Spatial Tactile Feedback Support for Mobile Touch-screen Devices

by

Koji Yatani

A thesis submitted in conformity with the requirements

for the degree of Doctor of Philosophy

Department of Computer Science

University of Toronto

Copyright © 2011 by Koji Yatani

Abstract

Spatial Tactile Feedback Support for Mobile Touch-screen Devices

Koji Yatani

Doctor of Philosophy

Graduate Department of Computer Science

University of Toronto

2011

Mobile touch-screen devices have the capability to accept flexible touch input, and can provide a larger screen than mobile devices with physical buttons. However, many of the user interfaces found in mobile touch-screen devices require visual feedback. This raises a number of user interface challenges. For instance, visually-demanding user interfaces make it difficult for the user to interact with mobile touch-screen devices without looking at the screen—a task the user sometimes wishes to do particularly in a mobile setting. In addition, user interfaces on mobile touch-screen devices are not generally accessible to visually impaired users. Basic tactile feedback (*e.g.*, feedback produced by a single vibration source) can be used to enhance the user experience on mobile touch-screen devices. Unfortunately, this basic tactile feedback often lacks the expressiveness for generating vibration patterns that can be used to convey specific information about the application to the user. However, the availability of richer information accessible through the tactile channel would minimize the visual demand of an application. For example, if the user can perceive which button she is touching on the screen through tactile feedback, she would not need to view the screen, and can instead focus her visual attention towards the primary task (*e.g.*, walking).

In this dissertation, I address high visual demand issues found in existing user interfaces on mobile touch-screen devices by using spatial tactile feedback. Spatial tactile feedback means

tactile feedback patterns generated in different points of the user's body (the user's fingers and palm in this work). I developed tactile feedback hardware employing multiple vibration motors on the backside of a mobile touch-screen device. These multiple vibration motors can produce various spatial vibration patterns on the user's fingers and palm. I then validated the effects of spatial tactile feedback through three different applications: eyes-free interaction, a map application for visually impaired users, and collaboration support. Findings gained through the series of application-oriented investigations indicate that spatial tactile feedback is a beneficial output modality in mobile touch-screen devices, and can mitigate some visual demand issues.

Acknowledgements

I would like to extend my utmost thanks to my advisor, Khai N. Truong, for his guidance and support over the past five years. This dissertation would not have been possible without his deep involvement and encouragement. I also would like to thank my thesis committee: Ravin Balakrishnan, Mark Chignell, Daniel Wigdor, and the external examiner, Stephen Brewster. Their feedback was simply invaluable.

This dissertation has much support from other researchers. Particularly, I would like to mention Darren Gergle's support on designing the experiment and performing analyses presented in Chapter 6. I also would like to thank Nikola Banovic, who helped me design the experiment and perform analyses on results described in Chapter 5. He also offered me great knowledge about interfaces for visually impaired users, which was very helpful.

Besides the dissertation work, I have been very fortune to work with many bright professors, researchers, and students. I learned many from all of them, and they are always my source of inspiration. A special thanks to Hrvoje Benko, Marshall Bern, Bill Buxton, Eunyoung Chung, Carlos Jensen, Nicole Coddington, David Dearman, Richard Guy, Ken Hinckley, Steve Hodges, Elaine M. Huang, Julie A. Kientz, Victor Kuechler, Frank Li, Mark W. Newman, Michael Novati, Michel Pahud, Kurt Partridge, Shwetak N. Patel, Jenny Rodenhouse, Jeremy Scott, Andrew Trusty, Nicolas Villar, and Andy Wilson.

I want to acknowledge faculties, postdocs, visitors and students I had the pleasure to meet in school, including Anand Agarawala, Seok-Hyung Bae, Xiaojun Bi, Nilton Bila, Simon Breslav, Xiang Cao, Fanny Chevalier, Gerry Chu, Mike Daum, Pierre Dragicevic, Carrie Demmans Epp, Dustin Freeman, Clifton Forlines, Tovi Grossman, John Hancock, Sam Hasinoff, Aaron Hertzmann, Justin Ho, Akitoshi Kawamura, Alex Kolliopoulos, Martin de Lasa, Shahzad Malik, Mike Massimi, Karyn Moffatt, Igor Mordatch, Nigel Morris, Tomer Moschovich, Peter O'Donovan, Matthew O'Toole, Gonzalo Ramos, Abhishek Ranjan, Mike Reimer, Alyssa

Rosenzweig, Ryan Schmidt, Eron Steger, Huixian Tang, Daniel Vogel, Jack Wang, Mike Wu, and Shengdong Zhao. I also thank to DAG members: Ablishek, Alex, Eron, Jack, Jean-Nicolas McGee, and Jenny Wang. It is always a very special place to me.

Finally, I am deeply indebted to my parents, Kenichi Yatani and Masae Yatani. I would not be able to complete my degree without their full support and confidence in me over years.

Copyright Notice and Disclaimer

Sections of this document have appeared in a previous publication or have conditionally accepted in a conference at the time of writing. In all cases, permission has been granted by the publisher for the work to appear here. Below, the publisher's copyright notice and disclaimer are given, with thesis chapter and corresponding publication noted.

Association for Computer Machinery

Copyright ©2011 by the Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page in print or the first screen in digital media. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted.

To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Send written requests for republication to ACM Publications, Copyright & Permissions at the address above or fax +1 (212) 869-0481 or email permissions@acm.org.

Portions of Chapter 4

Koji Yatani and Khai N. Truong. 2009. SemFeel: a user interface with semantic tactile feedback for mobile touch-screen devices. In *Proceedings of the 22nd annual ACM symposium on user interface software and technology (UIST 2009)*, ACM, New York, NY, USA, 111-120. DOI=10.1145/1622176.1622198

Portions of Chapter 6

Koji Yatani, Darren Gergle, and Khai N. Truong. 2012. Investigating effects of visual and tactile feedback on spatial coordination in collaborative handheld systems. *To appear in the proceedings of the ACM conference on computer supported cooperative work (CSCW 2012)*, ACM, New York, NY, USA.

Contents

1 Introduction **1**

1.1 Research Objective and Overview 3

1.2 Contributions 6

1.2.1 Spatial Tactile Feedback Hardware 6

1.2.2 Distinguishability of Spatial Tactile Feedback 6

1.2.3 Eyes-free Mobile Interaction with Spatial Tactile Feedback 6

1.2.4 Spatial Relationship Learning Support for Visually Impaired Users 7

1.2.5 Spatial Tactile Feedback Support in Remote Collaboration 7

1.3 Dissertation Outline 8

2 Background Literature **11**

2.1 The Hand 13

2.1.1 The Skin Senses 14
2.2 The Hardware 16
2.2.1 DC Motors 17
2.2.2 Voice Coil Actuator 19
2.2.3 Piezoelectric Bending Actuators 20
2.2.4 MR Fluid 21
2.2.5 Mechanical Pin Arrays 22
2.2.6 Ultrasound Transducers 23
2.2.7 Electrovibration 23
2.2.8 Force Feedback Hardware 25
2.2.9 Summary 26
2.3 Vibrotactile Feedback Parameters 26
2.3.1 Intensity 27
2.3.2 Frequency 28
2.3.3 Duration 29
2.3.4 Waveform 30
2.3.5 Rhythm 30
2.3.6 Locus (Spatial Patterns) 31
2.3.7 Spatio-temporal Patterns 32
2.3.8 Summary 34

2.4 User Interfaces with Tactile Feedback 34
2.4.1 Touch Screens and Touch-sensitive Surfaces 35
2.4.2 Handheld and Mobile Devices 40
2.4.3 Perception of Vibrotactile Feedback 43
2.4.4 Effects of Vibrotactile Feedback 44
2.5 Summary 46

3 Spatial Tactile Feedback Hardware 48
3.1 Design Requirements 48
3.2 First Prototype 49
3.3 Second Prototype 52

4 Eyes-free Interaction 53
4.1 SemFeel Concept 54
4.2 Related Work 55
4.3 SemFeel System 58
4.3.1 Hardware Configuration 58
4.3.2 Vibration Patterns 59
4.4 Experiment 1: Distinguishability of Patterns 60
4.4.1 Tasks and Stimuli 60
4.4.2 Variables 61

4.4.3 Apparatus 62
4.4.4 Procedure 62
4.4.5 Participants 62
4.5 Experiment 1 Results and Discussions 63
4.6 Experiment 2: User Performance on Eyes-free Number Entry Task 67
4.6.1 Tasks and Stimuli 67
4.6.2 Variables 71
4.6.3 Apparatus 71
4.6.4 Procedure 71
4.6.5 Participants 73
4.7 Experiment 2 Results and Discussions 73
4.8 Applications 77
4.9 Summary 79

5 Interface Design for Visually Impaired Users 81

5.1 SpaceSense Concept 82
5.2 Related Work 83
5.2.1 Spatial Learning with Physical Maps 84
5.2.2 Interfaces for Map Exploration 85

5.2.3 Interfaces for Navigation and Wayfinding 86
5.3 The SpaceSense Device 88
5.4 SpaceSense Interactions 88
5.4.1 Identifying Locations 89
5.4.2 Learning Place Details 91
5.4.3 Learning Directions 91
5.4.4 Bookmarking a Place 93
5.5 User Study 94
5.5.1 Experimental Design and Procedure 94
5.5.2 Participants 98
5.6 Results 98
5.6.1 Task Completion Time 99
5.6.2 Route Composition Accuracy 99
5.6.3 Drawn Places Accuracy 100
5.6.4 User Feedback 102
5.7 Discussion 105
5.7.1 Information Overload through the Audio Channel . . . 105
5.7.2 Errors and Limitations on Learned Routes and Spatial Relationships 106
5.7.3 Study Limitations 106

5.8 Summary 107

6 Collaboration Support 109

6.1 Related Work 110

6.1.1 Inter-personal Communication Systems 111

6.1.2 Collaborative Systems for Visually Impaired Users . . 112

6.1.3 Spatial Coordination 113

6.2 System 114

6.2.1 Game Design 114

6.2.2 System Architecture 117

6.3 Study1: Effects of Each Feedback Type 118

6.3.1 Conditions 118

6.3.2 Procedure 120

6.3.3 Participants 120

6.3.4 Utterance Analysis 121

6.4 Study1 Results 122

6.4.1 Score 122

6.4.2 Utterances 123

6.4.3 Strategies 125

6.4.4 Collaboration in the Audio-only Condition 125

6.4.5 Collaboration in the Visual Conditions 127
6.4.6 Collaboration in the Tactile Conditions 129
6.4.7 Study1 Summary 131
6.5 Study2: Effects of Combined Feedback 131
6.6 Study2 Results 132
6.6.1 Score 132
6.6.2 Utterances 133
6.6.3 User Feedback 134
6.6.4 Study2 Summary 135
6.7 Discussions 135
6.8 Summary 137

7 Conclusions **138**

7.1 Summary 138
7.2 Limitations 141
7.3 Future Work 145
7.3.1 More Complex Vibration Patterns 145
7.3.2 Navigation Systems for Visually Impaired Users . . . 147
7.3.3 Collaborative Systems for Visually Impaired Users . . 147
7.4 Final Word 148

Bibliography **148**

Appendix A Spatial Tactile Feedback Hardware Circuit Diagram **167**

Appendix B Study on Audio Review Spotlight in SpaceSense **170**

List of Tables

1.1 Research questions explored in this dissertation 5

2.1 Characteristics of the four mechanoreceptors 14

4.1 The confusion matrix for the recognition accuracies of the eleven vibration patterns tested in Experiment 1 63

5.1 An example of place details provided by SpaceSense 91

5.2 The combinations of the destinations and bookmarked places used in the study 95

5.3 The time participants spent in learning the route instructions and reproducing a route with acrylic pieces 99

5.4 The combinations of the destinations and bookmarked places used in the study 101

5.5 The ratings for all the combinations of two places in each map across the two conditions 103

6.1 The five conditions tested in Study 1 118

6.2 The coding scheme observed during the experiment 121

6.3 The number of utterances for the three audio-available conditions tested in Study 1 123

6.4 The number of utterances for the three audio-available conditions tested in Study 2 133

B.1 Examples of review presentations in SpaceSense used in the user study 171

List of Figures

1.1 The application space in spatial tactile feedback 4

2.1 Definitions of haptic feedback terminology 12

2.2 Displacement values of detectable vibrations across the frequency 15

2.3 A structure of a brushed motor 17

2.4 A brushless motor for a computer cooling fan 18

2.5 A voice coil motor 19

2.6 A structure of piezoelectric bending actuators 20

2.7 MR fluid activation 22

2.8 The mechanism of electrovibration 24

2.9 Examples of patterns used in the study by Cholewiak *et al.* (1995) 32

2.10 The T-Bar widget 37

2.11 The structure of a deformable surface 38

2.12 The structure of MudPad 39

2.13 The THMB prototype 41

3.1 The first hardware prototype 50

3.2 The special sleeve for my spatial tactile feedback hardware . . 51

3.3 The second hardware prototype 52

4.1 The SemFeel system concept 55

4.2 Eleven vibration patterns studied 58

4.3 The smoothing of the vibration strength for the right-left vibration 59

4.4 The screenshot of the software used in Experiment 1 61

4.5 The reaction time across the vibration patterns in Experiment 1 65

4.6 The reaction time across the four blocks in Experiment 1 . . . 66

4.7 The system used in Experiment 2 68

4.8 The mapping of the vibration patterns for the numeric keyboard used in the *Single Tactile* condition 69

4.9 The mapping of the vibration patterns for the numeric keyboard used in the *Multiple Tactile* condition 70

4.10 Experiment 2 setup 72

4.11 The typing error rate across the three feedback conditions in Experiment 2 74

4.12 The average performance time of typing a 4-digit number across the three feedback conditions in Experiment 2 75

4.13 The average performance time of typing a 4-digit number across the three feedback conditions in Experiment 2 76

4.14 SemFeel applications 78

5.1 SpaceSense concept 83

5.2 Interactions in selecting a category 90

5.3 Interactions in learning directions 92

5.4 Maps with four hypothetical locations used in the user study . 95

5.5 Acrylic pieces used in this user study 96

5.6 Route composition by a participant during the experiment . . . 96

5.7 A location point drawing example for Neighborhood A 97

5.8 Route composition examples 100

6.1 Screenshots of the game 114

6.2 Visual feedback used in the game 116

6.3 Spatial tactile feedback used in the game 116

6.4 Experimental setup 119

6.5 The average scores for the conditions tested in Experiment 1 . 122

6.6 An example of collaboration observed in the *Audio-only* condition 126

6.7 An example of collaboration observed in the *Visual-Audio* condition 128

6.8 An example of collaboration observed in the *Tactile-NoAudio* condition 130

6.9 The average scores for the conditions tested in Study 2 132

A.1 The circuit diagram for the first prototype 168

A.2 The circuit diagram for the second prototype 169

Chapter 1

Introduction

Mobile touch-screen devices have become increasingly more common in recent years. Touch-sensitive surfaces allow the user to input directly on the screen using a finger or a pen. Recent touch screens can also accept multi-touch input and whole-hand input, which extend the expressiveness of hand gestures. Thus, touch-screen devices typically do not need to include a large number of physical keys to support user input. Having fewer physical keys on touch screens also offers output advantages. For example, touch-screen devices can devote more of their surface area towards providing the user with a larger display screen. This benefit is significant for mobile devices because their screen size is limited. However, this also means that many of the interactive components in the user interface are designed as part of the graphical—and inherently visual—user interface.

Graphical user interfaces heavily rely on visual feedback, and this can cause many interaction problems. For instance, the user sometimes interacts with mobile touch-screen devices while she is in motion, the effects of which have already been studied (Kane *et al.*, 2008; Oulasvirta *et al.*, 2005; Yatani & Truong, 2007, 2009). This prior research indicates that visually-demanding user interfaces commonly found in today's graphical user interfaces require her to devote visual attention to the interface, and can impact the performance of the

user's primary task (*e.g.*, walking) or secondary task (*e.g.*, interacting with the device).

The high visual demand of user interfaces on mobile touch-screen devices also raises an accessibility issue for people with visual impairments. Visually impaired users would not be able to use touch screen devices without support of non-visual feedback. Auditory feedback can help them interact with these devices by providing textual information about what they are contacting in graphical user interfaces. But in mobile devices, this type of feedback is not always ideal or appropriate (*e.g.*, the user may want to interact with the device in silence when other people are around).

Besides auditory feedback, vibrotactile feedback has been used as an alternative way to provide tactile feedback on touch-screen devices (Fukumoto & Sugimura, 2001; Hoggan *et al.*, 2008a; Poupyrev *et al.*, 2002a). When the user touches an object on the screen, such as a button, a webpage link, or an item in a linear list, a vibration motor embedded in the device is activated to provide tactile feedback. This helps the user perceive whether she is touching an object or not. Previous studies have shown that this enhancement can improve user performance on different tasks using mobile touch-screen devices (Hoggan *et al.*, 2008a; Leung *et al.*, 2007; Poupyrev *et al.*, 2002a). However, unlike auditory feedback, this basic form of tactile feedback does not convey information which could help the user identify the object she is touching or gain additional semantic information about it. Although different vibration patterns (*e.g.*, different rhythms or different strength levels) can be used to convey some semantic information (Brown *et al.*, 2006a), additional ways to convey richer data that can be perceived and understood easily over the tactile channel remain to be explored.

This dissertation aims to improve the expressiveness of vibrotactile feedback on mobile touch-screen devices by using multiple vibration motors attached to different locations of a mobile device. It also validates the effectiveness of this spatially-enhanced vibrotactile feedback through three different applications.

1.1 Research Objective and Overview

I define *spatial tactile feedback* as tactile feedback using spatially-distributed vibration sources and encoding information in spatial patterns (*e.g.*, vibration at the top of a hand has a different meaning from one at the bottom). Spatial tactile feedback also includes patterns changing both spatially and temporally (*e.g.*, vibration flowing in one direction). This spatial aspect of tactile feedback has not been well explored in user interface designs of mobile devices, and the motivation of this work is to understand how spatial tactile feedback could support interactions in mobile touch-screen devices. Hence, the main contribution of this dissertation can be stated as follows:

> *Spatial tactile feedback produced by multiple vibration motors enhances the expressiveness of tactile feedback on mobile touch-screen devices, and enables the development of user interfaces which require low or no visual demand.*

I achieve this by taking an application-oriented approach. Figure 1.1 illustrates the application space in which spatial tactile feedback could be useful. I define two dimensions: *Level of visual demand* (how much visual demand needs to be reduced); and *Number of users* (how many users are involved in applications). These dimensions are not definitive and one can explore other dimensions, such as different contexts (indoor *vs.* outdoor). But I chose them because they cover application domains which highlight the effectiveness of spatial tactile feedback in mobile devices. They also represent a set of non-overlapping application features. Thus, exploration within each of the domains would investigate a unique type of use and benefits of spatial tactile feedback in mobile touch-screen devices.

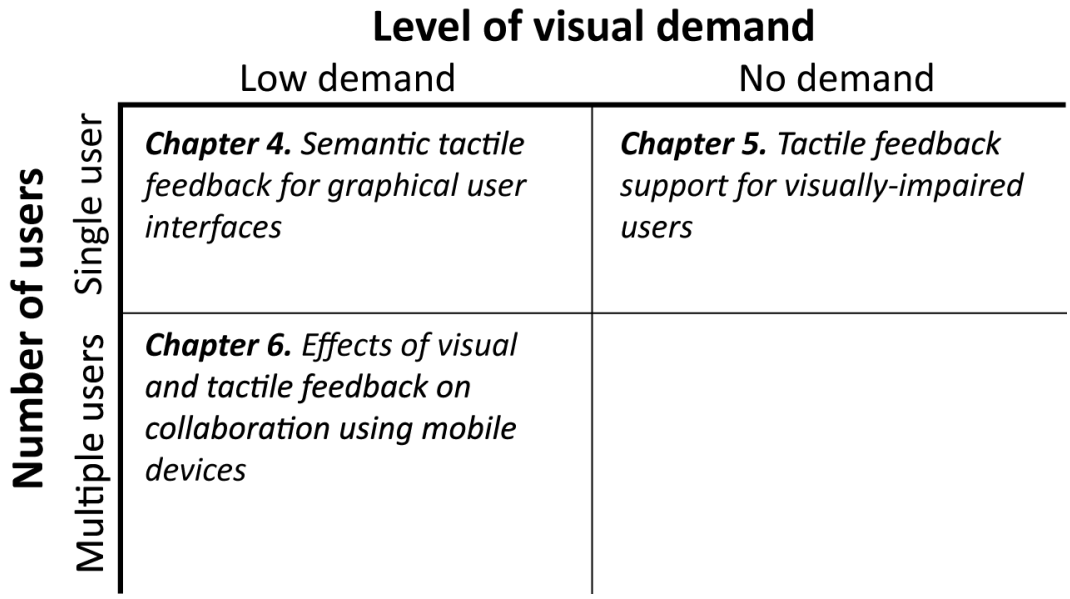

Figure 1.1. The application space in spatial tactile feedback.

For each dimension, I set up two levels. In *Level of visual demand*, *Low demand* means where user performance could be improved at least when visual demand is reduced. *No demand* means where visual demand needs to be completely removed to improve user experience. In *Number of users*, the two levels are whether multiple users are involved in interactions.

As shown in Figure 1.1, I focus on three application domains in this dissertation. With these application domain in mind, I developed research questions as summarized in Table 1.1. I answer these questions through developing interfaces with spatial tactile feedback and evaluating them.

	Motivation / Research question
RQ-1.	**Motivation**: The distinguishability of spatial vibration patterns are not well understood. **Question**: *How fast and accurately can users recognize various spatial vibration patterns generated by multiple vibration motors?*
RQ-2.	**Motivation**: Interfaces on mobile touch-screen devices are not often designed appropriately for eyes-free interaction. **Question**: *How fast and accurately can users perform eyes-free interaction with the support of spatial tactile feedback?*
RQ-3.	**Motivation**: Graphical user interfaces currently do not provide information about the spatial relationships between multiple objects in an accessible manner to visually-impaired users. **Question**: *What level of accuracy and detail can visually-impaired users achieve when they try to construct mental maps of spatial relationships between multiple objects with the support of spatial tactile feedback?*
RQ-4.	**Motivation**: Visual feedback may impact users' collaboration on mobile devices differently from desktop computer or tabletop displays because visual workspace is limited and the user's hand can occlude much of the screen. **Question**: *How does spatial tactile feedback affect the performance and communication of patterns in remote collaboration on mobile touch-screen devices?*

Table 1.1. Research questions explored in this dissertation.

1.2 Contributions

This dissertation offers five major contributions to the field of Human-Computer Interaction (HCI), particularly interface designs on mobile touch-screen devices.

Spatial Tactile Feedback Hardware

The hardware of spatial tactile feedback developed for mobile touch-screen devices is a unique contribution in this dissertation. Although many tactile feedback systems have been developed in mobile devices, systems which can generate spatially different vibration patterns have not been well explored. My hardware design is developed based on physiological and psychological requirements, and physical form factors of mobile devices. The hardware architecture can be easily transferable to other kinds of mobile touch-screen devices.

Distinguishability of Spatial Tactile Feedback

I examined how people can recognize different vibration patterns generated by my hardware. These pattern include vibration flowing from one point to another (spatio-temporal patterns) as well as vibration at one particular point (spatial patterns). In particular, the distinguishability of spatio-temporal patterns in the user's hand has not been studied well. This dissertation provides a basic understanding of how distinguishable spatial and spatio-temporal patterns are.

Eyes-free Mobile Interaction with Spatial Tactile Feedback

The first exploration in applications with spatial tactile feedback is eyes-free interaction. As described above, eyes-free interaction can be a beneficial interface design to users, but also challenging for mobile touch-screen devices. This work shows that spatial and spatio-temporal patterns can inform the user about what object she is touching without necessarily viewing

the screen. This dissertation validates it through a laboratory user study of eyes-free number entering tasks.

Spatial Relationship Learning Support for Visually Impaired Users

Accessibility to visually impaired users is another user interface challenge widely recognized in mobile touch-screen devices. Spatial tactile feedback can be used to provide visually impaired users with spatial or directional information. My design of map applications offers spatial relationships between multiple places through spatial tactile feedback, and helps visually impaired construct and maintain their mental map. My user study shows that participants could build and maintain spatial relationships between four places on a map more accurately with my system compared to a system without spatial tactile feedback.

Spatial Tactile Feedback Support in Remote Collaboration

Mobile and handheld devices have become a platform that supports remote collaboration. But, their small form factor may impact the effectiveness of the visual feedback channel often used to help users maintain an awareness of their partner's activities during synchronous collaborative tasks. I investigated how spatial tactile feedback affects collaboration on mobile devices, with emphasis on spatial coordination in a shared workspace. From two user studies, my results highlight that spatial tactile feedback has different benefits from visual feedback. They also show that visual and tactile feedback can complement each other, and systems using both feedback channels can support better spatial coordination than systems using only one form of feedback.

1.3 Dissertation Outline

I begin with my literature survey on tactile feedback. Chapter 2 describes the summary of my survey in physiology, psychology, hardware architecture, materials, and HCI research. I first discuss the human capability of sensing vibrotactile feedback in the hand. I then present various technologies used in HCI research to produce tactile feedback. Vibration patterns generated by vibrotactile feedback systems can be described by seven parameters. I explain these parameters along the distinguishability of vibration patterns varied in each parameter. Finally, I cover user interfaces with vibrotactile feedback, particularly emphasizing on interfaces in touch screens and mobile devices.

The first contribution of this dissertation is the development of spatial tactile feedback system for mobile touch-screen devices. I developed a tactile feedback system employing multiple vibration motors on the backside of a mobile touch-screen device. These multiple vibration motors are embedded in different locations on the device, and the system can produce various spatial vibration patterns, such as vibration from different sides of the device and flows of vibration (*e.g.*, vibration flowing from right to left). In Chapter 3, I present the implementation of two my tactile feedback hardware prototype.

I investigated how this spatial tactile feedback can enhance user interfaces on mobile touch-screen devices. I conducted two kinds of quantitative experiments to understand the capability of my spatial tactile feedback system (RQ-1) and the user performance of eyes-free interaction with it (RQ-2). The results revealed that with the system, called SemFeel, users can distinguish ten different patterns—including linear patterns and a circular pattern—at approximately 90% accuracy. I also found that SemFeel supports more accurate eyes-free interactions than interfaces without any tactile feedback or with tactile feedback using a single vibration motor. In Chapter 4, I present the details of these experiments and discuss possible applications for SemFeel.

I extend the idea behind SemFeel to design user interfaces for visually impaired users. Prior research shows that visually impaired users can construct mental maps of spatial relationships between multiple objects or locations through tactile sensation (Thinus-Blanc & Gaunet, 1997; Ungar, 2000). For example, a tactile map has been shown to be successful in offering visually-impaired users information about different locations on a map (Herman *et al.*, 1983; Ungar *et al.*, 1993, 1995). Although many mobile devices have map functionality to support user navigation, most of the components in these applications are visual, and they are not necessarily well-designed for visually impaired users.

In Chapter 5, I present a system to support visually impaired users in understanding spatial relationships between multiple objects by using the second prototype of my spatial tactile feedback hardware (RQ-3). The system, called SpaceSense, offers step-by-step instructions to get to a particular place through the speech feedback. SpaceSense also provides vibrotactile feedback about an approximate direction and distance between the user's current simulated location and places of interest in a map, such as the destination and places that the user already knows. Through this tactile cue, the user can have information about the spatial relationships between these places as she navigates the route instructions. Through a user study, I found that participants could build and maintain spatial relationships between four places on a map more accurately with SpaceSense compared to a system without spatial tactile feedback.

In addition to providing accessibility features to mobile touch-screen devices, spatial tactile feedback can also support remote collaboration. Mobile devices have become a platform for remote collaboration. People now use their mobile devices to email, instant message, and video conference. Tactile feedback may be particularly important for mobile touch-screen devices because of their small form factor, which may impact the visual feedback often used in synchronous collaborative tasks to help users maintain awareness about their collaborator's activities.

In Chapter 6, I present my investigation on how visual and tactile feedback affects remote collaboration on mobile devices to answer RQ-4. In particular, I examined how users would perform spatial coordination in a shared workspace when they need to perform an action collaboratively. I present my two experiments and analyses on the data gathered. The results highlight different benefits of each feedback channel in collaborative handheld systems. Visual feedback can provide precise spatial information for collaborators. But it degrades collaboration when the feedback is occluded, and sometimes can distract the user's attention. Spatial tactile feedback can reduce the overload of information in visual space and gently guides the user's attention to an area of interest. I also discovered that visual and tactile feedback can complement each other, and systems using both feedback channels can support better spatial coordination than systems using only one form of feedback.

In Chapter 7, I draw conclusions based on the findings I gained through this work. I also discuss limitations of this work and future research directions with spatial tactile feedback.

Chapter 2

Background Literature

Scientific research related to tactile sensation is recognized to have begun with the exploration by Ernst Heinrich Weber in the 1840s. His *De Tactu* (Concerning Touch, Weber (1995)) is probably the first major work exploring various aspects of tactile sensation. Particularly, his finding about the relationships between the magnitude of a stimulus and its perceived intensity is known as Weber's law (or Weber-Fechner law because Gustav Theodor Fechner developed an theoretical model based on Weber's findings). In his study, Weber gradually increased the weight a blindfolded person was holding and asked her to respond when she felt the increase. Weber found that the smallest noticeable difference in weight (the least difference that participants could perceive as a difference), was proportional to the initial weight value.

In 1925, David Katz reported his rigorous exploration on the role of temperature in touch perception, and vibration sensitivity in his book *The World of Touch* (Katz, 2009). Katz's exploration motivated many psychologists and physiologists to examine the human capabilities of different aspects of tactile sensation. The research community now shares a good understanding of how tactile stimuli can be sensed and perceived as I explain in the following sections although there are still many aspects that are not fully understood.

The early 20th century was also perhaps the time when the term "haptic" began to be used commonly in psychology. The two terms "tactile" and "haptic" may show very similar meaning in a dictionary:

tactile *adj.* (1645) 1. *perceptible by touch.*
2. *of, relating to, or being the sense of touch.*

haptic *adj.* (*ca.* 1890) 1. *relating to or based on the sense of touch.*
2. *characterized by a predilection for the sense of touch.*

(from Merriam-Webster's Collegiate Dictionary, 11th Edition.)

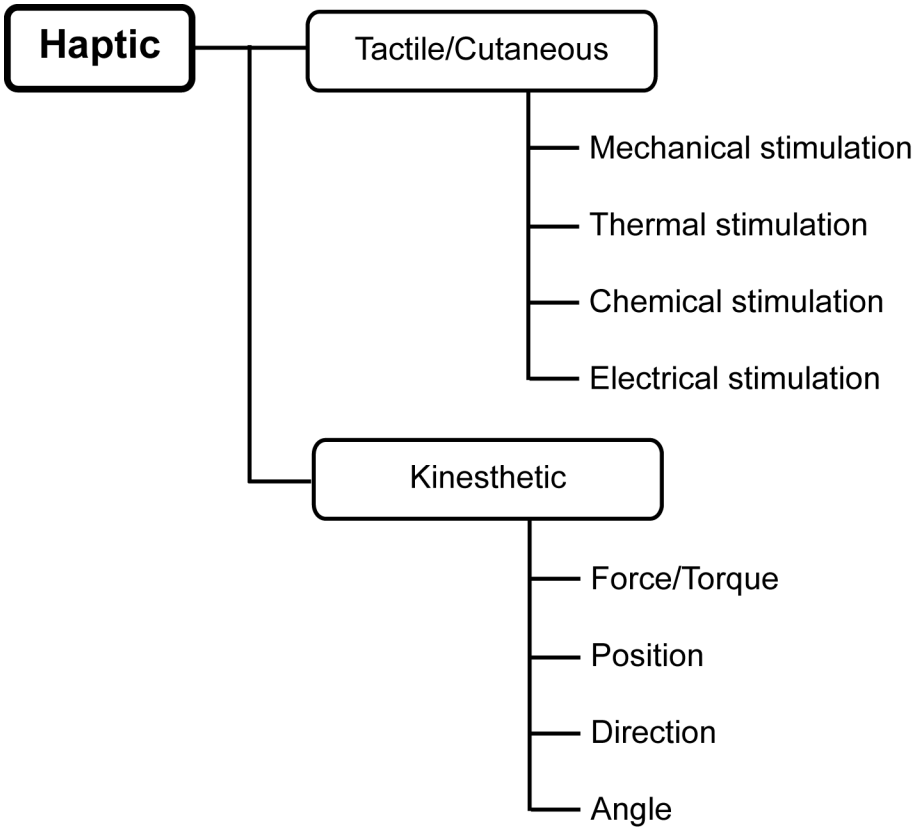

Figure 2.1. Definitions of haptic feedback terminology (adapted from ISO: Ergomics of human-computer interaction—Part910: Framework for tactile and haptic interaction (2009) and Hoggan (2010)).

However, there are multiple dimensions in haptic sensation, such as texture, stiffness, and temperature. Klatzky *et al.* (1985) showed that these dimensions are important for us to distinguish different objects through haptic sensation. Thus, "tactile" and "haptic" are distinct in the field of psychology and HCI to represent these dimensions. According to ISO: Ergomics of human-computer interaction—Part910: Framework for tactile and haptic interaction (2009), "haptic" covers the area of "tactile" (the sense of touch) and "kinesthetic" (the sense of body position or movements) as shown in Figure 2.1. This work focuses on tactile sensation, more specifically sensation caused by vibration stimuli. Thus, I use "tactile feedback" or "vibrotactile feedback" for my systems in this dissertation.

In this chapter, I first overview the physiological and psychological characteristics of the hand. I then review various technologies which can produce tactile feedback in computer systems. Next, I explain the design space of vibrotactile feedback. Finally, I overview tactile feedback systems in touch screens and mobile devices. Based on my literature survey, I conclude that although tactile feedback can be useful to support interactions in mobile touch-screen devices, its expressiveness is still limited. Particularly, spatial vibration patterns is an overlooked aspect of vibrotactile feedback.

2.1 The Hand

As Jones & Lederman (2006) illustrates in their conceptual framework of the hand ranging from tactile sensation to motor function, the hand has different functionalities and supports very accurate manipulation and perception. The hand is also one of the most sensitive parts in the human body to tactile sensation. In this section, I discuss physiological aspects of tactile sensation.

	Merkel discs	**Ruffini endings**	**Meissner corpuscles**	**Pacinian corpuscles**
Property	SA Type I	SA Type II	RA Type I	RA Type II
Sensation	Pressure	Stretch	Tap, flutter	Vibration
Frequency	0.4-1, 5-15	150-400	25-40	200-300
Field area	Small (12.6)	Large (101.0)	Small (11.0)	Large (59.0)
Adaptation	Slow	Slow	Rapid	Very rapid

SA: Slowly-adapting, RA: Rapidly-adapting

Table 2.1. Characteristics of the four mechanoreceptors.

The Skin Senses

The skin is the largest organ found in the human body. The skin area of the average adult human body is approximately 1.8 m^2. As a sensory organ, the skin has special nerve endings which form a large portion of the somatosensory system in the body. The somatosensory system produces various sensation as a reaction to stimulation, such as touch, temperature, proprioception (*i.e.*, the sense of the relative position of adjacent parts of the body), and nociception (*i.e.*, pain), by triggering sensory neurons.

Tactile stimulation is sensed by four mechanoreceptors and free nerve endings mostly found in the *epidermis*, the outer part of the skin. Although free nerve endings are sensitive to touch, pressure, stretch, and temperature, their primary role is a receptor for pain. This section, thus, discusses the four mechanoreceptors: Merkel discs, Ruffini endings, Melssner's corpuscles, and Panician corpuscles. Note that there is no definitive conclusion that a particular mechanoreceptor is activated by specific tactile stimulation; multiple mechanoreceptors usually yield tactile sensation at the same time. However, each mechanoreceptor has a different sensitivity bandwidth in frequencies of tactile stimuli, which relates to different kinds of

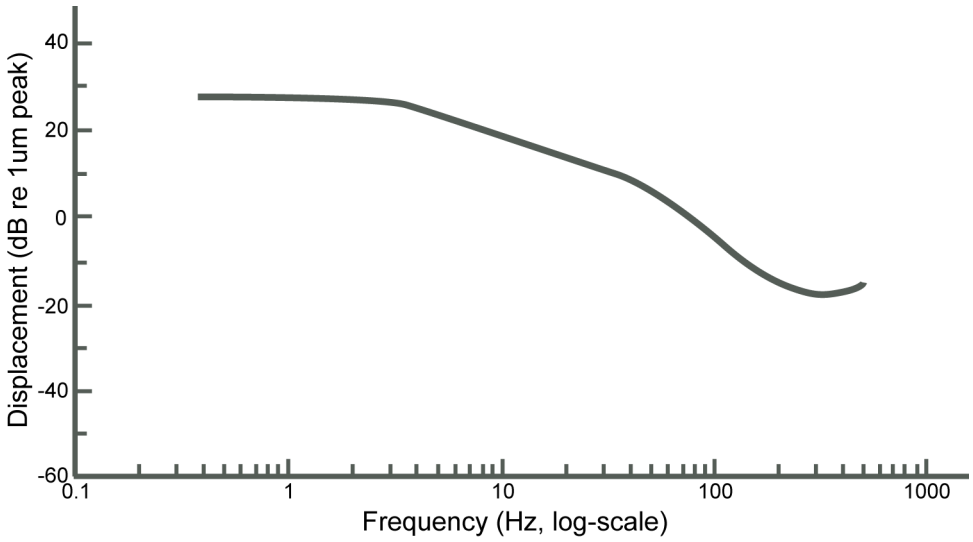

Figure 2.2. Displacement values of detectable vibrations across the frequency. A smaller displacement means smaller vibration (reproduced based on Sherrick (1985)).

perception. For example, a stimulus with lower than 1Hz, lower than 10Hz, or higher than 50Hz frequency produces a feeling of taps, fluttering, or vibration, respectively.

The sensation this dissertation focuses on is vibration; therefore, understanding the human capability of sensing vibration is important. Table 2.1 describes characteristics of the four mechanoreceptors. They can be categorized by their rate of adaption and sensory modalities. When a mechanoreceptor receives a stimulus, it fires impulses in the nerve system. But the receptor will adapt to a static stimulus, and impulses in the nerve system will become weaker. This is called adaption. Two of the four mechanoreceptors are known to demonstrate a rapid adaptation rate (impulses in the nerve system become weak quickly after the initial stimulus is provided). These rapidly-adapting mechanoreceptors contribute to the sense of rapidly changing simulation, such as texture or vibration.

Pacinian corpuscles, or Lamellar corpuscles, is a major rapidly-adapting mechanoreceptor responsible for vibration and pressure. A Pacinian corpuscle is a large oval-shaped

receptor—typically about 0.5 mm wide and 1.0 mm long. Each Pacinian corpuscles consists of 50 or more concentric layers of tissue and surrounding its sensory nerve. This layer structure of a Pacinian corpuscle allows this mechanoreceptor to be very sensitive to vibration. The optimal sensitivity is 250 Hz with the spatial vibration displacement being smaller than 200 μm. Figure 2.2 shows the minimum displacement of vibration across the frequency.

200–400 Pacinian corpuscles can be found in the human hand (Stark *et al.*, 1998). But Pacinian corpuscles do not have a very fine spatial resolution because the field area that one Pacinian corpuscle covers is relatively large. Palmer & Gardner (1990) conducted a study with adult monkeys to reveal that the nervous system fires two distinct nerve signal bursts when two stimuli are separated with the gap larger than 1 cm. This indicates that two adjacent vibration sources need to be separated by at least longer than 1 cm for the user to accurately distinguish their spatial difference.

