Tactile Display for Mobile Interaction

Jerome Pasquero

Department of Electrical & Computer Engineering
McGill University
Montreal, Canada

August 2008

A thesis submitted to McGill University in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Your file Votre référence
ISBN: 978-0-494-66671-5
Our file Notre référence
ISBN: 978-0-494-66671-5

To Tanya

Abstract

Interaction with mobile devices suffers from a number of shortcomings, most of which are linked to the small size of screens. Artificial tactile feedback promises to be particularly well suited to the mobile interaction context. To be practical, tactile transducers for mobile devices must be small and light, and yet be capable of displaying a rich set of expressive stimuli. This thesis introduces a tactile transducer for mobile interaction that is capable of distributed skin stimulation on the fingertip. The transducer works on a principle that was first investigated because of its potential application to the display of Braille. A preliminary study was conducted on an earlier version of the transducer. It concluded that subjects were able to identify simple Braille characters with a high rate of success. Then, a complete re-design of the transducer addressed the goal of integration in a handheld prototype for mobile interaction. The resulting device comprises a liquid crystal graphic display co-located with the miniature, low-power, distributed tactile transducer. Next, it was needed to measure the perceptual differences between the stimuli that the device could display. Our experiences with one evaluation approach raised questions relating to the methodology for data collection. Therefore, an analysis of the process was carried out using a stimulus set obtained with the device. By means of multidimensional scaling analysis, both the perceptual parameters forming the stimuli space and the evaluation technique were validated. Finally, two experiments were carried out with the objective to develop new mobile interactions paradigms that combined visual and tactile feedback. Both experiments modeled a list scrolling task on the device. The first experiment found a marginal improvement in performance when tactile feedback was employed. It also came at a higher attentional cost dedicated to operating the device. For the second experiment, the scrolling paradigm and the tactile feedback were improved. This lead to a decrease in the reliance on vision when tactile feedback was enabled. Results showed a 28% decrease in the number of key presses that controlled the visibility state of the scroll list.

Sommaire

Les interfaces des appareils portatifs sont criblés de problèmes d'ergonomie informatique, qui pour la plupart sont dus à la petite taille de leur écran. Dans ce contexte, la génération de sensations tactiles offre une piste de solution intéressante. Un afficheur tactile pour appareil portatif se doit d'être compact et léger, tout en étant en mesure de recréer une vaste gamme de sensations. Sa consommation électrique se doit aussi d'être faible. Cette thèse propose un transducteur qui répond à tous ces critères. Ce dernier fonctionne en appliquant un champ distribué de déformations locales à la surface de la peau du bout du doigt. Dans un premier temps, le transducteur est évalué afin de juger de son potentiel à afficher des points Braille. Par la suite, la conception du transducteur est repensée afin de rendre possible son insertion dans un prototype d'appareil portatif. Le prototype consiste en un boitier qui peut être tenu dans la main et sur lequel a été fixé un écran à cristaux liquides. L'étape suivante consiste en une exploration des différentes sensations tactiles que le prototype peut communiquer. L'identification des capacités d'affichage de l'appareil requiert une évaluation un peu plus approfondie de la technique utilisée. Cette dernière est éventuellement validée. Pour finir, deux études sont mises sur pied avec l'objectif de développer sur le prototype des applications qui associent un retour tactile à une information visuelle. Dans les deux cas, les sujets ont pour tâche de faire défiler une liste d'éléments qui est affichée sur l'écran de l'appareil et de s'arrêter sur une cible qui leur est pré-indiquée. La première étude ne trouve qu'un gain marginal à l'utilisation du retour tactile. De plus, le retour tactile semble attirer l'attention sur l'appareil au détriment d'une attention qui devrait plutôt être portée à l'environnement. Pour la deuxième étude, la métaphore de contrôle du défilement de la liste et le retour tactile sont tous deux améliorés. Les sujets sont présentés avec une liste qui est invisible sur l'écran par défaut et qui ne devient visible que lorsqu'une touche est appuyée. Les résultats obtenus pour la deuxieme experience montrent une diminution de 28% du nombre de fois que les sujets font appel à la touche lorsque le retour tactile est présent. Ceci semble suggérer que le retour tactile peut diminuer la dépendance à l'information visuelle.

Acknowledgments

Firstly, I would like to thank my supervisor Prof. Vincent Hayward. The quest for a Ph.D. is a journey that has its share of ups and downs, but the unconditional support and insight that Prof. Hayward provided me kept the downs short and bearable and the ups abundant and sweet.

Along my Ph.D. years, there are a handful of people who have provided expertise, encouragement and inspiration. My colleague and friend Vincent Levesque granted the right mix of constructive criticism and support that allowed me to better my work. Don Pavlasek communicated technical expertise in the art of mechanical design and machining with an enthusiasm that is contagious. On more than one occasion, Prof. Karon MacLean warmly welcomed me into her lab at the University of British Columbia, where I was kindly introduced to important tools such as experimental design. Prasun Lala, Steve Yohanan and Andrew "Goose" Gosline offered a combination of knowledge and wit that was both refreshing and a source of motivation. Joseph Luk was instrumental in opening up an entire new realm of possibilities for the technology and concepts that are developed in this thesis. Lastly, Cynthia Davidson offered a friendly smile and a rescuing hand every time I got lost in administrative burden. I am seriously indebted to all of them.

For their comradeship and advice, I would also like to thank the following current and former members of the McGill Haptics Lab (in alphabetical order): Omar Ayoub, Gianni Campion, Hanifa Dostmohamed, Akihiro Sato, Qi Wang, Hsin-Yun Yao and Mounia Ziat.

For their support and encouragement, I am grateful to my family and to the Nahorniaks.

Finally, I am ever indebted to Tanya Nahorniak for her unconditional love and kindness and for so much more that I unfortunately cannot list here because it would require doubling the number of pages of this manuscript...

Contents

List of Figures

List of Tables

List of Acronyms

CAD	Computer-Aided Design
CNC	Computer Numerical Control
DC	Direct Current
ECE	Electrical and Computer Engineering
FPGA	Field-Programmable Gate Array
HCI	Human Computer Interaction
IQR	Interquartile Range
LCD	Liquid Crystal Display
MDS	Multidimensional Scaling
MEMS	Microelectromechanical System
PC	Personal Computer
PCB	Printed Circuit Board
PDA	Personal Digital Assistant
PWM	Pulse Width Modulation
RA	Rapidly Adapting
SA	Slowly Adapting
SAW	Surface Acoustic Wave
SMA	Shape Memory Alloy
STRESS	Stimulator of Tactile Receptors by Skin Stretch
THMB	Tactile Handheld Miniature Bimodal
USB	Universal Serial Bus
VBD	Virtual Braille Display

Chapter 1

Introduction

Mobile Interactions

Mobile devices, such as PDAs, mobile phones and portable music players embody the essence of readily available and distributed computation that is embedded in our environment. Rich interactions with both information and people are now possible from anywhere, at any time, with almost no overhead and even while in motion. Mobile devices are constantly decreasing in volume and weight, but continue to provide access to ever-increasing computational power that is available on-the-move. While there seems to be no limit to how much mobile technology can connect us to both the digital and the physical worlds, mobile devices offer their share of challenges.

For one, mobile devices suffer from small display sizes. This situation is not likely to improve significantly as their popularity is precisely based on the fact that people do not like to carry around devices that are encumbering. The limited screen real-estate afforded by mobile devices means that information visualization is restricted [39, 80, 179].

Designing for appropriate levels of intrusiveness and interruption is a significant challenge for mobile interface designers. Sensorial overload and intrusiveness are known issues with modern human-computer interaction technology. Graphical and auditory information is incessantly assailing our senses and the channels for perceptual communication are becoming saturated. In a mobile context, users might be

walking at the same time they are talking on the phone or sending an email with their PDA. This adds to the already heavy cognitive burden they are confronted to [123]. Moreover, attention in a mobile environment also tends to be more fragmented than in its desktop equivalent [60, 109].

Mobile interaction also suffers from another important drawback of modern technology, though it is less often reported. An easy and readily available access to the digital world has distanced us from the physical world because our interactions are too often confined to the use of miniature screens and keypads [65]. Because they must remain small and cheap, mobile devices are usually made of miniature push-buttons and dials that display very limited haptic feedback. For the same reason, their ergonomics is often awkward. By contrast, the click of a latch (e.g., when closing a cupboard), the inertia of a knob (e.g., when opening a door) and even the feel of a texture (e.g., when choosing a fruit that is ripe) are integral parts of our daily interactions with the physical world. Mobile devices, while not its only instigator, contribute greatly to the loss of physicality experienced because they feel unnatural.

A Role for Tactile Feedback

The use of artificial tactile feedback offers solutions that seem particularly well suited in the context of mobile interactions. Touch is an intimate sense that fits well with the personal nature of mobile devices. More importantly, because users keep their mobile devices at hand or close to their body, they are in constant contact with the interface and are likely to feel any feedback provided by the device.

Conveying information through touch offers many potential advantages for mobile interactivity. For instance, it can:

- Diminish the need for frequent glances at the graphical information on the screen [131]. This can free part of the visual attention and therefore allow mobile users to carry out other tasks while constantly monitoring their device in the background.

- Be useful in noisy environments where audio signals cannot be heard or in very

quiet environments where they quickly become intrusive. Touch feedback is also significantly more discrete than its audio and visual counterparts when privacy is an issue.

- Grab a user's attention (e.g., to communicate a sense of urgency), as illustrated by how a tap on the shoulder can make us shift our focus to the interrupter.
- Reinforce or complement a message conveyed primarily by another modality, in the same way that a "new mail" icon is often accompanied by a short audio signal on a personal computer.
- Compensate for the downsides of miniaturization. For instance, it could guide text input on a small keypad by returning tactile sensations that correspond to what was just entered. This can be particularly valuable in environments where the screen is difficult to see or simply when entering text is an intricate operation because of the awkward size of the input buttons.
- Restore physicality to our experience with mobile devices. Our motor skills are deeply coupled to our sense of touch and the sensation that we are manipulating real and natural physical objects can be partially restored by providing artificial tactile feedback [96].

Problem Statement

While the advantages of using tactile feedback for mobile interactions are becoming apparent, attempts at making working prototypes are still quasi-inexistent (see [26, 45] for very rare exceptions). To be practical in a mobile context, the transducers for tactile interactivity must be small and light. They must require low power and must be safe and robust at the same time. Considering the state of the art (more on this in section 2.4.1), it is not surprising that only small rotary vibrating motors are currently used in commercial mobile devices [63, 94]. When compared to the full gamut of real-life tactile interactions, this approach can be limited in expressive capabilities. The full role that tactile feedback can play in mobile interaction can

only be revealed with the development of artificial tactile feedback that is rich in information content. Consequently, there is a clear need for miniature distributed tactile transducers that are capable of displaying large sets of tactile stimuli that are easy to differentiate.

Approach

One way to start fulfilling this demand consists of using an up-to-now unexplored mode of interaction that differs from the conventional approach of indenting the skin. When a wave of localized deformations travels tangentially on the fingerpad skin, one typically experiences the illusion of a small feature sliding under the finger. The phenomenon can be reproduced by electromechanical means by using an array of piezoelectric bending actuators.[1] The resulting sensations created this way can vary according to a multitude of parameters such as frequency, amplitude, trajectory and speed of the traveling wave. This allows for the display of a set of rich and diverse tactile stimuli.

Thesis Overview

This thesis tells the story about a tactile transducer for mobile interaction, from its inception to its initial validation. The transducer was integrated in a handheld case that constitutes, to the author's knowledge, the first realistic prototype of a mobile handheld interface capable of direct-contact distributed tactile feedback. The resulting device is referred to as the Tactile Handheld Miniature Bimodal (THMB) device.

[1] The underlying mechanisms that make this illusion work are beyond the scope of this thesis. The reader is referred to [171] for a partial explanation based on a contact mechanics analysis. In brief, Wang and Hayward suggest that, under certain conditions, an electromechanical display capable of lateral skin deformation can generate strain tensor fields that are undistinguishable from the fields created by a sinusoidal surface sliding on the fingertip. From the central nervous system's point of view, both stimuli are equivalent and they are therefore resolved as the same tactile sensation.

Chapter 2 reviews the sense of touch when seen from the point of view of a communication channel, because good design for tactile interaction should be guided by a good understanding of touch. It provides a background overview of different attempts at using the sense of touch to convey information to a human user by artificial means. The different mechanisms by which the human body encodes and decodes tactile information internally are also covered. Finally, state-of-the-art in fingertip tactile display technology is reviewed.

The three chapters that follow are reproductions of manuscripts that were published for publication. Taken together with the last chapter (Chapter 6), they make the core of this thesis which can be thought of as a chronological account of how the THMB device was motivated, designed, characterized and exercised respectively.

Chapter 3 describes an electromechanical prototype that generates lateral skin deformation on the fingertip to display Braille dots. The prototype, called Virtual Braille Display (VBD) is, in essence, the predecessor to the THMB transducer. It consists of a linear array of piezoelectric benders that can be programmed to apply patterns of local deformation to a user's fingerpad. This early device was used for a preliminary feasibility study that assessed the potential of using the technology to display Braille characters. Results and insight collected during the manufacturing of the VBD and the feasibility study that followed informed many subsequent design decisions during the development of the THMB device. Concepts that are introduced in this chapter are often revisited in some form or another in the later chapters.

