Understanding Interaction: Unraveling the Mysteries of the Mind Using Virtual Reality

David E. Caballero

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2020

Reading Committee:

Eric Rombokas, Chair

Blake Hannaford, Chair

Howard J. Chizeck

Program Authorized to Offer Degree:
Electrical & Computer Engineering

University of Washington

Abstract

Understanding Interaction: Unraveling the Mysteries
of the Mind Using Virtual Reality

David E. Caballero

Co-Chairs of the Supervisory Committee:
Dr. Eric Rombokas
Mechanical Engineering

Dr. Blake Hannaford
Electrical & Computer Engineering

Human robot interaction (HRI) is an important problem. As engineers we can design robots to match the requirements and desired capabilities as close as possible. However, the design targets are still unknown until we properly grasp how the human works. Specifically, human sensation is still poorly understood in the context of HRI. What I have contributed here are several firsts in our understanding of human perception by introducing experiments that study 1) how humans perceive sensory information from different sensory modalities, 2) how sensory information impact the way we perform tasks and 3) how we use such information to update the models of our own bodies.

By looking at this problem through the lens of *sensory conflict* and using recent advances and increase in accessibility of Virtual Reality, I investigate the effect of different sensory modalities in isolation. I pioneer a new method, the *Visuotactile Two-Point Discrimination Test*, that focuses on spatial conflict between visual touch and tactile touch to explore the connection between vision and taction. Additionally, by investigating the effect sensory information has on the expression of so called "dynamic primitives", I analyze for the very first time the effect of sensation on the switch between the primitives proposed in this theory. Finally, by introducing C/D-ratio manipulations and changing the relationship between the display (visual) and control signal (proprioception), I explore the role of sensation in updating our internal models.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I would like to express my most sincere appreciation to both of my advisors who guided me throughout my doctoral journey. Professor Blake Hannaford provided me with tremendous help and advice that proved priceless during my first years in the program. By allowing me to flex my muscles as an engineer, I learned to appreciate the complexities of the human-robot system. With Dr. Eric Rombokas I gained confidence to question the status-quo as a scientist. His incredible energy and insatiable passion for science will forever be instilled in my career as a researcher. Your guidance was instrumental for me reaching the finish line and for that I will always be thankful.

I wish to also thank the members of my Ph.D. committee, Professors Howard Chizeck and Valerie Kelly. Professor Chizeck was among the first members of the UW community who made me feel welcomed and helped me see through the fog of my first year as a Ph.D. student - thank you for the unwavering support when I needed it. Professor Kelly, thank you for always been available to provide me with feedback, suggestions or even just to chat. I will cherish your willingness to always be available and engage in discussions, even unplanned ones at the AMP lab.

I also want to thank all my collaborators in both of the incredible labs I had the honor to be a part of. Thank you to all the members of the UW BioRobotics Lab including Nava Aghdasi, Muneaki Miyasaka, Mohammad Haghighipanah, Kevin Huang, Andrew Lewis, Yangming Li, Jeffrey Herron, Niveditha Kalavakonda, Kyle Lindgren and Andrew Pace for creating a joyful environment. In RomboLabs I would like to thank all my fellow rombogoons including David Boe, Gaurav Mukherjee, Pornthep Preechayasomboon, Vijeth Rai, Astrini Sie, Nataliya Rokhmanova and, of course, Vlad for providing a creative environment to thrive in.

Finally, I would like to express my deepest and most sincere appreciation to my family. First, to my parents for teaching me the importance of working hard and continue pushing forward for the next generation. Thank you mom for your sacrifice during my first years of life to ensure we had a roof over our heads. Thank you dad for being the first and best teacher I ever had, instilling in me the passion for science and engineering. Thank you manita, my sister, for visiting me early on and bringing a sense of home to me. To my late grandpa, mi abuelo César, who I hope continues being proud of me from above. I am insanely grateful and forever indebted to my wonderful wife, biggest cheerleader and partner in life, Dr. Bożena

Caballero. I can say without a doubt that I could not have chosen a better person to help me weather the crazy and unpredictable game of life - without your unwavering support I truly would not be where I am today. To Oliver and Maya, the most amazing and incredible kids a father could have ever asked for. Thank you for helping bring me back to reality with your countless hugs, kisses, shakes and high-fives; being your "Duda" is the best thing that have ever happened to me.

DEDICATION

to my amazing family, my team in life

Chapter 1

INTRODUCTION

Human-Robot Interaction (HRI) is becoming ubiquitous and its potential to enhance human capabilities has never been greater. This is particularly true in the medical field. Doctors are increasingly relying on the use of surgical robots to perform complex procedures [1]. Systems such as the daVinci Xi provide surgeons with increasing levels of dexterity while minimizing the area they operate in. Amputees are slowly able to regain their mobility (even their lost sensation!) thanks to robotic prosthesis. By chronically implanting electrode arrays and mapping readings from sensors placed at the robotic prosthesis to appropriate stimulations at afferent sensory nerve endings, sensation can be restored [2]. Other procedures such as targeted sensory reinnervation (TSR) [3] and regenerative peripheral nerve interfaces (RPNIs) [4] are less invasive alternatives currently under investigation for sensory restoration in robotic prosthesis [5]. The elderly and people with chronic conditions are able to retain their independence longer due to exoskeletons.

When designing HRI systems, we should have a clear understanding on both the human and robot components and as engineers it is tempting to place most of our attention on the latter. As a matter of fact, much of my initial work focused on facilitating these interactions from the robotic side of the equation. By introducing enhancements to the RAVEN I [6] and expanding the compatibility of tools for both RAVEN I and II [7], the RAVEN platform now more closely matches the expectations of the human operating it. Additionally, through building a better understanding of the cables driving the RAVEN [8] and proposing a new control architecture [9], improvements were introduced in the precision of this platform. I also participated in the development of the kinematic model and initial controller development for the Roboscope [10], a flexible robotic manipulator that allows access to difficult to reach areas in the human anatomy. Through my contributions to the development of the Surgical Cockpit, a novel haptic and teleoperation system, I enabled more natural teleoperation of robotic systems.

Understanding how humans operate is key in properly designing human-robot systems. In many of the robotic platforms described above, the human-in-the-loop overcomes limitations on the robotic side of the interaction. This can clearly be seen in the case of the RAVEN. Figure 1.1 illustrates the positional error for one of the RAVEN II's end effectors. During teleoperation, the human compensates for these errors without effort. We observed a similar phenomenon during a rat dissection task under different time delays. Users were able to

compensate for a wide range of temporal delays.

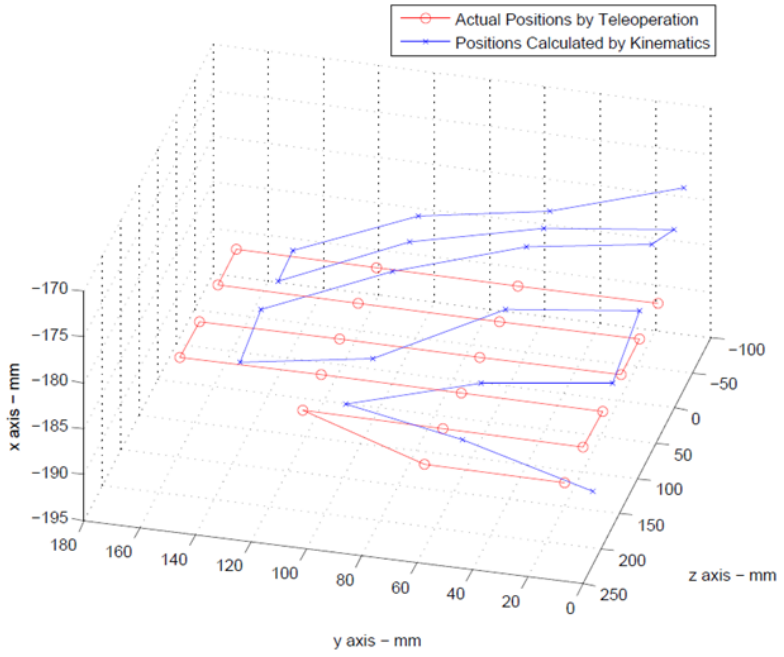

Figure 1.1: RAVEN II End Effector Positional Error

There are many open questions in how humans perceive information and generate behavior. Thus, the following chapters present my work towards applying a similar systems engineering approach as in my robotics work but to the human component instead. Humans are not simple input-output agents. We are equipped with a suite of extremely complex sensors, each providing sensory readings at varying intervals and with different delays. On top of this, our nervous system, due to its biological nature, introduces noise and uncertainty to these readings. To undertand how we can deal with all these issues and still generate a single sensory percept, a rich amount of literature exist looking at this problem from a sensor fusion perspective. A much less investigated paradigm, however, is to look at the problem when sensory readings are in complete conflict with one another. By looking at the problem from the lens of *sensory conflict*, I studied the contribution of each sensory modality in isolation. We proposed a new test we called *The Visuotactile Two-Point Discrimination Test* which explores the connection between vision and taction through spatial conflict between these two modalities [11]. Performing this test at areas of the body with wildly different tactile resolutions (thigh vs. index finger) provided us with interesting insights into the potential mechanisms at play.

We also need to understand how human movement is generated and how sensation plays

a role in how we decide to act on the world around us. To tackle this question, we look at Dynamic Primitive Theory [12], [13], an influential idea in neuroscience. It proposes that the brain uses distinct subsystems for different types of movements in order to mitigate noise, delay and uncertainty. In this literature, however, sensation is not closely taken into consideration. Thus, we introduce a sensory component to prior experiments that identify the point of transition between two of these primitives: discrete and rhythmic [14]. Understanding how primitive expression depends on sensory information will help us better understand our nervous system's ability to generate complex actions.

Humans are very active beings; we constantly sense and act on the world around us. We are equipped with a very rich and complex body at our disposal to do so. However, learning how to properly make use of such a complex system takes time and it has been shown that we create and adapt internal models of our body to help achieve such a difficult task. Thus, I expand on the work done in this field by investigating the role of sensation in updating these internal models.

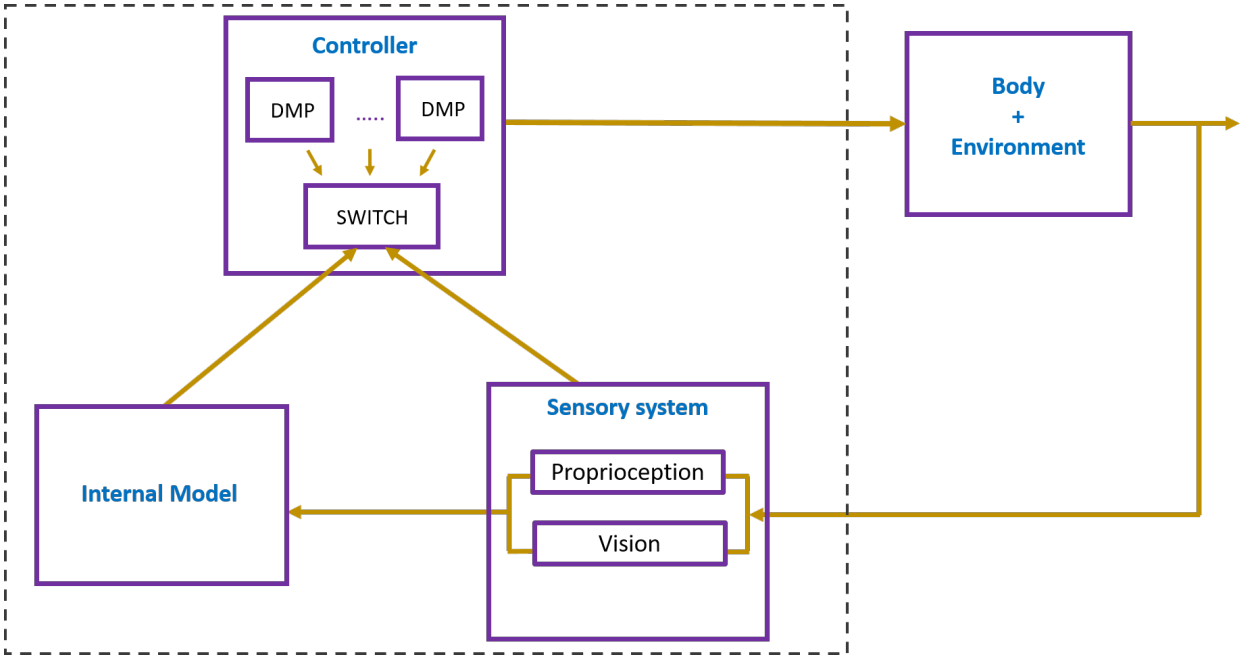

Figure 1.2: Proposed block diagram breaking down the human mind into subcomponents. Each block and their interconnections are explored in this thesis.

In the following chapters I will describe my contributions towards understanding the human mind using Virtual Reality. Specifically, I break down the problem into subcomponents by finding a balance between complexity and simplicity with an overview of our model presented in Figure 1.2. Chapter 2, focuses on the sensory system and studies how our brain

makes sense of conflicting visual and tactile information. By exploring different areas of the body, we observe very different results with more tactilely parts showing, unexpectedly, more reliance on vision. How this sensory information dictates the way our brain generates motor behavior is then looked at in Chapter 3 by introducing, for the very first time, sensory paradigms to the study of *Dynamic Primitives (DPs)*. Conclusions on the apparent insensitivity of DPs to sensation and potential explanations regarding the effect of *Internal Models* are then described. The last chapter puts *Internal Models* under the microscope and examines the effect of manipulating the ratio between tracked movements of the body and their rendered display (the control-to-display ratio). We demonstrate that our brain incorporates such a multisensory gain (which we refer to as *"proprioceptive gain"*) in the update of its internal model and observe interesting characteristics with respect to the directionality of the task and the time course of the aftereffects of such a recalibration.

Chapter 2

SENSORY CONFLICT

There are an increasing number of lower limb amputees and prosthesis designed for such a population usually fall short of augmenting the wearers' capabilities. If we could restore sensory feedback for amputees, we could inch closer to restoring their lost mobility. To this end, I have been working with the Center for Limb Loss and Mobility (CLiMB) of the VA Hospital in a project aiming to have a sensorized robotic limb that will relay back sensory information in a non-invasive way by cleverly stimulating the area where users had Targeted Reinnervation (TR) performed. However, in order to achieve this, we have to better understand how sensory integration work.

2.1 Introduction

Perception is created through the rich interplay of multiple sensory modalities. It has historically been difficult to quantify the number of modalities. For example, *touch* is a term that includes inputs from five types of cutaneous mechanoreceptors, sensors at the base of hairs, and sensors sensitive to the changes in pressure that result from deformation of soft tissues. Perception is also intertwined with internal models of the body and expectations of its movement and the outside world (e.g. [15]). To make sense of sensation, it is tempting to isolate modalities and study them singly, but this is made difficult by the unexpected ways that the neural and mechanical systems interact. There has been a heroic effort to document and explain how multiple sensory modalities work together, with sometimes contradictory conclusions. One of the most consistent observations produced by this research is that visual sensation appears to take perceptual preference over other senses in many contexts. Vision has been shown to be dominant when both stimuli have the same perceived intensity [16], but also even when the conflicting stimulus is presented with twice the intensity of the visual [17].

Instead of observing vision and other modalities under normal circumstances, we intentionally induce conflict between two senses to isolate the contributions of each. We hypothesize that in visuotactile perception, if vision is indeed dominant, the smallest noticeable intrapersonal spatial conflict between vision and taction will be different from a purely tactile one. There have been fewer studies of conflict than general multisensory perception

because creating precise errors between senses often requires elaborate experimental apparatus. The recent advances in "virtual reality" (VR) provide an opportunity to isolate vision from other senses more easily and accurately than ever before

2.1.1 Types of Sensory Conflict

The apparent dominance of vision decreases as temporal delays are introduced between vision and other sensory modalities. Complex temporal patterns, instead of simple delays, can also change the relative contribution of each modality [18], [19].

Spatial conflict may also be induced between senses. We make a distinction between spatial conflict in extrapersonal and intrapersonal space. Extrapersonal stimuli have apparent origin outside of the body, such as flashing shapes or sounds [20]. Stimuli in the intrapersonal space explicitly pertain to the body [20], [21].

Conflicting stimuli can be presented even within the same sensing modality, as in *binocular rivalry*, when different images are displayed to each eye [22], [23], or binaural rivalry, when disparate sounds are played to each ear [24], [25].

Figure 2.1 summarizes previous studies of conflict between vision and proprioception, audition, and taction, and categorizes that conflict into time, intrapersonal spatial, or extrapersonal spatial. The study we present in this manuscript is a study of conflict between vision and taction, performed in the intrapersonal space.

		Sensory Conflict Type		
		Time	**Intrapersonal Space**	**Extrapersonal Space**
Conflict Modalities	**Vision vs. Propioception**	Pausch et al. 1992 [44]	Hay et al. 1965 [7] Pick et al. 1969 [6] Brewer et al. 2005 [14] Brewer et al. 2008 [15] Petroni et al. 2015 [13]	Kim et al. 2012 [17] Kim et al. 2014 [18]
	Vision vs Audition	Thomas et al. 1941 [4] Radeau et al. 1987 [5]	Koppen et al. 2007 [22]	Pick et al. 1969 [6] Bertelson et al. 1981 [19] Mateef et al. 1985 [20]
	Vision vs. Taction	Botvinick et al. 1998 [29]	Hartcher-O'Brien et al. 2008 [2]	Rock et al. 1964 [23] Heller et al. 1982 [24] Ernst et al. 2002 [25] Riley et al. 1997 [27]

Figure 2.1: Summary of representative work of conflict between vision and other senses. The experiment described in this paper (VT Test) falls within the highlighted cell.

2.1.2 Vision and Proprioception

Early work showed that when the location of the hand was displaced by means of a prism, its felt location was displaced towards the location of the visualized hand [21], [20]. More recently, a proprioceptive illusion of finger elongation was shown in [26], and this illusion can be used to modify the estimation of reaching distance [27]. Visual feedback distortion, in which a visual representation of position or force is modified [28], [29], can influence reaching movements. Isolating vision and proprioception during reaching suggests that the two modalities are combined in fundamentally different ways during different phases of movement [30].

In [31], visual distortion of gait step lengths caused subjects to perceive their step lengths as being asymmetric during treadmill walking. The subjects were unaware of the visual-proprioceptive mismatch and modulated their step lengths away from symmetry as a response. Even when subjects were explicitly informed of the mechanisms in which the visual feedback distortion was being applied, subjects were unable to maintain a natural symmetric step pattern and had their gait step length systematically modulated away from symmetry [32].

2.1.3 Vision and Audition

By artificially generating discrepancies between the spatial location of stimuli provided visually and through audition, Pick et al. showed that vision biases audition [20]. This phenomenon has been coined the *ventriloquism effect* since it has the perceived effect of making the origin of the sound to be coming from a different location than its source [33], [19]. It has even been shown that motion of the source of such auditory stimulus can be induced by the movement of a visual target [34]. The quality of the visual target also plays an important role in this phenomenon. [35].

In [17], subjects were presented with visual-only, audition-only and audiovisual stimulus and were then asked to identify which stimuli (or stimulus) they perceived. Results showed a clear dominance of vision over audition, with subjects significantly providing vision-only responses in the multimodal case. This effect, later coined the *Colavita visual dominance effect*, was showed to be robust to the spatial directionality of the stimulus [36].

2.1.4 Vision and Touch

Early explorations suggested that vision was dominant over touch [37]. There is also evidence that the modalities are combined. For example, the perception of the texture of a touched surface is affected by both vision and touch [38]. More recently, such a combination of

visual and haptic information was found to follow the behaviour of a maximum-likelihood integrator [39], [40]. A study of gently touching a surface while standing showed that the haptic contribution to postural sway was similar to that of vision, but had little effect when both were available [41]. Additionally, the *Colavita effect*, shown to be present between vision and audition, was also proved to exist between vision and touch [16], with results following a similar trend as in the audiovisual studies. In [42], it was shown that the visual depiction of touch can induce the illusory tactile sensation of touch originating in that location. in the *Right Hand Illusion*, the dominance of vision over touch depends on their synchronization [43]. Visual stimuli have also been found to modulate the *cutaneous rabbit effect* [44]. These and myriad other experiments show that the fusion of touch and vision is complex and situational.

2.2 The Visuotactile Consistency Test

In order to better inform the design of haptic interfaces and understand the connection between vision and taction, we proposed the Visual-Tactile Consistency Test. Similar to the traditional Two-Point Discrimination Test, subjects were asked to distinguish between a one-point or two-point stimulation. The difference, however, was that now only one of the stimulation points was tactile while the other one was visual. Virtual Reality was used as the main tool to enable this test. This study concluded in a journal publication in Transactions on Haptics [16] and we have also been invited to present the results at the 2019 World Haptics Conference in July.

2.2.1 Visuotactile Two-Point Discrimination Test (VT Test)

In this paper we focus on spatial conflict between visual touch and tactile touch. We describe a protocol that is analogous to the tactile two-point discrimination test.

In the original test, the subject is asked to discern whether one or two stimuli are present while the experimenter touches the subject with either one or two stimulation instruments. Throughout the experiment, the subject's vision is blocked. An early comprehensive study [45] established the two-point discrimination resolution for different parts of the lower limb in a college age population of males and females with results ranging from 43.6 mm to 6.6 mm depending on the skin region. Simultaneous touches with two instruments closer together than the measured resolution were felt as single touches by the subjects. For most skin regions, no significant difference was shown between men and women. Since then, the test has been used in the clinic to assess tactile acuity for applications including diagnosis of diabetic neuropathy [46], [47] and peripheral nerve injury recovery [48]. However, it has faced criticism due to the report of contradictory results throughout the years, variation

across subjects and studies that remain unexplained, and the lack of a standardized testing procedure [49], [50].

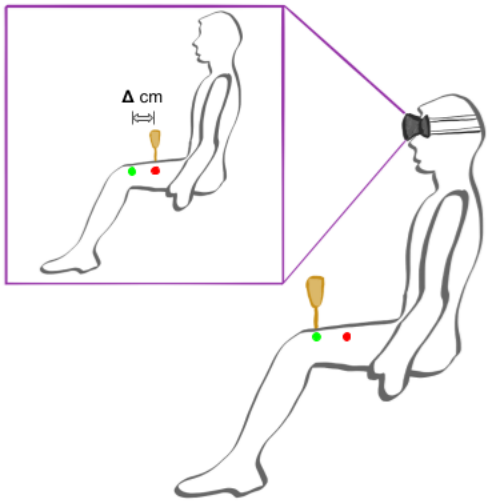

Figure 2.2: Schematic of VT test. Tactile touch presented to the subject at either one of the two predetermined points. A virtual offset of Δ cm is applied at random to the visual touch in the virtual environment.