2.2 The Hardware

Various technologies have been used to produce tactile sensation in user interfaces for computer systems. Tangible user interfaces involve physical objects to interact with systems and digital information, and inherently offer tactile sensation to the user (*e.g.*, a feeling of grasping a block). This can be called *passive* tactile feedback, where there is no need for electrical or mechanical components for producing tactile sensation as Harrison & Hudson (2009b) discussed.

This dissertation focuses on *active* tactile feedback. It involves electrical or mechanical activation which generates pressure, force feedback, or vibration. This section discusses some representative technologies to produce active tactile feedback in user interfaces.

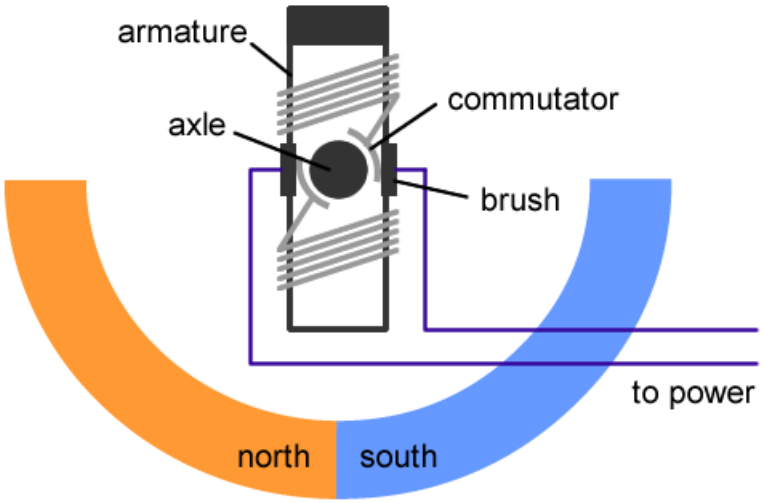

Figure 2.3. A structure of a brushed motor.

DC Motors

A direct current (DC) vibration motor is common hardware used for active tactile feedback. When it receives direct current, the motor inside is activated, and generates a rotation force. An imbalanced counterweight connected to the motor then starts to be rotated, and as a result, vibration is generated.

There are mainly two types of motors used in this type of hardware: Brushed motors and brushless motors. In both kinds of motors, a force for rotating the *armature* (the rotor) is generated by magnetism. As shown in Figure 2.3, the armature is an electromagnet made by coiling thin wires, and surrounded by permanent magnets, or *stators*. the north and south poles are placed on each side of the armature. The armature also has *commutators* and brushes, which play an important role to rotate the armature.

When the armature receives a current, it generates an electromagnetic field around. By the surrounding permanent magnets, one side of the armature is pushed away, and the other side is pulled. The combination of these forces causes a rotation of the armature. However, this

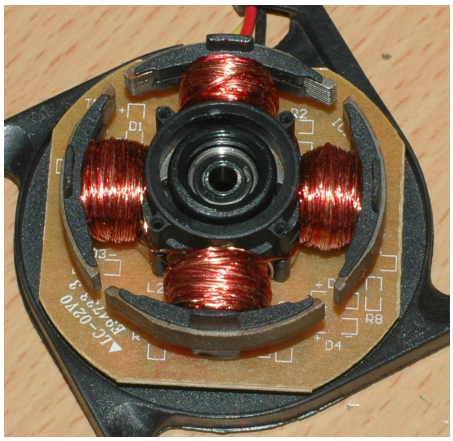

Figure 2.4. A brushless motor for a computer cooling fan. Four coil stators surround the rotor. The rotor has been removed from this figure. (This is a public domain figure.)

only causes a half a cycle of a rotation. When the armature becomes horizontally aligned, the rotation would stop. The commutators and brushes then reverse the direction of the current. In this manner, the armature continues to receive attracting and repelling forces, which generate a continuous rotation of the armature.

One major shortcoming of brushed motors is that brushes need to be physically contacted to the commutators to provide power to the armature. Thus, brushes are subject to wear by the friction against the commutators. This friction also can cause electrical noise which may impact other parts of the circuit.

Brushless motors can overcome some of these issues. In brushless motors, the armature consists of an permanent magnet, and it is surrounded by coil stators (Figure 2.4). Similar to brushed motors, by changing the direction of a current flowing to the coils, the armature receives attracting and repelling forces, and is rotated. Although brushless motors have advantages in terms of maintenance, they requires a precise control of the current. Thus, coils usually are controlled by computers, which increases the initial production cost.

DC vibration motors with different sizes have been developed, and small motors have been integrated into mobile devices. They are used to generate vibration to inform users an incoming call or email. This work also uses DC vibration motors to produce spatial tactile feedback. I explain the rationales behind choosing this technology in Chapter 3.

Voice Coil Actuator

Figure 2.5. A voice coil motor.

A similar actuator to DC vibration motors which is often used in HCI research is a voice coil actuator. Figure 2.5 is Audiological Engineering's VBW32, often known as a Tactaid actuator. The structure of a voice coil actuator is simple. In its center, a donut-shape permanent magnet is located. A coil (a voice coil) is then inserted in the hole. When an electrical signal passes through the coil's wire, the voice coil moves back and forth. The sides of this voice coil are attached to flat thin panels, and thus the actuator will produce vibration corresponding to the original input electrical signal.

This architecture can be found in some audio speakers (in these cases, a diaphragm is attached to the voice coil to produce the sound). Thus, in many of these voice coil actuators, a

sound signal can be used to produce different vibration patterns. Researchers easily can change the amplitude, frequency, and rhythm of vibration with voice coil actuators, and these aspects of vibration have been investigated in the context of user interface designs (Brewster & Brown, 2004). However, the physical form factor of these actuators is often relatively large for mobile devices. For example, the size of VBW32 is 26 × 18 ×12 mm. Thus, the integration of multiple actuators to mobile devices may be challenging.

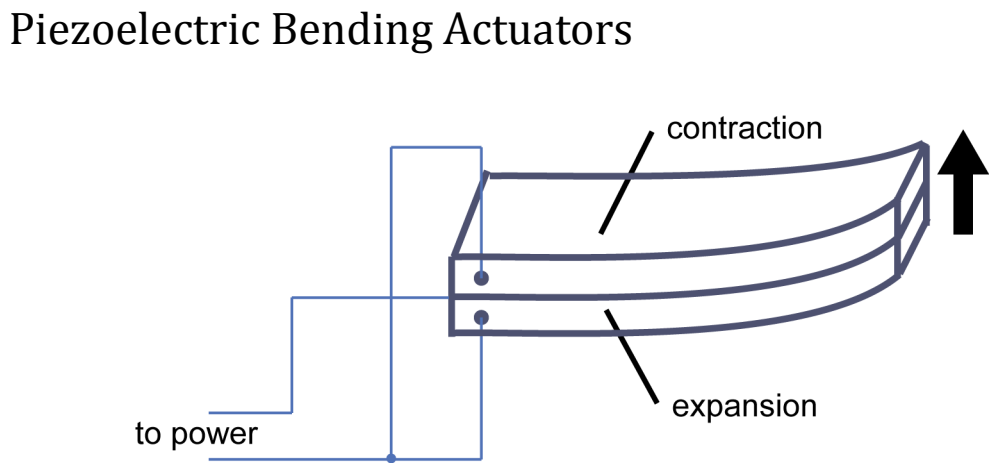

Figure 2.6. A structure of piezoelectric bending actuators. When voltage is applied, one side of the layers is contracted and the other is expanded. As a result, the entire structure is bended in one direction and generates a mechanical force.

A piezoelectric bending actuator is another common hardware used for generating vibrotactile feedback. There are mainly three kinds of structures in this type of actuators by the number of layers inside: Monomorph, bimorph, and multi-layer. In any of these types, the basic structure is two electrodes (*e.g.*, mental sheets) with piezoelectric materials in-between (*e.g.*, pizeoceramics).

Piezoelectric bending actuators are based on the phenomenon known as the piezoelectric effect. This effect is a product of interaction between the mechanical and electrical state in the material. More specifically, the piezoelectric effect means that an electrical charge appears

between the electrodes when a mechanical force is applied to the material. Thus, hardware can be used as a bending sensor by measuring the generated voltage.

The piezoelectric effect is a reversible process. When voltage is applied to the electrodes, the material generates a mechanical force (Figure 2.6). This is called the reverse piezoelectric effect. Therefore, hardware can generate vibrotactile feedback by applying alternating current to the electrodes.

The architecture of piezoelectric bending actuators used to be limited to a single-layer (monomorph) or double-layer (bimorph). However, some monomorph and bimorph actuators require high voltage (like 350V) to activate, and are not useful for mobile devices. This issue can be solved by having a multi-layer structure. For example, TouchEngine developed by Poupyrev *et al.* (2002b) has 37 layers (19 layers of electrodes and 18 layers of piezoceramics), and can be integrated into mobile devices. Thus, a multi-layer piezoelectric actuator is now often used in mobile devices (Luk *et al.*, 2006; Poupyrev *et al.*, 2002a).

MR Fluid

A magnetorheological fluid (MR fluid), initially discovered by Rabinow (1948), is another material which can be used to produce vibrotactile feedback with magnetism. It consists of ferromagnetic micrometer-scale particles (*e.g.*, magnetite or mangan zinc ferrite), surface active agents covering these particles, and carrier oil. When a magnetic field is applied to an MR fluid, its apparent viscosity is increased, and it starts to have a property of a viscoelasticity solid.

Figure 2.7 shows the mechanism of viscosity of an MR fluid variable by the strength of the applied magnetic field. When no magnetic field is applied, particles are distributed randomly within the carrier oil and are suspended (Figure 2.7a).

When a magnetic field is applied, the particles are aligned along the direction of magnetic flux. They form chains of particles, and constrain the movement of the fluid, perpendicular to

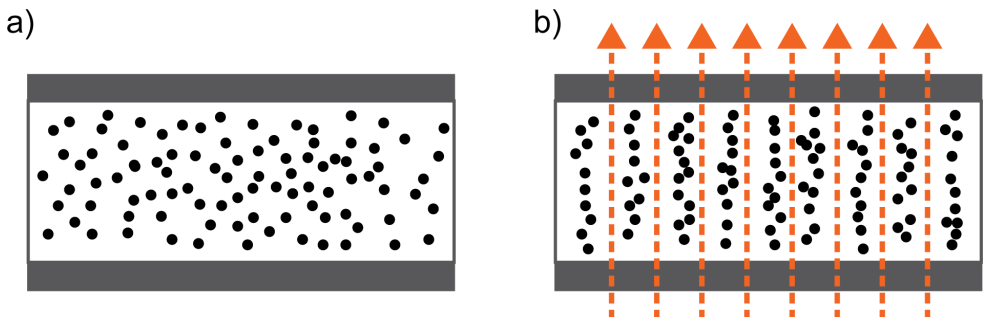

Figure 2.7. MR fluid activation. a) When no magnetic field is applied, ferromagnetic particles (black dots) are distributed randomly. b) When a magnetic field is applied, particles are aligned along magnetic flux.

the direction of magnetic flux (Figure 2.7b). This increases the effective viscosity of the fluid. A rapid change of the viscosity can be sensed as vibration when the user touches the surface of the material.

In many cases, the ability of changing the viscosity of an MR fluid is desirable for producing tactile feedback, and thus electromagnets are usually accompanied to control the strength of the magnetic field applied to the fluid. Jansen *et al.* (2010) developed a touch-sensitive surface with an MR fluid, called MudPad. Eighty-four electromagnets are placed underneath the fluid, and the system can control the viscosity of the fluid at different locations of the surface. In this manner, MudPad offers tactile feedback on tabletop interactions.

An electrorheological fluid (ER fluid) shares a similar structure to an MR fluid. Instead of an electromagnetic field, an ER fluid is activated by an electrical field.

Mechanical Pin Arrays

Pin arrays is another common device to produce tactile feedback. Each pin can be very small and the entire pin arrays can fit to a user's fingertip or a part of the hand. For example, in the VirTouch device used in the study by Wall & Brewster (2006), 4 x 8 arrays of pins are

placed to fit to a user's fingertip. Each pin can be controlled so that the surface of the device can form asperity or a pin can produce vibration.

Optacon is one of the early commercially available devices incorporating pin arrays (Linvill & Bliss, 1966). Optacon supports visually impaired users to read printed materials. Craig & Sherrick (1982) found that with Optacon, participants achieved 10–12 WPM after the nine-day training. Currently, similar hardware can be seen in many refreshable Braille displays.

Ultrasound Transducers

An ultrasound transducer is a device generating ultrasound waves from an electrical signal. Because a sound wave is a rapid change of air pressure, users can feel ultrasound with sufficient strength as tactile sensation. An advantage of this technology over simply using a rapid air flow or an air jet (Suzuki & Kobayashi, 2005) is that ultrasound can realize a finer output resolution because it is more robust to diffusion in the air than an air jet.

A system with ultrasound transducers is useful when the system wants to provide "mid-air" tactile feedback. Iwamoto *et al.* (2008) integrated the technology into a volumetric display to produce tactile feedback on a 3D image. Their system uses 91 ultrasonic transducers to produce mid-air single-point tactile feedback. Hoshi *et al.* (2009) further added an ability to change the focal point of vibration to this technology. Alexander *et al.* (2011) used a similar technology into a mobile TV device to enhance its user experience. However, this technology usually involves many transducers to produce perceivable air pressure .

Electrovibration

Electrovibration is a phenomenon initially found by Mallinckrodt *et al.* (1953), and started to be applied to haptic displays by Strong & Troxel (1970). The effect of electrovibration can be explained as follows. A conductive surface and a finger contacting the surface can be

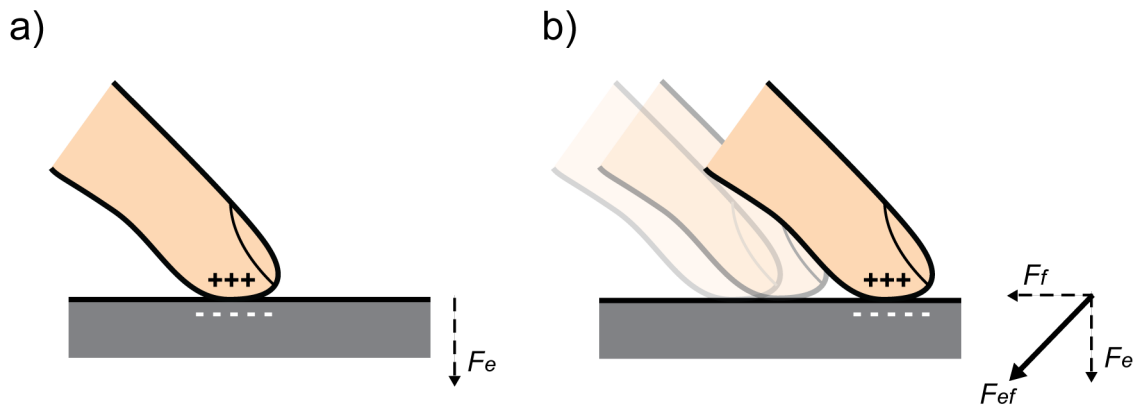

Figure 2.8. The mechanism of electrovibration. a) When the user's finger stays static on the surface, an electrically induced attractive force (F_e) is not strong enough to be perceived. b) When the user drags a finger on the surface, a shear force (F_{ef}) is generated and deforms the skin. As a result, the user feels friction underneath the finger.

considered as a capacitor; the outer skin works as the dielectric layer of this imaginary capacitor, and the conductive surface and tissues in the finger are two conducting plates. When alternating current is applied to the conductive surface, the imaginary capacitor is continuously charged and discharged, and an attracting force across the imaginary capacitor is generated (F_e in Figure 2.8a). At a power line frequency (50 or 60 Hz), the current between this imaginary capacitor is much smaller than the smallest perceivable current (Kaczmarek *et al.*, 1991). However, when the user is dragging the finger on the surface, a shear force is generated (F_{ef} in Figure 2.8b), and it deforms the skin. This deformation is perceivable even if it is very small, and as a result, the user receives a "rubbery" feeling.

TeslaTouch (Bau *et al.*, 2010), and E-Sense Technology from Senseg corporation are applications of electrovibration to touch-screen devices. Both devices offer sensation of different frictions controllable by a computer. Unlike other types of hardware, this technology can produce tactile feedback directly to the finger contacting on the surface. However, one shortcoming is that tactile feedback is perceivable only when the user is dragging the finger.

Force Feedback Hardware

The technologies discussed above are for generating tactile feedback. However, haptic feedback is not only limited to vibration, and many HCI projects use other types of feedback, such as force feedback. This work does not use force feedback, but I briefly review haptic feedback hardware, particularly force feedback hardware, because it is often used in user interfaces for mobile devices.

Robotic Arms

Many of force feedback devices use mechanical components similar to robotic arms. The part held by the user is connected to mechanical arms and joints. The system produces force feedback (*e.g.*, pulling the user towards a particular direction) by controlling these arms and joints. PHANTOM is one of the most common force feedback devices used in HCI research. The user holds the stylus part of the device. The system can sense the user's motion as well as generate force feedback when necessary.

Mechanical Weight Shifting

Mechanically changing the displacement of a weight in a handheld device can offer the user sensation as if an object inside the device is moving or the user is being pulled by the device. Amemiya & Sugiyama (2009) developed a handheld device which can control the vertical displacement of a weight inside. Slow movements of the weight are not clearly perceivable whereas its rapid moments cause a sensation of being pulled or pushed in a particular direction. With this, they developed a path navigation system for visually impaired users. Hemmert *et al.* (2010) extended this idea and developed a device which can control the 2D displacement of a weight. Their user study found that participants could identify the weight's location within a few centimeter error.

Virtual Springs

Gupta *et al.* (2010) developed a mobile device which can produce a repulsive force on its surface (a pressure plate). When the user grasps the device, the system senses the pressure applying to the device. The built-in motor starts to drive and pushes the pressure plate back. As a result, the user would feel as if she is squeezing a spring. The string constant of this virtual spring can be programmatically controllable. Their user study revealed that four different levels of the string constant are easily distinguishable.

Summary

I discussed a variety of technologies to generate tactile feedback in user interfaces. One objective of this work is the development of spatially-enhanced vibrotactile feedback. Thus, the system needs the ability to produce different spatial patterns. Furthermore, the hardware needs to be small so that it can be incorporated into mobile devices. DC vibration motors can be small and the circuit to drive the motor is straightforward. Therefore, I determined to use DC vibration motors; my custom hardware embeds DC vibration motors into different locations on the backside of a mobile device. In this manner, my hardware can produce vibrotactile feedback on different locations of the fingers and palm of the user's hand holding the device.

2.3 Vibrotactile Feedback Parameters

By far, vibrotactile feedback is a widely used modality in user interface designs. This work also uses vibrotactile feedback for interfaces on mobile touch-screen devices; thus, it is important to understand how expressive vibrotactile feedback can be. Geldard (1960) discussed parameters in vibration which can be controlled programmatically and can be sensed by humans. Brewster & Brown (2004) brought them to the context of user interface designs in their exploration of the Tacton system. Here, I explain the parameters defined by them.

I use the term "just noticeable difference" (JND) to explain the human sensitivity of each parameter. JND means the smallest detectable difference between an initial (l_i) and given level (l_g) of a stimulus; thus, $JND = |l_g - l_i|$. I also refer to $|l_g - l_i|/l_i$ as a JND ratio.

Intensity

Intensity means the strength of vibration. It is also often called amplitude because in most cases, the amplitude of electrical signals for generating vibrotactile feedback is correlated to the intensity of the feedback. This parameter is easy to control in most hardware; thus, it is probably the most frequently used parameter.

According to Geldard (1960), a useful range of stimulus amplitude is 50–400 μm in the chest area. 50 μm is the lower limit where a person can perceive stimuli at 100% accuracy. 400 μm is the upper limit where a person would not have any discomfort with stimuli. Verrillo & Gescheider (1992) found that the upper limit of the vibration intensity where people would not have discomfort is 55 dB.

Geldard's study also explained that people can detect fifteen different levels of intensity when an experimental procedure is carefully controlled. Craig (1972, 1974) reported that JND is close to 2 dB at the sensation levels of 20 dB regardless of the existence of background vibration. Bolanowski Jr. *et al.* (1988) found that JND can be about 2.5–0.7 dB when the stimulus intensity is 4–40 dB. It would be difficult to determine a precise value of an intensity JND because it depends on the stimulus intensity; however, 1–3 dB would be a good estimation.

However, it is also known that the perceivable resolution of vibration intensity is limited in a realistic setting. Geldard (1960) suggested that more than three different levels would become hard to distinguish. van Erp (2002) also discussed that although encoding information in this parameter is possible, the number of the intensity levels should be no more than four with sufficient differences between any two of the levels.

Although humans may not have a good ability to distinguish the absolute value of the vibration intensity, they may be sensitive to the change of the intensity. Brown *et al.* (2006b) investigated how accurately people can recognize different intensity changes (*e.g.*, vibration starting with low intensity and ending with high intensity). Their study revealed that participants could identify three different intensity changes (*i.e.*, increasing intensity, decreasing intensity, and no change in intensity) at 92–100% accuracy.

Frequency

Frequency is the number of vibration cycles happening in one second (the unit is Hz). As discussed in Section 2.1, the perceivable range of the vibration frequency is fairly wide; humans can sense up to 1000 Hz quite well with the sensitivity peak at 250 Hz (Verrillo & Gescheider, 1992).

However, investigations on the human sensitivity against the vibration frequency were not as straightforward as its definition. One reason is that the subjective intensity of vibration can change when the vibration frequency changes. Furthermore, hardware may have an internal "filter" against the vibration intensity in the frequency domain. This means that when the vibration frequency is controlled, the vibration intensity may increase or decrease even if the other factors are maintained to be the same (*e.g.*, the amplitude of an electrical signal). They are the aspects which HCI researchers should consider when they decide to use the frequency parameter.

Early psychological explorations on frequency JNDs did not consider the variation on subjective intensity, and as a result, the reported sensitivity in the frequency domain is not high. Knudsen (1928) found that an average JND ratio was 15–30% when the stimulus level was 34 dB at 64–512 Hz, and the stimulus was delivered to the index fingertip. Mowbray & Gebhard (1957) showed that a JND ratio was 2–8% when the stimulus frequency was 1–320 Hz.

Goff (1967) conducted an experiment controlling the subjective intensity, and found that a JND ratio was 18–36% with stimuli at 20 dB and 31–55% with stimuli at 35 dB. Sherrick (1985) also observed a low sensitivity of the frequency, and concluded that 3–5 different levels are reliably distinguishable within the range of 2–300 Hz. These results indicate that the human skin may not be very sensitive to absolute values of the vibration frequency.

However, humans may be sensitive to changes of vibration frequencies. The study by Rothenberg *et al.* (1977) implies that people demonstrate smaller JNDs when a stimulus changed its frequency over time than when two stimuli with different frequencies were given. However, these types of studies are not well-standardized yet, and further investigation is necessary to accurately estimate the sensitivity to relative values of the frequency parameter.

Duration

Duration means the temporal length of vibration. Geldard (1957, 1960) found that people can have approximately 25 JNDs between 0.1 and 2.0 sec with 50 msec difference at the low end and 150 msec difference at the high end. Thus, in general, the human sensitivity of the duration is considered high. However, if 100% accuracy needs to be achieved, four or five levels with larger differences would need to be chosen.

Gescheider (1966) investigated the human capability of perceiving the difference of activation timing of two stimuli at different parts of fingers. He prepared three settings: two stimuli on one index fingertip; one stimulus on one index fingertip and the other on the ring fingertip of the same hand; and one stimulus on each index fingertip. When the intensity of the second stimulus was 5 dB stronger than the first stimulus, people could differentiate the timings of stimulus activation with the minimum time difference of 10.0–12.5 msec.

These studies indicate that the duration parameter offers a higher output resolution than other parameters. However, note that the experiments in these studies removed external factors.

In a realistic setting, users may focus on other tasks and may not be attentive to feedback. Thus, short vibrotactile feedback can be easily missed or mis-recognized. As Geldard (1960) suggested, a limited number of well-separated levels should be used for user interface designs.

Waveform

Waveform means specific temporal vibration patterns. Examples of different waveforms are a sine wave, a square wave, and a triangular wave. Geldard (1960) suggested that different waveforms can be distinguishable if the basic frequency is low enough. However, this parameter has not been explored deeply, and it is still unclear how well humans can perceive and distinguish different waveforms through the tactile channel.

Rhythm

Similar to *waveform*, *rhythm* means temporal vibration patterns, but the information is encoded in the combination of other parameters. For example, one can use the duration parameter, and the number of vibration stimuli to encode a Morse code (*e.g.*, one short vibration followed by one long vibration means "A"). This parameter is not suggested by Geldard (1960). However, as computer devices and vibrotactile feedback devices are advanced, a precise control of temporal patterns becomes possible. Thus, this parameter has been used in user interface designs.

Summers (2000) used this parameter with changing the intensity and frequency parameter to encode speech information in vibrotactile feedback. They found that participants tended to use temporal patterns to obtain information rather than the other two parameters. Brewster & King (2005) encoded the progress of a file download in the temporal gap between one short pulse and four short pulses. The temporal length of the gap is proportional to the amount of the download remaining. They revealed that participants responded more quickly to the completion of a download with this tactile progress indicator than to the visual one. Brown

et al. (2005) used three different musical rhythms encoded in vibrotactile feedback to inform users different types of events (*i.e.*, an incoming voice call, an incoming text message and an incoming multimedia message). Their study found that participants could distinguish these three patterns at on average 92.7% accuracy.

Locus (Spatial Patterns)

Locus means locations on the body. I also refer to this parameter as *spatial patterns*. In this parameter, vibration at different body parts represents different information. This parameter has been deeply explored over decades. The sensitivity of the locus parameter are usually measured in two ways: localization accuracy and differentiation of patterns.

Localization accuracy measures how accurately people can recognize vibration at different points of the body. Howell (1958) and Geldard (1960) showed that people could differentiate vibration sources attached to seven different positions of the rib cage. Geldard (1960) also reported that the chest can accommodate five vibration sources. Cholewiak & Collins (2003) examined the human capability of vibration point localization on the arm. Their study with seven different points showed that participants were not accurate (about 30–40% accuracy) except the locations close to the wrist or elbow (about 65–70% accuracy). Brown *et al.* (2006a) tested the distinguishability of three vibration points on the user's arm. They found that participants could distinguish the vibration location at approximately 95% accuracy. This indicates that the distinguishability can be improved when the number of the vibration points is small.

The differentiation of patterns pertains to how accurately people can distinguish two different spatial patterns. This is impacted by a masking effect (Sherrick, 1964). The masking effect can be temporal (a vibration stimulus decreases the detectability of another which is happening right after the first one) or spatial (the vibration stimulus at one location decreases the detectability of the stimulus at another location). Cholewiak *et al.* (1995) examined how

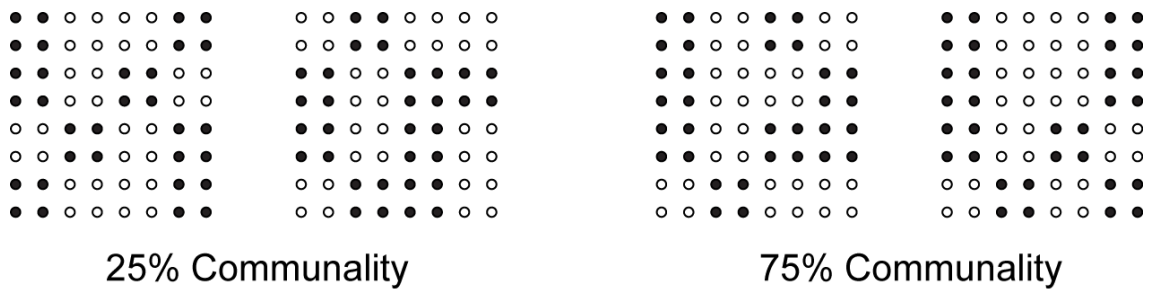

Figure 2.9. Examples of patterns used in the study by Cholewiak *et al.* (1995). This figure also illustrates the notion of communality—how similar two tactile patterns are to each other. Black dots represent pins generating vibration. (reproduced based on Cholewiak *et al.* (1995))

accurately people can distinguish two different patterns while the similarity of the patterns (referred as "communality" in their work) was controlled. They used an 8×8 pin array (each pin generated vibrations when activated) to generate spatial patterns. Figure 2.9 shows two sets of patterns with different levels of communality, which means how much the two patterns are overlapped. Their study indicates that accurately distinguishing two different patterns presented on the palm was difficult even if the value of communality was zero (*i.e.*, the two patterns were completely different in terms of spatial relationships). This suggests that people may not be able to easily distinguish spatial patterns generated by multiple simultaneously-activated vibration sources.

In summary, these studies indicate that people can only differentiate spatial patterns if two requirements are met. First, systems should avoid using multiple vibration sources simultaneously to encode information in spatial patterns due to the spatial masking effect. And vibration sources need to be well-separated.

Spatio-temporal Patterns

Geldard (1960) referred the spatio-temporal parameter as "spatially discrete loci," but

here I follow the terminology used by Brewster & Brown (2004) because it describes the characteristics of the vibration patterns in this parameter better. This represents patterns which changes its spatial property over time—patterns which can "draw" on the user's body according to Brewster & Brown (2004). This perception is related to a psychological illusion found by Geldard & Sherrick (1972), called "cutaneous rabbit illusion."

A few aspects of this parameter has been examined. For example, Stolle *et al.* (2004) tested how variations in a stimulus onset interval (a temporal interval between the end of one stimulus to the start of another; SOA) can affect user perception on the location of an illusionary vibration stimulus between the two actuators. They found that changing SOA can result in different perceived stimulus locations, but its effects are different depending on the part of the body. However, a comprehensive set of spatio-temporal patterns can be very large, and defining an objective unit like a JND is not straightforward. Thus, further efforts are necessary to understand various aspects of the spatio-temporal patterns, such as its distinguishability.

Nevertheless, the spatio-temporal parameter has started to be explored in HCI research. Hoggan *et al.* (2007) attached four voice coil motors on a mobile device, and used a circular pattern to inform the user the status of a file download. Israr & Poupyrev (2011) developed a chair with 12 vibration motors. Their system can produce various spatio-temporal vibration patterns on the back of the user sitting on the chair to enhance user experience in playing computer games.

Although this spatio-temporal parameter is attracting attention from researchers, it has not been explored well, particularly in the user's hand and fingers. Because previous research indicates that the spatial parameter is useful to encode information, further exploration on the spatio-temporal parameter will contribute to improvement on the expressiveness of vibrotactile feedback.

Summary

This section discussed what parameters vibrotactile feedback has to encode information and how well humans can sense and distinguish different patterns varied in each parameter. I found that the spatio-temporal parameter is still under-explored although prior studies about the spatial parameter imply that spatio-temporal patterns could be used to convey information to the user.

2.4 User Interfaces with Tactile Feedback

In the previous section, I discussed how expressive vibrotactile feedback can be. Although tactile feedback is not necessarily appropriate to convey complex information (*e.g.*, textual information or images), it can encode simple messages. This can be useful for a number of user scenarios in interactive systems (*e.g.*, providing a confirmation cue when the user touches a virtual object in a system). Thus, tactile feedback has been used in various interactive systems.

I review user interfaces incorporating tactile feedback in this section. User interface designs with tactile feedback were started in the field of virtual reality, such as Z-Glove developed by Zimmerman *et al.* (1987). Since then, tactile feedback has been integrated into a variety of systems, including interfaces for desktop computers, large displays, and wearable devices. Here, I focus on two interface properties closely related to my work: touch screens and touch-sensitive surfaces, and mobile devices. For touch screens and touch-sensitive surfaces, different input modalities with tactile feedback are available, such as pen input (Lee *et al.*, 2004; Poupyrev *et al.*, 2004). However, I mainly focus on interactions using direct touch input (*i.e.*, input with bare hands) as it is the input modality I used in this work.

Touch Screens and Touch-sensitive Surfaces

Fukumoto & Sugimura (2001)

Fukumoto & Sugimura (2001) perhaps were the first researchers to bring tactile feedback to touch-screen devices. Their system, Active Click, uses a voice coil actuator. When the user interacts with the device (*e.g.*, tapping a button), it generates vibration. They also demonstrated two different ways to convey vibrotactile feedback by attaching a vibration motor to the front side or backside of the mobile device. In this manner, Active Click can provide feedback to the finger contacting the screen or to the hand holding the device. In their preliminary user study, Active Click helped participants interact with a calculator application on the system faster in both quiet and noisy environments.

Poupyrev et al. *(2002a) and Poupyrev & Maruyama (2003)*

Even with only one actuator, it is possible to produce different vibration patterns by using different intensity levels and frequencies. Poupyrev *et al.* (2002a) developed the Ambient Touch system with a mobile device. It uses their TouchEngine actuator (Poupyrev *et al.*, 2002b) to change parameters of intensity, frequency, and rhythm. The TouchEngine actuator consists of layers of piezoceramic films and printed adhesive electrodes. The material on the top has an opposite polarity to that on the bottom. When a voltage is applied to the actuator, it bends towards one direction. This bending can be used to generate vibrotactile feedback.

They also showed that different vibration patterns can be used to convey information, such as the user's scrolling speed and position on the screen. In their user study, participants performed linear list menu selection tasks with a tilt-based interface (*i.e.*, scrolling the list in one direction by physically tilting the mobile device). Through their user study, they found that participants were able to select items in a linear list 22% faster than when no tactile feedback was provided.

Poupyrev & Maruyama (2003) further extended their interface to provide different vibration patterns for different gesture-based interfaces on mobile touch-screen devices. For instance, when the user initially contacts an object on a screen, the device generates a "springy" feeling by using a sine wave signal combining low frequency and high frequency signals. When the user is holding the object, the device generates continuous pulse vibrations. But when the user holds the object longer than the pre-defined time threshold, the system assumes that the user has been interrupted by other tasks, and offers low frequency vibration to inform the user that the contacted object is still active.

Kaaresoja (2006)

Kaaresoja (2006) developed a mobile touch-screen device with a piezo actuator, and demonstrated four interactions with tactile feedback: A numeric keyboard, text selection, scrolling and drag-and-drop interaction. These applications were also demonstrated or suggested by Poupyrev *et al.* (2002a), but Kaaresoja added several features to the applications. However, no informal or formal user study on the system is reported.

Hall et al. *(2008)*

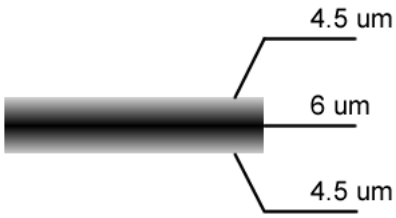

Figure 2.10. The T-Bar widget and its application. A darker color represents stronger vibration.

Hall *et al.* (2008) designed a graphical widget, called T-Bar, which represents vibration intensity. As shown in Figure 2.10, a color indicates vibration intensity. Their prototype system can change the intensity range between 4.0 and 6.5 μm with 0.5 μm increment. When the user crosses the T-Bar, she would have sensation like touching a round pipe. This widget can be used in various applications. Hall *et al.* demonstrated two applications of T-Bar: File-o-Feel (a linear list application) and Touch 'n' Twist (a launcher application). Their preliminary user study showed that although T-Bar potentially could reduce mental workload during interactions on mobile touch-screen devices, participants did not always feel the difference in vibration intensity.

Rantala et al. *(2009)*

Rantala *et al.* (2009) developed a method for presenting Braille characters on a mobile touch-screen device. Their system uses different peaks of the pulses to generate raised and lowered dots. Their experiment with three different presentation methods of Braille tactile feedback revealed that experienced Braille users could recognize letters at 91–97% accuracy.

Harrison & Hudson (2009a)

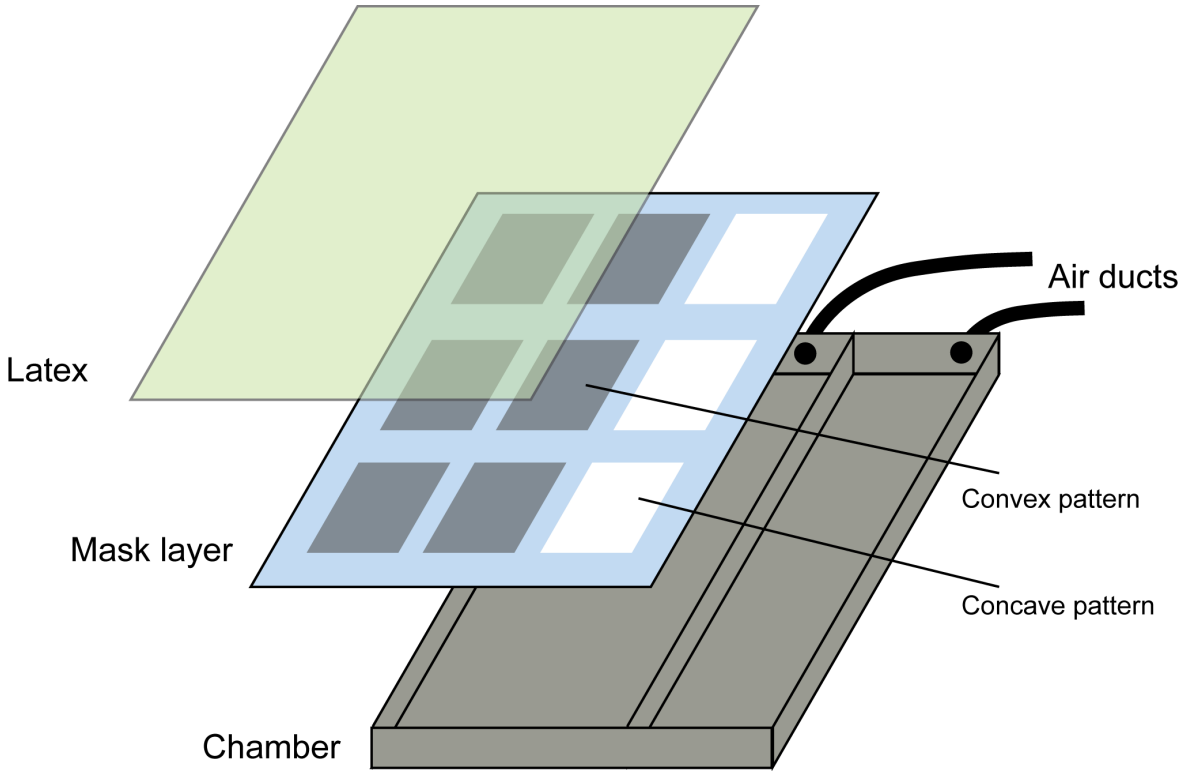

Figure 2.11. The structure of a deformable surface. The surface is dynamically deformable by changing the air pressure in the chamber covered by the latex and mask. Depending the pattern, a part of the surface over each mask pattern can be a convex or concave. (Harrison & Hudson, 2009a)

Harrison & Hudson (2009a) developed a system which can make a part of or the entire surface a convex or a concave. The surface is covered by a latex, and a pre-defined mask is placed underneath the latex. The back of the device forms a chamber, and an air duct is connected to control the amount of air in the chamber. When the air is injected, the latex over the mask is pushed and forms a convex. When the air is vacuumed, the latex is pulled and forms a concave. Thus, by changing the air pressure in the chamber, the system can deform the surface programmatically.

