2012

Expressive feedback from virtual buttons

Adam Joseph Faeth
Iowa State University

Recommended Citation

Faeth, Adam Joseph, "Expressive feedback from virtual buttons" (2012). *Graduate Theses and Dissertations*. 12719.
https://lib.dr.iastate.edu/etd/12719

Expressive feedback from virtual buttons

by

Adam Joseph Faeth

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Co-majors: Computer Engineering
Human Computer Interaction

Program of Study Committee:
Chris Harding, Co-major Professor
James Oliver, Co-major Professor
Stephen Gilbert
Julie Dickerson
Tien Nguyen
Alexander Stoytchev

Iowa State University
Ames, Iowa
2012

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGEMENTS

This research would not have been possible without the advice and support of colleagues and friends who challenged me to grow. I am very grateful for the direction of my Co-major Professors, Dr. Chris Harding, and Dr. James Oliver. I would also like to thank my committee members for their guidance: Dr. Stephen Gilbert, Dr. Julie Dickerson, Dr. Tien Nguyen, and Dr. Alexander Stoytchev. In particular, I would like to thank Dr. Harding for his active involvement in my research and graduate education. Dr. Oliver guided me in the program and encouraged me to make new connections. I would also like to thank Dr. Gilbert for helping to explore the human side of human computer interaction.

I was lucky to build friendships with so many smart students in the Human Computer Interaction program and Virtual Reality Applications Center. I would like to thank Pamela Shill, Karen Koppenhaver, Jean Bessman, and Beth Hageman for their tireless efforts to make the university and program work for the students. Michael Oren, Jeremy Bennett, and Michael Carter provided valuable feedback and support during this research.

This work is dedicated to my wife Erin. Her confidence in me and her tenacity helped me to quickly start again after the frustrating setbacks that occur during research.

I also appreciate all of the experiences that my parents, Loren and Arleen Faeth, helped create for me. Their encouragement helped to make me curious to explore the world.

ABSTRACT

The simple action of pressing a button is a multimodal interaction with an interesting depth of complexity. As the development of computer interfaces to support 3D tasks evolves, there is a need to better understand how users will interact with virtual buttons that generate feedback from multiple sensory modalities. This research examined the effects of visual, auditory, and haptic feedback from virtual buttons on task performance dialing phone numbers and on the motion of individual buttons presses. This research also presents a theoretical framework for virtual button feedback and a model of virtual button feedback that includes touch feedback hysteresis.

The results suggest that although haptic feedback alone was not enough to prevent participants from pressing the button farther than necessary, bimodal and trimodal feedback combinations that included haptic feedback shortened the depth of the presses. However, the shallower presses observed during trimodal feedback may have led to a counterintuitive increase in the number of digits that the participants omitted during the task. Even though interaction with virtual buttons may appear simple, it is important to understand the complexities behind the multimodal interaction because users will seek out the multimodal interactions they prefer.

CHAPTER 1. INTRODUCTION

1.1 Motivation

Our everyday familiarity with pressing mechanical buttons covers an interesting depth of complexity in what appears to be a simple multimodal interaction. When the user presses a button, the key travels a short distance before closing the electrical circuit to actuate the button. During this travel, the user can feel the resistance of the button and the click sensation, see the key traveling, and hear an actuation sound. This experience with physical buttons forms the expectations that users have when interacting with virtual buttons.

Figure 1.1 A stylus interacts with virtual buttons in a 3D interface

A virtual button has at least some of the feedback generated by a computer, rather than produced by inherent mechanical components. This means the designer can control the feedback experience, even though the design may be informed by experience with mechanical buttons.

However, virtual buttons do not always provide the same feedback modalities and fidelity of mechanical buttons. It is not well understood how tradeoffs in the design of virtual buttons, such as omitting certain modalities, violate expectations that users acquired from their interaction with mechanical buttons nor how this may affect task performance.

In both three-dimensional and touchscreen interactions, there is an increase in the prevalence of virtual buttons. Virtual buttons are used in 3D interfaces developed for applications in medicine, manufacturing, geoscience, and engineering (Figure 1.1). Gestural interfaces and direct manipulation are more appropriate than buttons for some interactions such as scrolling through a large area. However, buttons retain the advantages of discoverability and perceived affordance, which make them important components in an interface.

1.2 Definition of Terms

This research uses the following terms throughout the descriptions of the research questions, contributions, and study methods.

3D Interface A user interface where the user can interact with the controls in three dimensions.

Virtual Button A button in which the feedback generated for the user is not solely dependent on the movement of mechanical components.

Travel In a button, travel refers to the distance that the user presses the button away from the resting position.

Actuation Actuation refers to the threshold in the button travel where the button activates and records the press.

Deactuation When the user releases the button, the deactuation point is where the virtual switch of the button opens, ending the activation.

Touch Event The point where the user first makes contact with a button at the start of the press.

Release Event The point at the end of the press where the user releases the button.

Hysteresis Hysteresis refers to using a memory of the state in the evaluation of a new state. In a virtual button, hysteresis allows the button to use different thresholds for actuation and deactuation based on the prior actuation of the button.

Stiffness In the context of this research, stiffness refers to the hardness of the surface that the haptics device can render. Although the virtual environment might define a haptic surface as truly hard, the haptic device may give the surface a spongy feeling if it cannot bring enough force to perfectly counter the forces applied by the user.

Unactuated Press A button press where the user did not press far enough to travel past the actuation threshold.

1.3 Research Questions

The purpose of the user studies described in Chapters 4 and 5 was to gain a better understanding of how the combinations of visual, auditory and haptic feedback affect the performance of the user. This understanding is important because many systems make tradeoffs in the number of feedback modalities, and the effects of omitting modalities are not well understood.

- RQ1. Will user performance improve as the number of feedback modalities from virtual buttons increases?
- RQ2. Compared to trimodal feedback, how does reducing the number of feedback modalities affect task accuracy and completion time?

Another factor in the success of a 3D interface is the user's subjective rating of the feedback provided by the system. It is important to understand what the user perceives as effective feedback because designing a virtual button system based on combinations of modalities that the user finds ineffective should be avoided. The following research question addresses the users' subjective rating of the feedback in Chapter 4.

- RQ3. Will participants prefer virtual buttons that provide feedback using a higher number of modalities?

The motion of the button through space during a press can also lead to a better understanding of the way that users interact with multimodal feedback from virtual buttons. For example, examining the motion of the press can show how far past the virtual bottom of the button a user continued to press after actuation. This information may be important to explain how the different combinations of feedback affected task performance. The research presented in Chapter 5 addresses the following question.

- RQ4. How does the combination of feedback modalities affect the motion of the 3D button press?

Some of the participants from the first study anecdotally reported they were looking for cues from a particular modality before cues from the other modalities. If there is indeed evidence that one modality contributed more to the successful completion of the tasks, it would be important to the design of virtual button feedback. Therefore, part of Chapter 5 explores the following question.

- RQ5. When users press a button, are they focusing on a particular type of feedback such as a visual cue?

If users perceive the feedback cues from virtual buttons to be inconsistent or random, they will distrust the virtual buttons. Therefore, it is important to understand how sequences in the motion of presses may generate confusing feedback. The study presented in Chapter 6 measured how frequently events like extra threshold crossings and threshold dwell occurred in the motion of individual button presses and evaluated the effectiveness of touch hysteresis in preventing confusing feedback cues.

- RQ6. In a virtual button, what effect does hysteresis at the touch and actuation thresholds have on preventing confusing feedback signals?
- RQ7. What effect does the combination of feedback modalities used in virtual buttons have on the threshold crossings and threshold dwell?

1.4 Organization

Chapter 2 will detail the related work and background used to develop the methods for the user studies. Chapter 3 provides a framework for generating virtual button feedback. The study presented in Chapter 4 addresses the research questions about the effect of feedback on performance and the users' subjective rating of the feedback. In Chapter 5, we present the results of a follow-up study to evaluate the effect of the feedback conditions on the motion of individual presses and task success. Chapter 6 presents an evaluation of touch hysteresis to prevent confusing feedback cues and a model for virtual button feedback. Finally, Chapter 7 summarizes the findings and contributions of this research.

1.5 Contributions of this Research

This research presents a framework for virtual button feedback and a model for touch hysteresis in virtual buttons. The studies in the following chapters describe the evaluation of the effects of feedback modalities on performance, subjective rating, the individual motion of presses, and the effectiveness of touch hysteresis for preventing confusing feedback cues.

1. Evaluation of the effect of feedback modalities on task performance

   Although Chapter 2 presents the results of past studies of multimodal feedback from virtual buttons, prior studies were limited to 2D interfaces or bimodal feedback. The studies in Chapters 4 and 5 evaluate combinations of visual, auditory, and haptic feedback from virtual buttons in a 3D force-feedback environment. This research presents findings that were not suggested by prior studies, including a counterintuitive increase in errors with trimodal feedback and the shallower button presses that might explain that effect.

2. The effect of feedback modalities on the high-resolution motion of individual presses

   Several earlier studies detailed in Chapter 2 document the recording of physical button travel and the motion of virtual buttons. However, to the author's knowledge, this is the first research to measure the motion of the user pressing individual virtual buttons and compare this analysis with task performance (Chapter 5).

3. Evaluation of subjective rating of feedback modalities from virtual buttons.

   This research replicates the prior findings that users prefer multimodal interactions over unimodal interactions. The studies in Chapters 4 and 5 further investigate whether users prefer particular combinations of feedback modalities.

4. Evaluation of touch hysteresis for virtual button feedback

   Sequences of repeated threshold crossings during a button press could result in confusing feedback cues from a single button press. Chapter 6 presents the results of a study that investigated how often two of these sequences occurred and how effective touch hysteresis would be in preventing repeated feedback cues.

5. A framework for virtual buttons with multimodal feedback that includes touch hysteresis

   Chapter 3 presents a theoretical framework for virtual button feedback, and Chapter 6 presents a model for virtual button feedback that includes touch hysteresis. The studies in Chapters 4, 5, and 6 use this framework as a basis for creating virtual button feedback.

These contributions outline the complexity involved in transferring what appears to be a simple everyday interaction into a virtual experience. It is important to understand the complexities and potential pitfalls of implementing multimodal interactions because users tend to prefer multimodal interactions. This research details the intuitive and counterintuitive results observed during the study of virtual button feedback, and places it in the context of previous research.

CHAPTER 2. BACKGROUND

Pressing a physical button is a multimodal experience that provides visual, auditory, and haptic feedback. This work builds upon research in keyboard ergonomics, multimodal feedback, and studies of bimodal and trimodal feedback. The relationships between these areas is shown in Figure 2.1.

2.1 Related Work

2.1.1 Keyboard Size and Shape

Keyboards are a rich source for virtual button inspiration because the buttons are designed to provide clear feedback for presses in quick succession. The ideal characteristics of shape, size, travel, and force required to actuate a keyboard button are detailed in the literature (Clare, 1976; Kinkead and Gonzalez, 1969; Kinkead, 1975). Kromer (2001) provides a chronology of keyboard development between 1878 and 1999.

2.1.2 The Sense of Touch

The combination of tactile and kinesthetic sensory input make up the haptic system (Klatzky and Lederman, 2003). Touch is an exploratory sense, which requires the active exploration of objects to discover their properties (Klatzky and Lederman, 2008). This exploratory aspect also applies to the use of force-feedback haptic devices to render haptic properties of virtual objects, for example, the use of pressure to identify that the button travels, and contour following to identify the boundary of the button. Lateral motion may also help to identify the textures or shapes that distinguish a button from its surroundings.

The tactile sense resolves temporal differences in sequential data well and resolves limited

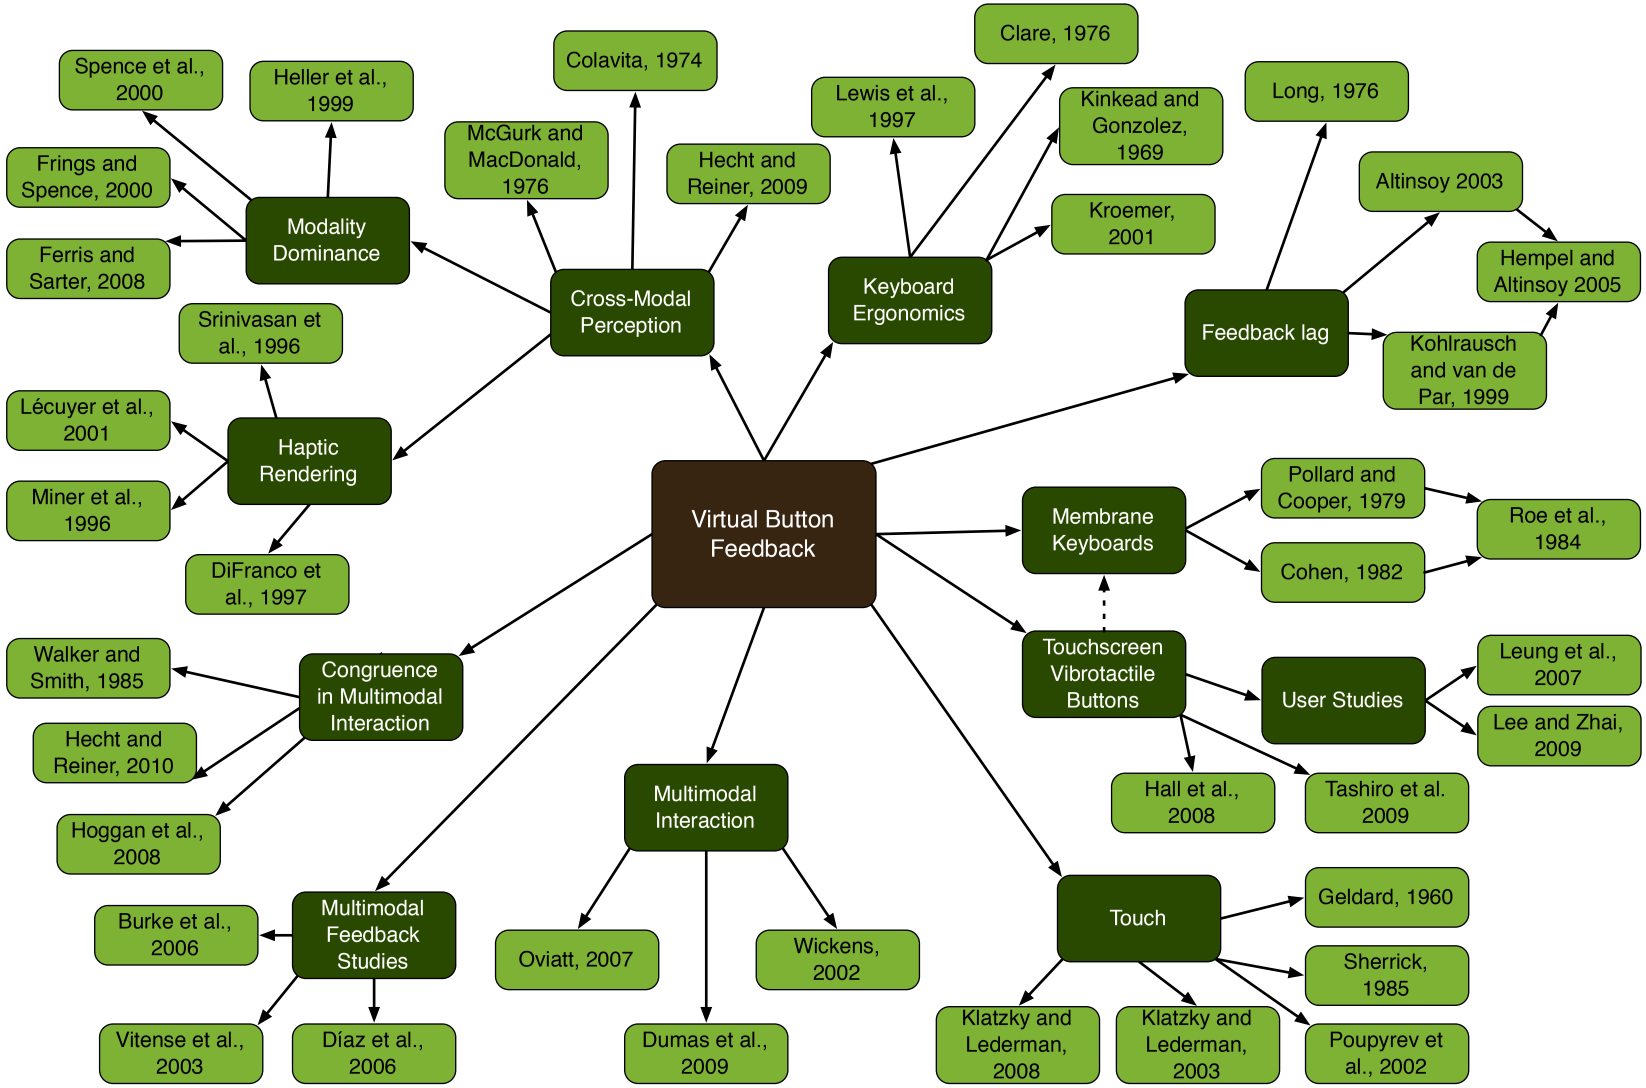

Figure 2.1 Relationships between related work including selected references.

differences in spatial data (Geldard, 1960; Sherrick, 1985). Jones and Sarter (2008) detailed the parameters of vibrotactile perception and its role in multimodal interfaces. In a practical application, ambient information conveyed through touch improved accuracy and completion time of tasks which require stopping on a target (Poupyrev et al., 2002).

2.1.3 Keyboard Travel and Typing on Flat Surfaces

Another desirable characteristic of a button is the haptic click sensation, a gentle rise in force as the button travels, and a sharp drop in force near the actuation point (Lewis et al., 1997). Typing on a flat surface, where they keys do not travel, decreased performance (Cohen, 1982; Pollard and Cooper, 1979). This negative effect may be reduced by adding auditory feedback or tactile features to the flat surface (Roe et al., 1984; Nashel and Razzaque, 2003; Lee and Zhai, 2009; Fukumoto and Sugimura, 2001). Hall et al. (2008) proposed that the user might also slide a finger along the surface of the flat surface to activate a button, an action analogous to the travel of a button.

2.1.4 Preventing Repeat Errors

It is also important to consider aspects of button design that reduce unintended multiple actuations from a single button press. Hysteresis refers to creating a distance between the point where a button actuates on the downstroke and where it deactuates on the upstroke (Lewis et al., 1997). Hysteresis helps to avoid repeated actuations for a single button press when the user's press repeatedly crosses the actuation point, or the sampled button travel fluctuates due to signal noise. When hysteresis is not practical for cost or technical considerations, interlocks also prevent unintended multiple actuations by ignoring subsequent actuations for a short period of time after an initial actuation.

2.1.5 Effects of Lag on Multimodal Feedback

A delay between the press and feedback from the press also hampers the use of keyboard buttons, and an irregular lag in feedback can even cause problems for skilled typists (Long, 1976). Humans are less sensitive to audio signals that lag behind visual or haptic signals. Kohlrausch summarized the minimum perceptible lag between visual and audio stimuli in various situations, including lag between visual and auditory feedback from a discrete event (Kohlrausch and van de Par, 1999). Altinsoy measured the minimum perceived lag between auditory feedback and haptic feedback (Altinsoy, 2001, 2003).

2.1.6 Multimodal Interaction

Many users prefer multimodal interactions, and in some situations, multimodal feedback can reduce errors and completion time (Oviatt, 2007; Dumas et al., 2009). Users also react faster to multimodal signals than to unimodal signals (Diederich and Colonius, 2004; Forster et al., 2002; Hecht et al., 2008a,b; Hershenson, 1962; Giard and Peronnet, 1999). The fast reaction times and increased accuracy also extended to recognition of short messages conveyed with multimodal feedback (Merlo et al., 2010). At a neurological level, there is also evidence that the superior colliculus integrates input from multiple modalities in a multiplicitive ratio (Meredith and Stein, 1986). This suggests that multimodal feedback produces more signals than the sum of unimodal signals at a neurological level.

Multiple resource theory provides one explanation for how humans process signals from multiple sensory modalities. According to the 4-dimensional multiple resources model, tasks which engage different modalities are less likely to interfere with each other (Wickens, 2002, 2008). Although the original model only addressed visual-audio signals, Boles et al. extended the model to incorporate haptic signals (Boles et al., 2007).

2.1.7 Studies of Multimodal Feedback

The addition of either auditory feedback or haptic feedback improved completion times and performance scores over visual feedback alone (Burke et al., 2006). The addition of vibrotactile feedback can also reduce mental workload (Vitense et al., 2003).

In a task that asked users to trace a path between two concentric circles, tactile feedback helped users stay within the lines (Sun et al., 2010). A similar task asked participants to trace a path through a 3D maze, and participants better remained within the walls when using bimodal and trimodal feedback compared to visual feedback alone (Díaz et al., 2006). Participants of a 3D multimodal assembly task also reported a greater the sense of presence when the task included haptic feedback (Petzold et al., 2012).

Several studies have examined the addition of haptic or multimodal feedback to 2D interfaces. For example, Miller and Zeleznik (1998) described the implementation of a haptic windowing system and reported that the authors completed tasks faster using the system. Multimodal feedback reduced targeting times (Akamatsu et al., 1995), but it could confuse the user if targets in close proximity caused overlapping feedback (Cockburn and Brewster, 2005). In addition to vibrotactile feedback, variable friction also demonstrated lower completion times in a targeting task (Levesque et al., 2011). The addition of active or passive force-feedback to a 2D button also decreased task-completion time (Rosenberg and Brave, 1996).

Leung et al. (2007) added haptic feedback to a mobile GUI that the user felt through a stylus and found that it reduced task completion times. Vibrotactile feedback has also been added to directly to a stylus (Lee et al., 2004). Tactile feedback from a stylus improved targeting time for pointing and crossing tasks (Forlines and Balakrishnan, 2008).

Multimodal interaction may provide a mechanism for incorporating affective communication into interfaces. Chu et al. (2009) suggested that interface widgets could sense the force of the user's input to allow the user to express their conviction with an action. The virtual feedback rendered by a control also demonstrated the ability to influence affective responses in users (Swindells et al., 2007).

Another area of research investigated the effects of multimodal feedback on older adults. Multimodal feedback improved the time-based performance for groups with and without age-related macular degeneration (Jacko et al., 2005). Users in both under-35 and over-55 age groups interacted with the same modality of a multimodal interface, but users in the over-55 group were less likely to switch interaction modalities after an unsuccessful attempt (Naumann et al., 2010).

2.1.8 Cross-modal Perception

However, when multiple modalities of feedback are combined, the effect is not always straightforward. In the McGurk effect, the strong visual cue of a person's lips repeating “fa, fa, fa" turns the sound "ba, ba, ba" into the perceived *fa* syllable even though the sound remains unchanged (McGurk and MacDonald, 1976; Green, 1996). When presenting auditory and visual cues together simultaneously, a user is more likely to miss the audio cue than the visual cue (Colavita, 1974). Similarly, one is more likely to miss a haptic cue presented with a visual cue, but this visual dominance diminishes when trimodal cues are provided (Hecht and Reiner, 2009).

Haptic rendering often employs an illusion based on cross-modal perception where the graphical representation of the stylus remains on the surface of an object, even when the haptic device cannot render enough stiffness to prevent the user's hand from penetrating the object (Figure 2.2). The strong visual cue makes the surface appear stiffer to the user than the haptic device can render (Srinivasan et al., 1996). Similarly, the visual displacement of a spring can influence the perceived stiffness of the spring (Lécuyer et al., 2001). An audio cue may also reinforce haptic sensation, but the effect may be weaker without carefully selected sounds (DiFranco et al., 1997; Miner et al., 1996). The viewing angle through which the user observes an interaction has also influenced the subjective perception of stiffness (Widmer and Hu, 2010). One paper proposed that this visual effect might be used to simulate tactile properties in the absence of haptic feedback (van Mensvoort, 2002).

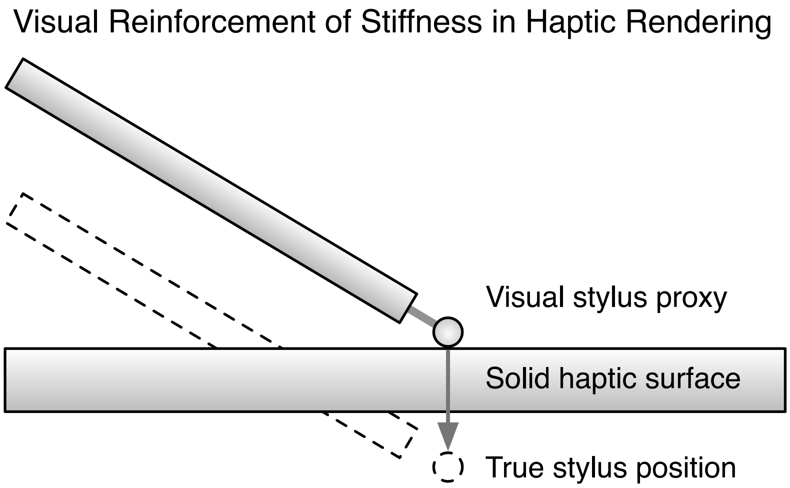

Figure 2.2 The visual representation of the stylus remains on the surface of the object, even when the haptic device cannot render enough stiffness to prevent the user from pressing through the surface.

However, the dominance of vision is not absolute, the task requirements and user's attention impact the reliance on one modality over another (Heller et al., 1999; Spence et al., 2000). For example, users were more accurate in identifying a rhythm presented to the auditory modality despite incongruent visual or haptic feedback (Frings and Spence, 2010). In another example, users responded faster to a visual target when it was preceded by a spatial audio cue (Ferris and Sarter, 2008).

2.1.9 Haptic Event Feedback

In an attempt to improve the realism of haptic contact with virtual surfaces, a haptic device might replay a complex prerecorded signal when the user makes contact with the virtual surface (Kuchenbecker et al., 2006). Haptic events such as a button click have also been rendered to the tactile sense using jets of air (Kim et al., 2006). Colton and Hollerbach (2007) created a model of mechanical button click that might provide a basis for rendering clicks as haptic events. Researchers recreated the feeling of a mechanical button click on a flat surface using a replay of vibrotactile event feedback (Tashiro et al., 2009b). User performance was similar using the ultrasonic vibration on a flat surface to pressing a mechanical pushbutton (Tashiro et al., 2009a).

2.1.10 Congruence in Multimodal Interaction

The congruence, or the intuitive match between visual, auditory, and haptic properties of a button may also affect interaction. Users performed tasks faster when the haptic size of a button key matched the pitch of an audio cue (Walker and Smith, 1985). Users also preferred buttons where the visual appearance matched the vibrotactile feedback of the button (Hoggan et al., 2008). Interference between symbols or language and haptic sensation, Stroop Interference, also affected accuracy and reaction time (Hecht and Reiner, 2010).

There exist several sets of guidelines to address the practical considerations when designing multimodal interactions. Stanney et al. developed guidelines for optimizing the multimodal presentation of information (Stanney et al., 2004). Hempel and Altinsoy detailed the properties and design consideration for the auditory and haptic modalities. The Multimodal Flexibility

Index indicates which modalities in a multimodal interaction the user's performance depends on and which might be available for other tasks (Oulasvirta and Bergstrom-Lehtovirta, 2010).

MacLean provided an overview of the attributes of haptic feedback and benefits of including the modality in an interface (MacLean, 2000). Sjöström's (2001) suggestions for designing haptic interfaces included using distinct texture on inactive buttons. Bodden and Iglseder (2002) developed guidelines for the design of sounds to augment the inherent sounds of a mechanical device.

2.1.11 Adapting Multimodal Feedback to the User's Environment

One advantage of multimodal interaction is that the system can provide feedback despite environmental interference with another modality (Hoggan and Brewster, 2010). Hoggan et al. (2009) measured noise and vibration thresholds for switching between auditory and vibrotactile feedback. With an identification of the user's context, the system might choose the appropriate modalities for the situation (Kong et al., 2011). When providing feedback to multiple modalities, designers may be able to identify patterns that communicate information in similar ways to the auditory and haptic modalities (Hoggan and Brewster, 2006).

2.1.12 Haptic Feedback from 3D Interfaces

In a 3D interaction, leveraging a user's awareness of the relative position of their own hands decreased errors when targeting menu items (Boeck et al., 2006). Another study found that the gesture of pushing through a 3D menu item also decreased selection time compared to clicking a stylus button while in contact with an item (Komerska and Ware, 2004). The specific interactions examined by these studies were similar to a button keypad because the users could target the menu items independently, without maintaining a hold on a parent item.

2.1.13 Virtual Buttons on Emerging Hardware

A number of other input techniques may also make use of virtual button feedback in interesting ways. The UnMousePad demonstrated a way to interpolate force sensing on a flat touchscreen film (Rosenberg and Perlin, 2009). This touchscreen technique could allow the force the user applied to the surface of the touchscreen to serve as a proxy for the travel of a virtual button. Skinput can project virtual buttons onto the user's skin and use the propagation of mechanical vibration to discern the location of finger taps on the user's skin (Harrison et al., 2010). The amplitude of those vibrations could provide a relative indication of the force applied to a virtual button to produce varied virtual button feedback. A pressure sensitive keyboard may provide an interesting method for pairing virtual button feedback with a physical keyboard button (Dietz et al., 2009). Although the physical keyboard would provide tactile, proprioceptive, and auditory feedback, the force of the keystroke could also be used to provide increasing visual and auditory reinforcement feedback from a virtual button. Snibbe and MacLean (2001) augmented physical buttons with a pressure sensor on the top to create a hybrid virtual-physical button which provided feedback about a button's action when the user touched the surface lightly. These emerging input techniques would benefit from a better understanding of the characteristics of effective virtual button feedback.