Chapter 4 provides rationale for building the THMB transducer, and integrating it into a handheld device with a Liquid Crystal Display (LCD). The THMB device is described in details so that it can be reproduced easily. The goal for building this prototype was to construct a platform that allows the exploration of application-specific studies of information manipulation in the small space available on a mobile device.

Chapter 5 validates an experimental technique that is used to identify the underlying perceptual parameters of a set of tactile stimuli. The technique, known as cluster-sorted Multidimensional Scaling (MDS), was popularized by MacLean and Enriquez as a rapid method to measure and characterize the perceptual distances

between a group of tactile stimuli [97]. It is useful in identifying which tactile stimuli are easy to discriminate for users. However, the assumptions and ramifications associated with using the technique had never been analyzed. This chapter takes a closer look at cluster-sorted MDS and points out its impact on the conclusions that can be drawn from it. Because the analysis is illustrated with sample data directly taken from a set of stimuli generated with the THMB, the expressive capabilities of the device are also explored.

Chapter 6 pursues the exploration of applications-specific studies with the THMB device. It describes the design and experimental evaluation of a multimodal scrolling interaction that combines visual and tactile feedback. This work fits under the overarching goal of finding how direct-contact tactile feedback can provide added-value in a mobile environment in terms of both performance and comfort. The chapter describes the results of two experiments that evaluate the effect of tactile feedback on user attention allocation and reliance on vision.

Chapter 7 concludes the thesis with a review of its main findings and states potential future work.

Summary of Contributions

This thesis describes the following contributions:

- The design and implementation of an electromechanical prototype capable of displaying Braille dots by means of varying patterns of localized tangential skin deformation on the fingerpad.
- The reduction to practice of a miniature distributed tactile transducer that is integrated in a prototype of a handheld mobile device. This device makes it possible to deploy a rich set of multimodal mobile interactions with co-located visual and tactile feedback.
- The formal validation of the cluster-sorted multidimensional scaling method for use in the perceptual characterization of new tactile feedback paradigms and devices.
- Experimental demonstration that artificial tactile feedback in mobile interactions can decrease the visual attentional load during scrolling tasks, comprising
  - the design of a scrolling interaction controller on a mobile device prototype that enables anticipatory control, and
  - an experimental method to unobtrusively measure subjects' attentional visual behavior.

Chapter 2

Background

2.1 Introduction

The tactile channel is adept at extracting information present in naturally occurring tactile sensations (e.g., appreciating the texture of a fabric). Information can also be conveyed by tactile stimuli that are artificially produced with a view to enhancing human-computer interaction (e.g., the "buzz" of a cell phone). Another possibility is to seek to replicate naturally occurring sensations for virtual reality applications as in a training simulator. Human-computer interaction technology (HCI) already leverages touch for information transfer. For instance, the click experienced when a character is entered on a keyboard informs the user of the occurrence of the event. However, it is only recently that HCI has attempted to use touch to close the communication loop from the computer to the human in a programmed manner. Typically, computer interfaces rely mostly on vision and audition to supply information [56], when, in fact, people spontaneously and unconsciously rely on touch to explore and experience their environment [130].

The design of artificial tactile feedback is a new research area with exciting applications. In HCI, artificial tactile feedback aims at communicating contextual information. Such interactions take advantage of the touch channel to provide information with minimal distraction to the awake individual. Consider how the sensation of a sharp edge experienced while holding a glass is an indication that it is likely cracked.

Similarly, devices able to supply tactile feedback artificially could communicate status information (e.g., to indicate the presence of a message) or instructions to navigate in an unknown environment (e.g., "turn right").

There is growing demand for communication through touch because of the increased number of opportunities in a wide range of areas such as entertainment (e.g., computer paddles), medical technologies (e.g., virtual surgery training, sensory substitution), research (e.g., study of perception), and many more. Research on tactile displays has been focused mainly on devices that apply normal indentation to a user's fingertip. In most cases, this is achieved through the vertical movement of miniature pins against the skin to reproduce small-scale shapes or textures. In general, however, artificial tactile feedback can be supplied by electromechanical devices, called *tactile displays*, which operate according to a wide variety of principles.

In this chapter, we first consider early and recent research on the development of artificial tactile communication. Then, we review current models of the encoding of tactile information in humans before examining the state-of-the-art for tactile displays.[1] While much remains to be discovered, we believe that these findings can guide the design of an artificial language for touch.

2.2 Tactile Communication

2.2.1 Tactile Languages

Nearly five decades ago, Geldard advocated that the sense of touch constituted a "neglected sense of communication" [51]. He noted that, while the visual and auditory systems were superior at spatial and temporal discrimination respectively, the somatosensory system was capable of both. Therefore, touch was particularly suited when hearing or sight were not available.

Others before Geldard had attempted to use touch as a communication channel. One example was Gault's Teletactor sensory substitution system that displayed the

[1] Our review is limited to cutaneous touch, i.e., related to the stimulation of the skin. We are mainly concerned with the mechanical deformation of the skin and leave out pain and temperature sensations.

mechanical equivalent of speech sounds onto the skin by means of an electromechanical device [49]. The goal was to use the skin to "hear". This approach, unfortunately, suffered from a disregard of some inherent properties of touch. The fundamental frequencies required for proper speech comprehension precisely lie outside the range where tactile frequency discrimination is the highest; thus, the communication device and the receptive system were simply not matched.

The example above illustrates how the development of an artificial tactile language demands a set of rules that is matched to the somatosensory system's capabilities and limitations. Following this premise, Geldard and his colleagues went on to develop "Vibratese", a tactile language based on both practical considerations and on results from a set of controlled psychophysical experiments on tactile discrimination [50]. Vibratese was composed of 45 basic elements – the tactile equivalent of numerals and letters – which were the intersection of three dimensions carefully chosen for their high tactile discriminability: signal amplitude, duration, and locus of interaction. The entire English alphabet and numerals 0 to 9 could be communicated this way. Geldard et al. reported that with proper training, legibility rates of more than 60 words per minute (wpm) were possible for common prose samples – i.e., reading rates approaching three times that of expert Morse code. The early successes of Vibratese in a laboratory context exemplify how careful design based on knowledge about touch is invaluable to the development of artificial tactile communication. Unfortunately, to the best of our knowledge, Vibratese seems to have completely vanished.

Throughout the years, Geldard's early work on the development of communication through touch has inspired various researchers, a large part of which were focusing on the development of tactile aids for deaf people from the 1960's to the 1980's. Tactile aids aim at substituting a defective hearing channel for tactile displays capable of communicating ambient sounds (e.g., alarm, door bell, telephone ringing) or human speech. The reader is referred to [138] for a comprehensive review of the systems that were developed in and prior to those times. Interestingly, the challenges that were identified by the community then are very similar to the ones we are faced with today when developing a tactile language – mainly, the limitations

of the technology and an ongoing unsatisfactory understanding of touch.

Noteworthy work from the 1980's on tactile communication for deaf people includes Brooks et al.'s tactile vocoder that they used to train subjects (both normal and deaf) to recognize tens of spoken common words (between 50 and 250) over training periods ranging from 24 hours to over 80 hours [17–19]. These successes can be explained partially by the use of an encoding mechanism better suited to human touch capabilities. Instead of directly converting audio frequency energy to mechanical skin stimulation (as for Gault's Teletactor covered above), Brooks et al. made use of a multi-channel encoding scheme by which speech frequency is displayed by the location of the stimuli on the skin and speech energy is communicated by stimulus amplitude.

Another tactile language developed for sensory substitution can be found in Braille. Invented more than 175 years ago, Braille still contributes to giving the blind access to the written word. The Braille alphabet consists of a series of tactile patterns that replace the sighted's printed characters. Each character (e.g., letters, numerals, punctuation) is composed of a 2-by-3 array of raised or of absent dots that were originally embossed on paper. Nowadays, Braille is also displayed by computer peripherals, called refreshable Braille displays, that make use of an extra 4^{th} row of dots to encode new characters such as "@". Despite the growing availability of other sensory substitution technology (e.g., speech synthesis), the Braille code remains an important access medium for the visually-impaired. Its success can be explained by a combination of factors. First, its physical characteristics (dot height, dot spacing, etc.) seem optimal for tactile discrimination [101, 120]. Second, it is relatively easy to produce. In fact, with a punch and portable template it is possible to use it to take notes. This played a great role in its early success. Finally, the existence of contracted Braille (made of abbreviations and contractions) makes it possible for expert users to read faster. With proper training, expert readers can reach reading rates of 100 to 150 wpm for normal Braille – about half the speed typically reported for print reading – and up to 190 wpm with contracted Braille [101]. Unfortunately, proficiency in Braille reading can only be achieved through extended amount of training (between 4 to 24 months).

Another example of a successful tactile communication system can be found in Tadoma, though it is much less widespread than Braille. Tadoma is a method used by a few deaf-blind to converse among each other and with those who can hear. It involves using the hands to monitor the lips movement and the vibrations from the vocal cords of the speaker by touching her or his face [127]. Tadoma experts are capable of high performance at understanding speech.

Haptic displays that communicate speech artificially, such as the ones mentioned above, are still limited at delivering information and require extensive training. One exception is the Tactuator device, designed by Tan et al., whose design goal was precisely to achieve high information transfer with a set of tactile stimuli that could be learned with minimum training [150]. The Tactuator is a haptic display capable of single-contact kinesthetic and tactile feedback to the thumb, the index finger and the middle finger. Tan et al. developed stimuli that were carefully designed for optimal discriminability. Based on results from two identification experiments, they estimated that the Tactuator was capable of achieving maximum information transmission rates of 12 bits/sec — i.e., comparable to that of Tadoma.[2]

2.2.2 Touch Iconography

Geldard's early efforts to develop Vibratese were partially motivated by the need to alleviate the cognitive load imposed on the visual and auditory channels by our environment. In his work, he also suggested using touch as a medium to grab a user's attention or to communicate a sense of urgency. Interestingly, his ideas are more pertinent than ever in today's modern world where we are constantly bombarded with visual and auditory information from a panoply of technological devices. In this computerized age, our attention is in constant demand. Whether we are working on our personal computer, trying to keep track of a busy schedule with a personal digital

[2]By considering the communication device and the human observer as connected systems, it is possible to make measurements about the amount of information that is transmitted. The amount of input information corresponds to the physical characteristics of the device's output stimuli. The amount of output information is related to what the human observer perceives. Finally, the amount of transmitted information is expressed by how many stimuli can be discriminated (i.e., the overlap between the input and output information) [102].

assistant (PDA) or driving a car, we are continually requested to shift our attention to information that must be attended to immediately: an icon on the computer suddenly appears to signify that a new email was just received; a PDA starts ringing to indicate that there are five minutes left before the next meeting; an icon on the dashboard is flashing to inform the driver that the gas tank is almost empty! With the explosion of the number of features and functionalities that has accompanied the increasing availability of visual and auditory interfaces, the amount and complexity of the information communicated through these sensorial channels have reached levels never attained before. Such pressure on our visual and auditory modalities saturates them to the point that important information does not get communicated anymore.

For the reasons stated above, Geldard's idea of using touch as a communication channel has recently been revisited with the development of artificial tactile patterns such as haptic icons [97] and tactons [15, 16]. Haptic icons are short haptic or tactile signals rich in contextual information that are typically delivered via simple electromechanical means. These signals may have varying degrees of structural complexity. They share with their graphical and auditory counterparts the function of communicating low-level, abstract information such as the state or function of an object or the occurrence of an event. For their part, tactons are defined by Brewster and Brown as "structured, abstract messages that communicate complex concepts non-visually [to the skin]" [12, 16]. They are, therefore, quite similar to haptic icons with the difference that they result from the general philosophy behind auditory icons ("earcons") and make use of concepts typically associated with music and speech synthesis (e.g., rhythm, vibration, pitch) [132].

The sought attributes for haptic icons are not different from those of their visual or auditory counterparts. First and foremost, haptic icons should be practical, reliable, quick to identify and pleasant to the tactile sense without being too distracting. Therefore, they should be designed with consideration for context and task: they should not get in the way of the goal but should support it. In order to bring added-value to an interaction, haptic icons must be easy to learn and memorize; they must carry evocative meaning or at least convey a discernible emotional content. Finally, they should be universal and intuitive, while, at the same time, support increasing

levels of abstraction as users become expert through repeated use.[3]

Recently, MacLean et al. have started applying and developing tools and methods to measure the discriminability of haptic icons in order to inform their design [25, 95, 97, 118, 152]. This research addresses questions such as: *What constitutes a meaningful artificial tactile signal? How should it be generated and delivered? What role does attention play on tactile perception? How many haptic icons can an average user remember?* Answers to these questions are starting to emerge. For instance, a recent study by Chan et al. found that seven vibrotactile icons could easily be learned in the absence of workload and with minimum training [25]. It also demonstrated how workload conditions, simulated with visual and auditory distractors, can significantly affect the time it takes for a user to detect a transition between two different icons. Taken together, these findings are of major importance to the development of haptic iconography.