Jansen et al. *(2010)*

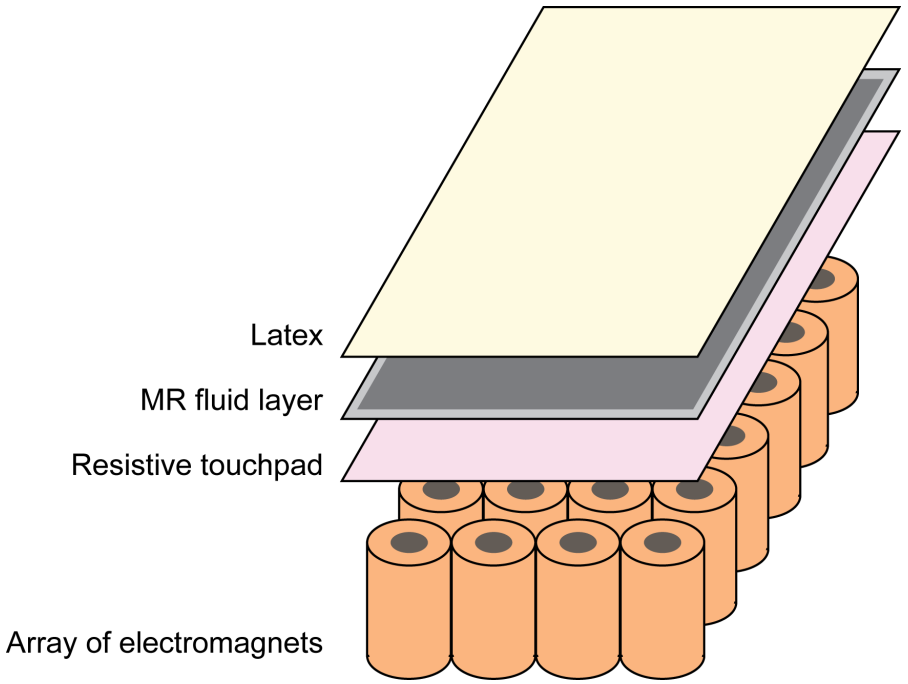

Figure 2.12. The structure of MudPad. A 2D array of electromagnets control the viscosity of a particular portion of the MR fluid. Visual information can be projected onto the latex from a projector on the ceiling. (Jansen *et al.*, 2010)

Jansen *et al.* (2010) developed a mid-sized (10-inch) touch-sensitive surface using an MR fluid, called MudPad. As shown in Figure 2.12, an MR fluid layer is placed over a resistive touchpad. An array of electromagnets (12 $\times$ 7) are located underneath the MR fluid layer and touchpad. The system can control the apparent viscosity of the MR fluid with this electromagnet array. Their system can produce localized vibration by rapidly turning electromagnets on/off as well as static haptic sensation (*e.g.*, solid or fluid feeling). They did not report any formal or informal user study about the system.

Bau et al. *(2010)*

Bau *et al.* (2010) developed a touch-screen device providing tactile feedback by electrovibration. Their system, called TeslaTouch, consists of a glass place (which is used as a projection screen for the projector underneath) and a transparent electrode and insulator over the plate. When the user is moving a finger on TeslaTouch, she can feel electrically generated friction. The intensity and texture of friction are controllable by changing the amplitude and frequency of an electrical signal to apply to the electrode, respectively. Their user study found that the average frequency JND varied from 11–25% at 400–120 Hz, which is relatively in line with findings reported by Goff (1967) (18–36% thresholds), and reported by Israr *et al.* (2006) (13–38% thresholds). They also found that the average amplitude JND was 1.16 dB and constant across all frequencies, which was slightly lower than what was reported previously (1.5–2.5 dB) by Israr *et al.* (2006) and Verrillo & Gescheider (1992).

Handheld and Mobile Devices

Oakley & O'Modhrain (2005)

Oakley & O'Modhrain (2005) used vibrotactile feedback on tilt-based interfaces. Poupyrev *et al.* (2002a) also incorporated this type of interactions in their Ambient Touch system; the user can scroll down or up a linear list by tilting the device. The angle of the tilt determines the scrolling speed. Vibrotactile feedback is provided when the user crosses an item while scrolling the list. However, this design often causes over-shooting (Poupyrev *et al.*, 2002a). Oakley & O'Modhrain (2005) instead used the absolute value of the angle; for example, having a device horizontally selects the first item, and having the device vertically selects the last item in the list. Their system has two settings of vibrotactile feedback. In one setting, vibrotactile feedback is provided when the user moves to the next or previous item. The other setting uses different vibration intensity levels based on the item's position in the list (*i.e.*, the lowest and highest

intensity for the first and last item, respectively). Their main focus on the evaluation was the comparison between absolute and relative list navigation, and it is not clear how effectively vibrotactile supported menu selection with tilt-based interfaces from their experiment.

Luk et al. *(2006)*

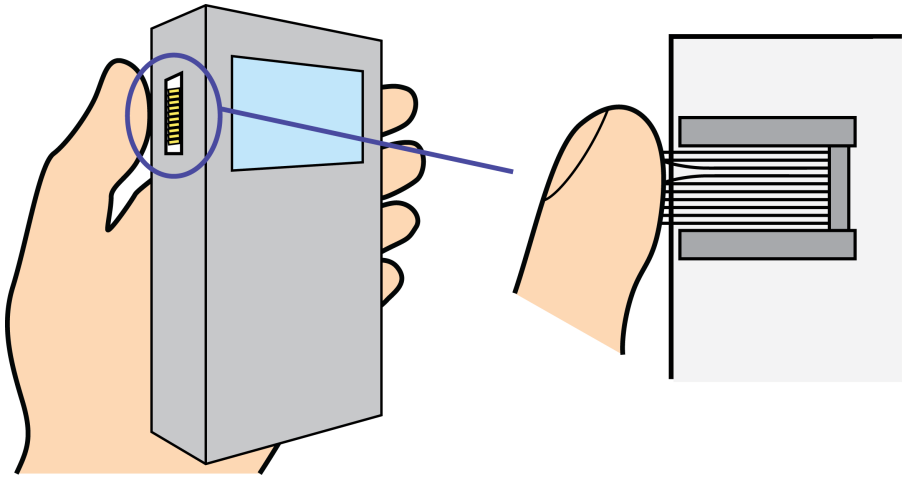

Figure 2.13. The THMB prototype. The piezoelectric actuators are exposed at the left side of the prototype mobile device. They contact the user's thumb and provide tactile feedback (Luk *et al.*, 2006)

Luk *et al.* (2006) explored how different stimuli generated by layered piezoelectric benders could produce different subjective perception. The eight layers of piezoelectric benders on their mobile device (Figure 2.13) can be controlled independently. Thus, their system can produce 1D spatio-temporal patterns as well as use different intensity, duration, and waveform patterns. Their user study examining the distinguishability of the patterns controlled by the four parameters revealed that those parameters could offer different sensation. In particular, they found that different levels in the intensity and duration parameters could be mapped orthogonally with respect to subjective perception.

Hoggan et al. *(2007)*

Multiple actuators can be embedded in mobile devices to produce a rich set of vibration patterns that could be used to provide users with tactile feedback in a "background" channel. Hoggan *et al.* (2007) demonstrated this concept by using tactile feedback to inform users about the progress of file downloads while they work on other tasks. Their device includes four voice coil actuators (two on the left side, one on the right side, and one on the backside of the device). These actuators can generate a circular vibration pattern on the user's hand holding the device by carefully choosing the timing of activating each actuator. Their user study revealed that participants could respond to the completion of a background task (a file download in their study) faster while they were involved in the main task of typing text in comparison to a visual status notification.

Ghiani et al. *(2008)*

Ghiani *et al.* (2008) developed a guide system which uses two strips with vibration motors worn by people with visual impairments on their thumb and index finger of the hand holding the device. They used vibration motors to present the direction to turn (*e.g.*, vibration from the right side means "turn right"). However, their user study with eleven visually impaired users did not find significant differences on performance of navigating tasks (reaching a specific artwork in a museum) in comparison to an audio guide.

Sahami et al. *(2008)*

Sahami *et al.* (2008) examined how accurately users can distinguish vibration generated by six vibration motors located along the left and right sides of a mobile phone (three motors on each side). Their device is a mobile phone with physical buttons, and their primary research objective was to understand how accurately users can recognize spatial and spatio-temporal patterns. Their experiment showed that the participants could distinguish eight vibration

patterns at 70–80% accuracy, but they had difficulty identifying the location of the vibration source when their system activated only one of the vibration motors at a time (36% on average).

Perception of Vibrotactile Feedback

Prior work also has investigated how differently users perceive various vibration patterns. In this section, I summarize the work and its findings.

Nashel & Razzaque (2003)

Nashel & Razzaque (2003) used the frequency parameter to convey information about which row in the numeric keyboard on a mobile touch-screen device the user is touching. They used 40, 200, 250, and 350 Hz to indicate four rows in the keyboard. They also provided a short pulse when the user contacts a button for the first time and when the user leaves the button. Their informal user study showed that most users were able to differentiate the rows through the tactile cue, but they did not report the actual accuracy data for the study.

Hoggan et al. *(2008b)*

Hoggan *et al.* (2008b) demonstrated that the texture of physical buttons could be mapped into vibrotactile parameters that are used later to be produced as tactile feedback for buttons on touch-screen devices. They used two different vibration sources (a DC vibration motor and voice coil actuator), and prepared two different patterns for each vibration source. Their study showed that different vibration sources can produce different perception of vibration, and this difference can be useful to represent different types of objects (*e.g.*, different button shapes).

Koskinen et al. *(2008)*

Koskinen *et al.* (2008) examined subjective preferences on vibration patterns with different intensity levels and durations. They also used two different types of hardware: a piezo motor

and DC vibration motor. Their investigation focused on button presses, where vibration was given when a user pressed a software button on the screen. They found that participants perceived the feedback with a piezo actuator slightly more pleasant than one with a DC vibration motor although no statistically significant difference was found.

Altinsoy & Merchel (2009)

Altinsoy & Merchel (2009) investigated the effects of different waveform vibration patterns on number entering tasks on touch screens. They prepared five different patterns: half a cycle of a sine wave, one pulse of a triangular wave, one pulse of a sawtooth wave, half a cycle of a $\sin^2$ wave, and a 50 Hz sine wave. Their study found the accuracy of number typing was improved when tactile feedback was provided, but they did not observe any remarkable difference across their waveform patterns.

Lylykangas et al. *(2011)*

Lylykangas *et al.* (2011) examined subjective preferences on vibration patterns varied in the duration parameter and the amount of delay. The range of their duration parameter was 10–198 msec, which was a wider range than ones studied by Hoggan *et al.* (2008b) and Koskinen *et al.* (2008). They found that the duration around 66 and 132 msec was preferred over 10 and 198 msec. They also found that a long output delay was not preferred independently of the durations.

Effects of Vibrotactile Feedback

Finally, I discuss prior work mainly investigating how vibrotactile feedback can improve user performance on various tasks in mobile devices.

Leung et al. *(2007)*

Leung *et al.* (2007) examined how tactile feedback can improve user performance on interactions with different widgets in graphical user interfaces on mobile touch-screen devices. They set up four different tasks in their user studies, and found that tactile feedback provided improvements on performance time on two kinds of their tasks. One task in which improvement was observed was a progress bar task, where participants expressed when a background task was completed while they performed a typing task. The intensity and frequency of the tactile cue indicated the percentage of the progress (1–10 Hz for 0–100% progress). The other was a scroll bar task, where participants had to find a sentence with a given keyword by scrolling the screen, and short pulse vibration was offered when the keyword appeared in the screen.

Hoggan et al. *(2008a)*

Hoggan *et al.* (2008a) studied the effects of vibrotactile feedback on typing tasks with the mini-QWERTY software keyboard. Their first comparative study against the physical mini-QWERTY keyboard and the software mini-QWERTY keyboard without tactile feedback showed that the keyboard with tactile feedback improved typing accuracy significantly, and its accuracy was comparable to the one with the physical keyboard. Their second study included an interface using two vibration actuators. One or both actuators were used to provide vibrotactile feedback depending on where the user contacted. For example, when the user touched "A," "G," or "L" in the mini-QWERTY keyboard, the system used the left actuator, both actuator, and the right actuator, respectively. Their results from the second study revealed that vibrotactile feedback by multiple actuators can improve text entry performance better than one by a single actuator.

McAdam & Brewster (2009)

McAdam & Brewster (2009) examined the effects of tactile feedback given at different body locations on typing tasks; thus, their tested parameter was locus. Unlike the experiment by Hoggan *et al.* (2008a), their experiment included vibration at the user's wrist, her upper arm, her chest, her waist (over a belt), and her front thigh (over a pocket in pants) as well as vibration on the device. They found that vibration at the wrist and upper arm improved text entry speed over the other conditions while the accuracy was comparable across all the conditions.

Pasquero & Hayward (2011)

Pasquero & Hayward (2011) investigated how frequently users would need to view visual feedback in a mobile device during linear menu selection tasks with and without tactile feedback. They used the THMB device (Luk *et al.*, 2006), and set up spatio-temporal tactile feedback patterns with two different durations. Vibration with a short and long duration was used when the user passes each item and each ten items in a linear list, respectively. Spatio-temporal patterns (moving upward or downward in their system) were also used to indicate the scrolling direction. Their user study showed that the number of the times when the screen was viewed decreased by 28% when tactile feedback was enabled while the performance time and accuracy did not show significant differences.

2.5 Summary

This chapter discussed tactile feedback from various perspectives. Section 2.1 describes the human capability of sensitivity to vibration. It highlights that the human skin has high sensitivity to vibration at 100–1000 Hz (with 250 Hz peak), and the limit of the spatial resolution is probably about 1 cm. Section 2.2 discusses various types of hardware to produce tactile feedback, and I draw the conclusion that a DC vibration motor is very appropriate hardware for producing vibrotactile feedback in mobile devices.

Section 2.3 then discusses the parameters of vibrotactile feedback. My survey along these parameters highlight that although some of these parameters have been examined deeply, the spatio-temporal parameter is under-explored. Finally, Section 2.4 describes tactile feedback in touch-screen devices and mobile devices. It shows that use of spatial and spatio-temporal patterns in mobile devices is still at an early stage.

In summary, spatial and spatio-temporal patterns can increase the expressiveness of vibrotactile feedback. This motivated me to investigate how spatial tactile feedback can be incorporated into mobile touch-screen devices and how it can support various interactions on these devices. The following chapters present how I investigated this with my custom hardware.

Chapter 3

Spatial Tactile Feedback Hardware

My literature survey indicates that although tactile feedback has been explored in various types of computer devices, its expressiveness is limited on mobile devices. Particularly, the spatial and spatio-temporal parameters has been under-explored. This motivated me to investigate how spatial and spatio-temporal vibration patterns can be incorporated into a mobile device, and how spatial tactile feedback can support various tasks on mobile touch-screen devices. In this section, I discuss my hardware prototype to produce spatial tactile feedback on mobile touch-screen devices.

3.1 Design Requirements

Based on the findings through my literature survey and physical designs of mobile devices, I have three important design requirements for my hardware.

- The size of hardware needs to be small.
- Vibration motors need to be well-separated (preferably longer than 1 cm).
- The hardware should be designed as an external accessory.

The last requirement is for maintaining the flexibility of the hardware. Instead of embedding vibration motors directly in a mobile device, this design allows users to attach or detach the hardware when they like. Furthermore, users would need to buy only this accessory to enjoy benefits of spatial tactile feedback rather than entirely replace their devices with a new one.

Based on these requirements, I decided to use small DC vibration motors for my hardware. The package of DC vibration motors are already small (12mm wide diameter in my hardware). A circuit to drive these motors does not need a complicated design, and it can be downsized greatly. This choice of the hardware may limit the expressiveness of the same vibrotactile parameters discussed in Section 2.3, such as frequency. However, my main objective is to investigate the effectiveness of spatial or spatio-temporal patterns, and my specific hardware design is intended to enable exploration of this through different applications.

I excluded other design requirements, such as power consumption. They are important for deployment of my systems in a realistic setting. However, the three requirements above highlight the most important aspects of hardware design. The hardware I prototyped allowed me to conduct in-depth laboratory user studies for understanding the effects of spatial tactile feedback described in the later chapters.

3.2 First Prototype

It is difficult with the current technologies to provide tactile feedback exactly at the contact point on the touch screen with vibration motors embedded inside the mobile device. Therefore, my design instead aims to provide feedback on the palm and fingers of the hand holding the device. Although it is known that the palm is less sensitive than the fingertip (or "fingerpad"), the study conducted by Craig & Lyle (2001) showed that a similar level of spatial distinguishability to that on the fingerpad can be achieved by sufficiently enlarging the stimulus space.

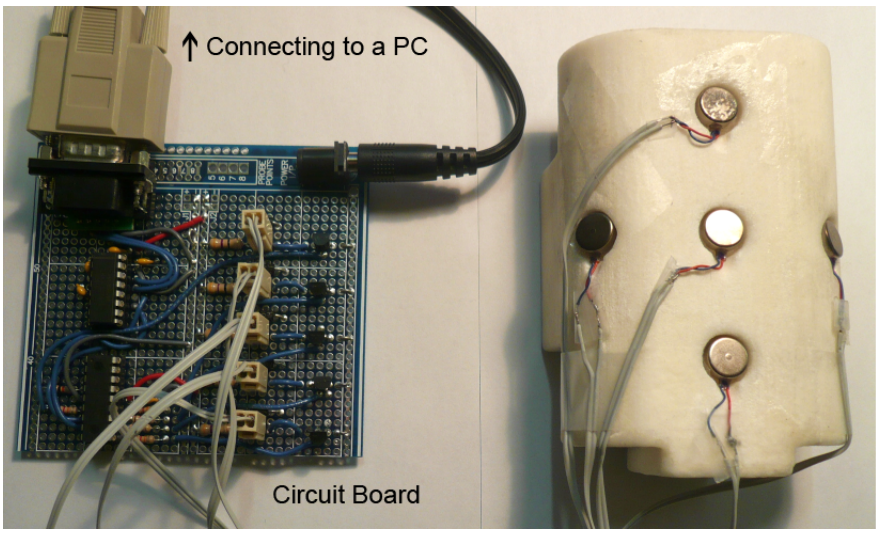

Figure 3.1. The first hardware prototype. The left side of this figure shows the circuit board, and the right side shows the backside of a mobile touch-screen device in which five vibration motors are embedded.

Figure 3.1 shows my SemFeel hardware prototype. I attach five vibration motors (Samsung Disk Coin-Type Vibration Motor APB108) on the backside of a mobile touch-screen device. The locations of these vibration motors were determined so that the hardware can produce basic spatial and spatio-temporal vibration patterns (*e.g.*, vibration flowing from the top of the device to the bottom). These motors are separated from other motors at least as far as 2 cm. This separation is sufficiently large for the user to distinguish different locations of vibration.

The vibration motors are connected to a circuit board, which contains two Integrated Circuit modules (PIC16F628 and MAX232A) for accepting signals from the computer and controlling each motor. The circuit board then communicates with a computer through a serial port connection. The circuit diagram of this hardware prototype is available in Appendix A (see Figure A.1).

In the study by Sahami *et al.* (2008), their participants had difficulty distinguishing some of the spatial vibration patterns produced by multiple vibration motors embedded along the left and

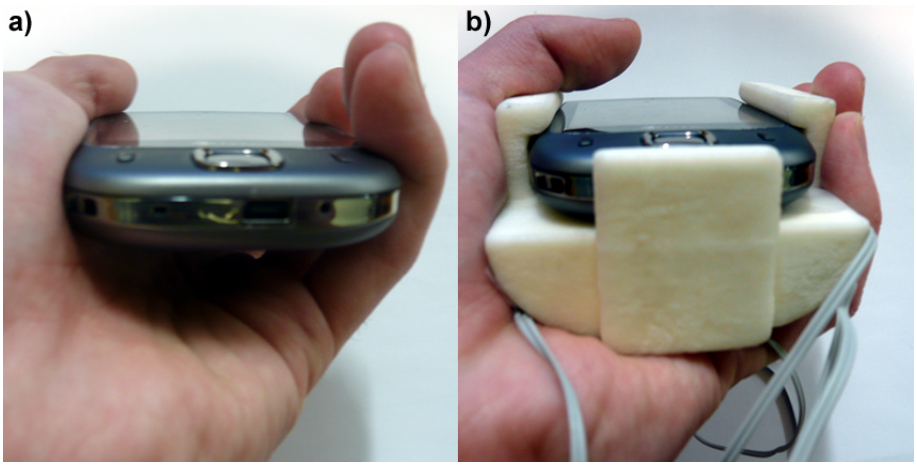

Figure 3.2. The special sleeve for my spatial tactile feedback hardware. a) Without the sleeve, a small gap between a user's palm and the device exists. b) The sleeve fills the gap and places vibration motors as close to the palm and fingers as possible.

right edges of a cellphone. From my informal observations of how users hold mobile devices, I noticed that there is typically a small gap between a person's palm and the device as shown in Figure 3.2a. This gap potentially could be one of the causes for the difficulties experienced by the users in the study by Sahami *et al.* (2008). As a result, when I manufactured a special sleeve that goes under a touch-screen device, I made it curve to fit the shape of the user's hand when it holds the device (Figure 3.2b). This sleeve allowed us to embed the vibration motors in the backside of the device while placing the motors as close to the palm and fingers as possible.

Each motor is controlled by a Pulse-Width Modulated (PWM) signal sent from PIC16F628. The signal is modulated at 10 kHz to eliminate audible noise. This hardware can change the intensity of vibration by changing the duty cycle of the PWM signal.

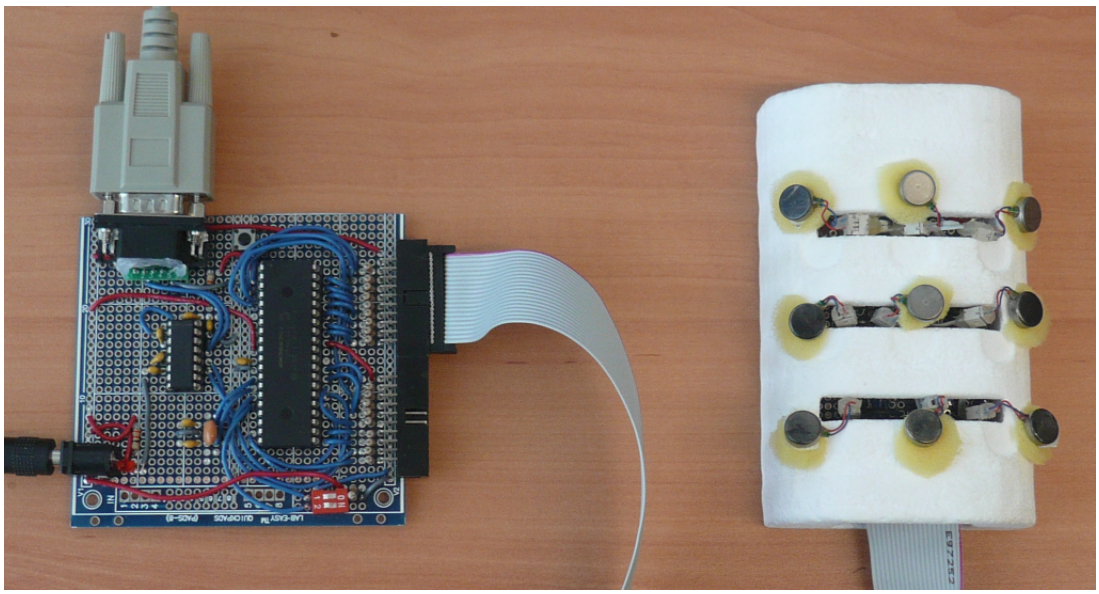

Figure 3.3. The second hardware prototype. The left side of this figure shows the revised circuit board, and the right side shows the backside of a mobile touch-screen device in which nine vibration motors are embedded.

3.3 Second Prototype

I further revised my hardware prototype to be able to accommodate nine vibration motors on the backside of a mobile touch-screen devices. This hardware allowed me to explore the effects of a finer output resolution of spatial tactile feedback in the interfaces described in Chapter 5 and 6. Figure 3.3 shows the second hardware prototype. I used the same vibration motors (Samsung Disk Coin-Type Vibration Motor APB108) as the first prototype. PIC16F877 in the circuit board controls these vibration motors based on commands sent from a computer through a serial communication. The circuit diagram of this hardware prototype is available in Appendix A (see Figure A.2).

Both prototypes were used in this dissertation to examine how spatial tactile feedback can support interactions on mobile touch-screen devices. I will explain which prototype was used in the later chapters.

Chapter 4

Eyes-free Interaction

As I discussed in Chapter 1, visually-demanding interfaces found in current user interfaces of a mobile touch-screen device degrade its user experience. The lack of tactile feedback on a flat touch screen hampers the user's ability to perceive the object that she touches on the screen when not looking at it. Thus, the user must view the screen regardless of how long or short her interaction might be. This produces a significant challenge with using mobile touch-screen devices because it is often hard for the user to devote visual attention to the devices, particularly in a mobile setting (Oulasvirta *et al.*, 2005).

Auditory feedback is one way to convey semantic information to the user about the object she is touching on the screen (Zhao *et al.*, 2007). However, auditory feedback is not always an appropriate output modality for mobile devices because the user may want to interact with applications in silence.

I argue that tactile feedback is a good alternative to mitigate some visual demand issues. Previous work focused on using basic patterns of tactile feedback as an enhancement of interactions in a sighted setting (Fukumoto & Sugimura, 2001; Hoggan *et al.*, 2007, 2008a; Poupyrev *et al.*, 2002a). Vibrotactile feedback systems are found to be useful for a variety of

tasks, such as item selection in a linear list (Poupyrev *et al.*, 2002a), and text entry (Hoggan *et al.*, 2008a).

Although vibrotactile basic feedback patterns used in the prior work above can help the user perceive whether she is touching any object or not, it is not necessarily sufficient to inform the user what object she is touching. Different vibration patterns in the intensity or rhythm parameter can be used to convey some semantic information (Brewster & Brown, 2004). But, additional ways to convey richer data that can be perceived and understood easily over the tactile channel remains under-explored.

I, therefore, looked into a design of eyes-free interaction with spatial tactile feedback hardware. Specific spatial vibration patterns can be used to offer detailed information, but it is not clear how accurately users can distinguish these patterns and how well spatial tactile feedback can support eyes-free interaction. I explored these questions with my SemFeel system.

4.1 SemFeel Concept

My eyes-free interface design for mobile touch-screen devices is called SemFeel, which was named after *Sematic Feeling*. It offers the user semantic information about the object which she touches on the screen as well as its presence. Figure 4.1a shows the SemFeel system concept, which has five vibration motors embedded in different locations in the back of a mobile touch-screen device, specifically the top, bottom, right, left, and center of the device. This placement allows the system to produce single-point vibration in specific locations as well as a "flow" of vibration (*e.g.*, a vibration moving from the top of the device to the bottom). These patterns can be designed to be easy for the user to perceive and to associate with a specific meaning based on the current application context. For example, in a music player application as shown in Figure 4.1b, when the user touches the "previous track" button, she feels vibration flowing from the right of the device to the left in the palm and fingers of her hand.

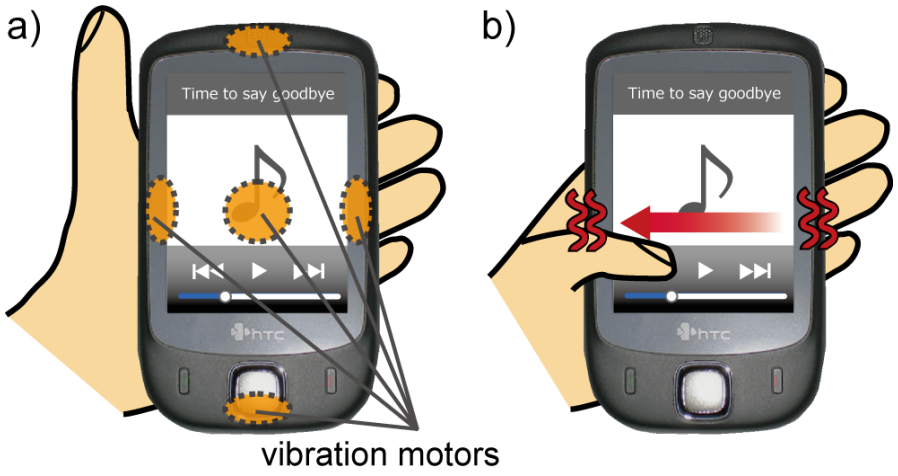

Figure 4.1. The SemFeel system concept: a) Multiple vibration motors are embedded in the backside of a mobile touch-screen device; b) The system generates vibration from right to left as feedback in response to the user's contact on "the previous track" button.

4.2 Related Work

Tactile feedback has been recognized commonly as an important user interface feature for touch-screen devices as I discussed in Section 2.4. My literature review showed that integration of spatial tactile feedback into a mobile touch-screen device and the effects of spatial and spatio-temporal patterns in such a device have not been well explored. However, spatial and spatio-temporal patterns have been used in wearable computers, and I discuss these systems here.

ActiveBelt is a wearable device that a user can wear around her waist (Tsukada & Yasumura, 2004). Eight vibration motors are evenly spaced in this belt-like device. Combined with a GPS module, the device can inform the user the direction to the destination by activating the vibration motor corresponding to that direction (*e.g.*, vibration from the right side of the device means that the destination is to the right). Their user study indicates that although participants could distinguish locations of vibration fairly accurately, they often mis-recognized or did not notice the change of vibration locations with a short duration (500 msec) while they were walking.

Heuten *et al.* (2008) examined how effectively a belt-like tactile feedback device similar to ActiveBelt (Tsukada & Yasumura, 2004) can support navigation in a physical space. Their experiment revealed that participants could follow the routes indicated by the device through the tactile channel within 15 m deviation in most cases.

Pielot *et al.* (2008) further investigated how an arbitrary angle can be presented by a belt-like device in which six vibration motors are evenly spaced. Their system changes the intensity of vibration proportion to the angle to be presented. For example, 40 degree can be presented by using two motors located at 0 and 60 degrees, and setting the 33% intensity on the 0-degree motor and 66% intensity on the 60-degree motor. Their user study found that participants could recognize angles at accuracy of approximately 20 degree.

Pielot *et al.* (2010) also studied the effects of using it to provide tactile feedback about where other team players are located in a 3D multiplayer game. In their device, the location of the vibration indicates the direction towards another player. They tested the intensity, duration, and rhythm parameter to encode the distance to that player. Their user study found that these three parameters were equally useful to convey the distance information.

Presenting directional information is a straightforward application of devices using multiple vibration motors. But they also can be used to offer other types of information or enhance visual stimuli the user is receiving. Wong *et al.* (2010) examined how spoken sounds can be mapped to the spatial parameter. Six equally spaced voice coil motors were attached to the user's left forearm. The formant transitions were mapped to the six discrete vibration points on the forearm. For example, if the formants decrease from initial-consonant to medial-vowel, this variation would be presented as a movement of vibrations towards the wrist. Their user study showed that their coding method was effective for encoding fricatives and affricated consonants, but needs improvements for plosive consonants.

Lee & Starner (2010) built a wearable device like a wrist band which contains three vibration motors aligned in a triangle. They examined the distinguishability of four parameters: 2 levels of intensity, 2 different rhythms, 2 different directions (*i.e.*, clockwise or counter-clockwise), and 3 different spatial points (starting point of vibration). Thus, they had 24 patterns in total. Their study found that intensity is the most difficult parameter to distinguish and the rhythm parameter is the easiest. They also conducted another study to understand users' abilities to perceive three different incoming alerts while performing visual search tasks. They used the three spatial patterns, and revealed that participants could respond to the alerts presented by their wearable device as quickly as they could do when they did not perform visual search tasks.

Israr & Poupyrev (2011) developed an algorithm, called Haptic Brush, for producing various spatial and spatio-temporal patterns with multiple vibration actuators. Their system uses a vibration motor array (12 motors in 4×3 alignment) attached to the back of a chair. Their system can produce various spatio-temporal patterns depending on the user's play of computer games. However, the effects of these spatio-temporal patterns in playing computer games were not tested in their work.

As seen in this section, spatial and spatio-temporal patterns have demonstrated effectiveness in wearable systems. Thus, their adaptation to mobile touch-screen devices could demonstrate benefits, particularly for mitigating visual demand issues. More specifically, I tried to answer the following research questions through an eyes-free interface design on mobile touch-screen devices using spatial tactile feedback.

H1 *Some vibration patterns are more accurate to recognize than others.*

H2 *Some vibration patterns are faster to perceive and react to than others.*

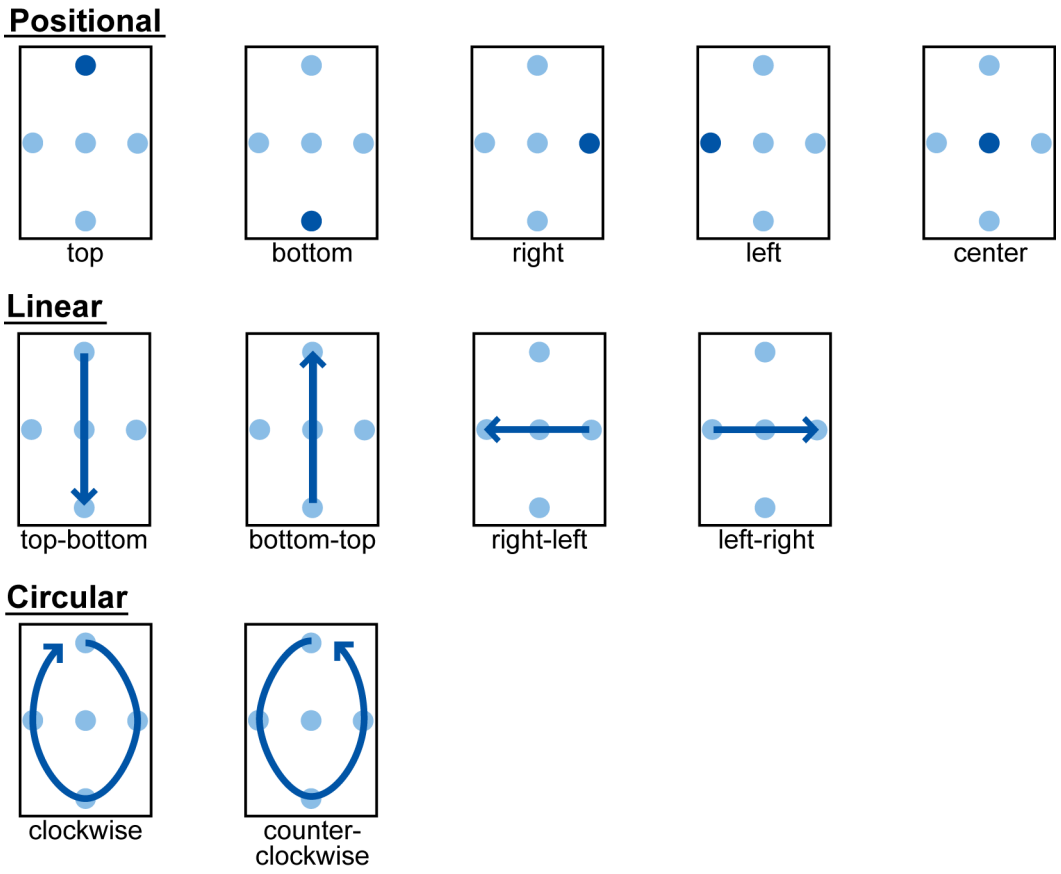

Figure 4.2. Eleven vibration patterns studied. The first hardware prototype was used (Figure 3.1).

4.3 SemFeel System

Hardware Configuration

SemFeel uses my first hardware prototype as shown in Figure 3.1. I conducted a pilot study to determine the range of vibration strength used in SemFeel. I found that more than 10% difference in the duty cycle is necessary for my pilot study participants to clearly distinguish between vibrations at two different strength levels, and weak vibration is harder to distinguish than strong vibration.

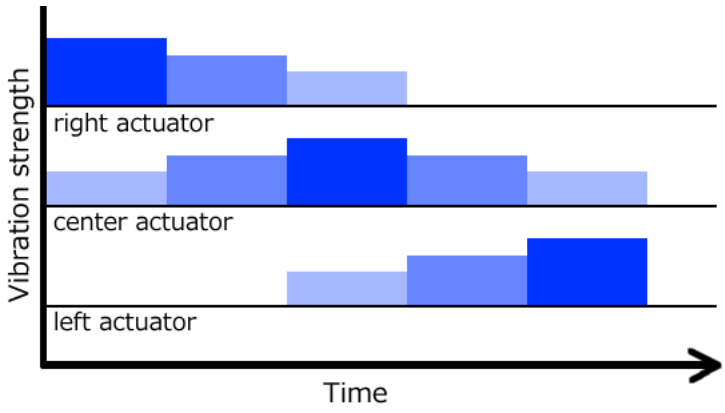

Figure 4.3. The smoothing of the vibration strength for the right-left vibration. The height of each square indicates the level of the vibration strength. The duration for each square is 200 [msec].

I designed my prototype to operate with four preset levels of vibration strength: 0% (completely off), 60%, 80% and 100% of the duty cycle. I set the smallest temporal duration to 50 msec because I found that my hardware often could not activate the vibration motors completely in a shorter duration through my pilot experiment. I also set the maximum duration of the vibration to 1 sec because longer durations are not practical.

Vibration Patterns

Figure 4.2 presents the eleven vibration patterns that I designed for the current SemFeel prototype. There are three types of patterns: positional (top, bottom, right, left, and center), linear (top-bottom, bottom-top, right-left, and left-right), and circular (clockwise and counter-clockwise). For the linear and circular vibration patterns, different levels of vibration strength are used to produce a smoother transition of the vibration. Figure 4.3 shows how the vibration strength is controlled in the right-left vibration pattern. Burtt (1917) developed this approach, and it is often used to produce sensation as if a vibration motor is moving smoothly in one direction. My second pilot study shows that participants preferred vibration with this smoothing over vibration without smoothing.

In the experiments described later, I used 100% strength for every vibration pattern except for the smoothing purpose. This setting allowed me to focus on evaluating how accurately users can distinguish vibration patterns generated by multiple vibration motors attached to different locations rather than the strength or rhythm of the vibration, which have been studied previously (Brewster & Brown, 2004). However, future work should examine how accurately users can distinguish spatial and spatio-temporal vibration patterns with different intensity levels, for example.

4.4 Experiment 1: Distinguishability of Patterns

I conducted a user study to examine how accurately users can distinguish different spatial and spatio-temporal vibration patterns, more specifically the patterns presented in Figure 4.2.

Tasks and Stimuli

In this experiment, I asked participants to determine which of the eleven vibration patterns shown in Figure 4.2 was being generated by the system at a time. Each time, the pattern was generated only once. After the system generated a pattern, a mouse cursor would appear in the small blue square on a computer screen (Figure 4.4). The participants would then move the cursor to the diagram representing the vibration pattern they thought the system had generated, and would click on the diagram. Each diagram was a square (70×70 pixels) placed the same distance (200 pixels) away from the initial cursor position. A dialogue would appear to show whether the response was correct. If the response was wrong, the correct answer also would be provided. During the experiment, the participants were asked to perform the task as quickly and accurately as possible.

