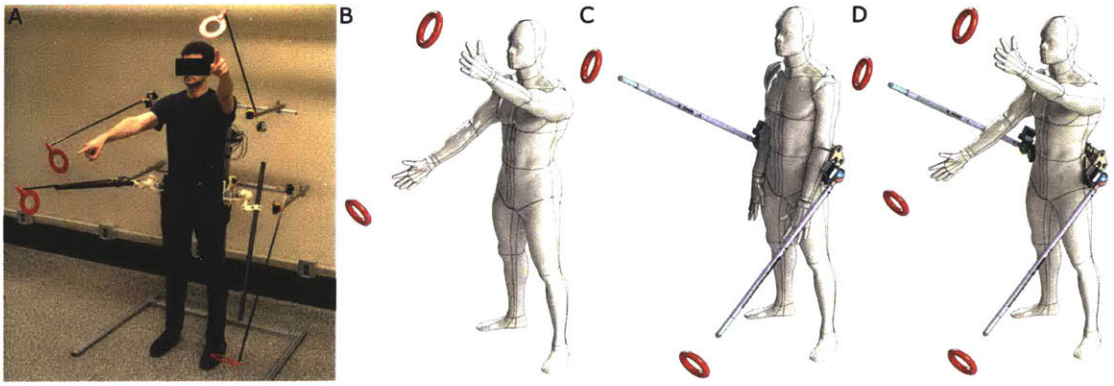

Figure 2-1: Experimental setup. (A) Subject showing the experimental setup and the Superlimb prototype. (B) Stage 1, where the subject moves to the targets with their natural limbs (NL). (C) Stage 2, where the subject moves to the targets with the Superlimbs. (D) Stage 3, where the subject moves to the targets with both the NLs and the Superlimbs.

We conducted human subject experiments based on the protocol approved by the Massachusetts Institute of Technology Committee on the Use of Humans as Experi-

mental Subjects (COUHES), 1604528486. All of the subjects (N = 11) were healthy, male (required because of the control input choice), right-handed subjects, in good physical shape ($20 \leq \text{BMI} \leq 30$) and between 23 and 37 years old. The participants wore a prototype of the Supernumerary Robotic Limbs while standing in front of a set of four targets. See Fig. 2-1A. The goal of the subjects was to minimize the position error between the red disk-shaped targets and the position of the two robotic limbs and the two natural arms as the targets moved from position to position. Each limb, robotic and natural, was constrained to rotate about its base in the sagittal plane. In the case of the natural arms, this motion corresponds to the flexion/extension of the shoulder joint in a sagittal plane. Subjects were asked not to move their arms outside of this plane, extending them and keeping their elbow and wrist joints fixed. Shoulder angle was tracked with an Inertial Measurement Unit (IMU) held in the hand of the subject.

The four targets were placed at the end of carbon fiber rods and moved with servomotors. The targets moved to a new, randomly generated position and held it for 15 seconds. The motions of the targets followed minimum-jerk trajectories [18] and lasted for three of the 15 seconds. The four target trajectories, having diverse step sizes and directions, are independent of each other. The experiment consisted of three stages. In Stage 1, subjects did not wear the robot but simply followed the motion of two targets with their two natural arms by pointing their fully extended arms at the targets (Fig. 2-1B). In Stage 2, subjects wore the robot and followed the motion of two targets with the two robotic limbs (Fig. 2-1C). They were instructed to relax their natural arms. In Stage 3, subjects wore the robot and followed the motion of the four targets using both natural and robotic limbs (Fig. 2-1D).

Each subject participated in two experimental sessions for Stage 1, three for Stage 2, and another three for Stage 3. Sessions consist of multiple trials of tracking, each lasting 3 minutes. Stage 1 and Stage 3 sessions contained 3 trials each, while Stage 2 sessions consisted of 5 trials. Subjects could rest between trials, and could not participate in more than one session per day. Thus, the experiment consisted of eight days of data collection over a two to three week period. Prior to beginning

data collection each day, the system was calibrated until the subject felt they could readily control their natural and robotic limbs. This daily calibration never exceeded 10 minutes. The end point for the experiment was pre-determined based upon a pilot experiment. In total, subjects spent roughly between 6-8 hours performing the experiment and 4.5-6 hours with the Superlimb; this is less than the total time required to perform the experiment as they did not control the Superlimbs during Stage 1.

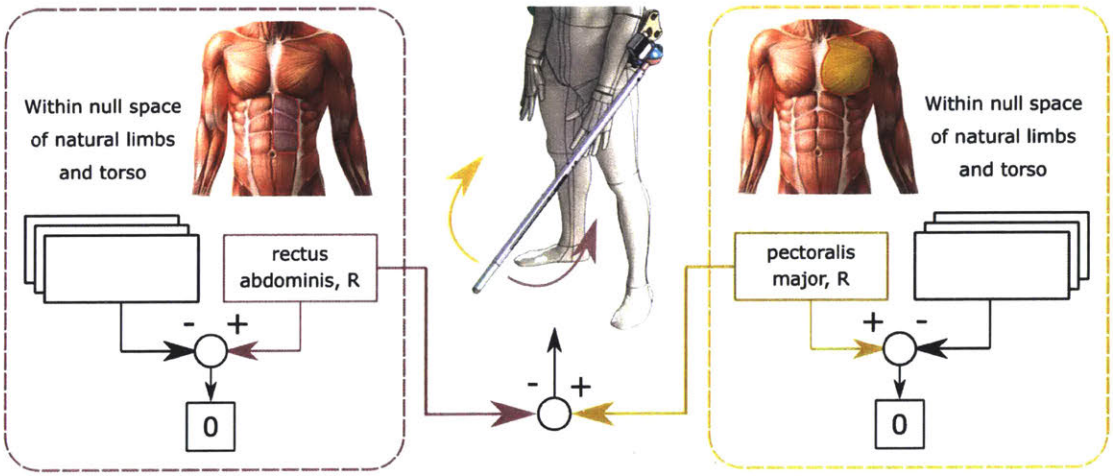

Figure 2-2: Control Architecture. The activation of the rectus abdominis on one side of the body rotates the corresponding robotic limb in the downward direction. The activation of the pectoralis major on the same side of the body rotates the corresponding robotic limb in the upward direction. This is analogous to an antagonistic arrangement of muscles.

While not the focus of this chapter, the control scheme has an impact on the control experience of the subjects. For the purposes of this study, we sought a control scheme that did not interfere with the movement of the natural limbs and had previously demonstrated good performance on a task similar to Stage 2 of this experiment–minimizing the position error between the Superlimb and a target. Given these requirements, we chose to use a previously studied control architecture [35]. In this previous work, the authors compared three different control strategies; position control, velocity control, and one based upon a muscle model. All of these control strategies used surface electromyography (sEMG) sensors placed on the pectoralis major and rectus abdominis muscles to get the subject's control input. The authors

asked their subjects to control two waist-mounted Superlimbs to minimize the error between targets and the Superlimbs as the targets moved from position to position. They found that the velocity control strategy lead to the highest control performance, lowest user effort, and greatest independence from the natural limbs. Thus, while we make no claims of optimality of this control architecture, it satisfies our conditions as a control scheme.

The control architecture consisted of four sEMG sensors (make: Delsys, model: Bagnoli) that were attached to the pectoralis major (left and right) and rectus abdominis (left and right) muscles as shown in Fig. 2-2. The muscle activation signals were acquired at a 250 Hz sampling rate, low-pass filtered (Butterworth filter, 2nd order, 10 Hz cutoff frequency), normalized (with maximum voluntary contraction values recorded at the beginning of every experimental session), rectified, clipped (between 0 and 1), and transformed into reference velocity commands for the robot's motors according to the velocity PI control law with gravity compensation

$$\tau = k_p(\dot{\theta}_{ref} - \dot{\theta}) + k_i \int_0^t (\dot{\theta}_{ref} - \dot{\theta})du + g(\theta)$$

$$\dot{\theta}_{ref} = k_{EMG}(\alpha_{pos} - \alpha_{neg})$$

Here α_{pos} is generated by the pectoralis major and α_{neg} is produced by the rectus abdominis on the same side of the body as the robotic limb. The ratio of the gains is $\frac{k_p}{k_i} = 10$, and their values have been chosen to yield peak time $T_p = 0.5$. The control scheme is applied to each robotic limb separately. Note that as this control strategy re-purposes the torso muscles for the control of the Superlimb, it is fundamentally a cannibalistic control strategy; it does not add degrees of freedom to the body. However, as it operates in the null space of shoulder flexion and extension, it is not functionally cannibalistic in this context.

2.2 Results

2.2.1 Raw Data

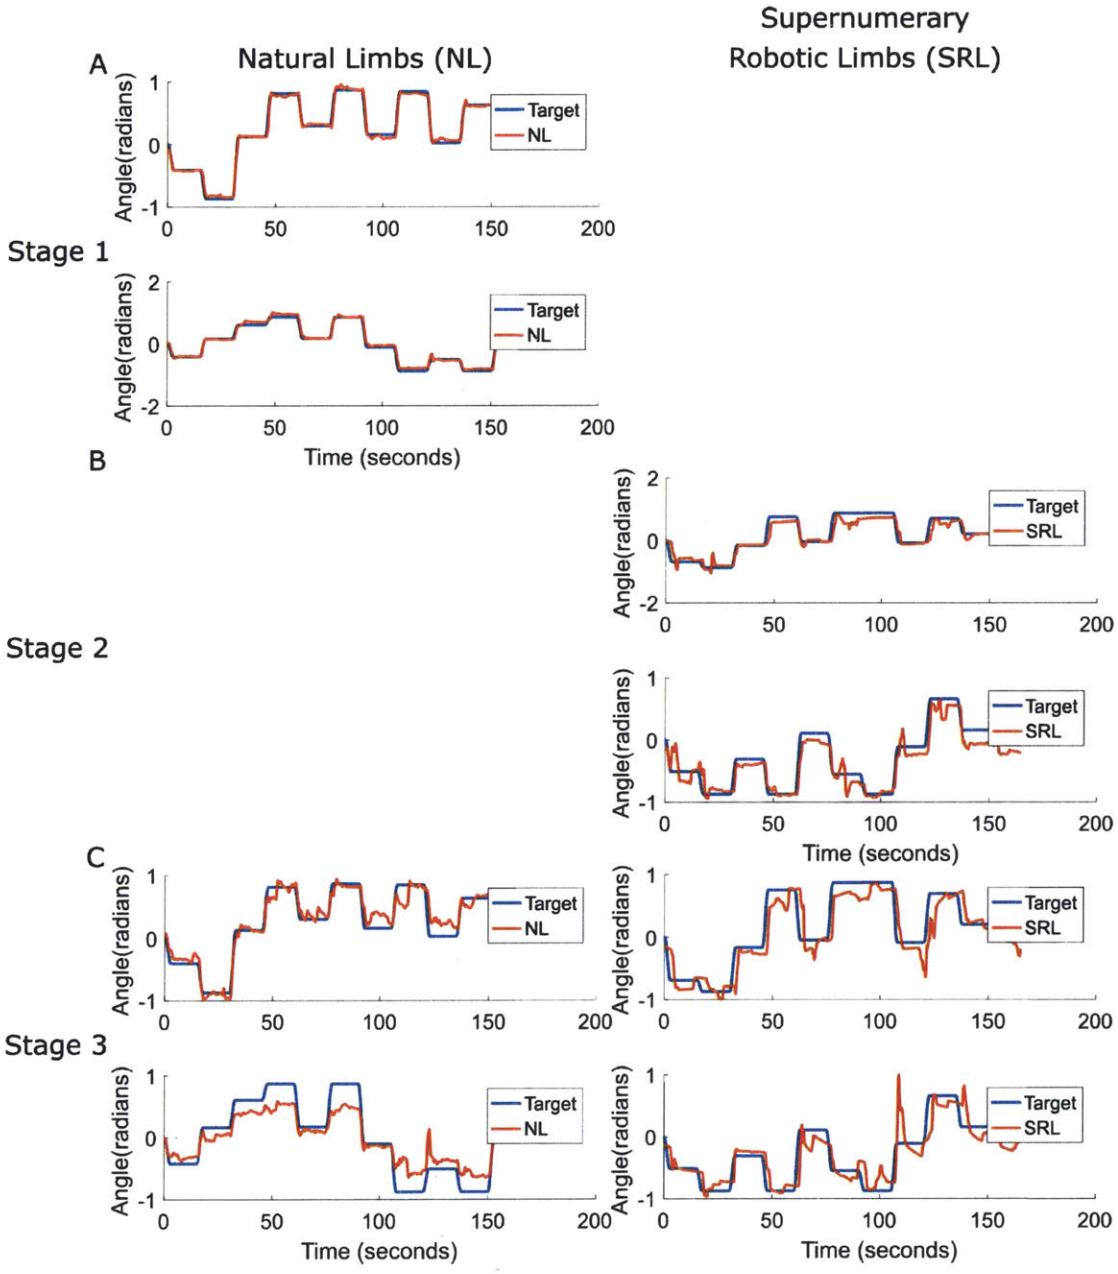

Figure 2-3: Individual Subject Sample Tracking. (A) Stage 1 trial: tracking of two targets with the human arms alone, (B) Stage 2 trial: tracking of two targets with the Superlimbs alone. (C) Stage 3 trial: concurrent tracking of four targets with the human arms and the Superlimbs.

All of the subjects completed the three stages of tracking experiments. A sample of one subject's performance during one trial for each of the three stages is shown in Fig. 2-3. Subjectively, the subject is able to quickly and accurately move their natural and robotic limbs as the targets move to each new position.

2.2.2 Results on Performance

Typical performance metrics for a step responses are the absolute error at steady state and the rise time and as such were applied here. The absolute error at steady state was found for each target position as the absolute value of the difference between the position of the limb and the position of the target at the 15 second mark. The results of all subjects were aggregated together and are shown with the mean and 95% confidence interval in Fig. 2-4A. These results were also aggregated by trial and are shown with the mean and 95% confidence interval in Fig. 2-4C and 2-4D for the natural limbs and Superlimbs respectively. The rise time was calculated as 5% to 95% of the subject's trajectory in response to each target position. The aggregated results are shown with the mean and 95% confidence interval in Fig. 2-4B. These results were also aggregated by trial and are shown with the mean and 95% confidence interval in Fig. 2-4E and 2-4F for the natural limbs and Superlimbs respectively.

Unsurprisingly, the subjects achieved good performance in Stage 1–using only the natural limbs–with a mean absolute error of 0.15 radians and a mean rise time of 2.78 seconds. During Stage 2–using only the Superlimbs–subjects achieved a mean absolute error of 0.18 radians and a mean rise time of 4.20 seconds. During Stage 3–using both sets of limbs–the subjects achieved a mean absolute error of 0.16 radians and a mean rise time of 3.79 seconds with their natural limbs and a mean absolute error of 0.18 and a mean rise time of 4.79 seconds with the Superlimbs. We think that the fairly wide confidence intervals shown in Fig. 2-4C-F, and thus the fairly wide variability, is caused by variability in the task difficulty. Because the target's motions were random, it is likely that some sets of motions were easier than others. For example, if during Stage 3 a subject was asked to move all four limbs up slightly, that would represent a relatively easy set of motions. In contrast, a set of motions

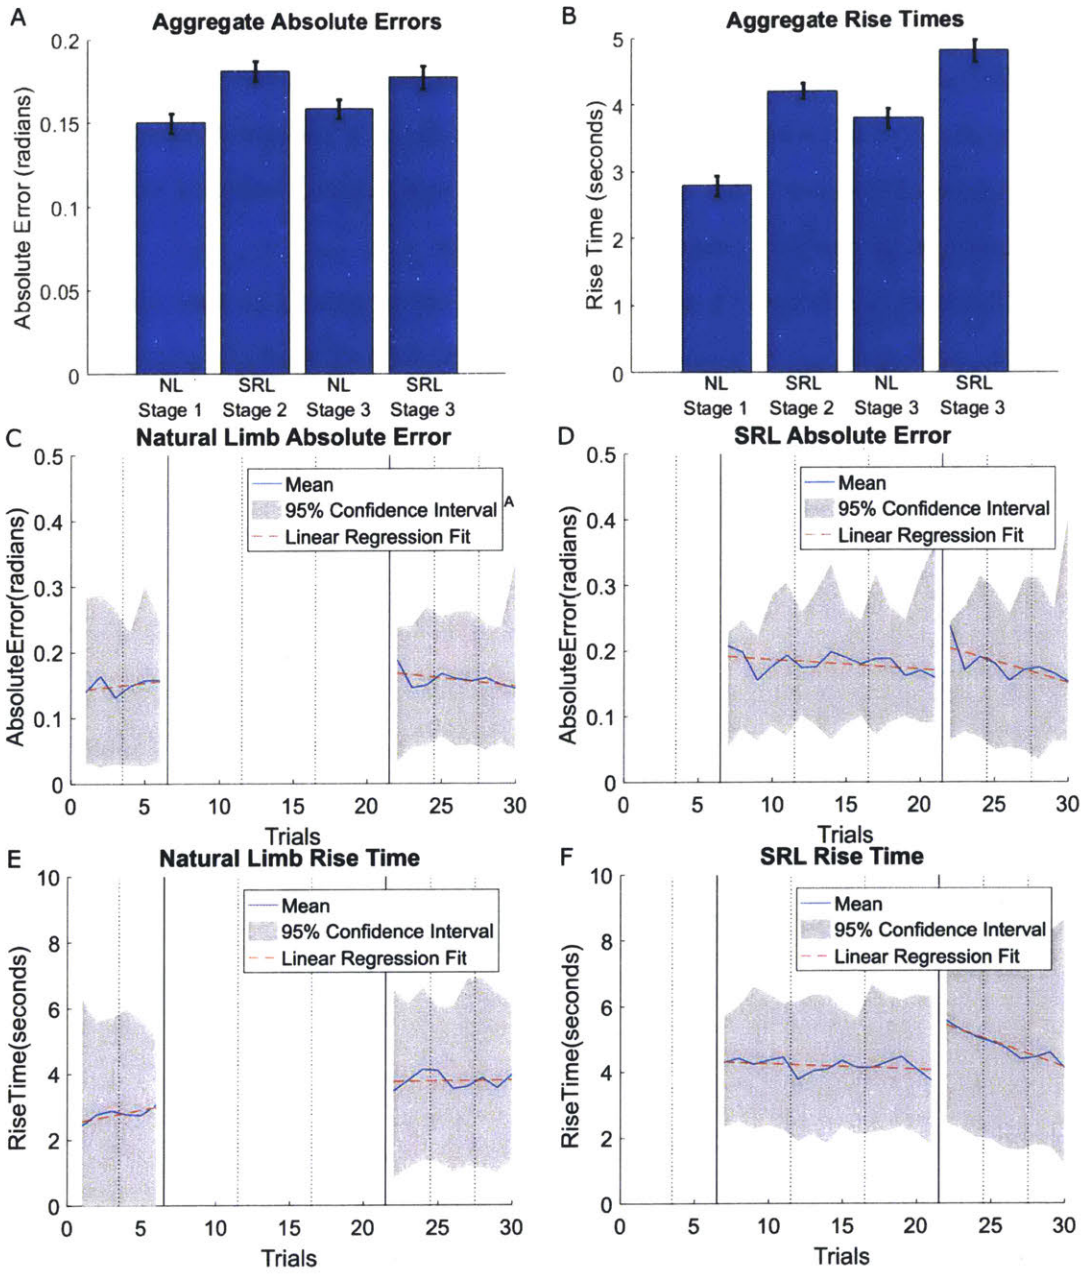

Figure 2-4: Performance. (A) and (B) Aggregate results. Both panels include the four bars which represent the natural limb (NL) stage 1, Superlimb stage 2, NL stage 3, and Superlimb stage 3 performance respectively. The error bars show the 5-95% confidence intervals. (C), (D), (E), and (F) NL absolute error, Superlimb absolute error, NL rise time, and Superlimb rise time respectively. Each panel is divided into 8 sections by the vertical lines, each representing a new session on a different day. The bolded vertical lines also split each panel into 3 sections, each representing a different stage of the experiment.

requiring two limbs to move up and two limbs to move down would be relatively more difficult. Despite this variability, at the aggregated full stage level shown in Fig. 2-4A and B, the confidence interval bars become relatively narrow; this is due to the large number of samples. This allows us to make meaningful statistical comparisons between stages on the performance of the subjects.

Additionally, while these performance results are promising in that the subjects performed well with the Superlimb, it is important to note that the purpose of this study was not to demonstrate the effectiveness of the control architecture used but rather to investigate performance when adding limbs. As such, the performance shown here confirms the control architecture works well enough for our purposes; the subjects can readily complete the tasks given to them.