2.3 Encoding of Tactile Information in Humans

From the time the skin is stimulated (e.g., by a pointy object) to the resulting perception (e.g., a localized prickly sensation), a variety of complex mechanical, perceptual and cognitive phenomena take place. At first, skin undergoes deformation that is mediated by diverse mechanoreceptors underneath the surface. Next, these receptors encode and transmit the stimulus to the central nervous system where it is integrated and relayed to increasingly higher levels of brain processing for interpretation. Psychological factors, such as attention and emotion also play an important role in the perceived sensation. Designing for meaningful tactile communication should be guided by a broad and integrated knowledge of how tactile information is encoded, transmitted, and processed at various stages of a tactile interaction.

[3] This process is referred to as "chunking" [22] and is well illustrated by the successes of contracted Braille with expert readers.

2.3.1 Modes of Skin Stimulation

A good understanding of the different modes of mechanical stimulation afforded by the skin sets the ground for a wide range of skin interactions; the skin can be tapped, vibrated, stretched, compressed, indented, heated, and more. While designing for a tactile display capable of multiple modes of stimulation, Kim et al. have argued that touch sensations should be restored in upper extremity amputees in the following order: (i) contact, (ii) pressure (iii) vibration (iv) temperature, (v) shear force and (vi) fine shape discrimination [83]. It is interesting to note, however, that the ordering of this list only applies to the field of rehabilitation. When designing for mobile interactions, a tactile display capable of reproducing fine shapes is likely to be more useful than a temperature display.

While not unique among all the perceptual systems in its property to convert external mechanical energy to internal neural impulses – the eardrum also carries out a similar function – the somatosensory system is the only one that offers such a wide surface for interaction. Skin is the largest organ of the body and covers almost 2 m^2 in an average adult. Therefore, in addition to the basic engineering attributes that are typically considered for the coding of artificial perceptual information (e.g., amplitude, frequency, duration, resolution and signal waveform), an extra dimensionality can be found in varying the locus of interaction [50, 73, 74, 162]. This is typically achieved by using a collection of tactile stimulators that are distributed across the human body or sections of it. Bach-y-Rita et al.'s tactile television apparatus is an early example of such an approach [2]. Their system consisted of a 2D array of vibrotactile actuators that was mounted to the rear of a wheelchair to allow the projection of visual information by tactile means. Visually-impaired users were submitted to various vibrotactile patterns projected onto the skin of their back. Complex patterns could be communicated this way to trained users.

The Optacon, a 1970s commercial sensory aid for the blind, is another example of an early tactile display capable of encoding information by varying the locus of interaction, but with an important distinction. Instead of stimulating the hairy skin of the back, the Optacon activates the receptors located in the glabrous skin of the

fingertip (more on this in section 2.4.1). Hairy skin and glabrous skin share many physiological and perceptual properties. However, the difference between their respective spatial acuity justifies considering the two perceptual sub-systems separately when designing for tactile communication. While the use of vibrotactile stimulation in its current technological form is likely to exploit most of the perceptual capabilities of hairy skin, it is not sufficient to encode all the different percepts that glabrous skin can detect (e.g., fine texture, small disparities, etc.). For the purpose of this thesis, we are mainly interested in the tactile stimulation of the glabrous skin of the hand which is in direct contact with the mobile device. Nevertheless, knowledge about tactile sensations in the hairy skin will be addressed whenever necessary. For now, the reader is referred to [75] for a review of the perception and generation of vibrotactile stimuli on the hairy skin.

2.3.2 Skin Anatomy and Neurophysiology

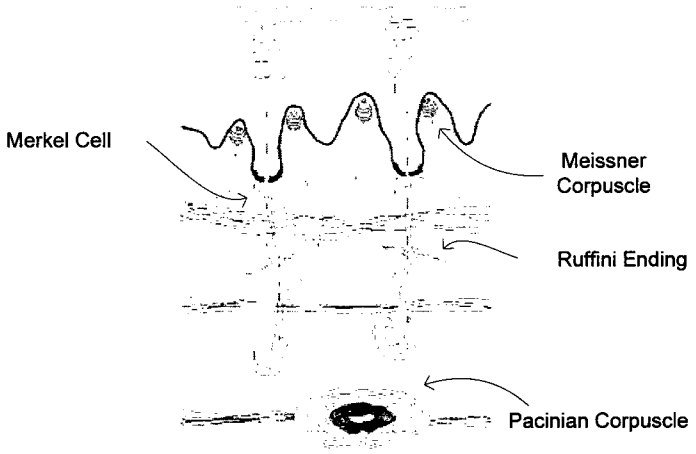

Fig. 2.1 Mechanoreceptors of the glabrous skin (drawing adapted from [1]).

Microneurographic studies in humans and monkeys have revealed the presence of four types of skin mechanoreceptor afferents in the glabrous skin (Fig. 2.1). These

are characterized by the size of their receptive field (type I for small and well-defined borders and type II for large and poorly-defined borders) and their adaptation rate to a stimulus (Slowly Adapting [SA] and Rapidly Adapting [RA]). Collected evidence indicates that Merkel cells (SA-I), which innervate the fingertip skin at 100 units/cm^2 in humans, are mainly responsible for the detection and identification of spatial patterns such as Braille dots and sharp edges [71, 72, 121]. Meissner corpuscles (RA-I), which are even more densely packed in the human finger (150 units/cm^2), are thought to be only, but highly, sensitive to dynamic skin deformation over a wide and uniform receptive field. Unlike Merkel cells, Meissner corpuscles poorly resolve spatial information but rather account for the detection and the neural encoding of skin motion. They detect low frequency vibrations and are responsible for signaling rapid state changes used for the accurate control of grip forces in prehension. Pacinian corpuscles (RA-II) are extremely sensitive to the smallest skin motion (in nanometers) and are mostly responsible for the perception of high frequency stimuli with peak responses between 200-300 Hz [71]. Therefore, they probably account for the remote perception of an object via a tool. Finally, the role of the Ruffini endings is still widely unknown, a puzzle that is accentuated by the fact that Ruffini endings were not found, neither in the glabrous skin of monkeys, nor in the human fingerpad skin [105, 111, 113].

Knowledge about the neural mechanisms that govern the peripheral encoding of tactile information is also limited. Phillips and Johnson listed four candidates for coding texture [122]: *spatial* codes, *temporal* codes, *spatiotemporal* codes and codes based on *intensity*. The results of psychophysical and neurophysiological experiments allowed them to confirm that SA-I units are most likely to resolve spatial detail [72, 183]. Smith et al. believe that both spatial and temporal codes are used to encode roughness [141]. Results from a roughness discrimination experiment with lubrication suggest that the rate of variation in tangential stroking force is important in the subjective determination of roughness. This finding is consistent with the presence of a temporal encoding mechanism. Nevertheless, Smith et al. also noted that any roughness estimation requires a minimum surface contact between the skin and the texture, an observation that supports spatial coding. Models of tactile in-

formation transmission, such as the ones above, are typically based on variations in mechanoreceptor firing rates. However, more recently, Johansson and Birznieks have put forward the idea that complex skin spatial events might be coded by the sequence in which different afferents initially discharge [69].

2.3.3 Skin Biomechanics

Results from a recent in vivo study of the skin under local tangential deformation indicate strong non-linear skin properties such as hysterisis and creep [170]. Different research groups have tried to develop models for the biomechanics of the skin to infer the resulting mechanistic behavior of the various mechanoreceptors located underneath the skin. One prominent model was put forward by van Doren who compared it to measured data obtained through a psychophysical experiment on tactile sensitivities [160]. The best match between his modeled and measured sensitivities was for predicted sensitivities based on *normal strain* at the mechanoreceptor levels. This is similar to Phillips and Johnson's finding of a model that is based on *maximum compressive strain* and that predicted spatial discharge-rate profiles measured in one type of afferents (SA-I) [121]. Others have suggested that *strain energy density* might be a better candidate for the encoding of shape, at least for SA-I receptors during static tactile sensing [144]. Despite the successes in matching models to psychophysical data, it should be stressed that biomechanical models of the skin make use of idealistic assumptions such as a linear visco-elastic, homogeneous, incompressible skin medium. In reality, skin is composed of different layers and elements (e.g., ridges, epidermis, dermis, papillae) that display a variety of complex biomechanical characteristics

Biggs and Srinivasan have compared tangential deformation to indentation of the skin [11]. They used a basic continuum mechanics model to predict the strain energy density at the mechanoreceptors level, for both the hairy skin of the forearm and the glabrous skin of the fingertip. In a complementary experiment, subjects were asked to adjust the magnitude of a skin tangential force so that it was perceived at the same intensity as a reference force applied normally to the skin. Interestingly,

the predictions from the model roughly matched data from the psychophysical experiment: results showed a higher sensitivity to normal forces than tangential forces at the fingerpad, and the opposite at the forearm. Based on these findings, Biggs and Srinivasan questioned the effectiveness of making use of tangential displacement of the skin for tactile displays. They concluded that, despite their interesting properties, fingertip tactile displays exhibiting tangential skin stimulation would suffer from serious drawbacks such as the inherent larger stiffness of the skin to tangential stimulation.

2.3.4 Psychology of Touch

Simple psychophysical experiments have unambiguously demonstrated the great subtleties and capabilities of the tactile system. Mechanoreceptors are capable of both detecting very fine tactile features [145] and conveying crucial kinesthetic information, such as finger joint position [40] or contact forces to the brain [55]. Similarly, the importance of skin tangential forces for object manipulation has long been demonstrated [112, 141]. It is well known that to prevent slips when holding an object, normal forces to the grip surface are applied in reaction to variations in the tangential forces sensed at the skin level.

One debate that persists, however, is about the perceptual and physiological differences between active and passive touch. Active touch refers to the exploratory action of touching, whereas passive touch describes a stimulation of the skin brought about by some outside agent [54]. On one side, a few experiments tend to support the superiority of active touch over passive touch (e.g., [58]). On the other, concerns have been raised as to whether these experiments did really constitute a fair comparison because they neglected to provide equivalent information in both modes [128]. Vega-Bermudez et al. found no significant difference in performance between the passive and active tasks of recognizing tactile letters, which suggests that the sensory neural mechanisms underlying both exploration modes are identical [164]. Surprisingly, contrary evidence suggests there does exist a phenomenon known as *gating* by which the transmission of tactile inputs to the primary somatosensory cortex is decreased

during active exploration. Chapman notes that the effects of gating during active touch are likely compensated by other mechanisms that can enhance performance, such as attention and hand movement [27]. This could explain why a superiority of passive touch is rarely reported in the literature (see [98] for an unusual exception).

Various space and time interactions among tactile stimuli, and their effects on perception, are commonly reported (refer to [52, 137] for overviews of the most important ones):

- *Masking* is a phenomenon by which the performance at identifying a target stimulus is decreased by the prior or subsequent presentation of a masker stimulus [30, 35–37]. To reduce the undesirable effects of temporal masking, Craig suggests increasing the interval between two successive stimuli [35]. On the other hand, he also notes that this can only result in lower rates of tactile communication since masking is related to the time interval between the onsets of the target and the masker. Similarly, increasing the spatial distance between the masker and the target, such as displaying them on two different fingers, will likely decrease the effects of masking; however, it will also introduce undesired outcomes due to the extra attentional load imposed by having to concentrate on both stimulation sites simultaneously.

- Vibrotactile *adaptation*, or the tendency for sensitivity to decline with prior exposure to a vibratory stimulation above threshold, is yet another example of tactile interaction. Adaptation is clearly reported by a handful of psychophysical experiments that either found an increase of the sensitivity threshold or a decrease of the perceived intensity following the exposure to a conditioning vibrotactile stimulus [6, 166, 167]. Fortunately, the effect is not permanent and proper time gaps between the conditioning stimulus and the target can avoid it all together. Accumulated evidence suggests that neural adaptation takes place both at the mechanoreceptors level (i.e., expressed by a decrease in firing rates), and at higher perceptual levels [6].

- Vibrotactile *enhancement* is a well-reported time interaction with an effect that is opposite to that of adaptation [53, 166]. It is expressed by an increase

in the magnitude estimation of a vibrotactile target stimulus following the presentation of a conditioning stimulus with significantly higher magnitude.

- The tactile equivalent of visual *change blindness* has also been recently observed with vibrotactile stimuli [46,47]. Change blindness is manifested by the failure to detect change in a tactile pattern that is presented repeatedly in between stimuli.

Interactions between stimuli have important implications for the future design and implementation of tactile displays, but they are not utterly drawbacks to the conception of an artificial tactile language. Some researchers have suggested taking advantage of well-known tactile phenomena to compensate for the technological limitations of tactile displays. Technological constraints make it difficult to pack miniature actuators in a tactile display densely enough to match the fingertip's spatial resolution. To this effect, it was suggested that the *saltation effect* (or "rabbit" effect) be used to convey the sensation of motion in between actuators of a tactile display with limited spatial density [52,67,151]. The "rabbit" effect is the sensation that a stimulus is progressively "jumping" between two stimulation points. It is generated by delivering a series of successive taps, first at the departure point and then at the arrival point. The resulting illusion of smooth motion between the two points is strong and convincing but requires that the stimuli be delivered with tight temporal control.