In our experiment, however, only one stimulation point is tactile while the other is visual as depicted in Figure 2.2. We achieve this by making use of Virtual Reality to render the visual touch to the subject. The recent breakthroughs in VR and its increasing accessibility makes these types of studies possible without the need of complex machinery and experimental setups [51].

We seek to use this information to restore sensory feedback in lower limb amputees who received the Targeted Muscle Reinnervation (TMR) procedure on their thigh region. Sensory feedback will be delivered to the user of a sensorized prosthetic limb, providing a "pass-through" for touch. To determine the required density of the tactor array, we must understand the spatial resolution of the skin. As described above, this resolution appears to be sensitive to vision, which in this application will appear to the user of the sensory pass-through prosthesis as touches on the prosthetic limb. Understanding visuotactile conflict is also important for the design of haptic hardware for other applications such as Virtual Reality haptic interfaces.

In Section 2.2.2 we describe the experiments and apparatus, in Section 2.2.3 the statistical methods used are outlined, in Section 2.2.4 we present the results, and in Section 2.2.5 we discuss the implications of the results, limitations of the study, and next steps.

2.2.2 Experimental Methods

The study was divided into two major experiments. The first was the "traditional" tactile two-point discrimination test, and the second was the visuotactile two-point discrimination test.

We focused on exploring the sensitivity to conflict between the visual and tactile channels by immersing the subjects in a virtual environment. This was done in order to attempt to shed some light on the role of vision in generating the overall picture of touch by comparing a condition in which only tactile information is present (two-point discrimination) to that of when the information provided through the visual channel is inconsistent to that of the sense of touch (VT test).

There are several experimental designs and procedures shared by both experiments. First, two locations, 12 centimeters apart, were marked on the subjects' left anterior upper leg using washable markers. The distal location, referred to as site 1, was 10 centimeters away from the knee, and the proximal location, site 2, was 22 cm away from the knee.

There were a total of twelve (seven men and five women), non-smoking, subjects for this study without any health issues that could affect sensation in their thigh. They were all college students with ages ranging between 20 to 30 years old and gave informed written consent before participating. The experimental protocols were approved by the University of Washington Institutional Review Board.

Tactile Two-Point Discrimination Experiment (TT Test)

Subjects were placed in a seated position with their view obstructed by means of a (turned off) VR headset throughout the experiment so that the elicited stimulation response was solely based on the sense of touch. We used a 2-point discriminator with adjustable points (aesthesiometer), which provides precise measurable distances, to accurately stimulate the targeted stimulation locations.

On each trial, one of the two stimulation points (site 1 or site 2) was randomly selected, followed by one of five possible offset conditions in either the proximal or distal direction to that site, or by the original point itself (zero offset condition). Offsets ranged from 2 to 6 cm (in 1 cm increments). This results in 11 possible stimulus offsets (including an offset of 0 cm) involving each stimulation point, with each offset repeated 5 times, resulting in a total of 110 stimulation instances (2 sites x 11 stimulus offsets x 5 repetitions). The choice of total repetitions per point to be explored was set to have the total experiment time below or around 25 minutes. This was done to keep the subject's fatigue level low and attention level high since his/her focus was essential for attaining reliable results.

During the experiment, the subject was stimulated for approximately 0.4 s and asked to report verbally how many points s/he perceived.

Visuotactile Two-Point Experiment (VT Test)

For this experiment, the two points drawn on the subjects' left thigh represented the locations where they were stimulated by the experimenter using a 300 g medical grade monofilament. By using this monofilament, we were able to guarantee a constant force was applied throughout the experiment[52] and ensure detectability of the tactile stimulus. This in turn enabled subjects to focus on spatial conflict between vision and taction.

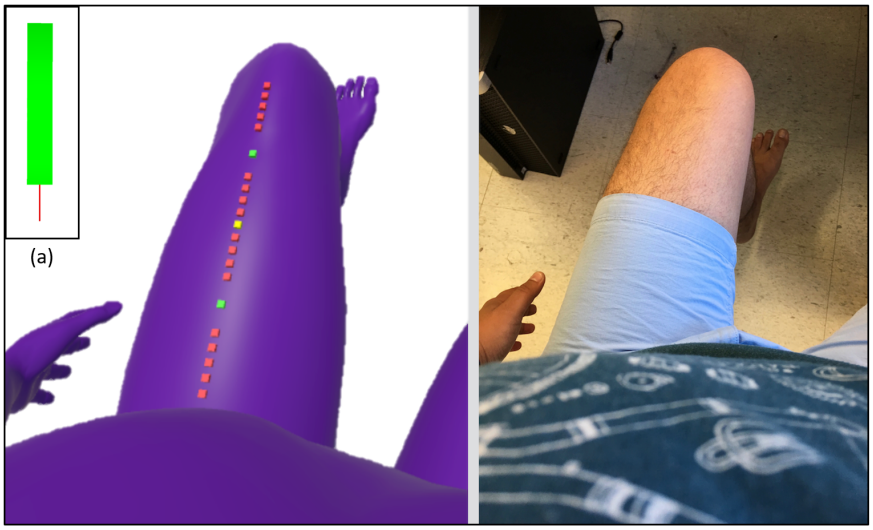

Figure 2.3: Virtual monofilament representation (a) and subject's point of view in virtual reality (b) and the real world (c). Cubes indicate the different visual "stimulation" points with green being those that overlap the real locations of the tactile stimulation points, red being the visual offset locations, and yellow the common visual point for both real tactile stimulation points. These cubes are not shown to the subject during the experiment.

Additionally, a virtual environment with an avatar, whose dimensions match the subject's anatomy as closely as possible, was built with the Unity game engine [53]. The "uncanny" valley, a phenomenon in which subjects can feel uncomfortable when presented with human-looking avatars or limbs, was considered during the design of the experiment. Due to the complexity of designing a personalized avatar for each subject which could avoid the "uncanny" valley, as well as the many open questions related to this phenomenon [54], the decision was made to make a single avatar with a realistic human shape but with a unicol-

ored texture (light purple). The avatar was in a sitting position to resemble the subject's pose during the experiment. For the present experiment, the HTC Vive [55], a virtual reality headset, was used to present a visual fake limb. The subject, upon wearing the device, was able to look from the avatar's perspective by moving and/or translating his/her head in real-time. The monofilament was attached to the bottom of one of the controllers by means of a custom-designed 3D printed adapter. A virtual rendering of this controller (with an attached monofilament) was also present in the virtual environment, providing a visual cue of touches (Fig. 2.3(a)).

Visual-tactile stimulation was presented in a random order with random offsets for both tactile stimulation points. Figure 2.3 shows the perspective as seen from the subject's point of view both in virtual reality (Fig. 2.3(b)) and in the real world (Fig. 2.3(c)). Please note that, for the VT experiment, only the virtual avatar is displayed to the subject. This figure also shows the different virtual stimulation points as colored cubes overlayed atop the avatar's left thigh. Green cubes depict the stimulation points in real-life on the subject (0 cm offset) and red cubes depict the visual location of touches occurring in virtual reality.

There was one virtual point common to both site 1 (representing a 6 cm offset in the proximal direction) and site 2 (representing a 6 cm offset in the distal direction) which is depicted as a yellow cube in Figure 2.3. Note that the cubes were not seen by the subject during the experiment but are shown in this figure for clarification purposes only.

The same procedure as in Section 2.2.2 was followed during this experiment, with the subject being presented a randomly generated offset for approximately 0.4 s and asked to report verbally if they perceived one or two points. The difference, however, is that now one of the touches is a visual one as seen in Figure 2.2. The visual location of touch was offset, while the tactile touch was always located in one of the two stimulation sites (green cubes in Figure 2.3). Thus, subjects were instructed to treat the visual and tactile stimuli as two points if they experienced them as coming from two separate locations.

In total, touches at each virtual point were delivered five times throughout the experiment for a total of 110 stimulation instances per subject.

2.2.3 Statistical Methods

Generalized linear mixed effects regression with binomial errors and logit link was used to estimate the proportion of responses that detected 1 touch (vs. 2) by site or experiment. In all models, response (1 or 2 touches) was the dependent variable and study participant was modeled as a random effect to account for each study participant having a different predisposition to detect stimuli. In order to look for possible differences in responses by site, site was treated as the independent fixed effect. A random component was also modeled for site, (site by study participant interaction) to account for variability in differences in

Table 2.1: Summary of statistical results (probability of one touch averaged across all offsets and odds ratio with 95% CI) for hypothesis concerning differences in response by site.

	Site 1	Site 2	OR for Site 1/Site 2	p
TT	0.51 (0.33, 0.68)	0.88 (0.78, 0.94)	0.14 (0.07, 0.28)	<0.0001
VT	0.64 (0.51, 0.75)	0.80 (0.68, 0.88)	0.44 (0.25, 0.80)	0.0049

responses by site across study participant. Separate regressions were carried out for the two-point discrimination and for the Visuotactile experiments (Table 2.1). Similarly, differences in responses by experiment were also explored (Table 2.2). To do so, experiment was treated as an independent fixed effect, and experiment by participant interaction was modeled as random, with separate regressions carried out for site 1 and for site 2. All of these models had offset as a fixed effect covariate (modeled as a 10 category dummy variable) to adjust for differences in responses by offset.

Another item of interest was to look at changes in response across offset in the proximal direction and how these differed from those in the distal direction. To do so, a different model was run both including and excluding the 0 offset, in case it overly influenced the estimation of the slopes.

Additionally, we were interested in possible differences in responses in the point common to both Site 1 (located at - 6 cm offset) and Site 2 (located at 6 cm offset) which is represented as a yellow cube in Figure 2.3. To do so, a generalized linear mixed effects regression model was carried out with response as the dependent variable, site as the independent fixed effect of interest and experiment type as an independent fixed effect covariate. A site by experiment fixed effect interaction term was added to the model in order to estimate separate differences by experiment and test if the differences in response by site differed by experiment.

Type 1 error was set at $p = 0.05$. Mean responses are presented as the probabilities of detecting 1 touch. Differences in responses by site or experiment are presented as odds ratio (ORs) for the odds of detecting 1 vs. 2 touches for site 1 or two-point experiment over the odds for detecting 1 vs. 2 touches for site 2 or Visuotactile experiment.

2.2.4 Results

Proportion of Perceived One-Touch

The key measure for the two experiments is the *proportion of perceived one-touch*, which is the proportion of the number of times a subject reported feeling one touch to the total number of trials presented. In the tactile two-point discrimination experiment, a one touch event

Table 2.2: Summary of statistical results (probability of one touch averaged across all offsets and odds ratio with 95% CI) for hypothesis concerning differences in response by experiment.

	TT	VT	OR for TT/VT	p
Site 1	0.50 (0.36, 0.65)	0.66 (0.52, 0.78)	0.52 (0.26, 1.06)	0.072
Site 2	0.82 (0.72, 0.90)	0.81 (0.70, 0.89)	1.06 (0.58, 1.96)	0.85

Table 2.3: Probability of one touch response (with 95% CI) at Site 1 with -6 cm offset and Site 2 with 6 cm offset for both experiments

	Site 1	Site 2	OR for Site 1/Site 2	p
TT	0.19 (0.06, 0.44)	0.43 (0.29, 0.58)	0.30 (0.08, 1.20)	0.092
VT	0.38 (0.16, 0.67)	0.70 (0.54, 0.82)	0.26 (0.07, 1.01)	0.052
Combined	0.27 (0.11, 0.54)	0.57 (0.44, 0.68)	0.28 (0.08, 0.96)	0.043

occurred when subjects were either given a one-point stimulation or when the offset between stimulation points in the two-point condition was not discernible. For the Visuotactile test, one touch happened when the tactile and visual touch were collocated (emulating a one-point stimulation) or when the offset between the tactile stimulation and its visual representation was indiscernible.

Thus, for small offsets, we expect subjects to report one touch more frequently since the two tactile stimuli (in the Tactile Two-Point Experiment) or the visual and tactile stimuli (in the Visuotactile test) are well aligned. When the offsets are large, we expect the *proportion of perceived one-touch* to be low, as subjects recognize the spatial disparity between the two tactile stimuli (in the Tactile Two-Point Experiment) or between vision and touch (in the Visuotactile test).

Tactile Two-Point Discrimination Experiment

Figure 2.4 presents the *proportion of perceived one-touch* across all subjects, arranged so that the most distal data point (Site 1 with offset of + 6 cm) is in the rightmost portion of the plot while the most proximal data point (Site 2 with offset of - 6 cm) is on the leftmost portion of the plot.

The biggest proportion of reported one touches occurred during the no offset condition. This is as expected since it is the only condition in which a one touch stimulation is presented to the subjects. As can be seen in Table 2.1, the probability of the subjects perceiving the stimulation as one touch was significantly higher for Site 2 (0.88) than for Site 1 (0.51), with

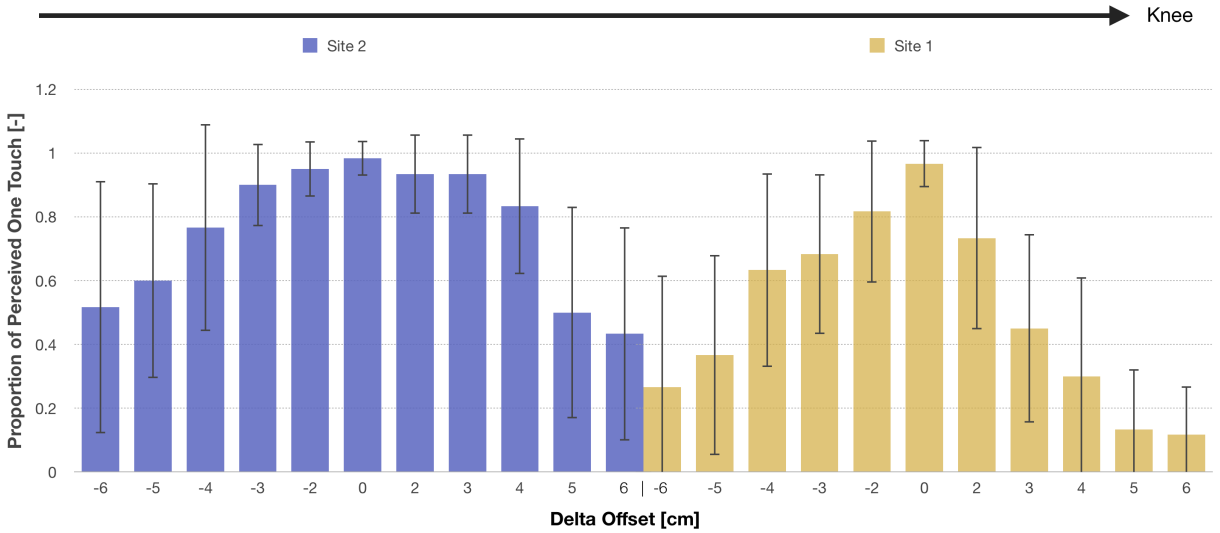

Figure 2.4: Tactile Two-Point Experiment: Proportion of perceived one-touch with their respective standard deviation for all offsets explored with Site 1 on the right and Site 2 to the left averaged across all subjects.

odds ratio of detecting one touch in Site 1 vs. Site 2 of 0.15 and p-value smaller than 0.0001.

Additionally, we found that the linear slope of change in responses as the offset was increased distally was slightly steeper than proximally for Site 1, but not by a significant margin ($p = 0.067$). For site 2 the difference was weaker. This result holds even when excluding potential biases introduced by answers at the 0 cm (single touch) condition.

For the two-point touch condition in which both sites share the same point (- 6 cm for Site 1 and + 6 cm for Site 2, shown as a yellow cube in Figure 2.3), the probability of perceived one touch was higher for Site 2 (0.43) than for Site 1 (0.19), with odds ratio of 0.3 and $p = 0.092$.

Visuotactile Two-Point Discrimination Experiment

The average *proportion of perceived one-touch* across all subjects is presented in Figure 2.5, with the rightmost point (representing an offset of + 6 cm in Site 1) being the most distal while the leftmost data point (offset of - 6 cm in Site 2) reflecting the results for the most proximal point.

Similar to the Tactile Two-Point Discrimination portion of the experiment, the highest proportion of perceived single touches for both stimulation sites occurred when no offset was

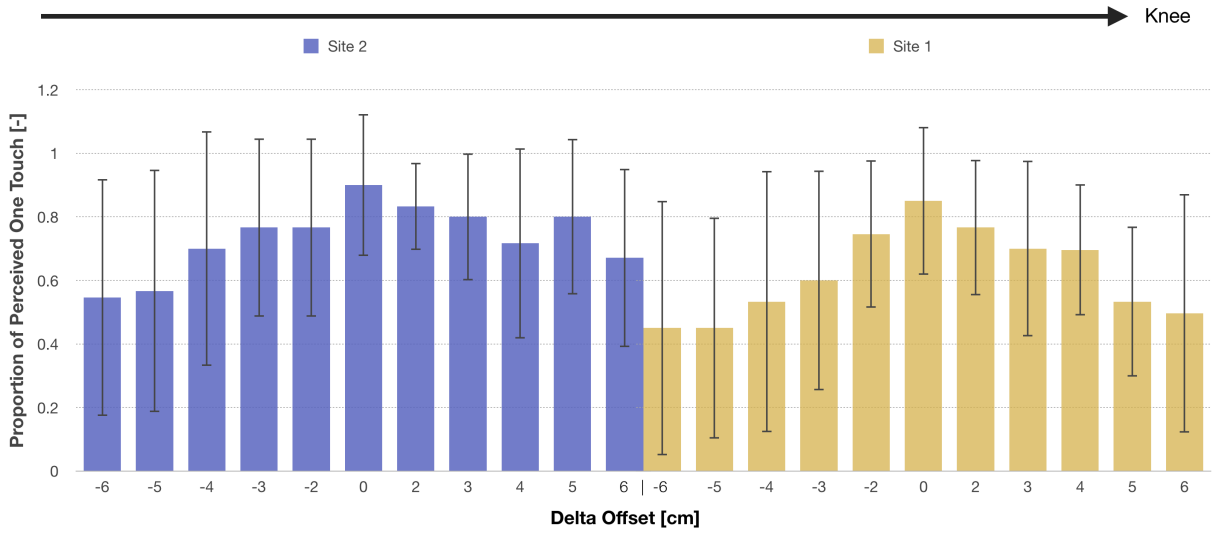

Figure 2.5: Visuotactile Experiment: Proportion of perceived one-touch with their respective standard deviation for all offsets explored with Site 1 on the right and Site 2 to the left averaged across all subjects.

present. This was expected as it is the only condition in which both the visual and the tactile touches are collocated.

From Table 2.1, we can also see that the probability of the subjects perceiving one touch was significantly higher for Site 2 (0.80) than for Site 1 (0.64), with odds ratio of detecting one touch in Site 1 vs. Site 2 of 0.44 and p-value of 0.0049. There was no significant difference in the linear slope of change in responses as the offset was increased either in the distal or proximal direction.

For the experimental condition in which the visual offset was the same for both sites (-6 cm for Site 1 and + 6 cm for Site 2, shown as a yellow cube in Figure 2.3), there was a higher probability of detecting one touch stimulation at the virtual offset corresponding to Site 2 (0.70) than that corresponding to Site 1 (0.38), with odds ratio of 0.26 and $p = 0.052$.

Subjects' Response Variability

Fig. 2.6 presents an overview of the proportion of perceived one touch for all subjects in both experiments. In order to determine if the probability of perceived one touch response differed across subjects, a logistic regression was carried out with response as the dependent variable and subject ID as a fixed effect, adjusted for offset, site and experiment type. Responses differed significantly across subjects ($p < 0.001$) with average responses across offset, site

and experiment varying from a minimum probability of 0.42 with standard error of 0.04 for subject 2 to a maximum of 0.89 with standard error of 0.02 for subject 1.

2.2.5 Discussion

There is some evidence that, for humans, vision contributes more than taction to unified visuotactile perception. Following this line of thought, we hypothesized that the visuotactile two-point discrimination distance (the smallest intrapersonal spatial conflict between vision and taction that is noticeable) would be different from the tactile-tactile counterpart. For instance, the multisensory fusion could place more weight upon the evidence provided by vision, resulting in greater precision and sensitivity, and a smaller two-point discrimination distance. There is also an argument for the opposite happening; increased reliance on vision could result in apparently conflicting tactile information to be disregarded, resulting in a larger two-point discrimination distance. In either case, the relative size of the distance is a clue about the mechanisms at play in fusing the visual and tactile touch sensations.

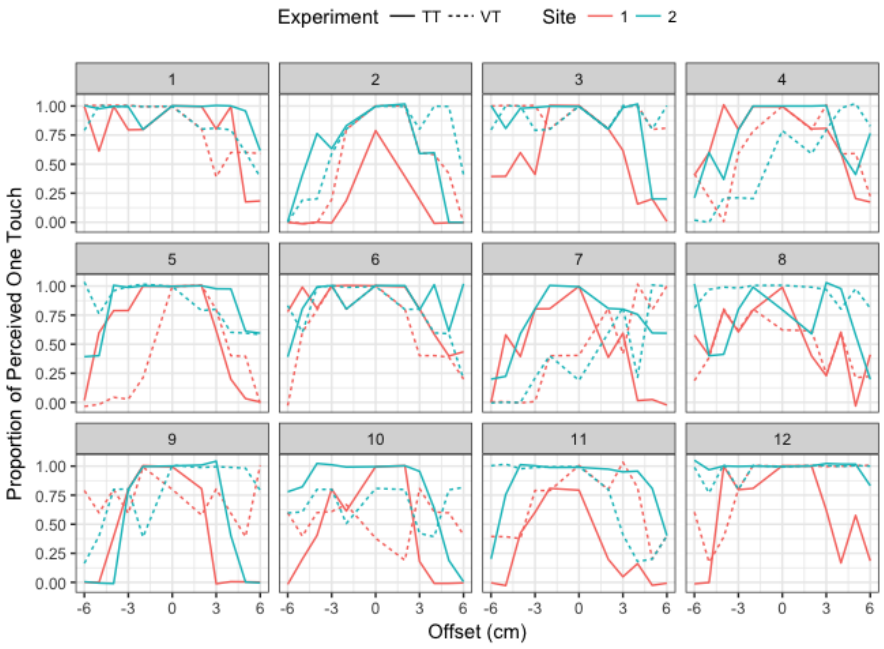

Figure 2.6: Overview of proportion of perceived one touch for all subjects across all sites and experimental conditions.

Surprisingly, as described in Section 2.2.4 and by comparing Figures 2.4 and 2.5, we can see that neither of these possibilities appears to have happened. We do not observe a signif-

icant difference in the rate of two-point discrimination variation by stimulation location for the Visuotactile compared to the tactile two-point discrimination test (Table 2.2). Specifically, if we look at Site 1, we see that, even though the probability of subjects perceiving one touch was higher for the Visuotactile (0.66) than the tactile (0.50) two-point discrimination test, it didn't reach statistical significance (odds ratio of 0.52 and $p = 0.072$). When looking at Site 2, this is even more pronounce since there is no difference in the probability of the stimulation being perceived as one touch between the tactile (0.82) and Visuotactile (0.81) test with odds ratio of 1.06 and $p = 0.85$. The proportion of perceived one-touch declines as conflict increases, as expected, but that decline is very similar for tactile-tactile and visual-tactile touch.