With these data, we can assess both performance within stages–to see if subjects improved over time, suggesting learning–and between stages–to see if adding additional limbs with their own required actions affected performance of the original limbs on the original positioning task. First, to assess improvement within stages, we will compare the slope of the means of the data over all the sessions within a stage, found by linear regression, to a slope of zero. Our null hypothesis is that there is no difference between the slope and zero while the alternate hypothesis is that the slope is less than zero. Using a one-sample, single-tailed t-test, we found that subjects showed no evidence of improvement during Stage 1 for either end error or rise time(4 degrees of freedom, $p = 0.77$ and 4 degrees of freedom, $p = 0.96$ respectively). During Stage 2, subjects showed statistically significant improvement in terms of absolute error (13 degrees of freedom, $p = 0.05$) but not in terms of rise time (13 degrees of freedom, $p = 0.09$). Finally, during Stage 3, subjects showed statistically significant improvement with the Superlimbs in terms of absolute error (7 degrees of freedom, $p = 0.02$) and rise time (7 degrees of freedom, $p < 0.01$) but no statistically significant improvement with the natural limbs (7 degrees of freedom, $p = 0.09$ and 7 degrees of freedom, $p = 0.56$ respectively).

In aggregate, these results suggest that subjects' performance with their natural limbs did not improve within either Stage 1 or Stage 3. This is expected; learning

is unlikely to take place given the subjects' familiarity with positioning their natural limbs. Modest improvements were shown in Stage 2, which again is to be expected; subjects' performance improved, at least in terms of rise time, as they became more familiar with the control scheme. Finally, subjects showed a fast rate of improvement during Stage 3 with the Superlimbs, suggesting that, while not the purpose of this study, a longer study would have shown even better performance with the Superlimbs before plateauing.

It is important to note that if we apply the same statistical methodology with the null hypothesis being that there is no difference between the slope and zero while the alternate hypothesis is that the slope is greater than zero, we find that the subjects showed statistically significantly worse performance with the natural limbs as judged by rise time during Stage 1 (7 degrees of freedom, $p = 0.04$). Given that the subject performed 3 trials per day with each trial lasting 3 minutes, we think that the reason for the performance degradation for natural limb rise time during Stage 1 is a result of vigilance decrement, which has been shown to slow reaction times and increase error rates during tedious monitoring tasks and has been ascribed to under-arousal caused by an insufficient workload [40].

Next, we can assess how additional limbs affected performance. To do so, we will compare the two performance metrics, absolute error and rise time, between Stage 1 and 3 for the natural limbs and between Stage 2 and 3 for the Superlimbs. To do so we used a one-sided two-sample t-test where the null hypothesis is that there is no difference in performance between the stages while the alternate hypothesis is that subjects performed worse in Stage 3 as compared to Stage 1 for natural limbs and Stage 2 for the Superlimbs. Subjects performed significantly worse with their natural limbs in Stage 3 as compared to Stage 1 with respect to both performance metrics (absolute error: DOF = 3628, $p = 0.03$, rise time: DOF = 3628, $p < 0.01$). Subjects also performed significantly worse with the Superlimbs in Stage 3 as compared to Stage 2 for rise time (DOF = 5806, $p < 0.01$) but not for absolute error (DOF = 5806, $p = 0.77$).

On the surface, these results suggest that subjects performed worse when at-

tempting to position all four limbs as compared to having to position only two limbs. However, it seems likely that, had the experiment continued, the subjects' performance in Stage 3 with the Superlimbs may have not degraded as compared to Stage 2 given the rapid improvement in Stage 3 with the Superlimbs. This is in contrast to the subjects' performance with their natural limbs, which remained reasonably constant across Stage 3. Thus, we can only safely conclude that adding additional robotic limbs degraded the performance of the subjects with their natural limbs over the course of the eight day experiment.

2.2.3 Results on Concurrency

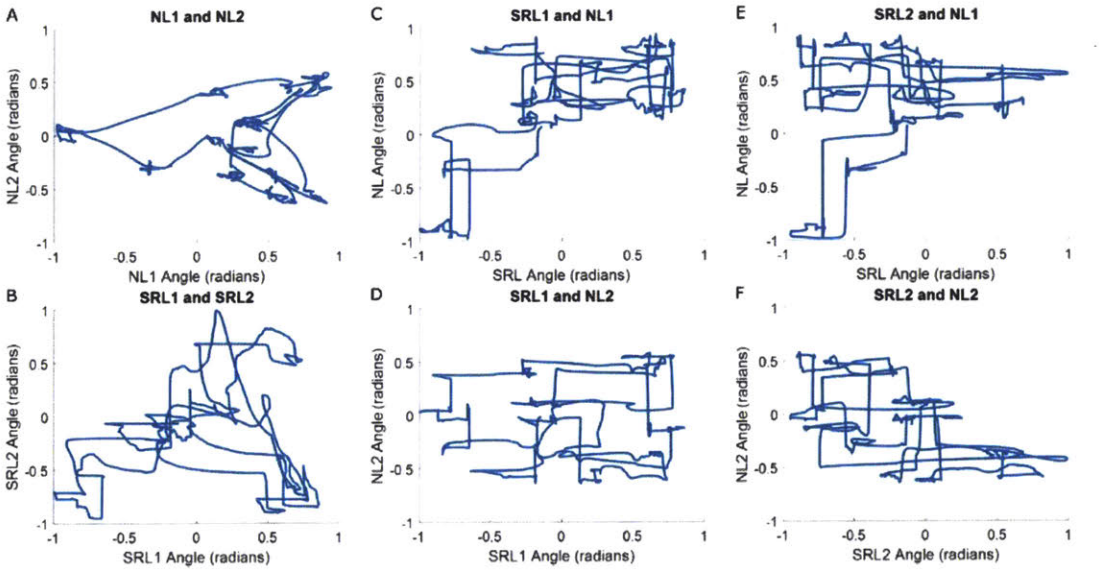

Figure 2-5: An individual subject's paired concurrency chart. Each pair of limbs' trajectories plotted against each other. (A) and (B) show the two NLs and the two SRLs, respectively. If motion of the two limbs is simultaneous, as it is in these two plots, the path will be oblique. (C), (D), (E), and (F) show the four other combinations of limbs. If the motion of these two limbs is not simultaneous, as it is in these four plots, the paths will be either horizontal or vertical.

Looking closely at the trajectories from Stage 3 in Fig 2-3, it appears that there is an offset between when the natural limbs begin to move and the Superlimbs begin to move. To investigate this more closely, we plotted each pair of limbs against one another for a single trial during Stage 3, shown in Fig. 2-5. Given the paired

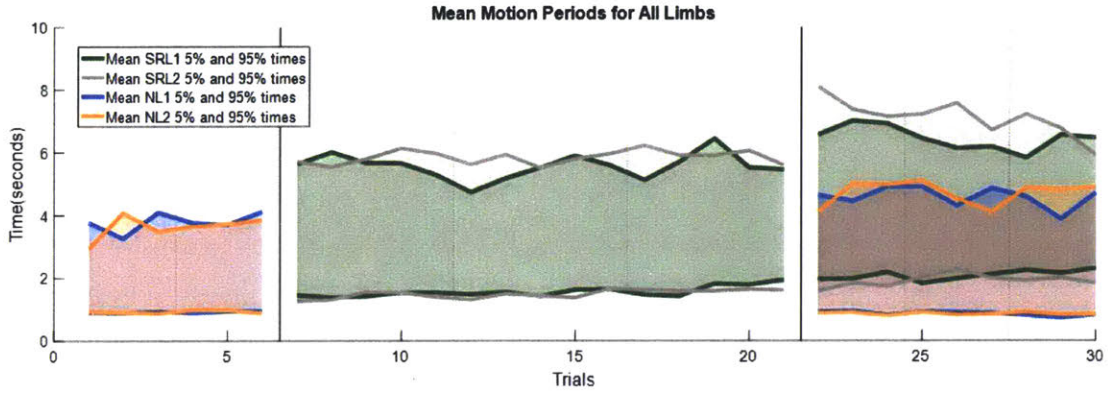

Figure 2-6: Motion periods. NL1 and NL2 serve to distinguish between the two natural limbs while Superlimb1 and Superlimb2 serve to distinguish between the two robotic limbs. Each limb has two lines associated with it. The earlier (lower) line represents the average time that that limb reached 5% of its final position while the later (higher) line represents the average time that limb first reached 95% of its final position. The earlier line represents the beginning of the motion for each limb. On average, the natural limbs were activated first and concurrently followed by the Superlimbs being activated concurrently.

concurrency chart, if two limbs moved simultaneously, we would expect to see oblique lines from point to point. This is in fact what we see in Fig. 2-5A and B. However, if the two limbs were moved in series, we would expect to see either vertical or horizontal lines, as we see in Fig. 2-5C, D, E, and F. Thus we can see that this representative trial is consistent with our hypothesis that there is an offset between when the natural limbs move and when the Superlimbs begin to move.

Fig. 2-6 shows an aggregated view of this data. In it, we plotted two lines for each limb. The first and second line represent the mean time at which a limb reached 5% and 95% of its final position respectively (Fig. 2-6). Roughly, the 5% line represents when the subjects started a movement with a limb. However, the 95% line does not necessarily represents when the subjects ended a movement as many movements had a much longer settling time as shown by the sample trajectories in Fig. 2-4. Within this figure, Superlimb1 and Superlimb2 serve to distinguish between the two robotic limbs while NL1 and NL2 serve to distinguish between the two natural limbs. During Stage 1, subjects' mean starting time with their natural limbs was 0.92 seconds for natural limb 1 and 0.93 seconds for natural limb 2. During Stage 2, subject's mean

starting time with the Superlimbs was 1.56 seconds for Superlimb1 and 1.49 seconds for Superlimb2. Finally, during Stage 3, subjects' mean starting time for natural limb 1, natural limb 2, Superlimb1, and Superlimb2 was 0.86, 0.85, 2.08, and 1.91 seconds respectively. Using the same statistical test as we used to assess improvement in the performance metrics (one-sample, one-tailed t-test with the null hypothesis being that there is no difference between the slope and zero while the alternate hypothesis is that the slope is greater than zero), we found no evidence that the starting time for any limb changed significantly over the course of a stage. However, using a one-sided two-sample t-test, we did find evidence that the starting time of the Superlimbs was later than the starting time of the natural limbs in Stage 3 (DOF = 4354, $p < 0.01$). In aggregate, these results suggest that, on average, in Stage 1 and 2, subjects tended to start the motion with their natural limbs and Superlimbs concurrently. In Stage 3, on average, subjects first moved their natural limbs together followed by moving the Superlimbs together. Because we found no statistical evidence that the slope of any of these starting times differed from 0, we can conclude that all of these results were stationary (i.e. no learning occured).

Finally, Fig. 2-7 shows the fraction of the time that all the different limb groups were moving concurrently during the first eight seconds of each new target position. This cutoff was chosen as it, on average, captures the limbs' movements during the time it takes to rise to 95% of the final settling value, as shown in Fig 2-6. During Stage 1, the two natural limbs move concurrently 24% of the time, which is higher than the percentage of time that either natural limb was moving by itself. During Stage 2, the two Superlimbs move concurrently 26% of the time. This is also higher than the percentage of time that either Superlimb was moving by itself. Finally, during Stage 3, the two natural limbs move concurrently 20% of the time–accounting for times when one of the Superlimbs may have been also moving concurrently–and the two Superlimbs move concurrently 30% of the time, again accounting for times when one of the natural limbs may have been also moving concurrently. Interestingly, both the natural limb-natural limb grouping and the Superlimb-Superlimb grouping are more common than any natural limb-Superlimb grouping. However, all together,

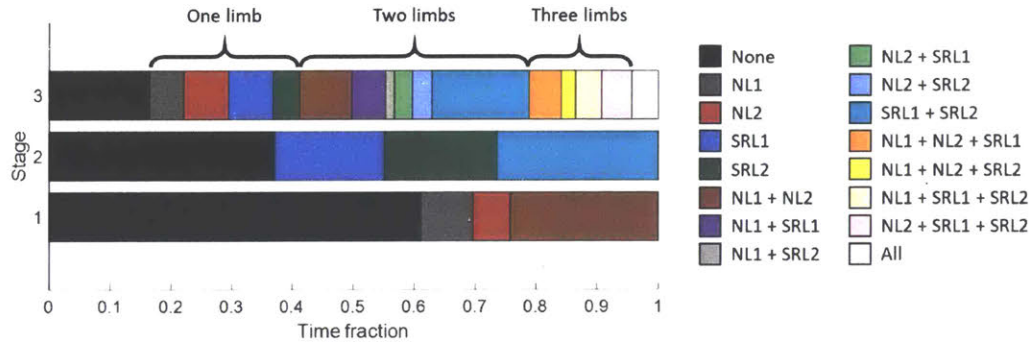

Figure 2-7: Average limb group movement in each stage. Each color within the bars represents a different group of limbs that were moving simultaneously.

at least one natural limb and one Superlimb were moving concurrently 34% of the time. All four limbs were moved concurrently 4.3% of the time.

2.3 Discussion

While Supernumerary Robotic Limbs (Superlimbs) have been successfully applied to tasks that only require perception internal the human-Superlimb system, such as bracing or as an assistive technology for people with disabilities, Superlimbs show theoretical promise in tasks that require perception external to the human-Superlimb system such as adding seasoning while cooking or operating secondary tools on the manufacturing floor. One attractive way to expand the task set that Superlimbs can be applied to is to insert humans into the loop and allow them to manually direct the Superlimb, thus leveraging the human's superior perception. However, limited effort has been put forth to date to understanding the effect of adding additional limbs on user performance. This study seeks to fill that gap by quantifying how subjects use physical Superlimbs during multi-limb concurrent and independent tasks.

First, we found evidence that subjects' performance with their natural limbs degrades when subjects are asked to control all four limbs on concurrent, independent tasks as compared to when the subject performed the task with only their natural limbs. Further, this effect did not diminish over the course of the multiday experiment. This is perhaps unsurprising; much work has suggested that multitasking leads to in-

creased error and longer execution times [39]. In adding additional concurrent actions required of the subjects, as was done during Stage 3, less of their focus can be directed at the two natural limb targets. We might expect that, given sufficient practice, the performance degradation shown here would dissipate; a singing piano player provides evidence that humans can perform complicated manual control tasks concurrently with other task. However, this requires year, and perhaps decades, of practice. Thus it is significant that the performance degradation exists, and was stationary, given the relative simplicity of the task–controlling four degree-of-freedoms–and length of the eight-day experiment. We desire that Superlimbs offer immediate performance improvements rather than performance improvements with significant practice. Thus it appears that Superlimbs would benefit from strategies developed by the shared control community to combat the difficulty of controlling many DOFs concurrently by attempting to seamlessly assist the user in completing the desire action with less control input.

It is unclear why the performance degradation that was found in the natural limbs during Stage 3 was not also found in the Superlimbs. One hypothesis is that the improvement from learning to operate the Superlimbs counteracted the performance degradation typically associated with increasing the DOFs that had to be controlled, thus disguising the effect. Future work could investigate this by lengthening the study to reach a plateau in subject's performance with the Superlimbs.

Second, we found evidence that, at least within the context of this experiment, subjects responded first and simultaneously with their natural limbs followed by moving their Superlimbs together. The reasoning for this is not a physical incapacity. As demonstrated in previous work, the EMG-based agonist-antagonist control scheme allows the subjects to operate the Superlimbs without causing motion of the natural limbs [35]. Further, we found that subjects did move all four limbs concurrently–4.3% of the time–during Stage 3.

The observation of how subjects chose to respond to four concurrent and independent tasks can be leveraged in two ways. First, it can help inform the choice of control scheme for the Superlimbs. Interestingly enough, as noted in the introduc-

tion, many of the previous examples of Superlimbs have used a leader-follower control paradigm [56], [25]. Placing the human as the leader in this control paradigm is a natural application of this observation.

Beyond this though, it informs which principles from shared control may be most successful for Superlimbs. In our experiment, we observed a time offset between when motion was initiated with the natural limbs as compared to the Superlimbs. Thus one naive control scheme may be to divide the control of the Superlimb up temporally by initially attempting to automate the actions of the Superlimbs as the subjects move their natural limbs. After a brief time, control of the Superlimbs could be transferred to the human who could complete the fine positioning of the Superlimbs.

Second, this observation can help inform how to divide labor between the natural and robotic limbs during concurrent and independent tasks. While not a perfect analogy, past work has shown that people tend to initiate movement with their dominate hand when presented with concurrent tasks, especially when the hands are not cooperating on the same object [17], [22]. Thus, in tasks that require concurrent robotic-natural limb action, the more time sensitive task should be assigned to the natural limbs.

Finally, we found evidence that while the natural limbs were privileged, moving first and simultaneously followed by the robotic limbs, at least one natural limb and Superlimb were moved concurrently 34% of the time during Stage 3. This suggests that the subjects did not simply break the task down into two sequential tasks to be operated on by each pair of limbs but rather partially combined the use of both the natural and robotic limbs.

We note that all of the findings of this chapter must be taken within the context of the controller choice. The particular muscles chosen here may not work well for women. This certainly opens the door for other muscle choices; the frontalis muscle on the forehead has been used to control a robotic sixth finger [44]. Further, the choice of a velocity-controller is less bio-mimetic than other potential choices as our limbs don't stop moving and hold position when our muscles stop firing. However, the control architecture chosen here allowed the subjects to perform sufficiently well with

the Superlimbs on this particular task to allow comparisons in performance during multi-limb concurrent and independent tasks. Future studies could look at whether the findings here hold true for alternate control schemes. Additionally, future studies could investigate the effect of providing proprioceptive feedback to the subjects on the state of the Superlimbs; proprioceptive feedback has shown value in the prosthetics community [19].

2.4 Conclusion

In conclusion, we first found evidence that while subjects were able to move all four limbs concurrently, within the context of this experiment, there is a limit to exploiting wearable robotic limbs for additional concurrent tasks. Namely, we found that the subjects' performance with their natural limbs degraded as additional robotic limbs were added. While future work could investigate whether this performance degredation dissipates with practice, we desire that Superlimbs offer immediate benefits rather than requiring significant practice. This suggests that Superlimbs may benefit from drawing on control strategies from the shared control or other communities. Second, we found evidence that, within the context of this experiment, subjects generally moved their natural limbs prior to their robotic limbs. This suggests a leader follower paradigm of control, where the Superlimb measures the actions of the human and responds accordingly. Further, it informs how labor could be divided between the natural and robotic limbs. Specifically, it suggests giving the primary, more time-sensitive task to the natural limbs.

Chapter 3

Exploiting the Human Operator's High Degrees of Freedom to Design and Control a Reduced-Actuator Supernumerary Robotic Limb

A number of Supernumerary Robotic Limbs (SuperLimbs) attached to various portions of a human body have been developed in the last several years. These include SuperLimbs attached around the waist for assisting a worker in holding an object [36], mounted on the shoulders of a worker for lifting an object while a worker is affixing it onto the ceiling [8], supernumerary robotic fingers attached to a wrist for grasping odd-shaped objects along with the five natural fingers of a wearer [56], and ones attached to the back for a variety of tasks, including picking up objects, soldering electric wires, or body support [45] [33] [25]. In these applications, the Superlimbs play a supportive role; while the human performs the primary task, the SuperLimbs assist the human by holding an object, fetching an object, or supporting the human body.

While a human performs the primary task, the human can still move other parts of the body without interfering with the execution of the primary task. A human

body has an enormous number of degrees of freedom (DOFs). This flexibility and redundancy of a human body has been used for controlling wearable robotic systems. Prosthetic arms can be controlled with sensors detecting the patient's shrug and shoulder movements [6] while Superlimbs have been controlled with foot switches when the human operator is seated [45] and electromyograms (EMG) sensors attached to abdominal and pectoral muscles for controlling two robotic arms [35]. These body movements or signals utilized for controlling and communicating with SuperLimbs and prosthetics are orthogonal to or separated from the body movements required for executing the primary task. In this current work, we aim to exploit the high redundancy and flexibility of the human body for solving some of the key challenges in the design, communication with, and control of SuperLimbs.

One of the fundamental difficulties in designing SuperLimbs is that SuperLimbs containing many active joints tend to be too heavy to wear for a long time. By exploiting the high redundancy of a human body, this problem can be solved or alleviated. Since SuperLimbs are attached to a human body at their base, the limbs can be controlled directly by moving the base with movements of the human body. No active DOF are needed for the SuperLimbs in the direction of motion that can be generated by the human body movements. The number of actuators can be reduced and thereby the weight of SuperLimbs can be reduced. This strategy, however, requires quantitative analysis of body movements to examine whether a human can comfortably or naturally generate the required movements while adhering to the constraints dictated by the primary task. A methodology for quantifying which active DOFs the human can supply will be developed in this chapter.