2.4 Distributed Tactile Displays

Numerous research groups around the world have tackled the problem of building useful and cost-effective tactile displays that allow for rich tactile communication. The challenge is not a trivial one. Simplified tactile displays already play a role in tactile communication, as seen by the widespread use of vibrating hand-held devices (e.g., cell phones and PDAs). These devices rely on a single actuator (e.g., on/off eccentric rotating motor, voice coil);[4] consequently, their expressive capabilities remain

[4]See Audiological Engineering Corp. Tactaid^{TM} – www.tactaid.com.

limited. Rich tactile communication, on the other hand, requires *distributed* displays comprising high-density arrays of high-performance miniature actuators. There is no alternative to reproducing the complex patterns of skin deformation that occur at the fingertip when we touch an object (e.g., [91]). This situation is analogous to visual signal communication. While simple abstract communication can be achieved with a single light source, e.g. a flashing light to signal Morse code, a larger amount of information can be transmitted with a two dimensional array of pixels (i.e., a field). Given the current state of the technology and the scale of the actuation mechanisms, packing a high density of individually controlled actuators on a small surface remains a challenge. By comparison, screen technology is much more mature.

2.4.1 State of the Art

Most attempts to date at building distributed tactile displays have focused on devices that stimulate the fingertip because of its high tactile acuity.[5] Some have considered other regions (e.g., the tongue and mouth [3, 154], the back [2], the torso [134, 163], the thighs [32]). Generally, distributed tactile displays operate by indenting the skin with arrays of pins that raise out of a surface in order to create a discrete representation of a texture or a small-scale shape [84, 119, 161]. Other techniques make use of actuators that vibrate [23, 62, 147], that heat up [59, 64], that blow pressured air [1], that change shape when submitted to an electric or magnetic field [8, 156], that create small currents through the skin [78, 79], and more (refer to Fig. 2.2). Among the most employed actuator technologies, we find shape memory alloys (SMA), piezoelectric ceramics, motors, pneumatic valves, Peltier elements, rheological fluids, pistons, electrodes, and others. Most of the time, these actuators are used to apply direct deformation of the skin. They are, therefore, subject to high force and size requirements because the skin is a tough medium. For this reason, it has also been suggested to make use of actuator mechanisms that do no active work against the fingertip skin, but instead lock skin contactors into passive mechanical states. This approach offers

[5] The lips and tongue are more sensitive to tactile stimuli than the fingertip but, by far, they do not have its spatial resolution. Moreover, using these organs for tactile communication is impractical and can even be inappropriate in some situations.

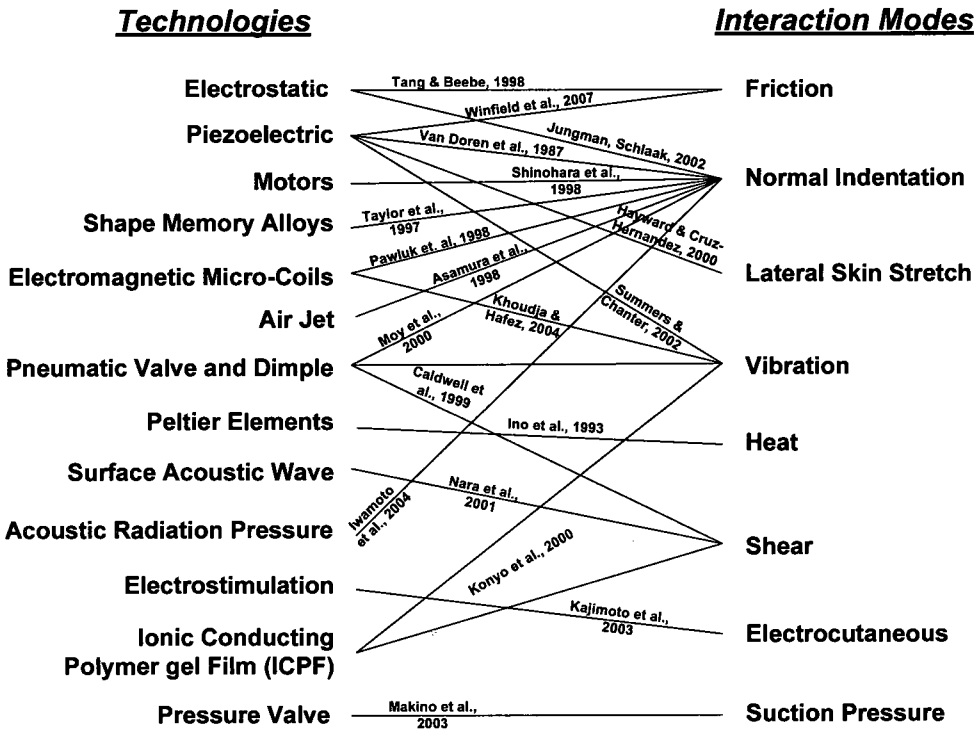

Fig. 2.2 Examples of combinations of modes of interaction and actuator technologies for the design of tactile displays.

potential for smaller actuators but it also has the drawback of requiring users to actively explore the actuated surface [155].

A number of technologies and modes of interaction have been prototyped, each combination yielding its own set of properties (The reader is referred to table A.1 of the appendix for a summary of the advantages and drawbacks of each technology). Recently, hybrid systems that make use of more than one technology and/or interaction mode have started to emerge [4, 83, 180, 185]. Unfortunately, most of the distributed tactile displays for the fingertip built to this day fail to convey meaningful tactile information and to be practical at the same time. The devices are too bulky or do not provide enough force to deform the skin. They are often constrained to a low bandwidth or are simply limited to a small number of actuators with low-spatial density. Lastly, they require constant maintenance or are too complex to operate most of the time.

For the reasons stated above, distributed tactile displays have seldom made it to the commercial market. An exception is the Optacon, a device manufactured by Telesensory Corporation. The Optacon was composed of a 24-by-6 array of vibrating pins on which users would lay down their left index fingertip in order to read printed material [14, 34]. Each pin of the array could be made to vibrate at a fixed frequency (around 230 Hz) or kept idle by the device's control system. The system also included a camera to allow the real-time conversion of optical information into an equivalent vibrotactile pattern. Typically, visually-impaired users would scan printed text with the camera probe in their right hand and feel the resulting tactile image under their left fingertip. Reading required dozens of hours of practice, but the device was embraced enthusiastically by a large part of the blind community. With proper training, reading rates could reach 50-100 wpm. Unfortunately, manufacturing of the Optacon was eventually discontinued because it was not commercially viable.

Another early inspirational example of a tactile display can be found in [31], where Cholewiak and Sherrick noted that "[f]or a given waveform (e.g., sinusoidal bursts, rectangular pulses, or haversine pulses), there exists no qualitative difference between tangential and perpendicular stimulation". Their device was composed of a vibrotactile matrix made of small piezoelectric benders that stimulated the palm

of the hand at a fixed frequency (250 Hz). The system provided for a simultaneous refresh of all of the elements in the array with each one being set to one of 64 possible activation intensities. Cholewiak and Sherrick's ingenious device allowed them to experiment with the delivery of complex spatiotemporal patterns to the skin, but only vibrational patterns could be displayed.

The transducer that is described in this thesis also makes use of an array of vertically mounted piezoelectric benders (bimorphs) to generate lateral skin stretch. It is capable of more than vibrotactile stimulation. Users typically lay their index finger directly on the top of the bimorph array. Applying individual control voltages across the piezoelectric bimorphs makes them bend at their tip, which in turn creates a distributed tangential deformation of the users' fingerpad skin. The bending of the bimorphs is controlled by a personal computer, which makes it possible to deliver programmable dynamic patterns of lateral deformation to the skin.

2.5 Summary

In 1960, Geldard noted that "almost any certain fact about somesthetic functioning is likely to prove valuable [to the development of tactile means for communication]" [51]. Unfortunately, as illustrated in section 2.3, the exact fundamental mechanisms governing the sense of tactile touch remain open for debate. This reality makes it difficult to infer practical considerations for the design of tactile stimuli and tactile displays. Designing an optimal tactile display requires a precise knowledge of the somatosensory system, which in turn, is one of the aspects we are trying to study with the device. Hence, we are caught in a vicious circle. Nevertheless, observations and insights collected from both the touch literature and personal experience at building tactile displays for the past few years should still guide our efforts.

Touch can be an effective means for communication. This is demonstrated by the long lasting successes of Braille (and to some extent Tadoma). Over the years, Braille has proved invaluable to the blind community. Unfortunately, the years of training that are required to master Braille make it a language that cannot be accessible to everyone. Undoubtedly, there is room and potential for other means of artificial

tactile communication that do not require as much training and that can appeal to the masses. As noted by Geldard when he was developing the Vibratese system: "coding to letters and numerals is really a quite pedestrian way of getting meaning into tactile patterns" [51]. This suggests that the design of an artificial tactile language should be inspired by the wide diversity of rich tactile interactions that we experience with the world on a daily basis. To this effect, some researchers have started the study and development of a universal iconography for touch that aims, among other things, at alleviating the current cognitive load imposed by modern technology.

Skin is a highly non-linear medium that acts as a filter between the stimulus and the mechanoreceptors. While the exact roles and functions of the different skin mechanoreceptors remain unclear, something is known about the psychology of touch; for instance, the effects of some complex interactions between tactile stimuli presented successively have been identified (e.g., masking, adaptation). These phenomena can significantly impair the performance of tactile communication if they are not well understood or are simply ignored. Conversely, they can also be exploited to compensate for some limitations of the technology. Therefore, knowledge about the limits and capabilities of the somatosensory system should go hand in hand with the design of tactile displays.

Rich and natural tactile communication that goes beyond single-point stimulation can only come from tactile displays that are capable of distributed interaction. Over the years numerous designs and technologies have been prototyped and some show potential but to this day none has succeeded at being both useful and practical. Failure to date can be partially explained by the technological challenge of having to pack a high number of fragile electromechanical actuators onto a small surface the size of a fingerpad.

Overcoming the limitations of actuator technology is not sufficient to guarantee the achievement of a an artificial tactile language that is both usable and practical. Work on artificial tactile feedback also has to address the challenge of matching the tactile interactions to the tasks they are trying to augment.

Chapter 3

Display of Virtual Braille Dots by Lateral Skin Deformation: Feasibility Study

Preface to Chapter 3

This chapter appeared in

> Levesque, V., Pasquero, J., Hayward V., and Legault, M., Display Of Virtual Braille Dots By Lateral Skin Deformation: Feasibility Study, *ACM Transactions on Applied Perception*, Vol. 2, No. 2, pp. 132-149, 2005

which is itself an extended version of

> Pasquero, J., Levesque, V., Hayward V., and Legault, M., Display of Virtual Braille Dots by Lateral Skin Deformation: A Pilot Study, *Proc. Eurohaptics 2004*, Munich, Germany, June 5-7. pp. 96-103, 2004.

This chapter lays the ground for the rest of the thesis. It introduces an early electromechanical prototype capable of displaying tactile sensations to users by means of a unique skin stimulation method. Results obtained from a preliminary study conducted with the prototype bear fundamental significance as they will repeatedly inform design decisions throughout the remaining work described in this thesis.

Contributions of Authors

The work described in this chapter is the result of a tight collaboration between Levesque and Pasquero. They were both responsible for the experimental design and the characterization of the device. They also contributed equally to running the experiments, analyzing the data and generating the figures and tables. Most of the work on the design and development of the hardware can be attributed to Pasquero. The software and tactile rendering was mostly the responsibility of Levesque. Prof. Hayward supervised the work and edited the manuscript. Finally Legault provided expertise on Braille reading and on designing for the visually impaired.

Display of Virtual Braille Dots by Lateral Skin Deformation: Feasibility Study

Abstract

When a progressive wave of localized deformations occurs tangentially on the fingerpad skin, one typically experiences the illusion of a small object sliding on it. This effect was investigated because of its potential application to the display of Braille. A device was constructed that could produce such deformation patterns along a line. Blind subjects' ability to read truncated Braille characters ('oo', 'o•', '•o', and '••') using the device was experimentally tested and compared to their performance with a conventional Braille medium. While subjects could identify two-character strings with a high rate of success, several factors need to be addressed before a display based on this principle can become practical.

3.1 Introduction

3.1.1 Braille Displays

Louis Braille's reading and writing system has given the blind access to the written word since the early 19^{th} century. Braille characters replace the sighted's written letters with tactile equivalents. In the Braille alphabet, each character consists of an array of two columns and three rows of raised, or absent dots. Traditionally embossed on paper, Braille has more recently also been provided by refreshable Braille displays that generally add a fourth row of dots. Refreshable Braille displays were initially the only type of computer interface available for the blind. Despite the growing popularity of more affordable speech synthesis hardware and software, refreshable Braille remains a primary or secondary access medium for many blind computer users.

Commercially available refreshable Braille displays have changed little in the past 25 years. Today's displays do not differ substantially from what is described by [158]. Typical systems use cantilevered bimorph piezo-actuators (reeds) supporting vertical pins at their free end. Upon activation, a reed bends, lifting the pin upward. Braille characters are displayed by assembling six or eight of these mechanisms inside a package called a cell (see Figure 3.1(a)). A basic system includes 40 or 80 cells to display a line of text (Figure 3.1(b)), plus switches to navigate in a page (Figure 3.1(c)) [146].