2.2 Methods Background

2.2.1 Evaluating virtual button performance

To assess the performance of the virtual buttons, the studies presented in Chapter 4 and Chapter 5 asked participants to dial random seven-digit phone numbers. Lewis et al. (1997) suggested that a wider audience of participants would be comfortable with a phone number task than with a calculator task. Typing on a full virtual keyboard with a PHANToM haptic device would be similar in practicality to typing on a keyboard using a single finger because the haptic device renders feedback from a single point in the virtual workspace.

2.2.2 Counting Errors

To calculate a fair measure of the errors that the participant made during the study, we used a string matching algorithm to compare the phone number assigned by the task to the digits entered by the participant. We started with an algorithm from Python's difflib library (Python-difflib, 2012), which uses a modified version of the Ratcliff/Obershelp algorithm (Ratcliff and

Metzener, 1988). For our short sequences of phone numbers in the first study, we found that the sum of differences reported by this algorithm was identical to the Levenshtein distance between the task and entry sequences (Levenshtein, 1965). However, in the second study, the Levenshtein distance reported a lower number of errors for three of the 700 tasks. In these three cases, the difflib library scored an insert and a delete operation instead of a single replace operation. Therefore, we changed the error-counting algorithm for both Study 1 and Study 2 to use the Wagner-Fisher algorithm for calculating Levenshtein distance (Wagner and Fischer, 1974).

2.2.3 Apparatus

To evaluate the performance of virtual button feedback, a test platform was constructed in a 3D environment. The user held a PHANToM haptic device in their dominant hand to collect 3D position and render force-feedback (Massie and Salisbury, 1994). The virtual buttons were implemented in H3DAPI, a scenegraph for haptic rendering (H3DAPI, 2012). The Pure Data framework rendered the auditory feedback from the buttons (Pure Data, 2012).

2.2.4 Monoscopic and Stereoscopic Display

The study presented the 3D button keypad on a monoscopic display. Stereoscopic displays provide additional depth cues and advantages in discriminating between objects in complex spatial scenes (Kim et al., 1987; Van Orden and Broyles, 2000; Wickens et al., 1994). However, the phone keypad presented to users in the studies in Chapters 4 and 5 used large, non-overlapping buttons arranged in a fixed location orthogonal to the user. There is little evidence that stereoscopic display provides a benefit in this simple planar arrangement because the user did not need to discriminate between overlapping buttons at different depths (Van Orden and Broyles, 2000; Steiner and Dotson, 1990; Stanney et al., 2012).

2.3 Summary

Although there are many studies of bimodal feedback, there are a limited number of studies that examined trimodal feedback (Díaz et al., 2006; Hecht et al., 2008a; Vitense et al., 2003).

Of the studies that examined trimodal feedback, even fewer investigated multimodal feedback in the context of three-dimensional interaction (Díaz et al., 2006). Díaz et al. only investigated multimodal combinations which included visual feedback. To help fill this gap, this study examined the effect of visual, auditory, and haptic feedback from virtual buttons on the 3D motion of the button press and on task performance.

CHAPTER 3. A THEORETICAL FRAMEWORK FOR GENERATING COPIOUS MULTI-SENSORY FEEDBACK FROM VIRTUAL BUTTONS

Modified from a paper published by Faeth, A. and Harding, C. (2010) In *Proc. ASME 2010 World Conference on Innovative Virtual Reality (WINVR2010)*. pp 119–125.

3.1 Introduction

The structure of a mechanical button provides insight into its feedback properties. The button key is the component that the user presses, and its travel characteristic defines how far it can be pushed. The *actuation point* marks the position on the travel where the electric circuit closes and the system responds to the press (Figure 3.1). In a buckling spring mechanism, the actuation point is close to the point where the spring buckles, creating a drop in force feedback. Below the actuation point is the cushion where the force feedback increases before reaching the end of the travel. When the user releases the button, a spring pushes the button key opposite to the travel direction. When the button reaches the *deactuation point*, the switch opens. If the deactuation point is above the actuation point, hysteresis prevents the switch from immediately closing again. When the user fully releases the button, it returns to the starting position.

A button provides continuous feedback as it slides along the travel, informing the user about its position relative to the end of the travel distance. The increasing force from a spring provides continuous *haptic* feedback. Friction from the button assembly moving along the travel generates continuous auditory feedback. The mechanical button also provides continuous visual feedback as it recedes relative to its surroundings.

Buttons provide event-based feedback at certain points along the travel, such as the actuation and deactuation points. A button may also provide visual feedback by turning on a light

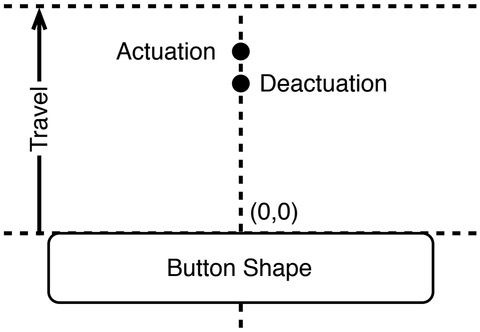

Figure 3.1 Characteristics of a button

at the actuation point, and turning it off at deactuation. The distinct buckle at the actuation point of a buckling spring mechanism provides haptic feedback. Most mechanical buttons emit a characteristic sound when the mechanical switch closes at the actuation point and opens at the deactuation point.

3.2 Actuation parameter

A framework for generating feedback for user interaction with a virtual button must support at least one actuation parameter to determine when the button actuation event occurs. Table 3.1 lists examples of actuation parameters. The actuation parameter may also serve as the input for transfer functions that generate feedback for the virtual buttons.

Table 3.1 Actuation parameters for a virtual button framework

Parameter	**Type**
Contact	Binary
Force	Threshold
Time	Threshold
Travel	Threshold

A straightforward approach would be to actuate the virtual button on contact with the user's avatar or virtual proxy. This method is also used by most implementations of touchscreen virtual buttons, though the actuation event may be delayed until the user lifts their finger. Although a contact actuation parameter is straightforward to implement, it may result in accidental activation if the buttons are located close to one another or within reach of the user's workspace.

If user interaction force can be measured, the button might also actuate when the force applied to the surface of the button exceeds a given threshold value. In a system lacking force feedback, contact velocity between the avatar and the button might provide a similar input. However, this approach may confuse users if there is no intermediate, or escalating feedback between touch and activation. If feedback only occurs upon actuation, the users may not realize that their contact had any effect on the button. Conversely, if the button provides feedback on contact, the users may sense they actuated the button when they failed to press it hard enough for actuation. Training may reduce errors, but it also has the potential to frustrate users interacting with virtual buttons.

A time parameter could also distinguish intentional actuations from unintentional or exploratory contacts. The button may require contact for a predetermined amount of time before actuating. This has the advantage of providing a parameter for intermediate feedback as time increases. However, the user must spend the actuation delay time in contact with the button for each press, reducing their potential clicks-per-second speed.

A better solution would be to model the travel of a mechanical button, designating the distance a user must push the button key between contact and the actuation point. A travel parameter allows continuous intermediate feedback between contact and actuation. The advantage of using travel instead of force as the actuation parameter is that the movement of the button provides strong visual feedback. As the button travels, the shape sinks or recedes in relation to neighboring objects, providing depth cues in a 3D virtual environment. A travel parameter also allows haptic feedback to simulate harder or softer button springs, and to generate the distinct click sensation of a mechanical button by using the force/travel model shown in Figure 3.2.

3.3 Properties of Virtual Button Feedback

Another requirement of a virtual button framework is the feedback generated when the button is pressed. The feedback characteristics define the user's experience while interacting with the virtual buttons.

There are several desirable characteristics for virtual button feedback. To make the button travel more discoverable, feedback should escalate between contact and actuation (Lewis et al., 1997). Actuation feedback should be discernible from the intermediate feedback, escalating to ensure the user understands when the actuation occurred. Initial contact feedback should be noticeable to help the user feel contact with the button. However, to allow for escalating feedback as the user presses the button, initial contact should barely be noticeable. Lag time between the user's action and button feedback must be minimal (Long, 1976; Clare, 1976).

There are two types of feedback a virtual button can generate for the user: event-based feedback and continuous feedback. Event-based feedback refers to predetermined sequences that may replay at any specific actuation value. Prominent events are actuation, deactuation, touch, and release. At these events a sound file could provide auditory feedback of a click, the color of the button shape could change until a subsequent event occurs, or a predetermined sequence of vibrotactile feedback could be replayed.

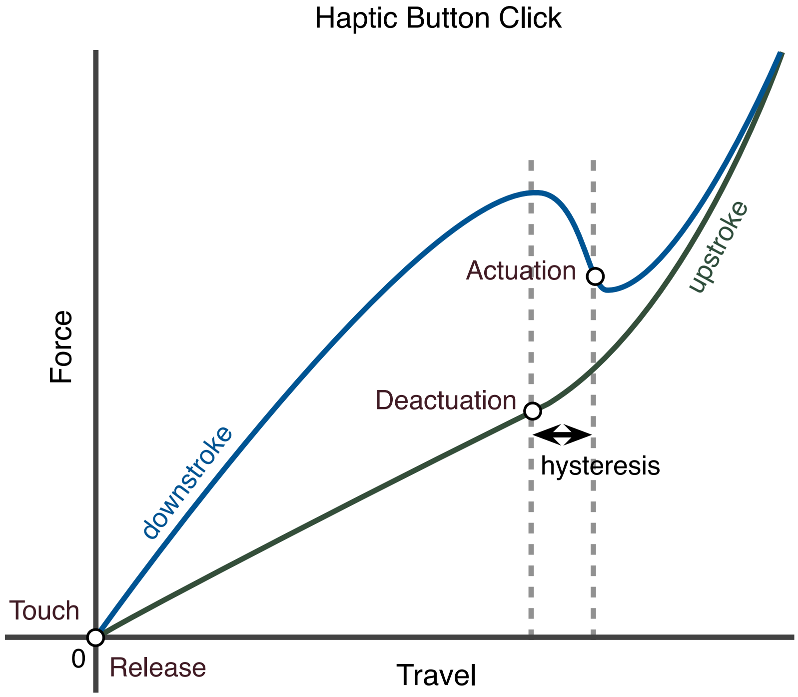

Figure 3.2 Force-travel transfer function for generating the feeling of haptic click

Continuous feedback is the output from a transfer function using the actuation parameter as the input. The force/travel graph from Figure 3.2 is an example of this type of feedback; the haptic click is not simply a replayed sequence—the force is calculated from the transfer function based on the current travel position. Auditory feedback could also provide continuous feedback by varying a sound parameter, such as amplitude or pitch, using a transfer function. Visual feedback indicating the current position of the button as it travels downward also provides continuous feedback as it recedes in size from the viewer. A stronger form of continuous visual feedback might be a two-color interpolation based on the actuation parameter: a base color and an actuated color.

Both continuous feedback and event-based feedback have advantages, and it is important to understand the tradeoffs between the two when designing virtual button feedback. Continuous feedback closely matches the actions of the user: a desirable feedback characteristic. It may be a greater challenge to design continuous feedback, as it must be effective when rendered for both quick and prolonged button traversals. This is primarily a concern for visual feedback, because people can resolve haptic and auditory feedback at higher rates (Geldard, 1960). If a button's travel is quick, complex visual feedback might be condensed to a brief flash.

The advantage of event-based feedback is that it replays at a constant rate, and thus it is simple to design. The disadvantage is that a feedback sequence may take longer to replay than the time available before the next event. This effect may be minimized by keeping feedback sequences short. Since both continuous and event-based feedbacks have advantages, combinations of continuous and event-based feedback may also be effective.

3.4 Virtual button events

A virtual button framework must consider what events will occur during a button press, and when to trigger those events. Applications employing haptic rendering use high-priority threads to approximate real-time systems. While the action that triggers an event is likely to occur in the haptic thread, processing the application's response to that event would cause severe problems. If haptic rendering were delayed even a few milliseconds while updating the application state to reflect the button press, it could cause sharp, intermittent increases in forces rendered to the user. Other input devices communicating at high rates could also produce jerky movements if the application state were updated in their update thread. Instead, a virtual button framework should cache any events until requested by the application's event processing thread. The virtual button feedback framework would then trigger any event when requested by the application. In an object-oriented paradigm, software that wishes to receive events could implement a delegate interface in one of their classes, and register that class to receive updates from a button. In a procedural programming paradigm, software could register callback functions for each button related event.

A virtual button framework should support a number of events, both for feedback and application developers. Actuation and deactuation events correspond to the switch closing and opening in a mechanical button. These events are likely to be important to the application developer who simply needs to know when the button was pressed. It is also important for implementing event-based feedback—indicating touched and released button events, as discussed in the previous section. The combination of these four events allows a feedback developer greater flexibility in designing event-based feedback during button actuation. Table 3.2 shows how these 3D events correspond to similar button events in 2D interfaces.

Table 3.2 Comparison of events in 3D and 2D environments

3D event	2D event
actuated	mouse down
deactuated	mouse up
touched	mouse over
released	mouse off

3.5 Result of implementing a software library based on the Virtual Button Framework

To test the framework presented in this chapter, we implemented a software library that generates button feedback. The software library was then tested by developing an extension for H3DAPI, an open-source scenegraph supporting haptic interaction (H3DAPI, 2012).

Figure 3.3 PHANToM force-feedback device

3.5.1 Setup

We developed the virtual button feedback library, and the haptic button node used to extend H3DAPI in C++. A PHANToM force-feedback haptic device, shown in Figure 3.3, was used to manipulate a stylus in the virtual workspace (Massie and Salisbury, 1994). H3DAPI provided the collision detection between the stylus and the button objects, and the virtual button feedback library was used to determine how the buttons responded to the interaction of the virtual stylus.

H3DAPI reads a scenegraph file from an X3D file, which is an XML format for defining 3D content. The scenegraph file defines nodes representing the objects in the scene, fields that store parameters of the nodes, and routes that make connections between fields. We wrote a prototype phone keypad with twelve virtual buttons, and used a Python script node to create the routes between buttons and a text node that updated with the number dialed. This virtual keypad scenegraph generated the 3D visual and haptic feedback using the software library we developed.

For prototyping auditory feedback from the virtual buttons, we used Pure Data (PD), a real-time processing environment with many sound generating tools (Pure Data, 2012). PD permitted the development of a visual network to process the actuation parameter into auditory feedback. The scenegraph sent continuous actuation parameter updates, and sent touched, actuated, deactuated, and released events to PD using the UDP network protocol. Based on the parameter and events, PD synthesized a real-time audio stream that played over computer speakers or headphones.

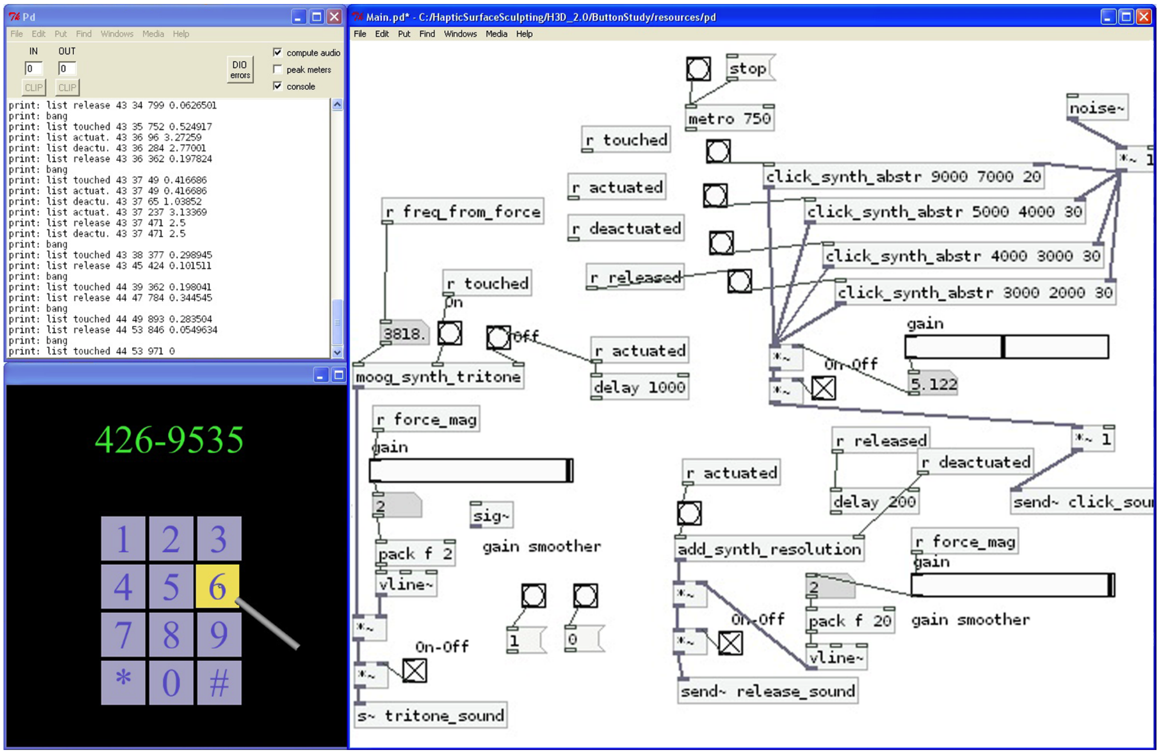

```
print: list release 43 34 799 0.0626501
print: bang
print: list touched 43 35 752 0.524917
print: list actuat. 43 36 96 3.27259
print: list deactu. 43 36 284 2.77001
print: list release 43 36 362 0.197824
print: bang
print: list touched 43 37 49 0.416686
print: list actuat. 43 37 49 0.416686
print: list deactu. 43 37 65 1.03852
print: list actuat. 43 37 237 3.13369
print: list release 43 37 471 2.5
print: list deactu. 43 37 471 2.5
print: bang
print: list touched 43 38 377 0.298945
print: list release 43 45 424 0.101511
print: bang
print: list touched 44 39 362 0.198041
print: list release 44 47 784 0.344545
print: bang
print: list touched 44 49 893 0.283504
print: list release 44 53 846 0.0549634
print: bang
print: list touched 44 53 971 0
```

Figure 3.4 An audio processing network in PD.

3.5.2 The Button Subgraph

The first challenge in implementing virtual buttons was to represent the button components in the scenegraph through a concise subgraph. The implementation required at least four components to represent the button in the scenegraph: the shape of the button key, the label text of the button, the shape of the shaft that bounds the button key as it travels, and a transform node which represents the travel distance of the button. While the user is pressing the button, the shaft enforces a boundary to prevent the stylus from slipping off the button. It is invisible in the implementation and only generates haptic force-feedback when the button is traveling.

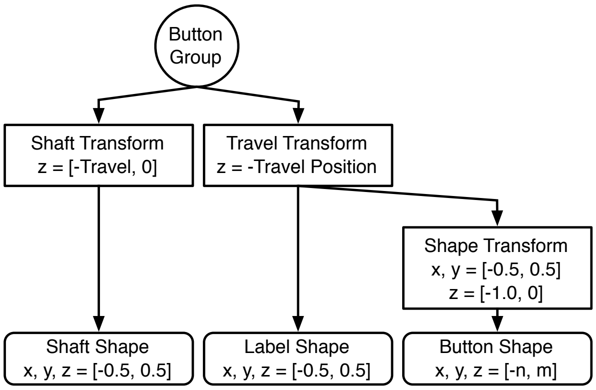

Figure 3.5 The sub-scenegraph used to represent a button.

The subgraph selected to represent the button shape added two additional transformation nodes, shown in Figure 3.5. The button shape can be an arbitrary size, so the first transform maps the button's coordinate space to a normalized one between [-0.5, 0.5]. The transform representing button travel then contains the button shape transform and the label text. A third shaft transform adjusted the shaft (assumed to have coordinates between [-0.5 and 0.5]) to the depth of the button's travel and the button's depth. The shaft transform and the travel transform were collected into a group to represent the subgraph of each button in the scenegraph. The implementer can then place a transform around this subgraph to adjust the button to the desired size and position in the virtual workspace, conforming to H3DAPI's convention.

3.5.3 Events

The events generated by the virtual button feedback library were condensed to conform to the event system of H3DAPI. The actuated and deactuated events were combined into a single *isActuated* H3DAPI field that sends a true or false event when the button actuates or deactuates respectively. The touched and released events were similarly combined into an *isTouched* H3DAPI field. These changes make the virtual button events more consistent with the conventions for events in H3DAPI.

The haptic thread from H3DAPI updates the travel parameter of the button to calculate any resulting force feedback. Events resulting from the updated position are cached until the next scenegraph update. H3DAPI uses a scenegraph traversal to update the values of nodes between rendered frames. Any events cached are triggered during this scenegraph update, and the state of the button is updated.

3.5.4 Feedback

The prototype includes buttons that generate visual, auditory, and haptic feedback when pressed. The button provides visual feedback from both the travel of the button, and from an increasing color change as it travels. When the user presses the button, it recedes away from the viewer as it travels. This effect is visible when rendered in stereo 3D, but the user can also perceive it on a 2D display as a relative change in button size. The color of the button also changes from neutral gray to bright yellow. The yellow is blended with the base gray according to the button's travel.

The button provides a combination of event-based and continuous auditory feedback through Pure Data (PD), shown in Figure 3.4. In PD, we initially rendered audio for each of the four events (touched, actuated, etc.) as distinct, predetermined notes. Each note corresponds to a pitched noise burst with an instant attack and preset decay time. We experimented assigning noise bursts with different apparent pitches, durations and characters to the four types of button events. While the short noise bursts could have simply played different prerecorded sounds, we used real-time subtractive synthesis methods that also take into account the user's applied force at the time of the event. We also experimented with expressing continuous changes of the actuation parameter over time as continuous audio signals.

Continuous auditory feedback uses a tension-relief scheme based on a classical music principle. As the user presses the button, a basic chord of five tritones continuously changes according to a transfer function, without changing the pitches of the chord. We currently use this sound only between contact and actuation while the haptic force gradually increases, generating a louder, more overtone-centric sound. This tritone chord is resolved into a major tritone chord as the actuation event is reached. Together, they are meant to build up tension and release in an auditory way that is comparable to the build up and release of haptic force feedback.

The prototype virtual button feedback library provides a force/travel function that outputs an increasing force, and click feeling similar to the one pictured in Figure 3.2. To compare the transfer function to the force rendered by the device, we graphed the position and force feedback detected by the fingertip in Figure 3.6. This graph shows a single button press with a long travel (4 cm). Positions beyond 4.0 cm were recorded because haptic force-feedback devices do not render hard surfaces with perfect stiffness. Users are typically unaware of this softness artifact unless their visual attention is focused on their hand, rather than the proxy representation in the virtual world, which does not move into the surface.

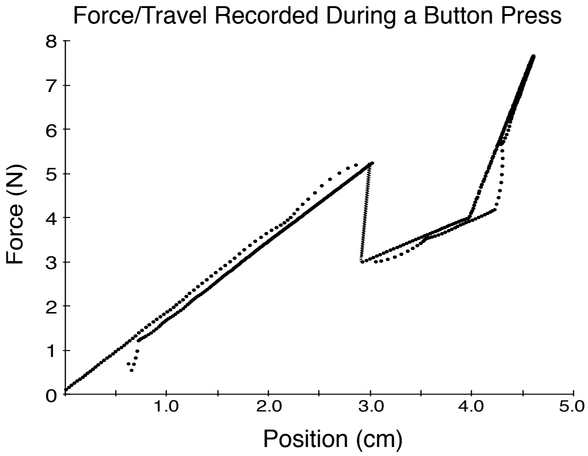

Figure 3.6 Haptic click experienced by a user of the prototype

3.5.5 Prototype Limitation

The prototype revealed an important consideration for future implementations. The sound feedback was prototyped in a higher-level programming language, a tradeoff to increase flexibility in prototyping at the expense of latency. As a result, a short lag was apparent in the prototype when pressing buttons that used event-based feedback for touched and actuated events in rapid succession. The feedback for touch was almost perceptible at the actuation event. There are two components to this lag. First, our H3DAPI prototype only sends audio parameter updates between graphic frames, which could result in up to a 25 ms delay. Second, events are sent over UDP to Pure Data for sound synthesis, which may also incur a smaller latency penalty. Future prototypes could send audio updates at a rate independent of the scenegraph update, and use a dedicated low-latency driver such as ASIO on Windows.

3.6 Development Outcome

This chapter presented a theoretical framework for virtual buttons to provide combinations of visual, auditory, and haptic feedback. These criteria were used in the development of a software library supporting virtual button feedback in 3D virtual environments. The software library was used in implementing a prototype virtual keypad by extending an open-source haptic scenegraph to support virtual buttons with more expressive feedback. The current software library development concentrated on travel position as a linear function of actuation parameter. Future implementations could incorporate a time-based actuation parameter, or a contact-time-to-travel mapping. In the future, further research could also characterize the design space of virtual button feedback and recommend types of feedback for each sensory modality.

CHAPTER 4. EMERGENT EFFECTS IN MULTIMODAL FEEDBACK FROM VIRTUAL BUTTONS

Modified from a paper submitted by Faeth, A. and Harding, C. to *ACM Transactions on Computer-Human Interaction (TOCHI)*. Submitted January 2012.

4.1 Abstract

This study investigated the effect of visual, auditory, and haptic sensory feedback modalities presented by a virtual button in a 3D environment on task performance (time on task and task errors) and user preference. Although we expected task performance to improve for conditions that combined two or three feedback modalities over a single modality, we instead found a significant emergent behavior that decreased performance in the trimodal condition. We found a significant increase in the number of presses when a user released the button before closing the virtual switch, suggesting that the combined visual, auditory, and haptic feedback led participants to prematurely believe they actuated a button. This suggests that in the design of virtual buttons, considering the effect of each feedback modality independently is not sufficient to predict performance, and unexpected effects may emerge when feedback modalities are combined.

4.2 Introduction

Virtual buttons have the potential to create more expressive interactions; the feedback they provide is not constrained by a mechanical design, and the button may adjust in response to the environment. For example, if a system detects that the user repeatedly presses the same button without completing the circuit, a virtual button could reduce the force used to resists pressure. Another type of virtual button might provide stronger visual feedback when noise levels in the room render auditory feedback inaudible. The designer of a virtual button may freely experiment with available feedback modalities to communicate both the state, and the transition between states, to the user.

Figure 4.1 Interacting with a 3D virtual button

A 3D virtual reality environment with force-feedback provides an interesting context for studying and prototyping virtual button feedback. Figure 4.1 illustrates a 3D force-feedback environment with a virtual button. In this environment, it is possible to isolate each feedback modality to examine it separately, and to study how multiple modalities of feedback combine and interact.

In the user interface design context, there is a common assumption that using a richer interface based on feedback from many sensory modalities is preferable to an interface for a single modality. In addition to the flexibility that multimodal feedback provides, the assumption is that increasing the number of feedback modalities increases performance. If this holds true, interfaces with bimodal or trimodal feedback should increase performance compared to unimodal interfaces.

However, before beginning to design more expressive interactions, we must first understand how virtual buttons communicate feedback to the user. Chapter 2 detailed previous work by researchers who attempted to characterize button feedback in mechanical keyboards and in touchscreen virtual buttons. There is, however, little research that investigates virtual button feedback when the limitations of flat surfaces and mechanical buttons are removed in a 3D environment. In this chapter, we present an overview of virtual button feedback, then describe a study that investigates how the presence of feedback modalities affects task performance, and finally discuss the results of our study and its potential design implications.

4.3 Feedback from Virtual Buttons

A virtual button generates feedback independent of its mechanical design. Feedback from a virtual button can occur in response to either an event such as the button traveling to a certain point, or a continuous interaction, such as force applied to the button.

4.3.1 Button Feedback Events

There are at least four events that comprise a button press: touch, actuation, deactuation, and release. When a user first makes contact with a button, touch signals the start of an interaction. The user continues to press the button until reaching the actuation point, which triggers the actuation event. The transition between touch and actuation is important because it allows the user to feel the button without triggering an actuation event, or to make a correction when realizing they mistakenly started to press the wrong button. As the user starts to release the button, it reaches the deactuation point, wherein a deactuation event signals the button is no longer active. The distance between actuation and deactuation points, or hysteresis, prevents multiple actuation-deactuation cycles from occurring during a single button press. The last event occurs when the user releases the button and loses physical contact. The buttons we used in this experiment provided auditory feedback at touched, actuated, and released events.

4.3.2 Continuous Actuation Feedback

The distance traveled by the button (travel) is another stimulus that can provide feedback to the user. For example, a system could use button travel to generate a resistive force through a transfer function, creating the feeling of a click (Lewis et al., 1997). The buttons used in this study generate both haptic and visual feedback from the button travel. Although the example above used button travel distance to illustrate button events, time and force parameters might also provide continuous feedback.

4.3.3 Visual Feedback

Visual feedback includes color or shape changes to the button itself. A change of color upon actuation is an overt example of visual button feedback. However, visual button feedback can also be subtle. As it travels, the movement provides a visual feedback cue that communicates how far the user has pressed the button, without the user feeling any resistance or friction. We chose to focus this study on the feedback the button itself provides, rather than including visual feedback from an ancillary source such as the partially-dialed number displayed on the LCD of a phone.

4.3.4 Auditory Feedback

One prevalent form of auditory feedback is a sound accompanying the actuation of a button. This sound might simulate the click sound of a switch closing on a mechanical button. Another common form of auditory feedback is a tone played upon actuation, such as the unique tone produced by each key on a telephone. Buttons may also generate auditory feedback to indicate how far a button has traveled, or express the force applied to it with an audio cue.