The chapter is organized as follows. First, the design concept of reduced-actuator SuperLimbs will be addressed. A methodology for quantifying the usable body movements while performing the primary task is presented in order to examine which of the SuperLimbs' active joints can be eliminated. Second, the communication from the human to the SuperLimb using finger forces will be discussed. An algorithm will be developed for quantifying the usable space of finger forces and generating patterns of coded finger forces without interfering with the primary task. Finally, both meth-

ods will be implemented on a proof-of-concept prototype, and a human subject will demonstrate the usefulness of the methods.

3.1 Reduced-Actuator Design

Within a limited range and direction, a human can directly adjust the position and orientation of SuperLimbs attached to the body simply by moving a particular part of the body. Direct positioning is simple and intuitive as it does not require the human to communicate their intent to the Superlimb to actuate the direct body-positioning DOFs. Thus, reducing the number of actuators is not only effective for reducing the weight of SuperLimbs, but can also make the control and communication task simpler.

The question is which required DOFs can be replaced by direct human body movements, and how to split the control space between direct human movements and active SuperLimbs DOF. Here, we propose the following three steps.

- Task analysis: Given task specifications, determine both the attachment point of the Superlimb on the body and the required DOFs for performing the task.

- Quantification of usable human body movements: Although a human has a high degree of flexibility and redundancy, the body can hardly be moved in certain directions at a particular posture. Namely, the movable range and direction of a DOF may be reduced or affected if some part of the body is constrained. For example, a human is engaged to execute a task using both arms, the arms must not be disturbed by the movement of other part of the body. Given constraints and conditions, the usable body movements that a human can naturally generate must be quantified.

- Splitting the motion space: Based upon the task analysis and quantification of usable human movements, the motion space is split such that the human supplies the DOFs they can provide while the Superlimb provides all remaining

DOFs required to perform the task. From this, the kinematic structure of the Superlimb is determined.

The key step in the above design procedure is the quantification of usable human body movements. One approach to determining this would be to look at hierarchical models of human kinematics [5]. However, the human body has over 600 muscles, 200 DOFs, and complex joint structures. Further, in our application the human body is subject to various constraints, which make the analytic modeling complex and less accurate. Thus while human body models are available, we instead take an experimental approach to quantifying available body movement when the human is constrained.

3.1.1 Experimental Quantification: A Case Study

This chapter demonstrates the above methodology by designing a lightweight, reduced-actuator Superlimb that can assist a human operator by opening a door when the human operator's hands are busy carrying a large box. This task was chosen both because it represents a potential task set out in the Superlimb community and because it represents a problem robots are often applied to in the broader robotics community [32] [52]. Further, it represents a problem that is readily relatable and broadly applicable. While in this particular instance we have chosen to have the user hold a box, the box could be swapped out with groceries, furniture, or stock materials.

Many previous works within the human factors community have used the desired task set to inform the design of the robotic system [48]. For the Superlimb community, the task set is also used to motivate the placement of the Superlimb on the human body. For overhead tasks such as holding a panel in place, the Superlimb was placed on the shoulder [8]. For body support during manufacturing tasks like welding, the Superlimb was placed on the back [25]. For helping the human operator open a bottle, the Superlimb was placed about the wrist [56]. All of these works have chosen to mount the Superlimb close to the eventual task space as it both reduces weight and limits the Superlimbs ability to collide with the human. In our case, the chosen task

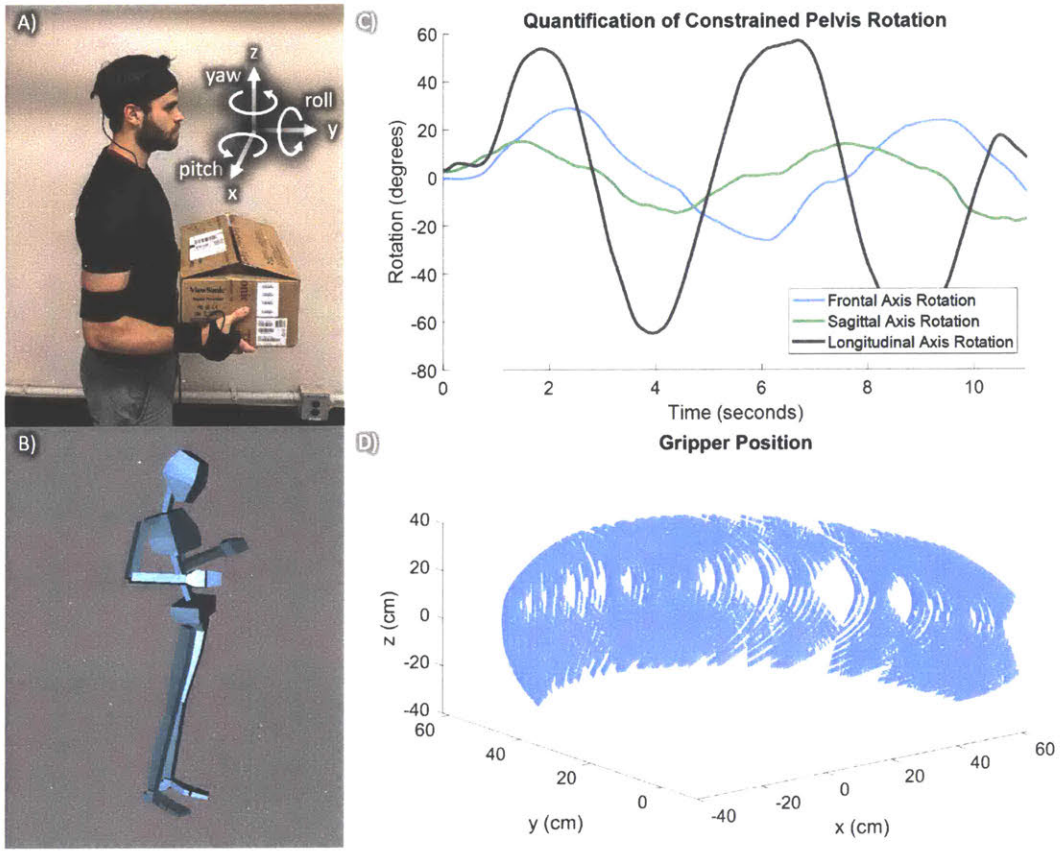

Figure 3-1: Experimental Quantification of Human Body Pelvis Motion With Constraints: (A) Shows a subject wearing the Shadow motion capture system while holding a large box. (B) Motion capture rendering of the human operator rotating about the pelvis while constrained by holding a box. (C) The movement of the pelvis about the three rotational degrees of freedom. (D) The rotational DOFs of the pelvis' effect on the position of the Superlimb gripper, derived through forward kinematics with the offset between the pelvic coordinate frame and the gripper coordinate frame defined as dx = 28 cm, dy = 55 cm, and dz = 0 cm.

occurs at roughly waist height. Given this, we can choose the waist as a mounting point for the Superlimb.

Next we must assess the minimum required active DOFs in order to perform the tasks. We will define our coordinates to be placed on the gripper, with y pointing forward, parallel to the sagittal axis and z pointing up, parallel with the longitudinal axis, as shown in Fig 3-1A. Further, we will define roll as rotation about the y axis, pitch as rotation about the x axis, and yaw as rotation about the z axis. In order to reach the door handle, the end effector on the Superlimb must be able to be positioned

in x, y, and z. To turn the door, the end effector on the Superlimb must be able to close and roll about the arm axis (which should be aligned with the door handle axis). We will refer to the end effector actuation as a gripper DOF. Next, to open the door, the Superlimb end effector must once again be able to be moved in x and y. Finally, to release the door handle, the Superlimb end effector must roll and release the gripper. Thus, for the door opening task, our required DOFs are x, y, z, and roll with an additional DOF for the gripper.

We can next look at what DOFs the human is capable of supplying for these tasks based upon the Superlimb being mounted on the waist. Fig. 3-1A shows a human whose body is partially constrained (holding a large box) and is wearing the Shadow motion capture system (Motion Workshop, Seattle, WA 98104, USA). The human operator was asked to move slowly and to comfortable positions, thus limiting the range to rotations the subject can achieve quasi-statically. From this, we can extract the three rotations about the pelvis, in terms of rotation about the frontal, sagittal, and longitudinal axes (Fig. 3-1B). Pelvic rotation about the sagittal axis represents roll at the gripper. Note that while the rotations are shown overlapping in time, they were collected during individual trials and thus don't overlap in time.

From this, we can see that the person can only contribute roughly +/- 16 degrees of roll. This is not sufficient to open a door handle. Though the operator shows the ability to contribute roughly +/- 60 degrees of yaw and roughly +/- 27 degrees of pitch, we require neither of these DOFs. Finally, Fig. 3-1C shows the pelvic rotations transformed to the Superlimb gripper position. This was achieved by assuming that the Superlimb is mounted rigidly to the pelvis and using the yaw, pitch, roll rotation matrix that has been augmented to include the translational offset between the pelvis and the gripper, defined as $\boldsymbol{T} \in \mathbb{R}^{4x4}$ [26]. We can find the translation of the gripper due to the rotation of the pelvis with

$$P_{gripper} = \boldsymbol{T} * \left[[d_x \quad d_y \quad d_z \quad 1]^T \right] \tag{3.1}$$

Assuming the Superlimb is held straight out from the person, parallel to the

ground, we used d_x = -28 cm, d_y = 55 cm, and d_z = 0 cm. Thus the rotation about the pelvis can provide roughly 90 cm of travel in x, roughly 70 cm of travel in y, and roughly 60 cm of travel in z, shown in Fig. 3-2 as a semi-transparent partial ellipsoid. Because the person can also translate in x and y by walking, the reasonably large magnitude of the x and y travel of the gripper due to pelvic rotation is redundant for this particular task. The ability to translate the gripper in x, y, and z from pelvic rotation will likely be useful as the human operator can then use pelvic rotation for some fine positioning of the gripper at the door handle while using the ability to walk to generate larger magnitude changes in x and y.

Clearly, the Superlimb must actuate the roll DOF as well as the gripper. However, as the human only has minimal comfortable control over the z height–neither squatting nor rotating about the frontal axis at the pelvis for prolonged periods is comfortable, especially with a large box–and because we want the Superlimb to be able to move out of the way, we will additionally choose to actuate z. Thus the Superlimb will have three DOFs; z, roll, and the gripper.

Importantly, the design methodology laid out here has taken a task that would have previously required the Superlimb to have at least five active DOFs and limited the Superlimb to needing only three active DOFs. This significantly decreases the complexity and weight of the Superlimb. Fig 3-2 shows a CAD rendering of the Superlimb prototype as designed through this methodology.

3.2 Control Concept

We next propose a novel control input methodology for communicating a rich variety of commands to the Superlimbs while both hands are busy. This is achieved by exploiting a type of redundancy in force at the fingertips. Namely, while the sum of the forces on the fingertips must often satisfy a certain condition, the distribution of forces on the fingers can vary. The fingers are an attractive means for communicating with the Superlimb as they represent arguably our highest bandwidth channel. By performing an analysis of natural behavior when performing a task, we can identify

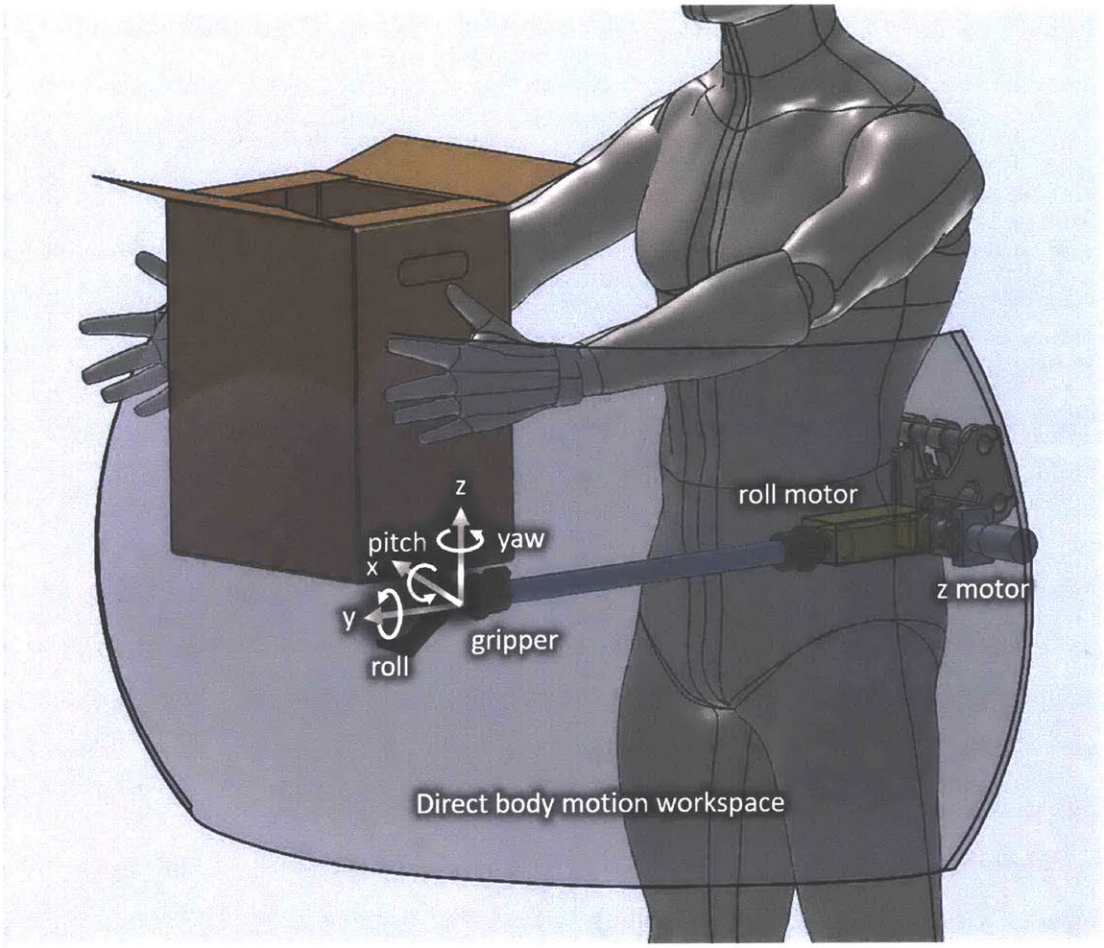

Figure 3-2: CAD rendering of the Superlimb prototype: Based upon the design methodology, the Superlimb only requires active degrees of freedom for the roll, z, and gripper axes. The rotational of the pelvis' effect on the position of the Superlimb gripper is shown as the transparent semi-ellipsoid.

commands that the human operator can generate but that also exist in the null space of the task. To do so, we propose the following three steps.

- Task Analysis: Given a task, we can quantify normal task performance by measuring the force on the fingertips as a human performs the task. Additionally, we can quantify the available variations in force on the fingertips by asking the human to explore different force combinations while still achieving the primary task goal.

- Identify Command Input Subspace: In order to find a command subspace that is

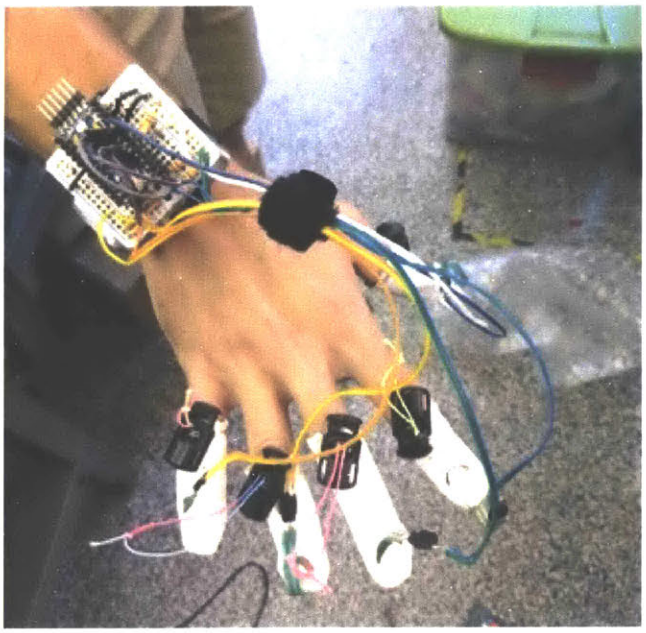

Figure 3-3: Force Sensing Resistors mounted to each of the person's fingertips.

both achievable given the constraints of the primary task and does not interfere with the primary task, we can look for the intersection of the null space of the task dataset and the the range space of the exploration dataset found in step 1.

- Parse Input Command To Superlimbs: Assuming we have identified the command subspace, we can manually place codes (i.e. particular force combinations on the fingertips) within this subspace. Then, when a new measurement is taken, we can project the new measurement into the command subspace and compare it to the manually defined codes to determine if the human intended to send a command to the Superlimbs.

Based upon this methodology, and in combination with an intermittent control structure where the control authority is alternated at the human operators' discretion, we controlled the reduced-actuator Superlimb described above.

The intermittent control architecture operates such that whenever the human generates a non-zero command input, the command authority is completely handed over to the human. However, when the human is not generating any command inputs, the Superlimb is free to perform actions. We propose this structure first as it ensures

safety; if the Superlimb behaves unpredictably, the human operator can take control to change or stop the actions of the Superlimb. Second, allowing the human operator to step in and explicitly direct the actions of the Superlimbs allows the Superlimb-human system to handle challenging problems by leveraging the combined strengths of the Superlimb-human system [47]. This control structure falls under the supervisor role in Human Robot Interaction literature [46].

Given this control structure, we propose to have the human operator have the ability to issue four different commands to the Superlimb; move the Superlimb's gripper up and level with the door handle, move the Superlimb's gripper down and out of the way, grip and rotate the door handle, and release the door handle. These four action primitives represent all that is required for the Superlimb to assist in opening a door when the human operator's hands are full. By breaking the control inputs into discrete action primitives, we allow the human to adopt a "move-and-wait" strategy, where open loop command inputs to the robot are followed by a wait of typically one loop delay time such that the human operator can get confirmation and feedback. This "move-and-wait" control strategy has been shown to be effective in the presence of large delay [49].

3.2.1 Input Via Redundant DOFs on the Human Hand: A Case Study

We desire a command input mechanism that doesn't interfere with the human operator's task and can be readily generated by the human operator. To identify this mechanism, we must first take data on the person performing the task naturally. In order to do this, we use the prototype shown in Fig. 3-3 to measure force at the fingertips. We then ask them to hold the box naturally; this data serves as the baseline for how one would normally perform the task. Fig. 3-4A shows three different trials of the person holding the box. Next, we ask the person to explore different ways of holding the box by varying the force distribution on the fingertips (shown in Fig. 3-4A in blue). Note that while the human operator wore all 10 sensors, only the data

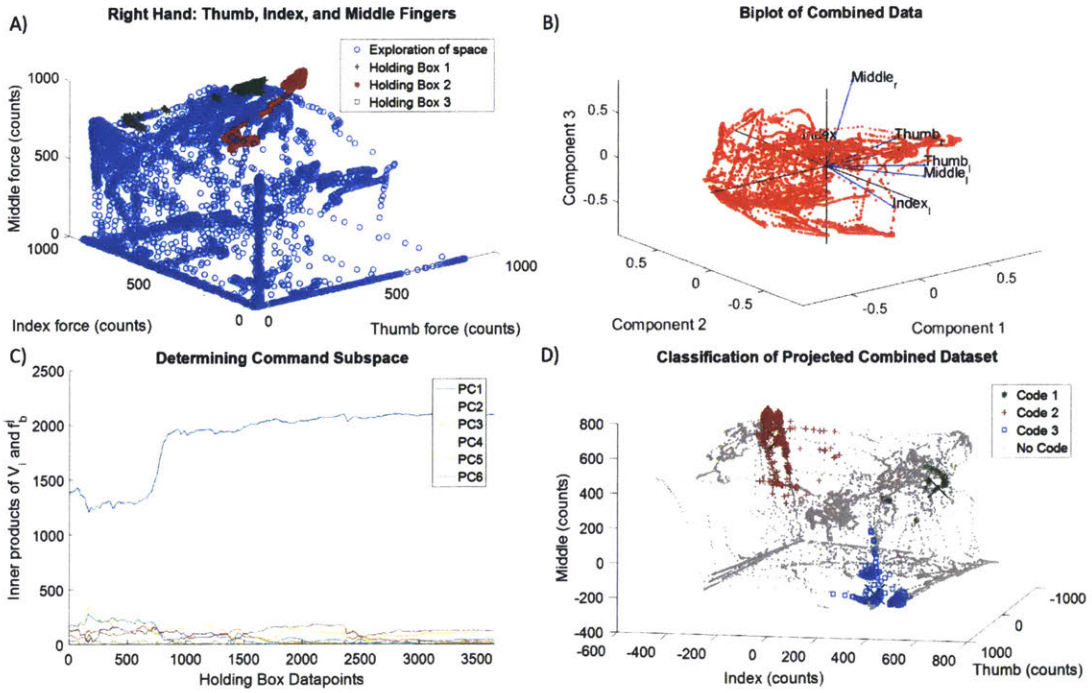

Figure 3-4: All data was obtained from a single subject. (A) Holding a box and exploration datasets for the right thumb, index, and middle fingers. Despite being constrained to hold a box aloft, the human operator shows that they can vary the distribution of forces at the fingertips while still holding the box in the exploration dataset. (B) Biplot showing the first three principal components of the combined dataset. (C) The inner product of the range space of the combined dataset with each point in the holding box dataset. From this, we can determine that the second through sixth components are orthogonal to the holding box dataset and exist within the range space of the combined dataset. Therefore, the second through sixth components span the subspace where codes can be communicated. (D) Classification of the projected combined dataset into three manually defined codes using Euclidean distance and a maximum distance threshold.

from the two thumbs, index fingers, and middle fingers was used within this section for ease of visualization.