By comparison, if we look at the tactile-tactile two-point discrimination literature we see that, in [45], the two-point discrimination values for the anterior thigh ranges from about 2.8 cm in the proximal location to around 4.8 cm in the distal location. Additionally, [56] reported a summary of results for the mid portion of the anterior thigh with values ranging from about 2.7 cm to over 6 cm. From Figure 2.4, we see that the results from this experiment fall within those reported in literature. By defining threshold as the spatial separation corresponding to 75% correct responding, we see that for Site 1, it lies around 4 cm away for the distal (positive) direction and about 6 cm in the proximal (negative) direction. For Site 2, we see the threshold both distally and proximally at over 6 cm. Overall, these results follow the expected trend of the threshold value becoming smaller distally. It is possible that this has more to do with sensor richness, rather than central multisensory integration; the intermediate cutaneous nerve of the thigh branches out towards the knee, so there are simply more spatially diverse sensors there.

From Figure 2.5, we see that the Visuotactile two-point discrimination thresholds are still within those reported in the Tactile-Tactile literature. For both sites 1 and 2, the threshold lies at over 6 cm in either direction.

As mentioned in Section 2.2.2, there was one common virtual stimulation point for both sites 1 and 2 which is present at an offset of - 6 cm for Site 1 and + 6 cm for Site 2 and represented as an yellow cube in Figure 2.3. Even though this point is exactly the same for both sites, we see that there is a higher probability of perceiving one touch for Site 2 than for Site 1. This result holds true for both Visuotactile and Tactile-tactile versions of the protocol, with a combined probability of detecting one touch at Site 1 of 0.27 and a probability of 0.57 for Site 2. This is particularly striking considering how different both experiments are. While in the tactile-tactile experiment this can potentially be attributed to the way the intermediate cutaneous nerve is arranged along the thigh, this reasoning doesn't explain why there would be greater sensitivity approaching the knee when one of the stimulations is visual. This effect might be addressed using Bayesian models of perceptual inference and prior expectations for the reliability of the sensory channels [57]. An ideal-observer, knowing it has greater uncertainty at Site 2 due to diminished tactile spatial acuity

there, would be more likely to fuse the two stimuli perceptually.

Mental models of the body and alterations to the perceived size of body parts have an impact on tactile perception [58], [59]. In our study, this could manifest as participants' perceived thigh dimensions of the avatar. Fig. 2.7 shows, however, that there is not a strong relationship between thigh length and the probability of perceived one touch for each participant. A correlation test for the TT experiment alone gives a correlation of 0.31 ($p = 0.32$), VT alone a correlation of 0.34 ($p = 0.28$), and with both experiments combined a correlation of 0.38 ($p = 0.23$).

The traditional two-point discrimination test is usually reported as a threshold. Here we see a quite gradual decline of the likelihood of reporting a single touch. Reporting these results as a single threshold would fail to tell the entire story because we are asking subjects to integrate information from two sensory modalities (vision and touch) and to distill their answer into a binary choice (if they sensed a unified single touch or not). It is clear from the feedback provided by the subjects after the experiment and by observing their behavior during the experiment, that in most cases the choice was not trivial.

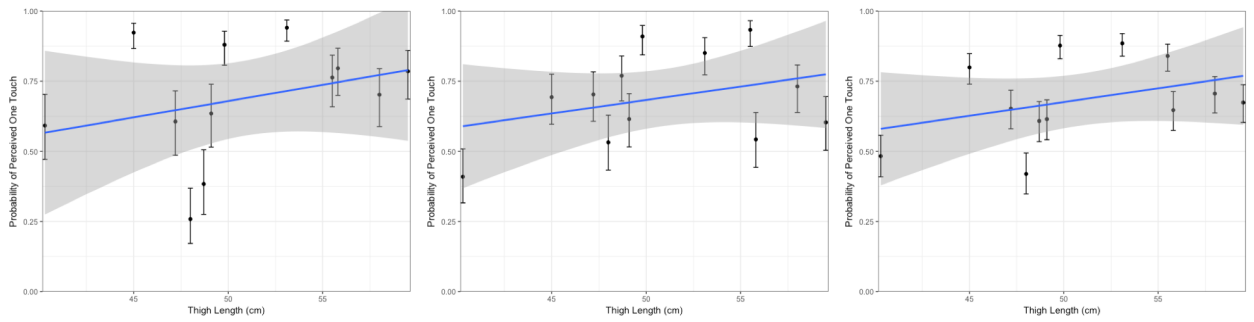

Figure 2.7: Probability of perceived one touch vs. participants' thigh length averaged across offsets and sites for (a) TT Test, (b) VT Test, and (c) both tests combined. Shaded area represents the 95% CI for the linear smooth across all 12 datapoints.

Limitations of the study

One limitation was the lack of control of the force applied for the TT portion of the experiment. In contrast with the VT experiment in which a monofilament was used, this experiment used an aesthesiometer instead. This practice is standard, but allows for the possibility of error or variability introduced by variable touch force.

Additionally, in spite of keeping the length of the experiment to a minimum and offering

breaks during and in between experiments, some subjects ended up feeling tired by the end of the session. This could potentially have impacted their level of attention which in turn could be reflected in their results. As might be expected for any task near the edges of perception, there was also a cognitive load associated with making decisions about what was felt. It is clear that, for various offsets in both stimulation sites, subjects had difficulty deciding their response. It should also be noted that the "threshold" from which to provide a one or two point answer was left up to the discretion of each subject.

Even though subjects were given the opportunity to try Virtual Reality either before or after the experiment (about half of the participants tried it once the experiment was over with the rest declining the offer), their degree of familiarity with this technology was not taken into consideration. This could be important to look into in order to assess potential biases they might have built regarding the reliability of what was being rendered to them visually during the Visuotactile experiment. The degree to which the virtual body was incorporated in the neural body representation could affect sensory integration.

Another factor that could have affected participants' expectations is the effect of order of presentation. All subjects were presented with the VT experiment first, followed by TT, rather than a counterbalanced design. It could be argued that, by being presented VT first, subjects developed an expectation of unreliability to tactile stimulus. This could have led subjects to become less trustworthy of their sense of touch once they were presented with the TT version of the experiment.

2.2.6 Conclusion

We performed a visual-tactile two-point discrimination test, analogous to the now-classic tactile-tactile test. This test fills a gap in the literature of intrapersonal spatial conflict between vision and taction. Understanding this sensory fusion is critical for the design of tactile feedback devices that work in concert with visual stimuli, such as in sensory-feedback prostheses or virtual reality. The test was performed on the thigh of twelve healthy adults with no sensory impairments. We observed two-point discrimination behavior similar in character and scale to those observed for the classic tactile-only test.

2.3 Visuotactile Two-Point Discrimination Test Along the Index Finger

As seen in Section 2.2, we discovered that the visual-tactile results behave similarly to its tactile-tactile counterpart. Due to this unexpected result, we proceeded to perform the Visuotactile Two-Point Experiment on the human hand. By exploring a region of the body where the tactile resolution is higher, we hoped to better unearth the mechanisms at play across these two senses. The procedure was similar to the thigh experiment, so that results

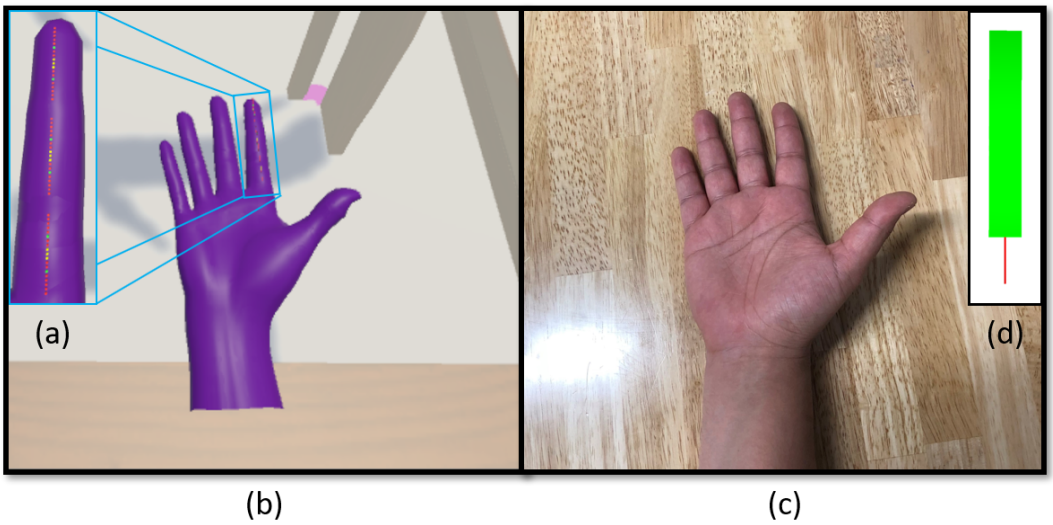

Figure 2.8: Subject's point of view in virtual reality (b) and the real world (c) and representation of the virtual monofilament used (d). Cubes indicate the different visual "stimulation" points with green being those that overlap the real locations of the tactile stimulation points, red being the visual offset locations, and yellow the common visual point for both real tactile stimulation points. These cubes are not shown to the subject during the experiment.

would be comparable. This test, however, focused on the palmar side of the right index finger.

2.3.1 Experimental Methods

Following up on our previous study [11], this experiment consisted of two separate parts. In one of them, we introduced subjects to the "traditional" tactile two-point discrimination test while the other consisted of the visuotactile two-point discrimination test. For this study, however, we performed the experiments along the index finger to analyze how our previous results hold up in a body part with higher tactile resolution.

Our main objective in this study was to continue the exploration on how vision affects our perception of touch by comparing a condition in which only tactile information was presented to the subject (two-point discrimination) to one where the visual information was inconsistent to the one given through the tactile channel (VT test).

Both conditions share some experimental designs and procedures. First, six sites were selected along the subject's palmar side of the right index finger so that two lied on the proximal phalanx, two on the intermediate phalanx, and two on the distal phalanx. Each

pair of sites within a phalanx were 8 mm apart from each other and 4 mm away from the middle portion of their respective phalanx. Sites were numbered in increasing order with the most proximal site referred to as Site 1 and the most distal site (closest to the finger tip) as Site 6. The study followed a counterbalanced design where the order of presentation of conditions were selected at random for each subject.

There were a total of fourteen (eight men and six women), non-smoking, subjects for this study without any health issues that could affect sensation in their right index finger. They were all college students with ages ranging between 19 to 32 years old and gave written consent before participating. The experimental protocols were approved by the University of Washington Institutional Review Board.

Tactile Two-Point Discrimination Experiment (TT Test)

To ensure only tactile stimulation was presented at each instance, subjects were placed in a seated position with their view obstructed by means of a (turned off) VR headset. We used a 2-point discriminator with adjustable points (aesthesiometer), which provides precise measurable distances, to accurately stimulate the targeted stimulation locations.

On each trial, one of the six stimulation points (one of sites 1-6) was randomly selected, followed by one of five possible offset conditions in either the proximal or distal direction to that site, or by the original point itself (zero offset condition). Offsets ranged from 1 to 5 mm (in 1 mm increments). This results in 11 possible stimulus offsets (including an offset of 0 mm) involving each stimulation point, with each offset repeated 5 times, resulting in a total of 330 stimulation instances (6 sites x 11 stimulus offsets x 5 repetitions). The choice of total repetitions per point to be explored was selected so that a high enough number of instances per offset condition was explored while minimizing the experiment time and, consequently, subjects' fatigue level. In average, this section of the study lasted about 45 minutes with some subjects opting to take a break during the experiment.

Throughout the experiment, the subject was stimulated for approximately 0.5 s and asked to report verbally how many points s/he perceived using their best judgment.

Visuotactile Two-Point Experiment (VT Test)

For this portion of the study, a 1.4 g medical grade monofilament was used to stimulate one of the six sites marked on the palmar side of subjects' right index finger. By using a monofilament, we can guarantee a constant force was applied throughout the experiment [52]. Furthermore, by selecting a 1.4 g monofilament, we can ensure detectability of the tactile stimulus while being able to compare our results with those in the traditional two-point

stimulation literature [56].

Additionally, a virtual environment with a virtual hand was built with the Unity game engine [53]. The virtual hand was modeled using Autodesk Maya [60] to make the finger dimensions match those of the subject's hand as closely as possible. The "uncanny" valley, a phenomenon in which subjects can feel uncomfortable when presented with human-looking avatars or limbs, was considered during the design of the experiment. Due to the complexity of designing a highly realistic personalized hand for each subject which could avoid the "uncanny" valley, as well as the many open questions related to this phenomenon [54], the decision was made to make a hand with a realistic human shape but with a unicolored texture (light purple). The virtual hand was placed with its palmar side facing up on top of a virtual table, resembling the subject's hand pose during the experiment. The HTC Vive [55], a virtual reality headset, was used to present the virtual environment to the subject. A custom-designed 3D printed adapter attached the monofilament to the bottom of one of the HTC Vive's controllers. A simplified virtual representation of the controller-monofilament combo was also present in the virtual environment as seen in Fig. 2.8(d).

Visual-tactile stimulation was presented in a random order with random offsets for all six tactile stimulation points. Figure 2.8 shows the perspective as seen from the subject's point of view both in virtual reality (Fig. 2.8(b)) and in the real world (Fig. 2.8(c)). Please note that, for the VT experiment, only the virtual hand is displayed to the subject. This figure also shows the different virtual stimulation points as colored cubes overlayed atop the avatar's right index finger (Fig. 2.8(a)). Green cubes depict the stimulation points in real-life on the subject (0 mm offset) and red cubes depict the visual location of touches occurring in virtual reality.

There were three virtual points common to each pair of sites on each phalanx. These represented a 3, 4 and 5 mm offset in the distal direction for the proximal site and a 3, 4 and 5 mm offset in the proximal direction for the distal site. There were a total of 9 common virtual points (3 per phalanx x 3 phalanx) and are depicted as yellow cubes in Figure 2.8(a). Note that the cubes were not seen by the subject during the experiment but are shown in this figure for clarification purposes only.

The same procedure as in Section 2.3.1 was followed during this experiment, with the subject being presented a randomly generated offset for approximately 0.5 s and asked to report verbally if they perceived one or two points. The difference, however, is that now one of the touches is a visual one. The visual location of touch was offset, while the tactile touch was always located in one of the six stimulation sites (green cubes in Figure 2.8(a)). Thus, subjects were instructed to treat the visual and tactile stimuli as two points if they experienced them as coming from two separate locations. This section of the study had an average duration of 50 minutes.

In total, touches at each virtual point were delivered five times throughout the experiment

for a total of 330 stimulation instances per subject.

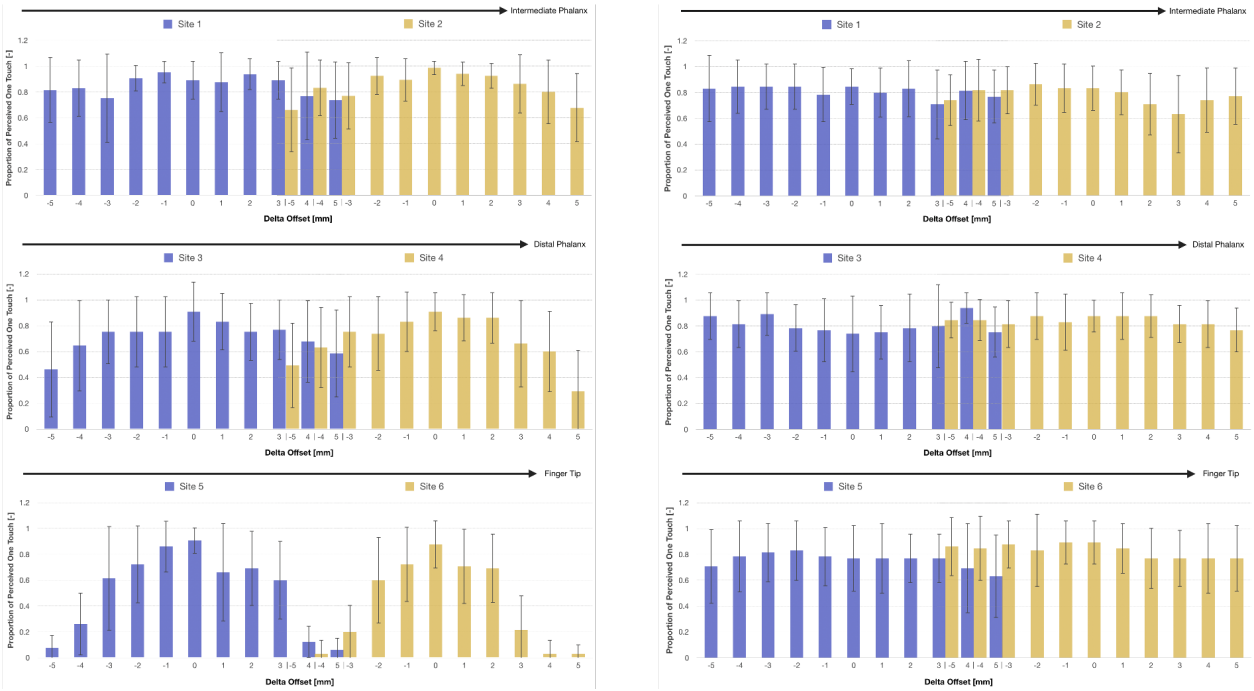

Figure 2.9: Proportion of perceived one-touch with their respective standard deviation for all offsets explored averaged across all subjects in the Tactile Two-Point Experiment (left) and Visuotactile Experiment (right).

2.3.2 Statistical Methods

Two sets of analyses were carried out corresponding to the TT and the VT experiments. Generalizedlogistic mixed effects regression was used to estimate the proportion of perceived one-touch responses (the dependent variable) by site by offset interaction (the fixed effects, both variables modeled as categorical, using 5 dummy variables for site, and 10 dummy variables for offset) and to test for differences in perceived one-touch at proximal sites compared to the most distal site (site 6), at each offset. Site by study participant interaction was modeled as a random effect (i.e., differences in perceived one-touch responses across site were allowed to vary by participant). First, the symmetry of the proportion of perceived one-touch responses from the zero offset in either direction was tested by comparing the model above (the full model) to a reduced model where offset is parameterized as the absolute value for the distance from the zero offset (using 5 dummy variables) using a likelihood ratio test and the BIC (Bayes Information Criterion). If no difference was found, then the model based

on absolute value distance from the zero offset was used to estimate the proportions. Then, the site by offset interaction was tested for significance. If the interaction was not found to be significant, then models were carried out with site and offset as independent main effects with the focus on differences across site adjusted for offset. Mean differences in response are presented as odds ratios (OR) and 95% confidence intervals (CI) for a one touch response per for each site compared to the most distal site (site 6). To compare common locations of points due to overlapping sites (9 pairs of points, for example, site 1, offset 3 is at the same location as site 2, offset -5), comparisons of the proportion of perceived one- touch responses within each pair of points was carried out using estimates from full model regression described above. All pair-wise comparisons were adjusted using simultaneous inference [61].

Response change across consecutive trials was computed (e.g., a change in response would be coded if a participant perceived a one touch response for a given trial and then perceived a two-touch response in the next trial), to assess if fatigue influenced the response to stimuli, where we hypothesized fatigued participants would choose the same response as the experiment progressed regardless of offset or site. Logistic mixed effects regression was used to assess the association between consecutive response change (the dependent variable) by trial number (the independent fixed effect modeled as a linear slope of change) with site and offset as model covariates and random effects for trial number by participant interaction (i.e., intercepts and slopes were allowed to vary across participant). Analyses were carried out using R 3.5.1 [62], using packages glmmTMB [63], multcomp [61], emmeans [64] and tidyverse [65] with type 1 error set at .05.

2.3.3 Results

Proportion of Perceived One-Touch

As in [11], the key measure for the two experiments is the *proportion of perceived one-touch*, which is the proportion of the number of times a subject reported feeling one touch to the total number of trials presented. As explained in Section 2.2.4, we expect this proportion to be higher when the offset between the two tactile stimuli (in the tactile-tactile condition) or the tactile and visual stimuli (in the visual-tactile condition) is small.

Tactile Two-Point Discrimination Experiment

The *proportion of perceived one-touch* across all subjects for the TT experiment is presented in Fig. 2.9 (left) with the most proximal data point (offset of -5 mm at Site 1) presented on the top-left side of the figure while the most distal one (offset of +5 mm at Site 6) is located to the bottom-right side of it.

As expected, the highest proportion of perceived one-touches occurred at the zero offset condition. This was the only condition in which a one touch stimulation was presented to the subjects. Additionally, there was no significant difference in the slope of change in responses as offsets increased either in the distal or proximal direction.

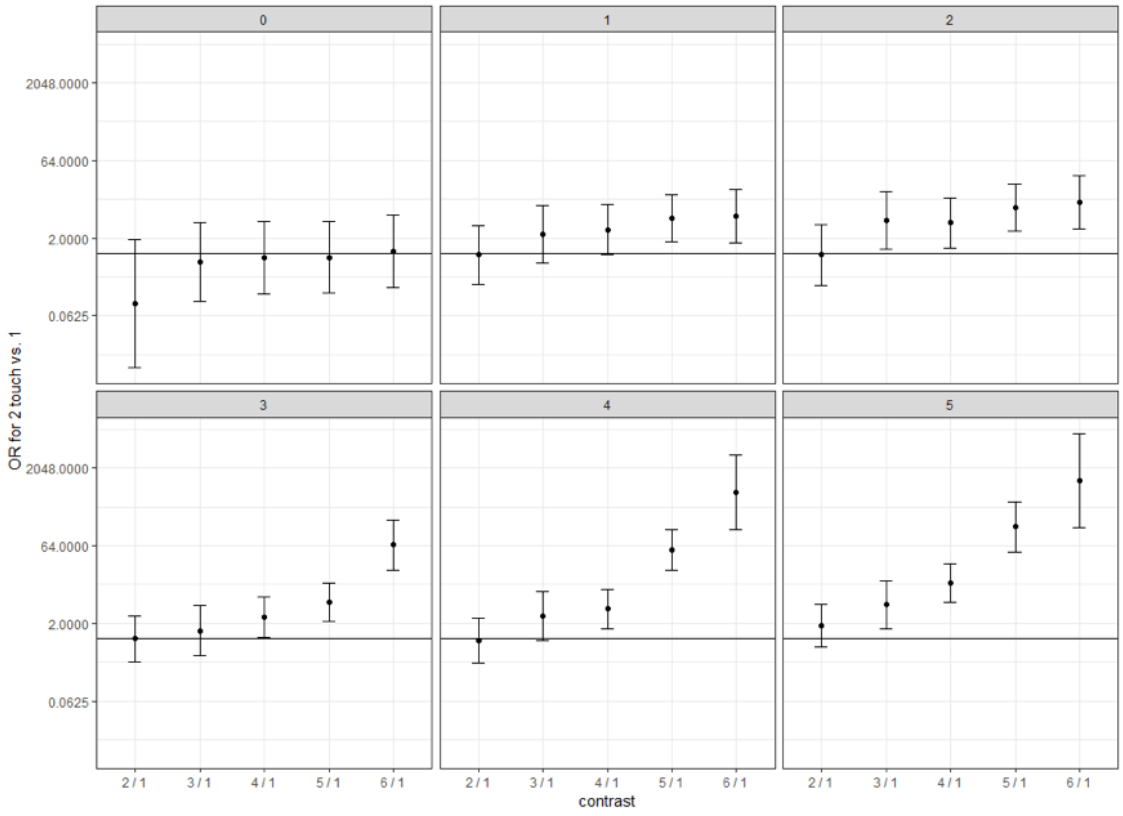

Figure 2.10: Summary of odds ratios for two vs. one touch for Sites 2 to 6 compared to Site 1 in TT experiment

Figure 2.10 summarizes the odds ratios (ORs) for the probability of one touch response for Sites 2 to 6 compared to Site 1. It is clear that the probability of perceiving one touch decreases as both site and offset increases, with ORs decreasing towards zero.