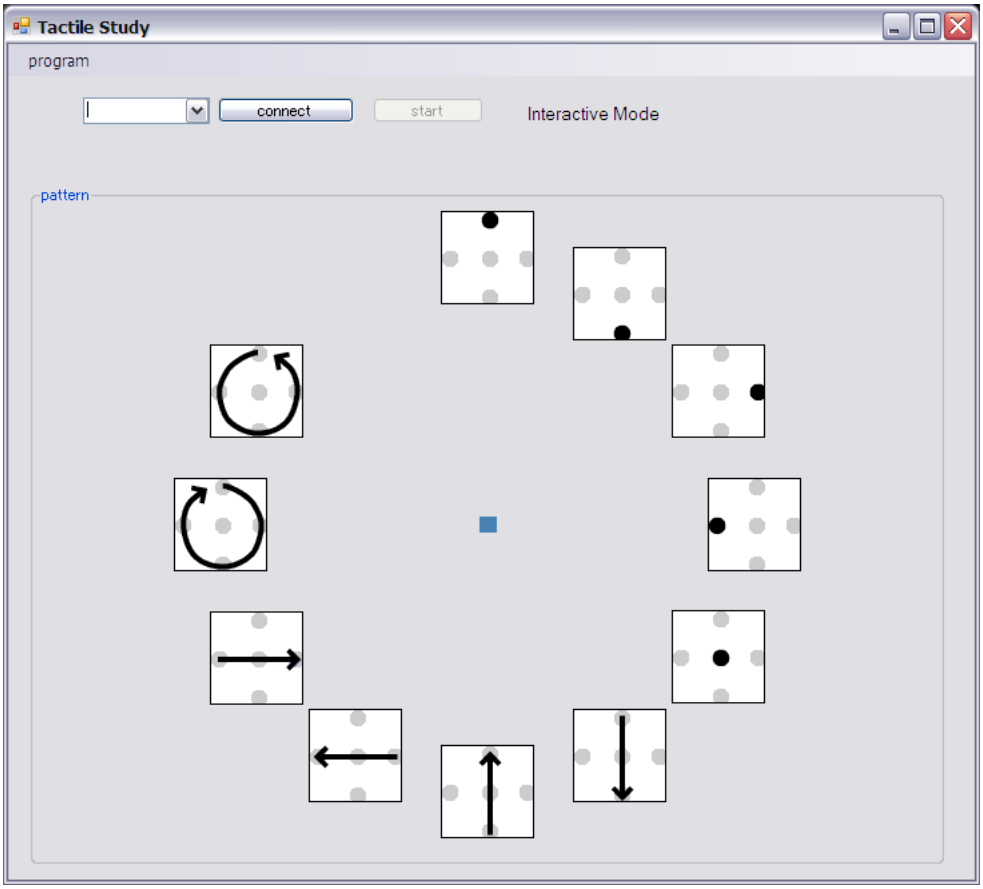

Figure 4.4. The screenshot of the software used in Experiment 1. Participants were asked to click one of the figures illustrating vibration patterns after they received feedback on the mobile device holding their hand.

Variables

The independent variable in this experiment was *Pattern* (five positional, four linear, and two circular patterns). Its presentation order of *Pattern* was randomized. Each vibration pattern was repeatedly presented three times in one block, and the experiment contained four blocks. Therefore, there were 11 (*Pattern*) × 3 (repetition) × 4 (blocks) = 132 trials per participant. I measured the reaction time as how long the participants took to click one of the diagrams shown in Figure 4.4 after a vibration pattern was presented completely. I also recorded the given vibration pattern and participants' response to measure the number of errors.

Apparatus

I used the same prototype and hardware configuration described in Section 4.3 in this experiment. A Windows Mobile 6 device (HTC Touch) was embedded in the custom sleeve. The application shown in Figure 4.4 was written in C# and ran on a Windows XP computer. The computer was connected to the circuit board. In this experiment, the durations of all the patterns were set to 1 sec.

Procedure

Participants were given the explanation of the system and instructed to hold the prototype mobile device with their non-dominant hand and to use a mouse with their dominant hand to interact with the application on the computer. They were then asked to complete a practice session in which the same tasks were used as the test sessions. They could continue to practice until they felt comfortable with the tasks and system. On the average, the participants practiced for about five minutes. After each block, the participants were allowed to take a short break. In total, the entire experiment took about 45 minutes.

Participants

Ten people (five male and five female, aged 18 to 50) with different professional backgrounds (students, law careers, business consultants, a physician, and a programmer) were recruited for this experiment. One male and one female were left-handed, and the others were right-handed. Three participants regularly used mobile touch-screen devices. The participants were compensated for their time and effort with $20 cash after the study.

		User response											
		T	B	R	L	C	TB	BT	RL	LR	CW	CCW	Recall (%)
Stimulus	T	112		4	1	1		2					93.3
	B		103	8		9							85.8
	R			107		13							89.2
	L	2			112	6							93.3
	C	1	2	3	1	109	2	2					90.8
	TB						110		2	3	2	3	91.7
	BT	4						110	1	3	1	1	91.7
	RL					1	3	2	107	1	6		89.2
	LR					2	1	1	2	106	2	6	88.3
	CW						3		1		100	16	83.3
	CCW										41	79	65.8
Precision (%)		94.1	98.1	87.7	98.2	77.3	92.4	94.0	94.7	93.8	65.8	75.2	

Table 4.1. The confusion matrix for the recognition accuracies of the eleven vibration patterns tested in Experiment 1. The numbers represent the numbers of occurrences of the user responses.

4.5 Experiment 1 Results and Discussions

I mainly examined the two hypotheses through the results I gained.

H1 *Some vibration patterns are more accurate to recognize than others.*

H2 *Some vibration patterns are faster to perceive and react to than others.*

Table 4.1 shows the confusion matrix for the eleven patterns. The error count data did not show a tendency to be a normal distribution: Shapiro-Wilk tests showed significant results on the error rates of the nine vibration patterns (all except the bottom and counter-clockwise patterns; $p < .05$). Therefore, I used Friedman non-parametric tests for the analysis. A Friedman test did not show a statistically significant differences in the error counts (recall) across *Patterns* ($\chi^2_{(10)} = 13.8$, $p = .18$). I then examined the error counts for each pattern category (positional, linear, and circular). However, Friedman tests for each category did not

show any significant difference at the 95% confidence level. Nevertheless, as seen in Table 4.1, the counter-clockwise pattern tended to be more inaccurate than the other patterns, and was mostly confused with the clockwise pattern. Therefore, H1 (*Some vibration patterns are more accurate to recognize than others.*) is not well supported although the accuracy drop in the counter-clockwise pattern indicates that improvements are necessary if both the clockwise and counter-clockwise pattern are used.

Figure 4.5 shows the mean reaction time for *Pattern*. I ran a one-way repeated-measure ANOVA test for performance time against the three pattern categories (positional, linear, and circular). The mean performance time for these categories was: 2.19 sec ($SD = 0.39$), 2.43 sec ($SD = 0.46$), and 2.95 sec ($SD = 0.77$) for positional, linear, and circular patterns, respectively. A Mauchly's sphericity test did not show a significant result ($p = .27$); thus, the results of an ANOVA test can be interpreted without a correction. A repeated-measure ANOVA test shows the existence of statistically significant differences ($F_{(2,18)} = 12.4$, $p < .001$, the effect size $\eta_p^2 = 0.58$). To accommodate the unbalanced sample sizes across the pattern categories, a Tukey HSD multiple comparison was used in the post-hoc test. It showed that the reaction time for the positional patterns was significantly faster than those for the linear patterns and circular patterns ($p < .05$ for both), and the reaction time for the linear patterns was significantly faster than for the circular patterns ($p < .05$).

I also tested the reaction time across the eleven patterns. A repeated-measure ANOVA test showed a significant difference ($F_{(10,90)} = 6.13$, $p < .001$, the effect size $\eta_p^2 = 0.40$). A post-hoc pairwise comparison with paired t tests using the Holm correction found a significant difference between the center and counter clockwise pattern ($p = .035$), but did not find any other significant difference at the 95% confidence level.

These results suggest that H2 (*Some vibration patterns are faster to perceive and react to than others.*) is not supported clearly. However, the positional patterns tended to be faster to recognize than the linear and circular patterns.

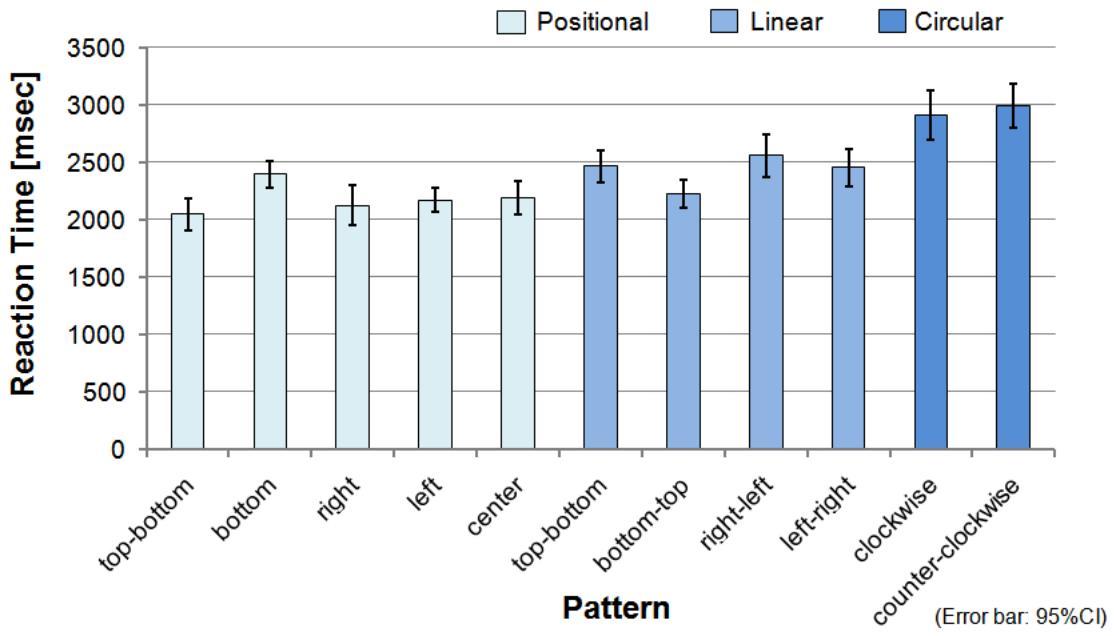

Figure 4.5. The reaction time across the vibration patterns in Experiment 1.

Figure 4.6 shows the reaction time across the four blocks. A Mauchly's sphericity test showed a significant result ($p < .05$, Greenhouse-Geisser's $\epsilon = 0.51$). A one-way repeated-measure ANOVA test with the Greenhouse-Geisser correction indicates the existence of statistically significant differences in the reaction time against the blocks ($F_{(1.51,13.7)} = 11.3$, $p < .001$). The post-hoc Tukey HSD multiple comparison discovered significant differences between the first block and the other blocks ($p < .05$). I did not find any statistical difference in the error rate across the four blocks with a Friedman test ($\chi^2_{(3)} = 5.90$, $p = .12$).

The results of this first experiment show that the participants could distinguish the eleven patterns except for counter-clockwise at 83.3–93.3 % accuracy (89.6 % on average) in spite of a short amount of practice. Although the results indicate that modification on circular patterns is necessary (*e.g.*, using only one of the circular patterns, or starting one of the circular patterns at a different location) to avoid user confusion, this is a large improvement from the system studied by Sahami *et al.* (2008). One possible reason is that the duration of the vibration in this

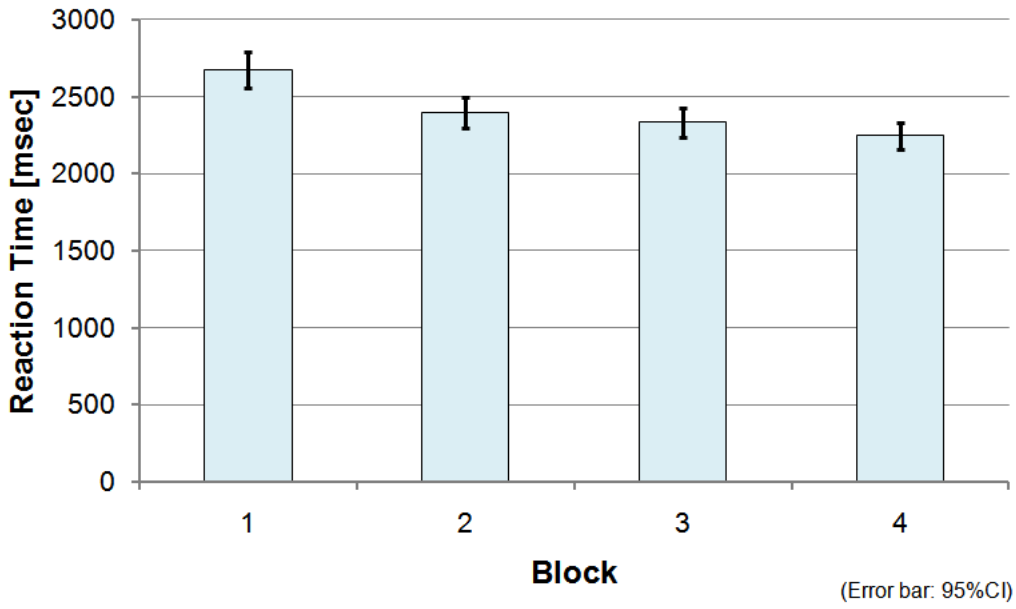

Figure 4.6. The reaction time across the four blocks in Experiment 1.

experiment was 1 sec whereas the patterns they used were between 300 msec and 900 msec. Additionally, my custom sleeve to fill the gap that normally exists between the user's hand and device might have helped the participants sense the vibrations better. Further investigation on exactly why my prototype achieves higher accuracy is needed, but the results indicate that vibration generated by multiple vibration motors can be accurately distinguishable by users.

I also found a slight improvement in the reaction time across the blocks. However, this reaction time included the selection time of the icons in the software running on the computer (Figure 4.4) as well as the actual response time to the stimulus. Thus, it is difficult to draw a conclusion that there is a learning effect on the response time to the vibration patterns tested in this experiment, and further investigation is necessary.

4.6 Experiment 2: User Performance on Eyes-free Number Entry Task

The first experiment shows that users can distinguish ten of the vibration patterns that my current SemFeel prototype can generate (five positional, four linear, and a clockwise circular vibration patterns) at 83.3–91.7% accuracy. I designed the second experiment to examine user performance in a realistic application with the SemFeel technology. In particular, I wanted to compare the accuracy of user input when using the SemFeel prototype against user interfaces that offer no tactile feedback or tactile feedback using only a single vibration source.

In this experiment, I asked participants to perform a number entering task. I chose this input task with the numeric keyboard because this is a commonly performed action on cellphones. Furthermore, this interaction can be extended to other applications (*e.g.*, text entry with the multitap method). In this experiment, I focused on an eyes-free setting where a clear difference could be expected between SemFeel and the reference systems.

Tasks and Stimuli

The system first presented participants with a 4-digit number in blue font on the computer screen (Figure 4.7a). The participants then typed that number on a mobile touch-screen device using the numeric keyboard shown in Figure 4.7b. Each key was a square (9.2 $\times$ 9.2 cm). This size was chosen to allow the participants to interact comfortably with their thumb based on the findings reported by Parhi *et al.* (2006). The participants could commit the typing by releasing the thumb from the screen, and then the entered number would appear on the computer screen. The character 'X' would be shown when the participant released the thumb outside any of the keys.

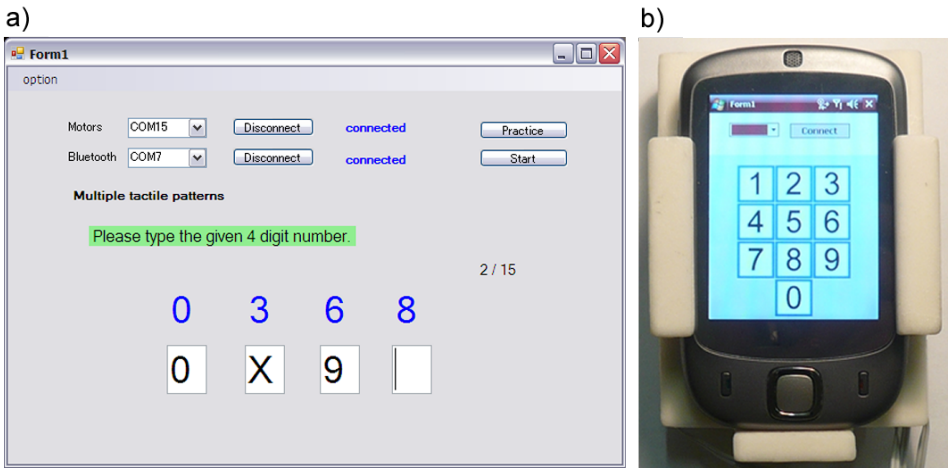

Figure 4.7. The system used in Experiment 2. a) A screenshot of the application running on a computer. This application presented participants a 4-digit number. The character "X" would appear when participants released the thumb outside the numeric keyboard. b) The numeric keyboard on the prototype mobile device.

There were three tactile feedback conditions studied in this experiment: none (*No Tactile*), tactile feedback provided through a single vibration motor (*Single Tactile*), and tactile feedback provided through the SemFeel technology (or tactile feedback provided through multiple vibration motors, which will be referred to as *Multiple Tactile*). Although the main objective of this experiment was a comparison between *Single Tactile* and *Multiple Tactile*, I decided to include that *No Tactile* condition because it allowed me to estimate how visually-demanding this task would be by examining the error rate.

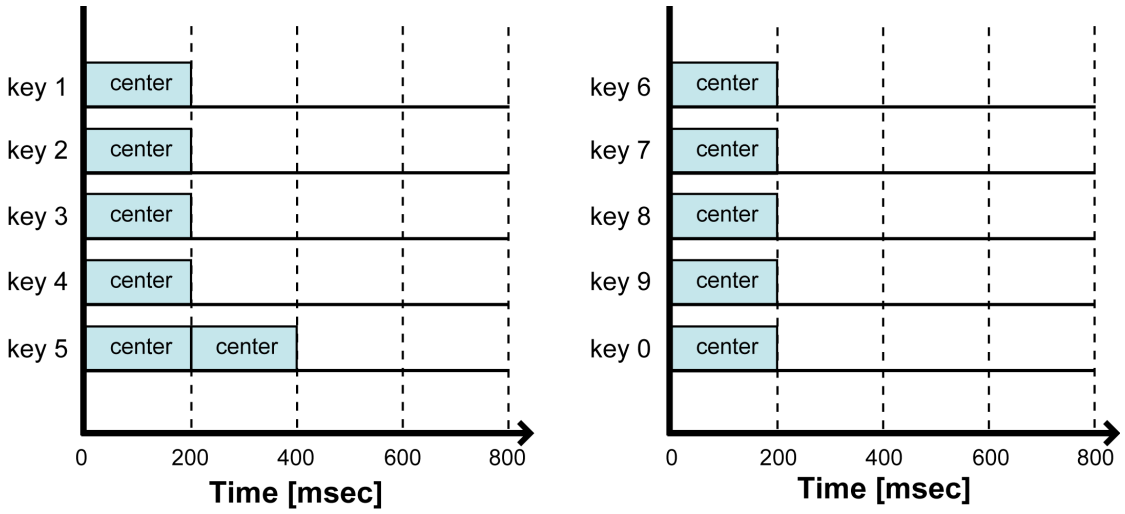

Figure 4.8. The mapping of the vibration patterns for the numeric keyboard used in the *Single Tactile* condition. Only the center vibration motor is used in this condition, and the '5' key has vibration with a different duration.

In the *Single Tactile* condition, the center vibration motor was used to provide tactile feedback when the participants were touching any of the keys on the keyboard (Figure 4.8). For the number '5' key, the system turned the center vibration motor on for 400 msec and off for another 400 msec. For the other keys, the system turned on the center vibration motor for 200 msec and off for 600 msec. These patterns repeated while the participant's finger continued to touch the button. This implementation was designed to emulate a physical numeric keyboard (*i.e.*, every key except for the key for '5' has the same texture, and the '5' key has a slightly different texture to indicate the home position).

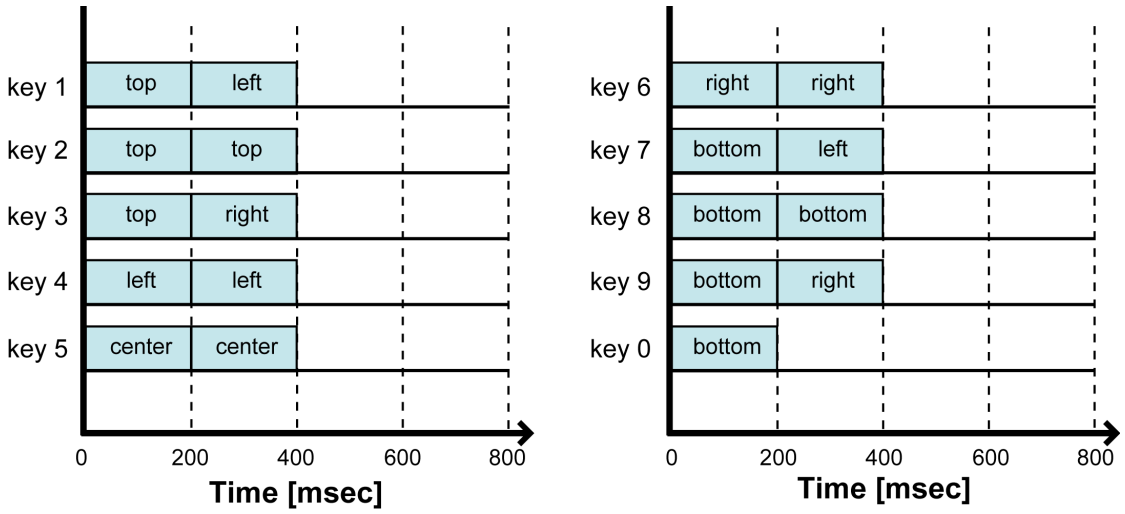

Figure 4.9. The mapping of the vibration patterns for the numeric keyboard used in the *Multiple Tactile* condition. Please note that the assignment of the vibration patterns is based on the spatial relationship of the keys (*e.g.*, the combination of top and left vibration is assigned to key 1).

In the *Multiple Tactile* condition, the vibration patterns were designed with simple combinations of the positional patterns, and assigned to match the spatial relationship as shown in Figure 4.9. Each pattern lasted 800 msec, consisting of two blocks of 200 msec vibration generated by one of the vibration motors followed by 400 msec without any vibration. Only one of the motors was activated at a time. For instance, the system turns the left vibration motor on for 400 msec and off for another 400 msec when the user touched the '4' key. It activated the top vibration motor for 200 msec first, and then the right vibration motor for the next 200 msec, and stopped all the motors for the next 400 msec when the user touched the '1' key. Similar to the *Single Tactile* condition, these vibration patterns were repeated while the participant's finger continued to touch a button. I confirmed that this design of the vibration patterns used was easy to perceive and distinguish through an informal study.

Variables

The independent variable that I controlled in this experiment was *Feedback* (*No Tactile*, *Single Tactile*, and *Multiple Tactile*). The order of the presentation of *Feedback* was counter-balanced across the participants. In each block, a 4-digit number was randomly generated, but it always satisfied the two following conditions: 1) each digit was different from the others, and 2) the frequency of each number was equal within the block. The experiment contained two blocks for each Feedback condition. Therefore, there were 3 (*Feedback*) × 2 (blocks) × 15 (trials) = 90 trials per participant.

I measured the performance time as the time from after a 4-digit number was shown to when the participants released their thumb from the screen to enter the fourth digit. The error rate was calculated for each block (*i.e.*, 100 * [the number of the wrong entries] / [the number of the digits in one block = 60]).

Apparatus

I used the same devices used in the first experiment. The applications used in this experiment were written in C#. All the touch events on the numeric keyboard were sent to the computer via Bluetooth. The application on the computer then sent a signal to the circuit board to generate the pattern corresponding to the key the participants were touching.

Procedure

Before the experiment, participants were given the explanation of the system, and instructed to hold the prototype mobile device with their dominant hand and to use the thumb of that hand to interact with it. This instruction was included because it is highly likely that the user would interact with a mobile device using only one hand in an eyes-free setting. However, due to the fairly large size and heavy weight of the prototype, I allowed the participants to support their

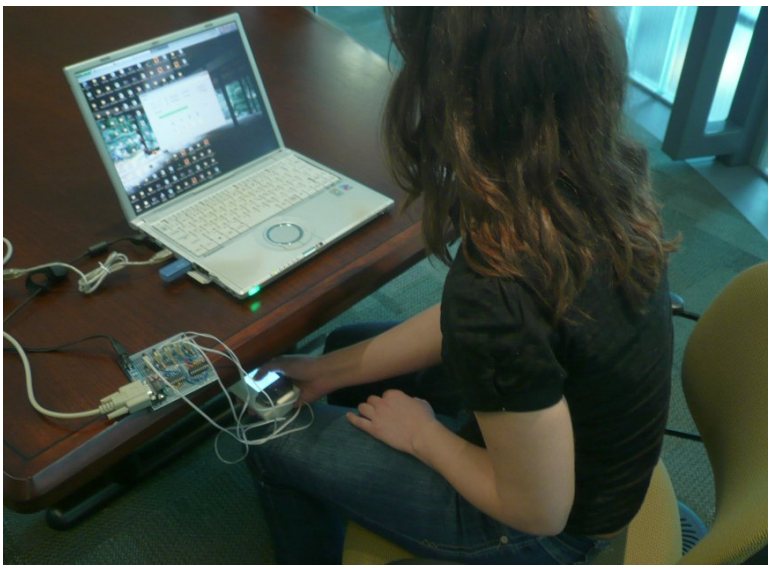

Figure 4.10. Experiment 2 setup. Participants were asked to type a 4-digit number presented in the screen of a laptop with the numeric keyboard on a mobile device. During the experiment, they were instructed to hold the mobile device under the table and were not allowed to view its screen.

dominant hand with the other hand. The participants were also instructed to touch the screen of the mobile device with their finger tip or nail due to the weak responsiveness of the resistive touch screen. During the experiment, I asked the participants to hold the mobile device under the table (as shown in Figure 4.10) so that they could not see its screen.

The participants were then asked to perform a practice set that used the same tasks as the test sessions to become comfortable with all the conditions at the beginning of the experiment. They could continue to practice until they felt comfortable with the tasks and system. On average, the participants practiced for about ten minutes. During the test sessions, the participants were asked to complete each trial as quickly and accurately as possible. After each block, participants were allowed to take a short break. In total, the entire experiment took about 45 minutes.

Participants

Twelve right-handed people (eight male and four female, aged 18 to 50) with a variety of backgrounds (university students, accountants, a health worker, a technician, a car dealer, a waiter, and an executive assistant) were recruited for this experiment. Eight of the participants regularly used mobile touch-screen devices. None of them participated in my first experiment. All the participants were compensated for their time and effort with $20 cash.

4.7 Experiment 2 Results and Discussions

I had the following main hypothesis for this second user study.

H3 *The Multiple Tactile condition is more accurate than the No Tactile and Single Tactile conditions.*

Figure 4.11 shows the mean error rates for the *Feedback* conditions: 34.7% ($SD = 15.8$, $Median = 35.0$); 22.3% ($SD = 13.7$, $Median = 19.2$); and 12.4 % ($SD = 0.10$, $Median = 0.10$) for *No Tactile*, *Single Tactile*, and *Multiple Tactile*, respectively. Similar to the first experiment, I used a Friedman test for the error count data and used Wilcoxon signed rank tests with the Bonferroni-Holm correction as a post-hoc non-parametric test. A Friedman test for the number of errors against *Feedback* reveals the existence of significant differences ($\chi^2_{(2)} = 32.0, p < .001$). A post-hoc pairwise comparison revealed that statistically significant differences exist between any of the two conditions ($p < .01$ for the differences between *No Tactile* and *Single Tactile* (the effect size $r = 0.53$), and between *Single Tactile* and *Multiple Tactile* ($r = 0.64$); and $p < .001$ for the difference between *No Tactile* and *Multiple Tactile* ($r = 0.98$)). Thus, H3 (*The Multiple Tactile condition is more accurate than the No Tactile and Single Tactile conditions.*) was supported.

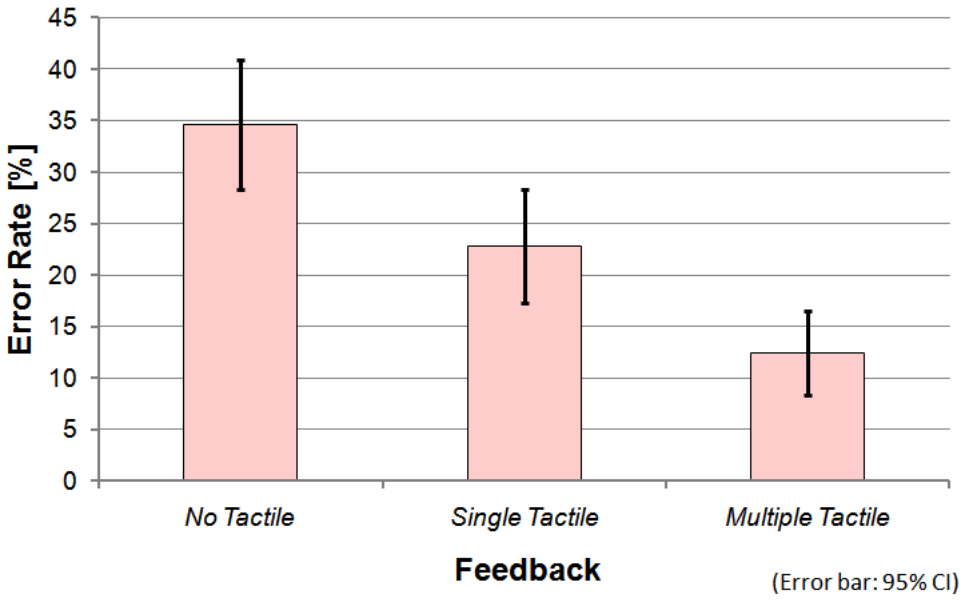

Figure 4.11. The typing error rate across the three feedback conditions in Experiment 2.

I then examined the performance time in each condition. Figure 4.12 shows the mean performance time of typing a 4-digit number for *Feedback*: 11.3 sec ($SD = 6.84$), 14.1 sec ($SD = 8.16$), 15.0 sec ($SD = 7.56$) for *No Tactile*, *Single Tactile*, and *Multiple Tactile*, respectively. A Mauchly's sphericity test did not show a significant result ($p = .75$); thus, the results of an ANOVA test can be interpreted without a correction. A one-way repeated-measure ANOVA test for performance time against *Feedback* indicates no significant difference ($F_{(2,22)} = 3.21$, $p = .06$, the effect size $\eta_p^2 = 0.23$).

I also examined the performance time along the error count in each trial. Figure 4.13 shows the performance time across the conditions and the error count in each trial. Note that some conditions had only a few samples, and direct comparisons shown in this graph may not be appropriate. However, the graph suggests that participants tended to spend more time when they encountered more errors, particularly in the *No Tactile* and *Single Tactile* conditions.

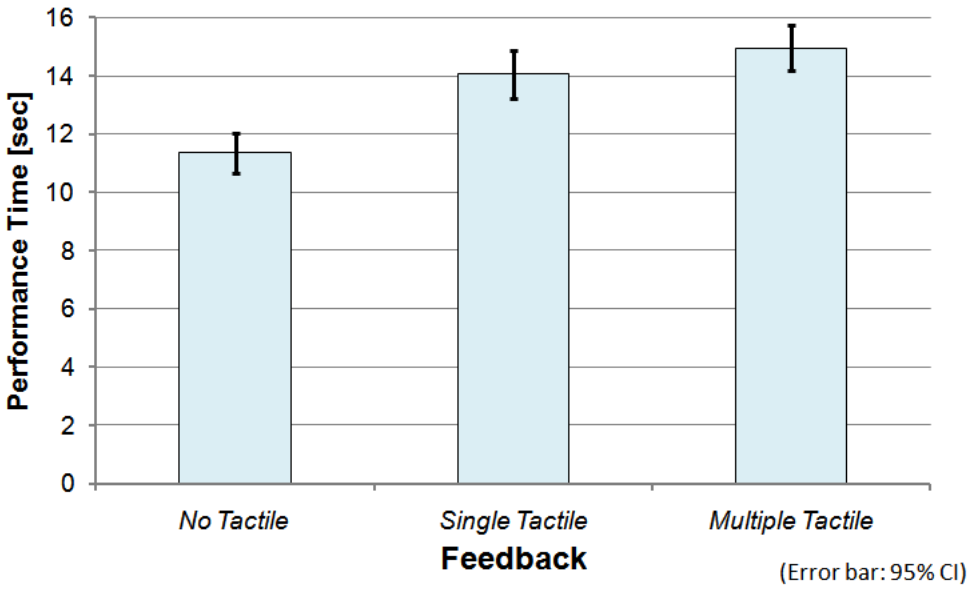

Figure 4.12. The average performance time of typing a 4-digit number across the three feedback conditions in Experiment 2.

The second experiment shows that SemFeel can support significantly more accurate eyes-free interactions on number entering tasks than a user interface without any tactile feedback and one with tactile feedback using a single vibration motor. Further investigation is necessary on how SemFeel could improve user performance in other applications or other situations (*e.g.*, while users are walking), but the results suggest that SemFeel has the potential to help the user accurately interact with mobile touch-screen devices without looking at the screen.

My second study also shows that the interfaces with tactile feedback was slower than the one without any tactile feedback. Participants often adjusted their contact point on the screen based on the tactile feedback to hit the right key. However, without tactile feedback, participants were unable to make such adjustments. This may be one reason that a difference in performance time between the *No Tactile* condition and the conditions with tactile feedback was observed. The results did not show a clear trade-off between performance time and accuracy although another user study is necessary to draw a conclusion on the relationships between them.

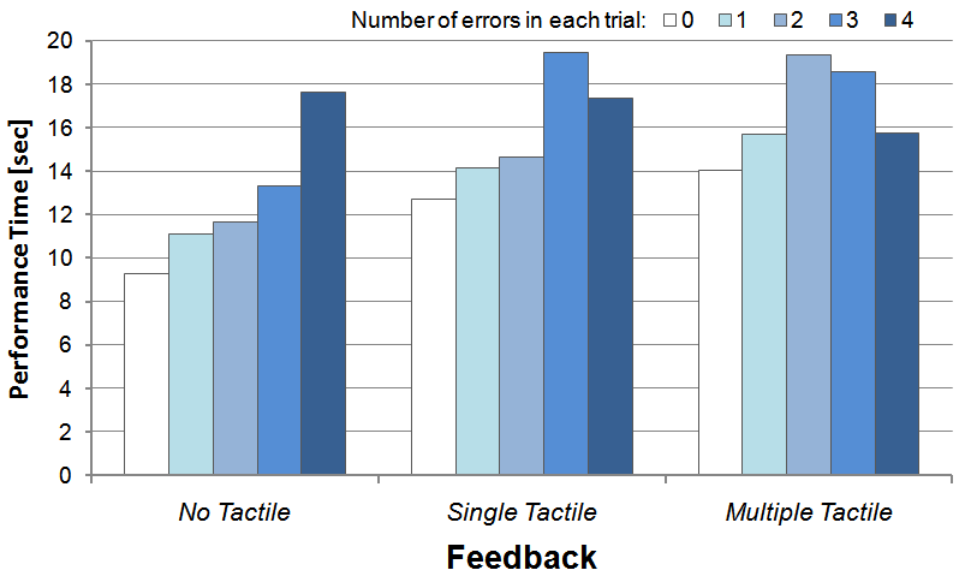

Figure 4.13. The average performance time of typing a 4-digit number across the feedback conditions and the number of errors in each trial. Note that some conditions included only a few samples. The graph suggests that participants tended to spend more time when they encountered more errors, particularly in the *No Tactile* and *Single Tactile* conditions.

The results from this experiment indicate the initial user performance on number entering tasks. By repeatedly exposed to the interface, the user would develop the muscle memory of the numeric keyboard, and the performance time and accuracy would be improved. An additional study with more sessions than this study is necessary to draw a conclusion on the effects of SemFeel for long-term use. However, SemFeel can support accurate eyes-free interaction even if the user has not developed the muscle memory completely.

There are several limitations in this user study. During the user study, participants could hear the noise caused by the activated vibration motor. This might have helped the participants know if they were touching a key. However, the noise did not vary by the vibration motors, and the observed difference on accuracy between the *Single-Tactile* and *Multiple-Tactile* is regarded as a result of using spatial tactile feedback. Additionally, the performance reported in this document

is for a number-entry task. The results may need careful interpretation or an additional user study when they are extended to other tasks (*e.g.*, text entry). The output was available visually to the participants during the experiment. This is one of the possible eyes-free conditions (the input space is not visible, and the typed letters are fully visible) discussed by Clawson *et al.* (2005). Further investigation is needed to understand the performance of SemFeel in different eyes-free conditions (*e.g.*, the typed letters are not available to the user).

4.8 Applications

There are other applications which also could leverage my SemFeel technology to support less visually-demanding or eyes-free interactions beyond the music player application (Figure 4.1) and a numeric keyboard that have already been discussed above. Figure 4.14a shows one possible design of an alphabetic keyboard which has three large keys. The left, center, and right keys activate the left, center, and right vibration motor, respectively. Each key contains nine alphabetical letters, and the user can enter one of those letters by touching the appropriate key and the makes a gesture to specify which letter she wants to input. For example, Figure 4.14a shows that a user is entering 'r'. In this case, she presses the center key first, and then moves the thumb towards the bottom-right direction. This keyboard can be less visually-demanding than a normal mini-qwerty keyboard because the tactile feedback tells the user which key she is touching, and does not require precise adjustment of the contact point to select the specific key for the desired letter. Thus, the user may be able to devote her visual attention to the text rather than the keyboard.