With this data, we want to identify codes that won't be accidentally sent to the Superlimb during normal task performance. To do this, we can place the commands within the null space of the task; this ensures that when the human operator generates the commands it does not interfere with the performance of the task. However, a command input existing in the null space of the task does not guarantee that it is achievable. For example, imagine if a command input required that a person remove

the entire left hand from the box. This would certainly exist within the null space of the normal task performance but it would also result in the person dropping the box. Thus we require that the commands we choose also exist within the range space of the combined exploration and task dataset (the exploration dataset is shown in Fig. 3-4A in blue). That is,

$$\boldsymbol{S} = \boldsymbol{R} \cap \boldsymbol{N_t} \tag{3.2}$$

where $\boldsymbol{S}$ is subspace where we can place stored codes, $\boldsymbol{R}$ is the range space of the combined exploration and task dataset, and $\boldsymbol{N_t}$ is the null space of the task dataset.

$\boldsymbol{R}$ can be found with Principal Component Analysis (PCA). Let $\boldsymbol{f_b} \in \mathbb{R}^{nx1}$ be the force at the fingertips for a single timepoint during the holding box trial and $\boldsymbol{f_e} \in \mathbb{R}^{nx1}$ be the force for a single timepoint during the exploration trial. Given that we collect N_b holding box datapoints and N_e exploration datapoints, we can concatenate the data from the two trials to get the combined data matrix $\boldsymbol{X} \in \mathbb{R}^{nx(N_e+N_b)}$. Let $\bar{f}$ be the mean of f^1 through $f^{N_e+N_b}$, then the covariance matrix $\boldsymbol{C} \in \mathbb{R}^{nxn}$ is given by

$$\boldsymbol{C} = \frac{1}{N_e + N_b} \sum_{n=1}^{N_e+N_b} (\boldsymbol{f^i} - \mu_{\boldsymbol{f}})(\boldsymbol{f^i} - \mu_{\boldsymbol{f}})^T \tag{3.3}$$

$\boldsymbol{C}$ can be decomposed using singular value decomposition (SVD) such that

$$\boldsymbol{C} = \boldsymbol{V}\boldsymbol{D}\boldsymbol{V}^T \tag{3.4}$$

$$\boldsymbol{D} = \text{diag}\begin{bmatrix}\lambda_1 & \dots & \lambda_k & \lambda_{k+1} & \dots & \lambda_n\end{bmatrix} \tag{3.5}$$

$$\boldsymbol{V} = \begin{bmatrix}V_1 & \dots & V_k & V_{k+1} & \dots & V_n\end{bmatrix} \tag{3.6}$$

where $\boldsymbol{D}$ has the eigenvalues placed in descending order, $\boldsymbol{V} \in \mathbb{R}^{nxn}$ holds the eigenvectors, and $|V_i| = 1$. The first eigenvectors are shown in the biplot in Fig. 3-4B. The eigen decomposition can be truncated at k where $\lambda_k \geq \epsilon$ and $\lambda_{k+1} < \epsilon$, and epsilon is some suitable value. Then we can find $\boldsymbol{R}$ and $\boldsymbol{N}$, the null space of the combined dataset, by

$$\boldsymbol{R} = \begin{bmatrix} V_1 & \dots & V_k \end{bmatrix} \boldsymbol{N} = \begin{bmatrix} V_{k+1} & \dots & V_n \end{bmatrix} \tag{3.7}$$

In this particular case, our smallest eigenvalue represented over 2% of the variance of the data, so we set $k = n = 6$. This means our range space $\boldsymbol{R}$ is spanned by the first through sixth principal components. A biplot of the first three principal components is shown in Fig. 3-4B.

Next we must distinguish between force vectors intended to be commands to the Superlimbs and force vectors generated during execution of the primary task. To do so, we will determine in which subspace the vectors in the task data lie, where $\boldsymbol{f_b^j}$ represents the j^{th} of N_b total force measurements during the performance of the task. This can done by checking the orthogonality of every datapoint $\boldsymbol{f_b^j}$ to every eigenvector V_i. That is, if

$$\langle V_i, f_b^j \rangle > \delta \text{ for } 1 \leq j \leq N_b \tag{3.8}$$

then V_i exists within the range space of the task, $\boldsymbol{R_t}$. Else, V_i exists within the null space of the task, $\boldsymbol{N_t}$. Let m, $m < k$, represent the cutoff such that V_1 through V_m exist within $\boldsymbol{R_t}$ and V_{m+1} through V_k exist within $\boldsymbol{N_t}$. Importantly, V_{m+1} through V_k also exist within $\boldsymbol{R}$ and thus

$$\boldsymbol{S} = \boldsymbol{R} \cap \boldsymbol{N_t} = \begin{bmatrix} V_m & \dots & V_k \end{bmatrix} \tag{3.9}$$

Fig. 3-4C shows a graphical representation of (8). Setting δ appropriately, we argue that V_1 exists within $\boldsymbol{R_t}$ while V_2 through V_6 span $\boldsymbol{S}$, the command subspace. Any code placed within $\boldsymbol{S}$ will both not interfere with the performance of the primary task as it exists within $\boldsymbol{N_t}$ and will be physically achievable as it exists within $\boldsymbol{R}$.

Finally, assume that we have a new force measurement $\boldsymbol{f^*}$. In order to determine if this measurement contains a command input for the Superlimb, we can project $\boldsymbol{f^*}$ into $\boldsymbol{S}$ by constructing a projection matrix $\boldsymbol{P_S}$ such that

$$\boldsymbol{P_S} * \boldsymbol{f^*} = \boldsymbol{f_S^*} \tag{3.10}$$

$$\boldsymbol{P_S} = \boldsymbol{S}\boldsymbol{S}^T \tag{3.11}$$

where $\boldsymbol{f_S^*}$ represents the new force measurement projected into $\boldsymbol{S}$. Then, we can compare it with any suitable distance metric to predetermined codes $\boldsymbol{c_i}$ that exist within the $\boldsymbol{S}$ to see if it matches one of them. This takes the form

$$d(\boldsymbol{c_i}) = \sqrt{(\boldsymbol{c_i} - \boldsymbol{P} * \boldsymbol{f^*}) \cdot (\boldsymbol{c_i} - \boldsymbol{P} * \boldsymbol{f^*})} \tag{3.12}$$

$$\hat{c} = \begin{cases} \arg\min_c d(\boldsymbol{c_i}) \text{ for } d(\boldsymbol{c_i}) < d_{thresh} \\ \text{no code for } d(\boldsymbol{c_i}) > d_{thresh} \end{cases} \tag{3.13}$$

where $\hat{c}$ is the identified command and d_{thresh} is defined as the maximum distance a code can be away for a point for that point to be classified to that code. The smaller the d_{thresh}, the fewer the points that will be classified as that signal. This allows us to avoid accidental code triggers. In Fig. 3-4D, three codes have been defined and the exploration dataset has been projected into $\boldsymbol{S}$. The force measurements that are sufficiently close to the manually defined codes have been classified as those codes while the remainder of the force measurements have been classified as not containing command inputs to the Superlimb.

Fig. 3-5 shows what happens when a new time sequence of force measurements contains a code. Fig. 3-5A shows a green dataset that briefly leaves the range space of the task and enters the null space of the task. Fig. 3-5B shows the green dataset projected into the command input subspace $\boldsymbol{S}$. The distance between all the projected data and the manually defined code, shown as an "x", is then calculated and compared to a distance threshold as discussed above. The projected data that is close enough to the manually defined code is classified as containing a command to the Superlimb.

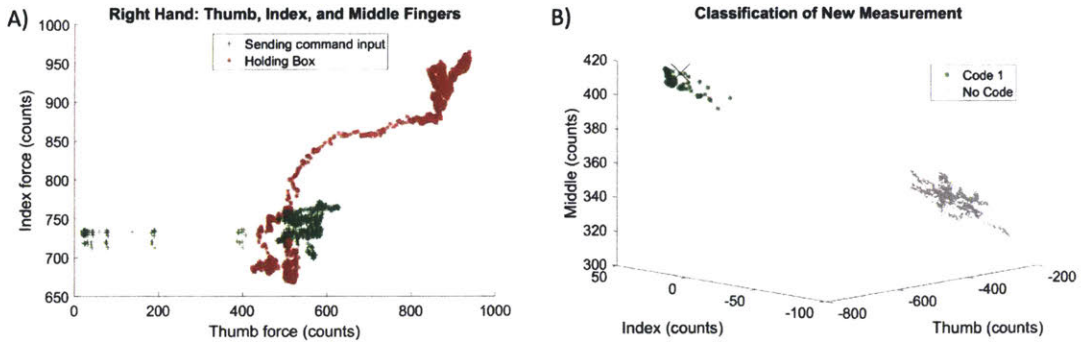

Figure 3-5: A force measurement dataset that contains a code. (A) A comparison between the nominal holding box dataset. The green dataset, which contains the code, briefly leaves the range space of the task data. (B) The green dataset projected into $\boldsymbol{S}$ and then classified using a distance metric. The data that is close to the manually defined code (the large x) is defined as containing a command to the Superlimb while the remaining data is defined as not containing a code.

3.3 Prototype and Demonstration

We built a prototype of the Superlimb and used it to open a door for a human operator who had her hands full. The prototype has two 118RPM, 6.766 Nm Robotzone planetary gear motors (Robotzone LLC, Winfield, Kansas, USA) to actuate the roll and z DOFs. The gripper (Makeblock, Shenzhen, China) is placed at the end of the Superlimb arm. The roll and z motors are driven using a Cytron MDD10A dual DC Motor Driver (Cytron, Pulau Pinang, Malaysia) while the gripper is driven with a Pololu Dual MC33926 Motor Driver (Polulu, Las Vegas, Nevada, USA). Finally, the encoders on the roll and z motors are read with a Dual LS7366R Quadrature Encoder Buffer (SuperDroid Robots Inc., Fuquay-Varina, North Carolina, USA). All signals are routed to an Arduino Uno (Arduino, Ivrea, Italy) which operates PI position controllers on the roll and z motors. The gripper is controlled through a velocity controller. The Arduino Uno communicates with a high level controller written in python through the Robotic Operating Systems (ROS) rosserial protocol. Additionally, the force on the human operator's fingertips is measured using force sensing resistors (Adafruit, New York City, New York, USA) as shown in Fig. 3-3.

For ease of implementation, we have chosen not to strictly follow the control input methodology described above. The FSRs mounted to the fingertips do not

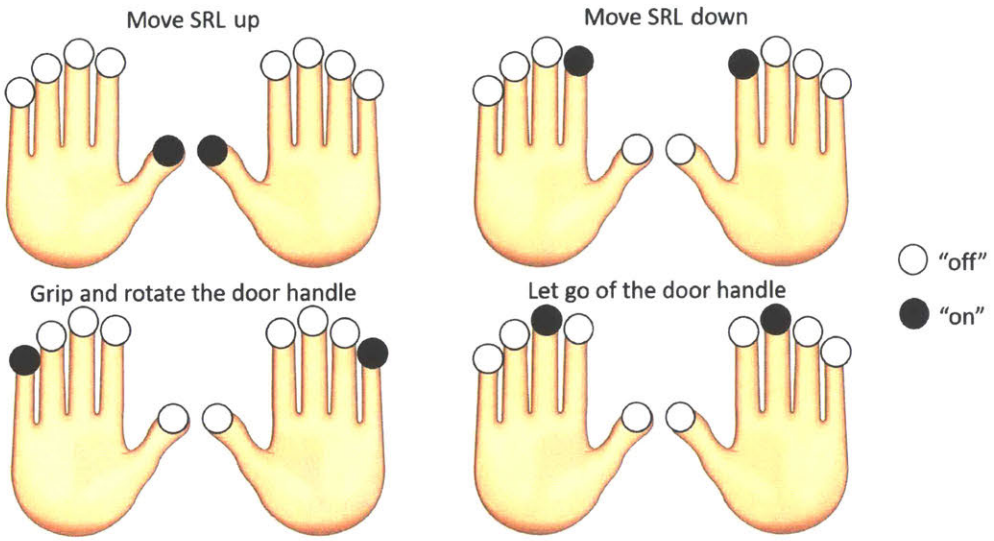

Figure 3-6: All the codes chosen to communicate the action primitives to the Superlimb. Each sensor is thresholded such that it is either deemed on–black–or off–white. Each code exists within $\boldsymbol{S}$

Figure 3-7: Demonstration of the Superlimb. In the first image, the human operator and Superlimb work together to position the end effector in x, y, and z. In the second image, the human operator commands the Superlimb to grab and twist the door handle with a single code. In the third image, the human operator backs up to open the door. Finally, in the last image the human operator commands the Superlimb to release the door handle. The human operator is then free to walk through the door.

sense force placed at other areas of the hand or fingers. Thus it is possible to hold the box without registering any force on the fingertips. Further, the FSRs do not provide robust signals as they respond to flexing of the sensor as well as force. Given these limitations, we have chosen to threshold each FSR sensor to either on or off.

We chose four codes, one for each action primitive outlined in the previous section. Each code will specify whether each of the ten sensors on the fingertips is on or off. The codes were found to exist within $\boldsymbol{S}$. Classification of the command input is done by requiring that the thresholded force sensor readings match the codes defined in Fig. 3-6 for at least 0.25 seconds.

Fig. 3-7 shows the prototype in action. First the human operator commands the Superlimb to raise its gripper toward the door. Next the human operator uses direct body motion to position the gripper at the door handle. The human operator then commands the Superlimb to grip and twist the door handle. Once the door handle is turned, the human operator backs up to open up the door. Finally, the human operator commands the Superlimb to release the door knob and drop back down out of the way. The human operator is then free to walk through the door. The human operator required less than fifteen minutes of training before being able to perform the task, with the primary difficulty resulting from ensuring that the force sensitive resistors on the human operator's fingers were picking up the codes the human operator was attempting to communicate to the Superlimb.

3.4 Discussion

While this study applied both the design and control input methodologies to having a Superlimb open a door when one's hands are full, these methodologies have broad applicability.

Namely, the design methodology proposed here can be applied to any task where the human operator and the Superlimb work collaboratively to perform some overall task. For example, if the task was instead to have the Superlimb hold a nail while the human hammers it in, we could once again take advantage of the human's ability to position the end effector to reduce the number of actuated degree of freedoms on the Superlimb. This holds true on tasks ranging from having the Superlimb add spices during cooking to holding a component that the human is going to solder.

One potential downside of the design methodology proposed here is that it in-

creases the workload on the user. Depending upon the task set, it might be best to allow the human to focus on their task while having the Superlimb have all the actuated DOFs it needs to perform its task without human collaboration. If an analysis of the taskset concludes that the human operator is likely to be overworked if they are asked to contribute active DOFs, then one should proceed toward a solution that limits the human operators' involvement.

While the design methodology applicability may be narrowed by the taskset, the control input methodology is not. Even in cases where it is determined that the Superlimb should perform its task without any collaborative help from the human operator, it is critical that the human operator still hold ultimate command authority. As the Superlimb is mounted to a person, the person must be able to stop or alter the actions of the Superlimb if they determine the Superlimb is acting unsafely. This control input methodology could be used to supply the person with an emergency stop command that does not interfere with the normal task operation.

There exist many avenues to improve upon the control input methodology proposed here. First, while in this work the force was only measured on the fingertips, the force could be readily taken from the whole hand or even other parts of the body. This would yield a more complete picture of a person's interaction with the object, likely leading to a larger null space and thus more room for command inputs. Second, improvements could be made on the diversity of signals that can be communicated to the robotic system. While we used discrete codes for the communication with the Superlimb, the force on the fingertips can be varied continuously and thus could be used to send continuous commands to the Superlimb. Third, improvements could be made to the choice of commands to ensure that they are intuitive for the human operator, perhaps by leveraging human factors research to understand which commands are easier to generate.

Fundamentally, both methodologies introduced here exploit the human operator to make the combined task of the Superlimb-human system easier. This points to a clear direction forward for the Superlimb research community; in what ways can we leverage the human operator? What are the limitations of mounting a robotic system

on a person?

3.5 Conclusion

In this chapter we first introduced a new design methodology that combines a taskset analysis with an analysis of the required degrees of freedom that the human operator can actuate to reduce the number of actuated degrees of freedom the Superlimb requires, thus reducing weight and complexity. Next, we introduced a novel control input methodology that also relies on a task analysis to identify input commands that do not interfere with the human operator's normal performance of the task and are physically achievable by the human operator. Finally, we used both of these methodologies, along with an intermittent control architecture, to build a Superlimb prototype that can assist the human operator by opening a door when the human operator's hands are full. Combined, these methodologies represent some of the first attempts to leverage the human operator within Superlimb research.

Chapter 4

Quantifying the Contributions of Inherent Haptic Feedback from Supernumerary Robotic Limbs

While still a relatively new field, Supernumerary Robotic Limbs (SuperLimbs) show promise both as an assistive technology for people with disabilities [56], [23], [54] and as an augmentation technology with applications ranging from body bracing during fatiguing work on a floor ([25]) to operating as an additional robotic helping hand during surgery [1] or assisting humans in daily chores at home [32]. Potentially, SuperLimbs allows a human to execute additional tasks on top of those ordinarily done by the human, but the challenge is how to communicate with and control the Superlimbs to perform the additional tasks.

Human-robot communication and control have been extensively studied in the collaborative robotics community in the last decade [45], yet SuperLimbs pose unique issues and challenges as well as opening up new opportunities. Since SuperLimbs are physically attached to the human body, forces and moments acting on a Super-Limb from the environment may be transmitted to the human body and, thereby, the human can recognize contact with the environment. Fig. 4-1 shows a pair of SuperLimbs sitting on the shoulders of a human who is working on the ceiling. As the SuperLimbs contacts the ceiling panel and pushes it upward, the reaction forces

Figure 4-1: A pair of Superlimbs sitting on the shoulder of a human operator. The human operator feels the reaction force from the Superlimbs holding the ceiling panel in place.

act on the shoulders of the human. The human can notice and confirm that the SuperLimbs are holding the panel; this allows the human to remove his hands from the ceiling panel in order to grab a screw and screwdriver to secure the panel. This is a type of haptic feedback differs from haptic displays as no explicit detection of physical contact between SuperLimbs and the environment is required. This haptic feedback is inherent; the force is transmitted directly to the human. This inherent haptic feedback can be exploited for human-SuperLimb coordination.

Prosthetics, which are also attached directly to humans, deliver a similar type of haptic feedback to the human. Specifically, passive prostheses powered by body movements are often preferred over active prostheses controlled with myoelectric signals

despite the active devices being the most efficient in terms of dexterity and interface intuitiveness. This is because body-powered prostheses have superior sensory feedback [14]. To illustrate one practical outcome of this superior feedback, one study found that the force feedback provided from body-powered prostheses allowed subjects to more accurately identify an object based upon stiffness as compared to visual feedback [10].