4.3.5 Haptic Feedback

One of the most recognizable forms of haptic feedback from buttons is a click feeling produced by increasing force until the user presses the button to the actuation point, when a sharp drop in force creates the click (Lewis et al., 1997). Instead of force feedback, buttons might generate feedback by modifying the frequency and amplitude of a sine wave to produce vibrotactile feedback.

4.4 Methods

The purpose of the study was to examine the contribution of visual, auditory, and haptic feedback modalities to the prevention of errors and to the time required to press buttons in a

3D virtual environment. The study only varied the presence or absence of each modality of feedback, the characteristics of the feedback were consistent between conditions. We designed the study to test the two research questions listed below:

- RQ1. Will user performance improve as the number of feedback modalities from virtual buttons increases?
- RQ3. Will participants prefer virtual buttons that provide feedback using a higher number of modalities?

4.4.1 Participants

Twenty participants (12 males, 8 females) between the ages of 19 and 59 (median age 26) volunteered to participate in the study. None of the participants reported previous experience with haptic devices, and all participants reported using a computer several times throughout the day. Participants also reported normal (corrected) vision and hearing. Eight participants reported using a device with a stylus, such as a PDA, tablet computer, or drawing tablet at least once a week. Participants received $10 USD in compensation for their time.

4.4.2 Equipment

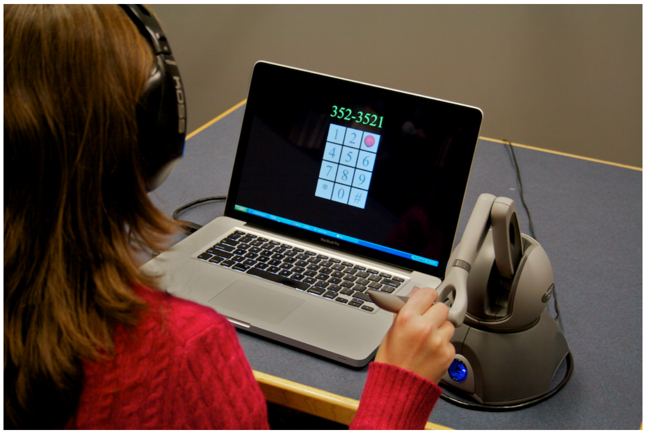

Figure 4.2 The experiment setup with PHANToM Omni.

Figure 4.2 shows the equipment used in the study. Participants held the stylus of a PHANToM Omni in their dominant hand to provide haptic force-feedback (Massie and Salisbury, 1994). A laptop displayed the virtual keypad and the stylus proxy of the haptic device. The participant also used the laptop's spacebar with their non-dominant hand to advance to the next task. This interaction followed Guiard's Kinematic Chain model: the non-dominant hand performed the macrometric role advancing the task, while the participant's dominant hand performed the micrometric role dialing the number (Guiard, 1987). The participants heard auditory feedback through a pair of closed headphones. We created the virtual buttons in H3DAPI, an open source scenegraph for haptic rendering (H3DAPI, 2012). Auditory feedback was created using Pure Data (2012).

4.4.3 Phone Number Dialing Task

We asked participants to dial seven-digit phone numbers using a PHANToM Omni, which provided haptic force-feedback and served as a 3D input device. Results of keyboard performance research summarized by Lewis et al. (1997) suggested the standard telephone layout is perhaps the most familiar pattern to participants. Lee and Zhai (2009) also reported that Bill Buxton suggested the phone number dialing task for their study. The software generated random phone numbers by picking a random starting digit and then moving to a random neighboring digit on the keypad. This limited the distance a stylus could travel between digits of the random phone numbers to adjacent digits.

4.4.4 Design

The study employed a within-subjects design. We varied which of the seven combinations of visual, auditory, and haptic feedback modalities a virtual keypad would present to the participant (Table 4.1). We randomly assigned the order in which the system presented conditions to each participant to mitigate learning effects. Participants repeated the tasks five times. To alleviate the effect of fatigue, all participants completed the study in 30 minutes, and participants were invited to take a break after training.

Table 4.1 Experiment conditions

Condition	**Abbreviation**
Visual	V
Auditory	A
Haptic	H
Visual + Auditory	VA
Visual + Haptic	VH
Auditory + Haptic	AH
Visual + Auditory + Haptic	VAH

4.4.5 Virtual Buttons

Virtual buttons were approximately 3.0 cm square on the screen and 2.5 mm thick with a 1.5 cm travel. We chose the travel distance because participants in the pilot study felt the feedback using a shorter travel distance was too jarring. This was likely caused by the haptic rendering compensating for the fast movements of the participants, and we extended the travel distance of the buttons to generate smoother feedback.

4.4.6 Feedback Generation

The method used to generate each feedback modality was consistent between conditions, we only varied the presence or absence of each modality. We used a force/position transfer function to generate a click sensation for the haptic feedback (Lewis et al., 1997). As the button traveled, the force gradually increased until sharply falling off at the actuation point, at which time a cushioning force rose until the button stopped at the end of its travel (Figure 4.3). Visual feedback employed an interpolation between the neutral gray of the keypad button and yellow, a color that contrasted well against the black background. Both haptic and visual feedback continued while the button was being pressed.

To signal initial contact, actuation, and release via audio cues, we experimented with simple noise-based sound schemes and with more complex, pitched (musical) tones—including a tension-release chord progression. Using real-time synthesis, rather than pre-recorded sound files, allowed us to change certain properties of the sound as a real-time response from the participant's input, similar to playing a virtual instrument. However, the more complex sound

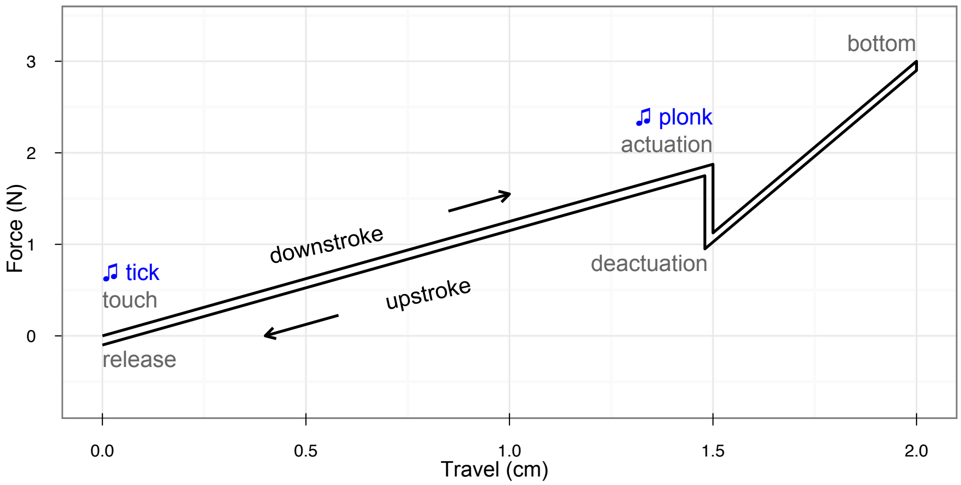

Figure 4.3 Transfer function to generate a haptic click

schemes failed to distinctly communicate actuation if several buttons were pressed in quick succession. Hence, we chose to reduce the auditory feedback to a very short *tick* sound at the initial contact, followed by a quickly decaying *plonk* sound, mimicking the sound of a dropped ball coming to rest. The *tick* sound lasted 30 ms, and the *plonk* sound decayed over a duration of 300 ms.

4.4.7 Measures

For each task, we assessed three components: task completion time, the number of errors made while dialing the phone number, and the participant's subjective rating of the button's feedback. For each task, we captured the times that the user touched, actuated, deactuated, and released each button press. The task completion time was measured from the time the participant first touched any button on the keypad until the time they released the last button on the keypad. During each task, we also counted the number of times that the user touched and released a button without actuating it, which we called the number of unactuated presses during the task.

To calculate a fair measure of overall error, we used the Wagner-Fisher string matching algorithm to compare the phone number assigned by the task to the digits entered by the participant (Wagner and Fischer, 1974). The algorithm calculated the minimum number of insert, delete, and replace operations required to transform one sequence into another. For example, changing 123-6547 sequence given to the participant into the 123-564 sequence entered by the participant requires replacing the 65 sequence with a 56 sequence, and deleting the 7, which results in two replacements and one deletion. We counted the sum of inserted, replaced, and deleted digits to obtain the total number of errors made during the task. This total is the Levenshtein distance between the two sequences (Levenshtein, 1965).

After each task, we asked participants to provide a subjective rating of the quality of feedback provided by the virtual buttons. We asked them to provide a response to the following statement on a 7-point Likert scale.

> *The previous buttons gave me enough feedback when I pressed them.*

The participant recorded the Likert rating on another virtual keypad that used the same combination of feedback modalities as the task (Figure 4.4). We asked for this subjective rating directly after each task because we anticipated that participants would have difficulty recalling and rating each combination of feedback modalities after all of the trials. Participants could not see if they dialed the number correctly until they provided the rating to prevent their task completion success from influencing their subjective impression of the combined effectiveness of the modalities available during the task.

4.4.8 Procedure

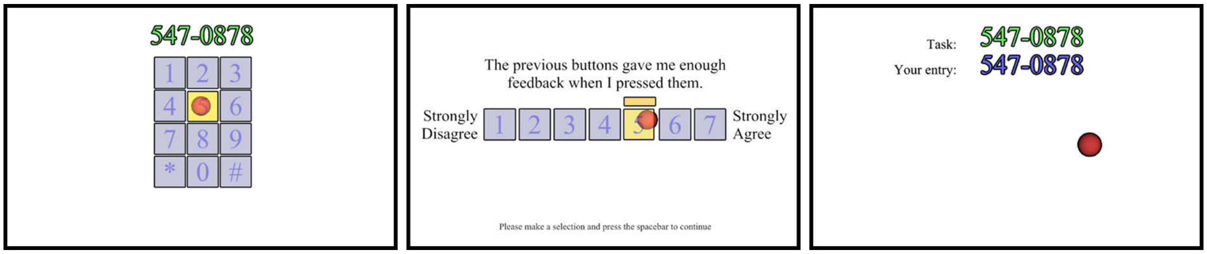

The previous buttons gave me enough feedback when I pressed them.

Please make a selection and press the spacebar to continue

Figure 4.4 Task sequence: The participant enters the number on the keypad (left), rates the button feedback (center), and views the task result (right).

Wherever possible, we kept the test environment consistent among participants. The same investigator interacted with participants in the same room, using the same materials, and a pre-written script.

After introducing the participants to the study and the equipment, we asked them to fill out a short background questionnaire. Participants then had a chance to practice dialing phone numbers under each condition three times. As shown in Figure 4.4, the software presented each phone number to the participant on the screen above the keypad, however the digits recorded by the system were not visible while the participant dialed. When the participant signaled the completion of the task by pressing the spacebar on the laptop, we displayed a second screen asking them to provide a subjective rating of the buttons. When the participant pressed the spacebar again, we displayed a third screen with the phone number assigned by the task and the phone number recorded by the buttons. After completing the training, participants performed five repetitions of the conditions. At the conclusion of the study, we gave participants an exit survey.

It is worth noting that participants only saw the phone number actually dialed after completing the task, and after giving their subjective rating of the feedback. This made sure that they received only direct information from the buttons themselves and not from ancillary sources such as the partially dialed number displayed on the monitor.

4.4.9 Controlling for Visual Feedback

Three of the seven conditions excluded visual feedback. Although a simple approach to controlling for visual feedback might not render any additional feedback (e.g., not flash a color upon actuation), it would be unsuitable for our study. The participant may still have received a subtle visual cue from the changing perspective and size of the button as it traveled away, even with reduced depth perception inherent in a monoscopic display. Likewise, it would be impractical to turn off all visual feedback; participants depended on their vision to locate buttons. Since this study is concerned with evaluating feedback rather than a search, we chose to show (visually render) the stylus and the keypad buttons even for conditions that lacked visual feedback; thereby enabling participants to locate the desired button on the virtual keypad and start pressing it.

Instead, during conditions not involving visual feedback, we did not show the button moving downward; we employed a perceptual illusion found in haptic renderings. The stylus graphic stops moving at the surface of the button, however, the participant's hand physically continues to move down through the stroke of the button, generating auditory or haptic feedback. The static visual feedback overrides the participant's proprioceptive sense of actual motion, allowing the participant to press a button with no visual feedback during the travel. Participants remained unaware of this sensory contradiction; we received no comments about this illusion from the participants.

4.5 Pilot Study

The pilot study exposed two issues we addressed before conducting the full study. The first issue was that the model representing the PHANToM stylus in the virtual environment afforded the ability to press the button with a cylinder, but force feedback was only calculated at the tip of the cylinder. We started the pilot study using the default stylus model included with H3DAPI. H3DAPI's stylus has a long cylindrical handle with a very small ball floating at the point where haptic interactions are calculated, shown in Figure 4.5.

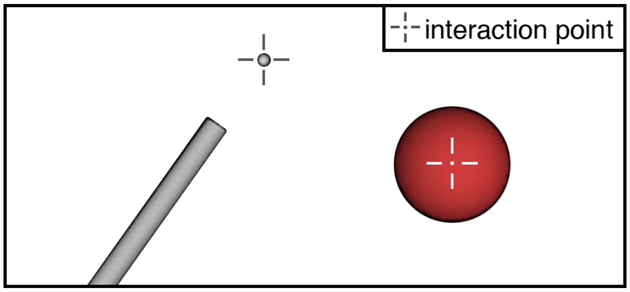

Figure 4.5 Default H3DAPI Stylus (left) and round stylus (right)

We used the OpenHaptics renderer in H3DAPI, which is a point-based haptic renderer (H3DAPI, 2012). Point-based haptic renderers calculate haptic collisions with objects in the virtual scene from a single point, rather than collision with a sphere, or full collision detection of the stylus geometry. With the cylindrical stylus, participants attempted to push the button with parts of the stylus handle that cannot provide haptic feedback. To prevent this mismatch between the mental model of the participant and the haptic renderer, we changed the stylus model to a spherical shape. The sphere is approximately 70% of the width of a button, red in color and semi-transparent. We did not observe any further confusion about how the stylus interacted with the button after changing the stylus shape.

The second issue raised by the pilot study was the type of button feedback used when recording the response on the Likert scale. We decided that the participant should keep one hand on the PHANToM stylus to minimize disruption of flow during the study. This permits the participant to keep their free hand on the physical spacebar to advance to the next task. In order to keep one hand on the PHANToM stylus and one hand on the task advancement key, the participant provided their subjective rating using the PHANToM. We created another virtual keypad for this subjective rating (Figure 4.4).

We initially thought that the virtual buttons for the subjective rating should use a consistent set of modalities throughout the study, and used buttons which provided VAH feedback. However, two of our pilot study participants commented that they thought we were encouraging them to assign higher ratings to buttons employing more feedback modalities as a result of this design decision. We changed the design so that the subjective rating keypad provided the same feedback modalities as the phone keypad from the preceding entry task.

4.6 Results

The study manipulated the presence or absence of each feedback modality during the phone-number dialing task, resulting in a single explanatory variable with seven levels: V, A, H, VA, VH, AH, and VAH. We measured three response variables: the number of errors during the task, the time to complete each task, and the participant's subjective rating of the button feedback. We combined the 700 task results from the within-subjects design. Twenty participants performed five repetitions of the phone-number-dialing task for each of the seven conditions. We compared the odds of completing the task successfully in each condition using a logistic regression. We also used analysis of variance to understand how the feedback type contributed to our observations.

We measured the number of errors committed during each task by counting the differences between the phone number assigned in the task and the number dialed by the participant (Table 4.2). Figure 4.6 shows the task results distributed over each condition and the number of errors made during each task. The number of errors during the VAH condition stands out when assessing the distribution of errors over each condition. The graph shows a comparatively large number of tasks with one or more errors during the VAH condition, far surpassing all other feedback modalities for tasks with two or more errors. Although we expected the number of errors to decrease or remain the same as the number of feedback modalities increased (e.g. V → VA → VAH), we instead found the number of errors increased in the trimodal condition.

4.6.1 Measuring Unactuated Presses

To further investigate the unintuitive increase in the number of errors in the trimodal condition, we examined how these errors occurred. Many of the phone numbers entered in the trimodal condition were shorter than seven digits, and appeared shorter than other conditions. This suggested that the errors could be a result of participants skipping digits rather than erroneously inserting extra digits in the phone number. To support this observation, we examined the number of unactuated presses in each task. Unactuated presses occurred when the participant touched a button, but then released it before pressing far enough to actuate the button. We measured the number of unactuated presses during each task by counting the number of presses where the participant touched and released the button without actuating it.

A Pearson's R test suggests a correlation between unactuated presses during a task and errors in the task ($R=0.76$, $p<0.001$). However, unactuated presses are not necessarily errors. If the participant touched the wrong button but realized the mistake before actuation, the correction would also appear as an unactuated press even though it was not counted as an error. If the unactuated presses largely occurred as a result of these error corrections, we

Table 4.2 Number of errors or unacutated presses within a single task

	0	1	2	3	4	5	6	7	8
Tasks with Errors	560	94	21	10	5	8	2	0	0
Tasks with Unactuated Presses	607	62	13	6	3	4	1	3	1

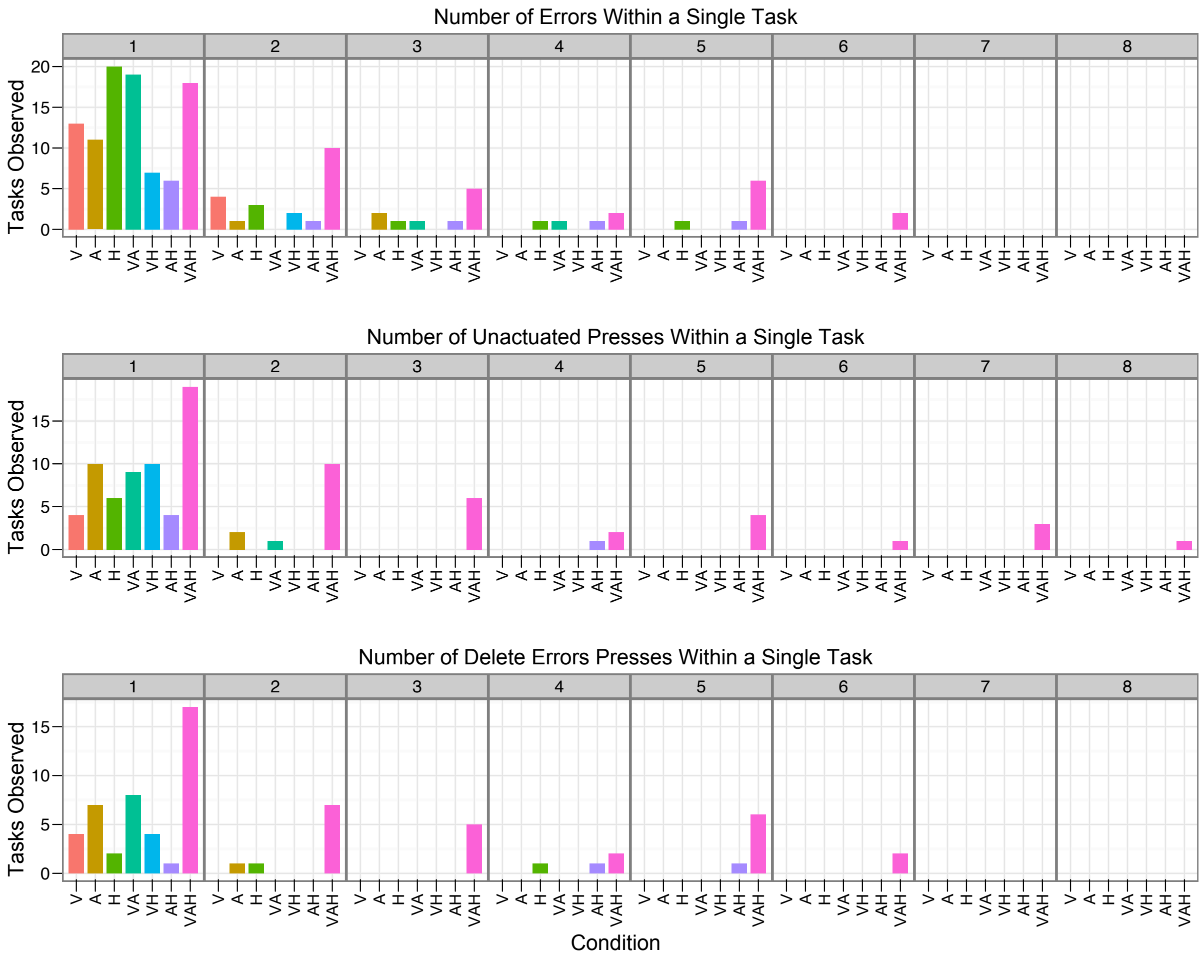
Number of Errors Within a Single Task

Number of Unactuated Presses Within a Single Task

Number of Delete Errors Presses Within a Single Task

Figure 4.6 Number of tasks with errors when dialing seven-digit numbers in each condition.

would expect them to be normally distributed across all conditions (Figure 4.6). Instead, the unactuated presses occurred more frequently in the trimodal condition (M=1.20, SD=1.87).

4.6.2 Success Between Modalities

If the participant did not make an error while completing a phone number dialing task, that task was considered a success. A logistic regression was performed to investigate the effect of feedback modalities on the dichotomous outcome of task success. The regression model used the condition as the explanatory variable and task success as the response variable (Equation 4.1). Each of the betas in the equation corresponded to an estimate of the log odds ratio between two conditions in Table 4.3. We obtained the odds ratio in Table 4.3 by finding the exponent of the log odds ratio.

$$logit(\pi) = \beta_0 V + \beta_1 A + \beta_2 H + \beta_3 VA + \beta_4 VH + \beta_5 AH + \beta_6 VAH \tag{4.1}$$

The goodness-of-fit tests for a logistic regression only compare the fit of the model relative to another model. We compared this model to a second one using the participant as the explanatory variable. A chi-square test of goodness-of-fit found that the model with condition as the explanatory variable had better fit (χ^2(6, N=700) = 48.80, p<0.001) than the model which used the participant to predict task success (χ^2(19, N=700) = 36.31, p=0.010).

Table 4.3 Estimation of mean accuracy with logit regression

Condition	**Est.**	**Std. Err.**	**Z**	**Pr** $> \lvert Z \rvert$	**Odds ratio**	**95% Conf. Limits**	
V	1.59	0.27	5.96	<0.001	–	–	–
A	0.23	0.39	0.59	0.558	1.26	0.58	2.75
H	-0.54	0.35	-1.54	0.124	0.58	0.29	1.15
VA	-0.26	0.36	-0.72	0.472	0.77	0.38	1.56
VH	0.73	0.44	1.66	0.097	2.07	0.89	5.1
AH	0.61	0.43	1.43	0.152	1.84	0.81	4.39
VAH	-1.30	0.33	-3.90	<0.001	**0.27**	0.14	0.51

Fit: Log Likelihood: -325.88 df=7. AIC: 665.76.

From the estimates in Table 4.3, we also calculated the odds that the participant would succeed at the task for each modality compared to visual feedback alone (the odds ratio). Although we expected a higher rate of success as the number of feedback modalities increased, the results suggest participants were almost 4 times more likely to succeed when using visual feedback alone, compared to the trimodal condition (*OR*=0.27, p<0.001).

With a categorical explanatory variable, the logistic regression uses one condition as the baseline for the comparison, called the intercept. The regression in Table 4.3 uses the visual condition for the intercept. We performed another six logistic regressions using each of the other conditions as an intercept. Table 4.4 shows the observed odds of success and the odds ratios for each of the seven conditions. The results show the odds of success improved when adding another feedback modality to haptic feedback (VH *OR*=3.55, p<0.01; AH *OR*=3.16,

Table 4.4 Odds ratios by condition

			Odds Ratio						
Intercept	Successes	Failures	V	A	H	VA	VH	AH	VAH
V	83	17	–	1.26	0.58	0.77	2.07	1.84	**0.27*****
A	86	14		–	0.46*	0.61	1.65	1.47	0.22***
H	74	26			–	1.32	**3.55****	**3.16****	0.47*
VA	79	21				–	**2.69***	**2.39***	0.35**
VH	91	9					–	0.89	0.13***
AH	90	10						–	0.15***
VAH	57	43							–

Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$

$p<0.01$). Looking at the bimodal feedback conditions, the results suggest the odds of success improve when switching from VA feedback to either VH (OR=2.69, $p<0.05$), or AH modalities (OR=2.39, $p<0.05$). Looking at the last row of Table 4.4, the results also show a significant increase in the odds of success from trimodal feedback and every other condition.

4.6.3 Effect of Feedback Type

We examined the correlation graph and performed a Pearson's R test to detect interactions between the response variables of time, and errors. Although we expected that the number of errors might increase task completion time, the correlation between errors and time was weak ($R=0.13$, $p<0.001$). Since we found no support for correlation between the response variables, we treated their analysis separately.

Table 4.5 summarizes the effect of the feedback modalities present in each condition on the response variables of time, errors, and subjective rating. The table also includes the effects of the condition on unactuated presses.

4.6.4 Errors and Unactuated Presses

Although we expected errors to decrease in bimodal and trimodal conditions, the highest frequency of errors occurred in the trimodal feedback condition ($M=1.03$, $SD=1.59$). A repeated-measures ANOVA yielded a significant effect of the condition on errors ($F_{6,114}=13.78$,

Table 4.5 Effect of feedback modalities on performance

		V	A	H	VA	VH	AH	VAH	F-value	p-value
Errors	Mean	0.21	0.19	0.38	0.26	0.11	0.20	1.03	13.78	<0.001
	SD	0.50	0.54	0.83	0.61	0.37	0.75	1.59	$F_{6,114}$	
Unactuated Presses	Mean	0.04	0.14	0.06	0.11	0.10	0.08	1.20	30.03	<0.001
	SD	0.20	0.40	0.24	0.35	0.30	0.44	1.87	$F_{6,114}$	
Time	Mean	9.89s	9.29s	9.46s	9.24s	8.73s	7.56s	9.19s	4.14	0.001
	SD	4.09s	3.97s	4.08s	4.17s	6.74s	3.03s	5.75s	$F_{6,114}$	
Subjective Rating	Mean	3.40	3.90	3.84	4.39	5.46	5.09	5.85	33.10	<0.001
	SD	1.61	1.53	1.49	1.60	1.16	1.33	1.27	$F_{6,114}$	

$p<0.001$). A post hoc Tukey HSD showed significant pairwise differences between the trimodal condition and every other condition ($\alpha = 0.05$). There were no significant differences between the other pairs of conditions.

The results did not show an improvement in the mean errors for the bimodal conditions compared to unimodal conditions. However, the results from Table 4.4 do suggest an improvement in task success for AH feedback compared to the haptic condition. This suggests that participants succeeded on more tasks with AH feedback, but when they failed, they made more severe errors during each task.

The highest mean number of errors occurred in the trimodal feedback condition. One possible explanation for the counterintuitive increase in errors under the trimodal condition was that the participants began pressing the correct button, but released it before the button traveled to the actuation point. We counted every time the user touched a button without actuating it during a task, and labeled such a movement an unactuated press.

The highest mean of unactuated presses during each task also occurred during the trimodal condition (M=1.2, SD=1.87). A repeated-measures ANOVA yielded a significant effect of condition on unactuated presses ($F_{6,114}=30.03$, $p<0.001$). A post hoc Tukey HSD again showed significant pairwise differences between the trimodal condition and every other condition ($\alpha = 0.05$). There were no other significant differences in the mean unactuated presses between the other pairs of conditions.

4.6.5 Time on Task

We investigated the effect of condition on Time on task as another indicator of performance. Participants were told to perform the tasks at their best, comfortable pace while still maintaining accuracy. Ideally, a short time would indicate this type of feedback was easy to use and enabled the participant to perform the task quickly. However, a short time might also indicate the task was performed hastily and the participant had skipped over digits in the task, increasing errors. Instead, a Pearson's R test did not show a negative correlation between completion time and errors ($R=0.13$, $p<0.001$). The results also showed a positive correlation between completion time and unactuated presses ($R=0.21$, $p<0.001$), which suggests that unactuated presses did not shorten completion times.

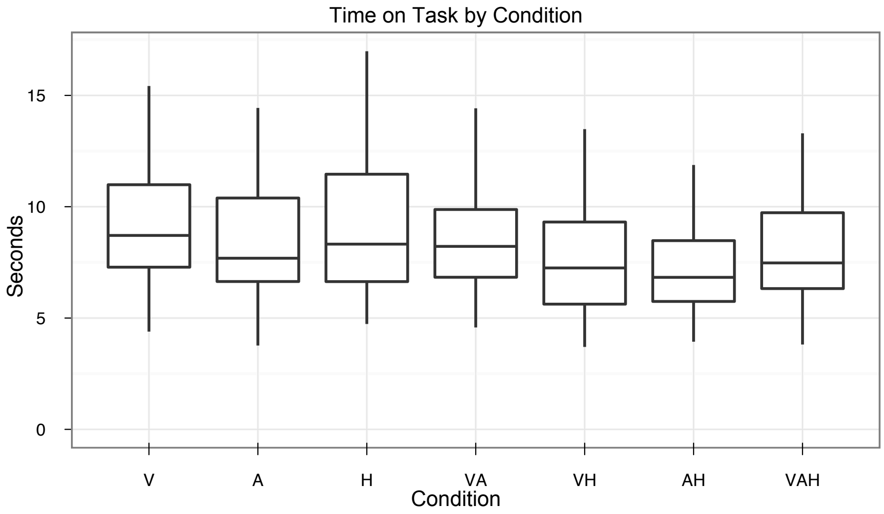

Figure 4.7 Median time on task for each condition

Figure 4.7 shows the effect of feedback modalities on completion time. Completion times were shortest in the AH feedback condition (M=7.56s, SD=3.03s), and longest for the visual condition (M=9.89s, SD=4.09s). A repeated-measures ANOVA showed that the condition had a significant effect on completion times ($F_{6,114}=4.14$, $p=0.001$). A post hoc Tukey HSD showed significant pairwise differences between the AH condition and every condition except the VH condition ($\alpha = 0.05$). There were no significant differences in the mean time on task between the other pairs of conditions.