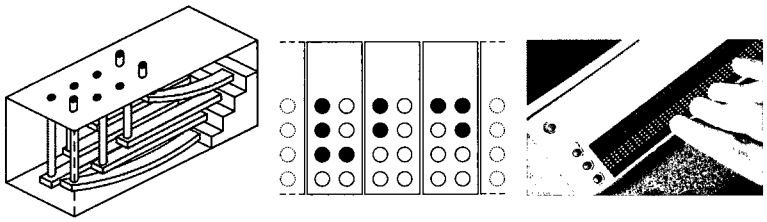

(a) (b) (c)

Fig. 3.1 Conventional Braille display: (a) cell actuation mechanism, (b) array of cells, and (c) picture of a commercially available Braille display (used with permission from Pulse Data International Ltd).

While the elements of these cells are simple and inexpensive, the cost is driven by the necessity to replicate the cell 40 or 80 times, or more if one contemplates the display of a full page. Typical Braille displays cost much more than a personal computer.

3.1.2 Alternative Technologies

In recent years many alternate designs have been proposed, all sharing the principle of raising individual pins, or dome shapes, out of a surface [13, 129, 182]. In 2004 alone, no less than six U.S. patents related to Braille cells have been granted, and many others are pending [e.g., [9, 125, 181]]. While most of the research focuses on reducing the cost of actuation, very little work is concerned with new approaches to the display of Braille.

Of note is a system proposed by [153] who sandwiched discrete electrodes in a dielectric. The application of high voltage to these electrodes causes the skin to adhere locally to a glassy surface, thereby creating small tactile objects. Patterns resembling Braille characters could presumably be displayed with this method, however it appears to suffer from sensitivity to environmental factors such as humidity or skin condition.

Several investigators proposed the idea of a single display moving with the scanning finger rather than the finger scanning over an array of cells. [44] mounted a single Braille cell on a rail and activated its pins with waveforms resembling "pink noise" in an attempt to imitate the effect of friction of the skin with a pin. [126] designed an experiment with a Braille cell used in conjunction with a planar "Pantograph" haptic device in an attempt to dissociate character localization from character recognition. The haptic device was programmed to indicate the location of the characters in a page, while the cell was used to read individual characters. Comparative tests were performed in different conditions with one or two hands. Again, the goal was to create an "array of Braille characters" with a single cell and reasonable reading performance could be achieved.

3.1.3 Overview

This paper reports on a feasibility study conducted to evaluate the potential of a new approach to the refreshable display of Braille. When the skin of the fingertip is locally deformed in the manner of a progressive wave, one typically experiences the illusion of objects sliding on the skin, even if the deformation contains no normal deflection [57]. An electromechanical transducer was designed to create such skin deformation patterns with a view to investigate the feasibility of displaying Braille dots. The novelty of this approach lies in that it creates a progressive wave of lateral skin deformation, instead of a wave of normal indentation [e.g., [161]] or localized vibration [e.g., [14]]. Our approach also relies on scanning motion, which is often mentioned as necessary to "refresh" the skin receptors and combat adaptation [44].

The transducer that we constructed was similar in principle to the 'STRESS' display [116], but had only one line of actuated contactors. This configuration allowed us to significantly increase the forces and displacements produced by the contactors. The Braille dots created by this device were "virtual" in that we attempted to recreate only the essential aspects of the skin deformation occurring when brushing against raised dots without actual physical dots.

The resulting system and the particular strain pattern — collectively termed "VBD" for Virtual Braille Display — were empirically designed with the assistance of the fourth author, a blind accessibility specialist, who also participated in the study in the capacity of "reference subject".

An experiment was conducted to tune the pattern's parameters to create a sensation as similar as possible to that experienced when brushing against physical Braille dots. The legibility of strings of truncated Braille characters — those comprising a single row of dots — was evaluated with five Braille readers on the VBD, and on a conventional Braille medium (embossed vinyl). The subjects' success rate and reading patterns were recorded and analyzed.

The study shows that reading with the VBD is possible with a high legibility rate given some personalization of the strain pattern. Reading is, however, more demanding and error-prone than on conventional media. More importantly, the

study helped identify the strengths and weaknesses of the current prototype, and indicates how the device could be improved to yield a workable system.

3.2 Virtual Braille Display

3.2.1 Device

The VBD device consisted of a tactile display mounted on a laterally-moving friction-less slider (see Figure 3.2) and interfacing control electronics.

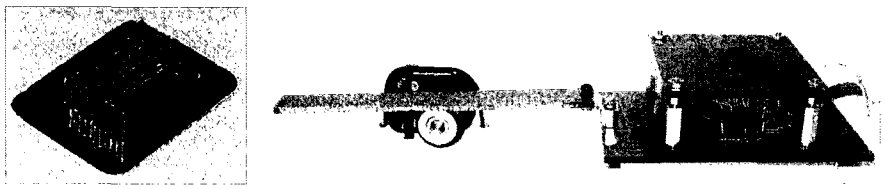

(a) (b)

Fig. 3.2 VBD device: (a) STRESS-type tactile display, and (b) display mounted on a slider with rotary encoder.

Reading virtual Braille was done by applying the tip of the index finger against the tactile display and sliding it laterally, as shown in Figure 3.3. When activated, the tactile display caused lateral deformation to the fingertip skin, that could be varied in response to slider movement. The finger remained in contact with the display and dragged it along the reading surface. Although this principle could allow reading with multiple fingers, the width of the display limited the reader to the use of a single finger.

Tactile Display

The tactile display was made of a stack of twelve 0.38-mm-thick piezoelectric bender plates,[1] sandwiched at their base between neoprene spacers and clamped between two rigid end-plates using four locating pins and four screws (see Figures 3.4(a)

[1] Y-poled, 31.8 mm x 12.7 mm, High Performance Bending Motors from Piezo Systems Inc., part number T215-H4CL-303Y.

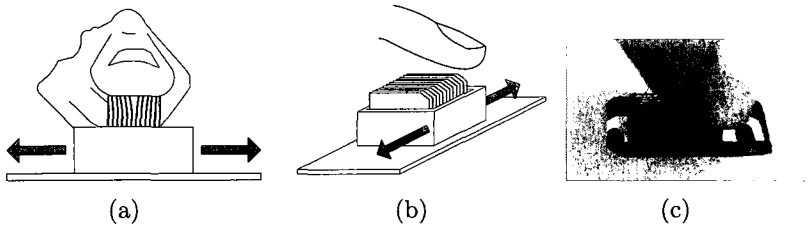

Fig. 3.3 Interaction with the VBD: (a) strain applied during exploration, (b) illustration and (c) picture of finger contact with the VBD.

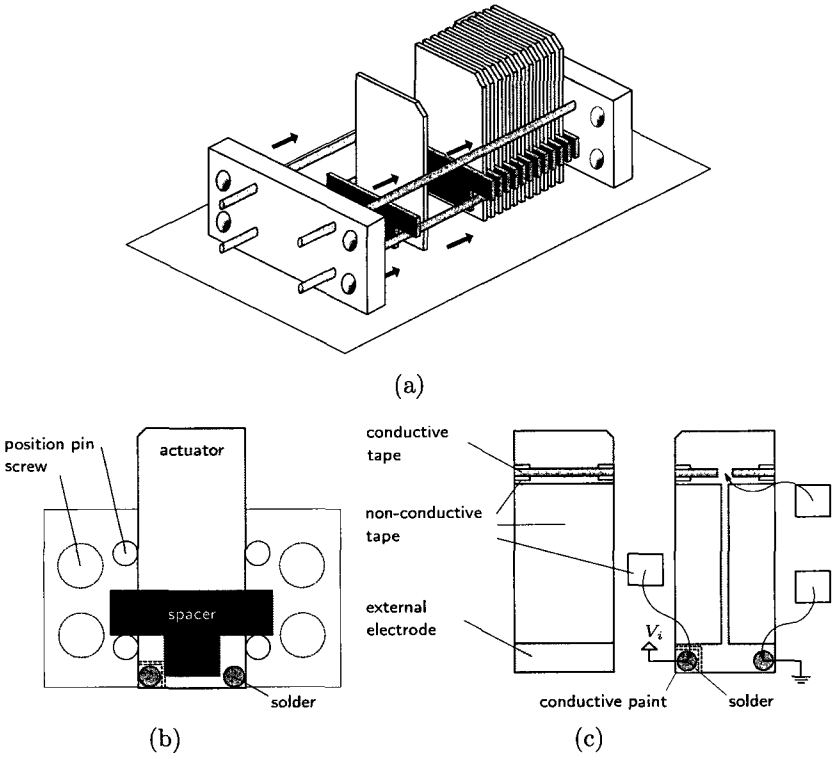

Fig. 3.4 Assembly of the VBD's tactile display: (a) perspective and (b) frontal views of stacked assembly, and (c) actuator fabrication process.

and 3.4(b)). The spacers were cut in a 12-mm-high T-shape so that they rested on the locating pins and allowed room for electrical connections (see Figure 3.4(b)). Once tightly secured, the spatial period ϵ, or contactor pitch, was approximately

0.7 mm. This assembly method was selected for the convenience of allowing the design parameters such as thickness, length, shape and material of the actuators and spacers to be changed. In the present study, however, only one configuration was used.

The actuators were driven by a bipolar voltage applied between their central electrode and their two electrically-connected external electrodes. Because of the small space between adjacent plates, the electrodes could not be connected using the methods recommended by the motor supplier. Therefore, the actuators were prepared as shown in Figure 3.4(c). The external electrodes were joined with adhesive electrically conductive tape[2] running over non-conductive tape on the sides to prevent shorting with the central electrode. One corner of an external electrode was soldered to a ground wire. The other corner was turned into a small electrode pad (isolated from the rest by grinding off the conductive layer) and connected to the central electrode with conductive paint. A wire used for the control voltage was then soldered to this pad. To prevent contact with the adjacent actuator, the solders were kept significantly thinner than the spacers (0.5 mm) and were protected with non-conductive tape. Traces of conductive paint were applied along the length of the electrodes to improve their reliability (not shown). The actuators were then isolated from the metallic locating pins using non-conductive tape. As illustrated in Figure 3.3(b), the top corners on one side of the blades were beveled to create a narrow linear array of skin contactors (around 0.2 mm^2 in area each). Finally, the tips of the actuators were coated with varnish to isolate them from the skin.

The display could be used by applying the finger either on the large horizontal contact surface or against the surface formed of the beveled corners of the contactors. The latter surface, as shown in Figure 3.3, provided a narrower contact area more appropriate for the display of dots and was the only one used in this study.

The deflection of the actuator tips was estimated with the help of a camera. Two sample measurements are shown in Figure 3.5. However, when not loaded by the finger, the deflection range was estimated to be 0.4 mm. As explained in the next section, the motion was limited in practice to a restricted range of approximately

[2]3M Corporation, EMI Copper Foil Shielding Tape 1181.

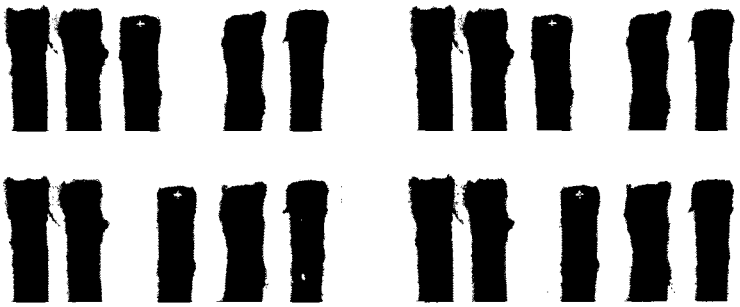

Fig. 3.5 Visual estimation of unloaded actuator deflection for (a) full range and (b) restricted range. Adjacent actuators were held deflected away from the actuator under study.

0.3 mm. The deflection when loaded with the fingerpad could not be quantified but appears to be significantly lower than when unloaded.

Control System

The position of the linear slider on which the tactile display was mounted was measured by an optical encoder with a nominal resolution of 17 μm. Interfacing electronics were constructed to permit the refresh of the actuators at 500 Hz according to patterns programmed on a personal computer. This enabled us to program the deflection of each actuator with arbitrary functions of space (see Section 3.2.2).

The interfacing electronics, adapted from a previous project, made use of a Field-Programmable Gate Array (FPGA) development board[3] with a Universal Serial Bus (USB) 1.1 interface. It was programmed to convert control frames coming from the computer, or "tactile images", into twelve bender voltages by means of 156-kHz pulse-width modulation (PWM). The same, however, could be accomplished by adopting a variety of other approaches, including the use of micro-controllers or dedicated logic, interfaced to the computer via parallel I/O or high-speed serial I/O.

The computer generated a set of 8-bit actuator control values based on the en-

[3]Constellation-10KE^{TM} from Nova Engineering Inc. operating an Altera FLEX 10KE^{TM} chip.

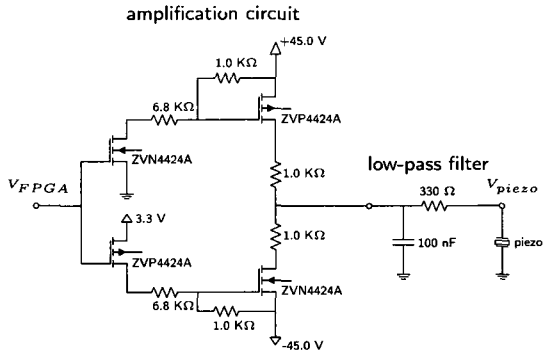

Fig. 3.6 Electronic circuits: amplification circuit (left) and low-pass filter (right).

coder readout every 2 ms on average. These tactile images were sent to the FPGA by packets of 5 through the USB channel where a FIFO buffer regulated the flow of tactile images to ensure a constant output rate. The logic-level signals were then amplified to a ±40 V range and low-passed by the circuit shown in Figure 3.6. In order to avoid non-linearities in the signal amplification at extreme PWM duty cycles, the control values were restricted to the range 10 (0x0A) to 250 (0xFA).