These prior works on prosthetic devices provide useful examples of the utility of inherent haptic feedback, which may also inform the importance of inherent haptic feedback in designing SuperLimbs. However, there is a significant difference between prosthetic devices and SuperLimbs. The former performs tasks in lieu of the lost limb, while the latter can allow the human to multitask by executing an independent or separate task in coordination with the natural arms and legs.

Based on this unique feature of SuperLimbs, we can envision three use cases of SuperLimbs where the inherent haptic feedback is useful.

- Case 1: When vision is either occluded or unavailable, the inherent haptic feedback from the Superlimb can fill the gap by allowing the human to sense the state of the Superlimb. For example, while a human is visually focusing on one task, the inherent haptics provides a separate, non-visual channel through which the human can sense the state of SuperLimbs.

- Case 2: When vision is neither the appropriate nor dominant modality for sensing a particular state, inherent haptic feedback may serve as an effective means for detecting and controlling that state, be it the stiffness of an object or the force output of the Superlimb. In the ceiling panel installation task shown in Fig. 4-1, the human confirms whether the panel is securely held with the SuperLimbs by detecting the reaction force, not the position of the SuperLimbs.

- Case 3: When vision is available and the dominant sensory modality, the inherent haptic feedback may still be useful to supplement visual observation. Multi-sensor integration may improve performance and usability. According to [24], the integration of haptic displays even in the presence of visual feedback

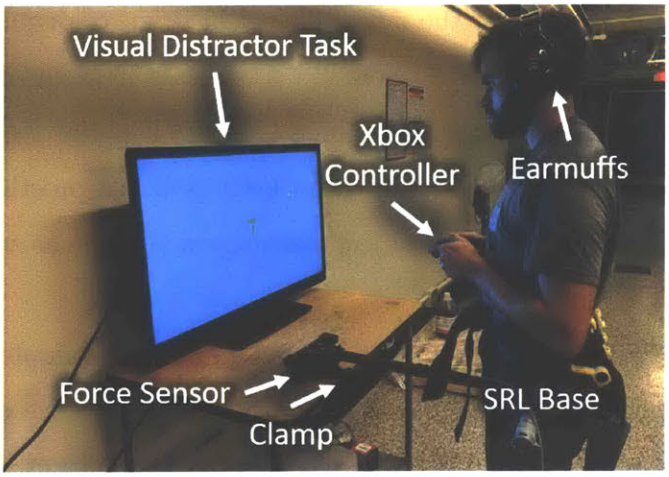

Figure 4-2: Experimental Setup: Shows a staged version of a subject performing both the visual and haptic task simultaneously. The subject uses an Xbox Controller to respond to the tasks. The Superlimb is clamped to the table so that there is no movement of the Superlimb. The visual task consists of a A series of yellow and white letters that were displayed on a monitor in front of the subject. The subject was asked to press âĂIJAâĂİ on the Xbox controller every time they saw a white letter.

will increase success rate and ease of use of a powered prosthesis but yield slower motions.

In the following sections, experimental studies using human subjects wearing SuperLimbs are conducted for two specific use cases; one is manual closed-loop control of force based on inherent haptic feedback and the other is supervision of SuperLimbs operation via inherent haptics. Both experimental results provide insights into inherent haptic feedback and argue for the importance of designing Superlimbs to exploit the inherent haptic feedback.

4.1 Methods

The inherent haptic feedback from Superlimbs is critical in Case 2, where forces and moments play the major role in performing a task. Thus, we aim to design experiments that feature this. Further, in many use cases, natural human limbs and SuperLimbs perform multiple tasks simultaneously. We aim to address whether

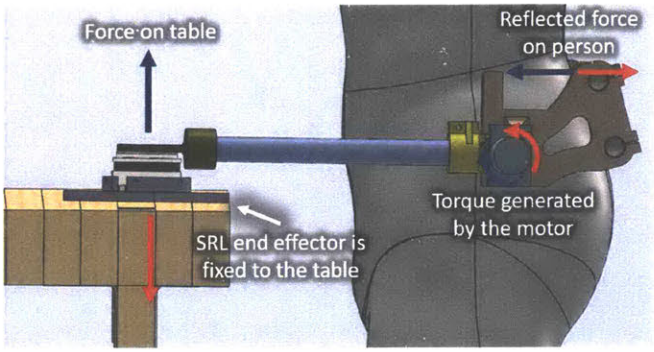

Figure 4-3: A rendering of the Superlimb. As the Superlimb exerts a torque on the table, an equal and opposite torque is applied to the subject, thus providing inherent haptic feedback on the state of the Superlimb to the wearer.

SuperLimbs with inherent haptic feedback allows humans to perform two tasks simultaneously. Typically, the natural limbs execute the primary task, while SuperLimbs perform secondary or supportive tasks. In such scenarios, vision is used mostly for the demanding, primary task, and is not available for the secondary task performed by SuperLimbs. This is, therefore, Case 1 where vision is not available.

4.1.1 Experimental Apparatus

Figure 4-2 shows a staged experimental setup, consisting of a prototype SuperLimb, an Xbox game controller, and a display. With this apparatus, two tasks can be executed simultaneously. One is detection and control of a force acting at the tip of the SuperLimb, and the other is a demanding visual task (Fig. 4-2). The human subject is prompted to respond to a series of images presented on the monitor display, while controlling and/or monitoring the force acting on the SuperLimb. The Xbox controller held in the subject's hand is used for executing both tasks.

For the force detection and control task, we devised the apparatus so that positional information plays no significant role. The end effector of the Superlimb was fixed to a table and only the force output of the Superlimb changes. This resulted in a situation where haptic feedback is clearly the dominant form of feedback as vision is incapable of assessing the force output of the end effector given no motion. This allows us to highlight just the contributions of the inherent haptic feedback to sensing

the state of the Superlimb.

We designed and conducted human subject experiments based on the protocol approved by the Massachusetts Institute of Technology Committee on the Use of Humans as Experimental Subjects (COUHES), 1903759075. All of the subjects (N = 9) were healthy, right-handed subjects in good physical shape ($20 \leq$ BMI ≤ 30) and between 23 and 37 years old. Experiments were performed in a single session not lasting more than 1.5 hours.

The experiment participants wore the prototype SuperLimb around their waist with a belt-like harness while standing in front of a desk. See Fig. 4-2. The robotic arm is made of simple composite tubes and has only one actuated degree of freedom, rotating in a plane parallel to the sagittal plane. The robot arm is actuated by a 12 rpm, 57 Nm stall torque brushed DC planetary gear motor (ServoCity, Winfield KS). Attached to the end of the arm is an Optoforce 6 axis force torque sensor (ONROBOT, Odense, Denmark).

As the SuperLimb's actuator generates a torque τ_{act}, its reaction torque $-\tau_{act}$ acts on the base of the harness. Furthermore, as the SuperLimb pushes down its endpoint to the table, its reaction force $-F_{table}$ acts on the harness base in the vertical direction. The harness base is a 3d-printed solid frame that transmits the reaction force and moment to a human subject's waist.

Using this apparatus, two human subject experiments were conducted.

4.1.2 Inherent Haptic Feedback for Low Autonomy Manual Control of a Superlimb

First, we looked at low levels of autonomy; we asked the subjects to continuously control the Superlimb to regulate the reaction force at a desired setpoint despite disturbances. This is similar to how the Superlimb may be used to help a person, perhaps elderly or recovering from injury, to transition from sitting to standing. As the amount of assistance one person needs is likely different from the assistance desired by other people and as an individual's desired assistance likely changes over time, it

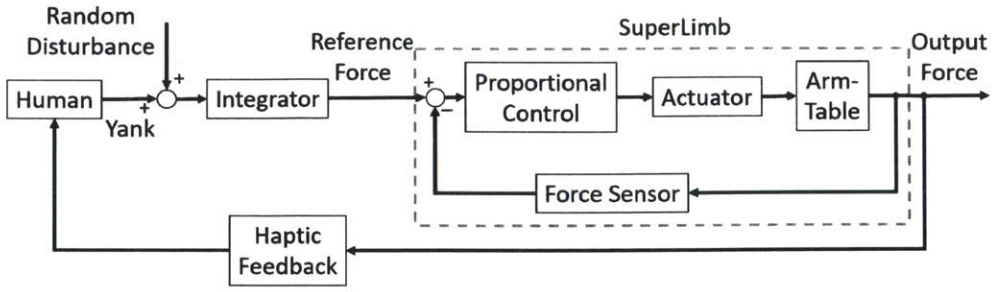

Figure 4-4: Block diagram of the manual force control system. The gray dotted line demarks the boundary between the Superlimb controller and the human operator. The human commands yank, the time derivative of force. A randomized disturbance is superimposed on this commanded yank and the output is fed through a integrator to generate a reference force for the Superlimb control system.

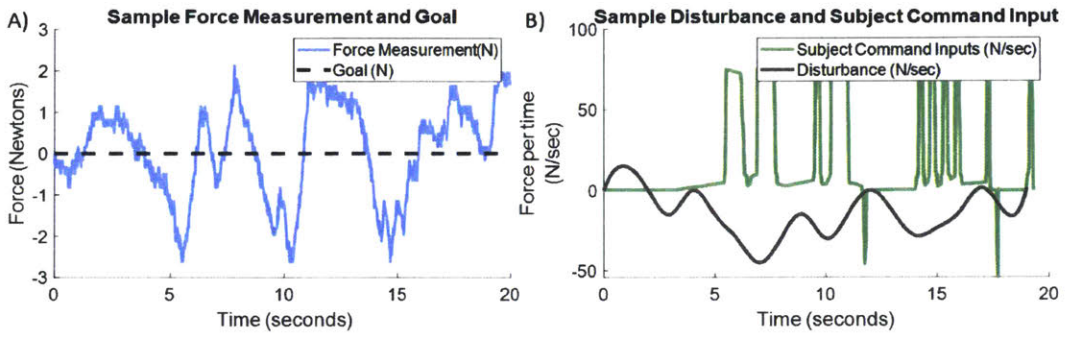

Figure 4-5: Sample Low Autonomy Manual Control Superlimb Task: (A) The subject's goal was to maintain zero force exertion as measured by the force sensor on the table. This figure shows a sample trial, including the actual force measurement (blue) and the goal (black dotted line). (B) Shows the randomized disturbance (gray) and the user command input that sought to cancel out the disturbance (green).

is difficult to pre-program a desired force trajectory for helping the person stand. A solution is to place the person in the loop and let them directly control the force output of the Superlimbs to regulate the amount of assistance they receive.

Figure 4-4 shows the block diagram of the manual force control system that includes a human in the feedback loop. The human detects the reaction force through the haptic feedback loop, compares it against a desired setpoint, and generates a control signal using a joystick of the Xbox controller. The human command is the yank, the time derivative of force, which is fed to the integrator to generate a reference force for the SuperLimb control system. A random noise signal is superposed

to examine how the human can regulate despite the disturbance. The output force at the endpoint of the SuperLimb is measured with the force sensor for closing the inner feedback loop so that the SuperLimb can follow the human command with high fidelity.

The subjects performed the manual force control under three conditions. First, the subjects were asked to perform only a visual task. The visual task consisted of displaying a series of yellow and white letters on the monitor. The subjects were asked to identify whenever they saw a white letter by pressing a button on the Xbox controller. Each trial lasted 20 seconds and there were five trials. In the second condition, the subjects were asked to reject randomized force disturbances applied to SuperLimb such that the force in the vertical direction at the table was regulated at 0 Newtons. Each trial lasted 20 seconds and there were five trials. Finally, during the third condition, the subjects were asked to reject randomized force disturbances applied to the Superlimb and perform the visual task simultaneously. Once again, each trial lasted 20 seconds and there were five trials. The order of the conditions was randomized for each subject. Prior to each condition, the subjects were allowed to familiarize themselves with the conditions and run a few practice trails.

4.1.3 Inherent Haptic Feedback for High Autonomy Supervision of a Superlimb

In the experiment of high autonomy supervision, the SuperLimb performs a programmed task automatically that generates a reaction force and moment acting on the human. While the SuperLimb increases and decreases the force, as shown in Fig.4-6, the human subject is asked to detect specific learned reaction force changes, referred to here as features, by promptly pressing a different button on the Xbox controller for each feature. These features, shown in Fig. 4-6, were taught to the subjects prior to beginning the experiments. The accuracy and delay of the feature detection were evaluated.

In practical applications, SuperLimbs may physically interact with the environ-

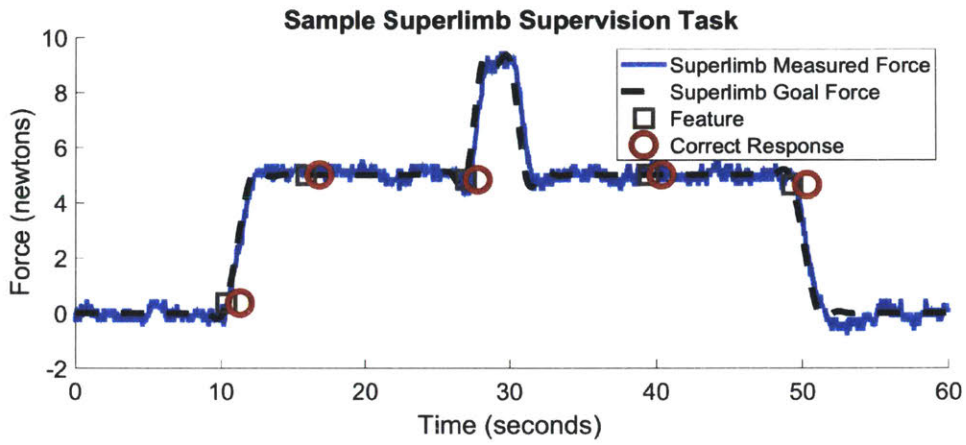

Figure 4-6: Sample High Autonomy Supervision Superlimb Task: The subject's goal was to supervise the state of the Superlimb and, when it performed a pre-described action primitive, to indicate which action primitive it was by pressing a button on the Xbox Controller. The figure on the left shows the Superlimb's force trajectory. Five total features are shown and identified with a red square. They are–from left to right–"start", "buzz", "up", "buzz", and "end". The start feature consists of the force going from 0 to 5 Newtons. The buzz feature consists of the piezoelectric vibration motors buzzing. The up feature consists of the force transitioning from 5 to 9 and back to 5 Newtons. Finally, the end feature consists of the force going from 5 to 0 Newtons. The red circles represent when the subject responded and whether they correctly identified the feature. It is important to note that the "buzz" feature is not reflected in the force measurement but is artificially added here for the purposes of illustration.

ment not only at the endpoint but also at other points along the arm link. SuperLimbs may hit objects unexpectedly, or an external object may hit the SuperLimbs. It is important that the wearer can notice such interactions with external objects in addition to the detection of the reaction force generated by the SuperLimbs. In order to emulate this, the prototype SuperLimb was equipped with a piezoelectric vibrator that emulates the physical interactions with external objects. Vibration forces are generated in the direction perpendicular to the vertical plane where the SuperLimb-generated reaction force acts. The strength of the vibration is small: lower than the noise level of the force sensor. Nonetheless, it can be felt by the human through the inherent haptic feedback. As indicated in Fig.4-6, these vibrations are applied to the SuperLimb while the SuperLimb's goal force is not varying.

Experiments were conducted under three different conditions. During the first

condition, the subjects were asked to perform only the visual task to detect while letters displayed on the monitor–the same one used in the previous experiment. Each trial now lasted 60 seconds, however, and there were five trials. In the second condition, the subjects were asked to detect the haptic features. Each trial lasted 60 seconds and there were five trials. In the third condition, the subjects were asked to simultaneously detect the haptic features while also performing the visual task. Each trial lasted 60 seconds and there were five trials. The order of the conditions was randomized for each subject and each subject was allowed to familiarize themselves with the tasks prior to beginning data collection.

4.2 Results

4.2.1 Low Autonomy Manual Control Experiment

All subjects completed all parts of the experiment. Fig. 4-5 shows a subject's performance during one trial of the third condition during which the subject performed both the visual and haptic task simultaneously. Subjectively, it seems as if the subject was able to control the force output of the Superlimb by closing the loop with direct haptic feedback.

This is backed up by the aggregated data shown in Fig. 4-7A-C. The subject's performance on the Superlimb low autonomy manual control task can be judged most readily by the root mean square error (RMSE) between the measured value, shown in blue in Fig. 4-5, and the goal, shown in black in Fig. 4-5, which stays steady at zero. By this metric, the subjects had a mean RMSE of 1.96 Newtons during the haptic only condition and a mean RMSE of 1.69 Newtons during the simultaneous haptic and vision condition. These data are shown in the form of a violin plot in Fig. 4-7A. A violin plot shows the probability density of the data, as estimated by a kernel density estimator, and thus is more informative than traditionally box plots.

Another way of judging a subjects' performance is to look at the command effort exerted by the subject during the Superlimb task. This is roughly a measure of how

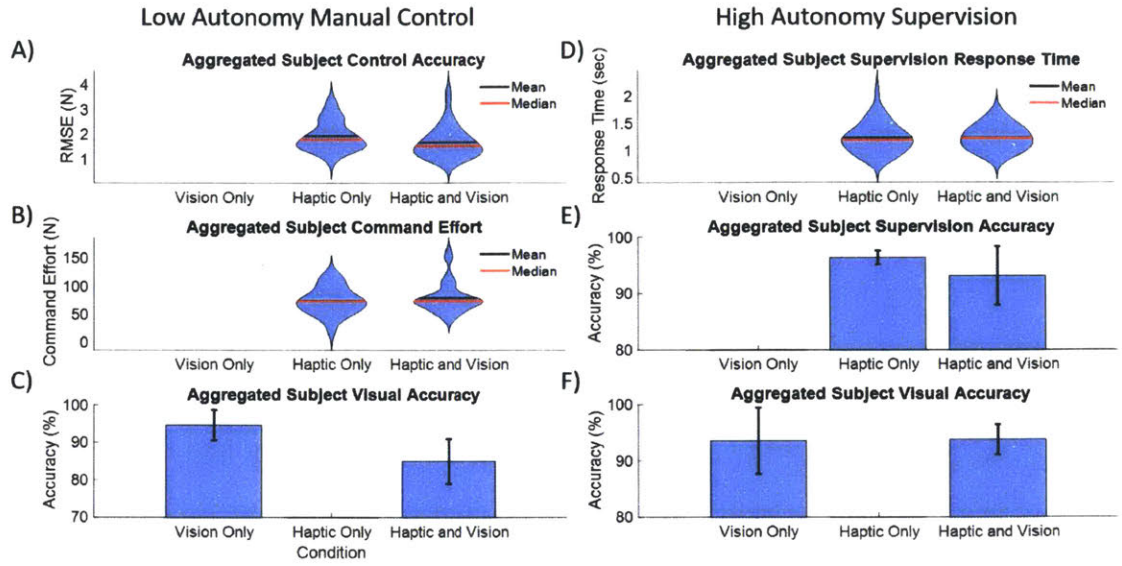

Figure 4-7: Aggregated Experiment Performance: (A) Low Autonomy Manual Control: The root mean square error (RMSE) between the goal–zero force–and the measured force throughout the trial. (B) Low Autonomy Manual Control: The user's command effort, as measured by the integral of all their command inputs. (C) Low Autonomy Manual Control: The subject's percent correct on the visual task. The error bars show the standard deviations. (D) High Autonomy Supervision: The response time to features as measured by the difference in time between when a feature begins and when the user presses a button on the Xbox controller. (E) High Autonomy Supervision: The user's percent correct on the Superlimb supervision task, a measure of accuracy. The error bars show standard deviation. (F) High Autonomy Supervision: The subject's percent correct on the visual task. The error bars show standard deviation.

hard a subject had to work in order to reject the disturbance. The command effort can be calculated by taking the time integral of the subject's command inputs, thus giving units of Newtons. The aggregated command efforts are shown in Fig. 4-7B. During the haptic only condition, subjects' mean command effort was 74.0 Newtons while during the haptic and visual task condition, subject's mean command effort was 73.8 Newtons.

Finally, Fig. 4-7C shows the subjects' performance on the visual task. Performance on the visual task was graded based upon the accuracy with which the identified each white letter. The data is displayed as a bar graph rather than as violin plots as the probability density of the data estimated by the kernel density estimator

resulted in probability densities that were impossible; for example, a percent correct over 100%. During the visual task only condition, the subjects achieved a mean accuracy of 94.6% while during the haptic and visual task condition, the subjects achieved a mean accuracy of 85.0%.