4.6.6 Subjective Rating

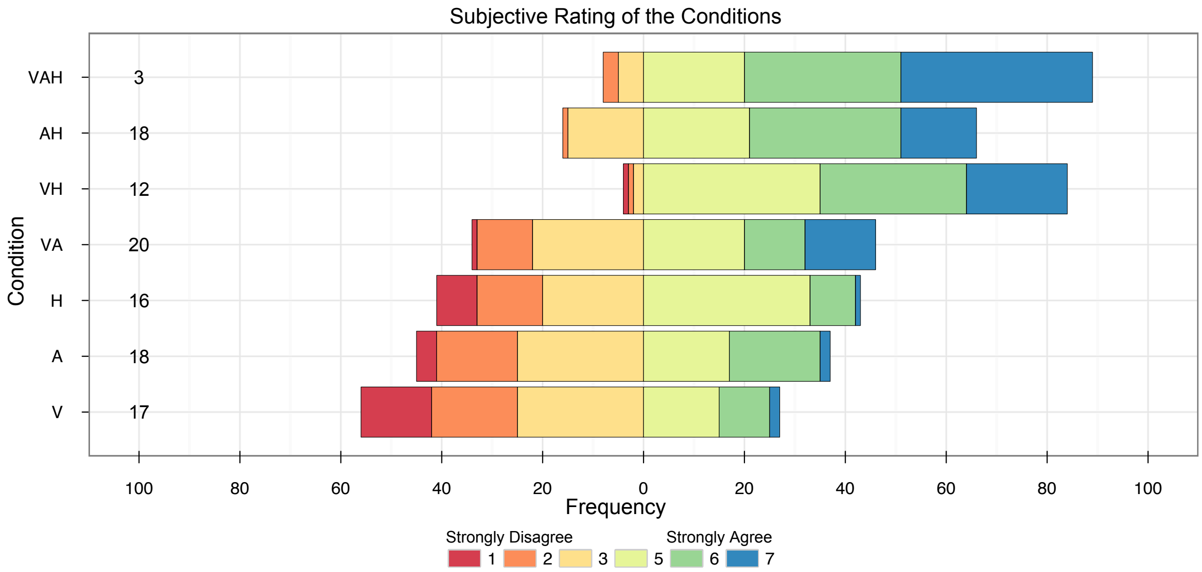

Figure 4.8 Subjective rating of each condition (neutral responses on the left).

After each task, participants were asked to respond to the following question on a 7-point Likert scale: "The previous buttons gave me enough feedback when I pressed them." We collected 100 responses for each condition. Figure 4.8 shows how strongly participants agreed or disagreed with this statement.

Table 4.6 Subjective preferences (Tukey HSD)

Preference		Over Conditions
VAH	>	{ V, A, H, VA, AH }
VH	>	{ V, A, H, VA }
AH	>	{ V, A, H, VA }
VA	>	{ V }

A repeated-measures ANOVA yielded a significant effect of the condition on subjective rating ($F_{6,114}=33.10$, $p<0.001$). A post hoc Tukey HSD showed significant differences between the pairs of conditions in Table 4.6. The differences between pairs of conditions which do not appear were not significant ($\alpha = 0.05$). The condition with the highest mean Likert rating was the trimodal feedback (M=5.85, SD=1.27). The condition with the lowest mean rating was visual feedback (M=3.40, SD=1.61). A Pearson's R test suggests a modest positive correlation between the number of feedback modalities (1, 2, or 3) and subjective rating (R=0.47, p<0.001).

4.6.7 Participants

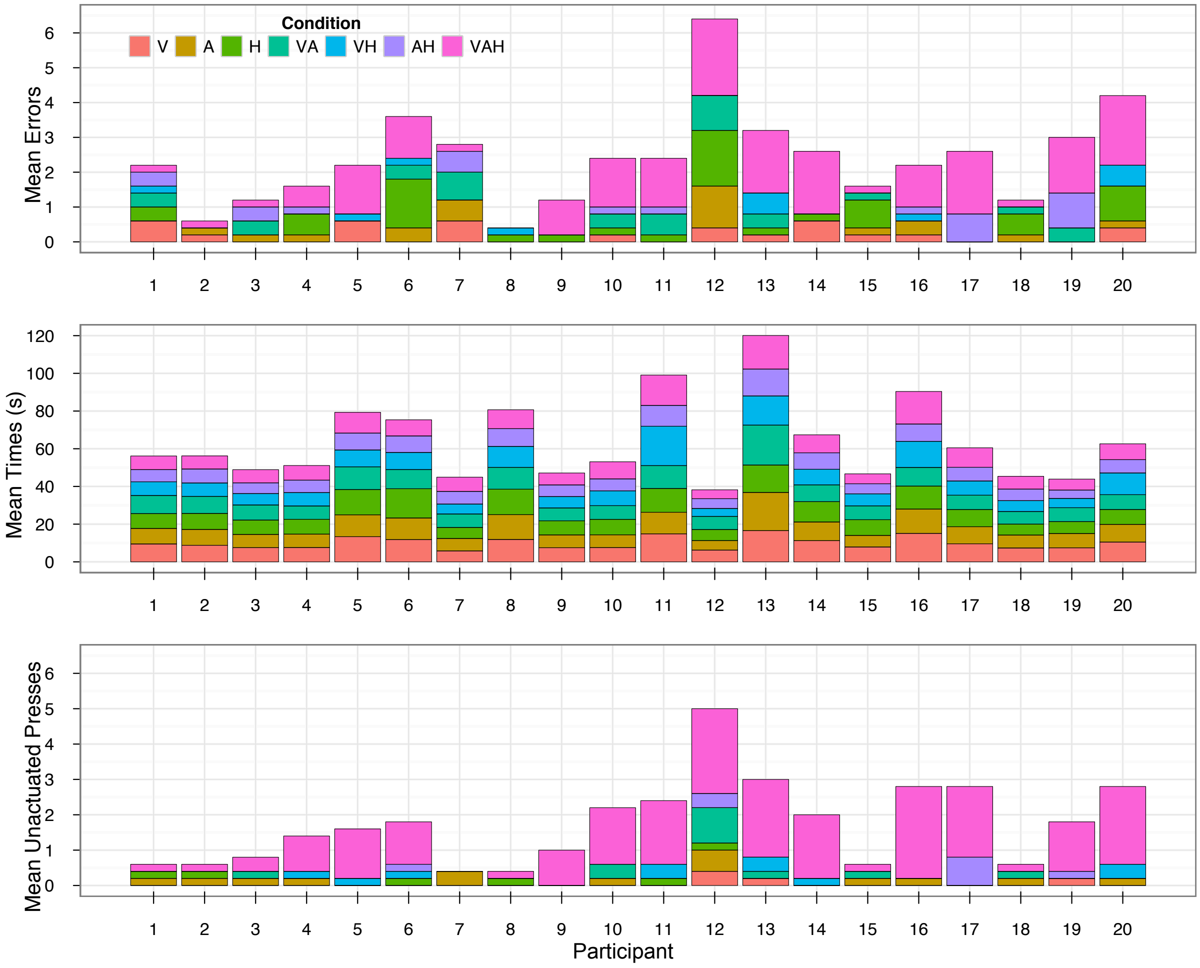

Figure 4.9 Mean performance by each participant during each condition.

We investigated the degree of variation between each participant's performance. Several trends are visible in the participant summary graphs shown in Figure 4.9. These graphs present the mean value of the repeated tasks for each participant during each condition. The participant with the highest mean errors across all conditions also had the shortest mean completion times

(Participant 12). However, the participants with next highest mean errors had a higher overall mean completion time (Participants 6 and 20). There is a clear trend of a higher number of errors occurring during the trimodal condition. There is also a similarity between the mean number of errors and the mean number of unactuated presses in the trimodal condition shown in the graph.

4.7 Discussion

The study explored whether the use of visual, auditory, and haptic feedback modalities had any effect on completion time and errors made by participants pressing virtual buttons in a 3D environment, and which modalities participants preferred. The results show a statistically significant effect on the set of feedback modalities used to communicate button feedback on participant errors, and on task completion time. We also found participants expressed a statistically significant preference for specific sets of feedback modalities over other sets. The results also offer further insight into unintuitive effects of feedback modalities on virtual button feedback.

4.7.1 Task Performance

> RQ1. Will user performance improve as the number of feedback modalities from virtual buttons increases?

Prior research suggests that multimodal interfaces improve both task completion time and reduce errors because the user distributed work across modalities (Oviatt, 2007; Burke et al., 2006). The feedback modalities employed by a virtual button had a statistically significant effect on both completion time and the number of errors made by participants. Although task completion time was similar among unimodal conditions, the results showed a decrease in completion time for trimodal condition compared to unimodal conditions. If the button did not present haptic feedback, the participant did not feel the “bottom” of the button press, and often continued to move the stylus past the button’s maximum travel distance. The decrease in completion time for these conditions containing haptic feedback paired with another modality suggests haptic resistance helped participants move onto the next digit more quickly than with other conditions. However, haptic feedback alone did not appear to provide this benefit, and completion times also slowed in the trimodal condition.

Of the unimodal conditions, it was not surprising that haptic feedback alone produced the highest number of errors because none of our participants had prior experience with haptic force-feedback. We did expect to see a decrease in errors when adding a second modality of feedback, and found that haptic feedback benefited from the addition of a second modality of feedback.

However, the significant increase in the number of errors in the trimodal feedback condition was unexpected. Further investigation suggested that participants left off more digits when dialing phone numbers in the trimodal condition. The results showed a statistically significant difference in the unactuated presses occurring during this condition. This suggests that participants started to press the correct digits, but released the button before pressing it far enough to actuate and record the digit.

One alternative hypothesis is that users were hampered by the feedback of a particular modality or by a lag between modalities. If this were the case, we would expect to see an increase in errors when a particular modality was present, or a particular combination of modalities in one of the bimodal conditions. However, the results do not show any single modality resulted in an increase in errors, or that any bimodal feedback increased errors.

Another alternative hypothesis is that multimodal feedback increased the cognitive load, causing participants to make more mistakes in the trimodal feedback condition. If this were true, we would expect an increase in errors during the trimodal feedback condition. However, we would expect these errors to be distributed over insert, delete and replace errors, and we would not expect a disproportionate increase in the number of unactuated presses. Neither alternative explains the significant increase in unactuated presses observed in the trimodal condition.

Therefore, we hypothesize that the combined visual, auditory, and haptic feedback gave participants a false degree of confidence that they actuated the button before reaching the actuation point. The same feedback that was effective for communicating button actuation singly and in pairs of modalities displayed an emergent behavior: when presented in a combination of visual, auditory and haptic modalities, it produced a signal that was stronger than the pairs of signals would suggest. This led participants to release the button before reaching the actuation point, believing they had already actuated the button.

The results do not suggest the increase in errors during the trimodal condition is inconsistent with multiple resource theory. If there were contention for common resources to process trimodal feedback, then we might expect the participant to continue pressing the button until processing the feedback, or to repeat the button press a second time. We would not expect the participant to release the button early after processing the feedback from the button travel. Since participants released the button early in the trimodal condition, the results indicate that errors were not caused by contention for shared resources to process trimodal feedback. Therefore, the results are consistent with multiple resource theory.

4.7.2 Subjective Preferences

> RQ3. Will participants prefer virtual buttons that provide feedback using a higher number of modalities?

The results of the study provide evidence that participants do express a preference for some combinations of feedback modalities over other feedback modalities provided by a virtual button. However, the results also suggest a disconnect between the feedback that participants subjectively preferred, and the feedback that enabled them to perform tasks well. Participants expressed a preference for the trimodal feedback condition, yet that condition produced the highest number of errors, and only produced average completion times. Since we asked participants to subjectively rate feedback type before revealing how well they performed on each task, it is not surprising to find this disconnect. However it is surprising that participants did not grow more frustrated with this condition over time.

One possibility is that participants were not able to distinguish the trimodal condition from bimodal conditions. Research into cross-modal perception suggests that in the presence of signals employing several modalities, participants are more unlikely to notice the audio or haptic signal (Colavita, 1974; Hecht and Reiner, 2009). It is possible that the participants could not distinguish the trimodal condition from similar conditions when they were successful, and may not have correlated a task's sensory experience with their prior errors.

4.8 Design Implications

When designing multimodal feedback, the results of the study suggest that feedback can potentially be stronger when using visual, auditory, and haptic feedback together than the feedback from pairs of modalities would suggest. The emergent effect observed in this study had a negative effect on the task, which caused the user to release the button before actuation. However, designers might also find this effect useful in situations where stronger feedback is desired.

The results also suggest some combinations of feedback were more successful than others. The odds of success improved when using auditory feedback alone instead of haptic feedback. Adding haptic feedback to visual or auditory feedback also increased the odds of success. Although the visual and auditory modalities are often paired together in practice, our results suggest that the odds of success improved when visual and haptic feedback are paired instead.

4.9 Conclusion and Future Work

This study examined what effect the different combinations of feedback modalities from 3D virtual button presses had on a user's task performance and user preference. We designed an experiment that allowed us to selectively isolate visual, auditory, and haptic feedback from a series of 3D virtual button presses and to test the task performance of these feedback modalities. The results suggest that participants expressed a significant preference for trimodal feedback over bimodal feedback, and for bimodal feedback over unimodal feedback conditions. We also found that the combination of feedback modalities presented by the virtual button has a significant effect on both completion time and errors, although time and errors do not necessarily decrease in conditions that presented more feedback modalities. Instead, we found a significant emergent behavior in button actuation feedback. Visual, auditory, and haptic feedback that was effective in unimodal and bimodal conditions caused participants to make more errors when combined in the trimodal condition; in the latter they released buttons before actuation. This suggests that virtual buttons presenting trimodal feedback gave the participants a false sense of confidence when actuating them.

As designers explore continuous multimodal feedback to convey transitions between events, we expect that emergent properties, such as the unintuitive effect we observed in this study, may play a larger role. The emergent property observed may be limited to continuous multimodal feedback, and may not have been observed yet in more common cases where designers limit feedback to discreet events only. For example, this property would not be easily observed in buttons lacking physical travel; those buttons only generate feedback when activated or released. We observed this emergent property when providing continuous, increasing feedback leading to actuation.

We selected the characteristics of each modality for their similarity. Future studies might investigate whether decreasing the intensity of the signal on one or more modalities reduces the emergent behavior that we noticed while remaining effective at communicating actuation when presented individually and in pairs. Further studies could also investigate situations where multimodal feedback communicates state and state transitions with greater intensity than single modalities of feedback alone.

We expect our findings will also pertain to more complex multi-modal interfaces beyond entry tasks. For example, if multimodal feedback were to convey the depth of a virtual drill, we might expect the emergent property to cause users to drill holes too shallow, were the feedback designed independently for each modality. We found that considering the effect of each modality independently is not sufficient for designing multimodal feedback; it is important to consider how the sum of the parts might emerge to create an unexpected feedback behavior.

CHAPTER 5. EFFECTS OF MODALITY ON VIRTUAL BUTTON MOTION AND PERFORMANCE

Modified from a paper submitted by Faeth, A. and Harding, C. to *Proc. ACM International Conference on Multimodal Interaction (ICMI2012)*. Submitted May 2012.

5.1 Abstract

The simple action of pressing a button is a multimodal interaction with an interesting depth of complexity. As the development of computer interfaces supporting 3D tasks progresses, there is a need to understand how users will interact with virtual buttons that generate multimodal feedback. Using a phone number dialing task on a virtual keypad, this study examined the effects of visual, auditory, and haptic feedback combinations on task performance and on the motion of individual button presses. The results suggest that the resistance of haptic feedback alone was not enough to prevent participants from pressing the button farther than necessary. Reinforcing haptic feedback with visual or auditory feedback shortened the depth of the presses significantly. However, the shallower presses that occurred with trimodal feedback may have led participants to release some buttons too early, which may explain an unexpected increase in mistakes when the participant missed digits from the phone number.

5.2 Introduction

Individuals' everyday familiarity with pressing mechanical buttons covers an interesting depth of complexity in what appears to be a simple multimodal interaction. When the user presses a button, the key travels a short distance before closing the electrical circuit to actuate the button. During this travel, the user can feel the resistance of the button and the click sensation, see the key traveling, and hear an actuation sound. This experience forms the expectations that users have when interacting with virtual buttons.

A virtual button has at least some of the feedback generated, rather than produced by inherent mechanical components. This means the designer can control the feedback experience, even though it may be informed by experience with mechanical buttons. Figure 5.1 illustrates a user interacting with a virtual button keypad.

Figure 5.1 Virtual keypad with PHANToM

In both three-dimensional and touchscreen interactions, there is an increase in the prevalence of virtual buttons. Virtual buttons used in interfaces are developed to support 3D tasks from medicine, manufacturing, geoscience, and engineering. Gestural interfaces and direct manipulation are more appropriate than buttons for some interactions such as scrolling through a large area. However, buttons retain advantages of discoverability and perceived affordance which make them integral components in an interface.

Compared to mechanical buttons, virtual buttons do not always provide the same feedback modalities and fidelity. It is not well understood whether design tradeoffs, such as omitting modalities, violates expectations users have when interacting with virtual buttons. This study examined the effect on task performance when some of the feedback modalities were deliberately omitted from a 3D interaction context.

5.3 Methods

The purpose of the study was to investigate the effect feedback modalities provided by a virtual button have on the motion of button press strokes, and on the task performance of the user. We designed the study to inform the following research questions:

- RQ4. How does the combination of feedback modalities affect the motion of the 3D button press?
- RQ2. Compared to trimodal feedback, how does reducing the number of feedback modalities affect task accuracy and completion time?
- RQ5. When users press a button, are they focusing on a particular type of feedback such as a visual cue?

5.3.1 Task

To assess the performance of virtual buttons with different feedback modalities, we asked the participants to dial random seven-digit phone numbers. Lewis et al. suggested a wider audience of participants would be more comfortable with a phone number task than a calculator task (Lewis et al., 1997).

The phone numbers used in each task were calculated randomly. The algorithm first picked a random starting digit, then chose a random adjacent digit until completing the seven-digit phone number. This maintained a similar distance between the buttons for all tasks.

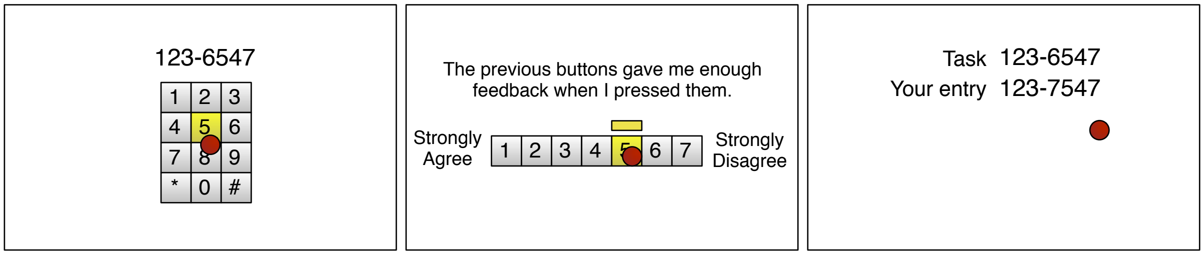

Figure 5.2 Phone number dialing task

The study divided the task into three parts, shown in Figure 5.2. The first part displayed the phone number to dial and the keypad, but did not show the digits dialed by the participant. Participants were instructed to dial the phone numbers quickly while maintaining accuracy. The participant could not see a partially dialed number on the screen, ensuring that the only feedback the participant received while dialing came from the virtual buttons.

After the participant dialed the number, the second part displayed a prompt to rate the feedback of the previous keypad, along with another set of virtual buttons for the response. The final part displayed the phone number assigned and the digits dialed by the participant.

5.3.2 Equipment

The setup of the study (Figure 5.3) used a PHANToM Omni to serve as a 3D input device and to generate haptic force-feedback from the buttons (Massie and Salisbury, 1994). The PHANToM Omni has a stylus attached to a mechanical arm, and as the participant moves the stylus through the 3D reach of the arm, the device renders forces that push back against the participant's movement, creating an illusion of touching objects in virtual space.

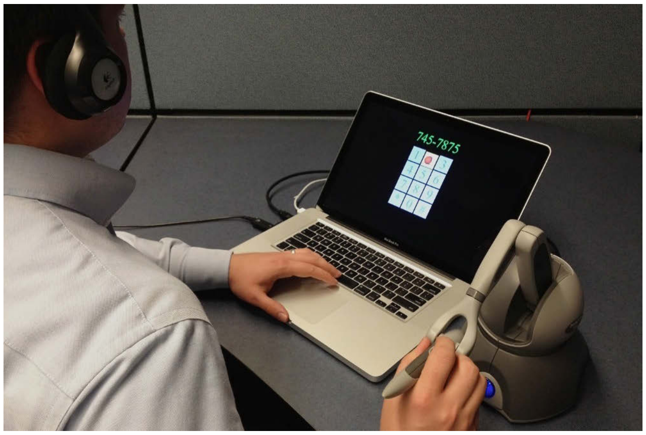

Figure 5.3 Study setup

The buttons were shown on a computer monitor, providing the participant with a means of targeting the buttons and any visual feedback from the buttons. The participant wore a pair of headphones to listen for auditory feedback from the buttons. Participants were instructed to hold the Phantom in their dominant hand to press the virtual buttons, and to rest their non-dominant hand on the keyboard spacebar to advance between the three screens of each task, similar to how one might write with their dominant hand and advance the paper with their non-dominant hand (Guiard, 1987).

5.3.3 Button Design

The entry keypad consisted of twelve virtual buttons arranged in a telephone layout. The buttons on the screen measured 2.5 cm square by 2.5 mm thick with a 2.0 cm travel. We chose the travel distance because participants in the pilot study reported that the haptic feedback felt too abrupt with a shorter travel. A longer travel allowed more time between the start of the feedback and reaching the bottom of the button.

The entry keypad did not have a bezel that surrounds the button keys on a physical keypad; this meant that if the participant missed the button key, the stylus would keep moving past the keypad surface. However, it would be confusing to have haptic feedback from a bezel when the buttons did not provide haptic feedback, or to have haptic feedback from a bezel only when the buttons provided haptic feedback.

Each virtual button used in the study had identical geometry and feedback characteristics. The only difference between the conditions was the presence or absence of the feedback modalities.

5.3.4 Button Feedback

Virtual buttons used in the study generated visual, auditory, and haptic feedback. The virtual buttons provided continuous haptic feedback to generate the feeling of a button click: a continuous rise in force, with a sharp drop at the actuation point, followed by another increasing force to cushion the bottom of the button (Figure 5.4). The bottom of the button travel was rendered as a solid haptic surface. When the participant touched a button with haptic feedback, the button also rendered four solid, invisible surfaces around the shaft of the button to prevent the participant's touch from inadvertently slipping beyond the button's bottom surface.

Auditory feedback from the button press played at three events during the button travel: a *tick* sound when the participant touched the button, a *plonk* sound at the actuation point,

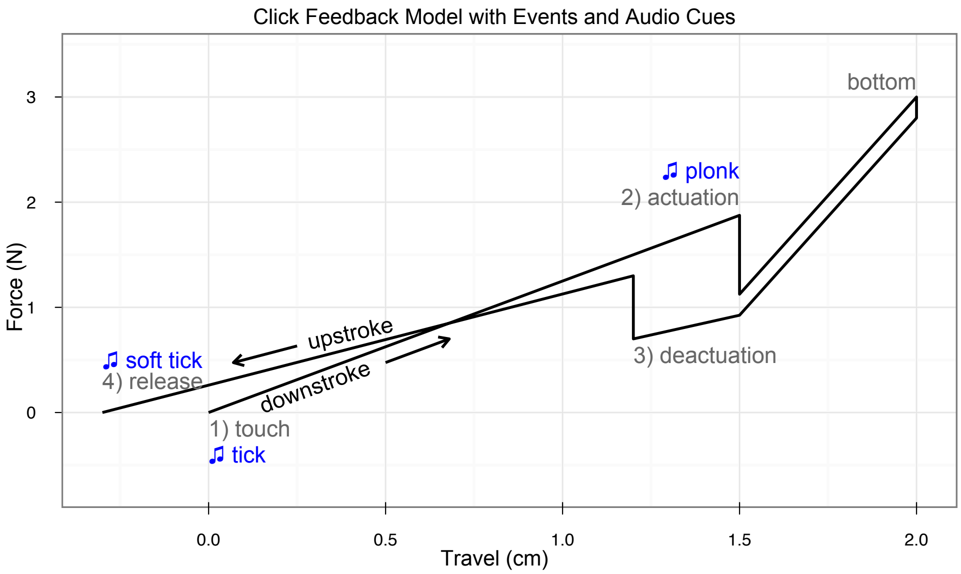

Figure 5.4 Force/travel for the haptic click

and a shorter, quieter *tick* sound when the participant released the button.

Visual feedback incorporated two components: the visible travel of the button, and a color change. As the button was pressed, the participant saw the button key recede relative to the rest of the keypad. The effect was visible without the aid of a stereoscopic display. As the button traveled to the actuation point, the color changed from gray to yellow using a linear interpolation.

Although we excluded visual feedback from some conditions, visual perception was still important for targeting virtual buttons. The participant used the keypad's graphic display and virtual stylus to target the next button to press. Once the stylus contacted the virtual button, it was important to control any visual feedback from that button.

When visual feedback was disabled, the virtual stylus appeared to stop at the button's surface while the participant's hand continued to move through the button's travel. This technique allowed virtual buttons to render haptic or auditory feedback from its travel when visual feedback was disabled. Participants were unaware of this effect and continued to press the button after the stylus appeared to stop on the button's surface.

5.3.5 Study Design

The study varied the presence or absence of visual, auditory, and haptic modalities to create the seven condition groups listed in Table 5.1. The characteristics of each modality's feedback were constant between conditions incorporating that modality.

Table 5.1 Experiment conditions

Condition	**Abbreviation**
Visual	V
Auditory	A
Haptic	H
Visual + Auditory	VA
Visual + Haptic	VH
Auditory + Haptic	AH
Visual + Auditory + Haptic	VAH

The study used a within-subjects design, each participant performing five repetitions of the task for each condition. The order of conditions were randomly assigned to each participant to control for learning and higher-order effects. To minimize the effect of participant fatigue, the study was designed to require no more than 30 minutes per participant, with a break provided between the training tasks and data collection.

5.3.6 Procedure

The same investigator conducted the study for all of the participants, and read from a script to aid in the consistency of the experience for each participant. After greeting the participant and explaining the study, the investigator asked participants to complete a short background questionnaire. Each participant then completed three training repetitions of the task for each of the conditions. After a short break, participants completed the five condition repetitions, and data was collected. After completing the tasks, the participants filled-out an exit survey.

5.3.7 Measures

For each task, we recorded the task duration, the number of errors in the dialed number, and the participant's subjective rating of the button feedback. We measured the completion time for each task between the time a participant initially touched a button after receiving the task, and the release of the last button pressed. The participant signaled completion of the task by pressing the spacebar of a physical keyboard with their non-dominant hand.

During each task, we counted the number of times a participant touched a button and then released it before the button traveled far enough to actuate the button (Figure 5.4). We called this measurement the number of unactuated presses during each task.

For each task, we calculated the number of digits that the participant inserted, deleted, or replaced from the assigned seven-digit phone number using the Wagner-Fisher algorithm (Wagner and Fischer, 1974). For example, to change the sequence 123-6547 into 123-564 requires a two-digit replacement of 65 with 56, and deleting the digit 7. We also calculated the total number of insert, delete, and replacement operations, which we called total errors. This is also called the *Levenshtein distance* between the two sequences of phone numbers (Levenshtein, 1965).

5.3.8 Subjective Rating

After dialing each phone number, but before seeing the phone number recorded by the virtual buttons, we asked the participant to rate the feedback on a 7-point Likert scale. The study prompted the participant with the statement: *"The previous buttons gave me enough feedback when I pressed them."* We used another virtual keypad with seven buttons arranged in a single line, and used the same feedback modalities from the task's condition. We chose to record the Likert feedback using another virtual keypad so that the participant could maintain one hand on the PHANToM and one hand advancing the task.

5.3.9 Measuring the Motion of Presses

During each button press, we also recorded the position of the stylus and any force rendered by the button click over the elapsed time of the press. Since humans can perceive variations in haptic feedback over very short periods of time, haptic rendering tries to achieve a minimum rate of 1000 Hz. To accurately capture the participant's haptic experience during each button press, we used a sub-millisecond precision timer (Microsoft, 2012). Since the haptic rendering was not performed by a real-time system, we observed that between 1500 and 2500 samples from the haptic thread were collected every second.

5.3.10 Participants

Twenty-one participants volunteered to take part in the study. One participant had a task success rate lower than the other participants (51%), and was also the only participant to report taking the tasks somewhat less than seriously on the exit survey. Therefore, this participant's results are not included in the analysis. The remaining 9 female and 11 male participants ranged in age between 19 and 42 years old (median 29).

One participant reported color blindness, but reported being able to distinguish between the gray and yellow colors used in the visual feedback. The remaining participants did not report conditions that might impair their ability to see, feel, or hear the button feedback.

Two of the participants reported their arm started to feel fatigued from holding their hand in the air to interact with the haptic device. These two participants opted out of taking the break offered between training and recorded trails.