SemFeel can also allow the user to interact with a calendar application and access its content without looking at the screen. Li *et al.* (2008) previously demonstrated an audio-based eyes-free interaction to access a calendar application through a cellphone. The user can use their system even when they are engaging in a conversation over the phone. However, the auditory feedback from the system could be a distraction to the phone conversation. With SemFeel,

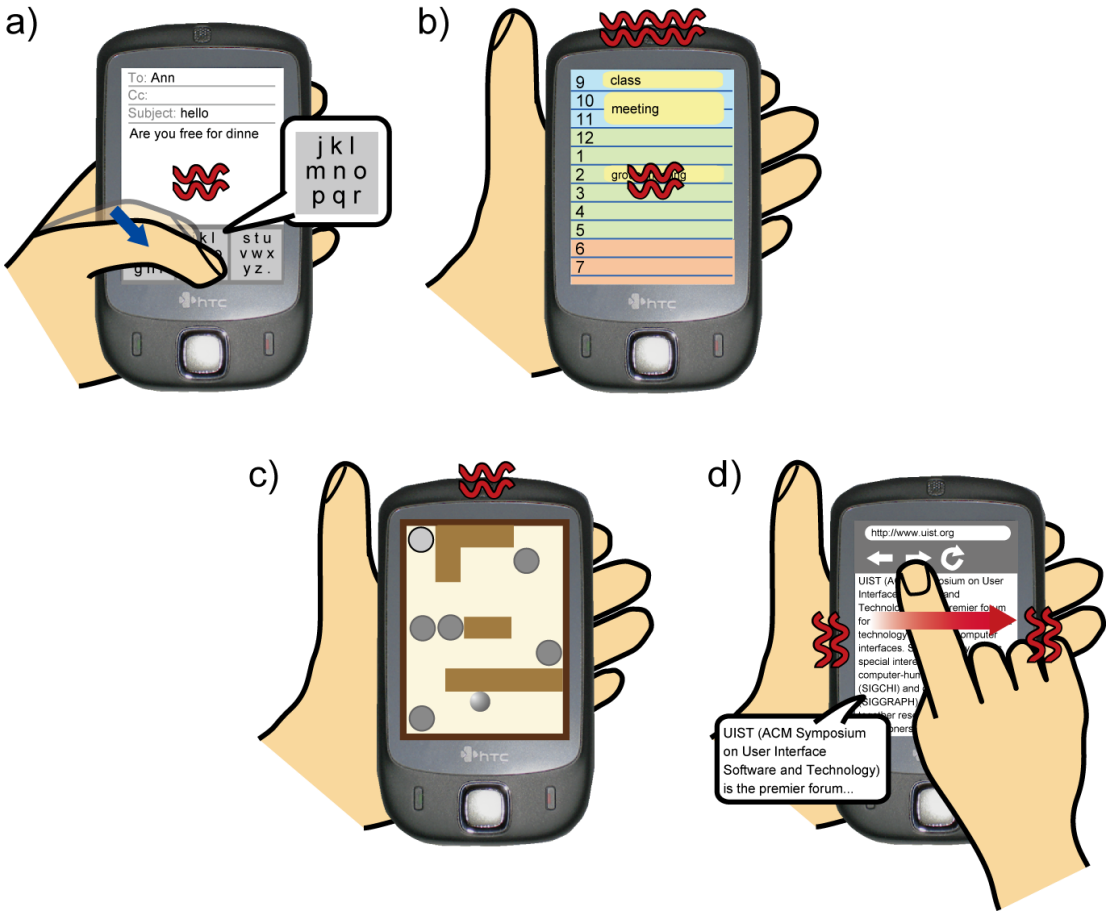

Figure 4.14. SemFeel applications: a) An alphabetic keyboard. The right, center and left vibration patterns are associated with three large keys which the user can touch and make a gesture to type a letter; b) A calendar application. The top, center, and bottom vibration motors are used for representing the morning, afternoon, and evening in a particular day, and the duration of the vibration generated by each vibration motor represents the availability of each time period (longer vibration means less available); c) A maze game. Users can interact by tilting the device, and when the ball hits the wall, the vibration is generated; and d) A web browser for people with visual impairment. Audio feedback is used for reading out the content of a webpage, and tactile feedback is used for providing information about the controls in the web browser application.

different time slots can be associated with vibration motors at different location. Figure 4.14b shows an example of mapping between time slots and vibration motors. For instance, the top, center, and bottom vibration motor are associated with the morning, afternoon, and evening in a particular day, respectively. The duration of the vibration can represent how busy the user is. Although the user will not be able to know the details of the schedule, SemFeel still can provide the user with a general idea of her availability.

SemFeel could be used to enhance the current user interfaces on mobile touch-screen devices. For example, SemFeel can be incorporated into game applications to provide a more entertaining user interface. Figure 4.14c shows an example game application, a maze. The goal in this game is to move the ball to a target by tilting the device. When the ball hits a wall, the system generates a vibration in the direction of the obstacle.

SemFeel also could improve the design of user interfaces for people with visual impairments. McGookin *et al.* (2008) studied accessibility issues on mobile touch-screen devices. One of their findings suggests that user interface designers should provide feedback for all the actions that have occurred on the screen. SemFeel could be used as an additional feedback channel for the user. Figure 4.14d shows a web browser application for people with visual impairments. The audio channel is used to read out the content of a Web page while tactile feedback is used to provide information about the control that the user touches on the screen. Auditory feedback can be used to provide information about the user's interaction, but it might be distracting because the read-out function has to be stopped or users would have two different kinds of information delivered through the audio channel at the same time.

4.9 Summary

Lack of tactile feedback can hinder the effective use of mobile touch-screen devices, especially when the user are unable to view the screen. I developed SemFeel, a tactile feedback technology for mobile touch-screen devices which provides the user with the semantic

information about the object she is touching. SemFeel encodes such information in spatial or spatio-temporal vibration patterns generated by multiple vibration motors embedded in the backside of the device. Through this SemFeel prototype, I wanted to answer the two following research questions as listed in Table 1.1.

RQ-1 *How fast and accurately can users recognize various spatial vibration patterns generated by multiple vibration motors?*

RQ-2 *How fast and accurately can users perform eyes-free interaction with the support of spatial tactile feedback?*

The results of my two user studies show that participants could distinguish ten vibration patterns, including linear patterns and a clockwise circular pattern, at 83.3–93.3% accuracy. They also revealed that SemFeel supports more accurate interactions in an eyes-free setting than systems that offer no tactile feedback or use only a single vibration motor. This chapter validates the effectiveness of spatial tactile feedback on eyes-free number entry interaction, and discusses other possible applications with spatial tactile feedback.

Chapter 5

Interface Design for Visually Impaired Users

The evaluation of SemFeel shows that spatial vibration patterns can be easily distinguishable and can support successful eyes-free interaction. This also suggests that the SemFeel system could be useful for visually impaired users. Although testing a similar system with visually impaired users would be interesting and has value, I believe that spatial tactile feedback would demonstrate its benefits strongly when it presents visual information which sighted users can easily glean but are not easily accessible to visually impaired users.

Acquiring knowledge about an environment can be challenging for visually impaired people. They often need support by sighted people to understand what is in the environment, where venues are positioned, and their spatial relationship with each other. They sometimes also need to physically navigate an environment multiple times before they can develop a familiarity with the environment. Although such activities can help visually impaired people increase their independence and confidence in their navigation, burdens associated with these activities often discourage them to explore and learn an environment (Yang *et al.*, 2011).

Braille maps (*a.k.a.* tactile maps) are an effective way for visually impaired people to develop an understanding of the spatial relationships between a set of places in addition to identifying directions between multiple points (Blasch *et al.*, 1997; Helal *et al.*, 2001). Prior work has shown that Braille maps can help visually impaired people prepare for their future trips to an area (Blasch *et al.*, 1997; Kane *et al.*, 2009). However, Braille maps use physical materials that are often costly to produce, may not always cover an area of interest with sufficient detail, and might not present updated information (Rowell & Ungar, 2005).

I explored a way to present geographical information to visually impaired users on a handheld device. My system makes use of two feedback channels—auditory (speech) feedback and tactile feedback—to present places of interests, provide detail information about those locations, and present routes to those points. With my map system, SpaceSense, I wanted to explore how these two feedback channels accessible to users with visual impairments can be integrated and how spatial tactile feedback can enhance the user's understanding of spatial relationships between multiple places.

5.1 SpaceSense Concept

Spatial tactile feedback is used to convey directional or spatial information in SpaceSense. For example, when the user browses route information, SpaceSense offers high-level information about the direction towards the destination (Figure 5.1). Furthermore, SpaceSense also provides directional information towards other locations that the user may know or like. In this manner, SpaceSpace can help visually impaired users maintain high-level spatial relationships between multiple locations.

SpaceSense also adapts the Review Spotlight system (Yatani *et al.*, 2011) for summarizing online user reviews about nearby locations. It reads out the most frequently-used adjective-noun word pairs extracted from reviews available online to offer an overview of what people often mention about a specific place. Through this summarized presentation of online reviews,

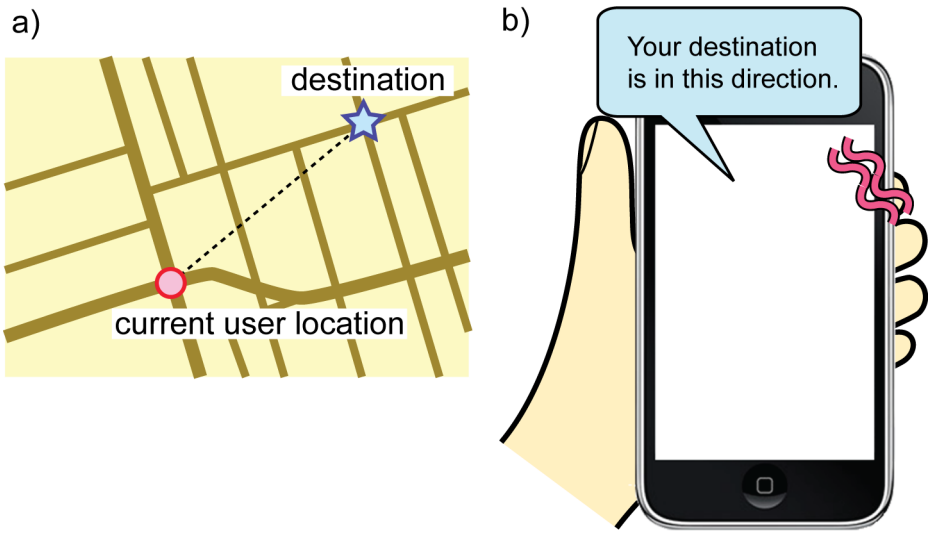

Figure 5.1. SpaceSense concept. a) An example of a map showing geographical information being browsed by the user (with the destination north-east from the simulated user location); b) Spatial tactile feedback conveys the high-level direction towards the destination.

the user can obtain information about a location before deciding whether she wants to further learn its directions and its spatial relationships to other locations.

5.2 Related Work

Understanding spatial relationships between multiple places is often difficult for people with visual impairments. Previous studies show that visually impaired people can learn spatial relationships through physical blocks or Braille maps (Thinus-Blanc & Gaunet, 1997; Ungar, 2000). My motivation for this work is to design a handheld touch-screen application enhanced by spatial tactile feedback to allow visually impaired users to obtain geographical information for places of interest. In this section, I first review systems that use audio and tactile feedback to enable visually impaired users to explore maps. I then discuss location-aware systems using audio and tactile feedback to provide visually impaired users with place-specific information *in situ*.

Spatial Learning with Physical Maps

Tactile maps have been recognized as an effective means of learning spatial relationships between multiple objects for the visual impaired. Herman *et al.* (1983) experimented with how accurately visually impaired people can locate objects in a space after they were exposed to a miniaturized version of the space layout. They found that participants could indicate the locations of objects from any point in the space fairly accurately. Ungar *et al.* (1993, 1995) found that children with visual impairments can also learn a spatial layout through tactile maps.

Espinosa *et al.* (1998) investigated how tactile maps can enhance visually impaired people's learning the spatial layout of urban environments (Madrid, Spain and Sheffield, United Kingdom in their studies). Their experiments showed that participants learned the spatial layout at a comparable accuracy with a tactile map or by direct experience (*i.e.*, walking freely along the routes between landmarks). Their study also revealed that participants demonstrated better learning with tactile maps while physically exploring the space in comparison to cases in which tactile maps were not provided.

The effectiveness of physical artifacts besides tactile maps also have been tested. Picard & Pry (2009) examined how a 3D small-scale model can help visually impaired users understand the spatial layout of three places. They found that participants' spatial knowledge was significantly enhanced through the repeated exposure to the small-scale model.

These studies highlight benefits of tactile maps in learning spatial relationships between multiple objects. However, these maps are physical, and their production cost is often considerable. These maps may sometimes not have sufficient detail or updated information of the space. Digital maps can address these issues; in the next section, I describe some previous explorations of how to interfaces can be designed to present digital map information through the tactile feedback channel.

Interfaces for Map Exploration

Auditory feedback is often used to make map information accessible to people with visual impairments (Heuten *et al.*, 2006; Sánchez *et al.*, 2010; Walker & Lindsay, 2006). To allow people with visual impairments to explore and learn geographical information, touch-based interaction can be integrated with such audible map systems. For example, the NOMAD system (Parkes, 1998) uses a touchpad placed under a paper Braille map to detect which part of the map the user is contacting. When the user touches an area on the map that contains information, the system generates speech feedback to describe the user's contact point.

Jacobson (1998) developed a similar system which replaces the use of static Braille maps with dynamically generated maps. It provides auditory feedback (either speech or non-speech sound) to inform what the user is touching and its surrounding content. Jacobson validated that participants with visual impairments could reconstruct map information after they used such a system. Parente & Bishop (2003) showed that vibrotactile feedback when combined with speech and non-speech audio feedback helps visually impaired users discover the boundaries of mapped elements (*e.g.*, boundaries of states in the United States map)

In the LoadStone project, Lahav and Mioduser developed a virtual environment system which combines speech audio and force feedback through a joystick. Their user study showed that the system enabled visually impaired users to build a cognitive map of an indoor location, and successfully navigate the corresponding physical space based on this cognitive map.

In addition to describing areas on a map, systems can also explain possible routes between different areas. Google Intersection Explorer is a handheld application which allows the user to drag her finger on the screen to explore walkable paths that a person can take from any intersection.

To help visually impaired users learn geometrical patterns of walkable path segments, Timbremap (Su *et al.*, 2010) uses non-speech sound cues, and helps the user trace routes on a handheld device. Its evaluation showed that participants with visual impairments could learn non-trivial geometries using only touch and simple audio cues. Similarly, Crossan & Brewster (2008) demonstrated that the combination of force feedback and non-speech auditory feedback can also facilitate the learning of trajectories.

The main focus of SpaceSense is to provide visually impaired users with a high-level understanding of spatial relationships between multiple locations instead of the exact shape of a street or a building. Thus, SpaceSense complements these existing systems by facilitating the acquisition of route information and high-level spatial information about multiple locations (*e.g.*, learning directions to the destination and spatial relationships between different locations with my system, and knowing the shape of a path to walk with one of the systems above).

Interfaces for Navigation and Wayfinding

Another opportunity for visually impaired users to explore and learn about places is *in situ* during navigation. There are a number of commercially available systems, such as Sendero GPS, Trekker, Trekker Breeze, and open-source systems (Loomis *et al.*, 1994) that provide speech navigation instructions and information about points of interest.

A system developed by Azenkot *et al.* (2011), ChattyEnvironment (Coroama, 2006), and Talking Points (Stewart *et al.*, 2008) are examples of research on location-aware navigation applications which tell the user about her environment. However, users may prefer to be able to limit the amount of information based on their *in situ* needs (Coroama, 2006); alternatively, a system can intrinsically limit it to what is immediately within Bluetooth range (Stewart *et al.*, 2008).

Past research has explored using auditory feedback to provide visually impaired users with trajectory information for wayfinding (Helal *et al.*, 2001; Marston *et al.*, 2006; McGookin *et al.*, 2008; Wilson *et al.*, 2007). Marston *et al.* (2006) showed that visually impaired users can navigate an environment faster using a continuously generated beeping 3D sound which encodes the direction to the next waypoint than Talking signs (using sound feedback when the user pointed the device towards the next waypoint). In addition, Wilson *et al.* (2007) showed that the user can successfully learn information about the environment (*e.g.*, obstacles in the way) encoded in different dimensions of the sound (*e.g.*, a sound type, pitch, and rhythm) or time-compressed speech.

Tactile feedback has also been used to aid visually impaired users in wayfinding (Amemiya & Sugiyama, 2009; Ross & Blasch, 2000; Zelek *et al.*, 2003). Ross & Blasch (2000) compared speech audio, non-speech audio, and tapping haptic feedback to indicate the direction in which a user should walk. They found that visually impaired users overall performed navigation tasks fastest and with fewest errors using the haptic interface, with non-speech audio coming in close second.

Zelek *et al.* (2003) showed that a haptic glove, which conveys the locations of obstacles in an environment through vibrations on different fingers, could help the user identify where they can walk. Amemiya and Sugiyama developed a mobile device using force feedback to provide the user with a sensation of being pulled towards the destination by the device, allowing her to navigate at her normal walking speed (Amemiya & Sugiyama, 2009).

These projects aim to help the user learn their environment in situ and reinforce their cognitive map of the geographic area. SpaceSense complements these systems by providing additional means that visually impaired users can leverage to understand the spatial relationship between multiple places and prepare for future trips to those environments. Thus, the main research question that I wanted to answer through the SpaceSense system was as follows:

RQ-3. *What level of accuracy and detail can visually-impaired users achieve when they try to construct mental maps of spatial relationships between multiple objects with the support of spatial tactile feedback?*

5.3 The SpaceSense Device

SpaceSense is a map application that runs on a handheld touch-screen device (an iPhone in my current prototype) enhanced by my custom spatial tactile feedback system. My motivation for choosing a mobile touch-screen device is that it allows for taps and flick gestures—interactions that are easy for visually impaired users to perform (Kane *et al.*, 2009; Paladugu *et al.*, 2010).

I used my second prototype described in Section 3.3. This hardware allows me to directly encode eight different directions (*i.e.*, north, north-east, east, south-east, south, south-west, west, and north-west) into eight different positions of spatial tactile feedback (*e.g.*, vibration at the top-right of the device means north-east).

For audio feedback, SpaceSense uses FliteTTS for iPhone to read out information through a synthesized voice. I chose this package instead of the iPhone VoiceOver functionality already built in some of the iPhone devices because I wanted the ability to precisely tune the timing of the speech feedback in SpaceSense so that the feedback will accompany spatial tactile feedback.

5.4 SpaceSense Interactions

SpaceSense supports three tasks related to acquiring spatial information about multiple locations and directions to each location. In this section, I describe interactions in these tasks supported by SpaceSense: identifying locations, learning place details and learning directions.

In SpaceSense, information is presented as a linear list. The user can perform simple flick gestures to navigate a list and select an item in the list. I selected flick gestures because they are easy for visually impaired users to perform (Kane *et al.*, 2011; McGookin *et al.*, 2008). The user can also repeat the current item by double-tapping the screen. I chose a double tap gesture because the user is less likely to perform a double tap accidentally than a single tap.

Identifying Locations

SpaceSense allows the user to select places of interests from a pre-defined list of categories ("restaurant," "café," "bar," "hotel," "attraction," and "bank" in the current prototype system). This list was determined based on the one in the map application available in Android mobile phones. Figure 5.2 shows an example of interactions through which the user is selecting a category. The user can perform a downward or upward flick gesture to navigate the linear list of the categories. When the user selects a category by performing a rightward flick gesture, the system retrieves up to 20 places (using Yelp API 2.0) sorted by distance within a 2 km radius centered on the user's current simulated location. SpaceSense then will start to present information about each place, starting with the closest one.

SpaceSense offers spatial information in an exo-centric manner similar to when a person views at a map. It reads the name and street address of each place (*e.g.*, "Alice's coffee shop, 100 John St., 0.4 km."), and uses spatial tactile feedback to indicate the place's distance and cardinal direction. For instance, if the place is to the west of a referenced location (e.g., the hotel where the user may be staying on an upcoming trip), the left side of the device vibrates (north is always set to the top of the device). I designed the current prototype to provide vibration at four different strength levels (100%, 80%, 60%, and 30% strength output of a vibration motor) to represent the distance (below 200 m, below 500 m, below 1 km, or farther than 1 km, respectively).

Interaction	Audio Feedback / Note
	"Welcome. What would you like? Select a category." (The system provides a welcome message.)
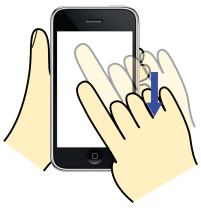	(The user performs a downward flick gesture to navigate the linear list of the categories.)
	"Restaurant." (The user is now on the first item.)
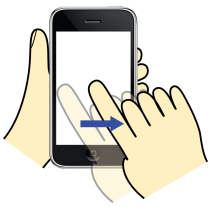	(The user performs a rightward flick gesture to select this category.)
	"Retrieving the data about restaurants." (The system searches restaurants nearby.)

Figure 5.2. Interactions in selecting a category. The interactions are similar when the user is selecting a place, browsing details of a place, and navigating step-by-step instructions of directions.

Learning Place Details

Item	Example
Name	*"Moonbean Coffee Shop"*
Address	*"30 St. Andrew St."*
Phone	*"416 595 0327"*
Distance	*"0.6 km"*
Review tags	*"Soy latte, Friendly staff, Great place, Back patio, Good coffee, Front patio, Helpful staff, Best coffee, Great person."*

Table 5.1. An example of place details provided by SpaceSense. The user can navigate these items by performing upward or downward flick gestures.

After the user selects a place, SpaceSense provides her with general information about the place, including the name, address, phone number, and distance. Table 5.1 shows an example of information SpaceSense offers. The user can navigate this information using upward or downward flick gestures.

SpaceSense also offers an overview of what people mention about a specific location on online review Websites (such as Yelp.com). It uses the adjective-noun word pairs extracted from the original review text that Review Spotlight (Yatani *et al.*, 2011) normally presents in a visual interface. Instead of visually displaying word pairs, SpaceSense reads out the ten most frequently-mentioned adjective-noun word pairs through the speech feedback instead of the original review text (see the "Review tags" item in Table 5.1).

Learning Directions

Finally, SpaceSense gives the user directions to the place she selects. The user can perform a rightward flick gesture to select a place while she is browsing its details (as described in the

Interaction	Map	Audio Feedback / Note
	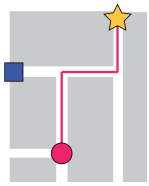	*"You will walk on three streets to get to this place. 1. Head north on University St., 0.3 km. Now you are at University St. and College Ave."* (The system provides the first step of the directions to the destination.)
	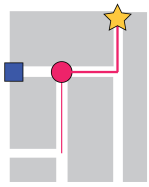	*"From this point, your destination, Alice's coffee shop, is in this direction."* (The top-right of the device vibrates to indicate the direction and distance towards the destination.)
	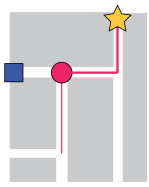	*"Your bookmared place, Jim's shop, is in this direction."* (The left of the device vibrates to indicate the direction and distance towards the place the user previously bookmarked.)
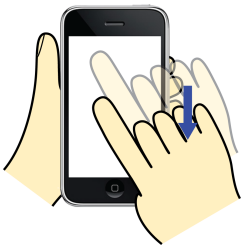	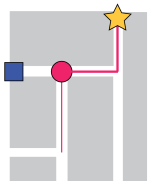	(The user is moving to the next instruction.)

Figure 5.3. Interactions in learning directions. The circle, star and square in the map represent the user's simulated location, destination, and bookmarked location, respectively.

previous section). The system then begins to provide step-by-step instructions (using speech feedback) to reach the destination. SpaceSense uses the Google Directions API to obtain these walking instructions. The system presents each step of the instructions along with its walking distance (*e.g.*, "head north on University St., 0.3 km") one at a time (see the first figure in Figure 5.3), followed by the intersection at which the simulated performance of the previous step would put the user.

After reading out the instruction and intersection, the system conveys the distance and direction to the destination through spatial tactile feedback. For example, if the destination is located to the north-east from the current simulated user location, the top-right side of the device will vibrate (see the second figure in Figure 5.3). Spatial tactile feedback is provided every time the system presents an intersection. In this manner, SpaceSense shows the user how the relationships of the destination and other locations change through the simulated movements of the referenced location similar to when reading a map.

The system also provides spatial tactile feedback to nearby bookmarked places (the bookmark functionality will be explained in the next section). The system will provide vibration in the direction of bookmarked places near the simulated route while the speech feedback offers the name of the location. For example, in the third figure in Figure 5.3, the system will vibrate on the left side of the device to indicate the bookmarked place (Jim's shop). I designed this feature to help the user build and maintain spatial relationships between multiple places of interest.

Bookmarking a Place

The user can add a place presented by SpaceSense to her bookmark list. This bookmark functionality allows the user to save and quickly access places that she likes. Most of the interactions related to bookmarking are two-finger gestures. For example, when the system is in the mode of providing the details of a place, the user can perform a two-finger rightward

swiping gesture on the touch screen to add the place to the bookmark list.

The user can also navigate her own bookmark list. By doing a two-finger downward gesture at any time, the user can move to her bookmark list. Similar to providing information about nearby places, SpaceSense offers the name of each bookmarked place and its distance from the current simulated user location. The user can also obtain its details and directions through the same interactions as described in the previous sections.

To remove a particular place from the bookmark list, the user can perform a two-finger leftward gesture while she is navigating the details or directions of the place. She can also perform a two-finger upward gesture to leave the bookmark mode.

5.5 User Study

To examine the efficacy of SpaceSense, I conducted a study with visually impaired users. My study consisted of two parts to independently evaluate how well the features in SpaceSense assist visually impaired people through: 1) an evaluation on the audio Review Spotlight presentation; and 2) an evaluation of the system's effect on the learning of routes and the spatial relationships between multiple locations. Here, I report the second part of the study as it is the main investigation related to spatial tactile feedback. Please refer to Appendix B for the experimental design and results of the first part.

Experimental Design and Procedure

I selected two Toronto neighborhoods with an area of approximately 4 km^2. For each neighborhood, I set the starting point near the center, and selected four locations which were on average 770 m from the starting point, and required 3 turns to reach. I labeled these locations as the stores of four hypothetical persons. Figure 5.4 shows the map areas and the locations used in the study.

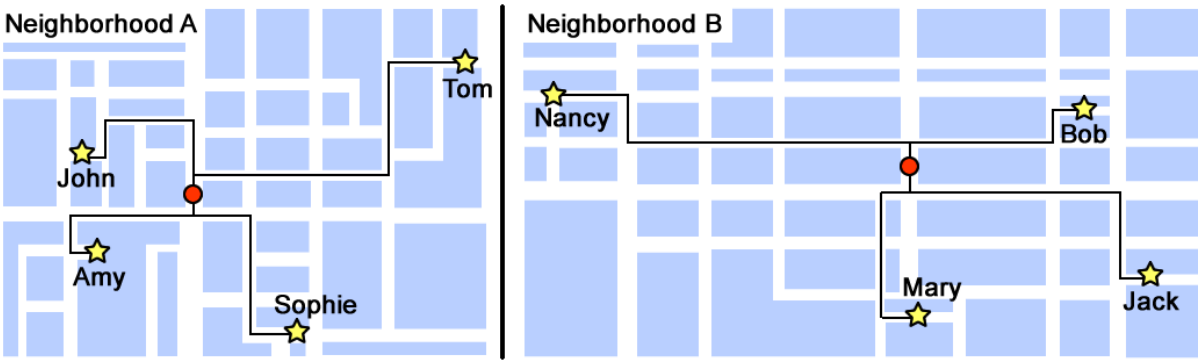

Figure 5.4. Maps with four hypothetical locations used in the user study, and routes to those places from a given starting point.

Neighborhood A		Neighborhood B	
Destination	Bookmark	Destination	Bookmark
Amy	John	Bob	Mary
John	Tom	Mary	Nancy
Tom	Sophie	Nancy	Jack
Sophie	Amy	Jack	Bob

Table 5.2. The combinations of the destinations and bookmarked places used in the study.

To examine whether spatial tactile feedback could help people with visual impairments develop a high-level understanding of spatial relationships between multiple places, I asked participants to learn the directions to four locations on a map, one at a time. For each destination, another place from the same map was set in the bookmark list beforehand; thus, the instruction provided to the user would always include directional information towards the destination and one other place. The combination of the destinations and bookmarked places used in the study was randomly determined (Table 5.2). This setup allowed me to examine how well SpaceSense could help participants learn spatial relationship between multiple place.

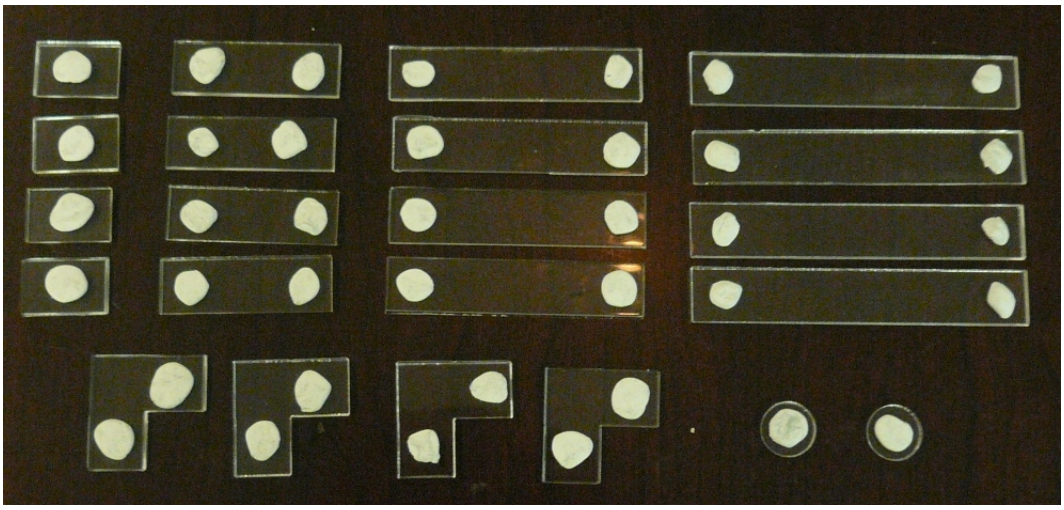

Figure 5.5. Acrylic pieces used in this user study. I-shape pieces with four different lengths, L-shape pieces and circular pieces were prepared for the map composition.

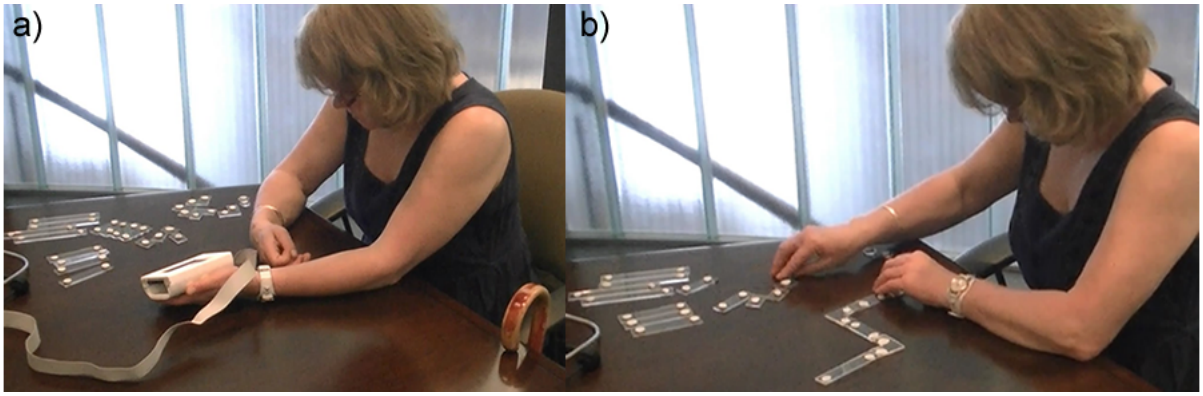

Figure 5.6. Route composition by a participant during the experiment. a) A participant was learning the direction to a destination and its spatial relationship to the bookmarked place; b) After she received all instructions, she constructed a map with acrylic pieces (Figure 5.5).

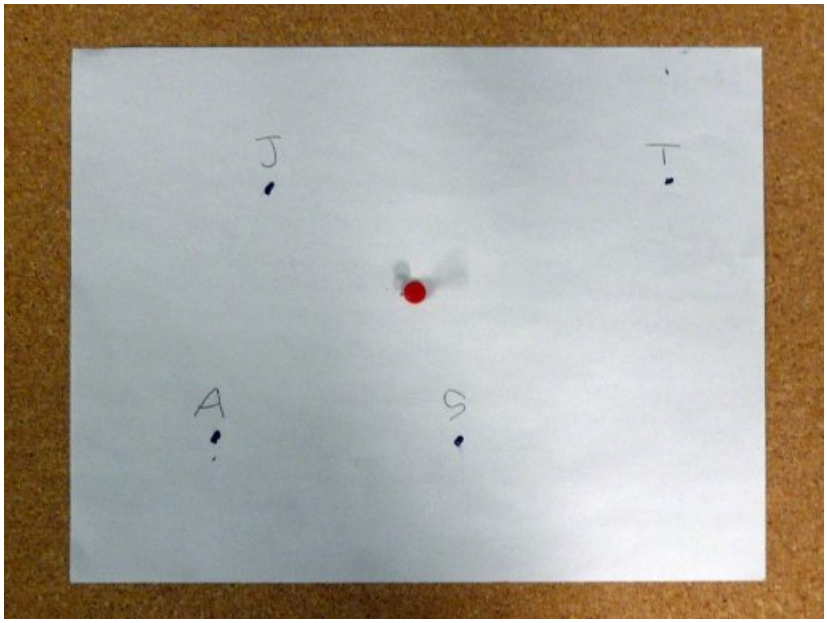

Figure 5.7. A location point drawing example for Neighborhood A. The red pin at the center of the paper indicates the starting point. The experimenter made annotations on the drawing for the later analysis.

After the participants went through the directions to one place, they were asked to reproduce the route with acrylic pieces and indicate the locations of the destination and the bookmarked place (Figure 5.5 and 5.6). This route reconstruction is a common task used to probe the visually impaired user's understanding of spatial knowledge (Kitchin & Jacobson, 1997). I prepared thin rectangular pieces with four different lengths for streets (3 cm, 6 cm, 9 cm, and 12 cm long for loosely representing from very short to very long street segments), L-shape pieces for corners, and circles for the destination and the bookmarked place. The experimenter did not correct the participant's route composition at any point.

After participants were exposed to the directions to all the locations, the experimenter asked participants to make a drawing of where all the places are located—indicating their perception of the positions of the four places on a blank sheet of paper with a blue marker (Figure 5.7). The experimenter made annotations for later analysis, but did not make any correction even if

the spatial relationships among the four places indicated by participants were incorrect.

I set up two conditions to compare: *Tactile* (the SpaceSense system design described in Section 5.4) and *Audio* (an interface which provided all information including the direction and distance for each location through only speech feedback). In the *Audio* condition, the system read out the direction and distance towards the destination and bookmarked place (*e.g.*, "Your destination, Amy's café, is to north-east, 0.5 km."). I tuned the speed of the speech feedback at a slower rate than the iPhone VoiceOver so that participants could follow the information provided by the system more easily. The presentation order of the two interface conditions and the neighborhoods (Figure 5.4) were counter-balanced across the participants. At the beginning of the task, I provided all participants with training and practice using each interface until they were comfortable with the procedure and system.

Participants

I recruited twelve visually impaired users (4 male and 8 female; P1–P12) on the word-of-mouth basis. The level of their visual impairments varied: 10 of them were blind, and 2 had low vision. Before the experiment, I asked them to describe familiar areas in the city, and none of them expressed a significant familiarity of the areas used in this experiment. The study took on average 80 minutes for participants to complete. Participants were compensated $50 cash for their participation after the study.

5.6 Results

I analyzed data on participants' learning of routes and spatial relationships between four places in a map and acquisition of information about places. I report task completion time, accuracy, and comments from participants.

	Neighborhood A		Neighborhood B	
Neighborhood	**A**	**B**	**A**	**B**
Tactile	*203* (64)	*195* (136)	*86* (16)	*97* (32)
Audio	*189* (99)	*220* (83)	*103* (45)	*95* (28)

Mean (SD)

Table 5.3. The time participants spent in learning the route instructions (*InstructionTime*) and reproducing a route with acrylic pieces (*CompositionTime*).

Task Completion Time

Participants used the system on average for 201 seconds to learn the route instructions (*InstructionTime*) and 95 seconds to reproduce a route with acrylic pieces (*CompositionTime*). Table 5.3 shows the results of these two time metrics across the conditions and neighborhoods. Participants were exposed to both conditions but with different neighborhoods. Thus, I ran a two-way mixed ANOVA on both time measures for the conditions and neighborhoods. It did not reveal a significant difference on either *IntructionTime* (Condition: $F_{(1,10)} = 0.04$; Neighborhood: $F_{(1,10)} = 0.29$; Condition $\times$ Neighborhood: $F_{(1,10)} = 0.26$; $p > .05$ for all) or *CompositionTime* (Condition: $F_{(1,10)} = 1.18$; Neighborhood: $F_{(1,10)} = 0.06$; Condition $\times$ Neighborhood: $F_{(1,10)} = 0.51$; $p > .05$ for all).

Route Composition Accuracy

I first examined how accurately participants could compose routes to destinations. Figure 5.8 shows examples of the routes created by participants. I adapted an evaluation approach used by Passini *et al.* (1990) to analyze the route compositions (*NumberElementsError*, *FormElementsError*, *PositionError*, and *PlacementError*). I included three additional

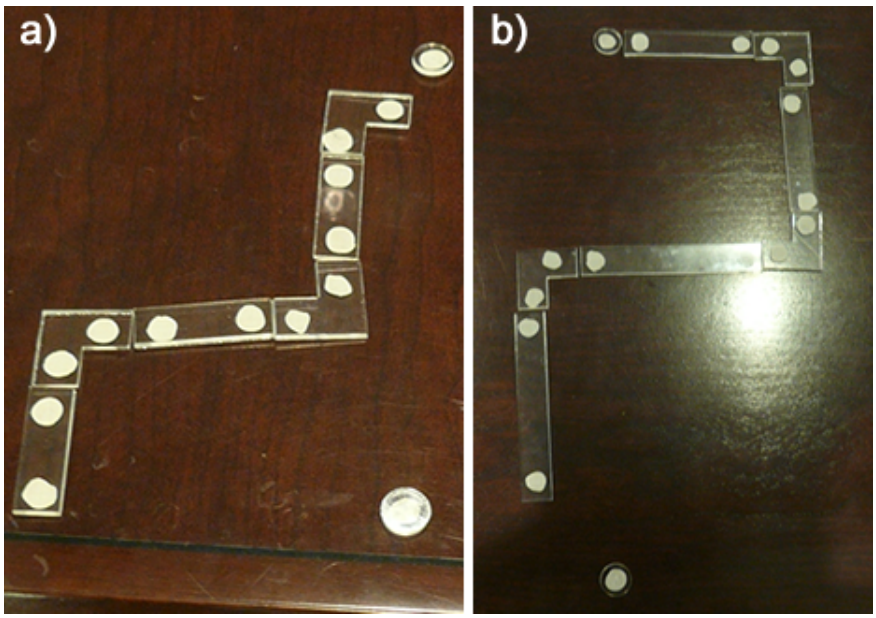

Figure 5.8. Route composition examples. They represent the route to Tom with Sophie as the nearby bookmarked place in Map A made by two different participants. a) A composition close to the correct route; b) An incorrect composition.

metrics to measure the accuracy of the positions of the destination and bookmarked place (*DestinationDistanceError*, *BookmarkDistanceError*, *TwoPointAngleError*). Table 5.4 illustrates the definition of each metric and its observed value for the two interface conditions. I ran a paired t test for each metric, and did not find a significant difference in any of the metrics at the 95% confidence level.

Drawn Places Accuracy

For the drawn places (Figure 5), I evaluated the spatial relationship between any two of the places from the starting point with three levels of ratings:

Metric and Description	*Tactile*	*Audio*
NumberElementsError The number of unnecessary acrylic pieces that participants used to recreate a route.	*0.80* (1.27)	*0.83* (1.18)
FormElementsError The Levenshtein distance (the minimum number of operations required to transform one sequence into the other) between the participant's route composition and the correct route composition.	*0.83* (1.28)	*0.87* (1.18)
PositionError The number of incorrect orientations of the L-shape blocks.	*0.77* (1.18)	*0.54* (0.93)
PlacementError The number of I-shape blocks used with the incorrect length.	*2.75* (1.16)	*2.79* (1.13)
DestinationDistanceError The absolute difference of the straight line distance between the starting point and destination from the one in the correct route composition.	*2.11* (1.47)	*2.25* (1.82)
BookmarkDistanceError The absolute difference of the straight line distance between the starting point and bookmarked place from the one in the correct route composition.	*2.67* (1.90)	*2.55* (1.73)
PlaceAngleError The absolute difference of the angle between the edge connecting the starting point and destination and the one connecting the starting point and the bookmarked place from the one in the correct route composition.	*53.0* (35.8)	*46.7* (33.8)
		Mean (SD)

Table 5.4. The map composition error metrics and the average and standard deviation of each measured metric. T-tests did not show a significant difference in any of the metrics.