With these data, we can first assess whether the inherent haptic feedback provided by the Superlimbs was sufficient to control the force output of the Superlimb. This can be done subjectively; it appears as if subjects were able to respond the disturbance and act to cancel it out, as shown in Fig. 4-5. Additionally, we can compare the subjects aggregated RMSE to what their RMSE would have been had they taken no action; this equates to the RMSE error generated by disturbances. Thus this can be found by taking the cumulative time integral of the disturbance and then finding the root mean square error of this trajectory. Using this, we find that the mean RMSE would have been 114 Newtons if the subjects had taken no action during the haptic only condition. Finally, we can test this using statistical analysis to ask there is a statistically meaningful difference in performance between the no action case and the subjects' performance as measured by RMSE. We used a two-sample t-test with the first population coming from the RMSE errors generated by the disturbances during the haptic only condition and the second population coming from the RMSE errors actually achieved by the subjects during the haptic only condition. The null hypothesis is that there is no difference between the means of these two populations while the alternate hypothesis is that the mean RMSE errors generated by the subjects was smaller. Applying this test, we found evidence ($t(16)=7.71$, $p < 0.01$) that the subjects performed better than if they had taken no action during the haptic only condition as measured by RMSE. This suggests that the subjects could successfully close the loop and control the force output of the Superlimb.

Next we can assess whether the addition of a simultaneous visual task caused a difference in performance on the haptic task. To do so, we can compare the performance data on the haptic task from the haptic only condition–RMSE and command effort–to the performance data on the haptic task from the haptic and vision condition. We used a repeated-measures multi-variate ANOVA and found no evidence

that the subjects performed differently on the haptic task as measured by RMSE and command effort despite the addition of a vision task in the haptic and vision condition ($f(1,16) = 0.43$, $p = 0.52$). It is important to note that this doesn't mean that we can conclude that there was no effect on performance when performing both the visual and haptic task simultaneously but simply that we found no statistical evidence of there being a performance degradation within the context of this experiment.

Finally, we can assess whether the addition of a simultaneous haptic task caused a difference in performance on the vision task as measured by visual task accuracy. To do so, we compared the performance data on the vision task from the vision only condition to the performance data on the vision task from the haptic and vision condition. We used a matched-pair t-test with the null hypothesis being that the means are equal and the alternate hypothesis being that the means are different. We found that we could reject the null hypothesis ($t(8) = -2.63$, $p = 0.03$). In other words, we found evidence that the subjects performed worse on the visual task when they were asked to do both the visual and haptic task simultaneously as compared to just doing the visual task. This follows our intuition; we expect it to be harder to perform two tasks simultaneously versus performing a single task.

Interestingly, we find there was a slight, though non-statistically significant, improvement in performance as measured by RMSE from the haptic only condition to the haptic and vision condition. This is counter-intuitive; we would expect performance to either stay the same or decrease when adding additional tasks. It is important to note that this performance decrease may not be operationally significant. That is, the difference in the means between the two conditions is ultimately 0.27 Newtons. At the subject, this would be felt as a force difference of roughly 1.25 Newton. While quite a small magnitude change, the just noticeable difference (JND) for human force sensing has been found to be 7% change in force regardless of test conditions [53]. The difference here is roughly a 15% change so it is likely a noticeable, though perhaps not significant, difference in haptic feedback to the subject.

4.2.2 High Autonomy Supervision Experiment

Once again, all subjects completed all parts of the experiment. Fig. 4-6 shows a sample performance of one subject on the haptic task during the third condition, performing both the haptic and vision task simultaneously. During this particular example trial, the subject correctly identified each feature.

We can judge performance on this task by looking at the response time and accuracy with which the subjects identified each haptic feature. These results are shown in Fig. 4-7D and Fig. 4-7E respectively. On average, subjects responded to a haptic feature in 1.24 seconds during the haptic only condition and 1.24 seconds during the simultaneous haptic and vision condition. Subjects displayed good accuracy as well, achieving 96.3% accuracy in identifying the haptic features during the haptic only condition and 93.1% accuracy during the simultaneous haptic and vision condition.

Fig. 4-7F shows the subjects' performance on the visual task as graded by accuracy. The subjects showed good accuracy in both the vision only condition and the simultaneous haptic and vision task, achieving mean accuracies of 93.5% and 93.8% accuracy respectively. Note that the same data was used for the vision only condition accuracy on both the low autonomy manual control and high autonomy supervision experiments as the task was identical in both cases.

Using these data, we can once again assess whether performance degraded when the subjects were asked to perform both the visual and haptic tasks simultaneously. Once again, because we are asking multiple questions with dependent data, we will first use a repeated measures multi-variate ANOVA to see if there was any difference in performance on the haptic task between the haptic only condition and the haptic and vision condition. Applying this statistical test, we found no evidence that the subjects performed differently on the haptic task as collectively measured by haptic task accuracy or haptic task response time ($f(1,16) = 0.13$, $p = 0.73$). Second, we can compare performance on the visual task during the vision only condition to the performance on the visual task during the haptic and vision condition using a matched-pair t-test. The null hypothesis is that there is no difference between the

Table 4.1: Aggregated performance on the supervision task by feature

Feature	Accuracy	Response Time
"Start"	96%	1.23 sec
"Up"	97%	1.17 sec
"Buzz"	97%	0.98 sec
"End"	89%	1.61 sec

performance in the haptic only and haptic and vision conditions and the alternate hypothesis is there is a difference in the means of the visual accuracy between the two conditions. Using this statistical test, we found no evidence that subjects performed differently when performing both the vision and haptic task as compared to just the vision task as measured by accuracy on the vision task ($t(8) = 0.13$, $p = 0.90$).

Finally, Table 1 breaks down the performance metrics on the haptic task during the high autonomy supervision task by feature. The results are aggregated across both the haptic only condition and the haptic and visual condition.

4.3 Discussion

Simply by wearing and operating a Superlimb, the wearer is given inherent haptic feedback on the current state of the Superlimb. We sought to demonstrate that this inherent haptic feedback was sufficient both for placing the wearer in the loop to control the force output of the Superlimb–a low Superlimb autonomy task–and for supervising the actions of an Superlimb–a high Superlimb autonomy task.

Within the low autonomy manual control experiment, which is similar to how the Superlimb would be used during standing assistance tasks, the haptic task only condition demonstrated that, when given manual control of the force output of the Superlimbs, the subjects were successfully able to control it. This suggests that Superlimbs may be a good fit for standing assistance or similar tasks.

Interestingly, adding a simultaneous visual task led to a small increase in performance on the haptic task as measured by RMSE on the haptic task. However, the

subjects showed a correspondingly large decrease in performance on the visual task when adding a simultaneous haptic task as measured by accuracy on the visual task. We think that this result supports the visuospatial sketchpad model of short-term memory proposed by Baddeley and Hitch. Namely, the visuospatial sketchpad is thought to incorporate both visual and tactile inputs [43] and has a limited amount of resources to incorporate new sensory information [28]. During the simultaneous visual and spatial task condition, we think that the subjects were overloaded with sensory input and thus their performance degraded [43]. For some tasks that do not have strict time or space constraints, such as body bracing during heavy industry work, the performance degradation may be acceptable. However, for time or space constrained tasks such as acting as a helping hand during robotic surgery, it points to either the need for additional forms of feedback on the actions of the Superlimb or an increase in the autonomy of the Superlimb. Future work could investigate whether this performance degradation disappears with additional practice with the Superlimbs and whether it was in fact due to overload in the first place.

During the high autonomy supervision experiment, we found no statistical evidence of a performance degradation on either the haptic or visual task when the subjects were asked to perform both tasks simultaneously. However, we did see a non-statistically significant decrease in supervision accuracy on the haptic task. Given these data, the results suggest that the inherent haptic feedback from Superlimbs is sufficient to supervise the autonomous actions of the Superlimb. This is important as it helps the wearer during normal operating conditions, such as helping the wearer to collaborate with the Superlimb were it acting autonomously to open a door, and during abnormal operating conditions, such as informing the user of Superlimb failures, thus allowing the wearer to take actions to ensure their and the Superlimb's safety. Further, these results agree with the results from the low autonomy manual control experiment; by increasing the autonomy of the Superlimb we were able to minimize the performance degradation associated with multitasking.

Diving a little deeper, Table 4.1 allows us to draw two conclusions. First, subjects seemingly performed roughly equivalently on the "start", "up", and "buzz" features

but performed worse on the "end" feature. It is unclear why this is the case and is perhaps an area of future work to understand which feature types subjects are better able to identify. Second, subjects identified the "buzz" feature both with a high accuracy and quick response time. This is despite the fact that the Superlimb was not explicitly measuring this state; the Superlimbs force measurement, the only sensor on the system, did not meaningfully pick up the buzz feature. While it is certainly not surprising that the subjects achieved good performance on the "buzz" feature as vibratory motors are often used in haptic displays, it does nicely illustrate a difference between the inherent haptic feedback from Superlimbs and traditional haptic displays. Because the Superlimb does not need to be explicitly aware of a state for the wearer to get haptic feedback on it, we can argue that the inherent haptic feedback from Superlimbs is more robust than traditional haptic displays to failure. Additionally, even during normal operation, the inherent haptic feedback from Superlimbs allows the wearer to get information on states that the system designer may not have recognized as important to functioning; when using an Superlimb to open a door, the ability to haptically sense that the Superlimb's end effector has made contact with the door handle can be useful feedback for the wearer.

Despite these advantages, it is important to note the disadvantages of inherent haptic feedback. Namely, given that the communication is purely through forces and torques about the base of the Superlimb, different inputs may result in the same haptic feedback; a change in position of the end effector versus a change in force at the end effector could yield identical feedback to the operator.

It is important to note that these results represent performance from subjects with very little experience using the Superlimb system. Between both experiments, the subjects wore the Superlimb for roughly 1.25-1.5 hours. We expect that performance would increase with further practice. However, as performance is also dictated by the task and haptic feedback setup, we chose not to seek a plateau in performance as this would only represent the performance plateau within this individual context rather than a global performance plateau. Thus, we chose to limit the study to one session.

While this study has introduced the concept of inherent haptic feedback and

argued for its usefulness during both low and high autonomy operations, there remains much work to be done. Beyond the future works already suggested, there is no doubt room to optimize the design of the Superlimb system for delivering inherent haptic feedback to the wearer. One could imagine a design process that takes into account the desired taskset the Superlimb will perform to inform the amount and type of contact points the Superlimb makes with the human body. In the current waist mounted system, all of the feedback was delivered to this one area of the body, thus making it difficult to resolve Superlimb state changes that resulted in identical forces and torques being reflected on the body. A design with more contact points, perhaps similar to the design of the Superlimb for overhead assembly, may allow the wearer to resolve these states [9]. Finally, these contact points could be tuned to the deliver forces that are tuned to the sensitivity and range of a particular body part, drawing on works from literature [53].

4.4 Conclusion

This paper offers the first investigation into the inherent haptic feedback provided by Superlimbs to the wearer of the system. We explore the value of haptic feedback under both low autonomy situations–where the wearer is manually controlling the actions of the Superlimb–and under high autonomy situations–where the wearer is supervising the actions of the Superlimb. First we show that, within the context of this experiment, the subjects were successfully able to manually control the force output of the Superlimb with haptic feedback alone. This suggests that the Superlimb could be applied to tasks like assisting the elderly or injured transitioning from sitting to standing as this task lends itself to manual control of the force output of the Superlimbs. Secondly, we show that the inherent haptic feedback from Superlimbs was sufficient for the subjects to supervise the autonomous actions of the Superlimb, even in the presence of a simultaneous visual task. This suggests that the inherent haptic feedback can be used by the wearer to track abnormal activities by the Superlimb and step in when necessary to ensure the safety of the wearer. While these results

demonstrate clearly the value of the inherent haptic feedback from Superlimbs, they also suggest further studies on how to design Superlimbs to optimally communicate the inherent haptic feedback to the wearer.

Chapter 5

Conclusion and Future Work

5.1 Additional Incomplete Studies or Unanticipated Outcomes

In addition to the work performed in completion of this thesis, and in an effort for full transparency, the following section presents some incomplete studies or studies with unanticipated outcomes to offer directions forward.

5.1.1 Inherent Haptic Feedback for Spatial Information

Vision is generally offers superior accuracy as compared to the somatosensory system to determine the spatial position of the natural limbs [15]. Similarly, we can hypothesize that vision is generally better at assessing the spatial position of Superlimbs as compared to haptic feedback. This is backed up by a wealth of haptic feedback for prosthetics research. However, similar to this wealth of research, spatial haptic feedback still can prove useful for Superlimbs when vision is otherwise occupied. Thus, in this experiment, we set up a situation where the person had two tasks, one that could only be accomplished with vision and the other that could be accomplished with both vision and haptic feedback. This would serve to naturally "occupy" the vision of the subject, thus incentivizing them to use haptic feedback to solve the second task. We varied the position of the haptic feedback task such that in some conditions the

person could rapidly visually assess both tasks while in other conditions the person was rate-limited by the need to turn their head, thus adjusting their field of view, in order to visually assess the haptic task.

We conducted human subject experiments based on the protocol approved by the Massachusetts Institute of Technology Committee on the Use of Human as Experimental Subjects (COUHES), number 1805387660. The subjects (N = 9) were healthy, right-handed, and in good physical shape ($20 \leq \text{BMI} \leq 30$).

Experimental Setup

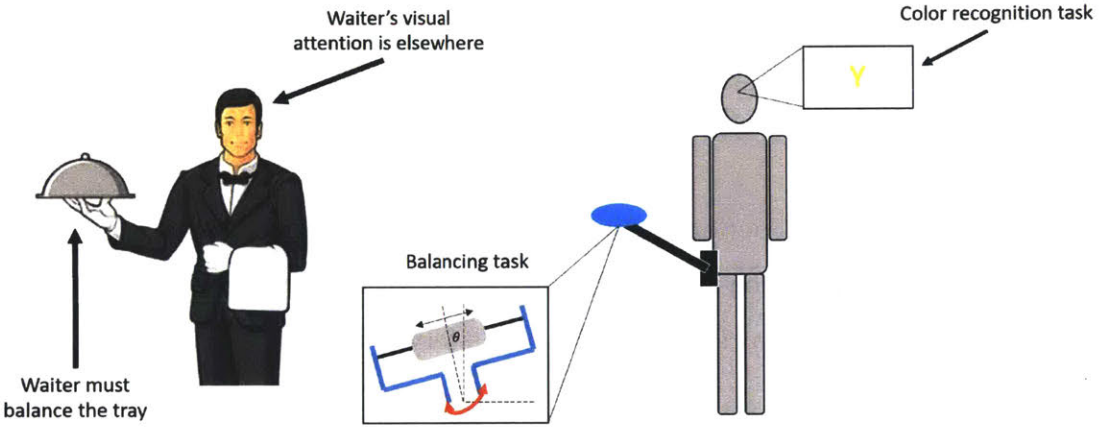

Figure 5-1: Experimental Inspiration: Drawing inspiration from a waiter, the subject is asked to perform two tasks. The Color Recognition Task asks the subject to identify when they see a white letter during a Rapid Serial Visual Presentation. The Balancing Task asks the subject to maintain the balance of a weight on the linear slide while small, computer-generated disturbances are applied to the weight that, without human intervention, would cause the weight to move toward the ends of the linear slide, thus unbalancing the Superlimb.

In this study we sought to determine the effectiveness of haptic feedback in communicating spatial information about the Superlimb. We provided the subjects with two tasks. The first task is the Color Recognition Task (CRT). In the color recognition task, the subject is shown a Rapid Serial Visual Presentation (RSVP). During the RSVP, the subject is presented with a series of rapidly changing yellow and white letters. The subject is asked to detect the white letter by pressing a button on an

Xbox gaming controller. By varying the ratio of white and yellow letters, as well the time for which each letter is presented, we can vary the difficulty level of the CRT.

The second task is the Balancing Task (BT). We designed and built a Superlimb that consisted of a weight mounted on a linear slide (Fig. 5-2A). We tasked the subject with keeping the weight balanced in the center of the linear slide. We provided randomized, computer-generated small disturbances to the weight that, if no action was taking by the user, would cause the weight the reach the end of the linear slide, thus unbalancing the Superlimb. These disturbances are shaped like Gaussian bells. The subject was given control of the velocity of the weight through the joystick on the Xbox gaming controller. The commanded velocity of the weight was the average of the randomized small disturbances and the subject's commands via the joystick. This gives the subject the ability to cancel out the disturbance and keep the weight centered on the linear slide. The difficulty of the BT can be varied by varying the amplitude and duration of the randomized small disturbances.

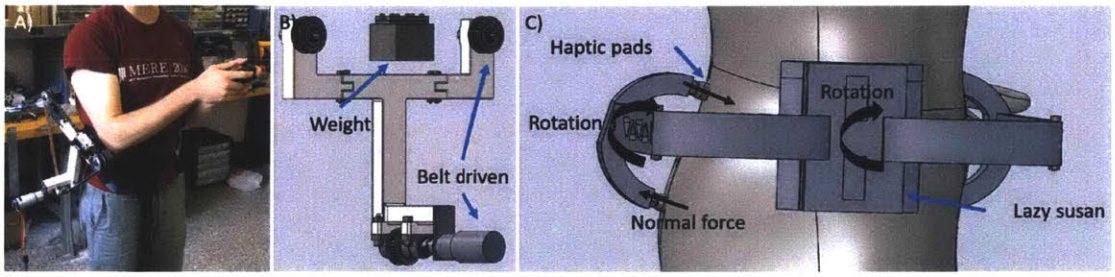

Figure 5-2: The balancing task mechanism. (A) A user wearing just the Superlimb without the haptic feedback system. (B) A CAD model of the balancing mechanism. The weight is driven along the linear slide with by a belt drive. (C) The haptic feedback setup. As the weight from (B) translates along the linear slide, the entire Superlimb rotates, as shown in (A). The haptic feedback setup turns the rotation of the Superlimb into a pair of normal forces applied to the subjects back via the haptic pads. The subject can determine the rotation of the Superlimb by the differential of the force between the two haptic pads on their back.

Fig. 5-2B shows a closer look at the Superlimb. The weight is mounted on a belt between two pulleys and the entire belt system is driven by a motor located at the base of the Superlimb (Fig. 5-2B). This entire setup is mounted to the subject via a 3D printed base. A lazy susan is placed between between the base of the Superlimb

and the rest of the Superlimb. The result of this is that as the weight moves along the linear slide, the entire Superlimb rotates about the lazy susan, thus simulating the tray becoming unbalanced.

We harnessed the rotation about the lazy susan to deliver haptic feedback to the subject on the rotation of the Superlimb. Fig. 5-2C shows the haptic feedback device. Through this device, the rotation of the Superlimb about the lazy susan is translated to a translation of two haptic pads placed on the subjects back. As the lazy susan rotates, the force differential between the two haptic pads changes, thus giving the subject haptic feedback on the rotation of the Superlimb. The haptic pads are spring-loaded into the user to provide constant feedback.

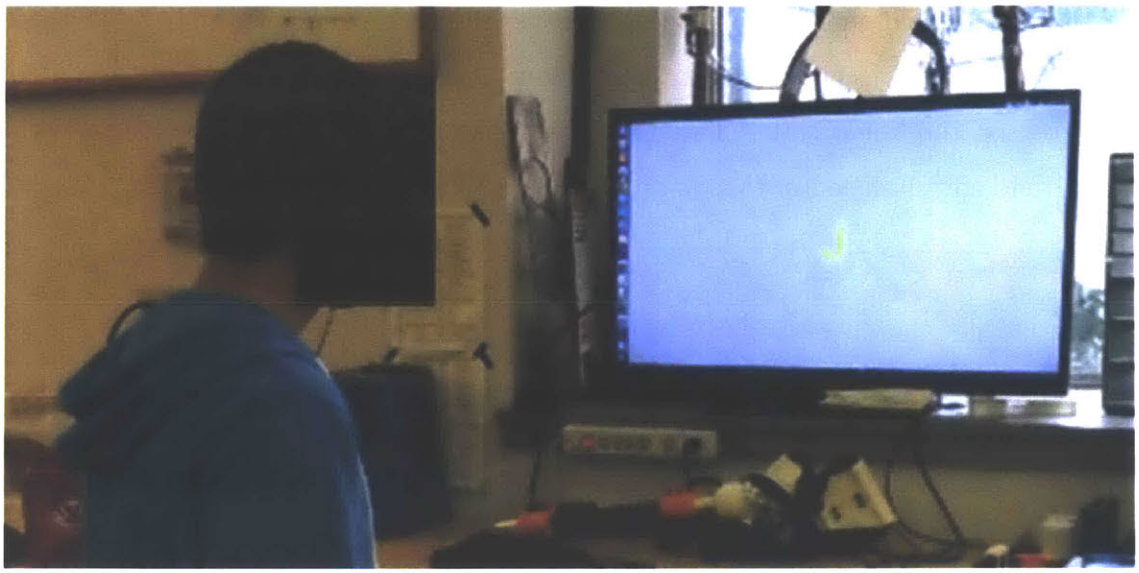

Figure 5-3: The Color Recognition Task is placed on a large monitor in front the user.