3.2.2 Skin Deformation Patterns

Trial and error led us to select a pattern solely on the basis of the resemblance of the sensation it provided compared to that of actual Braille, as felt by the reference subject. We are however unable to offer a principled explanation as to why this particular pattern creates sensations that resemble Braille dots more than others. The determination of the actual parameters is described in Section 3.3.

The deflection δ_i of each actuator i was a function of the slider position x_s obtained from the encoder. The actuators followed the same deflection function $\delta(x)$, where x was the actuator position along the reading surface, as illustrated in Figure 3.7. The physical configuration of actuators introduced a spatial phase difference of ϵ. The first actuator was given a position corresponding to the slider position.

$$\delta_i(x_s) = \delta(x_s + i\epsilon), i = 0, \ldots, 11 \tag{3.1}$$

What we selected was a pattern[4] such that the deflection of each actuator swept the first half-cycle of a sinusoid, starting from the left position, as it scanned a virtual dot, as shown in Figure 3.7. A small-amplitude, high-frequency sinusoid could also be added to the nominal waveform to enhance contrast. These representations were termed nominal and textured.

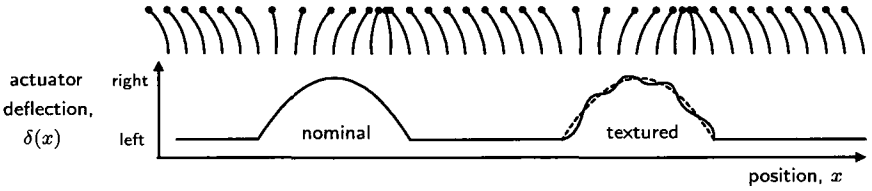

Fig. 3.7 Actuator deflection as a function of position. The curve shows the actuator deflection function with respect to actuator position for a nominal dot (left) and a textured dot (right). The deflection of actuators is illustrated at discrete points along the virtual reading surface. Texture was always applied either to all dots or none.

These patterns were found to better approximate the sensation of scanning over Braille dots than others that were experimented with, such as triangular or square waves, full-cycle sinusoids, or textured blanks.

The spatial phase difference between actuators resulted in the representation of dots as a traveling wave. Figure 3.8 and 3.9 illustrate the movement of actuators as a virtual dot traverses the length of the display. Moving the slider in one direction across a region containing a dot resulted in a wave of actuator deflections traveling at the same speed in the opposite direction on the tactile display, causing the illusion that the reading finger was scanning over stationary Braille dots. Since the deflection function was independent of direction, it caused actuator deflections in the direction of finger movement when reading from left to right, but opposing movement when reading from right to left. The resulting sensations, however, seemed to be similar.

This pattern had two distinct effects on the skin deformation. The first was to cause a net displacement of a skin region around each contactor. The second was a

[4]Movies of the VBD in action can be found online on the Haptics Laboratory's VBD web page, http://www.cim.mcgill.ca/~haptic/vbd.html.

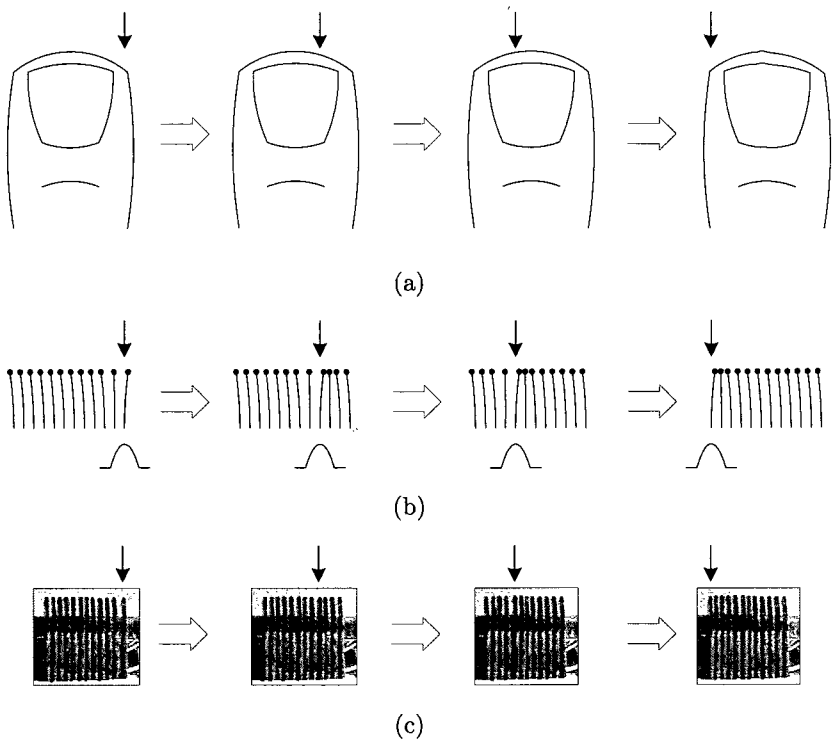

Fig. 3.8 Traveling wave representation of a Braille dot at four points in time: (a) finger-dot interaction, (b) depiction of the actuator deflection and corresponding deflection function, and (c) picture of the actuator deflection. The dot center is indicated by arrows.

pattern of compression and expansion of each small region of skin located between two contactors. Patterns of expansion and compression can actually be observed when a finger scans over small shapes [91]. The strain variations Δ_i caused by a pair of actuators is represented in Figure 3.10. The width ω of a virtual dot is shown relative to the spatial period ϵ. If $\omega < \epsilon$, then there was no overlap between the deflections of adjacent actuators. If $\omega > \epsilon$, then an overlap existed and there was a continuous transition from expansion to compression. If $\omega > 2\epsilon$ expansion and compression never reached their maximum values. It is not known whether local displacement or local variations in strain, or both forms of stimulation, caused the illusion of the dot moving under the finger.

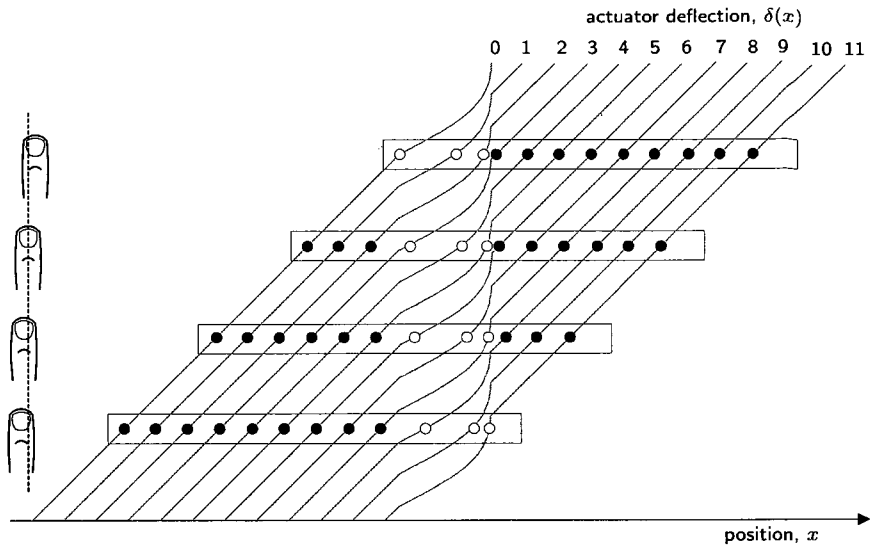

Fig. 3.9 Traveling wave corresponding to a single dot as the slider passes over it. Top-view of actuator deflections and corresponding finger-dot interactions are shown at four locations. Each of the twelve curves indicates the deflection pattern followed by an actuator as the slider moves from left to right.

The tactile display could only display a single row of Braille dots. From the Braille character set, the three characters that have dots in row 1 only, or a total absence of dots, could be displayed: 'a', 'c', and ' ', see Figure 3.11. The fourth possible combination, unused in Braille, was called 'dot #4'.

3.3 Parameter Tuning

Braille is normally produced according to strict geometrical specifications, but the manner in which these specifications translated into the VBD's parameters was not straightforward. For this reason, a first experiment was carried out to find the parameters that produced virtual Braille of appropriate dimensions. For the purposes of the feasibility study, only the width and the separation of dots were adjusted. The amplitude of the virtual dot sinusoid was set to the maximum that the system could provide. The amplitude and the wavelength of the texture were set empirically to

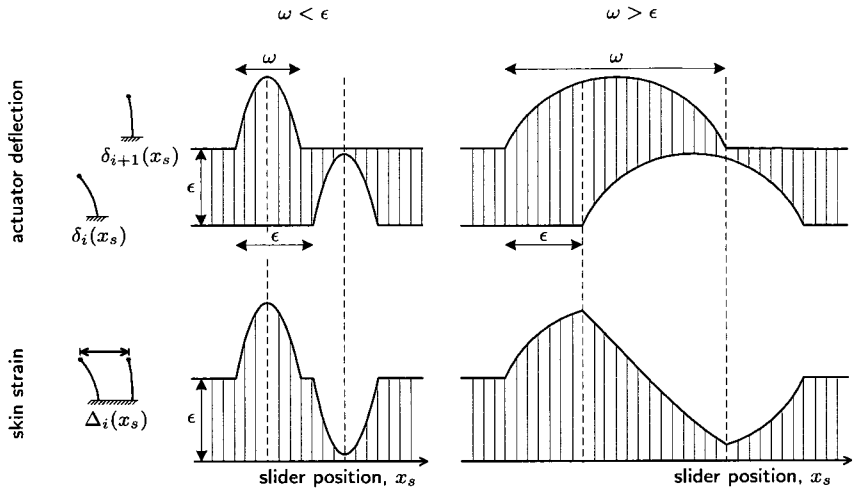

Fig. 3.10 Displacement of two consecutive actuators and corresponding skin strain patterns as functions of slider position, for width ω of virtual Braille dots smaller or greater than spatial period ϵ.

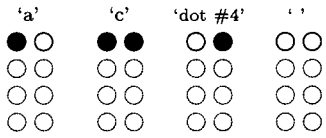

Fig. 3.11 Braille characters displayed by the VBD.

values equal to $1/8^{\text{th}}$ those of the virtual dot sinusoid.

3.3.1 Method

With the help of the reference subject, the width of a virtual dot was first adjusted to match the sensation caused by a single physical Braille dot. Then, the distance between dots within a virtual character was adjusted. The inter-character spacing value was inferred from the value found for intra-character dot spacing.

The tuning experiment was conducted following a 2-alternative forced choice protocol, using the two-hand method in order to speed up the process and facilitate comparisons. The subject was asked to touch a reference stimulus produced by a conventional refreshable Braille display with her left index. She then immediately

explored two stimuli on the VBD with her right index and selected the one that best matched the reference stimulus. Dots were always displayed with texture. After a short experimentation used to determine an appropriate range, the virtual dot width was varied among six equally spaced values from 0.5 mm to 3.0 mm. The intra-character dot spacing of the character 'c' (••) was similarly varied from 1.0 mm to 2.5 mm. The dot width found in the first step was used in the second.

3.3.2 Procedure

Both tuning experiments proceeded in the same manner. The reference subject moved the VBD toward the left, waited for an audible signal, explored the two different Braille dots or pairs of Braille dots, and verbally reported the stimulus that best matched the reference stimulus. Answers were logged by the experimenter. Each of the 30 possible ordered pairs of non-identical stimulus were presented to the subject 3 times, for a total of 90 trials. The different pairs of stimuli were presented in randomized order.

3.3.3 Results

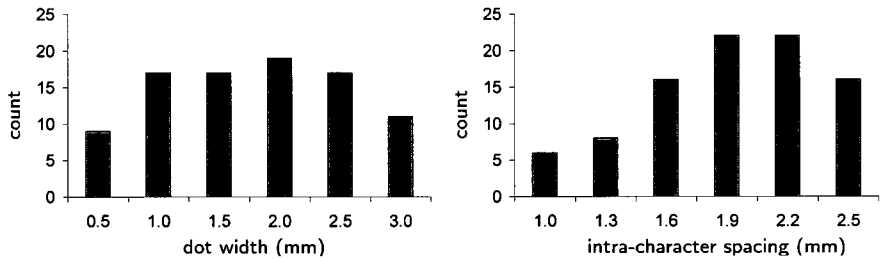

Fig. 3.12 Results of tuning steps: frequency distribution of (a) dot widths, and (b) intra-character dot spacings.

Figure 3.12 shows the results of the two tuning experiments. The preferred virtual dot width was found to be 2.0 mm, the most frequently selected parameter during

the first step. The virtual intra-character separation was also found to be 2.0 mm, between the two most frequently selected values during the second step.