2: Very close to the correct placement of the two places,

1: Neither completely correct nor incorrect, and

0: Not correct.

The second rater and I independently rated all the drawings after agreeing that this rating scheme can represent the accuracy of the drawings with several randomly chosen drawings. To measure the inter-rater reliability of this ordinal scale, I calculated the Cohen's κ with the squared weighting (Cohen, 1986). In this calculation, the disagreements of the ratings were weighted according to their squared distance from perfect agreement. As a result, the weighted Cohen's κ was .92 (95% CI: [.87, .96]), showing a strong agreement. The average rating was used to determine the accuracy of each drawing.

Table 5.5 shows the ratings for the correctness of the spatial relationship of all possible pairs of places in each map across the two conditions. Because these ratings should be considered as ordinal data, I decided to use a non-parametric test. A Mann-Whitney test found a significant difference in the correctness ratings between the *Tactile* and *Audio* conditions ($Z = 1.96$, $p < .05$, the effect size $r = 0.23$).

The drawings participants created after using SpaceSense were more accurate than those which provided map information using the speech feedback solely. The main difference was in what direction participants perceived the places to be from one to another. The results indicate that Neighborhood B was seemingly harder than Neighborhood A for participants to learn. But regardless of the neighborhoods, participants understood the spatial relationships between locations better in the *Tactile* condition than the *Audio* condition.

User Feedback

As one participant mentioned, the system provided information in a way that was very similar to what she experienced when orientation and mobility specialists taught her routes.

Map	Places	*Tactile*	*Audio*
A	Amy – John	*0.91* (0.73)	*0.50* (0.76)
	Amy – Sophie	*1.00* (0.58)	*0.67* (0.47)
	Amy – Tom	*1.08* (0.18)	*1.00* (0.81)
	John – Sophie	*0.83* (0.68)	*0.83* (0.89)
	John – Tom	*0.75* (0.50)	*1.00* (0.82)
	Sophie – Tom	*1.50* (0.50)	*0.75* (0.69)
	Average	***1.01* (0.59)**	***0.79* (0.78)**
B	Bob – Jack	*0.16* (0.37)	*0.58* (0.83)
	Bob – Mary	*0.50* (0.76)	*0.33* (0.55)
	Bob – Nancy	*0.58* (0.73)	*0.41* (0.60)
	Jack – Mary	*0.33* (0.47)	*0.33* (0.55)
	Jack – Nancy	*0.50* (0.76)	*0.17* (0.37)
	Mary – Nancy	*1.33* (0.74)	*0.25* (0.38)
	Average	***0.57* (0.75)**	***0.35* (0.59)**

Mean (SD)

Table 5.5. The ratings for all the combinations of two places in each map across the two conditions.

Participants expressed that the SpaceSense system allowed them to develop a rich cognitive map of an area in a way that is similar to and possibly even better than with Braille maps.

> *"[The vibration] was helpful because it gave me a very visual sense of where things should be… Instead of having to have someone label and make the map tactile, [the system] did that for me. And instead of having to take out a piece of paper with streets and stuff labeled and then have to look and see 'ok this is south-west,' for example, the vibration gave me that visual sense."* [P5]

Most of the participants noted that the strength of the SpaceSense system lies in its use of both the auditory and tactile channels. Particularly, they liked having directional information provided through the tactile channel.

> *"It gave me information about the direction more quickly. Because it takes more time to say 'north-east' or 'south-west.' But feeling the vibration in your hand gets the information to my brain more quickly.* [P2]

Participants felt that they were able to develop a mental map of the places more quickly with SpaceSense than the system using only speech feedback. The tactile feedback provided in each step of the directions helped to confirm their mental map.

> *"I could anticipate the next direction based on the vibration of the locations. It took longer [to do the same thing] with audio... I sort of knew which direction was next because the vibration was pointing me to a particular direction. So I could anticipate the audio instruction. I could anticipate that because of the vibration."* [P3]

In comparison, when only given audio interface, the participants described needing to work harder to construct a cognitive map.

> *"I had to abstractly think where we are going, and put the information provided by the system together. It is a little easier to put together in the map with the tactile stimulation with the combination of the sound as supposed to [audio-only]."* [P4]

5.7 Discussion

Information Overload through the Audio Channel

I found that the place drawings by the participants were more accurate in the *Tactile* than the *Audio* condition. The major reason participants found benefits with having a separate feedback channel for directional information was that spatial tactile feedback enhanced their memory of the spatial relationships between locations. P4 pointed out that it was difficult for her to maintain all the information in her mind in the *Audio* condition.

> *"(In the Audio condition), I had to listen to the audio over and over just to get the direction, right or left. And I had to keep track of Mary's store, Bob's store, the two names of the people. And then I had to keep track of the directions to get there."* [P4]

But she commented that having directional information through the tactile channel helped her concentrate on route information.

> *"With the tactile stimuli, you get the directions in your hand. So you don't have to worry (about the directions) because you can feel it. So you take it away from your memory. And now you just focus on how to get there."* [P4]

Their subjective impressions were corroborated by the difference in their drawing of spatial relationships between four locations across the two conditions. Thus, I conclude that spatial tactile feedback can help participants understand spatial relationships between multiple locations.

Errors and Limitations on Learned Routes and Spatial Relationships

As Table 5.4 and 5.5 indicate, the accuracy of the route and spatial relationships between places was not high. One reason might be that participants still often had to process information provided through the speech feedback even while the spatial tactile feedback provided the directional information. Future work is necessary to understand what an effective presentation of geographical information would be in order to support visually impaired users to maintain accurate route information and build a cognitive map of multiple places.

The current SpaceSense system only provides high-level spatial relationships between locations. When visually impaired users navigate in the actual space, they may also need other information, such as the shape of a street or intersection. I believe that integration with a system like Timbremap (Su *et al.*, 2010) would enable users to gain such information. Future work is to extend the SpaceSense system and investigate how it can support the user's acquisition of both high-level and low-level geographical information.

Study Limitations

There are several limitations to mention in this study. The user study included only four places in one neighborhood. During the presentation of route instructions, only two places were presented through the spatial tactile feedback (the destination and bookmarked place). Future work needs to investigate how the number of places in the space and the number of places presented in the route instructions can affect people's learning of spatial relationships.

The two neighborhoods used in the study did not include curved streets. Some cities have many non-straight streets, and improvements may be necessary to correctly provide route information and spatial relationships between multiple locations through the spatial tactile

feedback. But, we believe that the results would be transferrable to cities which have similar street configurations to what was tested in this work.

There are several aspects of the system which were not covered in this paper. For example, due to the large difference between the number of congenitally and after-birth blind participants, we did not examine the effect of this difference in learning spatial relationships. Our current implementation of SpaceSense uses the exo-centric presentation of directions. But the ego-centric presentation can benefit users better in some cases (*e.g.*, while the user is navigating the space physically). Our study shows that participants were able to learn routes and the spatial relationship between places through an exo-centric presentation of the map information similar to when a person reads a map before visiting a location; further research is necessary to investigate how to best present the spatial relationship of locations through a system like SpaceSense while the user is navigating *in situ*.

5.8 Summary

I developed SpaceSense, a handheld system using spatial tactile feedback to help people with visual impairments acquire details about places, learn the directions to a place, and understand high-level spatial relationships between multiple places. My research question with SpaceSense was as follows:

RQ-3. *What level of accuracy and detail can visually-impaired users achieve when they try to construct mental maps of spatial relationships between multiple objects with the support of spatial tactile feedback?*

My user study found that participants with visual impairments could maintain spatial relationships between four places on a map more accurately in the condition where directional

information was presented using spatial tactile feedback than only speech feedback.

I note that SpaceSense does not aim to support all aspects of learning the spatial layout or replace learning through actual explorations in a space. As Spencer & Travis (1985) discussed, a full understanding of a space needs actual explorations by the individual herself. SpaceSense will offer an ability of identifying locations of interest nearby and a high-level understanding of the spatial layout of the space, which can facilitate the actual exploration of the space.

Chapter 6

Collaboration Support

Mobile and handheld devices are becoming commonplace as devices to support remote collaboration. People now use their mobile devices to coordinate and interact with others through email, instant messaging, and video conferencing. They are using their mobile devices to communicate and collaborate in many similar ways to how people often do so with desktop computers.

However, the physical form factor of mobile devices can impact user collaboration. For instance, mobile devices offer a much smaller visual workspace in comparison to those available on desktop computers or tabletop displays. Furthermore, when a user interacts with a touch-screen interface, her hand can occlude much of the screen (Vogel & Balakrishnan, 2010). As a result, visual feedback often used in synchronous collaborative tasks to help users maintain awareness of their partner's activities (Gergle *et al.*, 2004) may impact collaboration differently on mobile devices.

I see many design opportunities for improving collaboration through mobile devices. This includes using different communication channels for providing feedback to collaborators such as vibrotactile feedback. Particularly, spatial tactile feedback can be very useful to

convey spatial information about collaborators and support successful spatial coordination. Furthermore, tactile feedback can be used to design more accessible user interfaces, particularly for visually impaired users as seen in Chapter 5. Understanding the effects of tactile feedback on collaboration can contribute to better design of collaborative systems for visually impaired users as well as sighted users.

I investigated how different feedback channels can affect collaboration on mobile devices. I specifically was interested in spatial coordination (Gergle *et al.*, 2004), an important aspect of collaboration. Such collaboration can happen when users are sharing screen views and discussing the content within them. To measure the effectiveness of feedback for supporting spatial coordination, I built a collaborative game played by two users on mobile touch-screen devices, and I observed how users played the game with and without receiving visual or tactile feedback about their remote game partner's action.

My two user studies discover the different benefits of each feedback channel, and also demonstrate better performance through the combination of visual and spatial tactile feedback for spatial coordination in collaborative handheld systems. Visual feedback can provide precise spatial information about collaborators, but degrades collaboration when feedback is occluded and sometimes can distract the user's attention. Spatial tactile feedback may not be appropriate for conveying precise spatial information; however, it can reduce the overload of information in visual space and can gently guide a user's attention to an area of interest. Furthermore, visual and spatial tactile feedback can complement each other, and systems using both feedback channels can offer better spatial coordination support than systems using only one form of feedback.

6.1 Related Work

Vibrotactile feedback has been widely used to provide tactile sensations on touch-screen devices. These vibrotactile feedback systems are useful for a variety of tasks, such as item

selection in a linear list (Poupyrev *et al.*, 2002a), text entry (Hoggan *et al.*, 2008a), and eyes-free interaction (SemFeel, Chapter 4). These systems mostly provide the user with feedback about her own interactions on a mobile device. But, I here focus on the use of tactile feedback for offering information about another user's interaction on a remote device to support the collaboration between users. In this section, I review collaborative systems using tactile or haptic feedback, and research studying communication strategies supported by this feedback channel.

Inter-personal Communication Systems

The use of tactile or haptic feedback in collaborative systems initially was explored in inter-personal communication systems. HandJive consists of two joystick-like devices enhanced by haptic feedback (Fogg *et al.*, 1998). One person's movement on the device is propagated onto another device as an orthogonal movement (*e.g.*, when a user moves a HandJive device forward, the partner will feel horizontal movement from the device). Their informal user study found that participants understood the concept of HandJive quickly, and enjoyed moving sensation conveyed through the devices.

InTouch developed by Brave *et al.* (1998) is a device with three cylindrical rollers, and the rotational velocity of each roller is synchronized with the paired device. Thus, remote users can feel each other's interaction with the rollers over a distance.

Chang *et al.* (2002) investigated the effects of vibrotactile feedback on remote voice communication. In their ComTouch system, they explored the use of vibrotactile feedback as a supplemental channel for voice communication. Their user study revealed that the participants used tactile communication for five purposes: emphasizing a voice message, turn-taking, duplicating a tactile message, responding "yes" or "no," and conveying integer numbers.

Brown *et al.* (2009) explored how couples would develop communication protocols with their mobile audio-tactile messaging system called Shake2Talk. They identified four purposes for audio-tactile messages: coordinating events and calls, maintaining awareness, sharing the fun, and expressing affection.

Chan *et al.* (2008) examined the effects of haptic icons on turn-taking in collaborative tasks. Their user study revealed that vibrotactile feedback can be useful to communicate some messages, particularly urgent requests in a visually-demanding situation (*i.e.*, the participants were instructed to solve puzzles while communicating through tactile feedback).

Pielot *et al.* (2010) built a belt-like device with multiple vibration motors and studied the effects of using it to provide tactile feedback about where other team players are located in a 3D multiplayer game. Their user study revealed that users could sense the other players' locations, but did not show clear evidence on whether collaboration between the players was improved through their tactile system.

Collaborative Systems for Visually Impaired Users

Tactile feedback is accessible to visually impaired users, and thus collaborative systems with tactile feedback often extend to this population. Plimmer *et al.* (2008) developed McSig, a system to support visually impaired users with learning how to write letters. Through a PHANTOM device held by a visually impaired user, the system haptically reproduces the trajectory of a letter written by a sighted user on a Tablet PC. Their user study showed that visually impaired users successfully could learn how to write some alphabets.

McGookin & Brewster (2007) explored the use of a collaborative tactile system to allow visually impaired users to explore a bar graph collaboratively. Through a user study with pairs of visually impaired users, they found that haptic feedback was useful for one user to guide the other user to the point of interest effectively.

However, collaborative systems for visually impaired users are still at the early stage, and understanding the effects of different types of feedback including tactile feedback could be useful for designing collaborative systems for visually impaired users as well as sighted users.

Spatial Coordination

My work focuses on how tactile feedback can support spatial coordination in collaborative tasks. Spatial coordination is an important aspect of interaction between users to successfully accomplish collaborative tasks (Gergle *et al.*, 2004), but it has not been deeply studied in the context of tactile feedback.

Oakley *et al.* (2001) implemented several haptic effects for supporting communication in tasks with a collaborative editor using the PHANTOM device. For example, one user can produce haptic feedback (called haptic gestures) on the other user's device to guide her to a specific point of the screen. They found that the participants frequently used haptic gestures to communicate with each other about the region of interest within the shared visual working space or objects they would want to discuss.

Jay *et al.* (2007) conducted a study about examining the effects of delay on haptic and visual feedback on the collaborative tasks which require strict spatial and temporal coordination. Their task was to move two cursors (each operated by one user) towards the target in a graphical user interface with maintaining the relative distance of the cursors within the threshold distance. Their results show that even a very small delay (25 msec in tactile feedback, and 50 msec in visual feedback) can impact the collaboration when strict spatial and temporal coordination is necessary.

My main interest is to investigate how spatial tactile feedback can support spatial coordination in mobile touch-screen devices instead of desktop or large displays tested in prior work. More specifically, I wanted to answer the following research question:

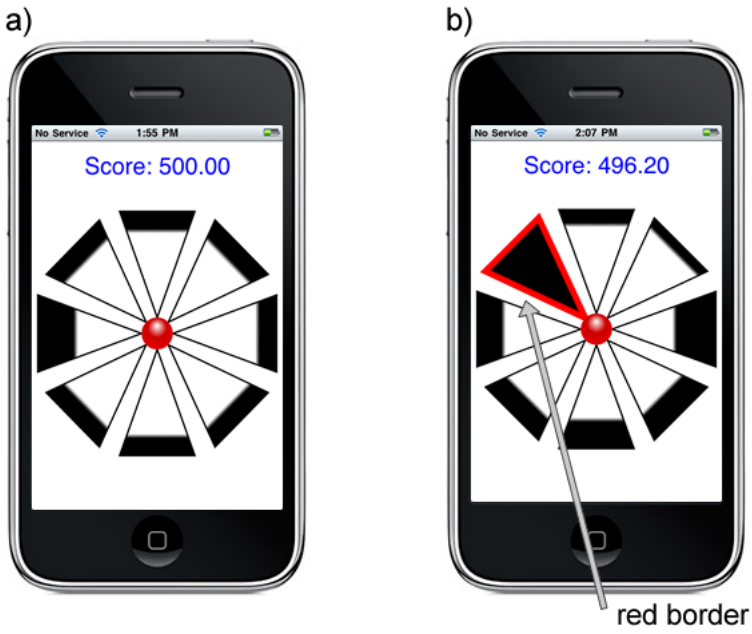

Figure 6.1. Screenshots of the game. a) Eight wedges fill with the black color from the bottom of the wedge towards its top; b) When a wedge becomes completely filled, the system shows the red highlight and the players start to lose the score.

RQ-4. *How does spatial tactile feedback affect the performance and communication of patterns in remote collaboration on mobile touch-screen devices?*

6.2 System

Game Design

The goal of the system is to measure the effectiveness of different feedback channels in spatial coordination between remote users. I decided to build a game involving abstracted spatial coordination tasks (*e.g.*, tapping the same region of the screen). This abstraction allows me to examine the effects of feedback channels while minimizing the effects of factors

which could affect spatial coordination between users. At the same time, it maintains the generalizability of the study results to systems including tightly-coupled spatial coordination within the working space (*e.g.*, a collaborative drawing or a system supporting the motor skill development of visually impaired users on gestures or handwriting (Plimmer *et al.*, 2008)).

Figure 6.1 shows a screenshot of the game used in my study. The game screen is shared on the two mobile touch-screen devices when users play the game. The screen shows eight interactive wedges and the score. The player pair initially start the game with 500 points. When the game starts, the system gradually fills the wedge (with black pixels) from the outside edge towards the center. The game fills each wedge at a different rate. When a wedge becomes completely filled, the system highlights that wedge with a red border, and the pair starts to lose points. One point is deducted for each completely filled wedge per second, and players do not have any way to gain points.

Success in this game depends on effective spatial coordination between the players. The objective of this game from the player's perspective is to keep as high a score as possible, given five minutes of play. To prevent a wedge from becoming completely filled, both must touch the same wedge, and at least one of them must perform a scrubbing gesture on that wedge. The game calculates the amount of black pixels to remove based on the length of the user's scrubbing gesture. The system ignores all scrubbing gestures when both game partners are not touching the same wedge. In this manner, the game requires frequent, tightly-coupled coordination between the players.

The game always provides the user with feedback about whether the touch screen has registered her touch by highlighting the selected wedge with a green border. Additionally, the game uses visual or tactile feedback to inform the user which wedge their game partner touches.

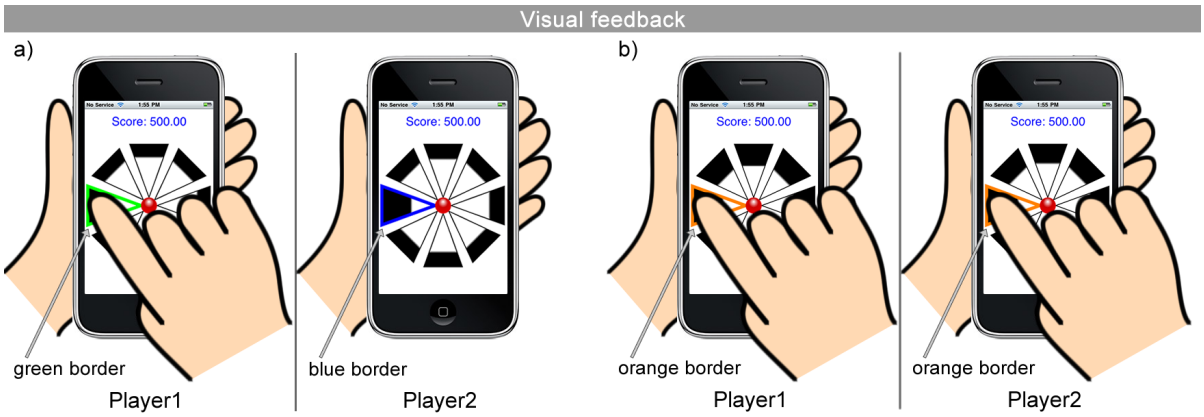

Figure 6.2. Visual feedback used in the game. a) When one of the players touches a wedge, it is highlighted with green, and the system also provides the blue highlight on the contacted wedge (in this example, the left wedge) in the game partner's screen; b) When both of the players touch the same wedge, the highlight turns orange.

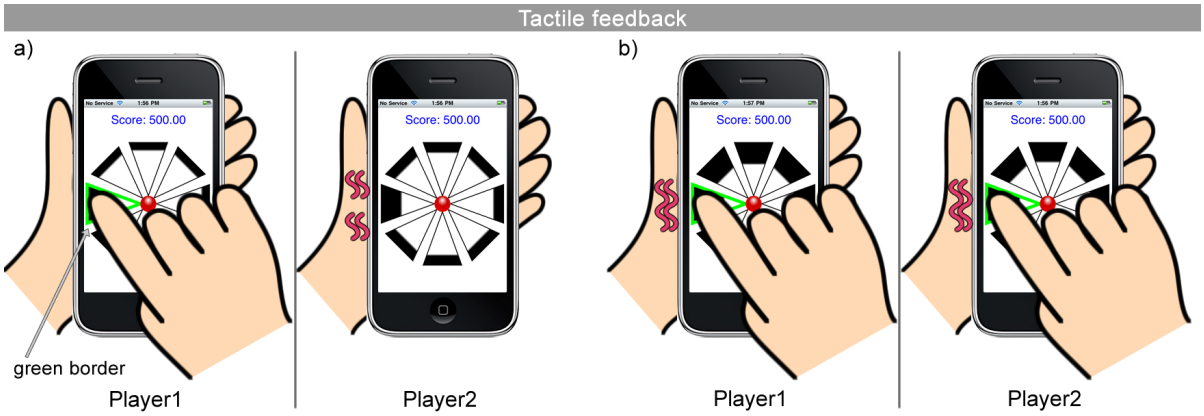

Figure 6.3. Spatial tactile feedback used in the game. a) When one of the players touches a wedge, it is highlighted with green, and the system provide discontinuous vibration from the vibration motor associated with the contacted wedge; b) When both of the players touch the same wedge, the vibration becomes continuous.

Visual Feedback

Figure 6.2 shows how my system provides visual feedback. When one player touches a wedge, that wedge is highlighted by a blue border on the game partner's screen (Figure 6.2a). When both players touch the same wedge, its border color turns orange on both devices (Figure 6.2b).

Tactile Feedback

Figure 6.3 shows how my system provides tactile feedback. In this part of the work, my prototype with nine vibration motors was used (Section 3.3). When one player touches a wedge, the game partner's device activates the vibration motor associated with that wedge to generate localized discontinuous vibration by turning it on/off every 200 msec (Figure 6.3a). When both players touch the same wedge, both devices generate continuous vibration with the motor associated with the contacted wedge (Figure 6.3b). The motor positioned at the center was not used in this study.

System Architecture

The mobile devices (iPhones in this system) connect to a server machine through wireless communication, and they report all touch events on the screens to this server. The server computes how fast each wedge gets filled based on the predefined game pattern. When the players are scrubbing the same wedge, it also computes how much of the fill needs to be removed based on their scrubbing. This information is sent back to each mobile device, which renders the game screen. When visual or tactile feedback is enabled, the server also sends each device the information about which wedge the game partner is touching. The mobile device then provides visual or tactile feedback as I described above.

	Feedback		
Audio	No	Visual	Tactile
Yes	*Audio-only*	*Visual-Audio*	*Visual-Tactile*
No	N/A	*Visual-NoAudio*	*Visual-NoTactile*

Table 6.1. The five conditions tested in Study 1.

6.3 Study1: Effects of Each Feedback Type

I conducted two laboratory studies to examine how spatial tactile feedback affects spatial coordination on mobile devices. In the first study, I focused on investigating the individual effects of visual and tactile feedback.

Conditions

I designed three types of feedback conditions: *No feedback* (the system did not provide any visual or tactile feedback about the game partner); *Visual feedback* (the system provided blue or orange visual highlight in response to the game partner's contact); and *Tactile feedback* (the system provided discontinuous or continuous vibration from the motor associated with the wedge contacted by the game partner's contact).

For each type of feedback, I also controlled the availability of the audio channel which allowed the participants to talk with each other: *Audio* (the audio channel was provided and the participants could talk with each other); and *NoAudio* (the audio channel was disabled). Table 6.1 shows the five conditions I studied. I excluded the condition of *No Feedback* without the audio channel because it does not allow the participants to communicate in any way and they would not be able to collaborate strategically. My main interest was to understand the effects of feedback on audio-enabled conditions as I was motivated by collaborative scenarios I can

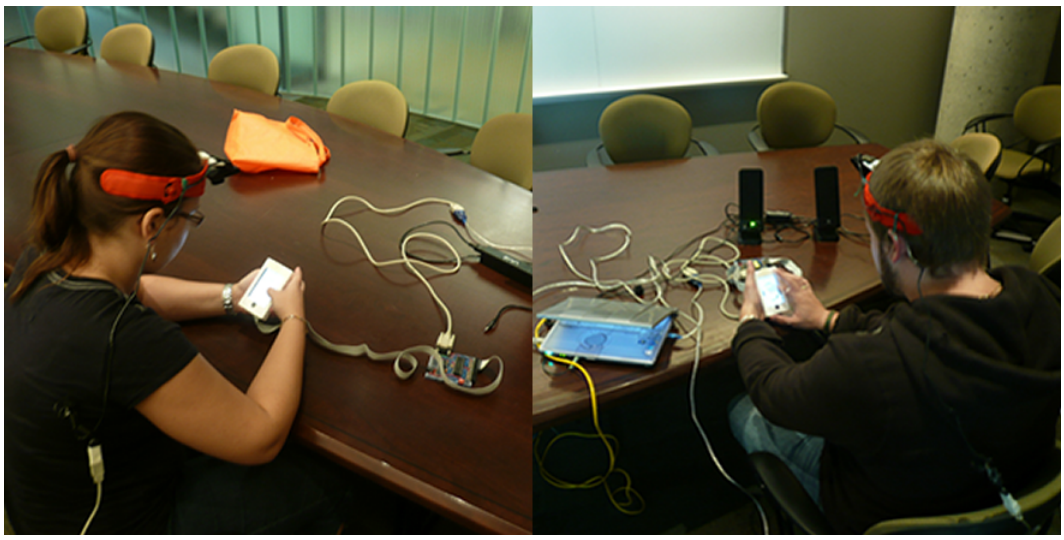

Figure 6.4. Experimental setup. Two participants were brought to different rooms. In the audio-enabled conditions, participants were allowed to talk with each other through microphones. They were asked to wear a headband with a Web camera to record their playing.

see frequently; however, this inclusion of audio-disabled conditions allows me to understand how well visual or tactile feedback could convey spatial information by comparing against the *Audio-only* condition.

The study was a within-subject design where each pair received exposure to all experimental levels, and presentation order of the system feedback was counter-balanced across the participants. The order of the audio channel availability was fixed to *Audio* followed by *NoAudio* within *Visual* and *Tactile* feedback conditions. All the conditions used a predefined script which specified changes in the rate at which the game would fill each wedge. Although the same script was used, the game rotated the script randomly for each condition. In this manner, I controlled the difficulty of the game to be the same across all conditions.

Procedure

Twenty-four participants were recruited in teams of two persons. Upon arrival to the laboratory, they were given an explanation about the system and the game. After this explanation, I separated the participants from each other into different rooms. I gave each participant a device to use during the study as well as a microphone and speaker so that they could communicate with each other during the experiment. I asked each participant to wear a headband with a mounted Web camera (Figure 6.4). I adjusted this Web camera so that it could record all interactions that each participant performed on the mobile devices. Participants then had a practice session to become comfortable with the system before starting the experiment.

All touch events generated by each participant were recorded. The state of all wedges and the point score were logged every 100 msec. The system also audio-recorded all conversations and stored videos recorded by the Web cameras attached to their headbands for analysis. At the end of the experiment, I conducted a short semi-structured interview to explore the difficulties with collaborating in each condition and their reasons.

Participants

The twelve pairs of participants (PA1–PA12) were recruited. They were between the ages of 18 to 39, and had a variety of backgrounds (such as students, teachers, engineers, and business persons). Three of the pairs were both male, one of them was both female, and the rest consisted of one male and one female. All the pairs knew each other before participating in this study. The study lasted approximately 70 minutes. Each participant was compensated with $30 cash after the study.

Theme	Description / Example
Targeting	Specifying an immediate target *"Seven", "Go to top"*
Confirmation	Agreeing with the previous *Targeting* or *Planning* *"OK", "Yeah"*
Clarification	Asking a question about the previous *Targeting* or *Planning* *"What?", "One?"*
Strategy switching	Changing play strategies *"Go clockwise", "Go this direction"*
Awareness	Asking or providing the status of a player *"I'm doing it", "Where are you?"*
Planning	Specifying a series of targets *"Seven, nine, and one"*
Prompting	Signaling a timing to move to the next target *"OK, go ahead"*

Table 6.2. The coding scheme for utterances observed during the experiment.

Utterance Analysis

All the conversations that the participants had during the experiment were transcribed with timestamps as faithfully as possible. The second rater and I conducted open coding of the quotes to identify seven themes pertaining to coordination, and developed the coding scheme as shown in Table 6.2. Another external coder and I independently categorized the recorded utterances along the scheme; they achieved high inter-rater reliability for every theme (higher than 95% agreement and Cohen's $\kappa > 0.8$).

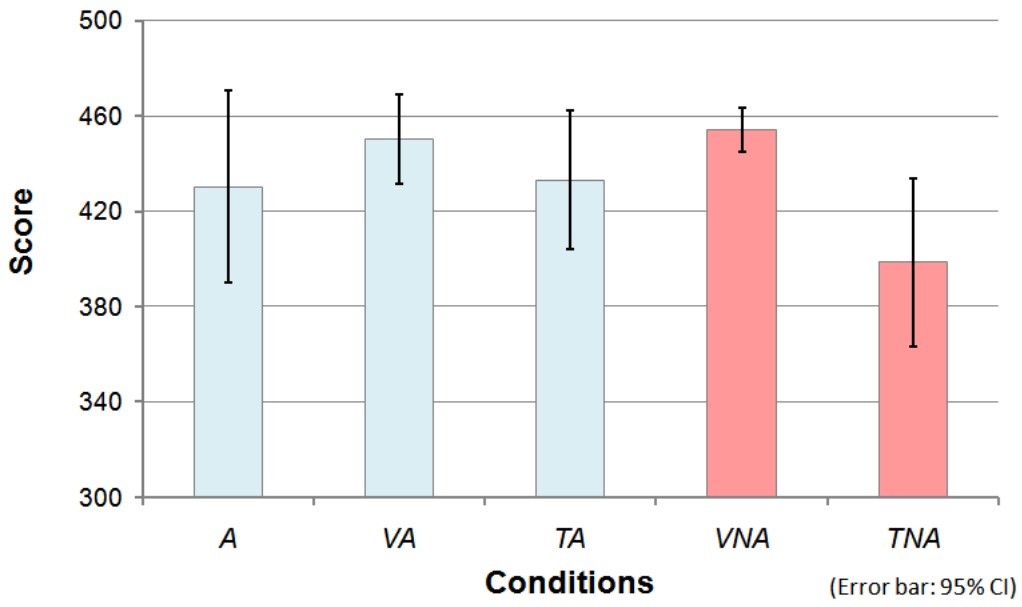

Figure 6.5. The average scores for the conditions tested in Experiment 1.

6.4 Study1 Results

Score

All pairs played the game for the full five minutes in each condition. Therefore, I was able to remove playing time from my analysis and focus on the performance scores. Figure 6.5 shows the average score for each condition. A Mauchly's test did not reveal a violation of sphericity, and therefore permits the direct interpretation of the ANOVA F-test results. A one-way repeated-measure ANOVA revealed a significant difference in the scores by condition ($F_{(4,44)} = 2.88$, $p < .05$, the effect size $\eta_p^2 = 0.21$). The post-hoc pairwise comparison with Tukey's HSD revealed that scores for the two visual conditions (*Visual-Audio* and *Visual-NoAudio*) were significantly higher than the *Tactile-NoAudio* condition ($p < .05$). The other conditions were not found to be significantly different.

Theme	*Audio-only*	*Visual-Audio*	*Tactile-Audio*
Targeting	*78.4* (16.0)	*37.8* (37.1)	*55.4* (36.0)
Confirmation	*12.2* (12.1)	*11.3* (12.5)	*14.0* (15.2)
Clarification	*1.4* (1.6)	*0.7* (0.8)	*0.7* (0.9)
Strategy switching	*2.1* (1.8)	*1.6* (1.8)	*1.3* (2.3)
Awareness	*3.3* (4.9)	*2.6* (2.3)	*2.5* (2.7)
Planning	*3.4* (5.6)	*1.8* (4.6)	*2.2* (6.3)
Prompting	*0.8* (0.9)	*2.4* (5.1)	*0.7* (1.4)
Total	*101.5* (24.0)	*58.3* (54.6)	*77.1* (54.0)

Mean (SD)

Table 6.3. The number of utterances for the three audio-available conditions tested in Study 1.

Utterances

In addition to the performance scores, it is telling to examine the communication processes that took place in the different feedback conditions, and understand their influence on performance. By doing so, I can gain insight into the type of coordinating information that each feedback mechanism provides. I classified the types of spoken content using the coding scheme described in Table 6.2. The results are shown in Table 6.3 which reports the average number of utterances by content type for the three audio-enabled conditions.

To examine how the number of these utterances affected the game score in each condition, I used a mixed effect linear regression model where *Condition*, *Content Theme* (the seven themes shown in Table 6.2), and the *Condition* × *Content Theme* interactions were included as independent variables. To avoid multicollinearity, I first examined the correlations among the possible independent variables, and chose those which were not strongly correlated with each

other. As a result, the variables for *Content Themes* were the utterance counts of *Targeting*, *Clarification*, *Planning*, *Prompting*, along with their interaction terms with *Condition*. Since the pairs participated in all conditions, observations were not independent and were therefore modeled as a random effect. The resulting model fit was moderate ($R^2 = .64$, *Adj-*$R^2 = .40$).

This model controls for the types of utterances that were generated in addition to the feedback provided by the condition, and describes their influence on game score. An ANOVA test can be used to examine fixed effects. However, the original ANOVA test does not necessarily produce reliable results when the sample size is small. I used an ANOVA test with the Kenward-Roger approximation (Kenward & Roger, 1997) for the analysis below. This approximation corrects the degrees of freedom which gives the closest match of an F-distribution to the distribution of the test statistics.

Controlling for language, I see a significant effect of *Condition* ($F_{(2,19.19)} = 3.72$, $p < .05$). I found that the game score increased from *Audio-only* to *Tactile-Audio* to *Visual-Audio*.

Of the content types examined, there is a main effect of *Targeting* content ($F_{(1,21)} = 8.37$, $p < .01$), controlling for condition. Higher-level interactions revealed an significant interaction effect of *Condition* and the number of *Targeting* utterances ($F_{(2,19.96)} = 4.63$, $p < .05$). Further examination of the interaction effects revealed that an increase in the production of *Targeting* utterances helped to improved game score in the *Audio-only* condition, but did not provide a similar improvement on *Visual-Audio* and *Tactile-Audio*. In other words, when the pairs did not have visual or tactile feedback, they had to compensate by increasing their production of spatial information regarding the targets.

I also found an significant interaction effect of *Condition* and the number of *Prompting* utterances ($F_{(2,20.32)} = 4.22$, $p < .05$). The increase of *Prompting* utterances was associated with a lower score in the *Tactile-Audio* condition, but such effects were not found in the other conditions.

Strategies

I analyzed the naming schema that the participants used for specifying the wedges on the screen. Understanding their naming schema is important because they are a part of the participants' strategies for effective coordination. I observed the three following naming schema:

- *Index*: Language using numbers the participants agreed. For instance, "one" meant the top wedge, "three" meant the right wedge, and "five" meant the bottom wedge.
- *Clock*: Language based on the clock metaphor. For example, "twelve o'clock" meant the top wedge, and "six o'clock" meant the bottom wedge. The use of numbers is similar to *Index*, but the wedges that the numbers correspond with here differ from *Index*.
- *Direction*: Language based on the direction or orientation. "North" or "top" meant the top wedge, and "south" or "bottom" meant the bottom wedge.

I observed that for most of the pairs, one participant would typically assume the responsibility for deciding the target and the other person would follow her. However, the participants often stop using this strategy when many wedges become almost completely filled. Both participants would then actively communicate about the next target. In terms of scrubbing, one person (who generally decides the target) would touch and continue to press the target while the other player scrubbed that wedge.

Collaboration in the Audio-only Condition

Participants used targeting utterances in the *Audio-only* condition most frequently as shown in Table 6.3. Because the system did not provide any feedback about the game partner's actions, participants often failed to figure out their game partner's location. There were different reasons

PA10-1	Time	PA10-2
“Nine.”	46	“Seven.”
Moving to the left wedge.		Moving to the bottom-left wedge.
“What are you doing?” 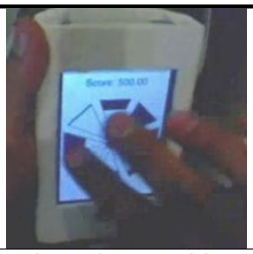	48	
Scrubbing the left wedge, but nothing occurred.		Holding the bottom-left wedge.
	49	“Seven, seven.”
Scrubbing the left wedge.		Holding the bottom-left wedge.
“OK”	50	
Moving to the bottom-left wedge.		Holding the bottom-left wedge.
	51	
Scrubbing the bottom-left wedge.		Holding the bottom-left wedge.

Figure 6.6. An example of collaboration observed in the *Audio-only* condition. This pair (PA10) mis-communicated about moving to the next target. PA10-1 then explicitly asked where PA10-2 was holding and re-coordinated the position.

for this. For example, both players may specify different wedges as the next target, and then fail to recognize the need to negotiate or clarify which wedge both should touch (*i.e.*, once a participant expresses the next wedge she will target, she assumes the partner will follow her verbal instruction even though both in the pair spoke at the same time). Figure 6.6 illustrates a case in which both players specified two different wedges as the next target, but did not make any clarification because both were focused on their own targets. This resulted in the pair falling out of synch and required explicit re-coordination.

Participants sometimes slipped and specified a wrong target, which led them to touch different wedges. In the post-experimental interview, one participant explained that the Audio-only condition was difficult because of possible slips in how to refer to the wedges.

> *"Probably the most difficult condition was audio-only… I keep getting confused by 'top-right' or 'top-left'."* [PA3-1]

Although participants performed well in the *Audio-only* condition, some participants commented that explicit audio communication increased the workload involved with playing the game compared to the *Visual* and *Tactile* conditions. As shown in Table 6.3, no visual or tactile feedback forced the participants to use more utterances. Although most of the utterances I observed were short, participants felt that frequent conversations often prevented them from collaborating efficiently in such a tightly-coupled coordination task.

Collaboration in the Visual Conditions

I observed that the participants generally used fewer utterances in the *Visual-Audio* condition than the *Audio-only* condition. As the participants mostly made utterances for targeting and confirmation in the *Audio-only* condition and the game provided visual feedback of where the game partner was touching on the screen; thus, they did not need to perform explicit coordination using voice in some cases.