The Color Recognition Task is shown in Fig. 5-3. A large monitor is placed in front of the subject. A rapid series of yellow and white letters are shown on the monitor. The subject is asked to identify when they see a white letter by pressing a button on an Xbox gaming controller.

In order to demonstrate that haptic feedback can compensate for a limited field of view, we must test the subjects on five distinct conditions as illustrated in Fig. 5-4.

During the first condition, the subject will perform only the Color Recognition Task (CRT). This condition is designated the color recognition only condition (CRO)

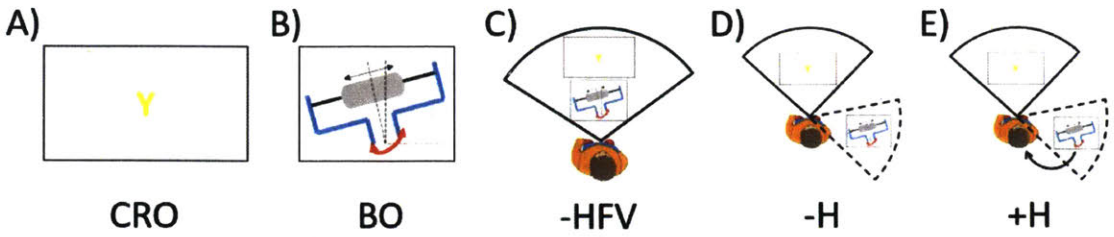

Figure 5-4: A pictorial representation of the five experimental conditions. The black outlined semi-circle represents the subjects field of view. In D) and E) the subject cannot see both tasks within the same field of view.

(Fig. 5-4A). During the second condition, the subject will perform only the Balancing Task (BT). This condition is designated the Balancing Only (BO) condition (Fig. 5-4B). During the third condition, the subject will perform both the CRT and the Balancing Task (BT) while not wearing the robot. The CRT and BT will be placed such that they are both in the subject's field of view (that is, the subject can see both tasks without moving their head). This condition is designated the field of view but no haptic feedback condition (-HFV) (Fig. 5-4C). During the fourth condition, the subject will perform both the CRT and BT without wearing the robot but the tasks will be placed such that the tasks are not both simultaneously in the subject's field of view. This is known as the no haptic feedback condition (-H) (Fig. 5-4D). Finally, during the fifth condition the subject will perform both the CRT and BT which are placed such that the tasks are not both simultaneously in the subject's field of view. However, during this final condition, the subject wears the robot and therefore receives haptic feedback on the position of the weight on the linear slide. This condition is designated the haptic feedback condition (+H) (Fig. 5-4E).

Results

An example trial from one subject is shown in Fig. 5-5. The figure shows the performance from the subject on both the Color Recognition Task and the Balancing Task for -H and +H conditions.

The subjects were graded along a number of different metrics for each condition.

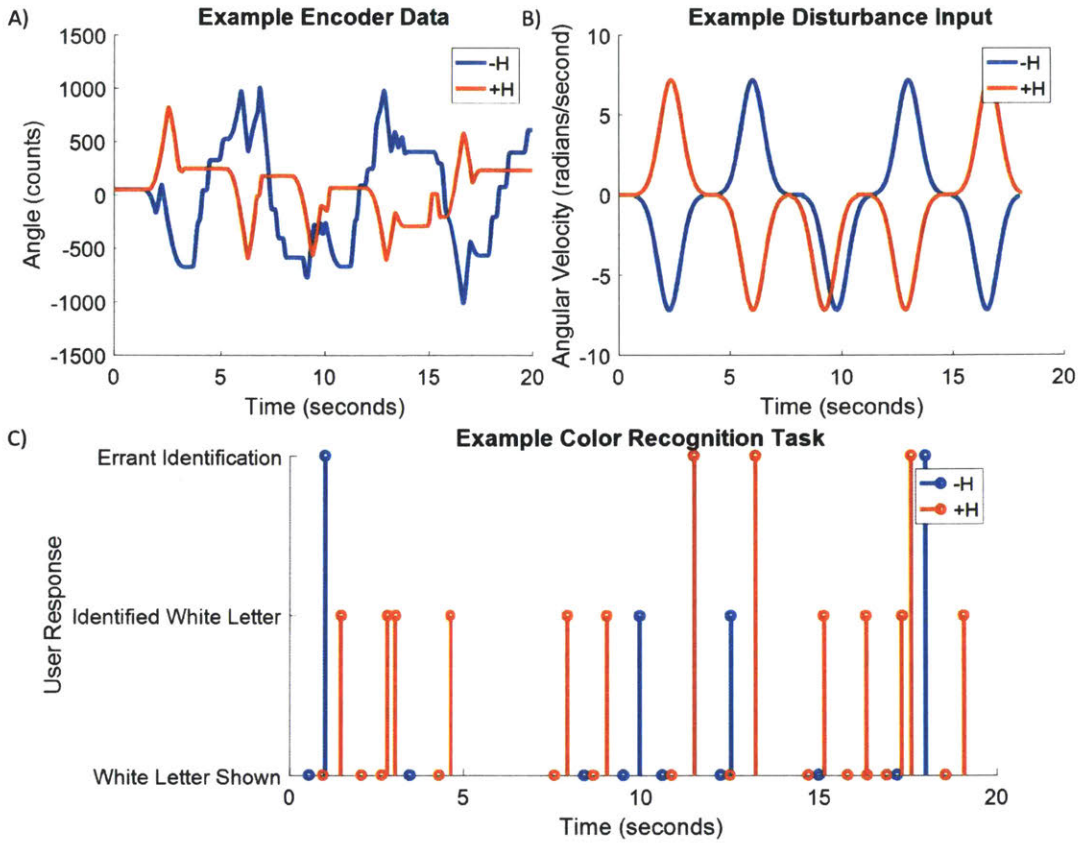

Figure 5-5: Example trial for both the no haptic (-H) and haptic (+H) conditions from the spatial feedback experiment. A) shows the encoder reading for both conditions. B) shows the disturbances associated with these trials. C) shows the color recognition task. The dots on the bottom show when a white letter was displayed. The middle represent when a white letter was correctly identified while the top line shows when there were errant button presses.

For the Color Recognition Task, errors were counted when the subject failed to press the button when there was a white letter or when the subject pressed the button when there was no white letter (Fig. 5-6E). For the Balancing Task, errors were counted whenever the weight reached either end of the linear slide, as determined by limit switches placed at the ends of the linear slide (Fig. 5-6B). This is denoted "dropping the tray". Additionally, we determined the sum of the absolute error from the center point for the balancing task (Fig. 5-6A), the user command effort on the balancing task (Fig. 5-6C), and the user response time to each disturbance on the balancing task (Fig. 5-6E). All of these results are shown in Fig. 5-6 using violin

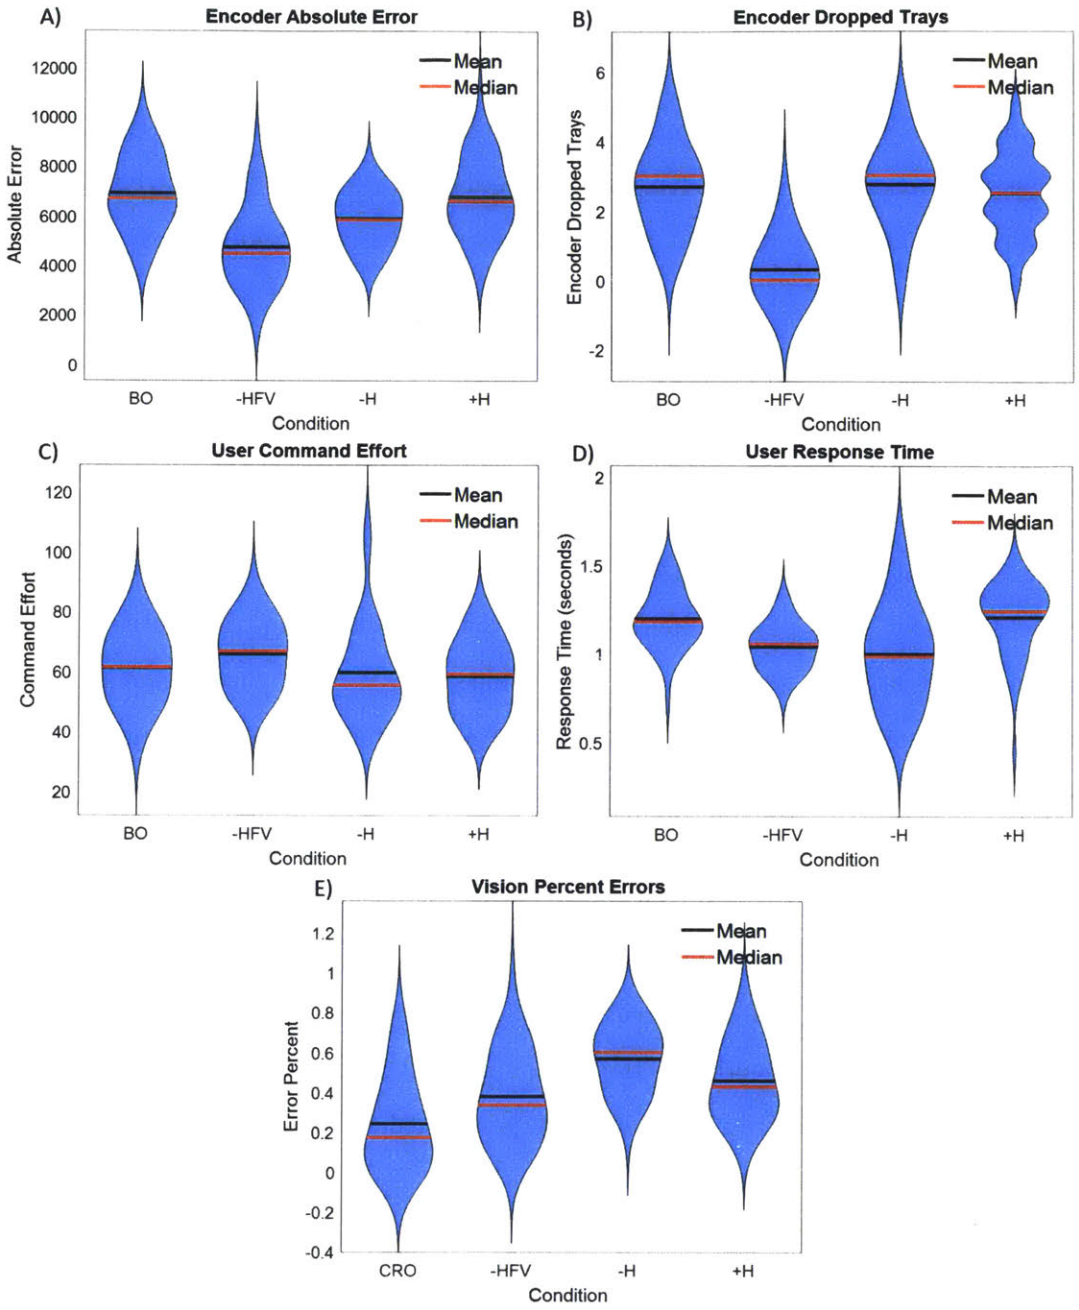

Figure 5-6: All results from the spatial feedback experiment. (A-D) show metrics taken from the balancing task while (E) shows metrics taken from the color recognition task.

plots. Violin plots are similar to box plots but they include a rotate kernel density plot on each side, thus showing an estimation of the probability density function. This shows significantly more information than a box plot, which often only include the

mean and standard deviation. However, because the probability density function is estimated, it may show impossible values, such as a negative number of errors for the balancing task. For each condition, the same difficulty parameters were swept. This allows us to compare errors in both the CRT and BT across the different conditions.

While there are many different pieces of data that we can highlight, we will focus on a few results. For all of the results that use p-values, we have applied a matched-pair, one-tailed, t-test.

Firstly, we can look at Fig. 5-6E to assess the subjects' performance on the Color Recognition Task. From the figure, its clear that users perform best when only performing the vision task, followed by when they had both tasks within their field of view, followed by the two conditions where the balancing task was placed out of their field of view. This pattern holds true for both the dropping tray metric and the absolute error metric on the balancing task. This suggests that, when vision is readily available, the subjects performed best (testing all of these hypotheses individually, they are all statistically significant with $p < 0.01$).

Next we can compare the performance specifically between the -H and +H conditions. First looking at the performance on the Color Recognition Task. In this we see when subjects had haptic feedback on the spatial position of the weight, they outperformed those who did not have haptic feedback ($p < 0.01$). However, subjects performed roughly equivalently on the balancing task when measured by number of times they "dropped the tray" and actually performed worse when measured by the sum of absolute encoder error ($p < 0.01$).

This is, on the surface, a counter intuitive result. Despite the additional information provided on the balancing task, the subjects performed better on the vision task. Further, they did not perform better on the balancing task. In fact, according to at least one metric, they performed worse. I hypothesize that these results stem from the strategy that the subjects adopted when given access to the haptic feedback. When the subjects did not have haptic feedback during the -H condition, the subjects rapidly moved their heads back and forth between the two tasks. By changing their field of view, the subject was able to get visual feedback on both tasks intermittently.

While this intermittent vision was enough to respond to the disturbances on the balancing task (which occurred over a couple seconds), it was not enough to keep errors low on the vision task (where the white letters only appeared for a fraction of a second). However, when the subjects were provided haptic feedback on the BT (+H), the subjects were able to direct their visual attention almost solely on the CRT. This decreased the errors on the CRT. But, as the subjects were relying almost solely on the haptic feedback to perform the BT, and as haptic feedback is far more coarse than visual feedback, the subjects performed either similarly or worse on the BT as compared to the no haptic feedback condition.

One final observation can be drawn from this experiment. Looking at Fig. 5-6D, we can see that, on average, subjects responded more slowly in the +H condition as compared to the -H and -HFV conditions ($p > 0.01$). This likely stems from an experimental design problem. The haptic feedback was delivered via the haptic pads on the subjects back shown in Fig. 5-2C. The mapping between the position of the weight on the linear slide and the force delivered by the haptic pads is sinusoidal. Thus as the weight moves toward the end of the linear slide, the sensitivity (the change in force per change in position) increases. Given this, and given the relative coarseness of the haptic feedback, it is advantageous for the subject to wait till the weight has moved significantly in one direction or another before responding. This hypothesis is supported by the fact that, on the first day of the experiment (and prior to subjects developing this strategy), subjects showed a significantly faster response time in the +H condition as compared to the -H condition. This suggests that haptic feedback may be useful for delivering temporal feedback.

In summary, while haptic feedback statistically significantly improves subjects' performance on the vision task and at best shows no improvement on the balancing task, it is not operationally significant. That is, within the context of this experiment, it does not improve performance across both tasks enough to make it a reasonable solution for someone to use. Interesting, and perhaps dishearteningly, while there exist works, including those highlighted here, that show promising results for delivering spatial information haptically for Superlimbs or prosthetics, to our knowledge no

study has ever expanded this to multiple degrees of freedom. This may be because, qualitatively, mapping the haptic feedback into spatial information required a good deal of attention for a single DOF. In the end though, haptic feedback for spatial information does not seem to be a promising direction, at least as presented here.

However, haptic feedback does show promise for temporal feedback. When vision is intermittent, as is the case in the -H case, the variance on the user response times is very large. Thus if it is critical that someone respond consistently on time, intermittent vision is not an acceptable solution. Haptic feedback can guarantee that the user will respond more consistently.

5.1.2 Gaze Tracking for Communication with Superlimbs

Direct teleopoeration of a high degree of freedom robot with a low dimension input device can lead to high cognitive and physical loads. Shared control will often be introduced to allow users to complete their desired action with less control input. As the robot does not know a priori the goal of the human, shared control requires a prediction of the user's goals. While many techniques have been applied to develop a prediction of the user's goals, one promising direction is eye gaze tracking [2]. It has been found that eye gaze during a task is rarely directed outside of the objects required for the task and that hand movements are delayed until the eye is available for guiding the movement [21] [41].

Given this, we proposed an architecture that relies on gaze tracking to determine when to switch control authority of the Superlimbs between the human and an autonomous algorithm (Fig. 5-7. This is built on an intermittent control architecture where the human has ultimate control authority. However, rather than simply having control authority whenever the human issues a non-zero input command to the Superlimb, the human only has control authority when they direct their gaze at the Superlimb.

In order to demonstrate this concept, we set up a situation where the Superlimb was assisting the human by moving parts from a side table, shown in Fig. 5-8, to the front table where the person was working. However, the Superlimb does not know

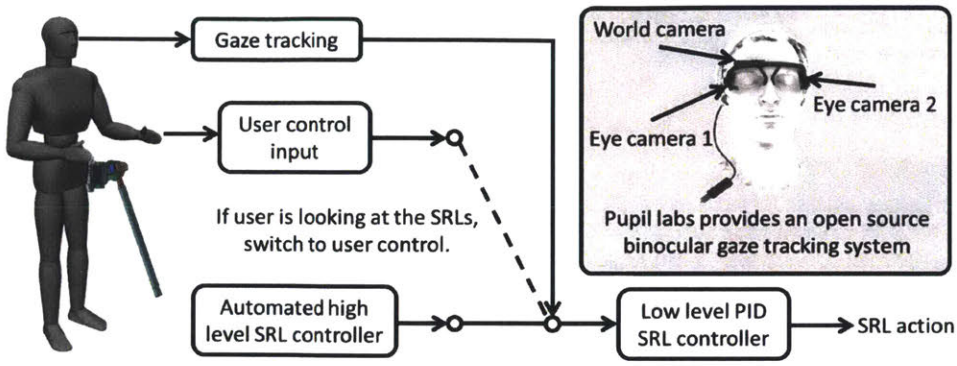

Figure 5-7: The Superlimb (Superlimb) will use open source software to track the gaze of the user. When the user turn their gaze toward the Superlimb, the high level control of the Superlimb will be turned over to the user. Otherwise, the Superlimb will act autonomously.

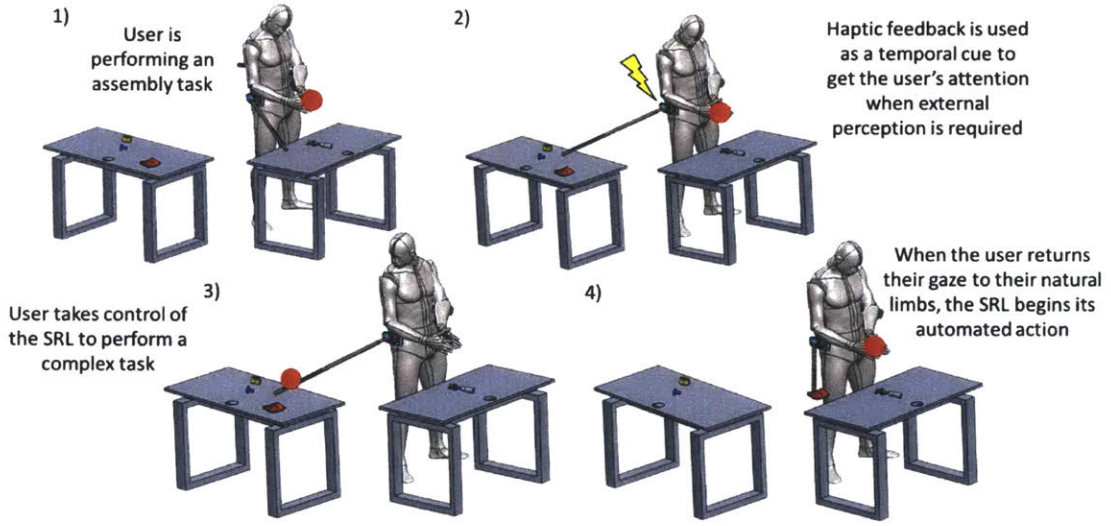

Figure 5-8: Gaze tracking for switching control authority: The red circle represents the location of the gaze of the human operator. When the human operator directs their gaze at the Superlimb, the human operator takes control of the Superlimb and can direct its actions.

which part the human operator needs next. Thus, it must seek input from the human operator on which part to grab next. When the Superlimb needs direction, it cues the human with a haptic signal and then waits for the human operator's input, as shown in Fig. 5-8.