For the two-point touch conditions in which each pair of sites within a phalanx shared three common points (denoted as yellow cubes in Figure 2.3), only Site 1 offset +3 mm differed from Site 2 offset -5 mm ($p < 0.001$) with a more borderline significant result between Site 3 offset +3 mm and Site 4 offset -5 mm ($p = 0.018$).

Visuotactile Two-Point Discrimination Experiment

Figure 2.9 (right) presents the average *proportion of perceived one-touch* across all subjects with the top-leftmost point being the most proximal (representing an offset of -5 mm in Site 1) and the bottom-rightmost data point being the most distal (offset of +5 mm in Site 6).

For this portion of the experiment, results are sharply different than in those in the TT experiment. Regardless of site number or offset amount, subjects mostly perceived a one-touch event throughout the experiment. The average *proportion of perceived one-touch* remained relatively constant, with no differences in the probability of perceiving one-touch by site across offsets. As can be seen from Fig. 2.11, the probability of perceiving one-touch independent of offset for all sites range from a low of 0.82 at Site 5 to a high of 0.886 at Site 4.

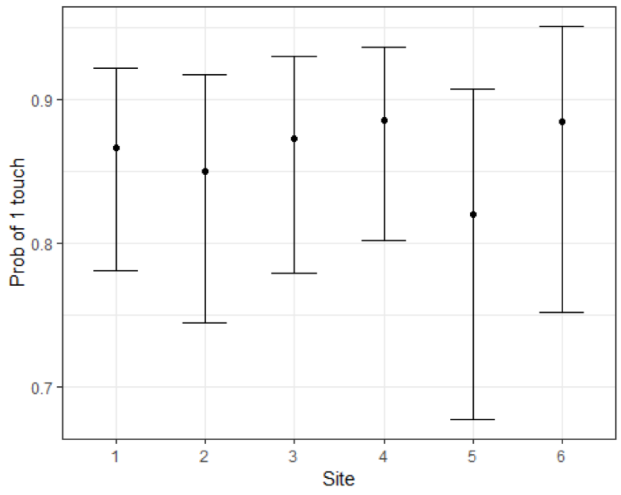

Figure 2.11: Probability of one touch independent of offset in VT experiment

For the experimental conditions in which the visual offset was the same for each pair of sites within a phalanx (yellow cubes in Figure 2.3), no differences were found.

Potential Effect of Fatigue on Observed Results

Due to the increased duration of this study compared to [11], as well as to the feedback provided verbally by several subjects, investigating potential effects of fatigue is important. One way this could manifest itself is by subjects becoming tired during the experiment and deciding to respond the same way. To check for such possibility, we looked at the probability of a response changing by increasing trial. After adjusting for site and distance from zero

offset, we found a slight decreasing trend (p = 0.033) in the tactile two-point experiment and no significant association between change in response and trial number (p = 0.73) in the visuotactile experiment. Overall, there was not a big difference with the probability of a change in response which went from 0.4 at the beginning of the experiment to 0.34 at the end in the tactile two-point experiment, and from 0.23 to 0.25 for the visuotactile experiment. An overview of the proportion of responses that differed from previous trial (per every 10 trials) for all participants can be seen in Figures 2.12 (for the tactile two-point experiment) and 2.13 (for the visuotactile experiment).

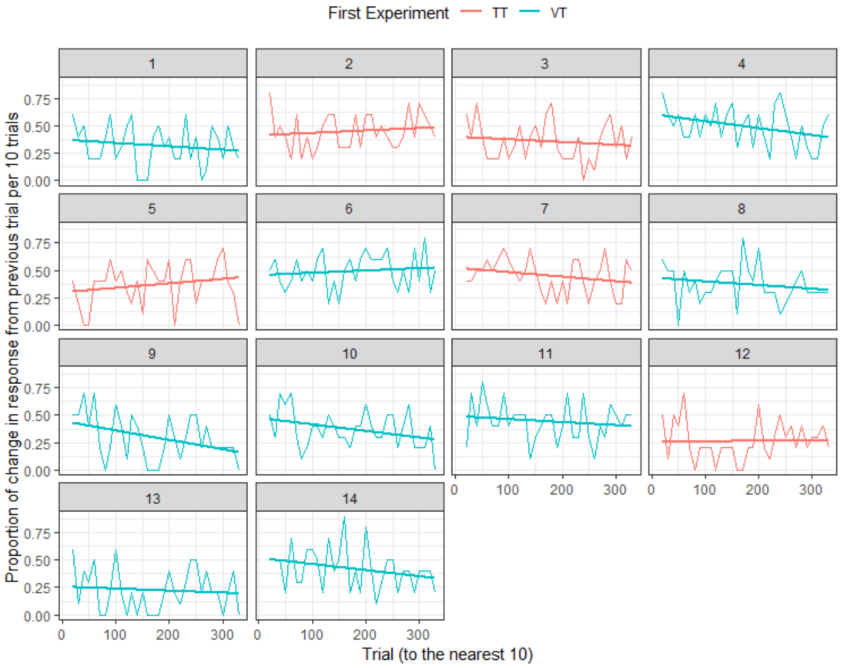

Figure 2.12: Tactile Two-Point Experiment: Proportion of responses that differed from previous trial per every 10 trials for all participants.

2.4 Visual-Tactile Consistency for Mapping of Targeted Reinnervation

Targeted Reinnervation (TR), a nerve transfer surgery, is becoming a common procedure to restore sensation in amputees [66]. When stimulated in their TR site, subjects report the feeling of being touched at their amputated (phantom) limb. Most of the work in this area has focused on characterizing this sensory recovery in the upper limb. We are currently building this understanding for lower limb [67]. One big issue with the current methodology, however, is its reliance on participants subjectively reporting the location of their perceived sensation.

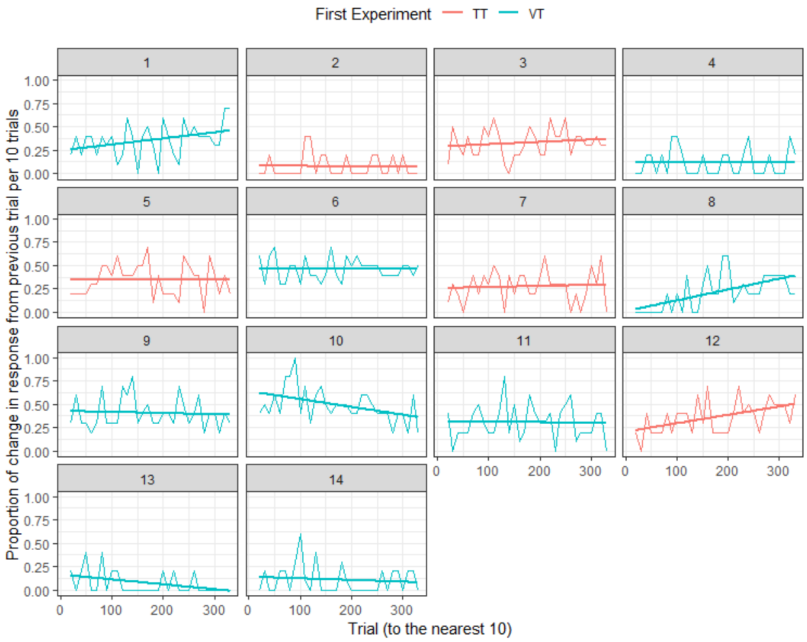

Figure 2.13: Visuotactile Experiment: Proportion of responses that differed from previous trial per every 10 trials for all participants.

We developed a protocol wherein visual touches are delivered to the subject using a Virtual Reality headset. At the same time, touches are delivered to the skin at the TR surgery site. Subjects respond whether the touches seem consistent (the touch appears visually where it feels like it is coming from). This replaces the current state of the art, subjective verbal response, with a forced-choice behavioral test. However, the surgeon who was performing these procedures in Seattle, Jason Ko, moved back to his previous institution so we have not been able to deploy this system on subjects with TR surgery. The experiments presented in Sections 2.2 and 2.3 have, however, been motivated by the work we developed for this area.

Chapter 3

EFFECT OF SENSORY CONFLICT IN THE EXPRESSION OF DYNAMIC PRIMITIVES

Dynamic Primitive Theory is an influential theory of neuroscience providing potential explanations on our brain's ability to generate complex behaviour while dealing with the noisy and slow signals from our sensory system. In this chapter, we investigate for the very first time what role sensation has in the way that these primitives are expressed. By methodologically performing experiments centered around vision and proprioception, we observe how the point of transition between two of these primitives (discrete movements and rhythmic actions) is affected by sensory conflict interventions.

If sensation does indeed play a role in how our brain chooses to switch between primitives, we expect the point of transition to change when errors between vision and proprioception are introduced. Although we do not see conclusive evidence of these senses influencing dynamic primitive expression, we conclude that they might still be doing so indirectly by influencing the updates of the internal models of our body.

3.1 Introduction

Humans are able to execute motor tasks with impressive speed despite the slow neural elements comprising the neuromotor system. It is hypothesized that humans overcome these performance limitations by relying on internal models that are based on compositions of dynamic primitives (DPs) [12], [13]. These primitives (discrete movements, rhythmic actions, and mechanical impedances) simplify control and require limited high-level supervision. Forcing the limits of a discrete movement by increasing temporal demand has been established as a method to observe the transition from a discrete to a rhythmic primitive during a reaching task [14]. This transition has been exposed by modulating a temporal feature of top-down neural control, but it is yet unclear whether these descending neural signals are in exclusive control of the transition. Understanding how primitive expression depends on sensory information, or bottom-up signals, will help build a full picture of the nervous system's ability to generate complex actions.

3.1.1 Objective

Using the existing methodology as a guide [14], we aim to measure the point of transition between discrete and rhythmic primitives under the influence of sensory conflict. If there is any reliance on visual or proprioceptive feedback, we expect the point of transition to be offset.

Figure 3.1: Overview of experimental setup in real life. Subject movement is tracked using a virtual reality controller and virtual targets are displayed through a headset

3.1.2 Experimental Design Considerations

The first step towards our goal is to reproduce the findings of [14] using a Virtual Reality first-person view. Instead of a 2-dimensional plane (as in [14]), here we introduce a 3-dimensional environment that is able to more accurately meet the expectations of movement and contribute to immersion. This approach removes potential artifacts that may have been introduced in the original study while also exploring the impact a proprioceptively aligned version of it has on the expression of dynamic primitives. Specifically, in [14], there were two details in the experimental setup which could have affected the results. First, the experimenters used a monitor placed 65 cm in front of the subjects to display the targets' position and a cursor indicating their movement. Due to such a setup, subjects were asked to perform the reaching task on a different plane as the one the visual feedback was provided on. Second, the display gain was set to 0.5 scaling everything in half. It is not clear why this

choice was made but the assumption is that this was dictated by hardware limitations of the display (anything larger would place the targets outside the boundaries of the display).

3.2 Methods

3.2.1 Participants

A total of twenty participants (age: {M = 21, SD = 4.3}, 8 females, all right handed) participated in this study. None of the participants had any health issues that could affect their motor performance or impact their sensory experience. One subject felt nauseated during the experiment and opted out. Most subjects reported having little to no virtual reality experience. They all were given an overview of the experiment before the start of the study and gave written consent before participating. Subjects received a twenty dollar gift card as compensation for their time. The experimental protocols were approved by the University of Washington Institutional Review Board.

3.2.2 Experimental Setup

An overview of the experimental setup can be seen in Figure 3.1. Subjects are seated in front of a table with their sternum touching the table edge. Chair height is adjusted for each participant so that their upper arm are position at roughly 45° to the horizontal with their forearms resting on the table. Their shoulders are fixated to the chair by means of a shoulder brace. To reduce the effect of static and dynamic friction while minimizing wrist motion, subjects wore a wrist brace throughout the experiment.

A virtual environment containing a table with similar dimensions as the one in real life was built using the Unity game engine [53]. Prior to the start of the experiment, a calibration procedure is performed to ensure that the real table is collocated with its virtual counterpart. The virtual scene is rendered to the subject by means of a virtual reality headset (Oculus Rift, 2160x1200 pixels, 110 degree field-of-view, 90 Hz refresh rate). In the virtual scene, a "neutral point", located 23 cm away from the edge of the table (in the anterior-posterior direction), is used to place each subject in a comparable position. Two targets are also presented to the subject. To minimize the importance of accuracy, these targets have a 5cm radius. An Oculus Rift controller is use to track the subject's hand in space and to render it in the virtual environment. A custom 3D printed adapter is attached to the bottom of the controller, making the contact point between the controller and the table flat in such a way that allows the controller to stand up when not in use. This makes performing the experimental task more comfortable for the subject. An overview of this setup can be seen in 3.2.

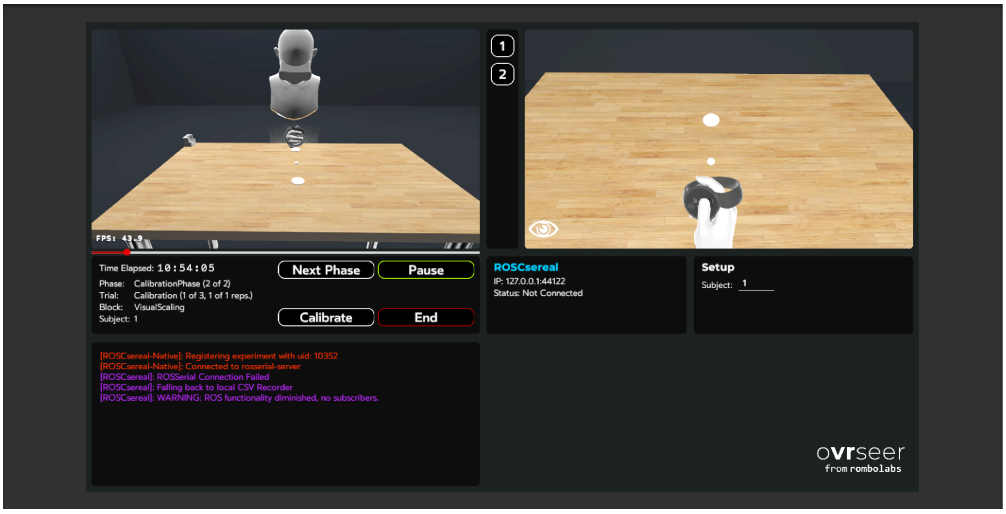

Figure 3.2: GUI with third person view of the virtual environment (left) and subject's point of view (right). Virtual targets located 28cm apart are displayed on top of the virtual table.

3.2.3 Experimental Design and Conditions

The trials in this experiment consisted of two phases: the washout phase and the test phase. In the washout phase, subjects are introduced to the virtual environment and are allowed to move at will without any sensory manipulation. After approximately ten seconds, subjects are then instructed to place their hand on the "neutral point" at which point the test phase begins. During the test phase, subjects are presented with one of five possible conditions at random as described below (with t_d representing the target distance used in [14] and having a value of 28 cm):

- **Condition 1:** Visual Target Distance = t_d, Proprioceptive Target Distance = t_d
- **Condition 2:** Visual Target Distance = $1.35 * t_d$, Proprioceptive Target Distance = $1.35 * t_d$
- **Condition 3:** Visual Target Distance = $0.65 * t_d$, Proprioceptive Target Distance = $0.65 * t_d$
- **Condition 4:** Visual Target Distance = $0.65 * t_d$, Proprioceptive Target Distance = t_d
- **Condition 5:** Visual Target Distance = $1.35 * t_d$, Proprioceptive Target Distance = t_d

In Conditions 1, 2 and 3, the visual representations of subjects' movement corresponds exactly to their real movements. In conditions 4 and 5, however, subjects visual representation of their hands are dynamically scaled by either 0.65 or 1.35. A value smaller than one has the effect of "compressing" the subject's movement (Figure 3.3) while a value larger than one has the opposite effect (Figure 3.4). This design enables us to explore the effect of visual and proprioceptive interventions on the expression of dynamic primitives. If there are differences in the outcome measures between Conditions 1, 4 and 5, it will provide evidence of sensory contributions to the expression of dynamic primitives. Specifically, we hypothesize that subjects will transition from discrete to rhythmic behavior later in Condition 5: perceiving their hands moving faster will enable subjects to maintain discrete movements longer. As such, we expect Condition 4 to show subjects transitioning between these primitives earlier.

Additionally, we introduced Conditions 2 and 3 as "control" conditions for Conditions 5 and 4 respectively. The washout phase allows us to avoid subjects carrying over any potential adaptations from previously experienced conditions.

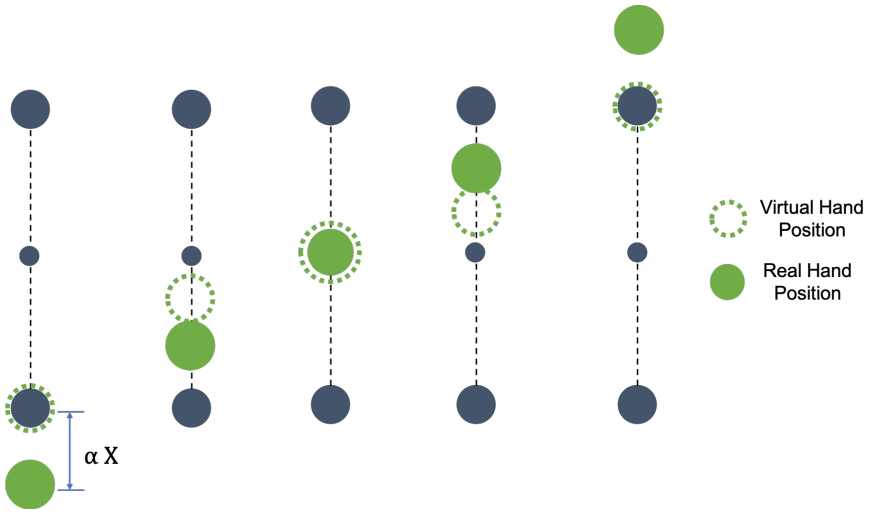

Figure 3.3: Representation of dynamic scaling condition (Section 3.2.3) with scaling factor smaller than 1. Visual representation of subject's movement is compressed.

3.2.4 Procedure

Subjects were seated at the table in the configuration described in Section 3.2.2. Once inside the virtual scene, they are presented with one of the five experimental conditions (see Section 3.2.3) chosen at random. Before each trial begins, subjects are given approximately ten seconds to get familiar with the virtual environment and encouraged to move freely while

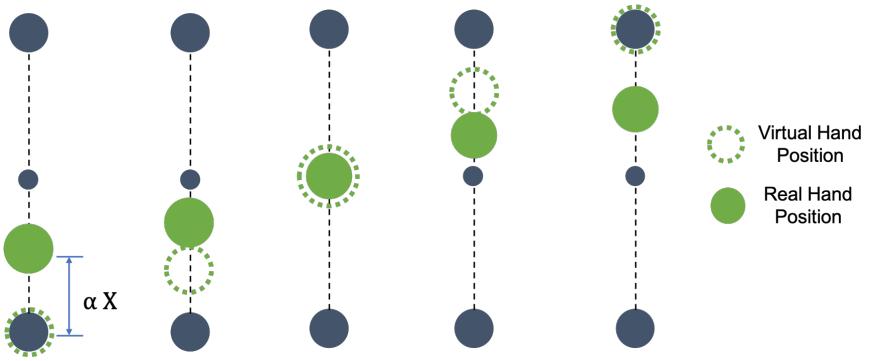

Figure 3.4: Representation of dynamic scaling condition (Section 3.2.3) with scaling factor greater than 1. Visual representation of subject's movement is enhanced.

holding the controller with their right hand. During this phase, no sensory manipulation is present and target distances in virtual reality are dictated by the chosen condition. They are then asked to place their right hand on the "neutral point" at which point the test phase begins. During the test phase, subjects are asked to move their right hand back and forth between the two targets in synchrony with the metronome sounds. They are instructed to perform these point-to-point movements with an explicit dwell time between them lasting for 50 % of the metronome interval as accurately as possible. Additionally, they were asked to maintain such discretization of movements for as long as possible, even for portions of the trial in which the metronome pace increases making the task difficult. It is during this phase that sensory intervention is introduced for those conditions that have them (Conditions 4 and 5). The exact script read to the subject was: *"The experiment consists of several trials in which you will be asked to move your hand back and forth between two targets while holding a VR controller. At the start of the trial, place your hand in the neutral position (smallest circle on the table in front of you). During the trial, please move your hand between the two targets (two biggest circles) in synchrony with the metronome sounds. Perform these point-to-point movements with an explicit dwell time between them lasting for 50 % of the metronome interval as accurately as possible. Try to maintain this discretization of movements for as long as possible, even for portions of the trial in which the metronome pace increases."*

In all variations of this study, the metronome profile lasted for 158s and contained a total of 140 moves broken down as follows:

- 10 sounds in 2 s intervals
- 50 sounds with decreasing intervals of 36 ms ending at an interval of 200 ms
- 20 sounds in 200 ms intervals

- 50 sounds with increasing intervals of 36 ms ending at an interval of 2 s
- 10 sounds in 2 s intervals

A visual overview of the metronome profile is presented in Figure 3.5. All conditions were presented twice, with order selected at random. The experiment lasted for about 50 min, with subjects allowed to take breaks at any point during the experiment (most took a break at the half-way point).