Seven of the participants reported previous experience with haptic devices, and three reported using a haptic device at least weekly. Five of the participants reported using a device with a stylus at least once a week. All of the participants used a cell phone several times throughout the day, and fifteen reported their cell phone had a touchscreen.

5.4 Results

We manipulated the presence or absence of three feedback modalities while participants dialed a phone number on a virtual keypad, resulting in one independent variable with seven combinations of feedback modalities. During each task, we measured the number of errors, the number of unactuated presses, the time on task, and the participant's subjective rating of the feedback. We also measured the motion of the stylus during presses and calculated the maximum depth and mean velocity for individual presses.

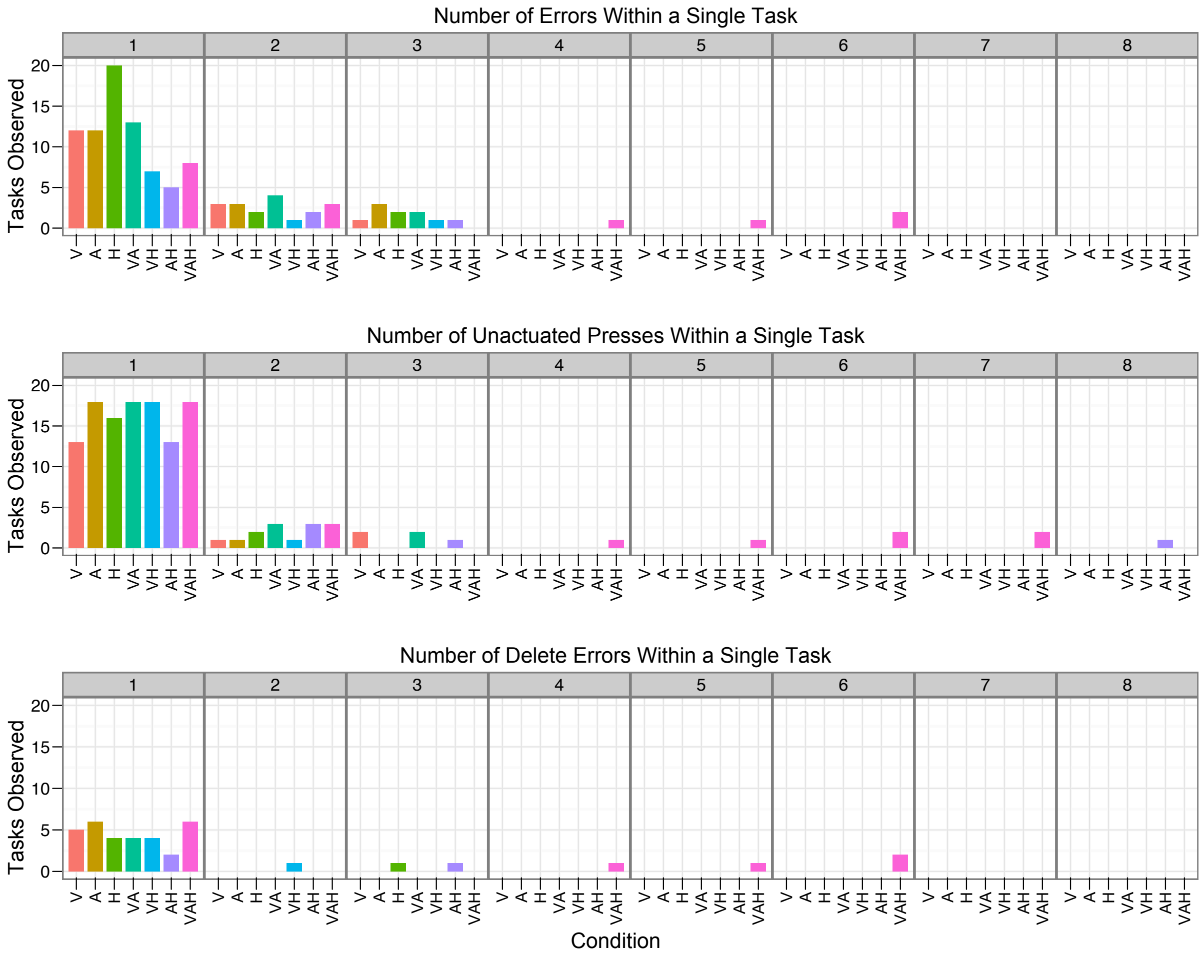
Number of Errors Within a Single Task

Number of Unactuated Presses Within a Single Task

Number of Delete Errors Within a Single Task

Figure 5.5 Errors and unactuated presses

5.4.1 Errors

Total errors for the conditions are summarized in Figure 5.5. A repeated-measures ANOVA revealed no significant effect of condition on total errors ($F_{6,114}$=1.289, p=0.268). Similarly, the effect of the condition on the insert and replace errors was not significant. However, a repeated-measures ANOVA did reveal a significant effect of the condition on the number of delete errors, when the participant missed a phone number digit ($F_{6,114}$=3.141, p=0.007). A post hoc Tukey HSD test revealed a significant increase in the delete errors between the trimodal feedback condition and each other condition, but no other significant differences between pairs of conditions ($\alpha = 0.05$).

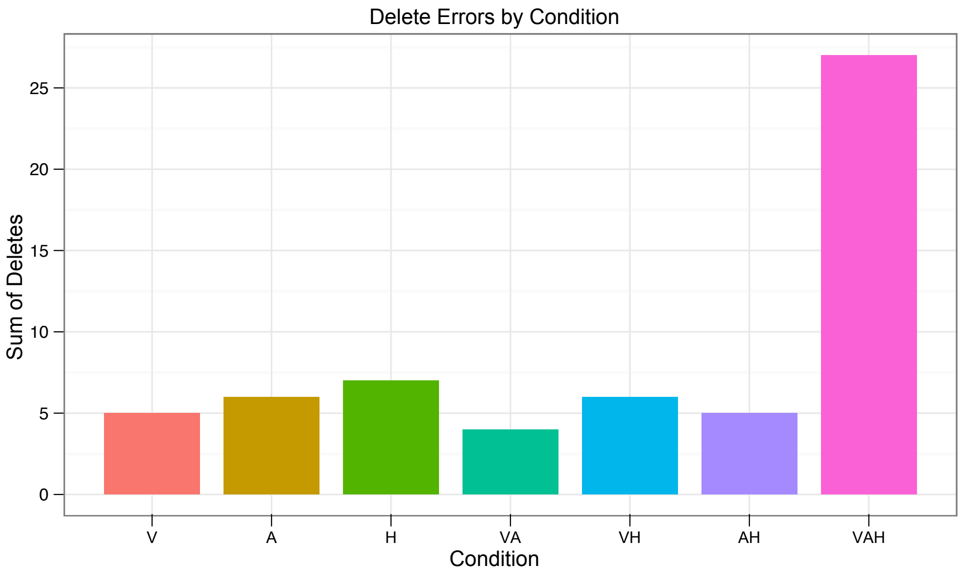

Figure 5.6 Sum of delete errors in each condition

This suggests that participants missed phone number digits more frequently when the buttons presented trimodal feedback than in any other condition (Figure 5.6). A Pearson's R test showed a significant correlation between deletion errors occurring during the trimodal condition and unactuated presses (R=0.69, $p<0.001$). The correlation between deletion errors and unactuated presses across all conditions was lower (R=0.55, $p<0.001$). The results suggest many of the deletion errors that occurred in the trimodal condition were unactuated presses and not simply errors where the participant skipped the digit completely.

5.4.2 Time on Task

The time on task for conditions is summarized in Figure 5.7. A repeated-measures ANOVA revealed a significant effect of the condition on completion time of the task ($F_{6,114}=5.097$, $p<0.001$). A post hoc Tukey HSD test showed time on task was shorter for VH compared to V and H; shorter for AH compared to V; and shorter for VAH compared to V and H ($\alpha = 0.05$). There were no significant differences between the other conditions.

A Pearson's R test shows an overall negative correlation between the number of modalities and completion time (R$= -0.16$, $p<0.001$). The bimodal combination of VH feedback reduced time on task compared to both unimodal V and H conditions. The trimodal condition also

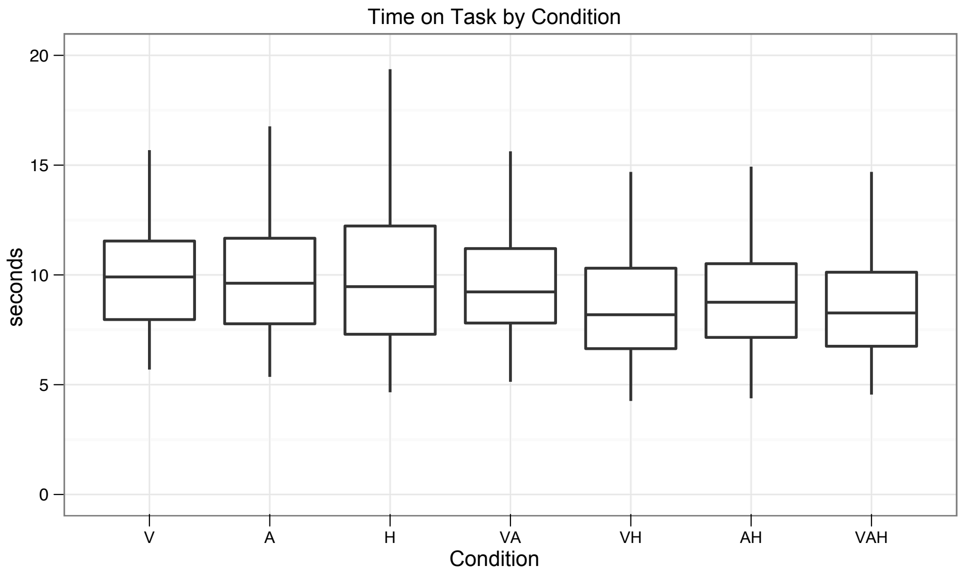

Figure 5.7 Time on task for each condition

showed faster completion times than either unimodal V or A feedback. The bimodal combination of AH feedback also reduced completion time compared to the unimodal V condition.

This suggests that although haptic feedback alone did not improve completion time, bimodal and trimodal conditions including haptic feedback did improve completion times over visual feedback alone. One possible reason is that the haptic feedback paired with other modalities helped to stop the participant's hand at the bottom of the button press, allowing them to move on to the next button quickly.

5.4.3 Subjective Rating

After each task, we asked participants to rate the feedback provided by the virtual buttons; a summary of their responses appears in Figure 5.8. A repeated-measures ANOVA revealed a significant effect of the condition on the participant's subjective rating ($F_{6,114}$=23.525, p<0.001). A post hoc Tukey HSD test revealed a significant preference for the bimodal and trimodal conditions shown in Table 5.2. Condition pairs that are not listed in table Table 5.2 were not significantly different ($\alpha = 0.05$).

Table 5.2 Subjective preferences (Tukey HSD)

Preference		Over Conditions
VAH	>	{ V, A, H, VA, AH }
VH	>	{ V, A, H, VA }
VA	>	{ V, A }
AH	>	{ V, A, H }

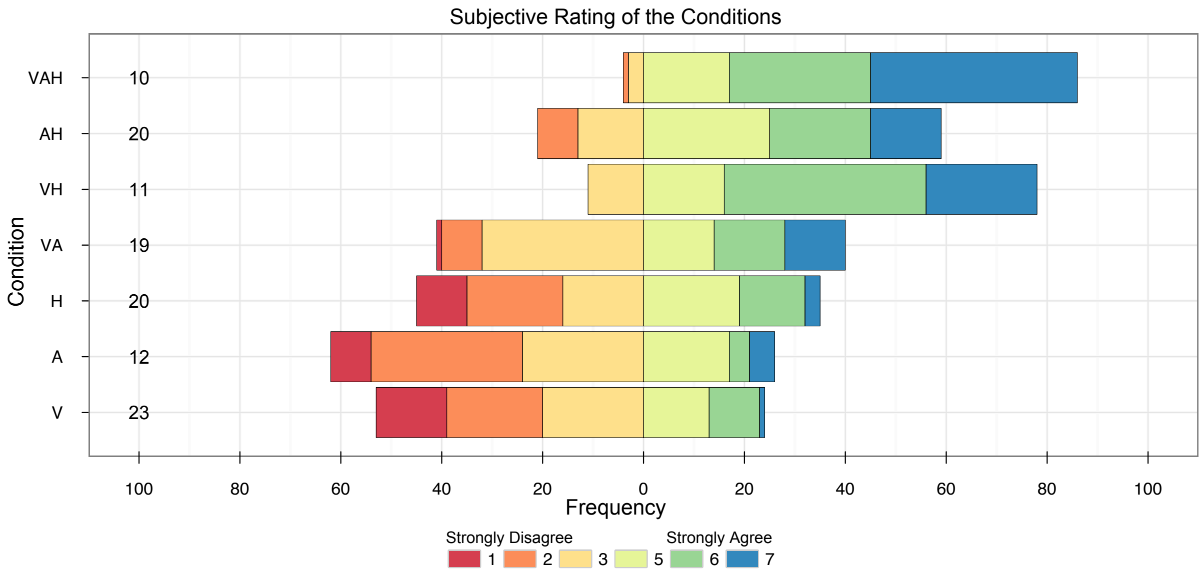

Figure 5.8 Subjective rating of the conditions. The numbers on the left show the frequency of neutral ratings.

5.4.4 Effect of Modalities on Task Success

The study also investigated whether a participant's focus on a particular modality affected the successful completion of a task. Past research suggests that attention has an effect on a user's reliance on one modality (Spence et al., 2000). We did not ask participants about attention directly, but five participants commented they were paying attention to one modality when provided an opportunity to comment on the tasks.

"I think that visual is more effective for me. When I see the color change, I know I have pressed the button I meant to press" (P14).

"I liked having the visual feedback—the ones that didn't do that, I was concerned whether I

pressed it, or I had slipped onto a different button" (S19).

"Visual seemed to be the only 'definitive' feedback" (P2).

"In the haptic-only ones, I wasn't sure how hard to press because there was no visual cue that you'd pressed the button" (P12).

"Midway through [the tasks], I figured out how far I needed to press the button using the audio and visual cues as a clarification" (P4). In this case, the participant reported paying attention to proprioception and using the audio and visual cues for reinforcement.

Despite these statements, the data for these participants shows no clear difference in task success between modalities (Figure 5.9). To investigate whether any of the participants were more successful with a single modality, we performed a logistic regression for each one. The logistic regression model used the presence of the modalities during the task's condition to predict task success. We performed a Chi-squared test on the logistic regression model for each participant, and rejected any model that failed to meet an α of 0.05.

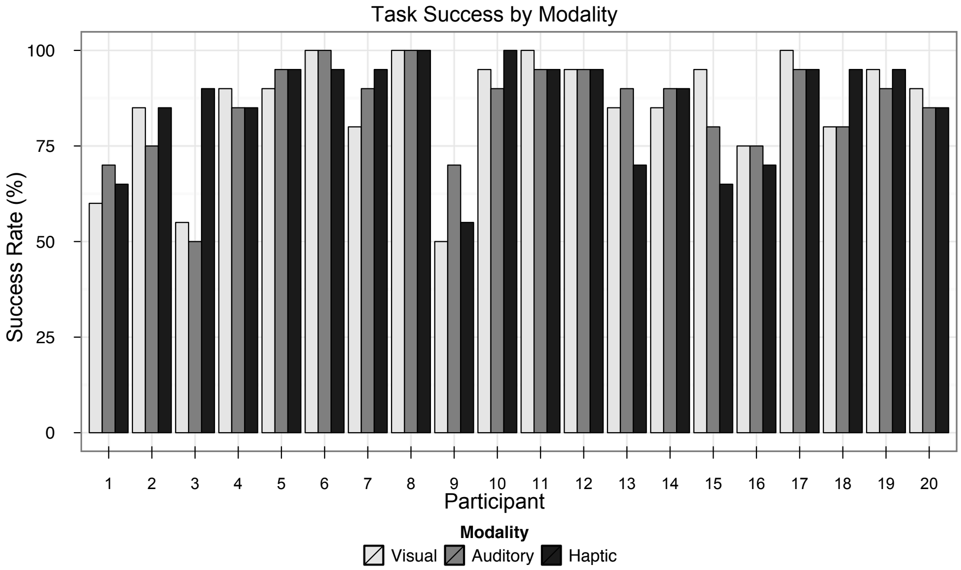

Figure 5.9 Success rates for each participant

We only found two incidences where the presented modalities predicted whether the participant would successfully complete the tasks (Table 5.3). P3 was more likely to succeed when haptic feedback was present than when either auditory or visual feedback occurred. The differ-

Table 5.3 Predicting success with modalities

	Modality	**Est.**	***SE***	**Z**	***p***
P3[1]	Haptic	2.197	0.745	2.948	0.003
	Visual	-1.997	0.870	-2.294	0.022
	Auditory	-2.197	0.869	-2.528	0.012
P15[2]	Visual	2.944	1.026	2.870	0.004
	Auditory	-1.558	1.168	-1.334	0.182
	Haptic	-2.325	1.128	-2.062	0.039

[1] χ^2(2, N=60)=9.44, p=0.009. Log Lik. -34.13 df=3.

[2] χ^2(2, N=60)=6.19, p=0.045. Log Lik. -26.93 df=3.

ence between visual-audio feedback was not significant. P15 was more likely to succeed when visual feedback was present than when haptic feedback was present. The difference between visual-audio and audio-haptic modalities was not significant for P15.

5.4.5 Press Motion Characteristics

We recorded the position and force of the stylus during a total of 5,015 button presses as the participants each dialed 35 phone numbers. From the motion of the stylus, we calculated the maximum depth the button traveled and the mean absolute velocity of the button press motion.

5.4.5.1 Maximum Press Depth

A summary of maximum press depth is shown in Figure 5.10. A repeated-measures ANOVA showed a significant effect of the condition on the maximum depth of the press ($F_{6,114}$=11.066, p<0.001). A post hoc Tukey HSD test showed significant differences between the VH, AH, and VAH conditions (Table 5.4). The remaining condition pairs did not show a significant difference ($\alpha = 0.05$).

Although the resistance of haptic feedback did not appear to constrain the depth of the button travel alone, the addition of visual feedback significantly shortened the maximum depth in both the VH and VAH conditions. The effect was also significant for auditory feedback in the AH and VAH conditions.

Table 5.4 Press depth comparison (Tukey HSD)

Shallower Press		Deeper Press
VAH	<	{ V, A, H, VA, AH }
VH	<	{ V, A, H, VA }
AH	<	{ H }

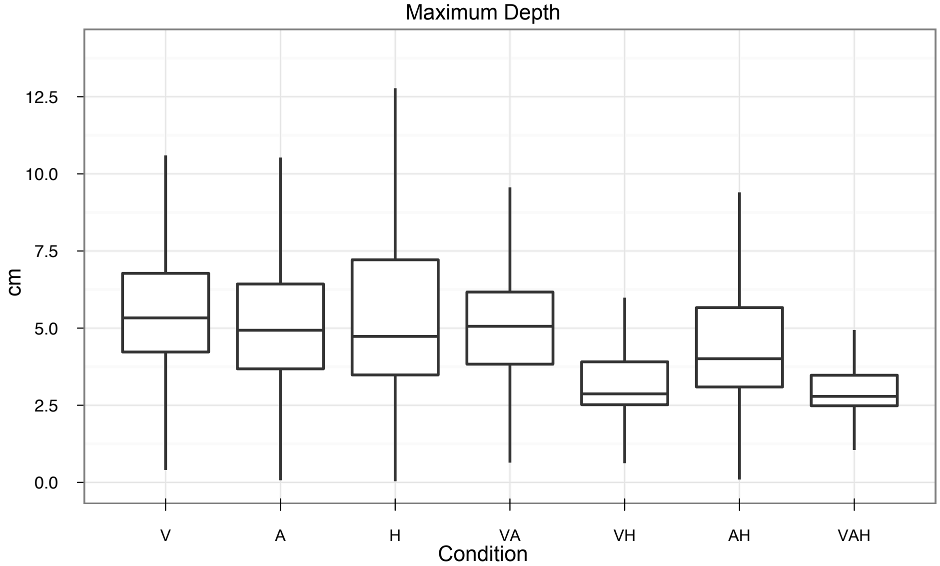

Figure 5.10 Press depth

Even though maximum travel of the button was 2.0 cm, participants continued to press some buttons past 5.0 cm. In conditions where haptic feedback was present, the participants felt the bottom surface of the button at 2.0 cm. Since the haptic device cannot render perfect stiffness, it was possible for the stylus to to sink into the bottom surface of the button, or for the participant to overpower the PHANToM and break through the bottom of the button.

5.4.5.2 Mean Press Velocity

The summary of the mean absolute velocity appears in Figure 5.11. A repeated-measures ANOVA showed a significant effect of the condition on the press velocity ($F_{6,114}$=20.492, p<0.001). A post hoc Tukey HSD test revealed the significant differences between conditions shown in Table 5.5. The remaining pairs of conditions were not significantly different from each other ($\alpha = 0.05$).

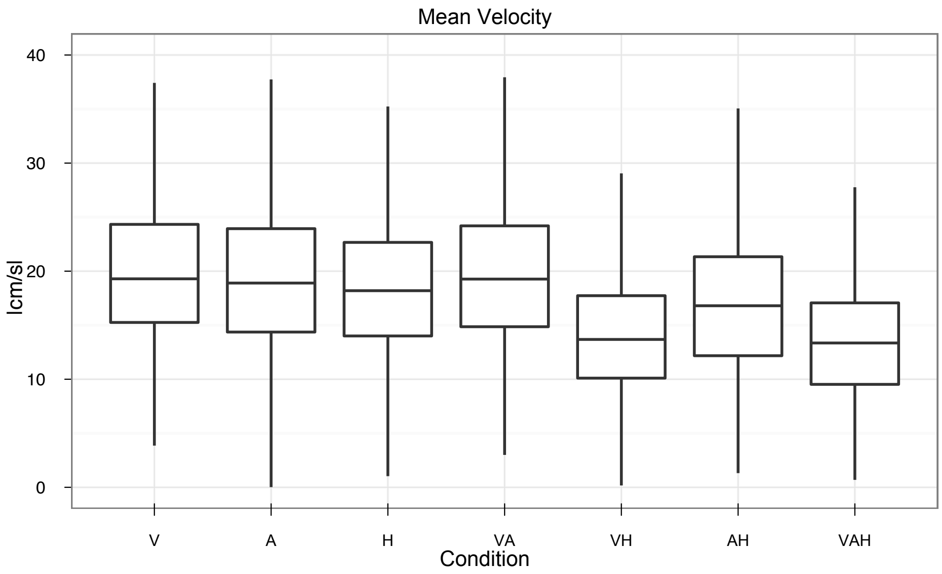

Figure 5.11 Press velocity

Table 5.5 Press Velocity Comparison (Tukey HSD)

Lower Velocity		**Higher Velocity**
VAH	<	{ V, A, H, VA, AH }
VH	<	{ V, A, H, VA, AH }
AH	<	{ V, A, VA }

Similar to the press depth, the results show a lower mean velocity in the VH and VAH conditions compared to the other conditions. If the press broke through the button bottom, the velocity of the stylus would increase when the resistance dropped away. Since the depth of presses during the H and AH conditions were higher than in the VH and VAH conditions, it is likely this break-through caused the higher mean velocity in the H and AH conditions.

5.5 Discussion

Participants expressed a significant preference for bimodal and trimodal feedback from virtual buttons. The following will examine how the results from press motion characteristics, task performance, and specific modalities fit into the following research questions.

5.5.1 Effect of Condition on Press Motion

RQ4. How does the combination of feedback modalities affect the motion of the 3D button press?

We expected the presence of haptic feedback would shorten the distance the user pressed the button past the bottom because the user would feel resistance. However, it was unexpected that haptic feedback alone did not significantly reduce the maximum depth of button presses. The addition of visual feedback significantly shortened button presses in both VH and VAH conditions. Auditory feedback also demonstrated a significant effect in shortening the button presses in AH and VAH conditions.

Visual and auditory feedback reinforce the subjective haptic perception of stiffness (DiFranco et al., 1997; Lécuyer et al., 2001; Miner et al., 1996; Srinivasan et al., 1996). However we would expect this effect to shorten the presses in the H condition where the button did not appear to travel, rather than in the VH and VAH conditions where the button did appear to travel. Díaz et al. (2006) observed a similar effect where participants stayed within the lines of a maze more often when haptic or auditory feedback reinforced visual feedback, but their study did not observe unimodal haptic or auditory feedback.

We found anecdotal evidence that some of our participants may have perceived a stronger haptic feedback when visual or auditory feedback was also present. Two participants commented that they thought the force-feedback presented to them changed between keypads, although it was kept constant throughout the study. *"I didn't expect you to change up the forces"* (P12). *"The pressure [force] changed a lot"* (P16). These comments may simply reference conditions without haptic feedback, however, in the context of the previous studies, participants may have perceived a change in force between keypads with identical haptic feedback.

5.5.2 Accuracy and Time on Task

RQ2. Compared to trimodal feedback, how does reducing the number of feedback modalities affect task accuracy and completion time?

The feedback condition did have a significant effect on time on task and on the number of delete errors occurring during the task, however it did not have a significant effect on the total number of errors. The results showed a decrease in time on task for the VH, AH, and VAH conditions, the same conditions under which we observed a decrease in the depth of the presses. This is consistent with past studies demonstrating that users react faster to multimodal feedback, because part of the task is reacting to the feedback from the buttons.

However, the results also showed a significant delete error increase in the trimodal condition, which was not consistent with past research showing decreased multimodal feedback errors. One explanation might be that the feedback from one of the modalities was poorly designed, or that lag between two modalities caused participants to make more errors in the trimodal condition. If that were true, we would expect to see increased errors in all conditions where one modality was present, or in bimodal conditions where the lagging modality was present. Instead, the results showed a significant increase in deletion errors in the trimodal condition, compared to the other conditions.

Another explanation might be that the participant never touched the button corresponding to the missing digit. However, we would then not expect to find the significant correlation between unactuated presses and delete errors in the trimodal condition.

The results therefore suggest that the participant released the button early in the trimodal condition and was able to perceive the feedback from all three modalities. This is consistent with multiple resource theory because there was no evidence of contention for a single resource in processing the feedback from multiple modalities.

5.5.3 Attention to a Particular Modality

RQ5. When users press a button, are they focusing on a particular type of feedback such as a visual cue?

Only two participants showed a significant increase in task success when a particular modality was present, perhaps indicating they focused on cues from a particular modality while other participants shifted focus between cues from multiple modalities as they performed the tasks. Another explanation might be that the cues from the modality that improved the odds of success for those two participants were simply more salient to them.

Although five participants reported feeling more successful when a particular modality was present, the results do not show that a single modality, such as vision, was necessary to complete the task. Instead, bimodal and trimodal conditions reduced time on task and shortened the depth of presses.

5.6 Conclusions and Future Work

The results showed that haptic feedback alone was not enough to constrain the button presses to the travel of the virtual button. Reinforcing the haptic feedback with visual or auditory feedback was effective in reducing the distance users pressed a button past the bottom. However, the shallower presses also suggest that participants released some buttons early in the trimodal condition. This may have led participants to make more errors by missing digits.

Some participants demonstrated more success with one particular modality. A future longitudinal study could better determine if the effect observed in this study changes over time for each participant. Another future area of study might investigate whether the color changes used with the buttons in this study have a similar effect on perceived haptic stiffness observed in other studies manipulating the visual displacement of springs and surfaces.

As virtual buttons become more prevalent in supporting 3D tasks, users will seek out the multimodal feedback they prefer. In this study, a difference in the way trimodal feedback was perceived caused users to make more errors. A better understanding of multimodal feedback will help improve future interactions with virtual interfaces.

CHAPTER 6. STUDY OF TOUCH HYSTERESIS AND PRESS MOTION CHARACTERISTICS

6.1 Introduction

When a user presses a mechanical button, the feedback from making contact with the button comes from touching the button key and the resistance to traveling. A virtual button computes feedback cues during the press instead of generating feedback though the movement of mechanical components. A virtual button communicates the initial contact with a touch event, and the loss of contact with a release event. As the user continues to press the button, it also renders feedback upon actuation and deactuation events.

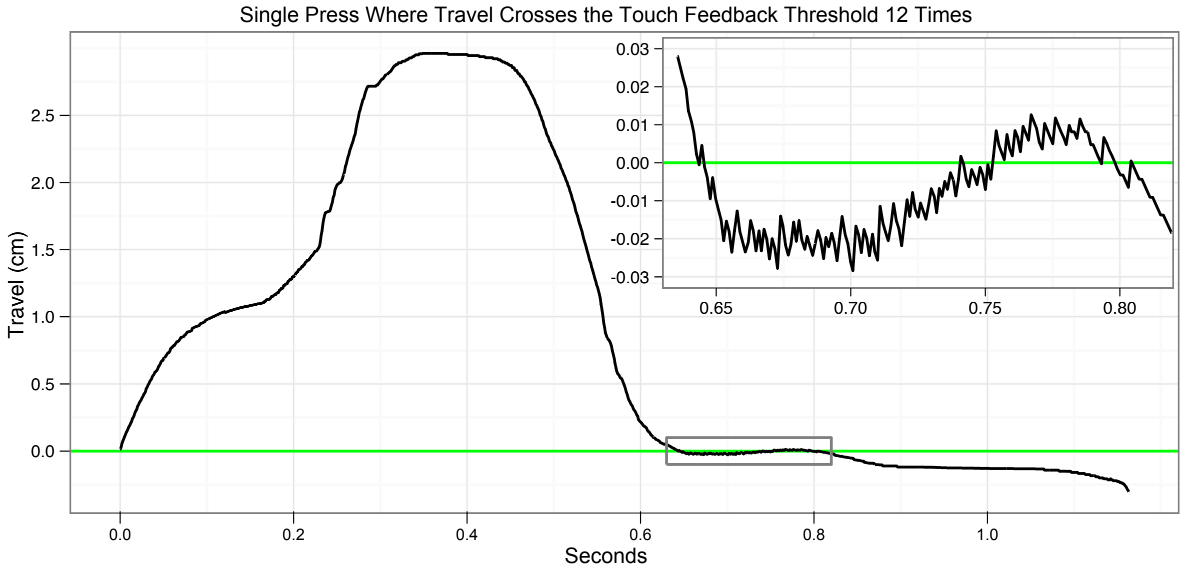

Figure 6.1 Without hysteresis at the touch threshold, the user may receive multiple touch and release feedback events from the same button press. The upper-right window of the graph shows the 10 extra threshold crossings in greater detail.