For ease of implementation we used an Xbox controller to send commands to the Superlimb when the human operator had control authority and April tags for tracking

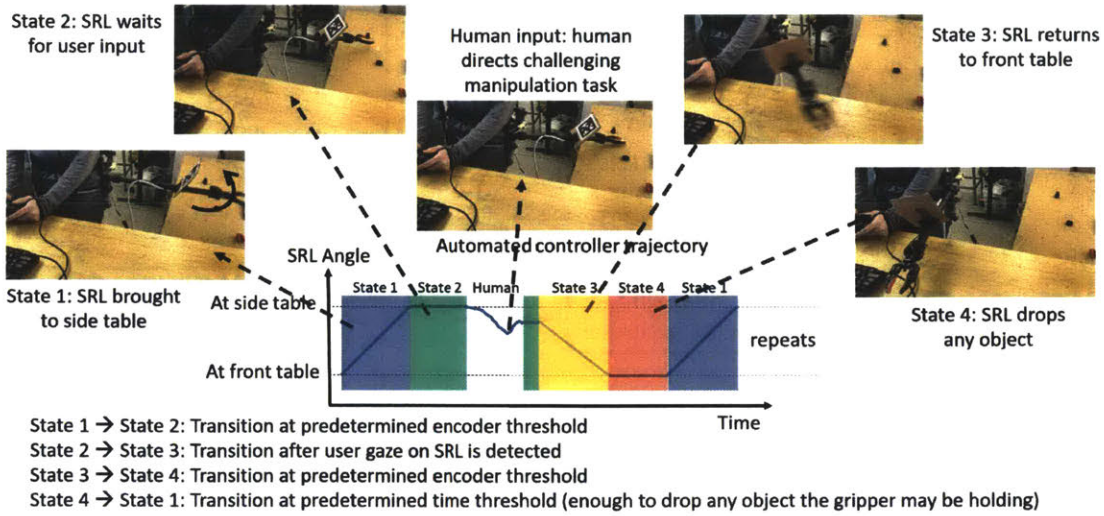

Figure 5-9: State machine controller for the Superlimb (SRL). The control authority is alternated between the human operator and the state machine controller at the human operator's discretion.

the position of the Superlimb position. We built a state machine controller for the Superlimb (Fig. 5-9). During State 1, the Superlimb moves to the side table. During State 2, the Superlimb cues the human operator with haptic feedback and then waits for the human operator to give a command input. In practice, this required that the human operator take command authority of the Superlimb for a minimum period of time. State 3 returns the Superlimb the front table. Finally, State 4 opens the gripper of the Superlimb to deliver any object it may have.

By using gaze tracking to switch control authority, the human operator can naturally take control of the Superlimb at their discretion. This allows the human to maintain their safety; if the Superlimb is acting incorrectly, the human can either alter or stop the Superlimb's actions. Further, it allows the Superlimb to get direction from the human operator during challenging portions of the task while limiting the distraction to the human operator during easy portions of the task.

While this direction appears promising, it is at best only a portion of the solution to intuitive communication with Superlimbs. First, because gaze tracking was used only as a switch to determine who has control authority of the Superlimb, it does not fundamentally solve how those commands get sent to the Superlimb. One could

imagine once again using gaze tracking to determine the human's desired action for the Superlimb; if the human operator looks first at the Superlimb and then at the red block, the Superlimb should go grab the red block. However, this potentially restricts the human operator's gaze unnaturally–what happens if they accidentally look at the orange block–and leaves the Superlimb to navigate picking up the red block with minimal help from the human. Second, in this implementation we used a distance threshold between the location of the gaze in the image and the location of the April tag in the image to determine who had control authority. This is obviously not robust to movements of the gaze around the image. Ideally, we would want the human operator to be able to maintain control authority of the Superlimb while they look at the red block just prior to commanding the Superlimb to go grab it rather than immediately handing control authority back over to the Superlimb's autonomous controller. Thus, a pure spatial threshold is not a satisfactory solution. Despite these limitations, gaze tracking as a means of dictating control authority seems to be worth pursuing.

5.2 Conclusion

Supernumerary Robotic Limbs, or Superlimbs for short, show great promise for augmenting the human operators capabilities by helping with holding, grasping, and manipulating objects as well as supporting or bracing the huma body. In the seven years since their invention, the Superlimb research community developed Superlimbs that are mounted to the upper back, lower back, wrist, shoulder, waist, and upper arm. Superlimbs have been built from traditional materials, 3D printed parts, or using soft robotic systems. Superlimbs have been controlled with foot motion, EMG signals from various muscles, and inertial measurement units or stretch sensors that measure the current state of the human operator. These advances have demonstrated clearly the promise and challenges associated with developing Superlimbs.

However, to date, little work has considered the human operator and their influence on the Superlimb-human system. And yet, as we have seen from other neigh-

boring fields including the exoskeleton, prosthetic, and collaborative robotics research communities, the human operator must be considered in order for Superlimbs to reach their potential. The work within this thesis, therefore, represents one of the first attempts to understand and, where possible, exploit the contributions of the human operator to the Superlimb-human system.

In Chapter 2, we explored whether fully manual control of the Superlimbs was a viable control strategy. We found evidence that when asked to perform four simultaneous and independent tasks, two with the natural limbs and two with the robotic limbs, the human operator performed worse as compared to when they were only asked to perform two simultaneous and independent tasks with their natural limbs. This is relatively unsurprising and yet impactful; given a demanding task, such as the one presented here, Superlimbs would benefit from increased autonomy. We also found evidence that the human operator preferentially moved their natural limbs before moving their robotic limbs. First, this suggests a leader-follower control paradigm where the Superlimb autonomy watches the human's actions to determine its own actions. Second, this suggests that the primary, more time-sensitive sub-task should be given to the natural limbs. Future work could explore whether fully manual control is viable for other types of tasks or with other command input mechanisms.

In Chapter 3, we introduced two new methodologies that leveraged the high degree of redundancy and flexibility of the human body to address key challenges in the design and communication with Superlimbs. First, we detailed a design methodology for Superlimbs that relied on a taskset analysis and empirical measurement of the human operator to split the control space such that the human operator provides the degrees of freedom that they can while the Superlimb actuates the remaining degrees of freedom required to perform the task(s). This results in a lightweight, reduced-actuator Superlimb. Second, we outlined a control input methodology that once again relied on empirical measurements from the human operator performing the task to identify a subspace that can be used to communicate the human operator's intent to the Superlimb without interfering with the human operator's primary task. Using these two methodologies, we realized a Superlimb prototype that assisted the human

operator by opening a door when the human operator's hands were full. Combined these methodologies show the potential for exploiting the human operator to allow Superlimbs to reach their potential as an assistive technology. Future work could explore additional ways to leverage the human operator or could expand upon the intermittent control structure used here, perhaps by leveraging work from the shared control community.

In Chapter 4, we offered the first exploration of the inherent haptic feedback provided by Superlimbs. As the state of the Superlimb changes, that change is often reflected as a change in force or torque about the base of the Superlimb. Thus the human operator receives haptic feedback on the state of the Superlimb. We first showed that this feedback was sufficient to allow the human operator to manually close the loop and control the force output of the Superlimb. This could be useful in applications like standing assistance for the elderly or injured as the desired amount of help in the form of supplemental force is not known and therefore should be manually controlled by the human operator. Second, we showed that the inherent haptic feedback allowed the human operator to supervise the autonomous actions of the Superlimb even when performing a separate visual task. This could be useful in many different applications as it allows the human operator to track the actions of the Superlimb and monitor for failures. This in turn helps keep the operator safe; given the detection of failure, the human operator can step in and alter the actions of or stop the Superlimb. Together, these two studies argue for the importance of inherent haptic feedback from Superlimbs. Further, they suggest future work on how to design Superlimbs that optimally communicate the inherent haptic feedback to the human operator.

This thesis represents one clear direction forward for the Superlimb community. In order to reach the potential of the Superlimbs, we must consider the contributions of the human operators.

Bibliography

[1] Elahe Abdi, Etienne Burdet, Mohamed Bouri, Sharifa Himidan, and Hannes Bleuler. In a demanding task, three-handed manipulation is preferred to two-handed manipulation. *Scientific Reports*, 6(1):21758, apr 2016.

[2] Henny Admoni and Siddhartha Srinivasa. Predicting user intent through eye gaze for shared autonomy. Technical report, 2016.

[3] Mohammed Al-Sada. Design Space of Multipurpose Daily Worn Snake-Shaped Robotic Appendages. In *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pages 455–456. IEEE, mar 2019.

[4] Rachid Alami. On human models for collaborative robots. In *2013 International Conference on Collaboration Technologies and Systems (CTS)*, pages 191–194. IEEE, may 2013.

[5] Paolo Baerlocher. *INVERSE KINEMATICS TECHNIQUES OF THE INTERACTIVE POSTURE CONTROL OF ARTICULATED FIGURES.* PhD thesis, 2001.

[6] Joseph E Barton and John D Sorkin. Design and evaluation of prosthetic shoulder controller. *Journal of rehabilitation research and development*, 51(5):711–26, 2014.

[7] Robert Bogue. Exoskeletons and robotic prosthetics: A review of recent developments. *Industrial Robot*, 36(5):421–427, aug 2009.

[8] Baldin Llorens Bonilla and H. Harry Asada. A robot on the shoulder: Coordinated human-wearable robot control using Coloured Petri Nets and Partial Least Squares predictions. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 119–125. IEEE, may 2014.

[9] Lawrence (Lawrence Zack) Bright. *Supernumerary robotic limbs for human augmentation in overhead assembly tasks.* PhD thesis, 2017.

[10] Jeremy D. Brown, Timothy S. Kunz, Duane Gardner, Mackenzie K. Shelley, Alicia J. Davis, and R. Brent Gillespie. An Empirical Evaluation of Force Feedback in Body-Powered Prostheses. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(3):215–226, mar 2017.

[11] Duncan Carter-Davies, Junshen Chen, Fei Chen, Miao Li, and Chenguang Yang. Mechatronic Design and Control of a 3D Printed Low Cost Robotic Upper Limb. In *2018 11th International Workshop on Human Friendly Robotics (HFR)*, pages 1–6. IEEE, nov 2018.

[12] Andrea S. Ciullo, Federica Felici, Manuel Giuseppe Catalano, Giorgio Grioli, Arash Ajoudani, and Antonio Bicchi. Analytical and Experimental Analysis for Position Optimization of a Grasp Assistance Supernumerary Robotic Hand. *IEEE Robotics and Automation Letters*, 3(4):4305–4312, oct 2018.

[13] Steven H Collins, M Bruce Wiggin, and Gregory S Sawicki. Reducing the energy cost of human walking using an unpowered exoskeleton. *Nature*, 522(7555):212–5, jun 2015.

[14] Francesca Cordella, Anna Lisa Ciancio, Rinaldo Sacchetti, Angelo Davalli, Andrea Giovanni Cutti, Eugenio Guglielmelli, and Loredana Zollo. Literature review on needs of upper limb prosthesis users, 2016.

[15] Frédéric Crevecoeur, Douglas P Munoz, and Stephen H Scott. Behavioral/Cognitive Dynamic Multisensory Integration: Somatosensory Speed Trumps Visual Accuracy during Feedback Control. 2016.

[16] Aaron M. Dollar and Hugh Herr. Lower extremity exoskeletons and active orthoses: Challenges and state-of-the-art. *IEEE Transactions on Robotics*, 24(1):144–158, feb 2008.

[17] Jacqueline Fagard. Bimanual Stereotypes. *Journal of Motor Behavior*, 19(3):355–366, sep 1987.

[18] Tamar Flash and Neville Hogan. The Coordination of Arm Movements: An Experimentally Confirmed Mathematical Model. *The Journal of Neuroscience*, 5(7):1688–1703, 1985.

[19] J. A. George, D. T. Kluger, T. S. Davis, S. M. Wendelken, E. V. Okorokova, Q. He, C. C. Duncan, D. T. Hutchinson, Z. C. Thumser, D. T. Beckler, P. D. Marasco, S. J. Bensmaia, and G. A. Clark. Biomimetic sensory feedback through peripheral nerve stimulation improves dexterous use of a bionic hand. *Science Robotics*, 4(32):eaax2352, jul 2019.

[20] Matthew L. Handford and Manoj Srinivasan. Robotic lower limb prosthesis design through simultaneous computer optimizations of human and prosthesis costs. *Scientific Reports*, 6(1):19983, apr 2016.

[21] Mary Hayhoe and Dana Ballard. Eye movements in natural behavior. *Trends in Cognitive Sciences*, 9(4):188–194, apr 2005.

[22] Hitoshi Honda. Eye movements and performance during bilateral tracing tasks. *Acta Psychologica*, 49(3):201–213, dec 1981.

[23] Irfan Hussain, Gionata Salvietti, Giovanni Spagnoletti, and Domenico Prattichizzo. The Soft-SixthFinger: a Wearable EMG Controlled Robotic Extra-Finger for Grasp Compensation in Chronic Stroke Patients. *IEEE Robotics and Automation Letters*, 1(2):1000–1006, jul 2016.

[24] Katherine J. Kuchenbecker, Netta Gurari, and Allison M. Okamura. Effects of Visual and Proprioceptive Motion Feedback on Human Control of Targeted Movement. In *2007 IEEE 10th International Conference on Rehabilitation Robotics*, pages 513–524. IEEE, jun 2007.

[25] Daniel A. Kurek and H. Harry Asada. The MantisBot: Design and impedance control of supernumerary robotic limbs for near-ground work. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5942–5947. IEEE, may 2017.

[26] Steven Michael LaValle. *Planning algorithms*. Cambridge University Press, 2006.

[27] Baldin Llorens - Bonilla, Federico Parietti, and H. Harry Asada. Demonstration-based control of supernumerary robotic limbs. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3936–3942. IEEE, oct 2012.

[28] Wei Ji Ma, Masud Husain, and Paul M. Bays. Changing concepts of working memory, mar 2014.

[29] Philippe Malcolm, Wim Derave, Samuel Galle, and Dirk De Clercq. A Simple Exoskeleton That Assists Plantarflexion Can Reduce the Metabolic Cost of Human Walking. *PLoS ONE*, 8(2):e56137, feb 2013.

[30] Noel Segura Meraz, Hiroshi Shikida, and Yasuhisa Hasegawa. Auricularis muscles based control interface for robotic extra thumb. In *2017 International Symposium on Micro-NanoMechatronics and Human Science (MHS)*, pages 1–3. IEEE, dec 2017.

[31] Junichi Nabeshima, MHD Yamen Saraiji, and Kouta Minamizawa. Prosthetic Tail. In *Proceedings of the 10th Augmented Human International Conference 2019 on - AH2019*, pages 1–4, New York, New York, USA, 2019. ACM Press.

[32] Koki Nakabayashi, Yukiko Iwasaki, and Hiroyasu Iwata. Development of Evaluation Indexes for Human-Centered Design of a Wearable Robot Arm. In *Proceedings of the 5th International Conference on Human Agent Interaction - HAI '17*, pages 305–310, New York, New York, USA, 2017. ACM Press.

[33] Pham Huy Nguyen, Curtis Sparks, Sai G. Nuthi, Nicholas M. Vale, and Panagiotis Polygerinos. Soft Poly-Limbs: Toward a New Paradigm of Mobile Manipulation for Daily Living Tasks. *Soft Robotics*, 6(1):38–53, feb 2019.

[34] Federico Parietti and H. Harry Asada. Supernumerary Robotic Limbs for aircraft fuselage assembly: Body stabilization and guidance by bracing. In *2014 IEEE*

International Conference on Robotics and Automation (ICRA), pages 1176–1183. IEEE, may 2014.

[35] Federico Parietti and H Harry Asada. Independent , Voluntary Control of Extra Robotic Limbs. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5954–5961, 2017.

[36] Federico Parietti and Harry Asada. Supernumerary Robotic Limbs for Human Body Support. *IEEE Transactions on Robotics*, 32(2):301–311, apr 2016.

[37] Federico Parietti and Harry H. Asada. Dynamic analysis and state estimation for wearable robotic limbs subject to human-induced disturbances. In *2013 IEEE International Conference on Robotics and Automation*, pages 3880–3887. IEEE, may 2013.

[38] Federico Parietti, Kameron C. Chan, Banks Hunter, and H. Harry Asada. Design and control of Supernumerary Robotic Limbs for balance augmentation. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5010–5017. IEEE, may 2015.

[39] H Pashler. Dual-task interference in simple tasks: data and theory. *Psychological bulletin*, 116(2):220–44, sep 1994.

[40] Nathalie Pattyn, Xavier Neyt, David Henderickx, and Eric Soetens. Psychophysiological investigation of vigilance decrement: Boredom or cognitive fatigue? *Physiology & Behavior*, 93(1-2):369–378, jan 2008.

[41] Jeff Pelz, Mary Hayhoe, and Russ Loeber. The coordination of eye, head, and hand movements in a natural task. *Experimental Brain Research*, 139(3):266–277, aug 2001.

[42] Christian Penaloza, David Hernandez-Carmona, and Shuichi Nishio. Towards Intelligent Brain-Controlled Body Augmentation Robotic Limbs. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1011–1015. IEEE, oct 2018.

[43] Robert W. Proctor and Trisha. Van Zandt. *Human factors in simple and complex systems*. CRC Press, 2008.

[44] Gionata Salvietti, Irfan Hussain, David Cioncoloni, Sabrina Taddei, Simone Rossi, and Domenico Prattichizzo. Compensating Hand Function in Chronic Stroke Patients Through the Robotic Sixth Finger. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(2):142–150, feb 2017.

[45] MHD Yamen Saraiji, Tomoya Sasaki, Kai Kunze, Kouta Minamizawa, and Masahiko Inami. MetaArmsS: Body Remapping Using Feet-Controlled Artificial Arms. In *The 31st Annual ACM Symposium on User Interface Software and Technology - UIST '18*, pages 65–74, New York, New York, USA, 2018. ACM Press.

[46] Jean Scholtz. Theory and Evaluation of Human Robot Interactions. *36th Annual Hawaii International Conference on System Sciences*, page 10 pp., 2002.

[47] Thomas .B. Sheridan. Function allocation: Algorithm, alchemy or apostasy? *International Journal of Human Computer Studies*, 52(2):203–216, feb 2000.

[48] Thomas B Sheridan. Human âĂŞ Robot Interaction : Status and Challenges. *Human Factors and Ergonomics Society*, 58(4):525–532, 2016.

[49] Thomas B. Sheridan and William R. Ferrell. Human control of remote computer-manipulators. In *Proceedings of the 1st International Joint Conference on Artificial Intelligence*, IJCAI'69, pages 483–494, San Francisco, CA, USA, 1969. Morgan Kaufmann Publishers Inc.

[50] Chang-Yeob Shin, Jangho Bae, and Daehie Hong. Ceiling work scenario based hardware design and control algorithm of supernumerary robotic limbs. In *2015 15th International Conference on Control, Automation and Systems (ICCAS)*, pages 1228–1230. IEEE, oct 2015.

[51] Samvrit Srinivas, Gurvinder Singh Virk, and Usman Haider. Multipurpose supernumerary robotic limbs for industrial and domestic applications. In *2015 20th International Conference on Methods and Models in Automation and Robotics (MMAR)*, pages 289–293. IEEE, aug 2015.

[52] Hong-Rui Su and Kuo-Yi Chen. Design and implementation of a mobile robot with autonomous door opening ability. In *2017 International Conference on Fuzzy Theory and Its Applications (iFUZZY)*, pages 1–6. IEEE, nov 2017.

[53] Hong Z Tan, Brian Eberman, Mandayam A Srinivasan, and Belinda Cheng. HUMAN FACTORS FOR THE DESIGN OF FORCE-REFLECTING HAPTIC INTERFACES. *Dynamic Systems and Control*, 55(1), 1994.

[54] Laura Treers, Roger Lo, Michael Cheung, Aymeric Guy, Jacob Guggenheim, Federico Parietti, and Harry Asada. Design and Control of Lightweight Supernumerary Robotic Limbs for Sitting/Standing Assistance. pages 299–308. Springer, Cham, 2017.

[55] F. Y. Wu and H. H. Asada. âĂIJhold-and-manipulateâĂİ with a single hand being assisted by wearable extra fingers. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6205–6212, May 2015.

[56] Faye Y. Wu and H. Harry Asada. Implicit and Intuitive Grasp Posture Control for Wearable Robotic Fingers: A Data-Driven Method Using Partial Least Squares. *IEEE Transactions on Robotics*, 32(1):176–186, feb 2016.