PA7-1	Time	PA7-2
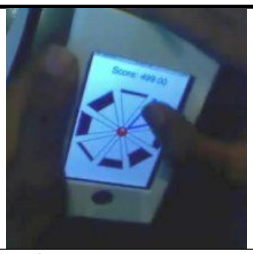	118	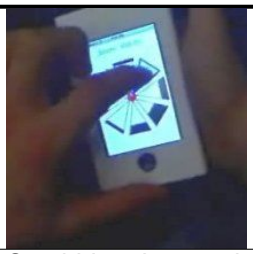
Holding the top-right wedge.		Scrubbing the top-right wedge. The filling was mostly removed.
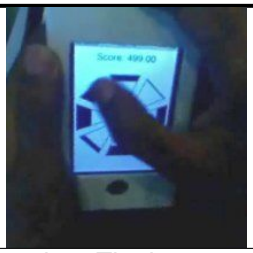	119	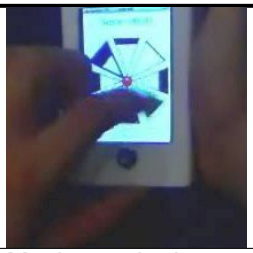
Moving to the top-left wedge. The bottom-right wedge is occluded.		Moving to the bottom-right wedge.
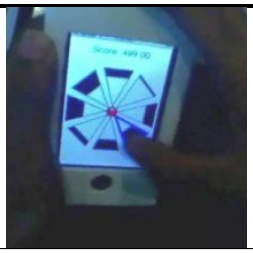	121	
Noticing feedback and moving to the bottom-right wedge.		Moving to the top-right based on visual feedback.
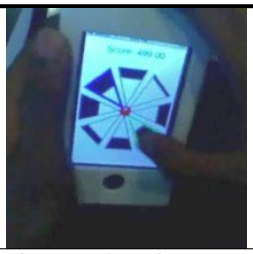	123	
Noticing that PA7-2 had moved to the top-left wedge.		Scrubbing the top-left wedge, but producing no effect.
“Five” 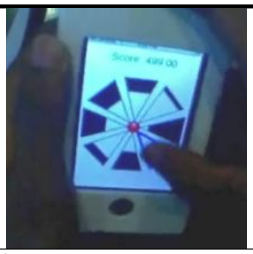	124	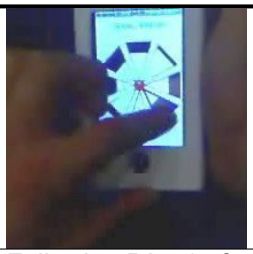
Explicitly mentioning the target.		Following PA7-1 after his instruction.

Figure 6.7. An example of collaboration observed in the *Visual-Audio* condition. PA7-1 did not notice PA7-2’s move because the visual feedback was occluded by his thumb. This caused another mis-coordination, and PA7-1 had to tell PA7-2 to come to the bottom-right wedge.

In both *Visual* conditions, the participant's hand sometimes occluded a significant portion of the screen. I found that the participants often failed to spot the visual feedback indicating which wedge their partner was touching. Figure 6.7 shows one of the instances in which the occlusion impeded collaboration. In this example, while PA7-1 moved to the top-left wedge, PA7-2 moved to the bottom-right wedge. However, PA7-1 did not notice the game partner's move because the bottom-right wedge initially was occluded by the thumb. PA7-1 then noticed the visual feedback, and tried to follow the partner. But PA7-2 also noticed that PA7-1 was on the top-left wedge, and tried to follow him. As a result, another mis-coordination happened, and P7-1 had to specify the target verbally.

Collaboration in the Tactile Conditions

Similar to the *Visual* conditions, participants generally used fewer utterances compared to the *Audio-only* condition. The tactile feedback helped the participants identify which wedge their game partner was touching. However, some participants, particularly those who had small hands, commented that they often had difficulty correctly identifying the location of the vibration. However, participants liked having a separate channel for knowing their game partner's location.

I found that the participants generally were able to use spatial tactile feedback to communicate location with their partner and often did not need to confirm it explicitly. Figure 6.8 shows an example of interactions that I observed in the *Tactile-NoAudio* condition. After the pair finished scrubbing the bottom wedge, PA3-1 moved to the bottom-left wedge whereas PA3-2 moved to the top-left wedge. PA3-2 immediately noticed the discontinuous vibration coming from the bottom-left. This caused PA3-2 to defer his location; PA3-2 then moved to PA3-1's location, and scrubbed the bottom-left wedge. Because PA3-1 was able to perceive that PA3-2's previous location through the tactile communication channel, after they finished the bottom-left wedge, they moved to the top-left wedge, which was PA3-2's previous target.

PA3-1	Time	PA3-2
	158	
Moving to the bottom-right wedge.		Moving to the top-right wedge.
	159	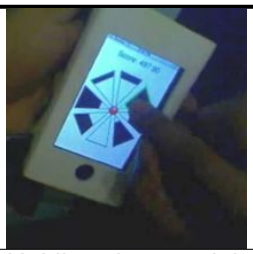
Scrubbing the bottom-right wedge, but the tactile feedback happened at the top-left.		Holding the top-right wedge, but noticed tactile feedback coming from bottom-left.
	160	
Still scrubbing the bottom-right wedge.		Moving to the bottom-left wedge.
	161	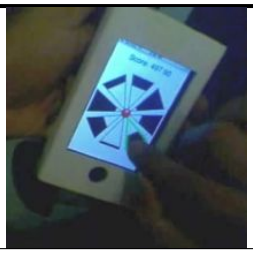
Scrubbing the bottom-right wedge. The black filling started to disappear.		Holding the bottom-right wedge.
	163	
Moving to the top-right wedge which PA3-2 was touching before.		Moving to the top- wedge.

Figure 6.8. An example of collaboration observed in the *Tactile-NoAudio* condition. This pair (PA3) made different moves initially, but PA3-2 followed PA3-1 based on the tactile feedback. After the bottom-left wedge, both participants moved to the top-left wedge which PA3-2 wanted to work on before.

Study1 Summary

I found that a system that uses visual feedback generally supports coordination better than one that uses tactile feedback or the system without any feedback. My analysis also revealed that the number of *Targeting* utterances affect game performance only in the *Audio-only* condition. The results imply that an additional feedback channel, either visual or tactile feedback, could help the participants perform spatial coordination without explicit audio communication.

The results also highlight different benefits of visual and tactile feedback in spatial coordination, and suggest that these two feedback channels could complement each other well. Visual feedback is beneficial to provide precise spatial information. This was highlighted clearly in the *NoAudio-Visual* condition. The results also suggest that tactile feedback can address some occlusion issues.

They motivated me to examine if combining visual and tactile feedback can improve spatial coordination because the combined feedback offers benefits provided by each feedback channel. My next user study, thus, examines the performance of spatial coordination with the combined feedback.

6.5 Study2: Effects of Combined Feedback

My second user study focused on comparing user collaboration with a system combining both visual and tactile feedback against systems using each feedback channel. I used the same system and experimental procedure as the first user study. I included three audio-enabled feedback conditions: *Visual*, *Tactile*, and *Visual+Tactile* (both visual and spatial tactile feedback were provided), and the presentation order of these conditions were counter-balanced across participant pairs. I recruited another twelve pairs of participants for this study (PB1–PB12). Their demographics were similar to the ones in my first user study. This study took 50 minutes

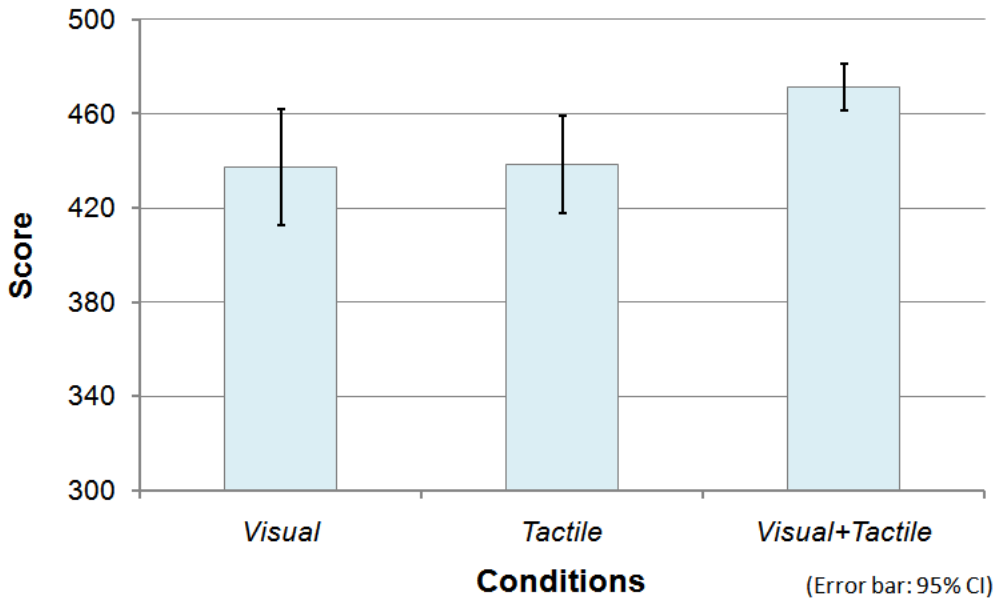

Figure 6.9. The average scores for the conditions tested in Study 2.

on average, and each participant was compensated with $30 after the study.

6.6 Study2 Results

Score

Figure 6.9 shows the game scores across the three feedback conditions. Unpaired Welch's t-tests did not show any significant difference in the scores of *Visual* and *Tactile* between my first and second study ($t_{(21.0)} = 0.81, p = 0.21$ for *Visual*; and $t_{(19.4)} = 0.28, p = 0.39$ for *Tactile*). A Mauchly's test of sphericity did not reveal a violation, so I once again report the results of the ANOVA F-tests. A one-way repeated-measure ANOVA revealed a significant difference in the scores by condition ($F_{(2,22)} = 6.01, p < .01, \eta_p^2 = 0.35$). Tukey's HSD revealed that the scores in *Visual+Tactile* were significantly higher than the scores in Visual and Tactile ($p's < .05$).

Theme	*Visual*	*Tactile*	*Visual+Tactile*
Targeting	*72.4* (26.0)	*77.3* (30.0)	*86.4* (12.2)
Confirmation	*10.5* (9.9)	*7.5* (6.0)	*5.8* (3.8)
Clarification	*0.9* (1.3)	*1.6* (2.0)	*1.8* (2.3)
Strategy switching	*1.8* (2.2)	*1.0* (1.6)	*0.9* (1.4)
Awareness	*4.3* (4.9)	*2.1* (2.6)	*1.6* (2.0)
Planning	*1.0* (1.2)	*1.7* (2.3)	*1.6* (2.1)
Prompting	*0.4* (1.1)	*0.5* (1.0)	*0.1* (0.3)
Total	*91.3* (27.0)	*91.8* (31.0)	*98.3* (11.0)

Mean (SD)

Table 6.4. The number of utterances for the three audio-available conditions tested in Study 2.

Utterances

Similar to the first study, I examined the communication processes that took place to better understand the coordinating role of each feedback condition. I used the same coding scheme presented in Table 6.2 and the average numbers of utterances by content type for the second study are shown in Table 6.4. I then compared differences in the number of utterances between the first and second study. As a result, unpaired Welch's t-tests did not show any significant difference in the scores of Visual and Tactile between the first and second study ($t_{(16.1)} = 1.69$, $p = .06$ for *Visual*; and $t_{(17.6)} = 0.80$, $p = .22$ for *Tactile*).

I analyzed the data using the same random effects linear regression model described in the first study with one exception being that interactions were removed from the model due to the fact that there were no significant higher order interactions. The resulting model fit was relatively high ($R^2 = .70$, *Adj-*$R^2 = .63$).

Controlling for language, I still see a significant effect of *Condition* ($F_{(2,21.37)} = 4.93$, $p < .05$) where performance for the *Visual+Tactile* condition was better than both the *Visual* and *Tactile* conditions. Of the content types examined, there is a positive main effect of *Targeting* content whereby increases in the use of targeting comments was associated with higher scores ($F_{(1,26.33)} = 14.0$, $p < .01$), controlling for condition. There was also a marginal negative effect of *Clarifications* where an increase in the use of clarifications was associated with lower scores ($F_{(1,28.93)} = 14.0$, $p < .10$), controlling for condition. While these findings suggest a link between the content of the discussions and the performance of the pairs, I found no strong evidence of differentiated influence of various discourse strategies in the conditions examined in the second study. This suggests that there may not be a difference in the need for additional information across the visual and tactile channels. Alternatively, it could be that my discourse coding scheme was not sensitive to the coordinating differences that are afforded by the visual and tactile channels; thus, I further analyzed the user feedback I gained during the post-experimental interview for evidence of differential use.

User Feedback

Most participants agreed that visual feedback was easy to understand and showed the location of a game partner accurately. But four pairs explicitly mentioned that occlusion by hands was a problem. In contrast, participants expressed different opinions on tactile feedback and used it in different ways. Four pairs mentioned that they needed more effort in associating tactile feedback to a particular wedge than visual feedback. However, tactile feedback could subtly guide a player's attention to the area of interest as one of the participants commented:

> *"The good thing [about vibration] for me was it was a more subconscious cue; wasn't something I had to pay attention to. But with color, I've got to pay attention to where the color is going, I had to process it. But vibration subconsciously pulled attention to the area."* [PB2-2]

Two pairs commented that tactile feedback offered a separate channel to maintain awareness about a game partner without disrupting the visual information space.

> *"The vibration is a much better safety net than visual. I think because I could feel that or maybe because so much was happening on the screen already. Tactile touch didn't add something extra on the screen... Changing colors was just a distraction, versus vibration wasn't another visual distraction."* [PB10-1]

Participants overall liked the *Visual+Tactile* condition, and often preferred it over the other conditions. One participant explained to us that she used tactile feedback as a redundant cue to ensure that she and her partner were contacting the same wedge.

> *"Vibration is confirmation of if you are doing the right thing. I'm hearing where to go, and vibration confirms me that we are going to the same spot."* [PB2-1]

Study2 Summary

My second user study revealed that a system using both visual and tactile feedback outperformed systems using only either type of feedback when verbal communication is available. Qualitative results also support the notion that visual and tactile feedback can complement each other to support users' spatial coordination.

6.7 Discussions

The results from the two user studies highlighted how the combined visual and tactile feedback improved game performance. The occlusion caused by the participant's hand often impeded smooth collaboration in the *Visual* conditions. Occlusion is a well-known problem in touch-screen devices, and different approaches have been developed to address this issue. Vogel

& Balakrishnan (2010) demonstrated a user interface which changes the locations of the objects depending on the position of the hand or arm over the screen. This technique could solve some problems caused by the occlusion observed in my study. But, rearranging the objects might introduce additional complexity on spatial coordination. Tactile feedback directly can mitigate some occlusion problems, and provide awareness for coordination particularly when the visual components are complex (*e.g.*, online network multiplayer games).

The system had delays in tactile feedback during my experiments due to my hardware and system limitation. Gergle *et al.* (2006) studied the effects of delays in visual feedback for collaborative tasks. They found that in rapidly changing dynamic environments, delays on the order of 200 msec can cause performance deficits in visual piece arrangement tasks. The study by Jay *et al.* (2007) also shows that even a very small delay can impact the collaboration when strict spatial and temporal coordination is necessary.

The tasks studied in my experiment include higher temporal demand (*i.e.*, the system fills wedges constantly) than those in the study of Gergle *et al.* (2006); thus, even a small delay might have caused a significant impact. This may be one reason why *Tactile-NoAudio* was the weakest condition in terms of the performance score because participants were unable to explicitly coordinate through the audio channel to compensate the delay in tactile feedback. However, as Jay *et al.* (2007) discussed, supporting strict spatial and temporal coordination is challenging in both visual and tactile feedback, and thus future research is necessary on how to overcome the delay on the feedback.

Finally, the output resolution of spatial tactile feedback is a limitation imposed by my hardware and game design. Israr & Poupyrev (2011) have developed a method to create a virtual vibration point where no physical vibration motors exist (*e.g.*, a mid-point between the two vibration motors). With their algorithm, a system would have a very fine output resolution, enabling the design of combined feedback to support more precise spatial coordination. My results highlight the different advantages for each feedback channel and benefits for combined

feedback in spatial coordination, and can generalize to future collaborative handheld systems with a higher output resolution of combined feedback than my system.

6.8 Summary

I investigated how visual and tactile feedback affects synchronous collaborative tasks on mobile devices. In particular, I examined how users would perform spatial coordination in a shared workspace. With a collaborative game on mobile touch-screens, I investigated the following research question:

RQ-4. *How does spatial tactile feedback affect the performance and communication of patterns in remote collaboration on mobile touch-screen devices?*

The results highlight the following findings in the context of collaborative handheld systems:

- Visual feedback can provide precise spatial information about collaborators, but can hamper collaboration when feedback is occluded and sometimes distracts the user's attention.
- Spatial tactile feedback can provide spatial information about collaborators as well, but improvements are necessary to convey precise spatial information. It can also reduce the overload of information in visual space and can gently guide the user's attention to an area of interest.
- Visual and spatial tactile feedback can complement each other, and systems with both feedback channels can offer better spatial coordination support than systems using only either of feedback.

Chapter 7

Conclusions

As I described in the introduction, the main contribution of this dissertation can be stated as follows:

> *Spatial tactile feedback produced by multiple vibration motors enhances the expressiveness of tactile feedback on mobile touch-screen devices, and enables the development of user interfaces which require low or no visual demand.*

The chapters of this dissertation cover my design of spatial tactile feedback hardware and investigation on effects of spatial tactile feedback on mobile touch-screen devices.

7.1 Summary

In Chapter 1, I introduce the area of tactile feedback on mobile touch-screen devices. User interfaces found in current mobile touch-screen devices are visually demanding. This causes various interface problems. For instance, the user may be interrupted to view the screen of a mobile touch-screen device while she is walking. This type of mobile devices is not generally accessible to visually impaired users. Tactile feedback can mitigate these issues,

but the expressiveness of tactile feedback in mobile devices is still limited. I also describe the following research questions which I explored in this dissertation.

RQ-1. *How fast and accurately can users recognize various spatial vibration patterns generated by multiple vibration motors?*

RQ-2. *How fast and accurately can users perform eyes-free interaction with the support of spatial tactile feedback?*

RQ-3. *What level of accuracy and detail can visually-impaired users achieve when they try to construct mental maps of spatial relationships between multiple objects with the support of spatial tactile feedback?*

RQ-4. *How does spatial tactile feedback affect the performance and communication of patterns in remote collaboration on mobile touch-screen devices?*

In Chapter 2, I present my literature survey on tactile feedback in the field of physiology, psychology, hardware architecture, materials, and HCI. My survey first highlights that humans have high sensitivity to vibration at 200–300 Hz although a spatial resolution is not particularly high (estimated to be approximately 1 cm). It also shows that a DC vibration motor is appropriate hardware for mobile devices, and spatial patterns have been left under-explored in tactile feedback systems. My survey on user interfaces with tactile feedback discovers that it is still unclear how spatial tactile feedback can be incorporated into interfaces in mobile touch-screen devices and mitigate high visual demand issues.

In Chapter 3, I describe my hardware prototypes of spatial tactile feedback systems for mobile touch-screen devices. These prototypes employ multiple vibration motors on the backside of a mobile touch-screen device. These multiple vibration motors are embedded in different locations on the device. The hardware can produce various spatial vibration patterns, such as vibration at different points of the device and flows of vibration (*e.g.*, vibration flowing

from right to left). My hardware conveys these vibration patterns to the user's hand holding the device. I also explain the design of my special sleeve to accommodate vibration motors and fill a gap between the device and the user's palm.

In Chapter 4, I present my design of SemFeel, an eyes-free interface on mobile touch-screen devices. SemFeel provides the user with the semantic information about the object she is touching through multiple vibration motors embedded in the backside of the device. I then design eleven vibration patterns: five positional (single-point) patterns, four linear patterns, and two circular patterns. I conducted two experiments to examine the efficacy of SemFeel. My first experiment discovers that users can distinguish ten vibration patterns, including linear patterns and a clockwise circular pattern, at around 85–90% accuracy (lower accuracy for the counter-clockwise circular pattern). My second experiment with number entering tasks shows that SemFeel supports more accurate interactions in an eyes-free setting than systems that offer no tactile feedback or use only a single vibration motor.

In Chapter 5, I describe my development and evaluation of a system to support visually impaired users in understanding spatial relationships between multiple places, called SpaceSense. It offers step-by-step instructions of directions to the destination the user selects. Additionally, it provides high-level directional information to the destination and bookmarked places through spatial tactile feedback. This feedback is intended to help the user maintain the spatial relationships between these points. My user study shows that participants could build and maintain spatial relationships between four places on a map more accurately with SpaceSense compared to a system without spatial tactile feedback. Participants pointed specifically to having spatial tactile feedback as the contributing factor in successfully building and maintaining their mental map.

In Chapter 6, I present my investigation on how visual and tactile feedback affects collaboration on mobile devices. In particular, I examined how users would perform spatial coordination in a shared workspace when they need to perform an action collaboratively. I

describe my experiment, as well as a quantitative and qualitative analysis of the results from my two user studies. The results I gained through these studies highlight different benefits of each feedback channel in collaborative handheld systems. Visual feedback can provide precise spatial information for collaborators, but degrades collaboration when the feedback is occluded, and sometimes can distract the user's attention. Spatial tactile feedback can reduce the overload of information in visual space and gently guides the user's attention to an area of interest. My results also show that visual and tactile feedback can complement each other, and systems using both feedback channels can support better spatial coordination than systems using only one form of feedback.

7.2 Limitations

I make a number of assumptions to make my investigation tractable. I summarize them to clarify that my results may need to be further tested before applying them to other applications or contexts.

The number of vibration motors and their alignment described in Chapter 3 are limitations in my systems. Although adding more vibration motors is possible, note that a spatial resolution to vibration stimuli is estimated to be approximately 1 cm. Thus, applications that require very precise spatial patterns may not be covered by my hardware or similar hardware. Israr & Poupyrev (2011) developed an algorithm to increase the spatial resolution of an vibration motor array without necessarily adding vibration motors. Future work should investigate how a system can have ability to produce more precise and more complex spatial patterns than my systems.

In this work, I used a specific type of hardware (*i.e.*, DC vibration motors). Using different hardware could lead to different results. For instance, one can use piezo actuators or pin arrays to build similar hardware to what I developed. Different DC vibration motors with different size could affect the results reported here as well. There are several points which researchers and

developers need to be careful about when they develop new hardware based on the findings in this dissertation. The physical size of the vibration point is one important aspect of the hardware. For instance, if a larger motor is used than the motor in my hardware (10 mm diameter), the area of vibration also becomes larger, and it can affect the accuracy of user determining the point of the vibration. I also fixed the frequency parameter, and only used a few intensity levels. If a new system uses different frequencies or different intensity levels, the distinguishability of the vibration patterns could be different.

The form factor of the prototypes used in this work is another hardware limitation. The prototypes are thicker than common mobile touch-screen devices due to the addition of the hardware. Thus, they are not optimal for people who have small hands. In future prototypes, a deformable material can be used for the sleeve to offer an improved grip of the device while vibration motors still contact the user's palm and fingers firmly.

Related to the hardware limitation, the vibration patterns used in my systems are also limited. My main objective of this work is to investigate how the spatial and spatio-temporal parameters of vibrotactile feedback can be used to improve interface designs of mobile touch-screen devices. The first experiment with the SemFeel system (Section 4.4) examined 5 spatial and 6 spatio-temporal patterns, and found that users can distinguish them accurately (except the counter-clockwise pattern). However, these patterns are very basic and many other spatio-temporal patterns are untested. Using different durations or different levels of intensity might show different results. Nevertheless, my results indicate that the spatio-temporal parameter can be useful to provide the user with distinct patterns.

My experimental designs also includes several limitations. the tasks used in the experiment may have biases. Although I carefully chose these task to reflect realistic tasks, I abstracted some tasks to remove compounding factors. In Chapter 6, I designed a collaborative game to examine the effects of spatial tactile feedback on spatial coordination. This game design required participants to perform frequent spatial coordination, and as a result, I was able to

perform an in-depth analysis as ultimately presented in Section 6.4 and 6.6. However, this game is not directly related to a particular application, and my results may need careful interpretation when they are applied to realistic spatial coordination tasks. Furthermore, my experimental designs do not necessarily cover all representative settings in which users would often wish to interact with mobile devices. For example, my investigation with SemFeel did not include any mobile setting. A follow-up experiment including walking scenarios like Yatani & Truong (2009) should be conducted to further validate the effects of spatial tactile feedback in settings which have not been tested in this work.

Not all of the participants volunteered for my user studies in Chapter 4 and 6 may not represent the target user population well. I recruited participants so that they formed a diversity in age and background. However, there are still some user groups which my user studies did not cover fully. For example, a user study with the older user population could show different results or subjective preferences on spatial tactile feedback. The interface designs presented in this dissertation may also need to be changed for children because they have smaller hands the adults. Thus, additional precautions may be necessary when extending the interface designs and findings reported here to user groups which have not been tested.

My user study did not have a large number of participants. I determined the target number of participants which can counter-balance the presentation order of the important experimental factors and offers the generalizability of the results. I stopped recruiting participants when the results became relatively consistent across the participants or new phenomena stopped emerging. The results reported in this document represent my observations, and the number of recruited participants was sufficient to derive the conclusions on the effects of spatial tactile feedback. Thus, the results would be likely to hold in a study with a larger number of participants, but such a study could show additional results which might lead to new interface design implications.

My experiments did not include workload measurements. There are several ways to measure physical and mental workload objectively or subjectively. Objective workload measurement can be done by using physiological sensors, such as electrocardiograms, electroencephalography (Haapalainen *et al.*, 2010) and functional near-infrared sensors (Solovey *et al.*, 2009). Another way to measure workload is to use subject metrics, such as NASA TLX (Hart & Staveland, 1988). However, these measurements often require careful interpretation. Although objective workload measurements can offer objective interval data, they are subject to noise and the measured values can be affected by various external factors, such as the position of the sensor, or the user's mental status before starting the experiment (*e.g.*, a participant might be stressed already for some reasons when she starts the experiment). Subjective methods can be generally easy to deploy because of no need to add sensors or special equipment. However, self-reported values gained by these methods can degrade the reliability of the data (Annett, 2002). Furthermore, objective metrics are often considered as ordinal data, and it is not straightforward to compare one value and another when they are measured in different experiments. Furthermore, these subjective methods are designed as post-hoc tests (*e.g.*, tests participants would perform after they complete all the tasks in one condition). Thus, they may only show the overall experienced workload, and may not reveal which particular interactions have caused increase or decrease in workload.

These limitations of methods for measuring workload suggest that workload which could have been measured in my experiments by the methods above may not be reliable or well generalizable. Thus, I decided not to measure workload. However, it might be possible that interfaces with spatial tactile feedback could cause additional physical or mental workload. Although this dissertation offers some qualitative insight on workload through interviews with participants, it does not offer any quantitative results, and it is one of the future work to quantitatively identify the workload of the interfaces presented in this document.

7.3 Future Work

This work opens up a number of future research with spatial tactile feedback on mobile devices.

More Complex Vibration Patterns

In this dissertation, I focus on the spatial and spatio-temporal parameters in the vibrotactile feedback modality. As discussed above, the patterns explored in this work are limited although I demonstrated that even a limited set of spatial patterns is useful to enhance user experience on mobile touch-screen devices. To expand this set, there are three possible directions. These directions are not mutually exclusive; thus, researchers may combine two or all of the following directions.

Improving the Spatial Resolution

The hardware used in this work has at most nine vibration sources (a 3×3 grid), which limits the spatial output resolution. Israr & Poupyrev (2011) developed an algorithm to realize a higher spatial resolution than the number of physical vibration motors. Their method offers the user a sensation as if a vibration point exists where no physical vibration motors exist. With their method, hardware used in my work would be able to increase its spatial resolution without adding motors. However, their investigation is on the human back (12 vibration motors were attached to the back of a chair). Thus, further exploration is necessary to determine whether their method would perform similarly for the hand.

As I discussed in Section 2.1, I expect that users would have difficulty distinguishing two vibration sources located closer at 1 cm. Note that my hardware has a 2 cm gap in the vibration motor grid. Therefore, even if the method developed by Israr & Poupyrev (2011) is applicable to a hand, it would be challenging to make the output resolution go beyond 5×5.

Designing More Complex Spatial Patterns

This work investigated positional, directional, and circular patterns, but designing more complex spatial and spatio-temporal patterns is also possible and an interesting future research direction. Examples of these patterns are: diagonal lines (and their combination, like a cross), triangles, rectangles, arcs. These patterns can be useful to enhance user experience of applications, such as computer games Israr & Poupyrev (2011).

Another potential application is presenting visual information to visually impaired users. Through SpaceSense described in Chapter 5, spatial tactile feedback is useful to help visually impaired users learn and maintain spatial relationships between multiple places. Thus, the system offers high-level spatial information about a space; however, it still lacks the ability to inform the user about lower-level information which is also important for successful navigation, such as the space of a street. A richer set of spatial and spatio-temporal patterns than what I used in this work could help visually impaired users understand finer geographical information.

Offering a tactile modality to access visual representations of data is often called haptic data visualization. According to the literature survey conducted by Paneels & Roberts (2010), haptic data visualization systems have focused on presenting *Charts* (scientific or statistical graphs), *Maps* (choropleth maps or visual representation of real environments), *Signs* (icons or simple messages), *Networks* (structures of trees, hierarchies, paths and nets), and *Diagrams* (visual representations of process, phenomenon, or concept). Spatial patterns have been used mainly for *Signs*, such as Brewster & Brown (2004); Brown *et al.* (2006a); Hoggan *et al.* (2007). SemFeel described in Chapter 4 falls into the same category (*i.e.*, presenting semantic information about the object a user is touching). Understanding how spatial tactile feedback could support the user to understand other types of applications would be an interesting direction upon this work.

Combining With Other Vibrotactile Parameters

Combining other vibrotactile parameters discussed in Section 2.3 is another way to extend the set of vibration patterns. For example, SpaceSense used the intensity parameter to indicate the distance to a place from the user's current simulated location. As I discussed with a possible calendar application upon SemFeel (Figure 4.8), the spatial parameter can represent a temporal dimension (*e.g.*, morning, afternoon, and evening), and the intensity or rhythm parameter can represent the user's availability or the importance of events. Future work should deeply investigate what type of information each parameter would be appropriate to convey.

Navigation Systems for Visually Impaired Users

Expanding SpaceSense to support the navigation in actual environments is an interesting future research direction in applications for visually impaired users. Although the current SpaceSense prototype is not designed to support visually impaired users during their actual navigation of a space, integration with a GPS module and digital compass can extend its usage scenarios. The features provided by the current SpaceSense offer people with visual impairments geographical information before they navigate within a space. Once they start exploring the space, the mobile device can track the user's location and orientation based on GPS and compass data. If the user gets lost during the navigation, the device can indicate which direction they need to walk towards using spatial tactile feedback. The direction will also match the current user's orientation; thus, vibration from the top side of the device tells the user to walk forward.

Collaborative Systems for Visually Impaired Users

Chapter 6 investigated how spatial tactile feedback can support spatial coordination between remote users. Results indicate that although spatial tactile feedback may not necessarily perform as well as visual feedback, it can still offer awareness about the collaborator's contact

point on the screen. This implies that spatial tactile feedback could be used to support collaboration between visually impaired users and sighted users. For example, sighted users can demonstrate visually impaired users how to interact with interfaces on a mobile touch-screen device. The interaction will be conveyed through tactile feedback, and visually impaired users will able to feel how to interact with a device remotely.

7.4 Final Word

In this dissertation, I focus on the high visual demand issue in user interfaces on mobile touch-screen device. As mobile touch-screen devices have become pervasive, understanding interface problems on these devices and improving interactions have strong value. My approach to addressing the high visual demand issue presented here is to improve the expressiveness of tactile feedback. For this purpose, I developed spatial tactile feedback systems—interfaces with vibrotactile feedback using spatially-distributed vibration sources and encoding information in spatial patterns. Through my series of investigations presented in this dissertation, I validated the effects of spatial tactile feedback on user interfaces for mobile touch-screen devices. This work offers researchers and developers user interface designs with spatial tactile feedback, and encourages them to further improve tactile feedback in mobile devices.

Bibliography

Alexander, J., Marshall, M. T., & Subramanian, S. (2011). Adding haptic feedback to mobile tv. In *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*, CHI EA '11, (pp. 1975–1980). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1979742.1979899

Altinsoy, M., & Merchel, S. (2009). Audiotactile feedback design for touch screens. In M. Altinsoy, U. Jekosch, & S. Brewster (Eds.) *Proceedings of Haptic and Audio Interaction Design*, vol. 5763 of *Lecture Notes in Computer Science*, (pp. 136–144). Springer Berlin / Heidelberg.
URL http://dx.doi.org/10.1007/978-3-642-04076-4_15

Amemiya, T., & Sugiyama, H. (2009). Haptic handheld wayfinder with pseudo-attraction force for pedestrians with visual impairments. In *Proceedings of the 11th international ACM SIGACCESS conference on Computers and accessibility*, Assets '09, (pp. 107–114). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1639642.1639662

Annett, J. (2002). Subjective rating scales: science or art? *Ergonomics*, *45*, 966–987.

Azenkot, S., Prasain, S., Borning, A., Fortuna, E., Ladner, R. E., & Wobbrock, J. O. (2011). Enhancing independence and safety for blind and deaf-blind public transit riders. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, (pp. 3247–3256). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1978942.1979424

Bau, O., Poupyrev, I., Israr, A., & Harrison, C. (2010). Teslatouch: electrovibration for touch surfaces. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, (pp. 283–292). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1866029.1866074

Blasch, B. B., Wiener, W. R., & Welsh, R. L. (1997). *Foundations of orientation and mobility*. American Foundation for the Blind, 2nd ed.

Bolanowski Jr., S. J., Gescheider, G. A., Verrillo, R. T., & Checkosky, C. M. (1988). Four channels mediate the mechanical aspects of touch. *Journal of the Aco*, *84*(5), 1680–1694.

Brave, S., Ishii, H., & Dahley, A. (1998). Tangible interfaces for remote collaboration and communication. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, CSCW '98, (pp. 169–178). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/289444.289491`

Brewster, S., & Brown, L. M. (2004). Tactons: structured tactile messages for non-visual information display. In *Proceedings of the fifth conference on Australasian user interface - Volume 28*, AUIC '04, (pp. 15–23). Darlinghurst, Australia, Australia: Australian Computer Society, Inc.
URL `http://portal.acm.org/citation.cfm?id=976310.976313`

Brewster, S., & King, A. (2005). An investigation into the use of tactons to present progress information. In *In Proceedings of INTERACT 2005*, (pp. 6–17). Springer.

Brown, L. M., Brewster, S. A., & Purchase, H. C. (2005). A first investigation into the effectiveness of tactons. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, WHC '05, (pp. 167–176). Washington, DC, USA: IEEE Computer Society.
URL `http://dx.doi.org/10.1109/WHC.2005.6`

Brown, L. M., Brewster, S. A., & Purchase, H. C. (2006a). Multidimensional tactons for non-visual information presentation in mobile devices. In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, MobileHCI '06, (pp. 231–238). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1152215.1152265`

Brown, L. M., Brewster, S. A., & Purchase, H. C. (2006b). Tactile crescendos and sforzandos: applying musical techniques to tactile icon design. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI EA '06, (pp. 610–615). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1125451.1125578`

Brown, L. M., Sellen, A., Krishna, R., & Harper, R. (2009). Exploring the potential of audio-tactile messaging for remote interpersonal communication. In *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, (pp. 1527–1530).

New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1518701.1518934

Burtt, H. E. (1917). Tactual illusions of movements. *Journal of Experimental Psychology*, *2*(5), 371–385.

Chan, A., MacLean, K., & McGrenere, J. (2008). Designing haptic icons to support collaborative turn-taking. *Int. J. Hum.-Comput. Stud.*, *66*, 333–355.
URL http://portal.acm.org/citation.cfm?id=1360720.1360937

Chang, A., O'Modhrain, S., Jacob, R., Gunther, E., & Ishii, H. (2002). Comtouch: design of a vibrotactile communication device. In *Proceedings of the 4th conference on Designing interactive systems: processes, practices, methods, and techniques*, DIS '02, (pp. 312–320). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/778712.778755

Cholewiak, R. W., & Collins, A. A. (2003). Vibrotactile localization on the arm: effects of place, space and age. *Perception & Psychophysics*, *65*, 158–1077.

Cholewiak, R. W., Collins, A. A., & Brill, J. C. (1995). Spatial factors in vibrotactile pattern perception. *Perception & Psychophysics*, *57*, 724–737.

Clawson, J., Lyons, K., Starner, T., & Clarkson, E. (2005). The impacts of limited visual feedback on mobile text entry for the twiddler and mini-qwerty keyboards. In *Proceedings of the Ninth IEEE International Symposium on Wearable Computers*, (pp. 170–177). Washington, DC, USA: IEEE Computer Society.
URL http://dl.acm.org/citation.cfm?id=1104998.1105301

Cohen, J. (1986). Weighted kappa: nominal scale agreement with provision for scaled disagreement or partial credit. *Psychological Bulletin*, *70*(4), 213–220.

Coroama, V. (2006). Experiences from the design of a ubiquitous computing system for the blind. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI EA '06, (pp. 664–669). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1125451.1125587

Craig, J. (1972). Difference threshold for intensity of tactile stimuli. *Attention, Perception, and Psychophysics*, *11*, 150–152.
URL http://dx.doi.org/10.3758/BF03210362

Craig, J. (1974). Vibrotactile difference thresholds for intensity and the effect of a masking stimulus. *Perception & Psychophysics*, *15*, 123–127.

Craig, J. C., & Lyle, K. B. (2001). A comparison of tactile spatial sensitivity on the palm and fingerpad. *Perception & Psychophysics*, *63*(2), 337–347.

Craig, J. C., & Sherrick, C. E. (1982). Dynamic tactile displays. In W. Schiff, & E. Foulke (Eds.) *Tactual perception: a sourcebook*, (pp. 209–233). Cambridge: Cambridge University Press.

Crossan, A., & Brewster, S. (2008). Multimodal trajectory playback for teaching shape information and trajectories to visually impaired computer users. *ACM Trans. Access. Comput.*, *1*, 12:1–12:34.
URL `http://doi.acm.org/10.1145/1408760.1408766`

E-Sense Technology. `http://www.senseg.com/`.

Espinosa, M. A., Unger, S., Ochaita, E., Blades, M., & Spencer, C. (1998). Comparing methods for introducing blind and visually impaired people to unfamiliar urban environment. *Journal of Environmental Psychology*, *18*, 277–287.

FliteTTS for iPhone. `https://bitbucket.org/sfoster/iphone-tts`.

Fogg, B., Cutler, L. D., Arnold, P., & Eisbach, C. (1998). Handjive: a device for interpersonal haptic entertainment. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '98, (pp. 57–64). New York, NY, USA: ACM Press/Addison-Wesley Publishing Co.
URL `http://dx.doi.org/10.1145/274644.274653`

Fukumoto, M., & Sugimura, T. (2001). Active click: tactile feedback for touch panels. In *CHI '01 extended abstracts on Human factors in computing systems*, CHI EA '01, (pp. 121–122). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/634067.634141`

Geldard, F., & Sherrick, C. (1972). The cutaneous □rabbit□: a perceptual illusion. *Science*, *178*, 178–179.

Geldard, F. A. (1957). Adventures in tactile literacy. *The American Psychologist*, *12*, 115–124.

Geldard, F. A. (1960). Some neglected possibilities of communication. *Science*, *131*(3413), 1583–1588.
URL `http://dx.doi.org/10.1126/science.131.3413.1583`

Gergle, D., Kraut, R., & Fussell, S. (2004). Language efficiency and visual technology: Minimizing collaborative effort with visual information. *Journal of Language and Social Psychology*, *23*, 491–517.

Gergle, D., Kraut, R. E., & Fussell, S. R. (2006). The impact of delayed visual feedback on collaborative performance. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '06, (pp. 1303–1312). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1124772.1124968

Gescheider, G. A. (1966). Resolving of successive clicks by the ears and skin. *Journal of Experimental Psychology*, *71*(3), 378–381.