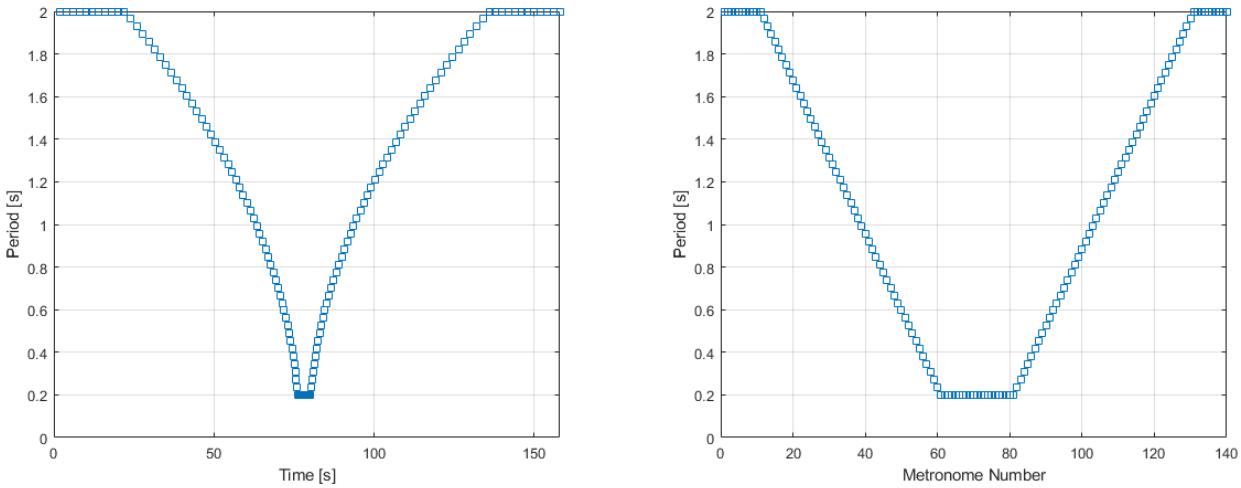

Figure 3.5: Metronome profile used for the experiment as a function of time (left) and metronome number (right).

3.3 Analysis

The controller's position and orientation was recorded throughout the experiment. However, only data corresponding to displacement in the saggital direction was processed. Before extracting quantitative markings, the positional data was smoothed using a five-sample moving average filter, with centered filtering that did not introduce lag. This same five-sample moving average filter was then used to smooth out the velocity signal calculated numerically from the two-sample difference of the position signal. Figure 3.6 shows the positional data for a full trial with metronome onsets as red vertical lines and onset/offset of each movements in green/red circles. It is clear from the figure that dwell times are present at the beginning of the trial but slowly fades away as the metronome pace increases until disappearing completely. It then reappears as metronome pace decreases. However, in order to properly find

the points at which dwell times appear/dissapear we need to first break down the kinematic profile of the trial into their individual movements.

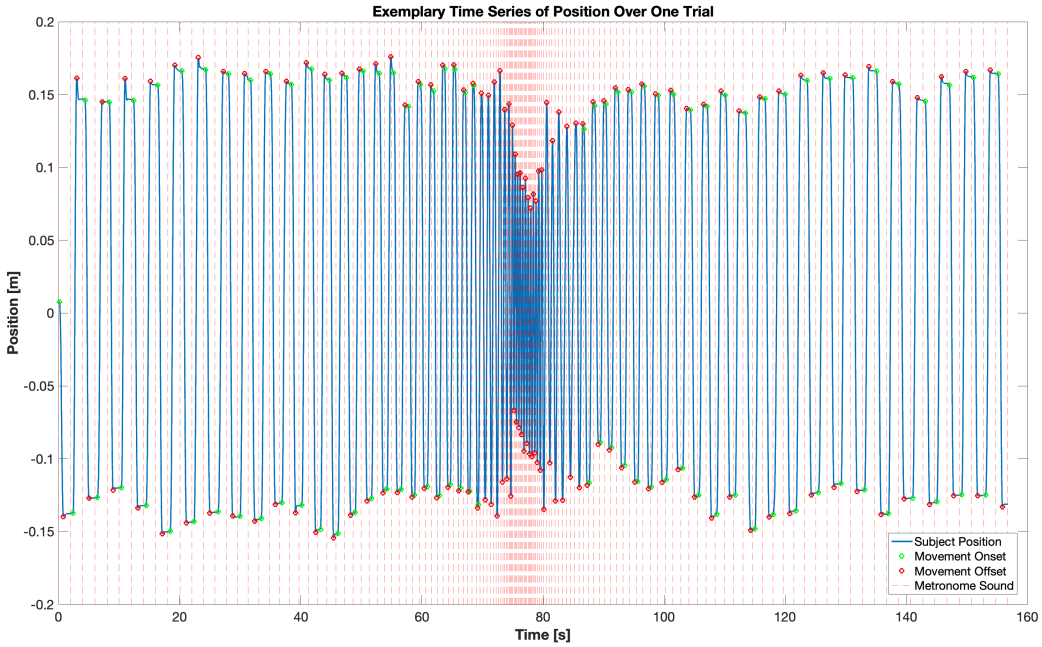

Figure 3.6: Time series of position data for one trial. Red vertical lines represent the onset of the metronome sounds, while green and red circles represent the onset and offset of each movement respectively.

Single movements are delimited by t_{onset} and t_{end} (green and red circles in Figure 3.6 respectively). These features are found by looking at the velocity profile of the trial. Specifically, onset and offset of a movement is defined as the points were each velocity peak crosses a threshold, defined as 3% of the peak velocity of the same movement. This parsing method, however, stops working reliably when the single movements merge and dwell time disappears as can be seen in Figure 3.7. To solve this problem and eliminate false movement detection, a linear fit is applied to the velocity data between each parsed movement (velocity samples below the 3% threshold). If the R^2 of the linear fit is greater than 0.99, the algorithm makes t_{onset} and t_{end} coincident.

In addition to the onset and offset of each movement, there are other features needed in order to identify the point of transition between discrete and rhythmic movements. Specifically, we calculate movement time MT_i, dwell time DT_i and duty cycle DC_i as:

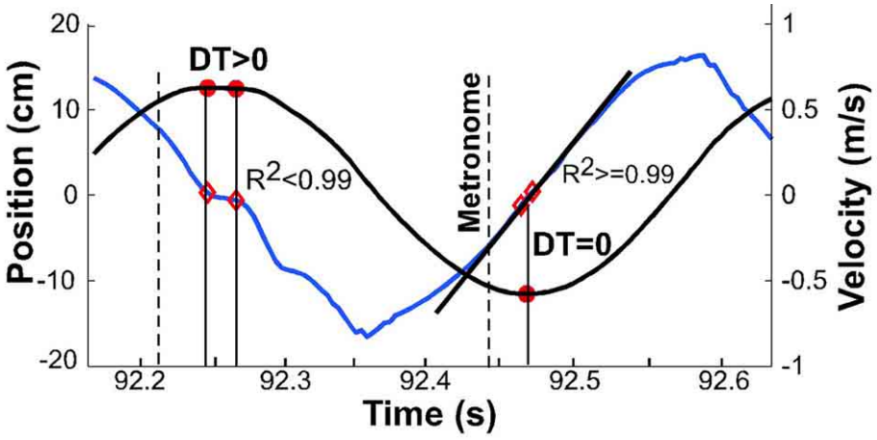

Figure 3.7: Exemplary kinematic profile with position represented by the black line and velocity by the blue line. As can be seen here, difficulty arises when dwell time disappears.

$$MT_i = t_{end,i} - t_{onset,i}$$
$$DT_i = t_{onset,i+1} - t_{end,i}$$
$$DC_i = MT_i/(MT_i + DT_i)$$

3.3.1 Transition between discrete and rhythmic movements

The point of transition between discrete and continuous rhythmic movements occurs when the dwell time disappears during the trial. In other words, the transition from discrete to rhythmic, $DT = 0_{Accel}$, is defined as the movement i at which dwell time was zero during the accelerating portion of the metronome and which was followed by at least two more movements with zero dwell time. Similarly, transition from rhythmic to discrete, $DT = 0_{Decel}$, was defined as the movement i when dwell time reappeared and which was preceded by two movements with zero dwell time.

3.4 Results and Discussion

As seen from the exemplary trial in figure 3.6, subjects started and ended trials with clear dwell times indicating that subjects were implementing discrete movements. The movements tended to merge around 70 seconds into the trial and they reappeared about 15 seconds after. Overall, movements were synchronized with the metronome. This synchronization tended

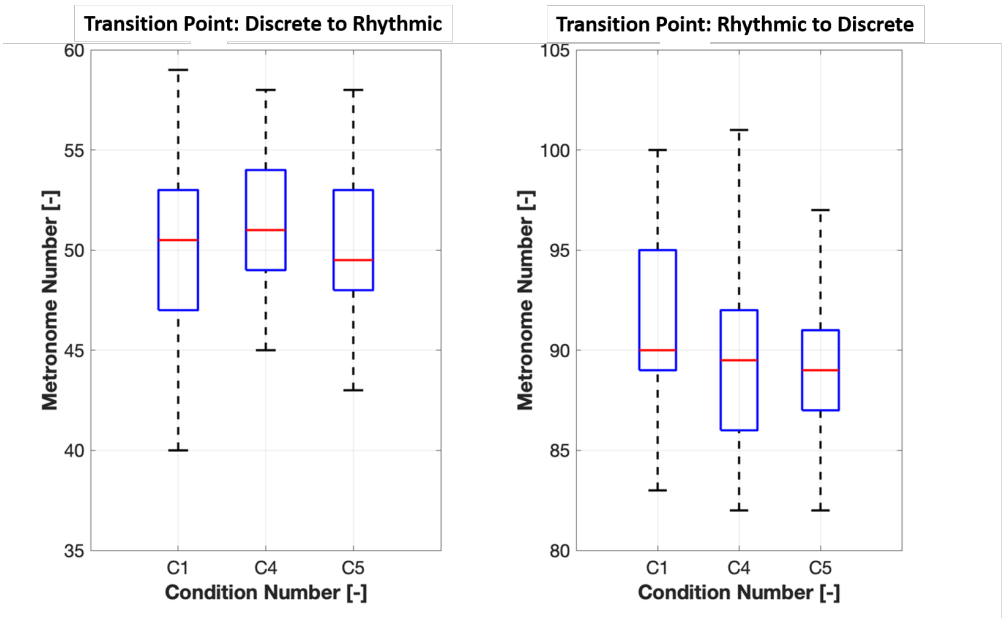

Figure 3.8: Metronome number at which transition from discrete to rhythmic (left) and rhythmic to discrete (right) occurred for all subjects in conditions where the distance traverse was proprioceptively the same.

to be lost, however, in the region of fastest metronome interval with most subjects missing one or more cycles or, in some instances, even introducing extra movements. Subjects also tended to report verbally how difficult it was for them to keep up with the metronome during the fast portion of the trial.

Movement time (MT), dwell time (DT), and duty cycle (DC) were calculated for all subjects. Data for one subject in condition 1 can be seen in Figure 3.9 with movement time in blue, dwell time in green, duty cycle in red and metronome interval in black. Dwell time steadily declines for the first portion of the trial until fully disappearing around metronome number 48, coinciding with duty cycle reaching its maximum. Duty cycle then starts to decrease once again at around metronome number 95, with dwell time reappearing at that point. Movement time, starts relatively steady value both at the beginning and end of the trial.

The relationship of interest to us, however, is that between movement time and dwell time. Subjects were instructed to perform the experiment in a way that would make their movement times and dwell times to decrease proportional to each other until they reach the minimum movement time they can achieve. What we see is that, although they both do decrease, they do so at different rates. More importantly, we see that dwell time decreases at a faster rate than movement time, with the former disappearing before the latter reached their

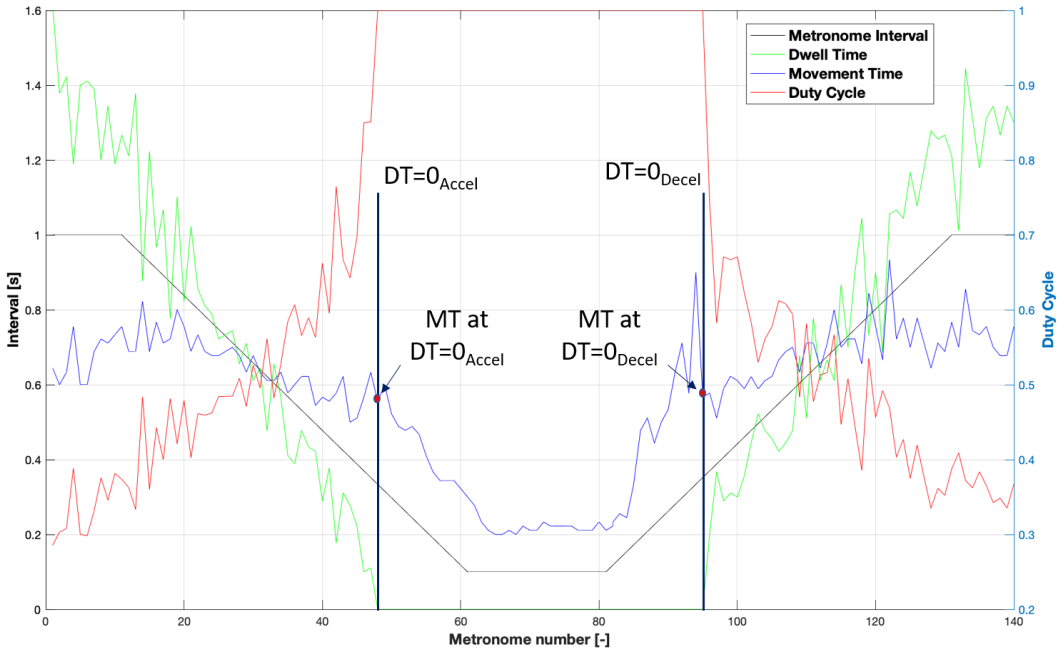

Figure 3.9: Movement Time (blue), Dwell Time (green) and duty cycle (red) for one trial. The point of transition are also shown as $DT = 0_{Accel}$ (for discrete to rhythmic) and $DT = 0_{Decel}$ (for rhythmic to discrete) with their respective movement times.

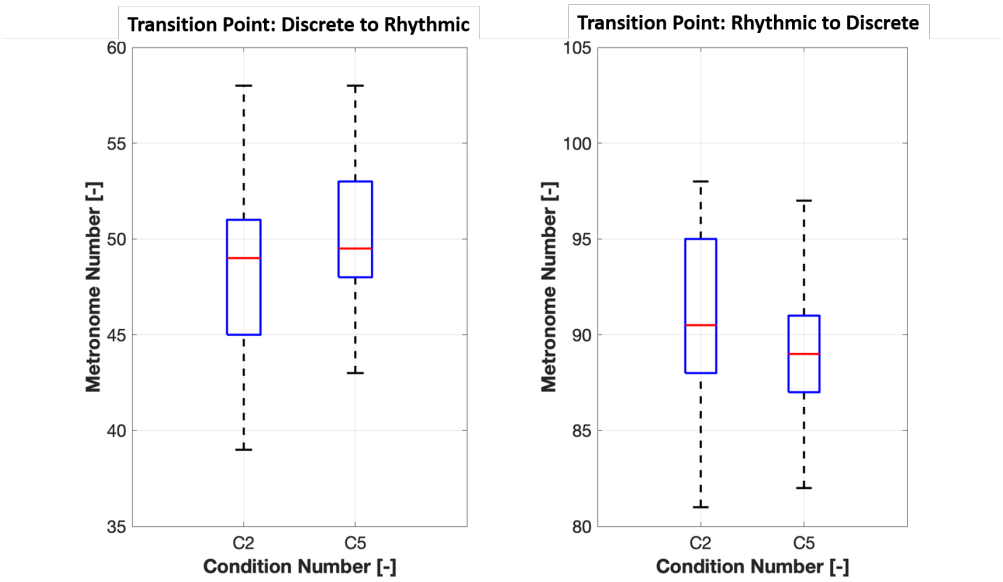

Figure 3.10: Metronome number at which transition from discrete to rhythmic (left) and rhythmic to discrete (right) occurred for all subjects in the first pair of conditions where the distance traverse was perceived visually to be the same.

minimum. In the accelerating portion of the trial, movement time starts to increase before dwell time reappears. These observations are important since they indicate that subjects stopped performing discrete movements, which become more challenging with speed during the accelerating portion of the metronome. Instead, they switch to rhythmic movements to comply with increasing timing demands.

The point of transition between discrete and rhythmic movements was found as described in section 3.3 for both the accelerating ($DT = 0_{Accel}$) and decelerating ($DT = 0_{Decel}$) portions of the trials. Figure 3.8 summarizes both of these transition points for all subjects in conditions where the distance traverse was physically the same. As can be seen from this figure, subjects transition from discrete to rhythmic movements before the fastest portion of the metronome was reached at metronome number 61 in all conditions. Similarly, all subjects switch back strategies (from rhythmic to discrete) after the fast metronome period was over at metronome number 81. The same was seen in the pair of conditions where the distance traveled by subjects was presented to be visually the same (Figures 3.10 and 3.11).

From a sensory perspective, we are interested in potential differences between conditions where one of the sensory modalities was modified. To do so, we performed a set of three two-way ANOVAs with condition and trial as factors. In the first ANOVA, we compared conditions 1, 4 and 5, where the proprioceptive distance across conditions was the same

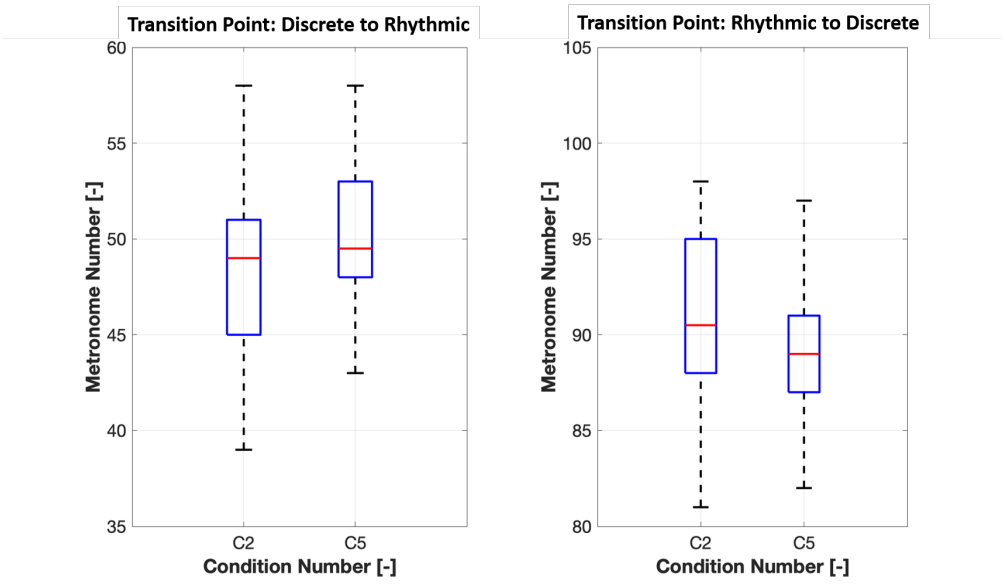

Figure 3.11: Metronome number at which transition from discrete to rhythmic (left) and rhythmic to discrete (right) occurred for all subjects in the second pair of conditions where the distance traverse was perceived visually to be the same.

but each of which had presented different visual distances to the subjects. Unexpectedly, there was no significant difference across conditions for either $DT = 0_{Accel}$ or $DT = 0_{Decel}$. Similarly, when comparing conditions where the visual distance is the same but proprioceptive distance changes (Conditions 3 vs. 4 and Conditions 2 vs. 5), the ANOVA analysis showed no effect of condition nor trial on either of the transition points. All these results are unexpected and contradicts our hypothesis that sensation matters for when the transition between discrete and rhythmic movements occur. They show that, regardless of the sensory modality providing erroneous information, subjects defaulted to a predetermined transition point.

The final measure of interest is the movement time at which subjects transition between discrete and rhythmic movements during the accelerating ($MTatDT = 0_{Acc}$) and decelerating ($MTatDT = 0_{Dec}$) portion of the trial. A summary of these values for all subjects and conditions are summarized in Figure 3.12. Here, we are interested to see if there are differences in the movement times when subjects switch their movement strategies. Specifically, for the conditions in which proprioceptive distance remained the same (Conditions 1, 4 and 5), we expect to see differences across them. However, when running a two-way ANOVA with condition and trial as factors, no significant effect is found. Similarly, when performing the same ANOVA but on the two sets of conditions with the same visual distance (Conditions 3 vs. 4 and Conditions 2 vs. 5), there are no significant effects of condition or trial.

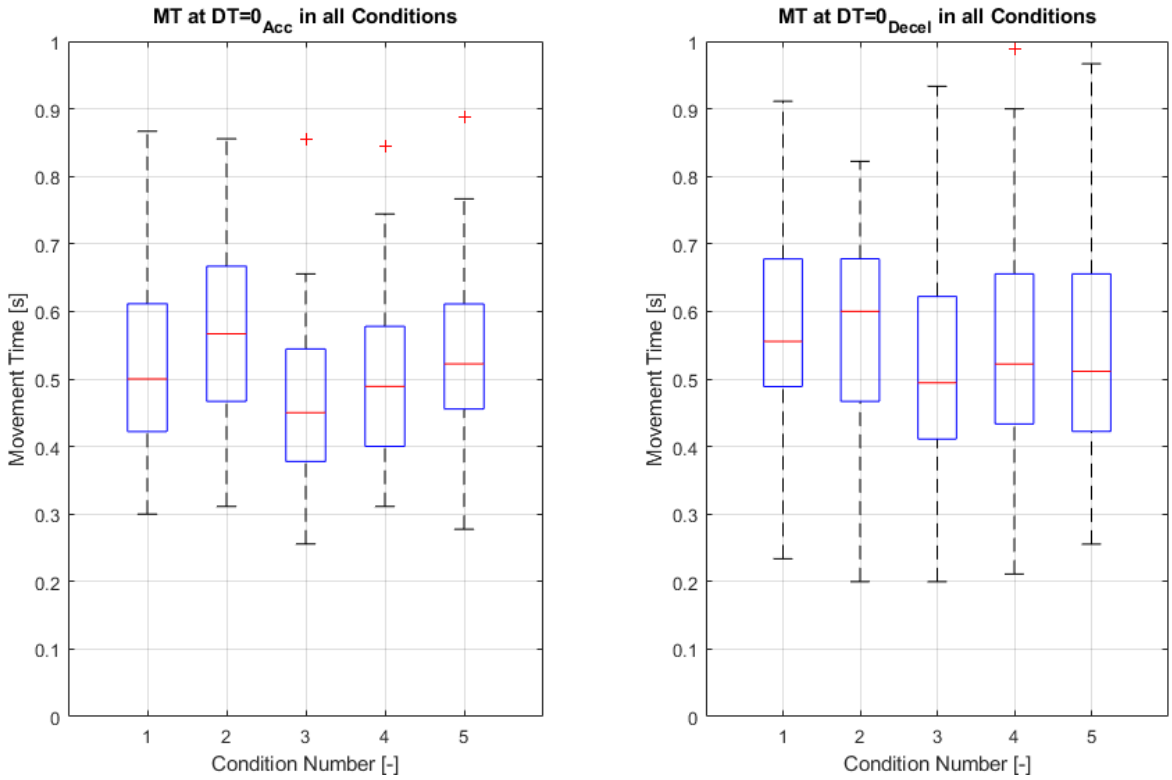

Figure 3.12: Subjects' movement time at point of transition from discrete to rhythmic (left) and rhythmic to discrete (right) in all conditions.