If a virtual button is not designed to prevent repeated feedback events from occurring in rapid succession, it may generate confusing sequences of repeated touch and release feedback events from what the user intends to be a single button press. Figure 6.1 shows an example of a press with 10 extra touch-threshold crossings in the upper-right window of the graph. Mechanical buttons use two strategies for preventing rapid changes between actuation and deactuation. Interlocks introduce a short timeout after an actuation that prevents another actuation from occurring until the duration expires. Hysteresis introduces a memory of the button state to create a separation between the actuation point on the downstroke of the button and the deactuation point on the upstroke (Lewis et al., 1997).

Mechanical buttons generate visual, auditory, and haptic feedback as the user presses the button. The user can see the displacement of the button as it travels, hear the click at actuation, and feel the resistance of the button during the press. Since the design of virtual buttons should be informed by experiences with mechanical buttons, the virtual buttons should also generate some combination of visual, auditory, and haptic feedback.

For the same reasons that hysteresis and interlocks benefit the actuation of a mechanical button, hysteresis and interlocks also benefit touch and release feedback events from a virtual button. However, previous research has not evaluated the effectiveness of hysteresis in virtual button feedback. Existing models of button feedback only prevent the repeated actuation of the button itself, rather than preventing the confusing cues to the user from repeated feedback events.

This research presents a model for virtual button feedback that includes hysteresis for touch feedback. We also evaluated the effects of touch and actuation hysteresis on virtual buttons in a 3D interface with a study of recorded virtual button presses.

6.2 A Model for Virtual Button Feedback

There are two categories of input events which could lead to repeated touch and release feedback in a virtual button: extra threshold crossings, and threshold dwell. In an ideal button press, the user pushes the button across the touch threshold once in each direction. Extra threshold crossings occur when the travel of the button repeatedly crosses the touch threshold more than twice. This has the potential to create confusing feedback cues as the button renders multiple touch and release feedback events from what the user perceives as a single press. Extra actuation crossings also occur when the user pushes the button across the actuation threshold more than twice.

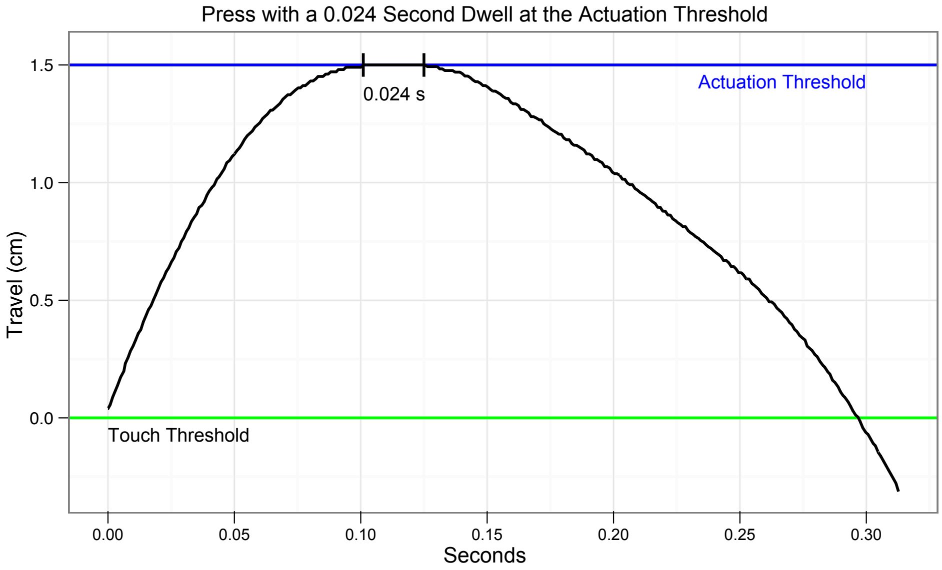

Figure 6.2 A press with dwell at the actuation threshold of the button.

Threshold dwell occurs when the user holds the button close to the touch or actuation threshold (Figure 6.2). The virtual button would observe the travel at the threshold and provide feedback from the touch event. If the travel remains at the threshold, the virtual button may record a release event at the next update. An intuitive solution would introduce a small hysteresis to record the button touch when the travel is greater-than-or-equal-to the threshold, and release it when the travel is less-than the threshold. While this would prevent threshold dwell from generating repeated feedback events, extra threshold crossings would still occur.

6.2.1 A Model of Virtual Button Feedback with Touch Hysteresis

Figure 6.3 presents a model for generating continuous and discrete virtual button feedback from the continuous input parameter of the button travel. The button travel describes how far the user presses the button from the starting point. The motion of the press from the touch to the actuation is the downstroke, and the return to the starting point is the upstroke. The continuous force feedback from the button renders a haptic click sensation at the actuation point. As the button travels, discrete Touch, Actuation, Deactuation, and Release events provide additional feedback. The model includes hysteresis for both the actuation and touch thresholds.

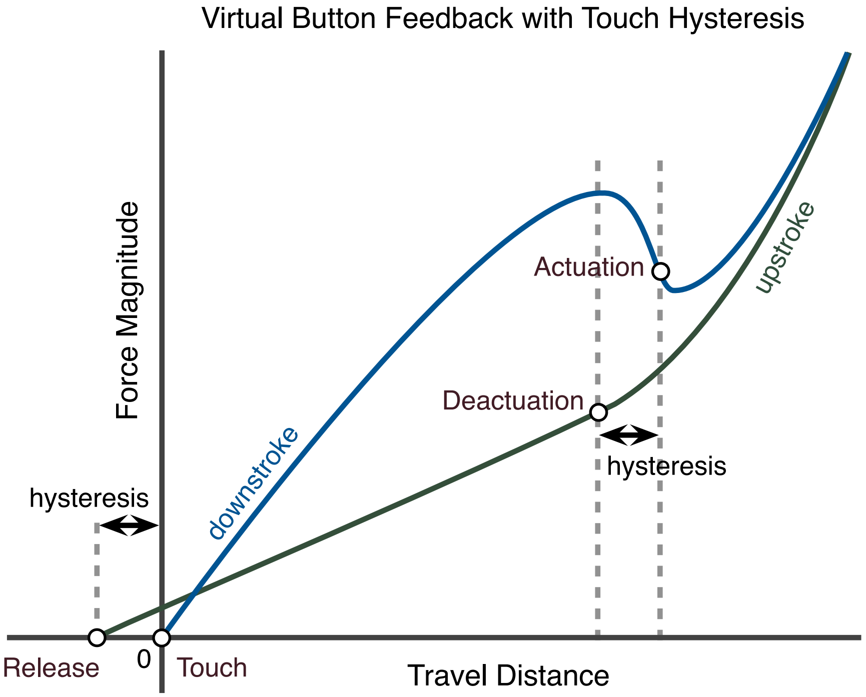

Figure 6.3 A model of Virtual Button Feedback with Touch Hysteresis. The graph shows the location of the Touch, Release, Actuation, and Deactuation events on the upstroke and downstroke.

6.2.2 Hysteresis vs. Interlocks in Virtual Button Feedback

In determining button actuation, hysteresis and interlocks are both effective in preventing repeated actuations from the same press. However, interlocks have a drawback that makes hysteresis more suitable for virtual button feedback. To prevent repeated touch and release feedback events with an interlock, the designer could use the same interlock for touch and release feedback. However, if the user made a short exploratory contact with the button, the feedback interlock might prevent the release feedback after rendering the touch event feedback.

A button with hysteresis would provide both the touch and release feedback events from short presses.

If the touch and release feedback used separate interlocks, it would be unlikely to prevent touch feedback event when there are extra touch threshold crossings at the end of the button press. For example, the press in Figure 6.1 would trigger a release feedback interlock around 0.65 seconds when the travel crosses the touch threshold. With separate touch and release interlocks, the button would render touch feedback when the travel parameter crossed the threshold again close to 0.74 seconds. The user may perceive this as touch feedback when they intended to release the button. For this reason, hysteresis is well suited to preventing repeated feedback events from the same button press.

6.3 Methods

This study addresses a new set of research questions using the high resolution motion data collected during the study in Chapter 5. Therefore, this study used the same methods and participants from the prior study, and this chapter presents a shortened description of the methods.

The purpose of the study was the evaluate the effect of touch hysteresis in preventing repeated feedback events from extra threshold crossings and threshold dwell. Since virtual button feedback may include a combination of visual, auditory, and haptic feedback, the study also explored whether the combination of feedback affected the threshold crossings or dwell observed. We designed the study to address the following research questions.

1. RQ6. In a virtual button, what effect does hysteresis at the touch and actuation thresholds have on preventing confusing feedback signals?
2. RQ7. What effect does the combination of feedback modalities used in virtual buttons have on the threshold crossings and threshold dwell?

6.3.1 Task

We asked participants to dial 7-digit phone numbers using a virtual keypad. The phone numbers were randomly generated by picking a random starting digit and then choosing adjacent digits to complete the number. The phone numbers did not contain any repeated digits, so any presses in which the participant actuated the button multiple times were unintentional.

6.3.2 Button Design

The entry keypad used twelve virtual buttons arranged in a telephone layout. The buttons on the screen measured 2.5 cm square by 2.5 mm thick with a 2.0 cm travel. The actuation point of the buttons occurred when the button traveled 1.5 cm. The hysteresis of both the touch and actuation thresholds was 3.0 mm. This meant that the deactuation point was 1.2 cm and the release point was -0.3 cm. To select the value for hysteresis, we measured the distance that the button travel deviated while we tried to hold the button steady at the bottom of the button. At the bottom, the user is pushing against the resistance of the button spring and the virtual surface of the button bottom. We calculated the maximum value observed and rounded up to 3.0 mm.

6.3.3 Study Design

We manipulated combinations of visual, auditory, and haptic feedback present to create seven feedback conditions. The study used a within-subjects design, with each participant performing five repetitions of the task for each condition. The study randomly assigned the order of the conditions to each participant in order to control for learning and higher-order effects. To minimize the effect of fatigue on the participants, the study was designed to take 30 minutes, with a break provided between the training tasks and data collection.

6.3.4 Measuring Threshold Crossings

For each button press, we counted the number of times that the button traveled across the touch threshold and the actuation threshold. We counted each time the button crossed the threshold in either direction. In a typical button press we would only expect the button to cross the touch threshold twice, once when the user touched it, and once when the user released it. Any additional crossings over the threshold were extra crossings. We counted the overall number of presses with extra threshold crossings, and recorded the number and time of each crossing in a press.

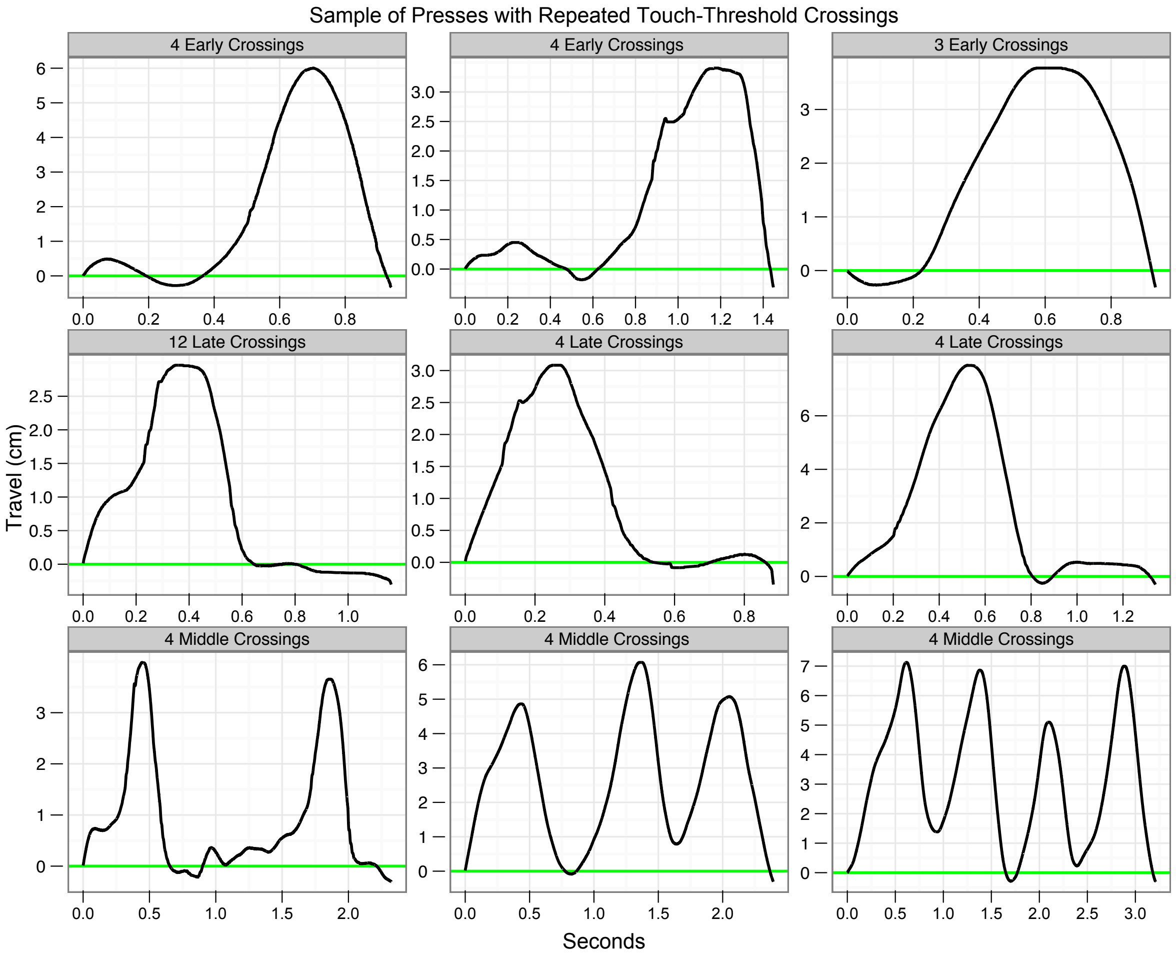

Figure 6.4 Sample of presses where the user repeatedly crossed the travel threshold for providing touch and release feedback.

6.3.5 Measuring Threshold Dwell

When the user held the button close to the touch or actuation thresholds (within $\frac{1}{1000}$ mm), we also counted the number of consecutive samples where the button had the same travel value. Since the haptic thread attempts to run at 1000 Hz, we counted any samples recorded within

0.0015 seconds of each other as consecutive. From the consecutive samples, we calculated the time that the user remained in close proximity to the threshold. We called this measurement the threshold dwell time. Figure 6.2 shows an example of a press with an actuation threshold dwell 0.024 seconds long.

6.3.6 Peak Button Travel

The point at which the user reached the peak of their press travel and started to release the button is the peak of the button travel. This indicates the proportion of the press duration that the user spent pressing the button compared to releasing the button. To determine the peak button travel, we calculated the global maximum of the button travel. In the rare cases where two or more peaks tied for the global maximum, we chose the later peak. Since there were no repeated digits in the task, if the user pressed the same button repeatedly, it was a mistake. We normalized the time of the peak travel in order to compare the time spent on the downstroke and upstroke across button presses with unequal durations.

6.4 Results

6.4.1 Participants

The same participants from Study 2 participated in the collection of these button presses. Twenty people between the ages of 19 and 42 years old (median 29) volunteered to participate in the study. We collected 5015 button presses from the 700 phone numbers that participants dialed.

6.4.2 Threshold Crossings

The summary for press threshold crossings appears in Table 6.1. A Chi-squared test found no evidence of a relationship between condition and total number of presses with multiple crossings ($\chi^2\{6, N=5015\}=11.5153$, $p=0.074$). We did not separate the actuation and touch crossings for the Chi-squared test because some of conditions had less than five occurrences of presses with extra crossings.

To determine whether extra crossings were more likely to occur at the touch or the actuation threshold, we performed a Chi-squared test on the proportion of presses with extra crossings recorded at each threshold. The results of the test provided support to reject the null hypothesis that extra crossings were equally likely at both thresholds ($\chi^2\{1, N=245\}=94.30, p<0.001$). The 95-percent confidence interval for the proportion of presses with actuation crossings compared to touch crossings was [0.14, 0.24]. The results suggest that when extra crossings occur in a press, it is more likely to be extra crossings of the touch threshold.

Table 6.1 Threshold crossings by condition

	Presses With Crossings			
Condition	**Touch**	**Actuation**	**Total**	**All Presses**
V	36	7	43	714
A	19	11	30	702
H	24	6	30	723
VA	35	13	48	713
VH	25	3	28	710
AH	29	1	30	722
VAH	31	5	36	731
	199	46	245	5015

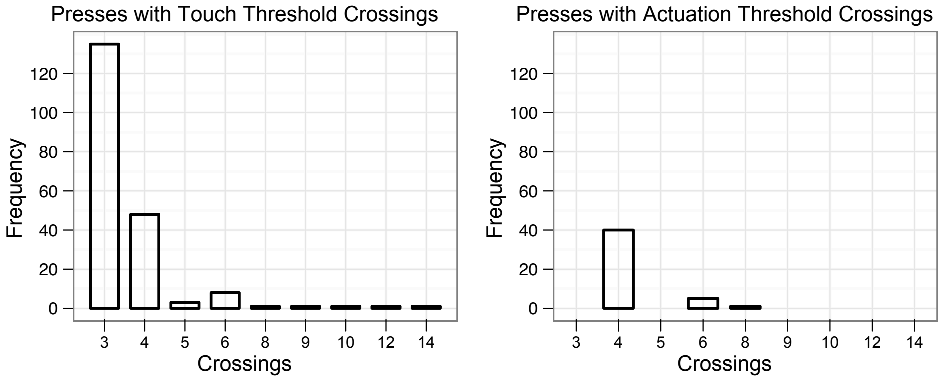

Figure 6.5 Frequencies of crossings across the touch and actuation thresholds.

Figure 6.5 shows the frequency of presses with extra threshold crossings. When a press had an odd number of touch-threshold crossings, the participant crossed the touch threshold and returned to the negative side of the threshold before the first sample was recorded. The top-right graph in Figure 6.4 shows this effect. The results only recorded odd numbers of crossings from at the touch threshold. The graphs show that most of the presses with extra crossings had small number of extra crossings.

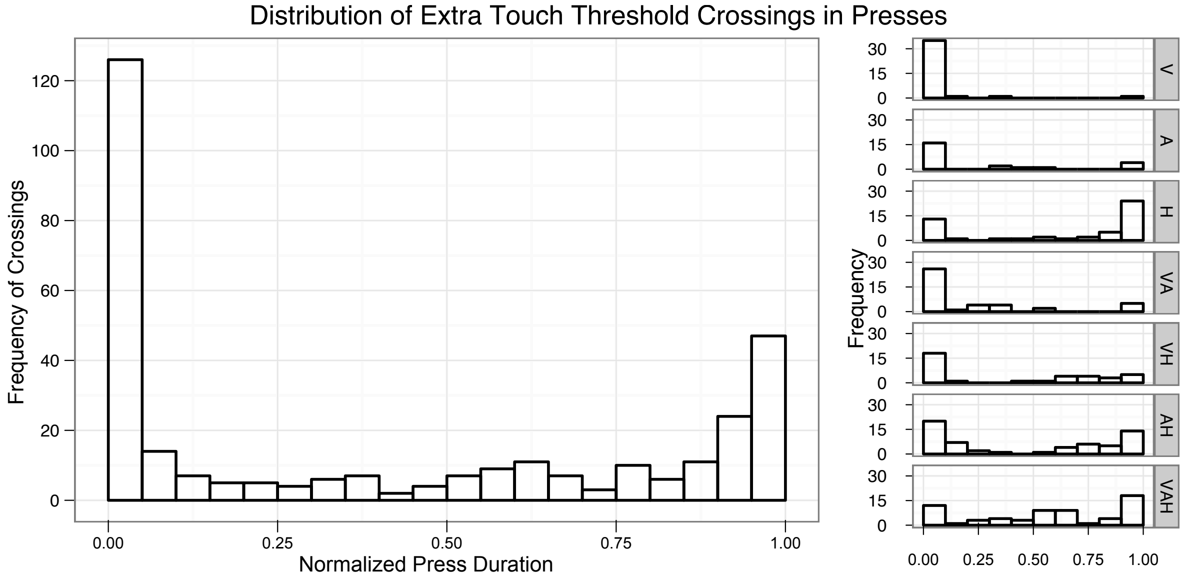

Figure 6.6 Distribution of the touch threshold crossings over the normalized press time

As expected, the majority of the extra touch threshold crossings occurred toward the start (Row 1 of Figure 6.6) or end of the button press (Row 2 of Figure 6.6). However, we observed smaller numbers of touch threshold crossings across the entire normalized press duration. The touch threshold crossings are similar between conditions. The graphs do show a small increase in touch crossings at the end of the presses for the H, AH, and VAH conditions, but not the VH condition.

The extra actuation threshold crossings tended to occur toward the middle of the press (Figure 6.7). The results did not record enough crossings for an effective comparison of actuation crossings between conditions.

6.4.3 Threshold Dwell

The summary for the presses with threshold dwell appears in Table 6.2. One of the presses had two separate occurrences of touch threshold dwell. We summed the dwell time for each

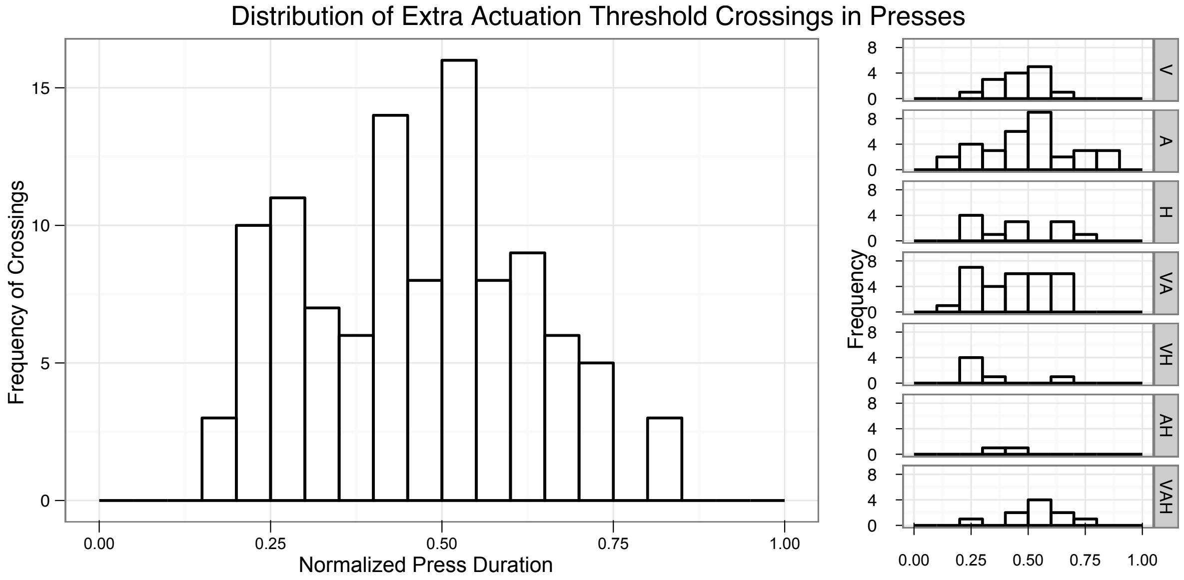

Figure 6.7 Distribution of the actuation threshold crossings over the normalized press time

occurrence to prevent the press from appearing in the table twice. None of the presses in the sample had both touch and actuation threshold dwell.

To determine whether threshold dwell was more likely to occur at the touch or the actuation threshold, we performed a Chi-squared test on the proportion of presses with dwell recorded at each threshold. The results of the test provided support to reject the null hypothesis that threshold dwell was equally likely at both thresholds ($\chi^2\{1, N=185\}=15.76, p<0.001$). The 95-percent confidence interval for the proportion of presses with actuation dwell compared to touch dwell was [0.57, 0.71]. The results suggest that when threshold dwell occurs in a press, it is more likely to be at the actuation threshold.

A Chi-squared test revealed a significant effect of the condition on the total number of presses with dwell (($\chi^2\{6, N=5015\}=19.0448, p=0.004$). A pairwise post hoc Chi-squared test with a Holm adjustment revealed a significantly higher number of presses with dwell in the H condition compared to the V condition ($\alpha = 0.05$). If the alpha was raised to 0.1, the VAH condition would also have a significantly higher number of presses with dwell compared to the V condition (p=0.066). The test revealed no other significant differences between pairs of conditions.

Table 6.2 Press dwell by condition

	Presses with Dwell			
Condition	**Touch**	**Actuation**	**Total**	**All Presses**
V	5	9	14	714
A	7	17	24	702
H	15	24	39	723
VA	3	13	16	713
VH	8	20	28	710
AH	13	15	28	722
VAH	14	22	36	731
	65	120	185	5015

The majority of the touch and actuation threshold dwell lasted less than one millisecond. Since there was no upper bound to on the rate of the samples recorded in the haptic rendering thread, this might be the result of two updates occurring before the haptic device updated its position information.

The results show only a few cases of touch threshold dwell (Figure 6.8). Most of the observed cases of touch threshold dwell occurred at the end of the button press instead of the distribution we expected at the start and the end of the press. One explanation might be that the participant was moving toward the next button in the task at this point, perpendicular to the surface of the button. Since we recorded dwell in the linear observation of button travel, this might result dwell near the touch threshold as the user moved toward the next button.

The graph in Figure 6.9 shows the actuation threshold dwell over the normalized duration of the button press. The graph shows two peaks where actuation dwell occurred more frequently. The first peak represents the actuation dwell on the downstroke, and the later peak represents the actuation dwell on the upstroke. The distance between the peaks of the distribution is not unexpected, because the user would continue to press the button past the actuation point to the bottom of the button (Figure 6.3). In the next section, we will examine why the valley this graph occurs more than halfway through the duration of the press.

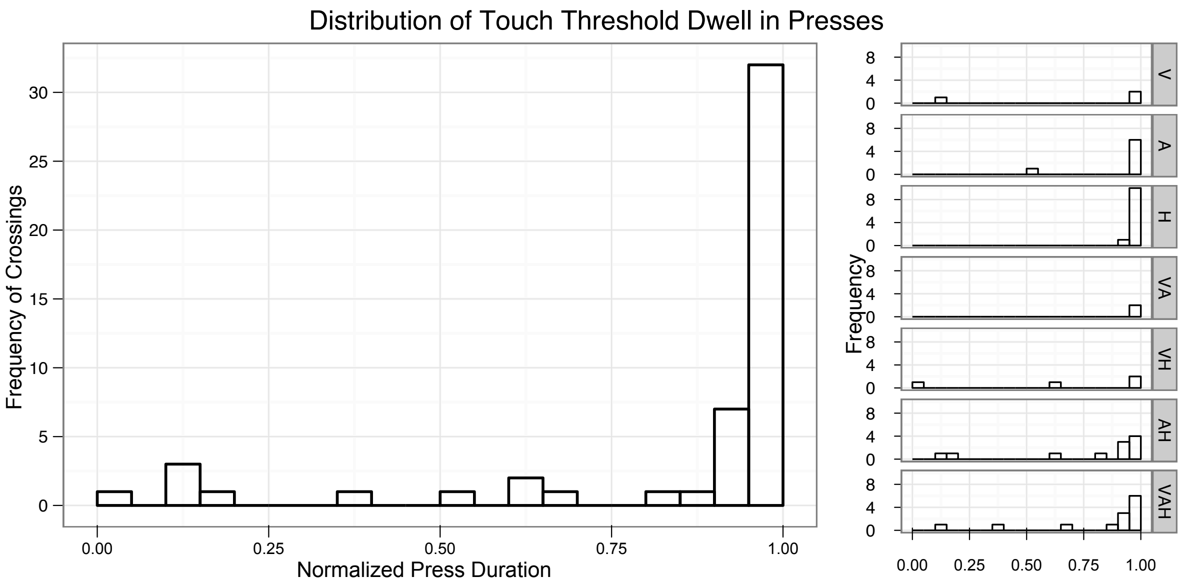

Figure 6.8 Distribution of touch threshold dwell over the button press duration

6.4.4 Peak Button Travel

The summary of the peak button travel appears in Figure 6.10. Since the presses had different durations, we normalized the time that the peak occurred. The vertical line in the graphs indicates the overall mean for peak travel time (m=0.612, sd=0.114) compared to the distribution of the peaks in each condition. The graph indicates that participants usually spent more of the press duration completing the downstroke of the button than the upstroke. This corresponds to the valley observed in the actuation dwell time distribution in Figure 6.9.

A repeated measures ANOVA yielded a significant effect of the feedback condition on the normalized peak location ($F_{6,114}$=18.257, p<0.001). A post hoc Tukey HSD test revealed significant differences between the pairs of conditions listed in Table 6.3. There were no other significant differences in the means of the normalized peaks between other pairs of conditions.