Ghiani, G., Leporini, B., & Paternò, F. (2008). Vibrotactile feedback as an orientation aid for blind users of mobile guides. In *Proceedings of the 10th international conference on Human computer interaction with mobile devices and services*, MobileHCI '08, (pp. 431–434). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1409240.1409306

Goff, G. D. (1967). Differential discrimination of frequency of cutaneous mechanical vibration. *Journal of Experimental Psychology*, *74*(2), 294–299.

Google Directions API. https://code.google.com/apis/maps/documentation/directions/.

Google Intersection Explorer. http://www.googlelabs.com/show_details?app_key=agtnbGFiczIwLXd3d3IVCxIMTGFic0FwcE1vZGVsGIHDtwIM.

Gupta, S., Campbell, T., Hightower, J. R., & Patel, S. N. (2010). Squeezeblock: using virtual springs in mobile devices for eyes-free interaction. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, (pp. 101–104). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1866029.1866046

Haapalainen, E., Kim, S., Forlizzi, J. F., & Dey, A. K. (2010). Psycho-physiological measures for assessing cognitive load. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, Ubicomp '10, (pp. 301–310). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1864349.1864395

Hall, M., Hoggan, E., & Brewster, S. (2008). T-bars: towards tactile user interfaces for mobile touchscreens. In *Proceedings of the 10th international conference on Human computer*

interaction with mobile devices and services, MobileHCI '08, (pp. 411–414). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1409240.1409301

Harrison, C., & Hudson, S. E. (2009a). Providing dynamically changeable physical buttons on a visual display. In *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, (pp. 299–308). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1518701.1518749

Harrison, C., & Hudson, S. E. (2009b). Texture displays: a passive approach to tactile presentation. In *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, (pp. 2261–2264). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1518701.1519047

Hart, S. G., & Staveland, L. E. (1988). Development of nasa-tlx (task load index): Results of empirical and theoretical research. In P. A. Hancock, & N. Meshkati (Eds.) *Human Mental Workload*. Amsterdam: North Holland Press.

Helal, A. S., Moore, S. E., & Ramachandran, B. (2001). Drishti: an integrated navigation system for visually impaired and disabled. In *Proceedings of the 5th IEEE international symposium on wearable computers*, ISWC '01, (pp. 149–). Washington, DC, USA: IEEE Computer Society.
URL http://portal.acm.org/citation.cfm?id=580581.856574

Hemmert, F., Hamann, S., Löwe, M., Zeipelt, J., & Joost, G. (2010). Weight-shifting mobiles: two-dimensional gravitational displays in mobile phones. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, CHI EA '10, (pp. 3087–3092). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1753846.1753922

Herman, J. F., Herman, T. G., & Chatman, S. P. (1983). Constructing cognitive maps from partial information: a demonstration study with congenitally blind subjects. *Journal of Visual Impairment & Blindness*, *77*(5), 195–198.

Heuten, W., Henze, N., Boll, S., & Pielot, M. (2008). Tactile wayfinder: a non-visual support system for wayfinding. In *Proceedings of the 5th Nordic conference on Human-computer interaction*, NordiCHI '08, (pp. 172–181). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1463160.1463179

Heuten, W., Wichmann, D., & Boll, S. (2006). Interactive 3d sonification for the exploration of city maps. In *Proceedings of the 4th Nordic conference on Human-computer interaction*, NordiCHI '06, (pp. 155–164). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1182475.1182492

Hoggan, E., Anwar, S., & Brewster, S. A. (2007). Mobile multi-actuator tactile displays. In *Proceedings of the 2nd international conference on Haptic and audio interaction design*, HAID'07, (pp. 22–33). Berlin, Heidelberg: Springer-Verlag.
URL http://portal.acm.org/citation.cfm?id=1775512.1775518

Hoggan, E., Brewster, S. A., & Johnston, J. (2008a). Investigating the effectiveness of tactile feedback for mobile touchscreens. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, (pp. 1573–1582). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1357054.1357300

Hoggan, E., Kaaresoja, T., Laitinen, P., & Brewster, S. (2008b). Crossmodal congruence: the look, feel and sound of touchscreen widgets. In *Proceedings of the 10th international conference on Multimodal interfaces*, ICMI '08, (pp. 157–164). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1452392.1452423

Hoggan, E. E. (2010). *Crossmodal audio and tactile interaction with mobile touchscreens*. Ph.d. thesis, University of Glasgow.

Hoshi, T., Iwamoto, T., & Shinoda, H. (2009). Non-contact tactile sensation synthesized by ultrasound transducers. In *Proceedings of the 3rd joint EuroHaptics conference and symposium on Haptic interfaces for virtual environment and teleoperator systems*, World Haptics '09, (pp. 256–260). Washington, DC, USA: IEEE Computer Society.
URL http://portal.acm.org/citation.cfm?id=1549825.1550293

Howell, W. C. (1958). *Discrimination of rate of amplitude change in cutaneous vibration*. Unpublished doctoral dissertation, University of Virginia.

iPhone VoiceOver. http://www.apple.com/accessibility/iphone/vision.html.

ISO: Ergomics of human-computer interaction—Part910: Framework for tactile and haptic interaction.

Israr, A., & Poupyrev, I. (2011). Tactile brush: drawing on skin with a tactile grid display. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI

’11, (pp. 2019–2028). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1978942.1979235

Israr, A., Tan, H., & Reed, C. (2006). Frequency and amplitude discrimination along the kinestheticcutaneous continuum in the presence of masking stimuli. *The Journal of the Acoustical society of America*, *120*(5), 2789–2800.

Iwamoto, T., Tatezono, M., & Shinoda, H. (2008). Non-contact method for producing tactile sensation using airborne ultrasound. In *Proceedings of the 6th international conference on Haptics: Perception, Devices and Scenarios*, EuroHaptics ’08, (pp. 504–513). Berlin, Heidelberg: Springer-Verlag.
URL http://dx.doi.org/10.1007/978-3-540-69057-3_64

Jacobson, R. D. (1998). Navigating maps with little or no sight: A novel audio-tactile approach. In *Proceedings of workshop on Content visualization and intermedia representations*, CVIR ’98, (pp. 95–102).

Jansen, Y., Karrer, T., & Borchers, J. (2010). Mudpad: tactile feedback and haptic texture overlay for touch surfaces. In *ACM International Conference on Interactive Tabletops and Surfaces*, ITS ’10, (pp. 11–14). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1936652.1936655

Jay, C., Glencross, M., & Hubbold, R. (2007). Modeling the effects of delayed haptic and visual feedback in a collaborative virtual environment. *ACM Trans. Comput.-Hum. Interact.*, *14*, 8.
URL http://doi.acm.org/10.1145/1275511.1275514

Jones, L., & Lederman, S. (2006). *Human hand function*. New York, USA: Oxford University Press.

Kaaresoja, T. (2006). Snap-crackle-pop: tactile feedback for mobile touch screens. In *Proceedings of Eurohaptics*, (pp. 565–566).

Kaczmarek, K., Webster, J., Pach-y Rita, P., & Tompkins, W. (1991). Electrotactile and vibrotactile displays for sensory substitution systems. *IEEE Transactions on Biomedical Engineering*, *38*(1), 1–16.

Kane, S. K., Jayant, C., Wobbrock, J. O., & Ladner, R. E. (2009). Freedom to roam: a study of mobile device adoption and accessibility for people with visual and motor disabilities. In *Proceedings of the 11th international ACM SIGACCESS conference on Computers and accessibility*, Assets ’09, (pp. 115–122). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1639642.1639663

Kane, S. K., Wobbrock, J. O., & Ladner, R. E. (2011). Usable gestures for blind people: understanding preference and performance. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, (pp. 413–422). New York, NY, USA: ACM. URL http://doi.acm.org/10.1145/1978942.1979001

Kane, S. K., Wobbrock, J. O., & Smith, I. E. (2008). Getting off the treadmill: evaluating walking user interfaces for mobile devices in public spaces. In *Proceedings of the 10th international conference on Human computer interaction with mobile devices and services*, MobileHCI '08, (pp. 109–118). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1409240.1409253

Katz, D. (2009). *The world of touch*. Psychology Press. Originally published in 1925 as "Der Aufbau der Taswelt".

Kenward, M. G., & Roger, J. H. (1997). Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*, *53*(3), 983–997.

Kitchin, R. M., & Jacobson, D. R. (1997). Techniques to collect and analyze the cognitive map knowledge of persons with visual impairment or blindness: issues of validity. *Journal of Visual Impairment & Blindness*, *91*(4), 360–376.

Klatzky, R., Lederman, S., & Metzger, V. (1985). Identifying objects by touch: An "expert system". *Perception & Psychophysics*, *37*(4), 299–302.

Knudsen, V. O. (1928). Hearing with the sense of touch. *Journal of Psychology*, *1*, 320–352.

Koskinen, E., Kaaresoja, T., & Laitinen, P. (2008). Feel-good touch: finding the most pleasant tactile feedback for a mobile touch screen button. In *Proceedings of the 10th international conference on Multimodal interfaces*, ICMI '08, (pp. 297–304). New York, NY, USA: ACM. URL http://doi.acm.org/10.1145/1452392.1452453

Lee, J. C., Dietz, P. H., Leigh, D., Yerazunis, W. S., & Hudson, S. E. (2004). Haptic pen: a tactile feedback stylus for touch screens. In *Proceedings of the 17th annual ACM symposium on User interface software and technology*, UIST '04, (pp. 291–294). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1029632.1029682

Lee, S. C., & Starner, T. (2010). Buzzwear: alert perception in wearable tactile displays on the wrist. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, (pp. 433–442). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1753326.1753392

Leung, R., MacLean, K., Bertelsen, M. B., & Saubhasik, M. (2007). Evaluation of haptically augmented touchscreen gui elements under cognitive load. In *Proceedings of the 9th international conference on Multimodal interfaces*, ICMI '07, (pp. 374–381). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1322192.1322258

Li, K. A., Baudisch, P., & Hinckley, K. (2008). Blindsight: eyes-free access to mobile phones. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, (pp. 1389–1398). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1357054.1357273

Linvill, J., & Bliss, J. (1966). A direct translation reading aid for the blind. *Proceedings of the IEEE*, *54*(1), 40 – 51.

LoadStone. http://www.loadstone-gps.com.

Loomis, J. M., Golledge, R. G., Klatzky, R. L., Speigle, J. M., & Tietz, J. (1994). Personal guidance system for the visually impaired. In *Proceedings of the first annual ACM conference on Assistive technologies*, Assets '94, (pp. 85–91). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/191028.191051

Luk, J., Pasquero, J., Little, S., MacLean, K., Levesque, V., & Hayward, V. (2006). A role for haptics in mobile interaction: initial design using a handheld tactile display prototype. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, CHI '06, (pp. 171–180). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1124772.1124800

Lylykangas, J., Surakka, V., Salminen, K., Raisamo, J., Laitinen, P., Rönning, K., & Raisamo, R. (2011). Designing tactile feedback for piezo buttons. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, (pp. 3281–3284). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1978942.1979428

Mallinckrodt, E., Hughes, A., & Sleator, W. (1953). Perception by the skin of electrically induced vibrations. *Science*, *118*(3062), 277–278.

Marston, J. R., Loomis, J. M., Klatzky, R. L., Golledge, R. G., & Smith, E. L. (2006). Evaluation of spatial displays for navigation without sight. *ACM Trans. Appl. Percept.*, *3*, 110–124.
URL http://doi.acm.org/10.1145/1141897.1141900

McAdam, C., & Brewster, S. (2009). Distal tactile feedback for text entry on tabletop computers. In *Proceedings of the 23rd British HCI group annual conference on People and computers*, BCS-HCI '09, (pp. 504–511). Swinton, UK, UK: British Computer Society.
URL http://portal.acm.org/citation.cfm?id=1671011.1671076

McGookin, D., & Brewster, S. (2007). An initial investigation into non-visual computer supported collaboration. In *CHI '07 extended abstracts on Human factors in computing systems*, CHI EA '07, (pp. 2573–2578). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1240866.1241043

McGookin, D., Brewster, S., & Jiang, W. (2008). Investigating touchscreen accessibility for people with visual impairments. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*, NordiCHI '08, (pp. 298–307). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1463160.1463193

Mowbray, G. H., & Gebhard, J. W. (1957). Sensitivity of the skin to changes in the rate of intermittent stimuli. *Science*, *125*, 1297–1298.

Nashel, A., & Razzaque, S. (2003). Tactile virtual buttons for mobile devices. In *CHI '03 extended abstracts on Human factors in computing systems*, CHI EA '03, (pp. 854–855). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/765891.766032

Oakley, I., Brewster, S., & Gray, P. (2001). Can you feel the force? an investigation of haptic collaboration in shared editors. In *In Proceedings of Eurohaptics*, (pp. 54–59).

Oakley, I., & O'Modhrain, S. (2005). Tilt to scroll: evaluating a motion based vibrotactile mobile interface. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, WHC '05, (pp. 40–49). Washington, DC, USA: IEEE Computer Society.
URL http://dx.doi.org/10.1109/WHC.2005.138

Oulasvirta, A., Tamminen, S., Roto, V., & Kuorelahti, J. (2005). Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '05, (pp. 919–928). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1054972.1055101

Paladugu, D. A., Wang, Z., & Li, B. (2010). On presenting audio-tactile maps to visually impaired users for getting directions. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, CHI EA '10, (pp. 3955–3960). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1753846.1754085

Palmer, C. I., & Gardner, E. P. (1990). Simulation of motion of the skin iv responses of pacinian corpuscle afferents innervating the primate hand to stripe patterns on the optacon. *Journal of Neurophysiology*, *64*, 236–247.

Paneels, S., & Roberts, J. C. (2010). Review of designs for haptic data visualization. *EEE Trans. Haptics*, *3*, 119–137.
URL http://dx.doi.org/10.1109/TOH.2009.44

Parente, P., & Bishop, G. (2003). Bats: the blind audio tactile mapping system. In *Proceedings of ACM South Eastern Conference*.

Parhi, P., Karlson, A. K., & Bederson, B. B. (2006). Target size study for one-handed thumb use on small touchscreen devices. In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, MobileHCI '06, (pp. 203–210). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1152215.1152260

Parkes, D. (1998). Nomad: an audio-tactile tool for the acquisition, use and management of spatially distributed information by partially sighted and blind persons. In *Proceedings of the Second International Symposium on Maps and Graphics for Visually Handicapped People*, (pp. 24–29).

Pasquero, J., & Hayward, V. (2011). Tactile feedback can assist vision during mobile interactions. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, (pp. 3277–3280). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1978942.1979427

Passini, R., Proulx, G., & Rainville, C. (1990). The spatio-cognitive abilities of the visually impaired population. *Environment and Behavior*, *22*, 91–118.

Picard, D., & Pry, R. (2009). Does knowledge of spatial configuration in adults with visual impairments improve with tactile exposure to a small-scale model of their urban environment? *Journal of Visual Impairment & Blindness*, *103*(4), 199–209.

Pielot, M., Henze, N., Heuten, W., & Boll, S. (2008). Evaluation of continuous direction encoding with tactile belts. In A. Pirhonen, & S. Brewster (Eds.) *Haptic and Audio Interaction Design*, vol. 5270 of *Lecture Notes in Computer Science*, (pp. 1–10). Springer Berlin / Heidelberg.
URL `http://dx.doi.org/10.1007/978-3-540-87883-4_1`

Pielot, M., Krull, O., & Boll, S. (2010). Where is my team: supporting situation awareness with tactile displays. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, (pp. 1705–1714). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1753326.1753581`

Plimmer, B., Crossan, A., Brewster, S. A., & Blagojevic, R. (2008). Multimodal collaborative handwriting training for visually-impaired people. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, (pp. 393–402). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1357054.1357119`

Poupyrev, I., & Maruyama, S. (2003). Tactile interfaces for small touch screens. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, UIST '03, (pp. 217–220). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/964696.964721`

Poupyrev, I., Maruyama, S., & Rekimoto, J. (2002a). Ambient touch: designing tactile interfaces for handheld devices. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, UIST '02, (pp. 51–60). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/571985.571993`

Poupyrev, I., Okabe, M., & Maruyama, S. (2004). Haptic feedback for pen computing: directions and strategies. In *CHI '04 extended abstracts on Human factors in computing systems*, CHI EA '04, (pp. 1309–1312). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/985921.986051`

Poupyrev, I., Rekimoto, J., & Maruyama, S. (2002b). Touchengine: a tactile display for handheld devices. In *CHI '02 extended abstracts on Human factors in computing systems*, CHI EA '02, (pp. 644–645). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/506443.506525`

Rabinow, J. (1948). The magnetic fluid clutch. *Transactions of the American Institute of Electrical Engineers*, *67*, 1308–1315.

Rantala, J., Raisamo, R., Lylykangas, J., Surakka, V., Raisamo, J., Salminen, K., Pakkanen, T., & Hippula, A. (2009). Methods for presenting braille characters on a mobile device with a touchscreen and tactile feedback. *IEEE Transactions on Haptics*, *2*, 28–39.
URL `http://portal.acm.org/citation.cfm?id=1525646.1525738`

Ross, D. A., & Blasch, B. B. (2000). Wearable interfaces for orientation and wayfinding. In *Proceedings of the fourth international ACM conference on Assistive technologies*, Assets '00, (pp. 193–200). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/354324.354380`

Rothenberg, M., Verrillo, R. T., Zahorian, S. A., Beracheman, M. L., & Bolanowski Jr., S. J. (1977). Vibrotactile freqeuncy for encoding a speech parameter. *Journal of the Acoustic Society of America*, *62*, 1003–1012.

Rowell, J., & Ungar, S. (2005). Feeling our way: tactile map user requirements – a survey. In *Proceedings of 22nd International Cartographic Conference*. International Cartographic Association.

Sahami, A., Holleis, P., Schmidt, A., & Häkkilä, J. (2008). Rich tactile output on mobile devices. In *Proceedings of the European Conference on Ambient Intelligence*, AmI '08, (pp. 210–221). Berlin, Heidelberg: Springer-Verlag.
URL `http://dx.doi.org/10.1007/978-3-540-89617-3_14`

Sánchez, J., Sáenz, M., Pascual-Leone, A., & Merabet, L. (2010). Navigation for the blind through audio-based virtual environments. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, CHI EA '10, (pp. 3409–3414). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1753846.1753993`

Sendero GPS. `http://www.senderogroup.com/products/shopgps.htm`.

Sherrick, C. E. (1964). Effects of double simultaneous stimulation on the skin. *American Journal of Psychology*, *77*, 42–53.

Sherrick, C. E. (1985). A scale for rate of tactual vibration. *Journal of the Acoustic Society of America*, *78*(1), 78–83.

Solovey, E. T., Girouard, A., Chauncey, K., Hirshfield, L. M., Sassaroli, A., Zheng, F., Fantini, S., & Jacob, R. J. (2009). Using fnirs brain sensing in realistic hci settings: experiments and guidelines. In *Proceedings of the 22nd annual ACM symposium on User interface software*

and technology, UIST '09, (pp. 157–166). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1622176.1622207

Spencer, C., & Travis, J. (1985). Learning a new area with and without the use of tactile maps: a comparative study. *British Journal of Visual Impairment*, *3*(1), 5–7.

Stark, R., Carlstedt, T., Hallin, R. G., & Risling, M. (1998). Distribution of human pacinian corpuscles in the hand. *The Journal of Hand Surgery: Journal of the British Society for Surgery of the Hand*, *23*, 370–372.

Stewart, J., Bauman, S., Escobar, M., Hilden, J., Bihani, K., & Newman, M. W. (2008). Accessible contextual information for urban orientation. In *Proceedings of the 10th international conference on Ubiquitous computing*, UbiComp '08, (pp. 332–335). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1409635.1409679

Stolle, A. M., Holzl, R., Kleinbohl, D., Mrsic, A., & Tan, H. Z. (2004). Measuring point localization errors in spatiotemporal tactile stimulus patterns. In *PEuroHaptics*, 512-515.

Strong, R., & Troxel, D. (1970). An electrotactile display. *IEEE Transactions on Man-Machine Systems*, *11*(1), 72–79.

Su, J., Rosenzweig, A., Goel, A., de Lara, E., & Truong, K. N. (2010). Timbremap: enabling the visually-impaired to use maps on touch-enabled devices. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, MobileHCI '10, (pp. 17–26). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1851600.1851606

Summers, I. R. (2000). Single channel information transfer through the skin: limitations and possibilities. In *In the Proceedings of International Sensory Aids Conference*, (pp. 23–26).

Suzuki, Y., & Kobayashi, M. (2005). Air jet driven force feedback in virtual reality. *IEEE Comput. Graph. Appl.*, *25*, 44–47.
URL http://portal.acm.org/citation.cfm?id=1042190.1042215

Talking signs. http://www.talkingsigns.com/.

Thinus-Blanc, C., & Gaunet, F. (1997). Representation of space in blind persons: vision as a spatial sense? *Psychological Bulletin*, *121*(1), 20–42.

Trekker. http://www.humanware.com/en-canada/products/blindness/talking_gps/trekker/_details/id_88/trekker.html.

Trekker Breeze. http://www.humanware.com/en-canada/products/blindness/talking_gps/trekker_breeze/_details/id_101/trekker_breeze.html.

Tsukada, K., & Yasumura, M. (2004). Activebelt: belt-type wearable tactile display for directional navigation. In *In Proceedings of the Sixth International Conference on Ubiquitous Computing*, UbiComp '04. Springer.

Ungar, S. (2000). Cognitive mapping without visual experience. In R. Kitchin, & S. Freundschuh (Eds.) *Cognitive Mapping: Past Present and Future*. London: Routledge.

Ungar, S., Blades, M., & Spencer, C. (1993). The role of tactile maps in mobility training. *British Journal of Visual Impairment*, *11*(2), 59–61.

Ungar, S., Blades, M., & Spencer, C. (1995). Visually impaired children's strategies for memorising a map. *British Journal of Visual Impairment*, *13*(1), 27–32.

van Erp, J. B. F. (2002). Guidelines for the use of vibro-tactile displays in human computer interaction. In *Proceedings of Eurohaptics 2002*.

Verrillo, R. T., & Gescheider, G. A. (1992). *Tactile aids for the hearing impaired*, chap. Perception via the sense of touch, (pp. 1–36). London: Whurr Publishers.

Vogel, D., & Balakrishnan, R. (2010). Occlusion-aware interfaces. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, (pp. 263–272). New York, NY, USA: ACM.
URL http://doi.acm.org/10.1145/1753326.1753365

Walker, B. N., & Lindsay, J. (2006). Navigation performance with a virtual auditory display: effects of beacon sound, capture radius, and practice. *Human Factors*, *48*, 265–278.

Wall, S. A., & Brewster, S. (2006). Sensory substitution using tactile pin arrays: human factors, technology and applications. *Signal Processing*, *86*, 3674–3695.
URL http://portal.acm.org/citation.cfm?id=1228582.1228593

Weber, E. H. (1995). *On the tactile senses*. Erlbaum, Taylor & Francis. Originally published 1846 as "De Tactu," In Handworterbuch der Physiologie, Volume 3, pp.481-588.

Wilson, J., Walker, B. N., Lindsay, J., Cambias, C., & Dellaert, F. (2007). Swan: system for wearable audio navigation. In *Proceedings of the 11th International Syposium on Wearable Computers*, ISWC '07, (pp. 1–8).

Wong, E., Israr, A., & O'Malley, M. (2010). Discrimination of consonant articulation location by tactile stimulation of the forearm. In *Proceedings of 2010 IEEE Haptics symposium*, (pp. 47 –54).

Yang, R., Park, S., Mishra, S. R., Hong, Z., Newsom, C., Joo, H., Hofer, E., & Newman, M. W. (2011). Supporting spatial awareness and independent wayfinding for pedestrians with visual impairments. In *Proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility*.

Yatani, K., Novati, M., Trusty, A., & Truong, K. N. (2011). Review spotlight: a user interface for summarizing user-generated reviews using adjective-noun word pairs. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, (pp. 1541–1550). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1978942.1979167`

Yatani, K., & Truong, K. N. (2007). An evaluation of stylus-based text entry methods on handheld devices in stationary and mobile settings. In *Proceedings of the 9th international conference on Human computer interaction with mobile devices and services*, MobileHCI '07, (pp. 487–494). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1377999.1378059`

Yatani, K., & Truong, K. N. (2009). An evaluation of stylus-based text entry methods on handheld devices studied in different mobility states. *Pervasive and Mobile Computing*, *5*(5), 496–506.

Yelp API 2.0. `http://www.yelp.com/developers/documentation/v2/overview`.

Zelek, J. S., Bromley, S., Asmar, D., , & Thompson, D. (2003). A haptic glove as a tactile-vision sensory substitution for wayfinding. *Journal of Visual Impairment & Blindness*, *97*(10), 1–24.

Zhao, S., Dragicevic, P., Chignell, M., Balakrishnan, R., & Baudisch, P. (2007). Earpod: eyes-free menu selection using touch input and reactive audio feedback. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '07, (pp. 1395–1404). New York, NY, USA: ACM.
URL `http://doi.acm.org/10.1145/1240624.1240836`

Zimmerman, T. G., Lanier, J., Blanchard, C., Bryson, S., & Harvill, Y. (1987). A hand gesture interface device. In *Proceedings of the SIGCHI/GI conference on Human factors in computing systems and graphics interface*, CHI '87, (pp. 189–192). New York, NY, USA:

ACM.
URL http://doi.acm.org/10.1145/29933.275628

Appendix A

Spatial Tactile Feedback Hardware Circuit Diagram

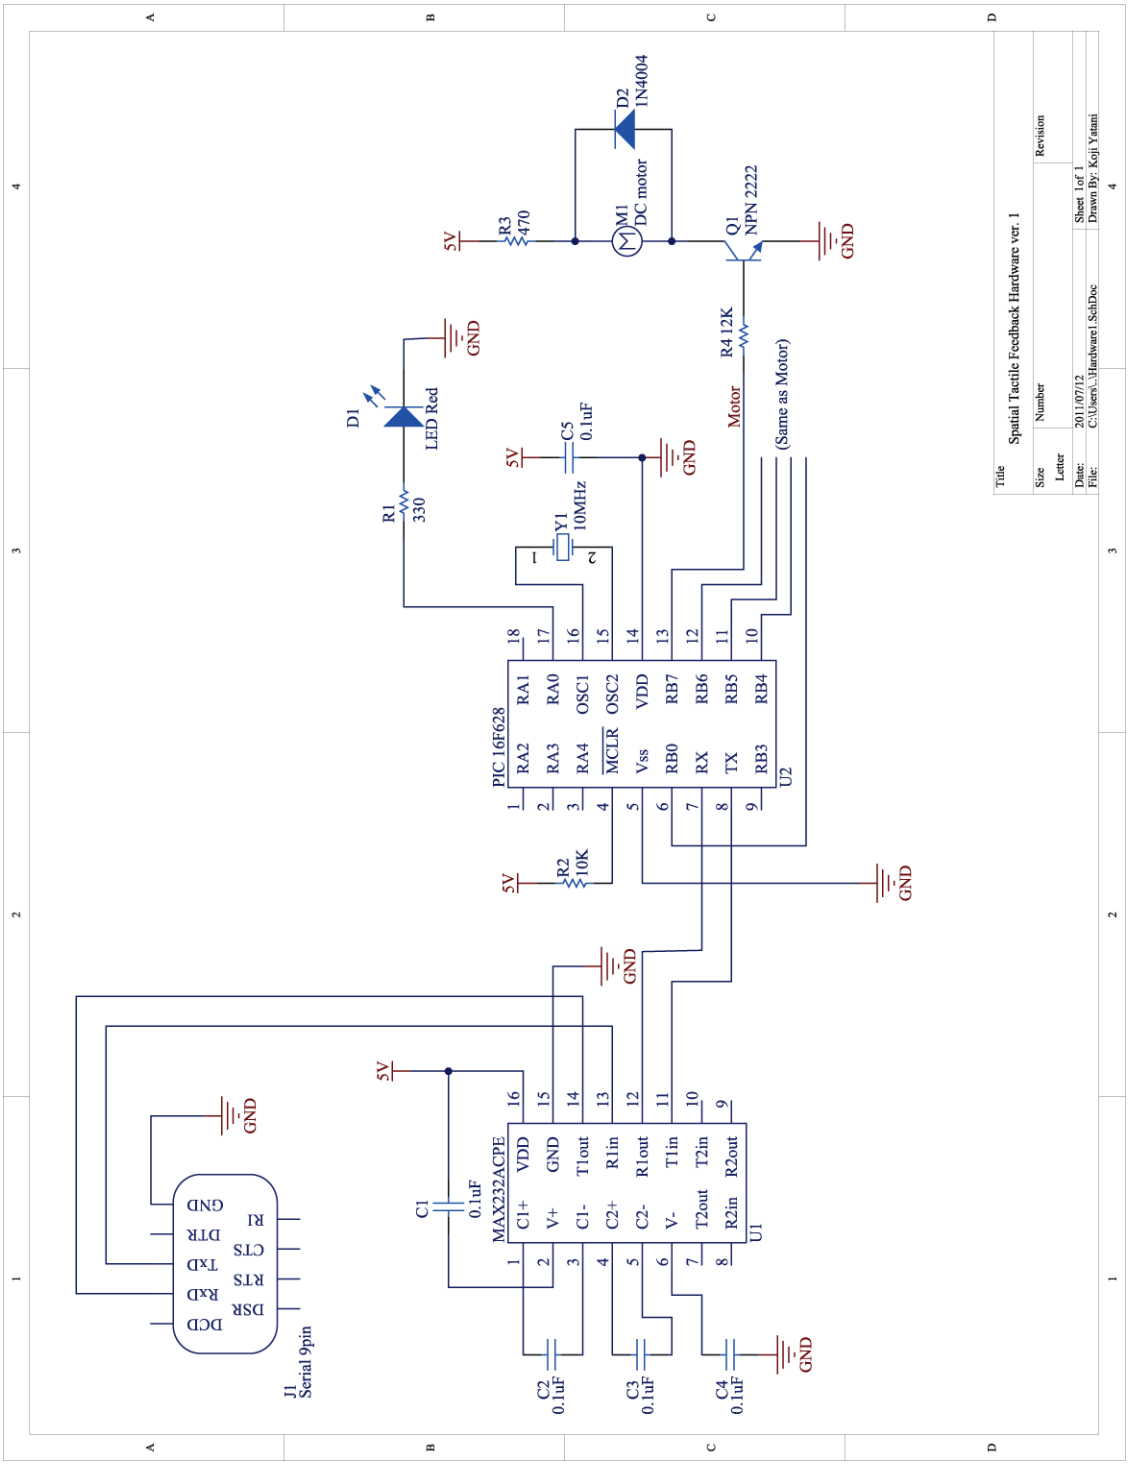

Figure A.1. The circuit diagram for the first prototype.

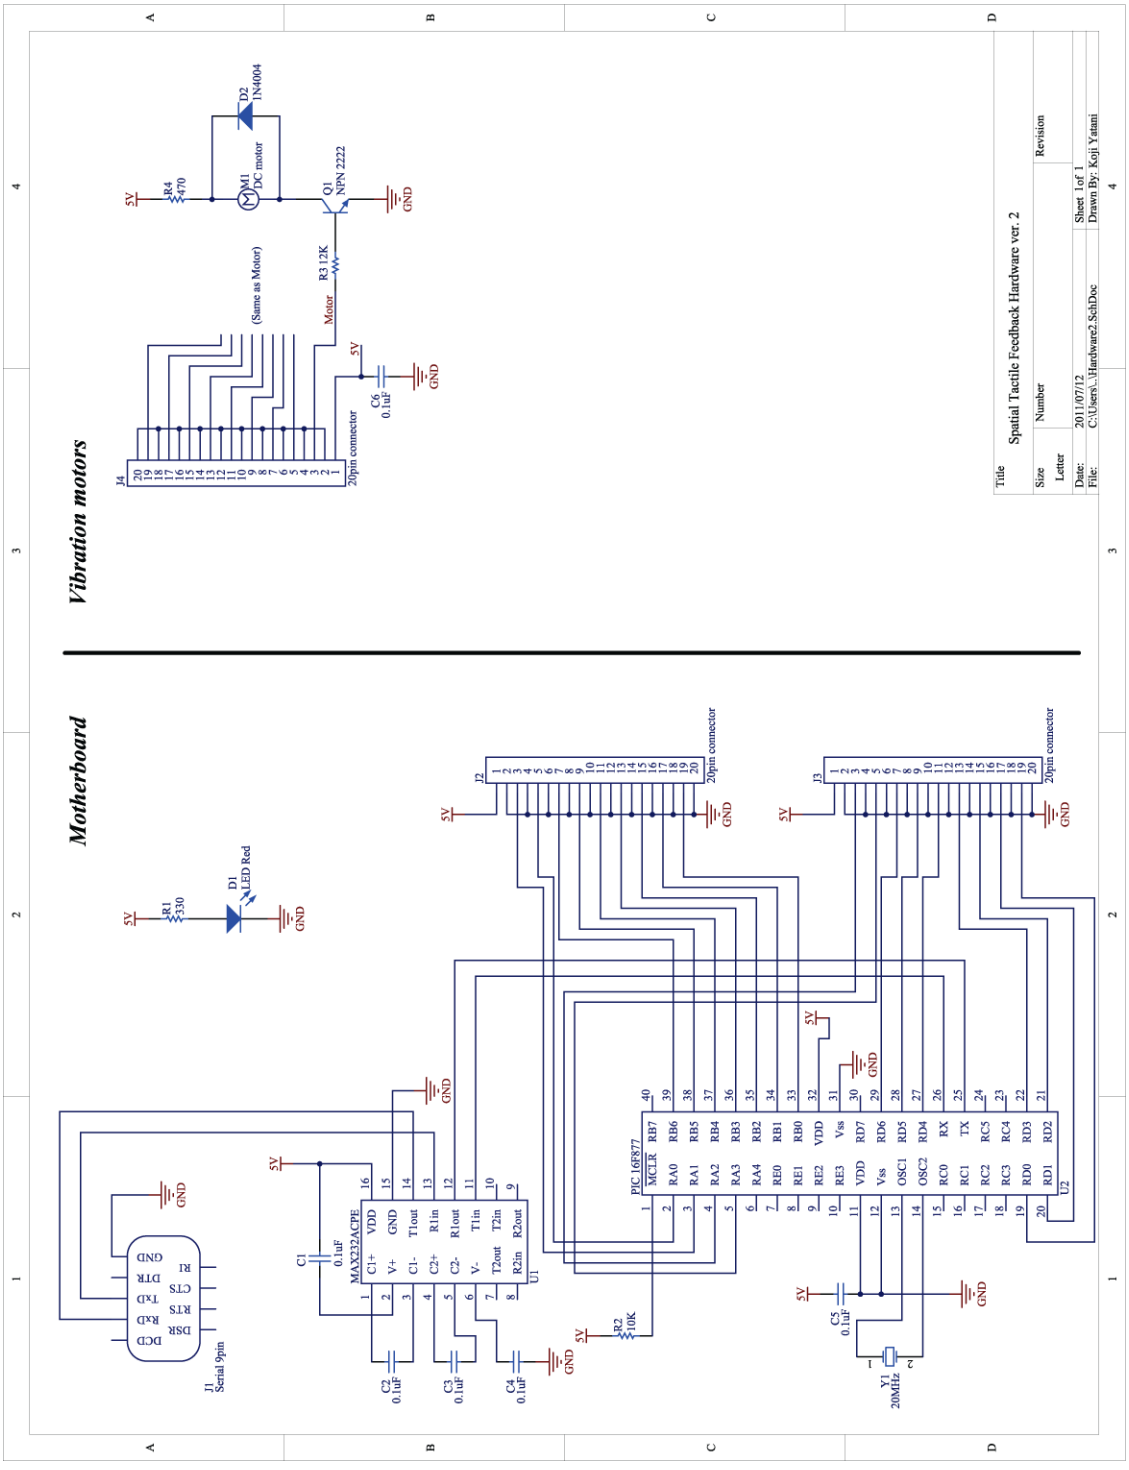

Figure A.2. The circuit diagram for the second prototype.

Appendix B

Study on Audio Review Spotlight in SpaceSense

Experimental Design

To examine whether an audio adaptation of the Review Spotlight system (Yatani *et al.*, 2011) could be useful in providing visually impaired users with descriptive information about a location, I set up a task for acquiring details of places. I prepared two places (a restaurant and a café) for this task. For each location, I prepared two presentations to provide the review information: *ReviewSpotlight* (reading out adjective-noun word pairs most frequently mentioned in Yelp.com) and *ReviewHighlights* (reading out five snippets extracted from the review text, provided in Google Map). Table B.1 shows an example of both presentations used in the study. Participants were exposed to both presentations to learn the details of the location, and were asked to express their opinions about them. The presentation order of *ReviewSpotlight* and *ReviewHighlights* was counter-balanced across the participants.

Presentation	Example
ReviewHighlights	*"Neatloaf tests pretty good, however the ensalada mexicana is nothing special. The little tea pots are cute. I was going to order the main entree neatloaf and replaced the mashed potatoes with salad. The waiter pointed out that it's cheaper to just order the sandwich and replace the..."* *"My boyfriend become a vegetarian last year so we have been trying some very unique restaurants to say the least. Ananda Fuara is painted in very calming colors and the servers wear saris. It has a bit of a cult like feel (with the pamphlets on the table) but everyone is so..."*
ReviewSpotlight	*"Good food, Vegetarian restaurant, Vegetarian food, Neatloaf sandwich, Great food, Reasonable price, Great place, Friendly staff, First time, Delicious food."*

Table B.1. Examples of review presentations in SpaceSense used in the user study.

Participants

All participants were the same for another user study on SpaceSense (Section 5.5). I asked them about their familiarity with the places used in the this part of the study; none were familiar with either place. Participants were asked to complete this task after they finished all tasks and post-experimental interview for SpaceSense described in Section 5.5.

Results

All of the participants preferred the *ReviewSpotlight* presentation over *ReviewHighlight*. They liked its succinct presentation. They also liked that the *ReviewSpotlight* presentation summarized all reviews instead of presenting one particular review.

> *"I am getting more opinions, and the information is presented more quickly (than ReviewHighlights)."* [P4]

Some participants explicitly commented that the *ReviewSpotlight* presentation was easier to follow than *ReviewHighlights*.

> *"The first one (ReviewSpotlight) is obviously a better one, much much better. It's clear in a sense, because of its form; and the other one is more jumpy… Your brain has to sort out… It stops randomly too much, and missing data. It (ReviewHighlight) would be something you have to replay."* [P3]

Two participants explicitly mentioned the benefits of having both presentations and wanted to use them in different ways. P10 explained to us that he would choose either of the presentations depending on how much time he can spend on listening to reviews.

> *"I think it would be better to have both because sometimes you don't want to hear full reviews and you don't want to hear keywords (ReviewSpotlight presentation)… If I have a lot of time, I would want it to read the full reviews, but if I'm in a rush, I just want it to read the quick words (ReviewSpotlight presentation)."* [P10]

Summary

The audio adaptation of Review Spotlight was received positively by participants mainly because of its succinct presentation of reviews. However, as the participants indicated, they may want to access to portions of the original review text to gain more detailed information. This was discussed in the original Review Spotlight work, which incorporates a hyperlink on the adjective-noun word pair to the sentences from which the clicked word pair was extracted (Yatani *et al.*, 2011). A faithful adaptation of the Review Spotlight system is out of the scope of this work; however, future work should investigate how a system can effectively support both a quick overview and exploration of details in online user reviews through the speech feedback.