The lack of effect of trial strengthens the expectation that no learning took effect during the experiment and conditions were properly presented at random.

As can be seen from the results above, regardless of the sensory experience, subjects switch from one control modality to the next at around the same time. This leads us to the surprising conclusion that sensory conflict, at least between vision and proprioception in the form presented in this chapter, appears to have no effect on the transitions between discrete and rhythmic movements. It appears that such a transition is mainly influenced by internal models of the body.

3.5 Conclusion

Dynamic Primitive Theory is an influential theory of neuroscience with a vast amount of research behind it. However, this is the first time in which the role of sensation has been methodologically investigated in this domain. Using virtual reality experiments, we explored two main channels: vision and proprioception. We hypothesized that introducing errors between these two senses will influence the point of transition between discrete and rhythmic movements. To our surprise, however, we did not see any effect of sensation on the generation of these dynamic primitives. We believe that the controller in charge of generating these primitives might be relying more on internal models. Then, perhaps sensation only has an indirect effect on dynamic primitives when they influence the update of these internal models. Our assumption is that, due to the drastic introduction of these sensory interventions, there was not enough sensory exposure for these internal models to update. In the next chapter we will further develop and explore the idea of sensation and its influence on internal model adaptation.

Chapter 4

MANIPULATION OF CONTROL TO DISPLAY RATIO PRODUCES PROPRIOCEPTIVE GAIN

Object manipulation in virtual reality is challenging because object interactions are not felt, only seen. Manipulation of the gain between movements of the body and how they are rendered visually (control to display, C/D, ratio) can be used as pseudohaptic cue of weight or effort.

The brain constantly updates its internal models to solve the correspondence problem between what is felt and what is seen. "Proprioceptive Drift" occurs when the internal models incorporate a spatial offset between the true body's location and that of the surrogate body. Body ownership over an offset virtual body has been shown to result in Proprioceptive Drift. Here we demonstrate that C/D-ratio manipulation produces a *"Proprioceptive Gain"* analogous to Proprioceptive Drift, with magnitude related to C/D-ratio. We additionally observe adaptation aftereffects lasting longer in azimuth versus depth direction. This has implications for the use of pseudohaptic C/D-ratio scaling in virtual reality.

4.1 Introduction

In virtual reality, it is difficult to create convincing interactions with objects without having haptic interactive hardware, because the interactions are only seen and not "felt" by the tactile and proprioceptive sensors in the body. However, manipulating the ratio between the tracked movements of the body and their rendered display (the control-to-display ratio) has been shown to create a pseudohaptic illusion of weight [68].

Given the promise of control-to-display ratio (C/D-ratio) manipulation as a tool for interaction design, it is important to understand how it affects the user. Specifically, we are interested in internal models of the body that combine information from the different sensory modalities. Since C/D-ratio manipulation presents a gain relationship between vision and proprioception to the user, we here aim to investigate whether the brain updates its internal model to incorporate that multisensory gain. We refer to this gain relationship as *"Proprioceptive Gain"*.

The central contributions of this study are:

- Demonstrating that C/D-ratio manipulation results in *"Proprioceptive Gain"*, a recalibration of the gain between vision and proprioception for reaching tasks.

- Quantifying the degree of recalibration observed in the *"Proprioceptive Gain"*. This recalibration is mainly driven by proprioception as hypothesized from the results of *Experiment 1* and later confirmed in *Experiment 2*.

- Characterizing the time course of washout and aftereffects of the recalibrations. We observe that recalibrations in the azimuth axis take longer to wash out than those in depth.

This work demonstrates that C/D-ratio manipulation for pseudohaptic perception results in brief but measureable aftereffects. This provides a guideline for designing these interactions.

In a similar way to vanBeers et al. [69], who demonstrated the weighting of vision and proprioception varies with direction, here we explore if directionality has any effect on how much *"proprioceptive gain"* recalibration occurs. As such, we instruct participants to perform reaching tasks along depth and azimuth. We hypothesize that since vision is more reliable in azimuth than in depth, the amount of recalibration will be greater along the former.

Additionally, we analyze the importance of proprioceptive error feedback in the amount of recalibration experienced by subjects by performing two variations of the study. In the first experiment, subjects perform reaches under different C/D-ratios but with the visual distance of the targets remaining the same across all conditions. In the second experiment, the proprioceptive distance across conditions is maintained for the second experiment. We expect that the in the second experiment where the proprioceptive error feedback is higher, participants will experience a higher *"proprioceptive gain"* recalibration. Finally, since it has been shown that information from both proprioception and vision are integrated in an optimal fashion to generate an estimate of the position of our hands [70], we expect participants to recalibrate their internal models back to their "normal" conditions almost instantly when presented back with a C/D-ratio of 1.0.

4.2 Related Work

The illusion of ownership over a virtual body can be induced by rendering sensory stimuli consistent with the motion and interactions of that body [43]. Simultaneous stimulation in multiple sensory modalities, such as vision, proprioception, and haptics, take advantage of the brain's natural tendency to recalibrate the relationships among those channels [71, 72, 73]. The internal neural models of the body are updated to best explain the sensory data.

"Proprioceptive Drift" refers to a recalibration that incorporates an offset, or translation in position. It has previously been used as an indicator of the presence of the body ownership illusion over a surrogate limb. In those studies [74, 43, 75], the relationship between the visual position of the limb and the proprioceptive one has been presented as a constant offset.

Pseudohaptics [76, 77, 78, 79] is emerging as a method to simulate force feedback, but typically presents a non-constant, or gain, relationship between vision and proprioception. By manipulating the Control-to-Display ratio, the visual signals of movement (the display) can be scaled relative to the physical movements of the user (the control). C/D-ratio manipulation preserves the illusion of body ownership as long as it stays within a modest range [68].

Our study is designed to determine if an analogous update to the internal model exists in C/D-ratio manipulation as well (*"proprioceptive gain"* instead of a *"proprioceptive drift"*).

4.2.1 Bayesian Sensory Integration

Multisensory integration has been modeled as a process of Bayesian inference that performs the inference based on the reliabilities of the signals (defined as the inverse of the noise in those signals), as well as their congruence in space and time [39, 80]. Thus, these models would predict that the signal with higher reliability (i.e., lower noise) would dominate the perceptual estimate.

C/D-ratio manipulation affects sensory integration between vision and proprioception because it changes the relationship between the display (visual) and the control signal (proprioception). Therefore, we investigated the relationship between the degree of recalibration and the presented C/D-ratio. We anticipated that we would observe partial, not complete, recalibrations. This is consistent with Bayesian models of sensory integration and internal model adaptation [81, 82]; updates are made in a way that accounts for the relative reliabilities of different sensory modalities. Since the intervention is one of scaling the visual scene, we expect updates only proportional to the visual contribution to the internal neural model.

This reasoning would suggest that depth and azimuth movements could result in different amounts of recalibration, because the reliability of the visual modality is higher for azimuth (left-right) judgements than depth (front-back) ones [69, 70]. In azimuth, movement is generated mostly from shoulder rotation while movements in depth are the result of shoulder and elbow flexion-extension (one vs. two channels of information for proprioceptive error feedback). In the case of vision, error feedback relies on differences between binocular images across time. The differences between these images changes more rapidly along azimuth than in depth [83, 84].

Figure 4.1: Overview of axes and directions explored in this study and example of the targets presented to participants during the test phase.

4.3 Methods

4.3.1 Apparatus

Participants performed the task on a table measuring 60 cm by 118 cm. They sat on a height adjustable chair to ensure they all started in the same pose. A virtual environment, built with the Unity game engine [53], contained a co-located virtual replica of the real table displayed using a virtual reality headset (Oculus Rift, 2160x1200 pixels, 110 degree field-of-view, 90 Hz refresh rate). Participants held the Oculus Touch controller in their right hand throughout the task which provided real-time movement tracking at 90 Hz. An overview of this setup can be seen in Figure 4.1 (left). A custom 3-D printed adapter attached to the controller provided a flat contact surface with the table. Participants received auditory timing cues through the headset's built-in headphones.

4.3.2 Experimental Design

The experimental trials consisted of two phases: an induction phase, and a test phase. During the induction phase, participants experienced a C/D-ratio manipulation while making rhythmic movements in order to adapt to the manipulation. During the test phase, they performed single reaches to targets. Importantly, these reaches are made without any visual feedback, and so must be done using proprioception to guide the movements. Thus, errors made during these reaches are a window into the recalibration of proprioception as they

would reflect that the internal model for proprioception had changed.

This design allows us to detect the presence of *"proprioceptive gain"* recalibration. If there are no systematic errors in reaching locations, then the internal model relating proprioception to vision is unchanged after the C/D-ratio manipulation. If the reaches show errors that are constant regardless of the reaching distance, then the internal model updated in the same fashion as reported in the *"proprioceptive drift"* studies (eg. [43, 74, 75]). If the errors are scaled exactly to match the manipulation (thus exhibiting a slope proportional to the manipulation), then the internal model has completely recalibrated. If the errors are scaled in-between, then vision and proprioception are combining together in some way. The last two scenarios represent a type of recalibration of the internal model we call *"proprioceptive gain"*.

This design also allows us to detect the differences in the degree of recalibration for each axis. To this end, half of the trials are performed in depth and the other half in azimuth in a random order. We expect to see greater recalibration for the azimuth condition, due to the improved reliability of vision in that axis as discussed in the Related Work.

Finally, the study consist of two distinct experiments using the aforementioned design. In *Experiment 1*, subjects performed the experiment under conditions in which they perceived the distance traveled across all conditions to be visually the same (we refer to this version of the study as the *Visual Distance Controlled Experiment*). Additionally, subjects performed a variation of the study (the *Proprioceptive Distance Controlled Experiment*) where distances traveled in the real world are the same regardless of the condition. The goal of these two variations on the main study is to better understand the contribution of vision and proprioception to the proprioceptive gain recalibration we expect to observe.

4.3.3 Procedure

Participants were seated at the table with the chair adjusted to a comfortable height and forearms resting on the table. They were presented with a virtual scene where a "neutral" point, located 23 cm away from the edge of the table, was used to place each subject in a comparable position. Two white virtual targets (2 cm radius), located 28 cm apart and equidistant from the "neutral" point, were presented to the subject. They were oriented along either the depth (toward or away from they body) or azimuth (left or right) axes as seen in Figure 4.1.

During the induction phase participants were instructed to reach, holding the controller, back and forth between the two targets, moving in synchrony with a metronome sequence presented at 1 Hz for 30 seconds, ensuring all subjects received the same exposure across all conditions. They were instructed to place higher importance on reaching accuracy than timing. For the duration of each trial, a control-to-display ratio of either 0.7, 1.0, or 1.3,

Figure 4.2: Virtual (left) and physical (right) representation of reach during induction phase under a C/D-ratio of 0.7 along the azimuth axis.

unknown to the subject, was present along the reaching axis. These C/D-ratios have been shown to be at the edge of preserving the illusion of body ownership (see [68]), an important feature when exploring the existance of an internal model recalibration analogous to that reported in *"Proprioceptive Drift"* studies. A snapshot of the induction under a C/D-ratio of 0.7 along the azimuth axis can be seen in Figure 4.2. C/D-ratio and reaching axis were presented in a random order. This induction phase ended with subjects returning to the "neutral" point after the last metronome instance.

To measure the effects of C/D-ratio manipulation, each induction phase was followed by a testing phase. Participants were instructed to perform six discrete reaches without visual feedback. The virtual representation of the controller and the participant's hand disappeared, along with the "neutral" point and both targets. Instead, a new target of different color (yellow) and size (1.25 cm radius) was presented along the same axis the induction task took place. The target was located at three different distances from the "neutral point", in either the front-back (for the depth condition) or left-right (for the azimuth condition) direction. Each reach began with the experimenter providing a verbal cue ("Reach"). Once participants performed the reach they were asked to hold their pose, at which point the virtual target disappeared. The experimenter then moved them back to the "neutral point" (still invisible to the participant) and a new target appeared at one of the remaining six positions. Movement by the experimenter to the neutral point ensured that each reach would truly begin at the neutral point, because the participants had no visual confirmation of their position. The same process was then repeated until all six target positions were presented once for each testing phase. Figure 4.1 (right) shows all the targets presented during the test phase along both axes.

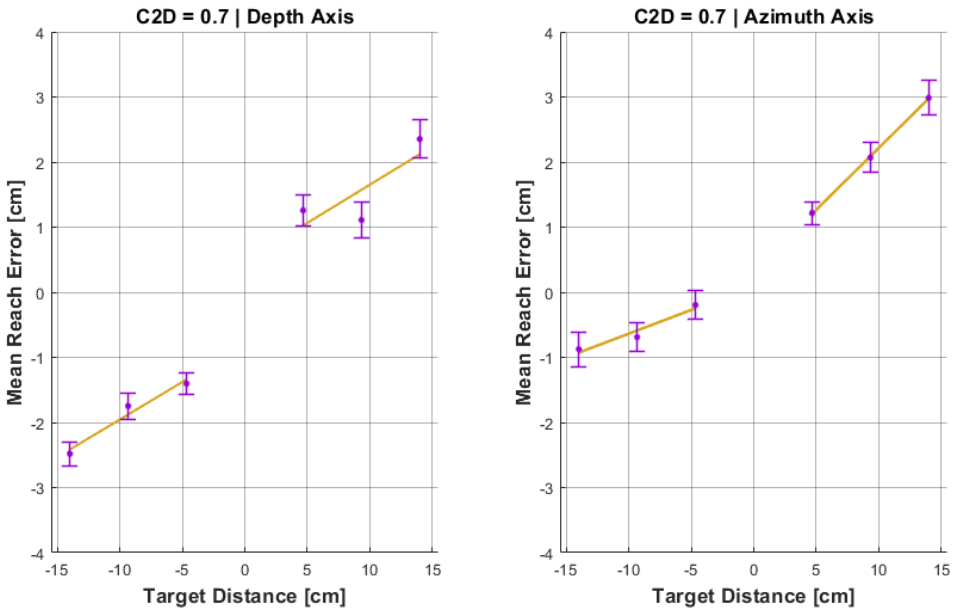

Figure 4.3: Normalized mean reaching error with standard error and resulting linear fit for C/D-ratio of 0.7 for *Experiment 1*.

In both experiments, there were 3 C/D-ratio conditions (0.7, 1.0, and 1.3) as well as the depth and azimuth axes. Each trial consisted of an induction phase using a single C/D-ratio and axis, followed by a testing phase wherein each of the six targets was presented in a random order. Trials for each combination of conditions was presented 10 times in random order. Each experiment lasted for about 65 min, with subjects allowed to take breaks at any point during the experiment. Most subjects took one break roughly halfway through the experiment.

4.4 Experiment 1: Visual Distance Controlled

4.4.1 Participants

A total of 14 participants (age: $\{M = 28, SD = 6.7\}$, 6 females, all right handed) participated in this study. None had any health conditions that could impact their motor performance and they all gave informed written consent before participating. Three of the 14 chose to terminate the experiment before completion, due to fatigue or boredom. The experimental protocols were approved by the University of Washington Institutional Review Board.

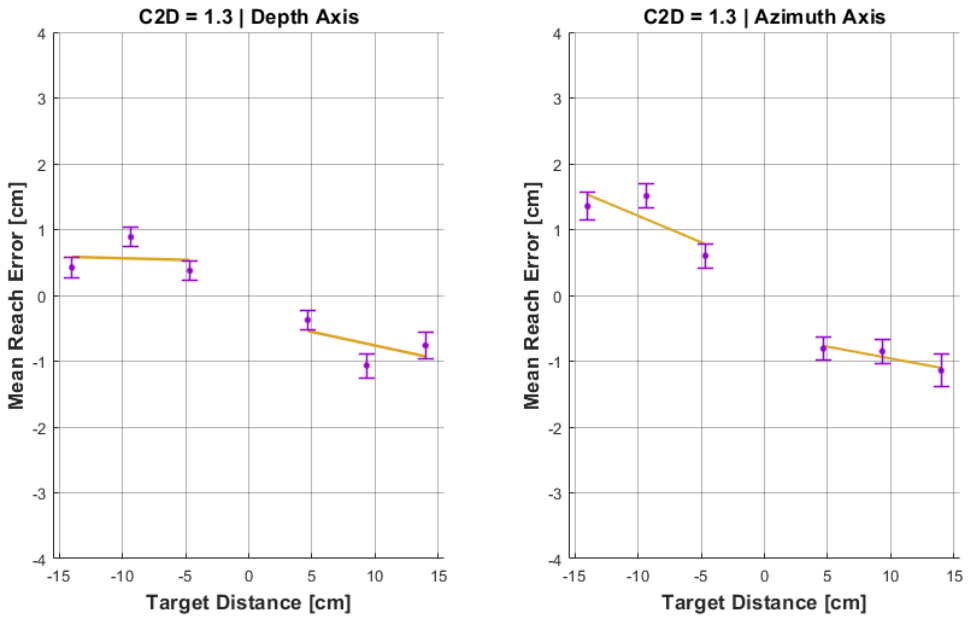

Figure 4.4: Normalized mean reaching error with standard error and resulting linear fit for C/D-ratio of 1.3 for *Experiment 1.*

4.4.2 Target Distances During Testing Phase

For the testing phase, targets were located 4.67, 9.33 or 14 cm away from the "neutral" point, equivalent to 1/3, 2/3, or equal to the target position during the induction phase, in either the front-back (for the depth condition) or left-right (for the azimuth condition) direction. The choice of total test targets to be explored was selected so that a high enough equally spaced points were tested while minimizing the experiment time and, consequently, subjects' fatigue level.

4.4.3 Analysis and Results

The objective of this study was to analyze how reaching errors relate to C/D-ratio manipulation. Reaching error is defined as the difference between the location that was reached to and the target location during the test phase. Our central hypothesis was that C/D-ratio manipulation would result in increased reaching errors for more distant targets. We anticipated that this effect would scale with the C/D-ratio. Specifically, we expected overshooting when exposed to a C/D-ratio of 0.7 and for this overshooting to increase for more distant targets. This is because in the 0.7 condition, the visual depiction of movement is scaled-

down compared to the physical movement. We expect the opposite direction of change for the C/D-ratio of 1.3.

We hypothesized that perceived errors would result in cross-sensory recalibration. If this recalibration was complete this would result in reaching errors that are fully predicted by the C/D-ratio. In the C/D-ratio of 0.7, this would translate to a predicted overshoot of 4.2 cm for reaches to the 14 cm target, and an undershoot of the same amount in the 1.3 condition. The predicted slopes between reaching errors and target distances would be 0.3 for the C/D-ratio of 0.7 and -0.3 for the C/D-ratio of 1.3 (dotted line in Figure 4.5). The observed slopes indicate the degree of recalibration. If slopes were flat, it would indicate no recalibration and if they match the prediction it would indicate full recalibration.

Each of the 6 reaches following the induction period constituted a single measurement of error. This resulted in a total of 110 data points per target and condition (11 subjects, 10 repetitions for each of the 6 targets). Baseline reaching error was measured for each subject by calculating the mean reaching error for each target in the C/D-ratio of 1.0 trials. This mean error accounts for individual differences in accuracy for each target, independent of the C/D-ratio manipulation. Reaching errors for the two experimental conditions (C/D-ratio of 0.7 and 1.3) were normalized by subtracting the baseline mean reaching error.

Linear regression analysis was performed using MATLAB R2018b, statistics and machine learning toolbox. Figures 4.3 and 4.4 depict the fits from this regression for each reaching direction in each axis.

Given the slope outcome measures, a 3-way repeated measures analysis of variance (ANOVA) was performed. C/D-ratio (0.7 or 1.3), axis (depth or azimuth), and direction (distal/proximal in depth or left/right in azimuth) were the factors. The dependent variable was the slope of the reaching errors. This analysis revealed a significant main effect of C/D-ratio ($F(1, 10) = 42.36$, $p < .0001$). Collapsing across axis and directions and running a 1-way ANOVA with C/D-ratio as the factor returned the same result. Figure 4.5 depicts the normalized reaching errors, independent of axis and direction, across all subjects.

Based on the ANOVA, regressions were performed to obtain summary slopes to compare to the "full recalibration" slopes, shown in Figure 4.5. These linear regression fits combined axis and direction, for each of the 0.7 and 1.3 C/D-ratio conditions. The least-squares linear regression slope for 0.7 and 1.3 were 0.1238 ($F(2, 1318) = 47.4$, $p < .0001$) and -0.0407 ($F(2, 1318) = 8.42$, $p = .0038$) respectively. The positive slope in the 0.7 condition indicates that participants showed an increasing trend of overshooting as the target distance increased, which aligns with our central hypothesis that C/D-ratio results in a *"proprioceptive gain"*. This is also true for the C/D-ratio of 1.3 where the negative slope reflected participants' tendency to increasingly undershoot targets as distance increased.

The observed slopes of 0.1238 for C/D-ratio 0.7, and -0.047 for C/D-ratio 1.3, correspond

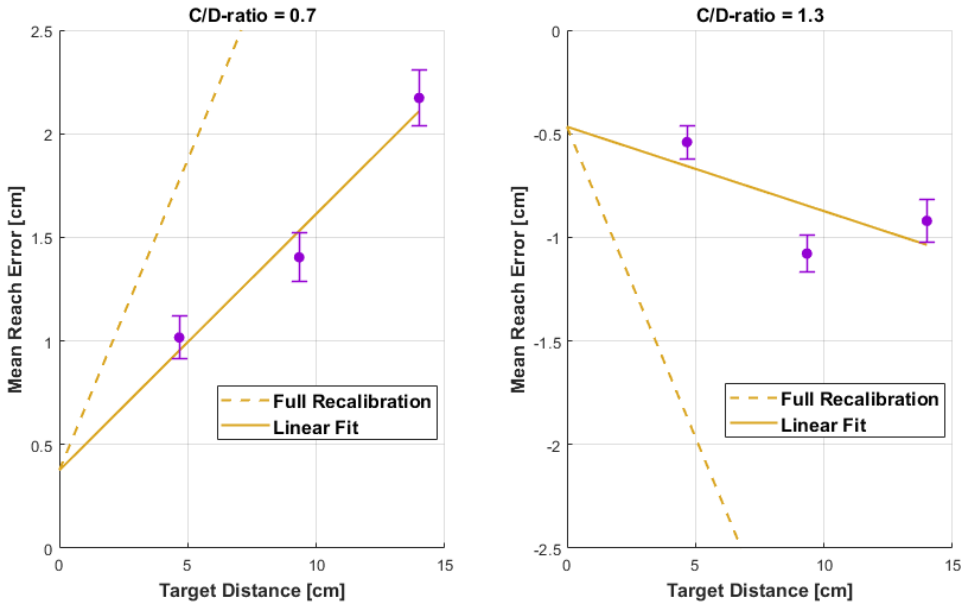

Figure 4.5: Linear regression using combined axis and direction based on ANOVA finding (only C/D-ratio was found to have a significant effect). Normalized mean reaching error with standard error and resulting linear fit for C/D-ratio of 0.7 (left) and 1.3 (right).

to 41.267% and 13.567% recalibrations relative to the "full recalibration" predictions.