Table 6.3 Press peak comparison (Tukey HSD)

Earlier Peak		Later Peak
{ V, A, VA, VH, VAH }	<	H
{ V, A, VA, VH, VAH }	<	AH
V	<	VH

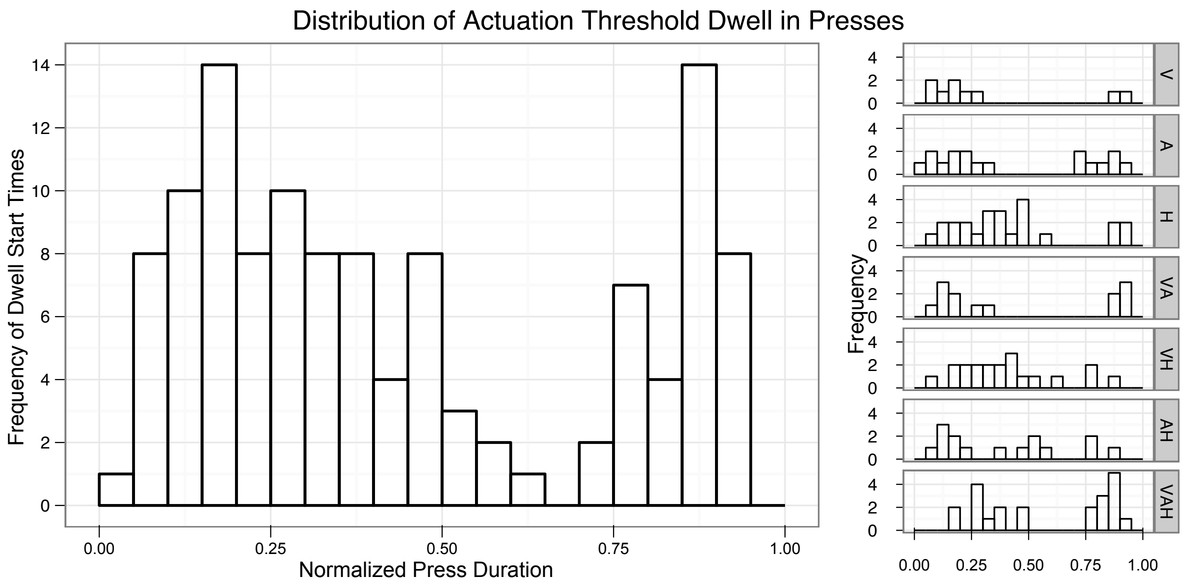

Figure 6.9 Distribution of actuation threshold dwell over the button press duration

6.5 Discussion

6.5.1 Preventing Confusing Feedback Cues

RQ6. In a virtual button, what effect does hysteresis at the touch and actuation thresholds have on preventing confusing feedback signals?

The results suggest that both threshold toggles and threshold dwell occur in virtual button feedback. The study found that touch threshold crossings occurred more frequently than actuation threshold crossings in our sample of presses. We also observed a higher occurrence of actuation threshold dwell compared to touch threshold dwell. The higher occurrence of touch threshold crossings suggests that some method of preventing repeated touch and release feedback at the touch threshold would be beneficial.

These results suggest that hysteresis is well suited to preventing confusing feedback from extra touch threshold crossings and threshold dwell. Although interlocks are also suited to preventing repeated actuations, it would be difficult to design interlocks to effectively control repeated feedback. A shared interlock for touch and release event feedback might miss the release event after an short exploratory press. Separate interlocks for touch and release events

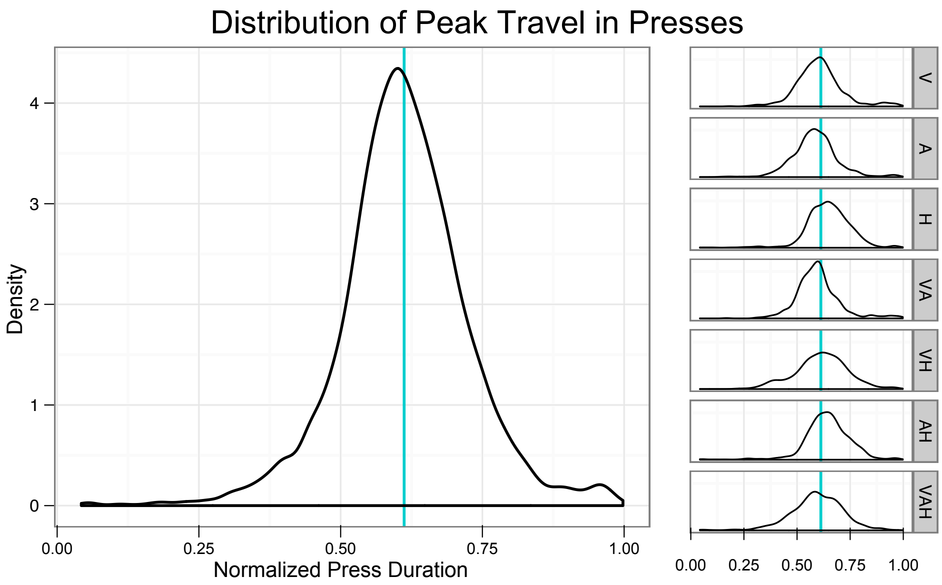

Figure 6.10 Distribution of the peak travel for all presses (left) and by condition (right). The vertical line indicates the overall mean compared to each condition.

might render a second touch event as the user releases the button after the timeout expires. Hysteresis allows both feedback events to be rendered with some separation provided by the users own motion.

6.5.2 Effect of Feedback on Threshold Crossings and Dwell

RQ7. What effect does the combination of feedback modalities used in virtual buttons have on the threshold crossings and threshold dwell?

The study did not discover strong effects of the feedback condition on threshold toggles and threshold dwell. The sample sizes of presses with crossings and dwell were limited in some conditions. However, the results did suggest differences in where the peak travel of the button press occurred between some conditions. This observation also helps to explain why the valley in the actuation dwell graph appears slightly past the middle of the press. Therefore, we might expect differences in the time actuation dwell occurs consistent with the later peak travel observed in some conditions.

6.5.3 Generalizability

Of the two sequences measured, threshold dwell is more likely to depend on the individual system rendering the virtual button feedback. The majority of the dwell recorded during the study occurred between updates that were faster than intended haptic update rate of 1000 Hz. Since this study ran on a system where the upper and lower limits of the update rate were not guaranteed, this is not surprising. However, we did record a press with threshold dwell that lasted 0.024 seconds, which would exceed the 0.016 seconds between 60 Hz updates.

6.6 Conclusions and Future Work

The results suggest that hysteresis at the touch threshold is beneficial for preventing repeated touch and release feedback cues that could cause confusion. This research also presented a model of virtual button feedback that included touch hysteresis.

It is important to prevent repeated feedback events when designing a system for virtual button feedback. Although the frequencies of adverse events were low in the individual presses, tasks would likely require users to interact with multiple virtual buttons. If the feedback appears random or inconsistent to the user, the user will grow to distrust the interface.

CHAPTER 7. CONCLUSION

This research presented a framework for virtual button feedback and the results of three studies of virtual button feedback. This section will summarize the differences between studies, revisit the research questions presented in the study, discuss limitations of the research, and summarize the contribution of this research.

7.1 Differences Between Studies

In the time between the first two studies, several adjustments were made to improve the methods for the follow-up study. The major change to the follow-up study was the recording of time, position, and force during the motion of each button presses. In the first study, no motion data was collected and the only information about each press was the timestamps of touch, release, actuation, and deactuation events.

7.1.1 Implementation Differences Between Studies

In addition to the change in data recording, Study 2 also used a button feedback model that included hysteresis at both the touch and actuation thresholds. The buttons in Study 1 used a very small hysteresis that checked if the travel was greater than or equal to the threshold on the downstroke, and less than the threshold on the upstroke. Study 1 also used the slower update of the main application thread as a short interlock. In Study 2, the buttons used the virtual button feedback model described in Chapter 6, which included a 3.0 mm hysteresis at the touch and actuation thresholds.

During Study 1, we observed that users kept the stylus toward the center of the workspace between tasks, where the buttons would appear. We added a 0.6 second animation between the phases of the task, so that the keypad and Likert rating buttons would appear toward the

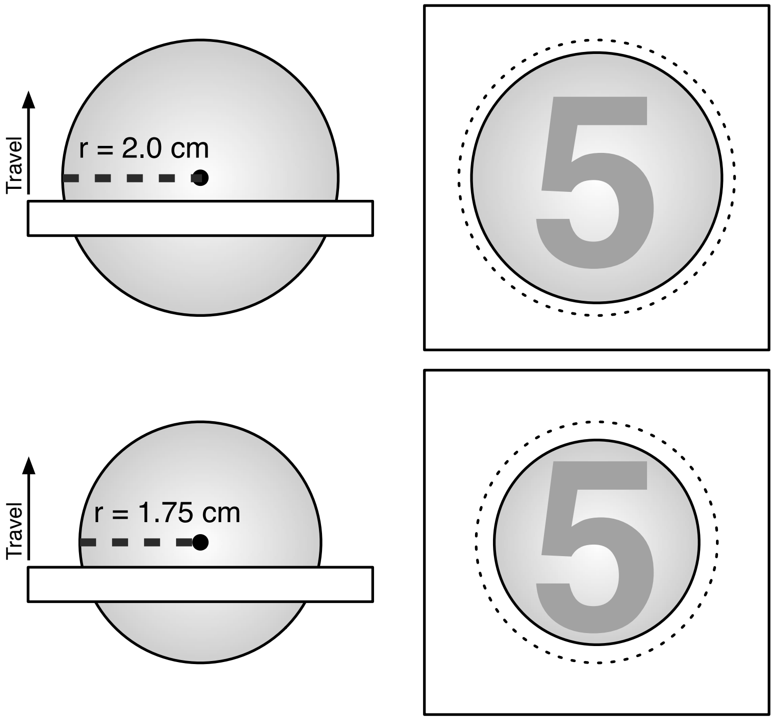

Figure 7.1 **Top**: Study 1 used a 2.0 cm sphere as the stylus. When the center of the sphere was behind the button key, a large area of the stylus would still be visible from the front. The dotted circle represents the full size of the stylus in front of the button key. **Bottom**: Study 2 used a 1.75 cm sphere, which reduced the visible portion of the sphere in front of the button key when the sphere was behind the button key while still keeping a finger-sized stylus.

back of the workspace and move into position. We did not observe participants accidentally activating a button because the keypad appeared next to the stylus in Study 1, but we thought the animation might help participants keep the stylus in front of the keypad when starting the task.

The travel distance of the buttons in Study 2 increased from 1.5 cm to 2.0 cm. We chose 1.5 cm travel for the first study because participants commented the feedback felt too abrupt in the pilot study. We observed that participants appeared to press the buttons farther past the bottom than we anticipated during Study 1, so we tried increasing the travel to 2.0 cm. A longer travel allowed more time between the start of the feedback and reaching the bottom of the button (see Figure 5.10). In Study 2, participants still pressed the button past the bottom, so it is unlikely that increasing the distance between touch feedback and the bottom of the button had any significant effect.

In the pilot study, we used the stylus included with H3DAPI, which is a cylinder with a small sphere at one end (Figure 4.5 Left). The haptic interaction would occur at the point in the center of the small sphere. This did not match the cognitive model that the users had when interacting with the virtual buttons because the users would try to press the buttons with the cylinder. In Study 1, we changed the stylus to a 2.0 cm sphere, where the haptic interaction would occur at the center of the sphere. This better matched the cognitive model of the users, but there was still a problem where it might appear that the stylus was in front of the button when interaction point was behind the button (Figure 7.1). For Study 2, we reduced the size of the stylus sphere to 1.75 cm. This maintained a visible stylus while reducing the portion of the sphere that would intersect the front of the button key.

7.1.2 Task Success between Studies

To determine the effect of the feedback condition on the successful completion of a task, we performed a logistic regression on the task success data from Study 2. The logistic regression model used the feedback condition presented by the buttons to predict the successful completion of the task, shown in Equation 7.1. We then rotated the condition used as the intercept, β_0, to obtain a comparison of success between conditions. The results from Study 1 are presented in Table 4.4, and the results from Study 2 are presented in Table 7.1. A Chi-squared test of goodness-of-fit failed to reject the model (χ^2(6, N=700) = 15.02, p=0.020), and the Log Likelihood of the model was -295.24 with 7 degrees of freedom.

$$logit(\pi) = \beta_0 V + \beta_1 A + \beta_2 H + \beta_3 VA + \beta_4 VH + \beta_5 AH + \beta_6 VAH \tag{7.1}$$

As in Study 1, the results suggest that the odds of successfully completing the task improve when switching from unimodal haptic feedback to either the VH or AH conditions. In both studies, the odds of success also improved when switching from the bimodal VA condition to either VH or AH conditions. Unlike Study 1, the results from Study 2 show an improvement in the odds of success when adding auditory feedback to the unimodal haptic condition (H → AH). The significant differences in task success for the trimodal condition observed in Study 1

Table 7.1 Odds ratios by condition

			Odds Ratio						
Intercept	Successes	Failures	V	A	H	VA	VH	AH	VAH
V	84	16	–	0.87	0.60	0.81	1.93	2.19	1.08
A	82	18		–	0.70	0.94	2.22	2.52*	1.24
H	76	24			–	1.35	3.19**	3.63**	1.79
VA	81	19				–	2.37*	2.70*	1.33
VH	91	9					–	1.14	0.56
AH	92	8						–	0.49
VAH	85	15							–

Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$

were not observed in Study 2. Although more tasks were completed successfully with unimodal auditory feedback than with haptic feedback in Study 2, the results do not show the significant difference in the odds of success observed in Study 1.

7.1.3 Success by Modality

In Study 2, several of the participants commented that they were looking for cues from a particular modality. We did not collect as many comments in Study 1, but one participant made a similar comment to the ones reported in Section 5.4.4. "I was always looking for visual cues first, then audio, then haptic" (P1). The results of the feedback modality on task success appear in Figure 7.2.

We performed a logistic regression for each participant to investigate the effect of the modality on task success for each participant. The logistic regression model used the presence of each modality during the task's condition to predict task success. We performed a Chi-squared test on the logistic regression model for each participant, and rejected any model which did not meet an α of 0.05. Although we observed two of the participant's models from Study 2 passed this significance test, all of the participant's models from Study 1 were rejected. Despite anecdotal reports from participants in both studies that they were looking for cues from a particular modality, the results suggest that it is rare for a participant to be more successful with a particular modality.

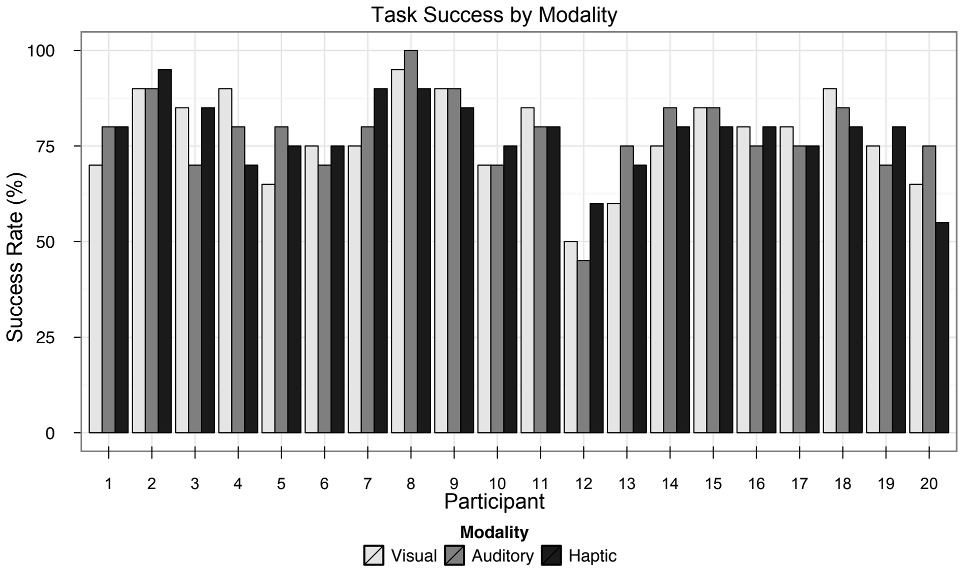

Figure 7.2 Task success by modality for each participant in Study 1

7.1.4 Participants

Between the groups of participants who volunteered for the first and follow-up studies, there were a few small differences. Study 1 had a more even distribution of ages because Study 2 had a larger number of 24-29 year olds volunteer. None of the participants had previous experience with haptics in Study 1, but the participants for the follow-up studies had a number of participants report previous experience with haptics. However, when we explored adding haptic experience to the models used in the evaluation of task success, errors, and time on task, we did not see significant results. This could be the result of the uneven group sizes, and it also might be that the survey question did not accurately capture the participants' familiarity.

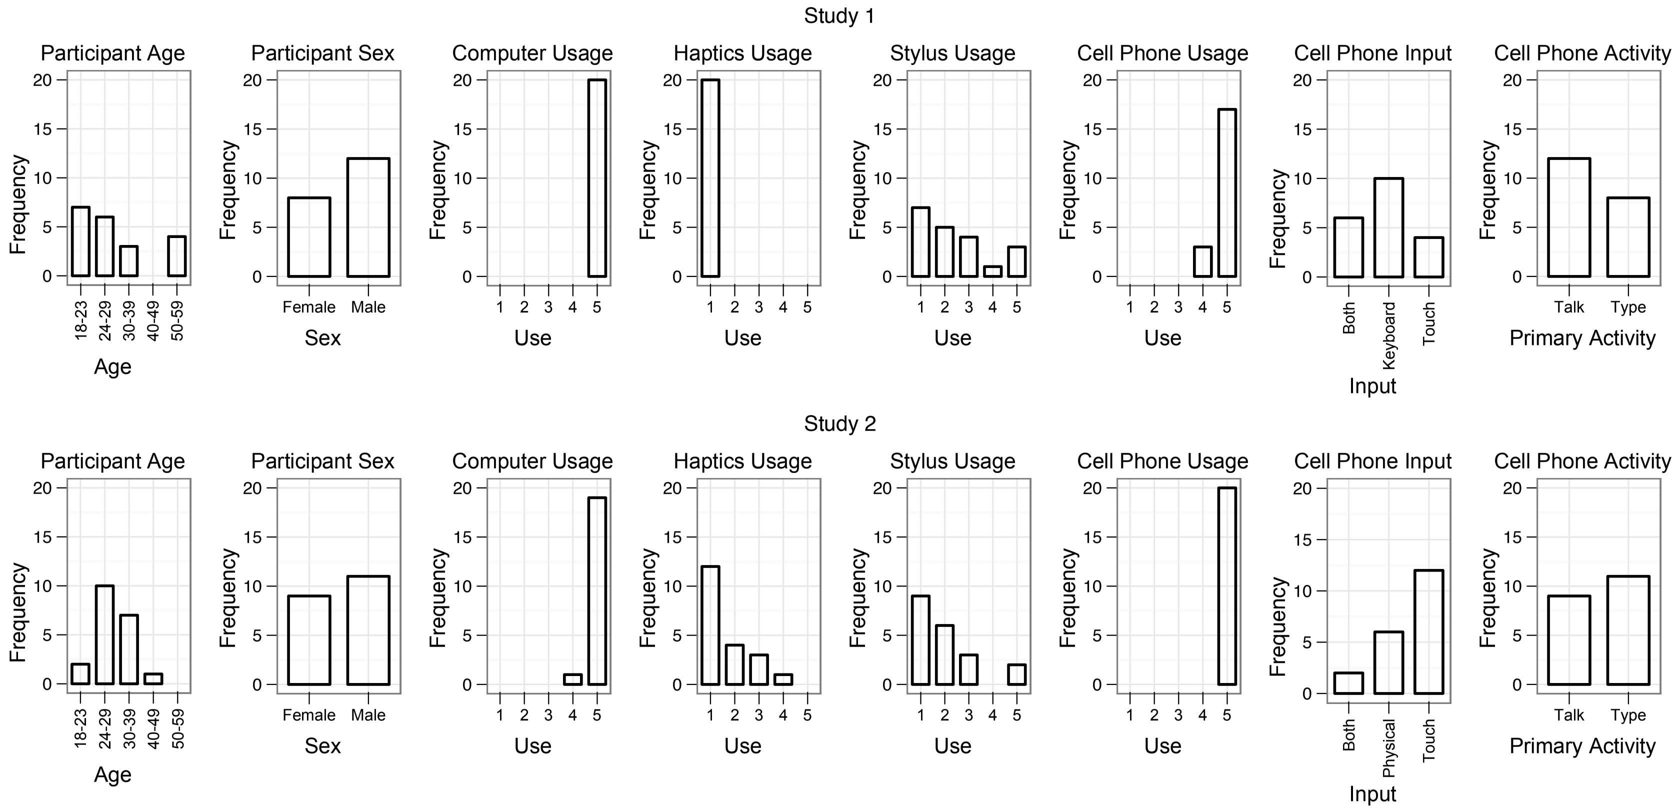

Figure 7.3 A comparison of participant demographics between studies.

7.2 Summary of Major Findings

7.2.1 Press Motion Characteristics

RQ4. How does the combination of feedback modalities affect the motion of the 3D button press?

We expected that the resistance provided by haptic feedback would shorten the maximum distance that the user pressed the button. However, these shorter press depths were not observed in the unimodal haptic condition. Instead, the presses were only significantly shorter in those bimodal and trimodal conditions that included haptic feedback. Visual feedback with haptic feedback significantly shortened the button presses, in both the VH and VAH conditions. Auditory feedback also demonstrated a significant effect in shortening the button presses in the AH and VAH conditions.

This is consistent with previous studies that found that visual and auditory feedback increased the perceived stiffness of a haptic spring or a solid haptic surface. However, those studies observed the increase in stiffness when manipulating the perceived displacement of the stylus (Figure 2.2). In this study we did not observe that this illusion led participants to shorten presses in the conditions that manipulated the displacement of the stylus (conditions without visual feedback). In fact, users continued to press the buttons well past the bottom in the H and AH conditions despite any potential increase in the perceived stiffness.

Instead, the results of Study 2 observed the shallowest presses in the VH and VAH conditions in which visual feedback was provided by a color change from gray to yellow. Future studies could investigate whether this type of visual feedback influenced the perceived stiffness of the buttons.

7.2.2 Effect of Feedback Modalities on Performance

RQ1. Will user performance improve as the number of feedback modalities from virtual buttons increases?

RQ2. Compared to trimodal feedback, how does reducing the number of feedback modalities affect task accuracy and completion time?

One might expect multimodal feedback to improve task performance over unimodal feedback and a reduction of feedback channels to adversely affect performance. Participants did show significant improvements in completion time with some combinations of multimodal feedback. However the results were not consistent between Study 1 and Study 2. In Study 1, the completion times in the VH and VAH conditions were indeed shorter than the separate V and H conditions. The AH condition also showed a significant improvement in completion time over the V condition. In Study 2, the AH condition showed an improvement over each condition except for VH.

The results of both studies do suggest that completion time improved with multimodal feedback, and the haptic modality was present in each of the multimodal conditions that showed a significant reduction in completion times. One explanation is that the resistance provided by haptic feedback helped the user stop near the bottom of the button and move onto the next press more quickly. However, in Study 2, the H and AH conditions did not result in shallower presses than the conditions without haptic feedback.

The overall results from both studies do not suggest clear improvements in completion time between specific combinations of multimodal feedback. This is consistent with prior studies that found multimodal feedback had a more significant effect in reducing errors rather than improving completion times (Dumas et al., 2009; Oviatt, 2007).

It was unexpected that the results showed multimodal feedback did not have a significant effect in reducing the total errors during the tasks. Instead, the results of Study 1 found a significant increase in total errors for the trimodal condition. Similarly, the results of Study 2 showed a significant increase in delete errors in the trimodal condition, but no significant differences in total errors. In both studies, the trimodal condition resulted in a significant increase in the number of missing digits for phone numbers dialed and showed a strong correlation between the missing digits and unactuated presses.

One explanation for the unexpected increase in errors might be that the feedback from one of the modalities was poorly designed and that lag between two modalities caused participants to make more errors in the trimodal condition. If that were true, we would expect to see increased errors in the bimodal conditions where the lagging modality was present. For example, if auditory feedback lagged, the VA, AH, and VAH conditions should all show significantly more errors. Instead, the results only showed increased errors during the trimodal feedback condition.

Another explanation is that multimodal feedback increased the task's cognitive load, causing participants to make more mistakes in the VAH condition. If this were true, we would expect the increase in errors observed during the trimodal feedback condition. However, we would expect these errors to be distributed equally over insert, delete and replace errors, and we would not expect a disproportionate increase in the number of unactuated presses. Instead, the results do show a disproportionate increase in both delete errors and unactuated presses in the trimodal condition. Neither explanation accounts for the significant increase in unactuated presses observed in the trimodal condition.

Therefore, we hypothesize that the combined visual, auditory, and haptic feedback gave participants an erroneous impression that they actuated the button before reaching the actuation point. The results of Study 2 suggest that the button presses in the VAH condition were significantly shallower than in each other condition except for VH. This may have led participants to release the button before reaching the actuation point in the trimodal condition, believing that they had already actuated the button. This explanation is consistent with the increase in unactuated presses and delete errors observed in the trimodal condition during both studies.

Our results are consistent with multiple resource theory because there was no evidence of contention for shared resources. If there was contention for resources to process trimodal feedback, we might expect the participants to continue to press the button until they processed the feedback, or repeat the button press a second time because they missed cues from the feedback. We would not expect a different result during the trimodal condition because it suggests the user perceived trimodal feedback as distinct from the other conditions. The results suggest that users were able to perceive multimodal feedback without contention for shared resources.

7.2.3 Subjective Preference for Multimodal Feedback

RQ3. Will participants prefer virtual buttons that provide feedback using a higher number of *modalities?*

Participants expressed a strong preference for multimodal feedback from virtual buttons in both Study 1 and Study 2. This is consistent with earlier findings that users prefer multimodal feedback over unimodal feedback (Dumas et al., 2009; Oviatt, 2007). In both studies, VA feedback was the only bimodal condition which did not score significantly higher than all of the unimodal conditions in the subjective ratings (Figure 4.8 and Figure 5.8). The VH and VAH conditions also scored significantly higher than the VA condition in both studies. This suggests that users do prefer some combinations of multimodal feedback over other combinations of multimodal feedback.

As both of the studies used a similar feedback design, it is possible the subjective ratings were influenced by the design of the cues for one of the modalities. If this were true, we would expect a significant difference in the ratings between the three unimodal conditions. However, the results did not show a significant difference in the ratings of the unimodal conditions.

7.2.4 Participant Success with a Particular Modality

RQ5. When users press a button, are they focusing on a particular type of feedback such as a visual cue?

Although participants in both studies reported looking for cues from a particular modality, this research found little evidence that the modalities present had a significant effect on the participants' success. Between the forty participants who took part in the two studies, only two succeeded more often when feedback from a particular modality was present. None of the participants who reported focusing on cues from a particular modality were more successful with that modality.

One explanation for this mismatch between reported attention and performance may be that these participants initially looked for cues from a particular modality, but quickly adapted to the cues from the modalities provided during the task. It might also be that the cues from a particular modality were more salient to those two participants who were more successful with that modality. The results do not suggest that a particular modality such as vision was vital to completing the tasks successfully. Instead, participants demonstrated shorter completion times and more compact button presses with multimodal feedback.

7.2.5 Preventing Confusing Repeated Feedback Cues

RQ6. In a virtual button, what effect does hysteresis at the touch and actuation thresholds have on preventing confusing feedback signals?

The results from Chapter 6 showed that extra touch threshold crossings occurred more often than extra actuation threshold crossings. This makes sense in a 3D force-feedback environment because touch is an exploratory sense. These results also suggest that virtual buttons benefit from some method of preventing repeated touch and release feedback signals from what the user perceives to be a single press. Hysteresis at the touch and actuation thresholds would help to prevent repeated feedback cues from single button presses.

RQ7. What effect does the combination of feedback modalities used in virtual buttons have on the threshold crossings and threshold dwell?

The feedback modalities rendered by a virtual button did affect the motion of the individual button presses in Study 2. However, the results of Study 3 did not find strong evidence that the feedback condition affected extra threshold crossings and threshold dwell. The peak of the button travel was significantly different in some conditions, and this might influence when these adverse events occur during the press.

7.3 Limitations

The results of this research are expected to apply to scenarios where users interact with buttons in a 3DUI at a single point, such as when using a stylus-based device. Although Klatzky and Lederman (2008) documented differences between interacting with a stylus and a whole hand, it is not clear how these differences would affect button presses in a 3DUI. The exploratory procedure of applying pressure plays a large role in interaction with buttons, and this procedure is possible with a force-feedback stylus. However, it is not clear how the availability of other exploratory procedures in a richer virtual environment might change the experience of interacting with a button.

Another limitation of this research is the longer button travel (1.5 and 2.0 cm) used in the studies. We chose the longer travel distance because participants in the pilot study felt the feedback was too abrupt with a more typical travel distance for a physical keyboard. This phenomena seemed related to the limited stiffness the haptic device used in the experiment could render. An improved haptic device may be able to render a smoother transition in forces over the shorter periods of time required by a shorter travel distance.

The visual and haptic feedback signals were not collocated during the studies presented in this research (Figure 4.2). Participants appeared to adapt immediately to the experimental setup and kept their attention on the laptop display during the tasks. It is not clear whether any differences would result from using a setup with collocated visual and haptic display.