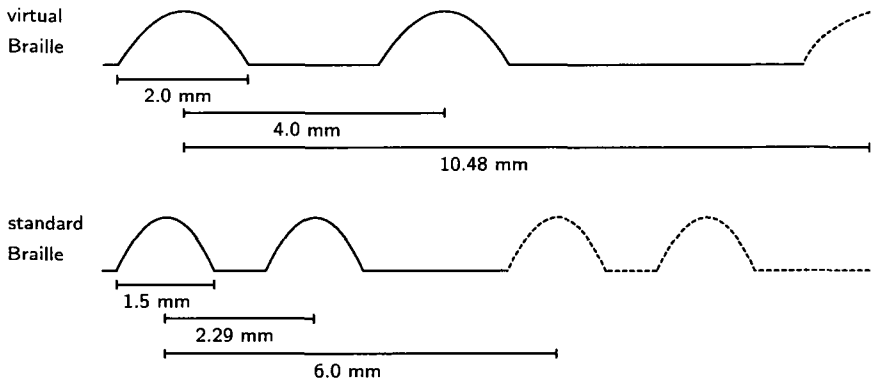

Fig. 3.13 Dimensions of virtual and standard English Braille.

These parameters corresponded to a intra-character dot spacing of 4 mm, compared to 2.29 mm for standard English Braille. The standard horizontal character-to-character distance of 6 mm was scaled accordingly to a virtual distance of 10.48 mm as illustrated by Figure 3.13.

3.3.4 Discussion

The preferred virtual dot width of 2.0 mm turned out to be greater than the spatial period ϵ of 0.7 mm, leading to a strain pattern similar to that shown on the right-hand side of Figure 3.10. Moreover, since the dot width was more than twice the spatial period, the peak strain was lower than the maximum achievable. Assuming that loaded actuator deflections are half those measured without load (Section 3.2.1), this pattern resulted in strains in the order of $\pm 20\%$.

Although the tuning experiments allowed us to find reasonable parameters, the tuning should ideally have been done either with a representative population of Braille readers or individually for each subject. Moreover, the use of texture on the dot may have contributed to an overestimation of the virtual dimensions. The two-hand comparison method may also have introduced errors, but constantly switching

from a conventional Braille display to the VBD was impractical. Finally, the inter-character distance should have been tuned too. These coarse results were found to be sufficient given the scope and the aim of this study.

3.4 Virtual Braille Legibility

The next step needed to evaluate the feasibility of the tangential skin deformation approach for reading Braille was to determine whether blind subjects could read the subset of characters that could be displayed by the VBD. Here, we hoped to also begin identifying the strengths and weaknesses of the concept.

3.4.1 Method

Participants

Two females and three males, experienced Braille readers, volunteered for the study. All subjects were blind from birth. Their ages varied between 22 and 55. The subjects' primary reading finger was the right-hand index. All subjects except the reference subject had never experienced the VBD or heard about our efforts.

Task

The reading task was designed to evaluate the legibility of sequences of first-row characters displayed on the VBD. Subjects were asked to read individual 4-character strings using their dominant reading finger. The first and last characters of each string was always 'c' (●●). The two middle characters could be any of the 16 combinations of the characters 'a' (●○), 'c' (●●), 'dot #4' (○●), and ' ' (○○): "●● ●○ ○● ●●", or "●● ○○ ●● ●●" for example. The character 'c' (●●) was chosen as the string delimiter to avoid confusion.

Procedure

The subjects were given written Braille instructions and had supervised practice trials until they felt comfortable with the task. They were presented with strings to

read in block trials. They placed the slider to the left, waited for an audible signal, read the string, and reported verbally the two middle characters. There was no time limit but they were strongly encouraged to answer quickly. In case of doubt, they were asked to give their best guess. Subjects could stop at any time if they no longer felt comfortable (e.g. loss of tactile sensation, fatigue). They were given the choice of doing the trials with texture, without texture, or in both conditions. Some subjects were tested in both conditions while others decided to experiment with only one type. A trial block comprised 80 strings with each of the 16 possible combinations appearing 5 times in randomized order.

Data Collection

The experimenter logged the subject's answer for each trial. The slider trajectory was automatically recorded by the system. It was analyzed off-line to compute the duration of trials. A trial was considered to begin when the rightmost actuator first arrived at the leftmost virtual dot, and to end when it crossed this dot again for the last time in the opposite direction. In other words, the leading and trailing parts of the slider trajectory for which no actuator was affected by the Braille string were discarded.

3.4.2 Results and Discussion

The main results of this legibility experiment as well as those of a control experiment with conventional Braille (Section 3.5) are summarized in Table 3.1.

Legibility

Legibility was defined by the proportion of correct identifications of 2-character strings. Results suggest that the effect of adding texture was idiosyncratic (see Table 3.1). A dramatic improvement in performance was seen in one subject while a loss was observed in two other subjects. Retaining the best conditions for each subject, the legibility rates were between 71.3% and 98.8%.

	number of trials			legibility (%)			average trial duration (s)		
subject	nominal	textured	control	nominal	textured	control	nominal	textured	control
CN	80	160	80	97.5	88.1	100.0	4.8 ± 2.7	4.6 ± 2.2	3.4 ± 0.7
RB	0	80	80		95.0	100.0		10.8 ± 7.0	7.4 ± 1.4
ML	80	0	80	98.8		100.0	4.6 ± 3.2		2.3 ± 0.6
AB	40	80	80	45.0	86.3	100.0	18.6 ± 12.2	10.5 ± 6.3	2.4 ± 0.7
MS	80	80	80	71.3	66.3	100.0	8.0 ± 2.7	9.1 ± 3.0	5.8 ± 1.7
average	56	80	80	78.1	83.9	100.0	9.0	8.8	4.3

Table 3.1 Summary of results from legibility experiments. Results from the first experiment (VBD, Section 3.4) are presented in columns 'nominal' and 'textured'. Results from the second experiment (conventional Braille, Section 3.5) are presented in the column 'control'. Trial durations are shown with their standard deviation.

Legibility rates were also plotted over time to assess the effect of fatigue. Without texture, no significant change with time could be noticed. However, for some subjects, performance tended to decrease noticeably after about 50 trials when texture was used (see Figure 3.14 for one of the worst-case examples).

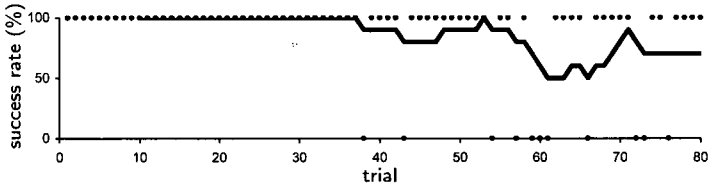

Fig. 3.14 Worst-case example of gradual decrease in performance over time (subject AB, with texture). Dots indicate individual trial results. The curve is a moving average over the past 10 trials.

Character Pairs Legibility

Regardless of the string, individual characters having one dot, '•○' or '○•', were harder to read than characters having no or two dots, '••' or '○○' (see Table 3.2). The legibility also varied with the 2-character string (see Table 3.3). Except for

special cases such as the pair "oo oo" which was read perfectly, no insight could be gained regarding the cause of variations in reading difficulty. It is not clear, for example, why the string "•o o•" has much lower legibility than "o• •o".

•o	o•	••	oo
83.3	88.9	92.2	98.1

Table 3.2 Average character legibility (%). The average was computed across all subjects using individual subject means to compensate for the unequal number of trials under the different conditions.

•o o•	•o oo	o• ••	•o ••	•• •o	o• o•	•• o•	•o •o
66.0	74.0	76.0	77.5	78.0	81.0	81.7	84.5
oo •o	o• •o	•• oo	o• oo	oo ••	oo o•	•• ••	oo oo
85.0	90.0	92.0	92.0	94.0	95.0	98.0	100.0

Table 3.3 Average 2-character string legibility (%). The average was computed across all subjects using individual subject means to compensate for the unequal number of trials under the different conditions.

Table 3.4 shows the confusion matrix for individual characters. No clear pattern emerged, except perhaps that '••' and 'oo' were rarely mistaken for one another. Similarly, Table 3.5 shows the confusion matrix for pairs of characters. Again there was no clear pattern. However, it does seem that the most frequent errors ($n = 3, 4, 5$) generally corresponded to the insertion of a single extra dot or to the incorrect localization of a dot within a character.

		answered			
		•o	••	oo	o•
presented	•o	284	17	23	20
	••	11	315	0	13
	oo	7	1	327	0
	o•	22	9	13	298

Table 3.4 Confusion matrix for individual characters. Answers from all trials were pooled together.

presented \ answered	●○ ●○	●○ ●●	●○ ○○	●○ ○●	●● ●○	●● ●●	●● ○○	●● ○●	○○ ●○	○○ ●●	○○ ○○	○○ ○●	○● ●○	○● ●●	○● ○○	○● ○●
●○ ●○	37	·	·	·	·	1	1	·	2	·	·	·	·	·	3	·
●○ ●●	1	34	·	3	·	1	·	1	·	2	·	1	·	1	·	·
●○ ○○	·	·	29	·	1	·	5	·	·	·	1	·	·	·	5	·
●○ ○●	2	1	2	25	·	·	·	2	2	1	·	4	1	·	1	1
●● ●○	·	·	1	·	35	2	3	·	·	·	·	·	2	·	·	·
●● ●●	·	·	·	·	1	39	·	·	·	·	·	·	·	·	·	·
●● ○○	·	·	2	·	2	·	37	·	·	·	·	·	·	·	·	·
●● ○●	1	·	·	·	4	1	2	33	·	·	·	·	1	1	·	·
○○ ●○	·	·	·	1	·	·	·	·	37	1	·	5	·	·	·	·
○○ ●●	·	·	·	·	·	·	·	·	·	41	·	2	·	·	·	·
○○ ○○	·	·	·	·	·	·	·	·	·	·	42	·	·	·	·	·
○○ ○●	·	·	·	1	·	·	·	1	2	·	·	38	·	·	·	·
○● ●○	·	·	·	1	·	·	2	·	·	·	·	·	37	1	·	1
○● ●●	·	1	·	·	·	·	·	·	1	2	·	·	4	35	·	2
○● ○○	1	·	1	·	·	·	·	·	1	·	·	·	·	·	37	·
○● ○●	1	·	·	·	·	1	·	·	1	·	·	2	2	1	1	35

Table 3.5 Confusion matrix for pairs of characters. Answers from all trials were pooled together.

Reading Patterns

Table 3.1 shows the average duration of trials. The reading speed was far from the expected Braille reading speed of 65 to 185 words per minute [89], but the conditions are so different that a direct comparison is not possible.

Correlations between reading speed and string legibility were also investigated but none could be found, even though there could be important duration variations between trials of a same subject. If it is assumed that the time taken to read a pair of characters is an indication of the confidence the subject has in her or his answer, the characters that the subjects thought were hard to read were not necessarily the ones they had difficulty reading.

The recorded trajectory of the slider was used to investigate the reading pattern used by the subjects. Three classes of patterns were identified (see Figure 3.15). Subjects often used one or two straight passes over the dots. On other occasions, they would explore the virtual Braille string with short back-and-forth motions. In all cases it appeared that subjects read from left to right since they moved slower in that direction.

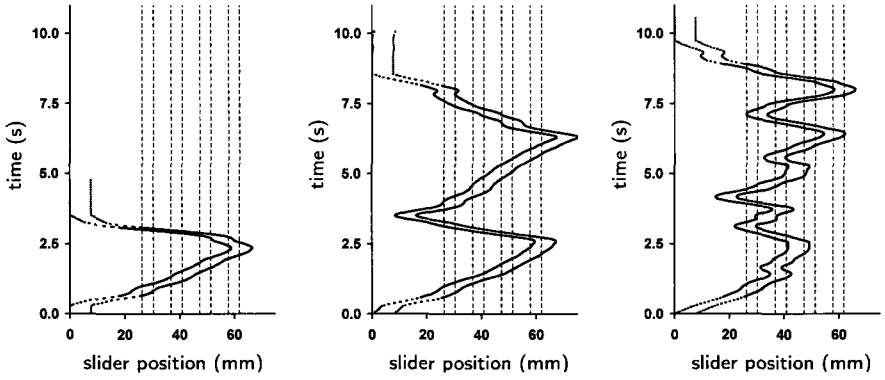

Fig. 3.15 Typical reading patterns: (a) one pass, (b) two passes, and (c) character re-scan. The band between the two curves indicates the span of the display. Dotted sections were not taken into account when computing trial durations. Vertical dashed lines indicate the location of the eight Braille dots.

Verbal Reports

All subjects reported that reading requires concentration, mostly because the dots were subtle (possibly due to the limited range of motion of actuators) and differed in perceived shape from physical Braille dots. Adding texture seemed to facilitate the perception of the dots for some subjects, while others found the sensation unpleasant. Some subjects also complained of loss of tactile sensation over time in both nominal and textured conditions. The occasionally-observed decrease in performance over time seems to confirm that a loss of tactile sensation was occurring with the textured representation and was likely due to the adaptation of tactile receptors.

Subjects also reported difficulties with locating the stationary Braille dots on the virtual Braille line. Most subjects mentioned that the display of characters with more than one row of points and having meaning would help reading.

Finally, contrary to our expectations, subjects reported that scanning constrained by a slider was beneficial because it guided their hand movement. They found it to be an advantage over paper Braille.