Finally, we analyzed the timing for washout after the manipulation. To do this, we use the reaching error during the induction phases following experimental conditions. Trials were included in this analysis if they presented a C/D-ratio of 1.0 and were preceded by a C/D-ratio of 0.7 or 1.3 along the same axis (depth or azimuth). Mean reaching errors over time during the induction phases are presented in Figure 4.6.

4.4.4 Discussion

We observed a *"proprioceptive gain"*, a recalibration of the gain between vision and proprioception, after C/D-ratio manipulation. Unexpectedly, ANOVA analysis indicated no significant effect of axis (depth or azimuth) and direction (toward or away), only C/D-ratio condition.

We expected that, due to the difference in the reliability of vision for depth and azimuth, there would be more recalibration in the latter. This stems from the biomechanical and kinematic differences between these kinds of movement. For example, movement in azimuth

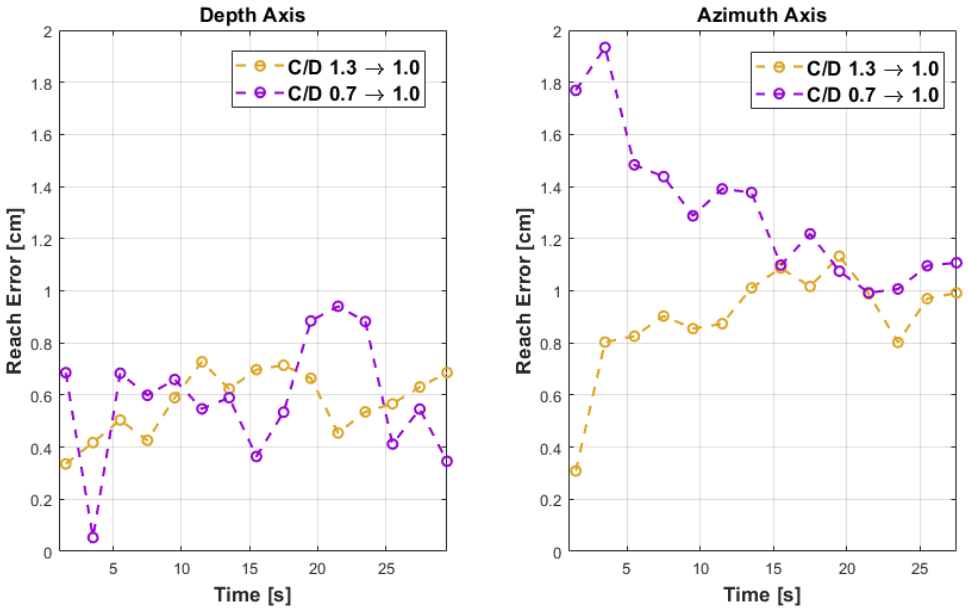

Figure 4.6: Mean reaching error during the induction phase for all C/D-ratio 1.0 trials preceded by either a C/D-ratio of 0.7 or 1.3 in *Experiment 1*. There appears to be a difference in washout timing for manipulations in the depth and azimuth axes.

can be mostly achieved using only shoulder rotation, but movement in depth requires both shoulder and elbow flexion-extension [69, 70]. This would have been reflected in our linear regression analysis as significantly steeper slopes compared to the depth axis, indicating a higher degree of recalibration. We did observe a numerical difference between these slopes, but the variance was too high for it to reach statistical significance. This suggests that future studies with higher statistical power may find a difference between these slopes, albeit with a small effect size.

We conjectured that there could also be an effect for direction, due to various contributions such as proximity to the body midline, potential effects of the edge of the table being felt on the forearm, the orientation of the wood grain in the virtual table depiction, etc. However, these appear not to have strongly affected the outcome, based on the lack of significance in the ANOVA analysis.

Degree of Recalibration

If participants updated their internal models to completely compensate for the visual component of the C/D-ratio manipulation, then their reaches during the test phase would terminate at locations predicted only by the C/D-ratio. For C/D-ratio 1.3, this corresponds to a slope

of -0.3, while 0.7 corresponds to a slope of +0.3. The observed recalibrations are smaller in magnitude than these predictions, corresponding to 41.267% for C/D-ratio 0.7, and 13.567% for C/D-ratio 1.3.

One way to interpret the disparity in degree of recalibration for the C/D-ratio conditions is to note that different amounts of physical movement were required under the two conditions. Although the targets appear visually the same distance away in both conditions, the physical movements required to reach the targets is scaled by the C/D ratio. Therefore in the 1.3 condition, there is substantially less physical movement being performed. There is evidence that movements have an important role in creating and updating internal models of the body, e.g. [85]. Adaptation is likely driven by perceived errors or inconsistencies across the sensory modalities [86, 87, 88]. It is possible, then, that the degree of recalibration is gated by the amount of movement of the limb in question. Increased movement yields richer error feedback, and increases the engagement of the motor generation system. This suggests a follow-up study analogous to this protocol, but in which the physical movements are identical in size while the distance between the visual targets differ. The next section (see *Experiment 2*) explores this hypothesis further.

Another interpretation is based on a Bayes Filter model for sensorimotor integration in the internal model [81, 89, 90]. In this framework, the evidence (experience of C/D-ratio) is only partially incorporated into the model, integrated with the prior (the existing internal model estimate, tuned during normal behavior). The update accounts for the noise in the sensory estimate and the uncertainty of the internal model. Differences in the apparent reliability of the visual and proprioceptive sensory channels under the two conditions could result in differences in the magnitude of the update.

Washout

We observed an interesting phenomenon in the reaching error during the induction phases following experimental conditions (Figure 4.6). As expected, we see that reaching errors decrease over the course of the induction phase, as the gain adaptation from the previous block is washed out to the 1.0 C/D-ratio. Note that the errors do not converge exactly to zero, but the targets themselves had a radius of 2 cm, so errors of this magnitude were likely not interpreted by the participants as error.

Interestingly, the time course of the adaptations is markedly different for the depth and azimuth axes. In depth, the errors immediately resemble the fully washed-out errors near the end of the induction phase; it appears that participants nearly instantaneously switch back to their natural 1.0 movements. In azimuth, however, the time course of washout is slower, taking some 10 to 15 seconds.

This result could arise from differences in the reliability of vision for the two axes. C/D-

ratio manipulation can be thought of as a visual scaling, and vision provides more information about movements in azimuth than it does in depth. The persistence of adaptation to azimuth movements could be greater because those signals are more reliable.

Finally, the timing of washout is important for designing interactions using pseudohaptic weight. If different objects are meant to be interacted with in quick succession, it could matter how they are manipulated. Internal model updates induced by interactions in depth appear to quickly fade, while those in azimuth persevere slightly longer. Even these, however, fade in a few (10-15) seconds.

4.5 Experiment 2: Proprioceptive Distance Controlled

4.5.1 Participants

A total of 11 participants (age: $\{M = 27, SD = 4.5\}$, 4 females, all right handed) participated in this study. None had any health conditions that could impact their motor performance and they all gave informed written consent before participating. Two subjects opted out during the experiment due to nausea. The experimental protocols were approved by the University of Washington Institutional Review Board.

4.5.2 Target Distances

Target distances for this experiment were dependent on the C/D-ratio manipulation presented to the subject. For the C/D-ratio condition of 0.7, targets were located 18.6 cm apart, in the 1.0 condition they were 28 cm apart, and in the 1.3 condition they were 36.4 cm apart. In doing so, the proprioceptive distance across all conditions remained the same (28 cm).

For the testing phase, targets were located 1/3, 2/3, or equal to the target position during the induction phase, in either the front-back (for the depth condition) or left-right (for the azimuth condition) direction. As in *Experiment 1*, this ensured data from equally spaced targets were used when analyzing the recalibration of the internal model.

4.5.3 Analysis and Results

As in the previous experiment described in Section 4.4, the objective in this study was to analyse how reaching errors relate to C/D-ratio manipulation. In this experiment, however, we also wanted to explore our hypothesis that the amount of recalibration generated by the C/D-ratio manipulation is related to the amount of physical movement.

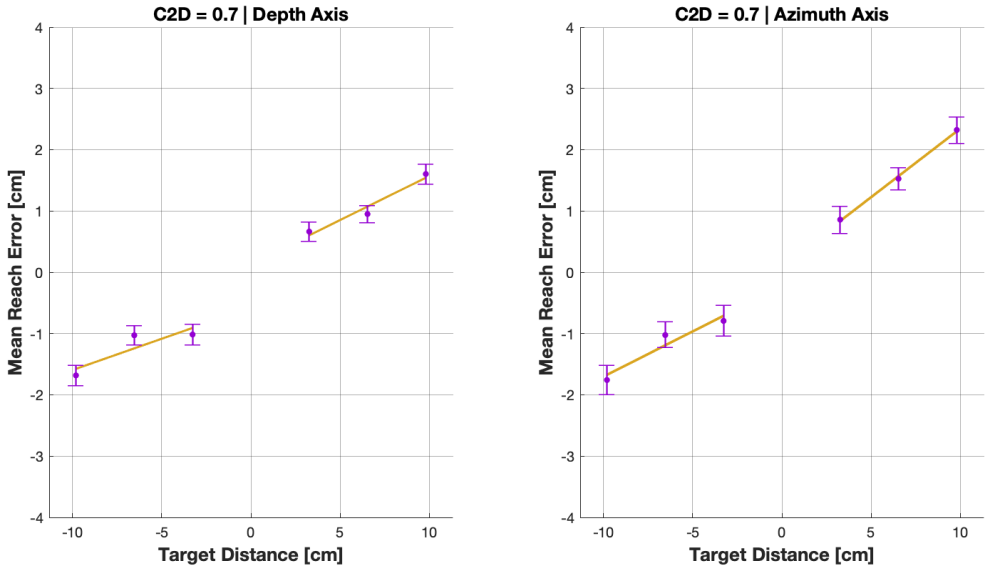

Figure 4.7: Normalized mean reaching error with standard error and resulting linear fit for C/D-ratio of 0.7 in *Experiment 2*.

We hypothesized we will still see cross-sensory recalibration resulting from the perceived errors. By having subjects experience the same proprioceptive target distances in all experimental conditions, we expected to see an increased in recalibration under a C/D-ratio of 1.3 compared to that seen in *Experiment 1* (Section 4.4).

As before, baseline reaching error was measured for each subject by calculating the mean reaching error for each target in the trials under C/D-ratio of 1.0. This was used to normalize the reaching errors for the other two experimental conditions (C/D-ratio of 0.7 and 1.3) so that to account for individual subject differences in accuracy for each target.

Linear regression analysis was performed using the statistics and machine learning toolbox in MATLAB R2018b with the fits of each regression for each reaching direction in each axis shown in Figures 4.7 and 4.8.

A three-way repeated measures analysis of variance (ANOVA) was performed with the slope of the reaching errors as dependent variable and C/D-ratio (0.7 or 1.3), axis (depth or azimuth), and direction (toward/away in depth or left/right in azimuth) as factors. This analysis revealed a significant main effect of C/D-ratio ($F(1, 9) = 69.61$, $p < .0001$). Collapsing across axis and directions and performing a one-way ANOVA with C/D-ratio as the factor return the same conclusion. The normalized reaching errors across all subjects, independent of axis and direction, can be seen in Figure 4.9.

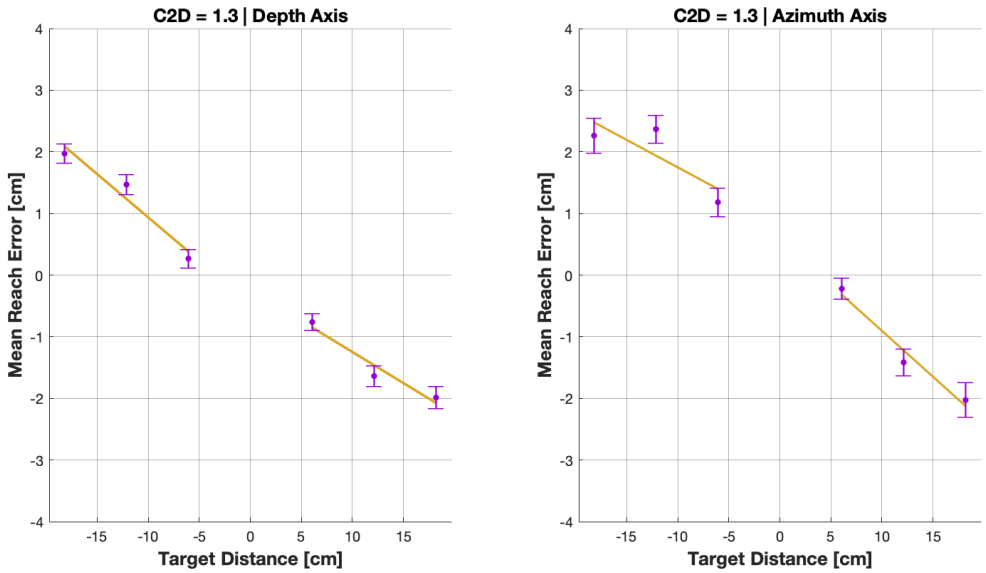

Figure 4.8: Normalized mean reaching error with standard error and resulting linear fit for C/D-ratio of 1.3 in *Experiment 2*.

Following the conclusions provided by the ANOVA, linear regression was performed with axis and directions combined for each of the 0.7 and 1.3 C/D-ratio conditions. The resulting linear fit as well as the predicted "full recalibration" slope can both be seen in Figure 4.9. The least-squares linear regression slope was 0.1544 ($F(2, 1318) = 54.1$, $p < 0.0001$) for 0.7 and -0.1198 ($F(2, 1318) = 101$, $p < .0001$) for 1.3. These results once again align with our central hypothesis that C/D-ratio results in a *"proprioceptive gain"*, with subjects increasingly overshooting as distances increased for the 0.7 condition and undershooting for the 1.3 condition.

The observed slopes of 0.1544 for C/D-ratio 0.7, and -0.1198 for C/D-ratio 1.3, corresponded to 51.47% and 39.93% recalibrations relative to the "full recalibration" predictions.

The timing for washout after manipulation was also analyzed by using the reaching error during the induction phases following experimental conditions. Trials considered for this analysis were those with a C/D-ratio of 1.0 and were preceded by either C/D-ratio of 0.7 or 1.3 along the same axis (depth or azimuth). Figure 4.10 shows the mean reaching error during the induction phase for the trials that satisfy the aforementioned condition.

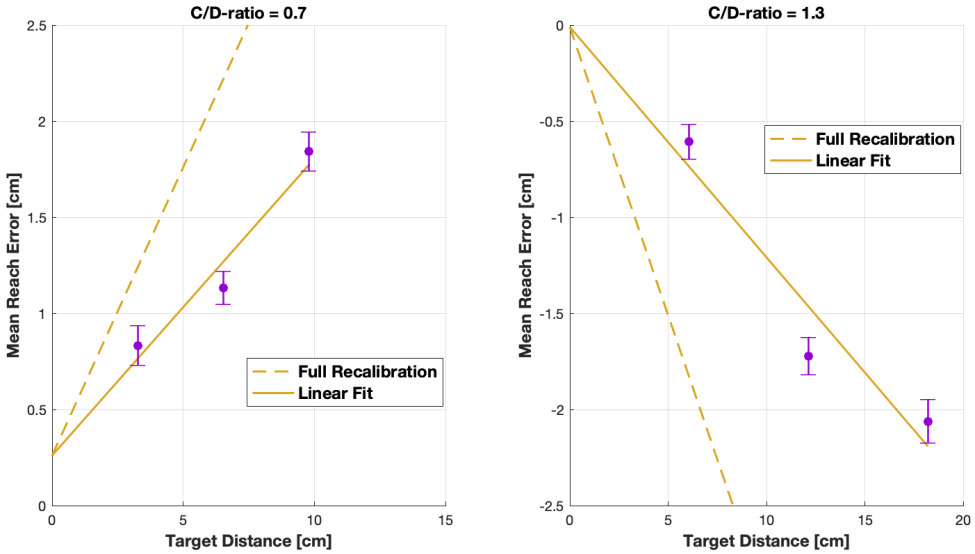

Figure 4.9: Linear regression using combined axis and direction based on ANOVA finding (only C/D-ratio was found to have a significant effect). Normalized mean reaching error with standard error and resulting linear fit for C/D-ratio of 0.7 (left) and 1.3 (right).

4.5.4 Discussion

In this experiment we once again saw a recalibration of the gain between vision and proprioception (a *"proprioceptive gain"*) following a C/D-ratio manipulation. However, we saw a similar lack of effect of axis (depth or azimuth) and direction (toward or away) as in *Experiment 1* (Section 4.4) where the visual distance was controlled across all conditions.

Although we expected a difference in the amount of recalibration between the explored axis conditions, this did not show up in our analysis. The linear regression slopes for both axes were not significantly different from each other.

Directionality of the task (toward or away) did not show an effect either as shown by the lack of significance in the ANOVA analysis. We expected that due to factors such as the table edge, orientation of the wood grain on the virtual table or proximity to the body midline, direction could have an effect on the *"proprioceptive gain"* experience by subjects.

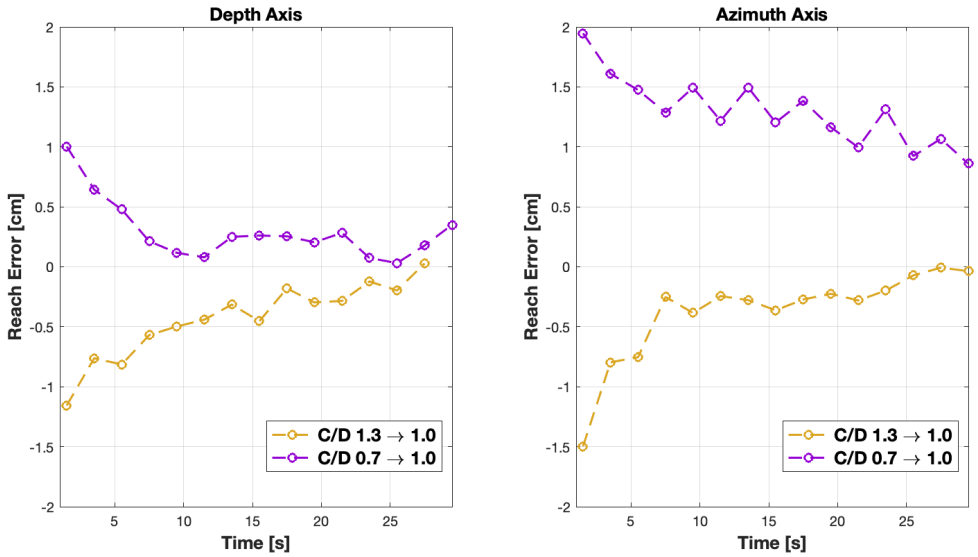

Figure 4.10: Mean reaching error during the induction phase for all C/D-ratio 1.0 trials preceded by either a C/D-ratio of 0.7 or 1.3 in *Experiment 2*. There appears to be a difference in washout timing for manipulations in the azimuth axis.

Degree of Recalibration

If subjects fully updated their internal models after receiving a C/D-ratio manipulation, we would have expected that their reaches during the test phase would be fully predicted by the C/D-ratio (shown as a dotted line in Figure 4.9). However, the resulting recalibration were smaller than those predictions: 51.47% and 39.93% for C/D-ratio of 0.7 and 1.3 respectively (solid golden lines in Figure 4.9).

Although we still observed a lower amount of recalibration for the 1.3 C/D-ratio condition compared to the 0.7 condition, the overall amount of recalibration in both conditions were higher than those observed in *Experiment 1* (Section 4.4). While the 0.7 C/D-ratio condition showed an increase in recalibration of approximately 24% (from 41.267% to 51.47%), the 1.3 C/D-ratio condition's recalibration showed close to a 300% increase (from 13.567% to 39.93%). This clearly validates our hypothesis that the small recalibration seen in the 1.3 C/D-ratio condition in *Experiment 1* was due to the reduced physical movement in that condition. In other words, increased movement resulted in a higher degree of recalibration as seen in this experiment.

Washout

Figure 4.10 shows the reaching error during the induction phases in a 1.0 C/D-ratio condition following either a 0.7 or 1.3 C/D-ratio manipulation. As expected, as the gain adaptation from the previous condition is washed out to the 1.0 C/D-ratio condition, the observed reaching errors decrease. This is true for both axes (depth and azimuth).

Although we see reaching errors in all conditions decrease over the course of the induction phase, all but one seems to completely washout. In the depth axis, the reaching error starts with either an overshoot (for those preceded by C/D-ratio 0.7) or undershoot (for those preceded by C/D-ratio 1.3) as expected and they completely disappear by the end of the induction phase. The azimuth axis shows a different picture: while both conditions start with either an undershoot or overshoot as expected, only the one preceded by a C/D-ratio of 1.3 dies out completely by the end of the induction phase. The reaching error for the conditions preceded by a C/D-ratio of 0.7 has a much slower washout, with subjects' still overshooting the targets by 1 cm towards the end of the induction phase. These observations could once again be partially explained by the difference in reliability of vision for the two axes.

Reaching errors also show a different behavior in both versions of the experiment. Although in both experiments subjects reaching errors decrease in all conditions, the type of errors are different in both experiments. In *Experiment 1*, as seen in Figures 4.6, subjects reaching errors show a tendency to overshoot targets when previously exposed to either a 0.7 or 1.3 C/D-ratio conditions in either depth or azimuth axis. However, in *Experiment 2*, subjects overshot targets only when coming from a C/D-ratio of 0.7 (see Figure 4.10 (left)). For those preceded by a 1.3 C/D-ratio condition, subjects started the trial by undershooting the targets (Figure 4.10 (right)).

From a timing perspective, we also saw differences in washout between the two variations of the study explored in this chapter. In *Experiment 1* (Section 4.4), internal model updates persisted for about 10-15 seconds along the azimuth axis and almost instantly disappeared in the depth axis. During this experiment we observed that in all conditions in the depth axis and one condition (C/D-ratio 1.3 $\rightarrow$ 1.0) in azimuth, the effects disappear in about 20-25 seconds. In the last condition (C/D-ratio 0.7 $\rightarrow$ 1.0 in azimuth), we see the effect not fully disappearing by the end of the induction phase.