7.4 Summary of Contributions

This research presented a framework for virtual button feedback and a model for virtual button feedback which includes touch hysteresis. The studies evaluated effects of feedback modalities on performance, subjective rating, the individual motion of presses, and the effectiveness of touch hysteresis, and resulted in the following contributions to the field:

1. Evaluation of the effect of feedback modalities on task performance

   The results of the studies in Chapters 4 and 5 show that some combinations of multimodal feedback improved completion times but also showed a counterintuitive increase in the number of errors during the trimodal condition which were caused by a missing digit. The shallower presses observed during the trimodal condition in Study 2 provide support for the hypothesis that this was caused by participants releasing the buttons early because the trimodal cues caused them to believe that they had already actuated the button.

2. Evaluation of the effect of feedback modalities on the motion of individual presses

   The resistance provided by haptic feedback alone did not prevent users from pressing well past the virtual bottom of the buttons. However, the addition of visual or auditory feedback to haptic feedback did result in shallower press motions, and in improvements to task completion times.

3. Evaluation of subjective rating of feedback modalities from virtual buttons.

   Participants expressed a preference for multimodal feedback from virtual buttons in both studies (Chapter 4 and Chapter 5). Participants gave higher ratings to some combinations of multimodal feedback than to others.

4. Evaluation of touch hysteresis for virtual button feedback

   The results of the study in Chapter 6 suggested that touch hysteresis would help to prevent some of the adverse events that could lead to confusing feedback from virtual buttons, and that virtual buttons should incorporate touch hysteresis into their design.

5. A framework for virtual buttons with multimodal feedback that includes touch hysteresis

   A thorough quantitative comparison of the effects of feedback from single and multiple sensory modalities requires an effective research framework. Chapter 3 presented a framework for designing virtual buttons that provide any combination of visual, auditory and haptic feedback. Chapter 6 added touch hysteresis to the framework. The studies in Chapters 4, 5, and 6 used this framework as a basis for creating virtual button feedback.

These contributions outline the complexity involved in transferring what appears to be a simple multimodal interaction to a virtual environment. The research revealed a pair of unintuitive results. First, haptic feedback alone was not enough to constrain the motion of a virtual button press, but it did shorten the presses when combined with visual or auditory feedback. Second, feedback that worked well in unimodal and bimodal conditions resulted in missing digits during the trimodal condition. These counterintuitive results contribute to the understanding of multimodal feedback in the design of virtual buttons.

7.5 Conclusions and Future Work

Virtual buttons are used in interfaces developed to support 3D tasks in medicine, manufacturing, geoscience, and engineering. Although other forms of 3D interactions are emerging, buttons retain the advantages of discoverability and perceived affordance which make them integral components in an interface.

This research identified several areas where the results observed in these studies were unexpected. Future studies could investigate whether the types of visual feedback used in this research influenced the user's perceived stiffness of the buttons. It would also be interesting to observe whether users are more successful with particular combinations of modalities over a longer period of time, or in situations when their attention is occupied with another task. Future studies might also investigate whether the results observed from virtual button feedback in these studies transfer to similar 3D interfaces, such as a menu system composed of virtual buttons.

As virtual buttons become more prevalent in supporting 3D tasks, users will seek out the multimodal feedback that they prefer. Compared to mechanical buttons, virtual buttons do not always provide the same feedback modalities and fidelity. In these studies, a difference in the way that trimodal feedback was perceived caused users to shorten their button presses and miss digits from the phone number. A better understanding of multimodal feedback will help to improve future interactions with virtual interfaces.

APPENDIX A. SURVEYS FOR SELF-REPORTED DATA COLLECTION

Pre-Experiment Questionnaire

You may skip any question that you do not wish to answer or that makes you feel uncomfortable. If you have any questions, please ask one of the researchers.

Age: ____________

Sex: Male or Female

1. How often do you use a haptic device such as the SensAble Phantom or the Novint Falcon?
   (a) Never
   (b) Once a month or less often
   (c) Once per week
   (d) Once per day
   (e) Several times throughout the day
2. How often do you use a computer?
   (a) Never
   (b) Once a month or less often
   (c) Once per week
   (d) Once per day
   (e) Several times throughout the day
3. How often do you use a stylus on a PDA, tablet computer, drawing tablet or other device?
   (a) Never
   (b) Once a month or less often
   (c) Once per week
   (d) Once per day
   (e) Several times throughout the day
4. If you use a cell phone, which do you use more?
   (a) Physical keypad or keyboard
   (b) Touch screen (flat surface that provides virtual buttons)
   (c) Both
   (d) I do not use a cell phone

5. If you use a cell phone, what is your main activity?
    (a) Talking
    (b) Typing (text messages, email, internet use)
    (c) I do not use a cell phone

6. How often do you use a device that provides vibration feedback such as the Nintendo Wii?
    (a) Never
    (b) Once a month or less often
    (c) Once per week
    (d) Once per day
    (e) Several times throughout the day

7. How many undergraduate programming courses have you taken?

    ____

8. Do you have any conditions that would impair your ability to clearly see the computer screen, hear sounds through the headphones, or distinguish between colors today?

    ____

Post-Study Exit Survey

You may skip any question that you do not wish to answer or that makes you feel uncomfortable. If you have any questions, please ask one of the researchers.

In this experiment, you experienced virtual buttons that provided visual feedback (by changing color), audio feedback (by playing a sound), or haptic feedback (by resisting the force you applied). For the next three questions, please consider the visual, audio, and haptic feedback separately.

1. The button changing color effectively communicated when I activated a button.

    Strongly disagree ○ ○ ○ ○ ○ ○ ○ Strongly agree

2. The button playing a sound effectively communicated when I activated a button.

    Strongly disagree ○ ○ ○ ○ ○ ○ ○ Strongly agree

3. The button resisting the force I applied effectively communicated when I activated a button.

    Strongly disagree ○ ○ ○ ○ ○ ○ ○ Strongly agree

4. Please rank which mode of feedback you feel was most effective with a (1), next most effective with a (2), and least effective with a (3).

    ______ Audio

    ______ Visual

    ______ Haptic

5. How seriously did you take the tasks in this experiment?
   (a) Not at all seriously
   (b) Not very seriously
   (c) Neutral
   (d) Somewhat seriously
   (e) Very seriously

6. Please share any additional comments or suggestions for improvements in the space below

APPENDIX B. EXIT SURVEY RESULTS MODALITIES AND CONDITION RANKINGS

Participants provided the following responses to rate the overall feedback modalities and rank the feedback conditions on the exit surveys. This data was consistent with the subjective ratings we collected after each task.

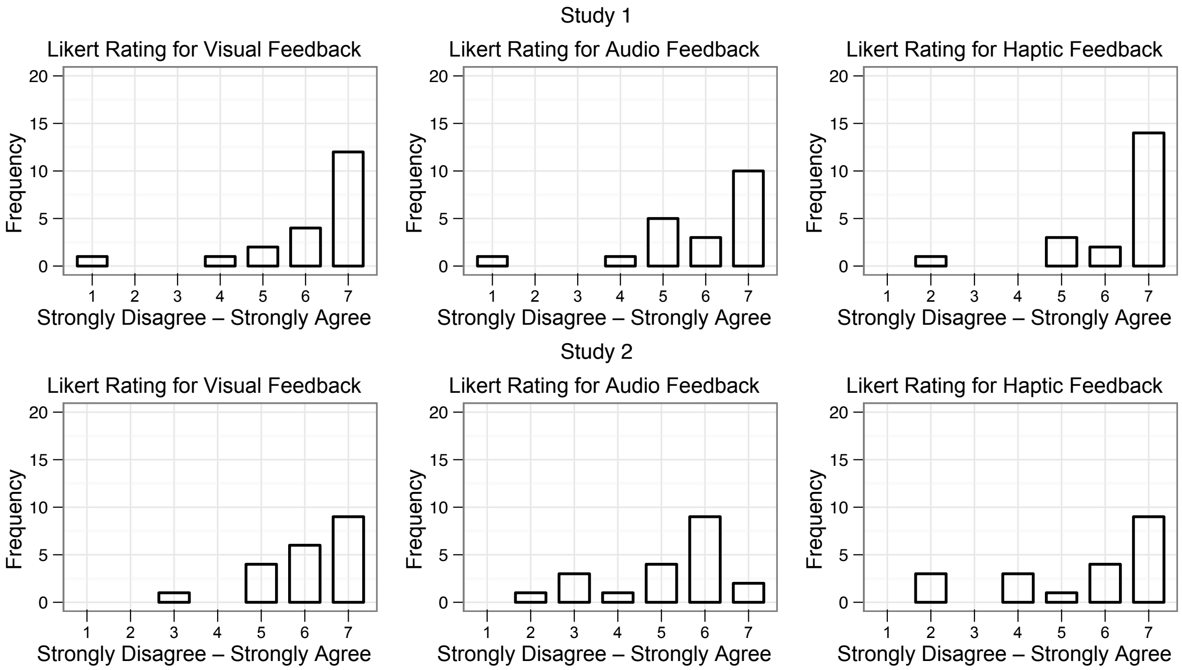

Figure B.1 Participants' overall ratings of each modality from the exit survey

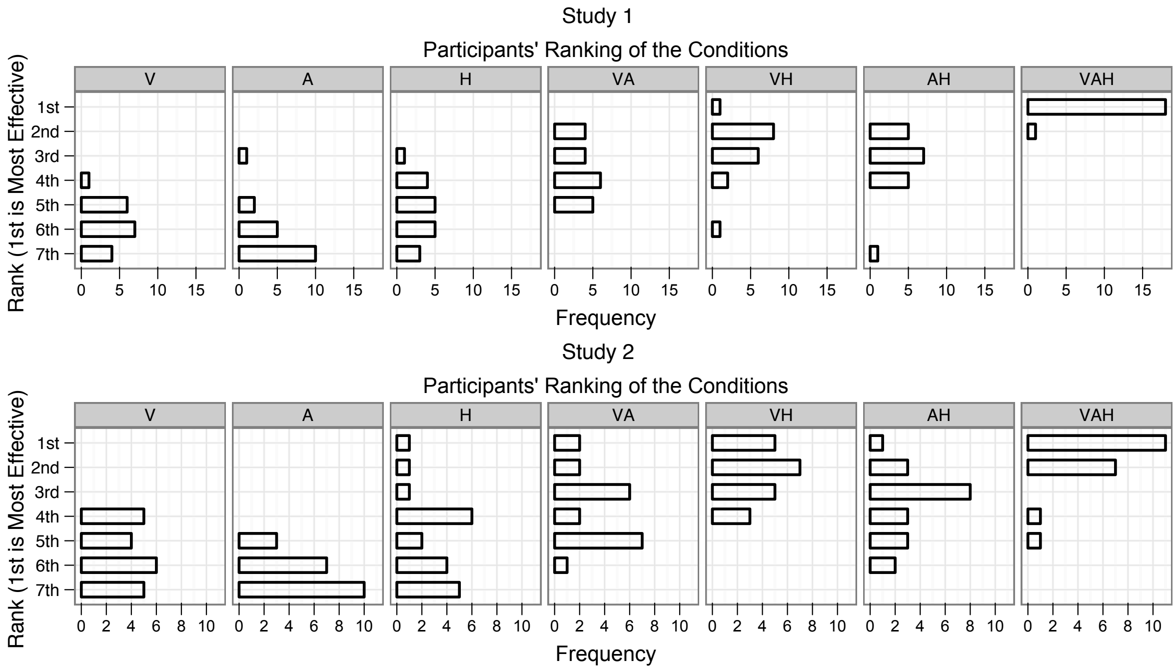

Figure B.2 Distribution of participants' forced rankings of feedback conditions from the exit survey

BIBLIOGRAPHY

Akamatsu, M., Mackenzie, I. S., and Hasbroucq, T. (1995). A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. *Ergonomics*, 38(4):816–827.

Altinsoy, M.E., B. J. . T. C. (2001). Inter- modal effects of non-simultaneous stimulus presentation. In Alippi, A., editor, *Proceedings of the 17 th International Congress on Acoustics*.

Altinsoy, E. M. (2003). Perceptual aspects of auditory-tactile asynchrony. In *Proc. ICSV 10*, pages 3831–3838.

Bodden, M., . I. H. (2002). Active sound design: Vacuum cleaner. *Revista de Acustica*, 33.

Boeck, J., Raymaekers, C., and Coninx, K. (2006). Exploiting proprioception to improve haptic interaction in a virtual environment. *Presence: Teleoperators & Virtual Environments*, 15(6):627–636.

Boles, D. B., Bursk, J. H., Phillips, J. B., and Perdelwitz, J. R. (2007). Predicting dual-task performance with the multiple resources questionnaire (mrq). *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 49(1):32–45.

Burke, J. L., Prewett, M. S., Gray, A. A., Yang, L., Stilson, F. R. B., Coovert, M. D., Elliot, L. R., and Redden, E. (2006). Comparing the effects of visual-auditory and visual-tactile feedback on user performance: a meta-analysis. In *Proc. ICMI '06*. ACM.

Chu, G., Moscovich, T., and Balakrishnan, R. (2009). Haptic conviction widgets. In *GI '09: Proceedings of Graphics Interface 2009*, pages 207–210, Toronto, Ont., Canada, Canada.

Clare, C. (1976). Human factors: A most important ingredient in keyboard designs. *EDN Magazine*, 21(8):99–102.

Cockburn, A. and Brewster, S. (2005). Multimodal feedback for the acquisition of small targets. *Ergonomics*, 48(9):1129–1150.

Cohen, K. M. (1982). Membrane keyboards and human performance. *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 26:424–424.

Colavita, F. (1974). Human sensory dominance. *Perception and Psycophysics*, 16(2):409–412.

Colton, M. and Hollerbach, J. (2007). Reality-based haptic force models of buttons and switches. In *IEEE Robotics and Automation 2007*, pages 497–502.

Díaz, I. n., Hernantes, J., Mansa, I., Lozano, A., Borro, D., Gil, J. J., and Sánchez, E. (2006). Influence of multisensory feedback on haptic accessibility tasks. *Virtual Reality*, 10(1):31–40.

Diederich, A. and Colonius, H. (2004). Bimodal and trimodal multisensory enhancement: Effects of stimulus onset and intensity on reaction time. *Perception & psychophysics*, 66(8):1388–1404.

Dietz, P. H., Eidelson, B., Westhues, J., and Bathiche, S. (2009). A practical pressure sensitive computer keyboard. In *UIST '09: Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 55–58, New York, NY, USA. ACM.

DiFranco, D. E., Beauregard, G. L., and Srinivasan, M. A. (1997). The effect of auditory cues on the haptic perception of stiffness in virtual environments. In *ASME Dynamic Systems and Control Division*, volume 61, pages 17–22. ASME.

Dumas, B., Lalanne, D., and Oviatt, S. (2009). Multimodal interfaces: A survey of principles, models and frameworks. In Lalanne, D. and Kohlas, J., editors, *Human Machine Interaction*, volume 5440 of *Lecture Notes in Computer Science*, pages 3–26. Springer Berlin / Heidelberg.

Ferris, T. K. and Sarter, N. B. (2008). Cross-Modal Links Among Vision, Audition, and Touch in Complex Environments. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(1):17–26.

Forlines, C. and Balakrishnan, R. (2008). Evaluating tactile feedback and direct vs. indirect stylus input in pointing and crossing selection tasks. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, pages 1563–1572, New York, NY, USA. ACM.

Forster, B., Cavina-Pratesi, C., Aglioti, S. M., and Berlucchi, G. (2002). Redundant target effect and intersensory facilitation from visual-tactile interactions in simple reaction time. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 143(4):480–487.

Frings, C. and Spence, C. (2010). Crossmodal congruency effects based on stimulus identity. *Brain Research*, 1354(0):113–122.

Fukumoto, M. and Sugimura, T. (2001). Active click: tactile feedback for touch panels. In *ACM CHI '01 extended abstracts on Human factors in computing systems*, pages 121–122.

Geldard, F. (1960). Some neglected possibilities of communication. *Science*, 131(3413):1583–1588.

Giard, M. and Peronnet, F. (1999). Auditory-Visual Integration during Multimodal Object Recognition in Humans: A Behavioral and.. *Journal of Cognitive Neuroscience*, 11(5):473–489.

Green, K. (1996). Studies of the McGurk effect: implications for theories of speech perception. In *Spoken Language, 1996. ICSLP 96. Proceedings., Fourth International Conference on*, pages 1652–1655.

Guiard, Y. (1987). Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behavior*, 19:486–517.

H3DAPI (2012). Open source haptics. http://www.h3dapi.org.

Hall, M., Hoggan, E., and Brewster, S. (2008). T-bars: towards tactile user interfaces for mobile touchscreens. In *MobileHCI '08*, pages 411–414.

Harrison, C., Tan, D., and Morris, D. (2010). Skinput: Appropriating the body as an input surface. In *CHI'10*.

Hecht, D. and Reiner, M. (2009). Sensory dominance in combinations of audio, visual and haptic stimuli. *Experimental brain research*, 193:307–314.

Hecht, D. and Reiner, M. (2010). Stroop Interference and Facilitation Effects in Kinesthetic and Haptic Tasks. *Advances in Human-Computer Interaction*, 2010:1–10.

Hecht, D., Reiner, M., and Karni, A. (2008a). Enhancement of response times to bi- and tri-modal sensory stimuli during active movements. *Experimental brain research*, 185(4):655–665.

Hecht, D., Reiner, M., and Karni, A. (2008b). Multisensory enhancement: gains in choice and in simple response times. *Experimental brain research*, 189(2):133–143.

Heller, M. A., Calcaterra, J. A., Green, S. L., and Brown, L. (1999). Intersensory conflict between vision and touch: the response modality dominates when precise, attention-riveting judgments are required. *Perception & psychophysics*, 61(7):1384–1398.

Hershenson, M. (1962). Reaction time as a measure of intersensory facilitation. *Journal of experimental psychology*, 63:289–293.

Hoggan, E. and Brewster, S. (2006). Crossmodal icons for information display. In *ACM CHI '06 extended abstracts on Human factors in computing systems*, pages 857–862.

Hoggan, E. and Brewster, S. A. (2010). Crosstrainer: testing the use of multimodal interfaces in situ. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 333–342, Atlanta, Georgia, USA. ACM.

Hoggan, E., Crossan, A., Brewster, S., and Kaaresoja, T. (2009). Audio or tactile feedback: which modality when? In *ACM CHI '09*, pages 2253–2256.

Hoggan, E., Kaaresoja, T., Laitinen, P., and Brewster, S. (2008). Crossmodal congruence: the look, feel and sound of touchscreen widgets. *Proc. IMCI '08*, pages 157–164.

Jacko, J., Moloney, K., Kongnakorn, T., Barnard, L., Edwards, P., Leonard, V. K., Sainfort, F., and Scott, I. (2005). Multimodal Feedback as a Solution to Ocular Disease-Based User Performance Decrements in the Absence of Functional Visual Loss. *International Journal of Human-Computer Interaction*, 18(1):183–218.

Jones, L. A. and Sarter, N. B. (2008). Tactile Displays: Guidance for Their Design and Application. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(1):90–111.

Kim, W., Tendick, F., and Stark, L. (1987). Visual enhancements in pick-and-place tasks: Human operators controlling a simulated cylindrical manipulator. *Robotics and Automation, IEEE Journal of*, 3(5):418–425.

Kim, Y., Kim, S., Ha, T., Oakley, I., Woo, W., and Ryu, J. (2006). Air-jet button effects in ar. *LNCS*, 4282:384–391.

Kinkead, R. (1975). Typing Speed, Keying Rates, and Optimal Keyboard Layouts. *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 19(2):159–161.

Kinkead, R. and Gonzalez, B. (1969). Human factors design recommendations for touch-operated keyboards — final report. Technical Report 12091-FR, Honeywell, Inc.

Klatzky, R. and Lederman, S. J. (2003). Touch. In Weiner, I. B., editor, *Handbook of Psychology*, pages 147–176. John Wiley & Sons, Inc., Hoboken, NJ, USA.

Klatzky, R. and Lederman, S. J. (2008). Object Recognition by Touch. In *Blindness and Brain Plasticity in Navigation and Object Perception*, pages 185–207.

Kohlrausch, A. and van de Par, S. (1999). Auditory-visual interaction: from fundamental research in cognitive psychology to (possible) applications. In *Proc. SPIE*, pages 34–44.

Komerska, R. and Ware, C. (2004). A study of haptic linear and pie menus in a 3d fish tank vr environment. In *Proc. 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS '04).*, pages 224–231. IEEE.

Kong, J., Zhang, W. Y., Yu, N., and Xia, X. J. (2011). Design of human-centric adaptive multimodal interfaces. *Journal of Human Computer Studies*, 69(12):854–869.

Kroemer, K. H. (2001). Keyboards and keying An annotated bibliography of the literature from 1878 to 1999. *Universal Access in the Information Society*, 1:99–160.

Kuchenbecker, K., Fiene, J., and Niemeyer, G. (2006). Improving contact realism through event-based haptic feedback. *Visualization and Computer Graphics, IEEE Transactions on*, 12(2):219–230.

Lécuyer, A., Burkhardt, J.-M., Coquillart, S., and Coiffet, P. (2001). "boundary of illusion:" an experiment of sensory integration with a pseudo-haptic system. In *Proc. VR'01*, VR '01, pages 115–122, Washington, DC, USA. IEEE Computer Society.

Lee, J. C., Dietz, P. H., Leigh, D., Yerazunis, W. S., and Hudson, S. E. (2004). Haptic pen: a tactile feedback stylus for touch screens. In *UIST '04: Proceedings of the 17th annual ACM symposium on User interface software and technology*, pages 291–294, New York, NY, USA. ACM.

Lee, S. and Zhai, S. (2009). The performance of touch screen soft buttons. In *ACM CHI '09*, pages 309–318.

Leung, R., MacLean, K., Bertelsen, M. B., and Saubhasik, M. (2007). Evaluation of haptically augmented touchscreen gui elements under cognitive load. In *Proc. ICMI '07*, pages 374–381, New York, NY, USA. ACM.

Levenshtein, V. (1965). Binary codes capable of correcting deletions, insertions, and reversals. *Problems in Information Transmission*, pages 8–17.

Levesque, V., Oram, L., MacLean, K., Cockburn, A., Marchuk, N. D., Johnson, D., Colgate, J. E., and Peshkin, M. A. (2011). Enhancing physicality in touch interaction with programmable friction. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 2481–2490, New York, NY, USA. ACM.

Lewis, J. R., Potosnak, K. M., and Magyar, R. L. (1997). *Keys and Keyboards*, chapter 54, pages 1285–1315. Elsevier, Amsterdam.

Long, J. (1976). Effects of delayed irregular feedback on unskilled and skilled keying performance. *Ergonomics*, 19(2):183–202.

MacLean, K. (2000). Designing with haptic feedback. In *Robotics and Automation, 2000. Proceedings. ICRA '00. IEEE International Conference on*, pages 783–788 vol.1.

Massie, T. H. and Salisbury, J. K. (1994). The phantom haptic interface: A device for probing virtual objects. In *ASME Winter Annual Meeting*.

McGurk, H. and MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264(5588):746–748.

Meredith, M. A. and Stein, B. E. (1986). Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration. *Journal of Neurophysiology*, 56(3):640–662.

Merlo, J. L., Duley, A. R., and Hancock, P. A. (2010). Cross-modal congruency benefits for combined tactile and visual signaling. *The American Journal of Psychology*, 123(4):413–424.

Microsoft (2012). QueryPerformanceCounter. `http://msdn.microsoft.com/en-us/library/windows/desktop/ms644904`.

Miller, T. and Zeleznik, R. (1998). An insidious haptic invasion: adding force feedback to the x desktop. In *UIST '98: Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 59–64, New York, NY, USA. ACM.

Miner, N., Gillespie, B., and Caudell, T. (1996). Examining the influence of audio and visual stimuli on a haptic display. In *Proc. IMAGE '96*.

Nashel, A. and Razzaque, S. (2003). Tactile virtual buttons for mobile devices. In *ACM CHI '03*, pages 854–855.

Naumann, A. B., Wechsung, I., and Hurtienne, J. (2010). Multimodal interaction: A suitable strategy for including older users? *Interacting with Computers*, 22(6):465–474.

Oulasvirta, A. and Bergstrom-Lehtovirta, J. (2010). A simple index for multimodal flexibility. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, pages 1475–1484, New York, NY, USA. ACM.

Oviatt, S. (2007). *Multimodal Interfaces*, chapter 21, pages 413–432. CRC Press.

Petzold, B., Zaeh, M. F., Faerber, B., Deml, B., Egermeier, H., Schilp, J., and Clarke, S. (2012). A Study on Visual, Auditory, and Haptic Feedback for Assembly Tasks. *Presence: Teleoper. Virtual Environ.*, 13(1):16–21.

Pollard, D. and Cooper, M. (1979). The effect of feedback on keying performance. *Applied Ergonomics*, 10(4):194–200.

Poupyrev, I., Maruyama, S., and Rekimoto, J. (2002). Ambient touch: designing tactile interfaces for handheld devices. In *ACM UIST '02*, pages 51–60.

Pure Data (2012). Pd community site. `http://puredata.info/`.

Python-difflib (2012). difflib – Helpers for computing deltas. `http://docs.python.org/library/difflib.html`.

Ratcliff, J. and Metzener, D. (1988). Pattern matching: the gestalt approach. *Dr. Dobb's Journal*, 13(7):46.

Roe, C. J., Muto, W. H., and Blake, T. (1984). Feedback and key discrimination on membrane keypads. *Human Factors and Ergonomics Society Annual Meeting*, 28:277–281.

Rosenberg, I. and Perlin, K. (2009). The unmousepad: an interpolating multi-touch force-sensing input pad. In *SIGGRAPH '09: ACM SIGGRAPH 2009 papers*, pages 1–9, New York, NY, USA. ACM.

Rosenberg, L. and Brave, S. (1996). Using force feedback to enhance human performance in graphical user interfaces. In *CHI '96: Conference companion on Human factors in computing systems*, pages 291–292, New York, NY, USA. ACM.

Sherrick, C. (1985). A scale for rate of tactual vibration. *The Journal of the Acoustical Society of America*.

Sjöström, C. (2001). Using haptics in computer interfaces for blind people. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 245–246, New York, NY, USA. ACM.

Snibbe, S. S., MacLean, K. E., Shaw, R., Roderick, J., Verplank, W. L., and Scheeff, M. (2001). Haptic techniques for media control. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 199–208, New York, NY, USA. ACM.

Spence, C., Pavani, F., and Driver, J. (2000). Crossmodal links between vision and touch in covert endogenous spatial attention. *Journal of Experimental Psychology: Human Perception and Performance*, 26(4):1298–1319.

Srinivasan, M. A., Beauregard, G. L., and Brock, D. L. (1996). The impact of visual information on the haptic perception of stiffness in virtual environments. In *ASME Winter Annual Meeting*, volume 58, pages 555–559.

Stanney, K., Samman, S., Reeves, L., Hale, K., Buff, W., Bowers, C., Goldiez, B., Nicholson, D., and Lackey, S. (2004). A Paradigm Shift in Interactive Computing: Deriving Multimodal Design Principles from Behavioral and Neurological Foundations. *International Journal of Human-Computer Interaction*, 17(2):229–257.

Stanney, K. M., Mourant, R. R., and Kennedy, R. S. (2012). Human Factors Issues in Virtual Environments: A Review of the Literature. *Presence: Teleoper. Virtual Environ.*, 7(4):327–351.

Steiner, B. A. and Dotson, D. A. (1990). The use of 3-d stereo display of tactical information. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 34(1):36–40.

Sun, M., Ren, X., and Cao, X. (2010). Effects of multimodal error feedback on human performance in steering tasks. *Journal of Information Processing*, 18:284–292.

Swindells, C., MacLean, K. E., Booth, K. S., and Meitner, M. J. (2007). Exploring affective design for physical controls. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 933–942, San Jose, California, USA. ACM Request Permissions.

Tashiro, K., Shiokawa, Y., Aono, T., and Maeno, T. (2009a). A Virtual Button with Tactile Feedback Using Ultrasonic Vibration. *Virtual and Mixed Reality*, pages 385–393.

Tashiro, K., Shiokawa, Y., Aono, T., and Maeno, T. (2009b). Realization of button click feeling by use of ultrasonic vibration and force feedback. In *Proc. EuroHaptics*, pages 1–6.

van Mensvoort, I. M. K. (2002). What you see is what you feel: exploiting the dominance of the visual over the haptic domain to simulate force-feedback with cursor displacements. In *Proceedings of the 4th conference on Designing interactive systems: processes, practices, methods, and techniques*, DIS '02, pages 345–348, New York, NY, USA. ACM.

Van Orden, K. F. and Broyles, J. W. (2000). Visuospatial task performance as a function of two- and three-dimensional display presentation techniques. *Displays*, 21(1):17–24.

Vitense, H. S., Jacko, J. A., and Emery, V. K. (2003). Multimodal feedback: an assessment of performance and mental workload. *Ergonomics*, 46(1-3):68–87.

Wagner, R. A. and Fischer, M. J. (1974). The String-to-String Correction Problem. *Journal of the ACM*, 21(1):168–173.

Walker, P. and Smith, S. (1985). Stroop interference based on the multimodal correlates of haptic size and auditory pitch. *Perception*, 14(6):729–736.

Wickens, C. D. (2002). Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science*, 3(2):159–177.

Wickens, C. D. (2008). Multiple resources and mental workload. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(3):449–455.

Wickens, C. D., Merwin, D. H., and Lin, E. L. (1994). Implications of Graphics Enhancements for the Visualization of Scientific Data: Dimensional Integrality, Stereopsis, Motion, and Mesh. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 36(1):44–61.

Widmer, A. and Hu, Y. (2010). Effects of the Alignment Between a Haptic Device and Visual Display on the Perception of Object Softness. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 40(6):1146–1155.