4.6 Conclusion

Manipulating the C/D-ratio for pseudohaptics is a promising means for inducing illusory perceptions of force in virtual interactions, but it causes cross-sensory recalibration after-effects. Interacting with a manipulated C/D-ratio causes a *proprioceptive gain* in reaching

tasks with no visual feedback. This is analogous to the well-known *proprioceptive drift* induced by interactions that include visual-proprioceptive offset. We see a stronger effect in the condition that required more movement of the body, and increased duration of the aftereffect for movement along the azimuth axis.

4.6.1 Limitations and Future Work

In conditions where the task was performed along the depth axis, subjects were looking down at the virtual table at an angle. This means that the information available to the visual system was not purely depth. Thus, a potential difference in the degree of recalibration for depth and azimuth axes should not be completely ruled out. Performing a follow-up study with the depth axis positioned at eye-level could provide a more conclusive result.

There have been numerous other studies of how vision and other sensory modalities contribute to adaptation, dating from early examinations of prism adaptation, e.g. [91]. More recently, other relationships beyond translation have been explored, from rotation [92] to complex nonlinear relationships [93]. Our work focuses on pseudo-haptic applications of C/D-ratio manipulation, and we measure a behavioral *"Proprioceptive Gain"* during voluntary movements. However, this does not definitively address the mechanisms that generate these movements—i.e., we cannot rule out the possibility that the adaptation is occurring in motor planning, rather than cross-sensory recalibration [94, 95].

This study focuses on interactions along a 2-D plane, but it would be interesting to see how these results extend to richer interactions. There could be important effects of C/D-ratio manipulation, which we have defined here simply using the global Cartesian coordinate frame, defined instead in joint coordinates or other representations [96].

BIBLIOGRAPHY

[1] K. H. Sheetz, J. Claflin, and J. B. Dimick, "Trends in the Adoption of Robotic Surgery for Common Surgical Procedures," *JAMA network open*, vol. 3, no. 1, p. e1918911, 2020.

[2] 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*, vol. 4, no. 32, pp. 1–12, 2019.

[3] J. J. S. Hebert, J. J. L. Olson, M. M. J. Morhart, M. Dawson, P. Marasco, T. A. Kuiken, and K. Chan, "Novel Targeted Sensory Reinnervation Technique to Restore Functional Hand Sensation After Transhumeral Amputation.," *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. PP, no. 99, p. 1, 2013.

[4] T. A. Kung, N. B. Langhals, D. C. Martin, P. J. Johnson, P. S. Cederna, and M. G. Urbanchek, "Regenerative peripheral nerve interface viability and signal transduction with an implanted electrode," *Plastic and Reconstructive Surgery*, vol. 133, no. 6, pp. 1380–1394, 2014.

[5] B. T. Nghiem, I. C. Sando, R. B. Gillespie, B. L. McLaughlin, G. J. Gerling, N. B. Langhals, M. G. Urbanchek, and P. S. Cederna, "Providing a Sense of Touch to Prosthetic Hands," *Plastic and Reconstructive Surgery*, vol. 135, no. 6, pp. 1652–1663, 2015.

[6] M. J. H. Lum, D. C. W. Friedman, G. Sankaranarayanan, H. King, K. Fodero, R. Leuschke, B. Hannaford, J. Rosen, and M. N. Sinanan, "The RAVEN: Design and validation of a telesurgery system," *International Journal of Robotics Research*, vol. 28, no. 9, pp. 1183–1197, 2009.

[7] B. Hannaford, J. Rosen, D. W. Friedman, H. King, P. Roan, L. Cheng, D. Glozman, J. Ma, S. N. Kosari, and L. White, "Raven-II: An open platform for surgical robotics research," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 4, pp. 954–959, 2013.

[8] M. Miyasaka, R. Gosalia, D. E. Caballero, M. Haghighipanah, and B. Hannaford, "Bending over Sheave Fatigue of Small Scale Wire Rope Used for Cable Driven Surgical Robot," *Robotics and Automation*, 2019.

[9] Y. Li, S. Li, D. E. Caballero, M. Miyasaka, A. Lewis, and B. Hannaford, "Improving control precision and motion adaptiveness for surgical robot with recurrent neural network," *IEEE International Conference on Intelligent Robots and Systems*, pp. 3538–3543, 2017.

[10] E. Rosen, D. Whitney, E. Phillips, G. Chien, J. Tompkin, G. Konidaris, and S. Tellex, "Communicating Robot Arm Motion Intent Through Mixed Reality Head-mounted Displays," pp. 1–16, 2017.

[11] D. E. Caballero and E. Rombokas, "Sensitivity to Conflict Between Visual Touch and Tactile Touch," *IEEE Transactions on Haptics*, vol. 12, no. 1, pp. 78–86, 2019.

[12] N. Hogan and D. Sternad, "Dynamic Primitives of Motor Behavior," *Biological Cybernetics*, vol. 106, no. 11, pp. 727–739, 2012.

[13] N. Hogan and D. Sternad, "Dynamic primitives in the control of locomotion," *Frontiers in Computational Neuroscience*, vol. 7, no. 71, pp. 1–16, 2013.

[14] D. Sternad, H. Marino, S. K. Charles, M. Duarte, L. Dipietro, and N. Hogan, "Transitions between discrete and rhythmic primitives in a unimanual task," *Frontiers in Computational Neuroscience*, vol. 7, 2013.

[15] D. Goldreich and J. Tong, "Prediction, postdiction, and perceptual length contraction: A bayesian low-speed prior captures the cutaneous rabbit and related illusions," *Frontiers in Psychology*, vol. 4, pp. 1–26, 2013.

[16] J. Hartcher-O'Brien, A. Gallace, B. Krings, C. Koppen, and C. Spence, "When vision 'extinguishes' touch in neurologically-normal people: Extending the Colavita visual dominance effect," *Experimental Brain Research*, vol. 186, no. 4, pp. 643–658, 2008.

[17] F. B. Colavita, "Human sensory dominance," *Perception & Psychophysics*, vol. 16, no. 2, pp. 409–412, 1974.

[18] G. J. Thomas, "Experimental study of the influence of vision on sound localization.," *Journal of Experimental Psychology*, vol. 28, no. 2, pp. 163–177, 1941.

[19] M. Radeau and P. Bertelson, "Auditory-Visual Interaction and the Timing Of Inputs," *Psychological Research*, pp. 317–329, 1987.

[20] H. L. Pick, D. H. Warren, and J. C. Hay, "Sensory conflict in judgments of spatial direction," *Perception & Psychophysics*, vol. 6, no. 4, pp. 203–205, 1969.

[21] J. Hay, H. Pick, and K. Ikeda, "Visual capture produced by prism spectacles," *Psychonomic science*, vol. 2, no. 8, pp. 215–216, 1965.

[22] C. Wheatstone, "Contributions to the Physiology of Vision. Part the First. On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision," *Philosophical Transactions of the Royal Society of London*, vol. 128, no. 0, pp. 371–394, 1838.

[23] C. L. E. Paffen and D. Alais, "Attentional Modulation of Binocular Rivalry," *Frontiers in Human Neuroscience*, vol. 5, no. September, pp. 1–10, 2011.

[24] B. Milner, L. Taylor, and R. W. Sperry, "Lateralized suppression of dichotically presented digits after commissural section in man.," *Science*, vol. 161, no. 3837, pp. 184–186, 1968.

[25] A. Brancucci and L. Tommasi, ""Binaural rivalry": Dichotic listening as a tool for the investigation of the neural correlate of consciousness," *Brain and Cognition*, vol. 76, no. 2, pp. 218–224, 2011.

[26] F. De Vignemont, H. H. Ehrsson, and P. Haggard, "Bodily illusions modulate tactile perception," *Current Biology*, vol. 15, no. 14, pp. 1286–1290, 2005.

[27] A. Petroni, M. J. Carbajal, and M. Sigman, "Proprioceptive Body Illusions Modulate the Visual Perception of Reaching Distance," *Plos One*, vol. 10, no. 6, 2015.

[28] B. R. Brewer, M. Fagan, R. L. Klatzky, and Y. Matsuoka, "Perceptual limits for a robotic rehabilitation environment using visual feedback distortion," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 13, no. 1, pp. 1–11, 2005.

[29] B. R. Brewer, R. Klatzky, and Y. Matsuoka, "Visual feedback distortion in a robotic environment for hand rehabilitation," *Brain Research Bulletin*, vol. 75, no. 6, pp. 804–813, 2008.

[30] R. A. Scheidt, M. A. Conditt, E. L. Secco, and F. A. Mussa-Ivaldi, "Interaction of Visual and Proprioceptive Feedback During Adaptation of Human Reaching Movements," *J Neurophysiol*, vol. 93, pp. 3200–3213, 2005.

[31] S.-J. Kim and H. I. Krebs, "Effects of implicit visual feedback distortion on human gait.," *Experimental brain research*, vol. 218, pp. 495–502, 5 2012.

[32] S.-J. Kim and D. Mugisha, "Effect of explicit visual feedback distortion on human gait.," *Journal of NeuroEngineering and Rehabilitation*, vol. 11, no. 1, p. 74, 2014.

[33] P. Bertelson and M. Radeau, "Cross-modal bias and perceptual fusion with auditory-visual spatial discordance," *Perception & Psychophysics*, vol. 29, no. 6, pp. 578–584, 1981.

[34] S. Mateeff, J. Hohnsbein, and T. Noack, "Dynamic visual capture: apparent auditory motion induced by a moving visual target.," *Perception*, vol. 14, no. 6, pp. 721–727, 1985.

[35] D. Alais and D. Burr, "Ventriloquist Effect Results from Near-Optimal Bimodal Integration," *Current Biology*, vol. 14, no. 3, pp. 257–262, 2004.

[36] C. Koppen and C. Spence, "Audiovisual asynchrony modulates the Colavita visual dominance effect," *Brain Research*, vol. 1186, no. 1, pp. 224–232, 2007.

[37] R. Irvan and V. Jack, "Vision and Touch : An Experimentally Created Conflict between the Two Senses," *Science*, vol. 143, no. 3606, pp. 594–596, 1964.

[38] M. A. Heller, "Visual and tactual texture perception: Intersensory cooperation," *Perception & Psychophysics*, vol. 31, no. 4, pp. 339–344, 1982.

[39] M. O. Ernst and M. S. Banks, "Humans integrate visual and haptic information in a statistically optimal fashion," *Nature*, vol. 415, no. 6870, pp. 429–433, 2002.

[40] S. Gepshtein and M. S. Banks, "Viewing Geometry Determines How Vision and Haptics Combine in Size Perception," *Current Biology*, vol. 13, pp. 483–488, 2003.

[41] M. A. Riley, S. Wong, S. Mitra, and M. T. Turvey, "Common effects of touch and vision on postural parameters," *Experimental Brain Research*, vol. 117, pp. 165–170, 10 1997.

[42] K. J. McKenzie, E. Poliakoff, R. J. Brown, and D. M. Lloyd, "Now you feel it, now you don't: How robust is the phenomenon of illusory tactile experience?," *Perception*, vol. 39, no. 6, pp. 839–850, 2010.

[43] M. Botvinick and J. Cohen, "Rubber hands' feel'touch that eyes see," *Nature*, 1998.

[44] T. Asai and N. Kanayama, ""Cutaneous rabbit" hops toward a light: Unimodal and cross-modal causality on the skin," *Frontiers in Psychology*, vol. 3, no. OCT, pp. 1–12, 2012.

[45] M. F. Nolan, "Limits of two-point discrimination ability in the lower limb in young adult men and women.," *Physical Therapy*, vol. 63, no. 9, pp. 1424–1428, 1983.

[46] E. S. Dellon, R. Mourey, and a. L. Dellon, "Human pressure perception values for constant and moving one- and two-point discrimination.," 1992.

[47] R. Periyasamy, M. Manivannan, V. Balakrish, and R. Narayanamurthy, "Changes in Two Point Discrimination and the law of mobility in Diabetes Mellitus patients," *Journal of Brachial Plexus and Peripheral Nerve Injury*, vol. 3, no. 1, pp. 23–40, 2008.

[48] C. Jerosch-Herold, "Assessment of sensibility after nerve injury and repair: A systematic review of evidence for validity, reliability and responsiveness of tests," *Journal of Hand Surgery*, vol. 30, no. 3, pp. 252–264, 2005.

[49] G. Lundborg and B. Rosen, "The Two-Point Discrimination Test - Time For A Re-Appraisal?," *Journal of Hand Surgery*, vol. 29, no. 5, pp. 418–422, 2004.

[50] J. Tong, O. Mao, and D. Goldreich, "Two-point orientation discrimination versus the traditional two-point test for tactile spatial acuity assessment," *Frontiers in Human Neuroscience*, vol. 7, no. September, pp. 1–11, 2013.

[51] B. Spanlang, J.-M. Normand, D. Borland, K. Kilteni, E. Giannopoulos, A. Pomés, M. González-Franco, D. Perez-Marcos, J. Arroyo-Palacios, X. N. Muncunill, and M. Slater, "How to Build an Embodiment Lab: Achieving Body Representation Illusions in Virtual Reality," *Frontiers in Robotics and AI*, vol. 1, no. November, pp. 1–22, 2014.

[52] J. Bell-Krotoski and E. Tomancik, "The repeatability of testing with Semmes-Weinstein monofilaments.," *The Journal of hand surgery*, vol. 12, no. 1, pp. 155–161, 1987.

[53] Unity Technologies, "Unity3D."

[54] J. Kätsyri, K. Förger, M. Mäkäräinen, and T. Takala, "A review of empirical evidence on different uncanny valley hypotheses: support for perceptual mismatch as one road to the valley of eeriness," *Frontiers in Psychology*, vol. 6, p. 390, 4 2015.

[55] HTC Corporation, "HTC Vive," 2015.

[56] F. Mancini, A. Bauleo, J. Cole, F. Lui, C. A. Porro, P. Haggard, and G. D. Iannetti, "Whole-body mapping of spatial acuity for pain and touch," *Annals of Neurology*, vol. 75, no. 6, pp. 917–924, 2014.

[57] K. P. Körding, U. Beierholm, W. J. Ma, S. Quartz, J. B. Tenenbaum, and L. Shams, "Causal Inference in Multisensory Perception," *PLoS ONE*, vol. 2, 9 2007.

[58] M. Taylor-Clarke, P. Jacobsen, and P. Haggard, "Keeping the world a constant size: Object constancy in human touch," *Nature Neuroscience*, vol. 7, no. 3, pp. 219–220, 2004.

[59] M. R. Longo and P. Haggard, "Weber's Illusion and Body Shape: Anisotropy of Tactile Size Perception on the Hand," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 37, no. 3, pp. 720–726, 2011.

[60] Autodesk Inc., "Autodesk Maya," 2019.

[61] T. Hothorn, F. Bretz, and P. Westfall, "Simultaneous inference in general parametric models," *Biometrical Journal*, vol. 50, no. 3, pp. 346–363, 2008.

[62] R. core Team, "R: A Language and Environment for Statistical Computing," *R Foundation for Statistical Computing, Vienna, Austria*, 2018.

[63] M. E. Brooks, K. Kristensen, K. J. van Benthem, A. Magnusson, C. W. Berg, A. Nielsen, H. J. Skaug, M. Mächler, and B. M. Bolker, "glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling," *R Journal*, 2017.

[64] R. Lenth, H. Singmann, J. Love, P. Buerkner, and M. Herve, "emmeans: Estimated Marginal Means, aka Least-Squares Means," 2019.

[65] H. Wickham, M. Averick, J. Bryan, W. Chang, L. McGowan, R. François, G. Grolemund, A. Hayes, L. Henry, J. Hester, M. Kuhn, T. Pedersen, E. Miller, S. Bache, K. Müller, J. Ooms, D. Robinson, D. Seidel, V. Spinu, K. Takahashi, D. Vaughan, C. Wilke, K. Woo, and H. Yutani, "Welcome to the Tidyverse," *Journal of Open Source Software*, 2019.

[66] T. A. Kuiken, P. D. Marasco, B. a. Lock, R. N. Harden, and J. P. a. Dewald, "Redirection of cutaneous sensation from the hand to the chest skin of human amputees with targeted reinnervation.," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 104, no. 50, pp. 20061–20066, 2007.

[67] N. Rokhmanova, D. Boe, D. E. Caballero, L. Palve, and E. Rombokas, "Mapping of Referred Sensation after Targeted Reinnervation Surgery in the Lower Limb," in *Congress of the International Society of Biomechanics (IRB)*, 2019.

[68] M. Samad, E. Gatti, A. Hermes, W. Redmond, H. Benko, and C. Parise, "Pseudo-Haptic Weight: Changing the Perceived Weight of Virtual Objects By Manipulating Control-Display Ratio," *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, p. 13, 2019.

[69] R. J. Van Beers, D. M. Wolpert, and P. Haggard, "When Feeling Is More Important Than Seeing in Sensorimotor Adaptation," *Current Biology*, vol. 12, no. 02, pp. 834–837, 2002.

[70] R. J. Van Beers, P. Baraduc, and D. M. Wolpert, "Role of uncertainty in sensorimotor control," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 357, no. 1424, pp. 1137–1145, 2002.

[71] M. Samad and L. Shams, "Recalibrating the body: Visuotactile ventriloquism aftereffect," *PeerJ*, vol. 2018, p. e4504, 3 2018.

[72] F. E. van Beek, I. A. Kuling, E. Brenner, W. M. Bergmann Tiest, and A. M. L. Kappers, "Correcting for Visuo-Haptic Biases in 3D Haptic Guidance," *PLoS ONE*, vol. 11, 7 2016.

[73] I. A. Kuling, F. E. Van Beek, W. Mugge, and J. B. Smeets, "Adjusting Haptic Guidance to Idiosyncratic Visuo-Haptic Matching Errors Improves Perceptual Consistency in Reaching," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 6, pp. 921–925, 2016.

[74] H. H. Ehrsson, N. P. Holmes, and R. E. Passingham, "Touching a rubber hand: feeling of body ownership is associated with activity in multisensory brain areas," *Journal of Neuroscience*, vol. 25, no. 45, p. 10564–10573, 2005.

[75] M. Tsakiris and P. Haggard, "The rubber hand illusion revisited: visuotactile integration and self-attribution.," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 31, no. 1, p. 80–91, 2005.

[76] L. Dominjon, A. Lécuyer, J. M. Burkhardt, P. Richard, and S. Richir, "Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments," in *Proceedings - IEEE Virtual Reality*, pp. 19–318, IEEE, 2005.

[77] D. A. G. Jauregui, F. Argelaguet, A. H. Olivier, M. Marchal, F. Multon, and A. Lecuyer, "Toward 'Pseudo-haptic avatars': Modifying the visual animation of self-avatar can simulate the perception of weight lifting," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 4, pp. 654–661, 2014.

[78] K. L. Palmerius, D. Johansson, G. Höst, and K. Schönborn, "An analysis of the influence of a pseudo-haptic cue on the haptic perception of weight," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 8618, pp. 117–125, Springer, Berlin, Heidelberg, 2014.

[79] M. Rietzler, F. Geiselhart, J. Gugenheimer, and E. Rukzio, "Breaking the tracking: Enabling weight perception using perceivable tracking offsets," *Conference on Human Factors in Computing Systems - Proceedings*, vol. 2018-April, 2018.

[80] L. Shams and U. R. Beierholm, "Causal inference in perception," *Trends in Cognitive Sciences*, vol. 14, no. 9, pp. 425–432, 2010.

[81] K. P. Körding and D. M. Wolpert, "Bayesian integration in sensorimotor learning," *Nature*, vol. 427, no. 6971, pp. 244–247, 2004.

[82] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan, "An internal model for sensorimotor integration," *Science*, pp. 1880–2, 1995.

[83] J. M. Foley and R. Held, "Visually directed pointing as a function of target distance, direction, and available cues," *Perception & Psychophysics*, 1972.

[84] J. M. Foley, "Successive stereo and vernier position discrimination as a function of dark interval duration," *Vision Research*, 1976.

[85] D. Burin, A. Livelli, F. Garbarini, C. Fossataro, A. Folegatti, P. Gindri, and L. Pia, "Are Movements Necessary for the Sense of Body Ownership? Evidence from the Rubber Hand Illusion in Pure Hemiplegic Patients," *PLOS ONE*, vol. 10, p. e0117155, 3 2015.

[86] M. Berniker and K. Kording, "Estimating the sources of motor errors for adaptation and generalization," *Nature Neuroscience*, vol. 11, no. 12, pp. 1454–1461, 2008.

[87] R. Shadmehr, M. A. Smith, and J. W. Krakauer, "Error Correction, Sensory Prediction, and Adaptation in Motor Control," *Annual Review of Neuroscience*, 2010.

[88] Y. W. Tseng, J. Diedrichsen, J. W. Krakauer, R. Shadmehr, and A. J. Bastian, "Sensory prediction errors drive cerebellum-dependent adaptation of reaching," *Journal of Neurophysiology*, vol. 98, no. 1, pp. 54–62, 2007.

[89] C. Tin and C. S. Poon, "Internal models in sensorimotor integration: Perspectives from adaptive control theory," *Journal of Neural Engineering*, vol. 2, no. 3, 2005.

[90] D. M. Wolpert, "Probabilistic models in human sensorimotor control," *Human Movement Science*, vol. 26, pp. 511–524, 8 2007.

[91] B. Craske and S. J. Gregg, "Prism After-effects: Identical Results for Visual Targets and Unexposed Limb," *Nature*, vol. 212, pp. 104–105, 10 1966.

[92] S. Abeele and O. Bock, "Sensorimotor adaptation to rotated visual input: different mechanisms for small versus large rotations," *Experimental Brain Research*, vol. 140, pp. 407–410, 10 2001.

[93] J. R. Flanagan and A. K. Rao, "Trajectory adaptation to a nonlinear visuomotor transformation: Evidence of motion planning in visually perceived space," *Journal of Neurophysiology*, vol. 74, no. 5, pp. 2174–2178, 1995.

[94] D. Y. P. Henriques and E. K. Cressman, "Visuomotor Adaptation and Proprioceptive Recalibration," *Journal of Motor Behavior*, vol. 44, pp. 435–444, 11 2012.

[95] C. S. Harris, "Adaptation to displaced vision: Visual, motor, or proprioceptive change?," *Science*, vol. 140, no. 3568, pp. 812–813, 1963.

[96] M. Berniker, D. W. Franklin, J. R. Flanagan, D. M. Wolpert, and K. Kording, "Motor learning of novel dynamics is not represented in a single global coordinate system: Evaluation of mixed coordinate representations and local learning," *Journal of Neurophysiology*, vol. 111, no. 6, pp. 1165–1182, 2014.