3.5 Control Experiment

The reading task performed in the legibility experiment was not representative of typical Braille reading. It was hypothesized that the reading difficulties experienced by some subjects were inherent to reading a single row of Braille dots. The dots found in the bottom rows of most Braille characters could facilitate the localization of the dots within the cell. Without this extra information, locating a dot completely depends on evaluating the length of the spaces between dots.

A follow-up experiment was designed to test the subjects' ability to read single-row Braille characters on a conventional Braille medium: Braille embossed on vinyl tape.

3.5.1 Method

Participants

The experiment was conducted with the same five participants, one year after the original experiment.

Materials and Tasks

The reading task was identical to that used in the original experiment. Subjects were asked to read sequences of four single-row Braille characters starting and ending with '••'. The Braille strings were embossed on 1/2" adhesive vinyl tape using a Braille labeler. The non-existent 'dot #4' character was produced by sanding down the extra dot on character '.' of the embossing wheel.

The resulting Braille has smaller, sharper dots than paper Braille but is still easily readable and commonly used. Vinyl was preferred over paper because it afforded better control on the uniformity of the test plates.

Sixteen Braille labels (one per string) were produced. Each was affixed to a thin, right-angled metallic plate. The placement of the tape was carefully controlled to minimize differences between the plates and allow sufficient space for the finger. During trials, the plates were held down by a switchable magnetic clamp that allowed us to change the strings quickly, see Figure 3.16(a). A single flattened dot was printed on the extreme-left of the tape to serve as a starting point for reading.

Procedure

The subjects were read written instructions and had supervised practice trials until they felt comfortable with the task. They were presented with strings to read in block trials. They placed their finger to the left of the plate, waited for an audible signal, slid their finger over the flattened positioning dot, read the string, reported verbally the two middle characters, slid their finger back over the tape, and removed it from the plate. They were instructed to read only with their right-hand index finger and to keep their other fingers against their palm. There was no time limit

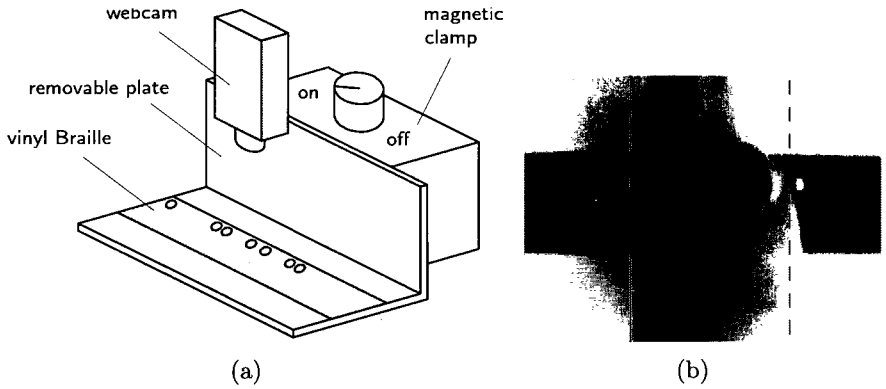

Fig. 3.16 Control experiment with conventional Braille: (a) apparatus, and (b) example of image processing. The location of the leftmost Braille dot is indicated by a solid line. The right edge of the finger is indicated by a dotted line. A green dot was affixed to the nail but not used for processing.

but they were strongly encouraged to answer quickly. In case of doubt, they were asked to give their best guess. The experimenter logged the result of each trial. A trial block comprised 80 strings with each of the 16 possible combinations appearing 5 times in randomized order.

Data Collection

A camera was positioned above the reading surface and was used to record color movies of the finger movements at a rate of 30 frames per second. The movies were compressed and stored for later analysis. A single, uncompressed image of the plate was also taken. The reading patterns and trial durations were analyzed from the movies captured during the experiment. Simple image processing operations were applied on the plate image to extract the position of the leftmost dot of the 4-character string. The absolute difference between the saturation levels of each frame with the background frame was then used to locate the fingertip in the image sequence. In order to approximate the definition of trial duration used in the original experiment, the trial was considered to begin as the rightmost part of the index

crosses over the leftmost dot, and to end when it crosses it again in the reverse direction for the last time. The automated measurements were inspected visually and corrected for 75 of the 400 trials. The precision was estimated to be within 3 frames (±0.1 s).

3.5.2 Results and Discussion

Legibility

All five subjects read the 80 strings presented to them with 100% accuracy. It is thus clear that the reading difficulty experienced by the subjects with the VBD cannot be explained solely by the inherent difficulty of the task.

Reading Patterns

Table 3.1 shows the timing measurements. The subjects read faster on vinyl Braille than on the VBD, more than twice as fast on average. However, one of the slowest subjects on the VBD (AB) is also one of the fastest on vinyl Braille. The reading patterns were also inspected visually to assess their naturalness. Some subjects clearly slowed-down when sliding over dots and frequently returned to previous dots, or to the beginning of the string. This suggests that the reading strategies used by the subjects were different from those used in normal Braille reading [7, 43, 101].

Verbal Reports

Most of the subjects mentioned spontaneously that the reading task was easier on embossed plastic than it was on the VBD. Upon questioning, however, some acknowledged that the reading task was more difficult than typical Braille reading due to the lack of context (meaningless strings) and the absence of cues on the bottom rows. Subjects were also uncomfortable reading with a single finger, and particularly with keeping the other fingers in a fist.

3.6 Conclusion and Future Work

This study showed that experienced Braille readers could read sequences of first-row Braille characters using the VBD with a legibility ranging from passable to excellent. This is encouraging considering that most subjects had little prior training with the device and that the character strings were meaningless.

Reading with the VBD was nevertheless difficult. The control experiment showed that subjects could read faster and with perfect accuracy when the strings were presented on embossed vinyl tape. This observation is confirmed by the subjects' verbal comments during both experiments.

Adding texture to the dots was found to increase performance for some subjects but was rejected by others. Prolonged use also seemed to cause tactile fatigue (numbness in the reading finger) as reported by subjects for both nominal and textured stimulus. This phenomenon was sometimes confirmed by an increase in the number of reading errors over time in the case of textured stimulus.

While this study was conducted with too few subjects to make it possible to draw final conclusions, it suggests nevertheless that reading Braille characters with devices based on the principle of the VBD could be possible. The study also helped identify the strengths and weaknesses of the current prototype and, more importantly, provides indications as to how it could be improved.

The strength of the dot sensation must be increased to realistically convey the illusion of a Braille dot. This issue involves the improvement of the actuators used and of their configuration. On-going work concerned with the micro-mechanical properties of the skin and the means to deform it at a very small scale is expected to yield an improvement in the performance of piezoelectric benders for this application [172]. Further experimentation with deflection functions could also lead to better approximations of the sensation of scanning over Braille dots. Designing deflection functions for maximum strain variations at skin mechanoreceptors may, for example, increase the strength of the dot sensation. Indeed static mechanical models of the skin designed by [120] and [160] suggest that deformation of receptors, as opposed to their displacement, is likely to determine sensation. The strain experienced by

slowly adapting (SA) receptors may be particularly important to resolve the spatial details of scanned Braille [122].

The cause of tactile fatigue must also be addressed. It is not clear what causes it and how it can be avoided. It is likely however that replacing the contact line by a more uniform contact surface could significantly delay the onset of tactile numbness.

Reading with the VBD requires a scanning movement. While this allowed the VBD to render adequately the dynamic sensation of sliding over dots, it prevented the user from stopping over a region as is sometimes done when reading physical Braille. Display without net movement could be possible if the magnitude of the strain produced by the device was made to depend on the force applied by the subject. This not only could make the sensation of sliding over a dot more realistic, but also allow the subject to stop and press against virtual dots.

The results of the tuning experiment and of the legibility experiment conspire to indicate that perhaps the greatest problem with the current display design is the difficulties experienced by subjects in evaluating the distance between dots. This is suggested by the distortions in perceived Braille dimensions introduced by the VBD when compared to standard English Braille. This is also consistent with observed reading errors and verbal reports of the legibility experiment.

Finally, the VBD should be extended to display complete Braille characters. Packing 4 rows of actuators capable of displaying forces and displacements similar to those of the VBD within the height of a Braille cell will be a significant technical challenge.

Acknowledgement

The authors would like to thank the reviewers for their contribution to the clarity and completeness of this paper. This project was mostly done under contract with VisuAide inc (Longueuil, Canada). The authors wish to thank Pierre Hamel and Jean-Michel Gagnon of VisuAide Inc. for management and technical support respectively, and the subjects who participated in the study. Others aspects of this project are supported by the E. (Ben) & Mary Hochhausen Fund for Research in Adaptive Technology For Blind and Visually Impaired Persons. Other contributions from the

Canadian National Institute for the Blind (CNIB), the Institut Nazareth et Louis-Braille (INLB), and the Centre de Recherche Interdisciplinaire en Réadaptation du Montréal métropolitain (CRIR) are gratefully acknowledged. Vincent Lévesque would like to thank the Natural Sciences and Engineering Council of Canada (NSERC) and the Fonds Québécois de la recherche sur la nature et les technologies (FQRNT) for postgraduate fellowships. Jérôme Pasquero would like to thank FQRNT for a postgraduate fellowship. Vincent Hayward would like to thank NSERC for an operating grant.

Chapter 4

Haptically Enabled Handheld Information Display with Distributed Tactile Transducer

Preface to Chapter 4

The following chapter appeared in

> Pasquero, J., Luk, J., Levesque, V., Wang, Q., Hayward, V., and MacLean, K. E., Haptically Enabled Handheld Information Display with Distributed Tactile Transducer, *IEEE Transactions on Multimedia*, 9(4):746-753, 2007.

This chapter describes the design of a miniature tactile transducer similar to the one introduced in Chapter 3. At the time of writing, there did not exist any other tactile feedback transducer of similar expressive capabilities that could successfully be integrated into a handheld device. The miniature transducer opened up an entire new class of possibilities which are well suited to the context of mobile interaction.

Contributions of Authors

Pasquero contributed to most of the work described in this chapter by writing over 80% of the paper and generating all the figures. He is responsible for designing and assembling the tactile transducer for the THMB device. The device itself is the result of a close collaboration between Pasquero and Don Pavlasek of the machine shop in the ECE department at McGill University.[1] With the help of Levesque, Pasquero also devised the scroll-and-feel application described at the end of the chapter. Most of the software used for this application is the work of Levesque. The concept for the device was co-developed with Luk who also ran the user evaluation and reported on the findings. Wang provided expertise on the design of the electronics. Professors Hayward and MacLean supervised the work and contributed in the form of discussions and guidance. They also edited the manuscript.

[1] Refer to appendix B for more details about this collaboration

Haptically Enabled Handheld Information Display with Distributed Tactile Transducer

Abstract

This paper describes the design, construction, and initial evaluation of a handheld information device that supports combined tactile and graphical interaction. The design comprises a liquid crystal graphic display co-located with a miniature, low-power, distributed tactile transducer. This transducer can create electronically-controlled lateral skin deformation patterns which give the sensation of sliding over small shapes. It is integrated within a slider mechanism to control scrolling. It also functions as a detent when pushing on it. Tactile feedback and the combination of visual and tactile feedback in a mobile context enable the development of new functions, such as multimodal navigation within large graphic spaces.

4.1 Introduction

Artificial tactile feedback provides an underutilized channel which can be used to gain information about the state of a device, the occurrence of an event, or its content. In mobile devices, these possibilities have been exploited so far in a very limited way for want of adequate transducers.

In general, diverse approaches have been proposed as operating principles for tactile transducers, the electromechanical subsystem of a complete tactile display. These approaches differ by the actuator technology employed, by the manner in which tactile signals are delivered to the skin, by the type of tactile signals they produce, by practical factors such as bulk and power, as well as by the range of sensations that they can mediate. Many principles able to produce movement in a small space have been explored, including miniature solenoids that vibrate steel membrane elements by variable reluctance [5], surface acoustic waves that modify the sliding interaction between a surface and a plate [149], electrogel lamellae that push against the skin [85], and pneumatically-activated dimples rising out of a surface [103]. A more complete survey of tactile display technologies can be found in [115].

The list of designs attempted to date is truly great. The most common approach to making tactile transducers is to pack a dense array of actuated pins; an early example of which can be found in [161]. At rest, the pins resemble a flat surface. The display is activated by moving the pins normally to the skin to approximate a shape, much the way a graphical display approximates an illumination field with an array of pixels. However, this type of transducer has significant limitations including the difficulty of achieving a high density array, high power consumption, and high manufacturing cost. Due to these challenges, the resulting devices are typically complex (e.g. [136]) or not portable.

Compact designs are required for tactile displays to enter the mobile device arena. This is of particular importance for the development of synthetic, abstract tactile signals, such as *haptic icons* [97], and *tactons* [16], which aim at alleviating, through touch, the attentional and cognitive load imposed on users by human-computer technology. For instance, portable devices typically have a multitude of functionalities

and are often used in demanding multitasking situations. Distributed skin stimulation delivered by integrated tactile displays can help create expressive and pleasant percepts capable of delivering substitutive or complementary information to users.

4.1.1 Tactile Displays in Portable Devices

The Optacon, a sensory-substitution device for the blind developed in the 70s, is a key example of an early portable system. It is made of a dense array of 144 skin contactors configured in a gutter geometry to maximize the amount of skin being stimulated [14]. Each contactor is on-off in that it can either vibrate at full amplitude or remain still. When connected to a camera held in one hand, it is possible — with training — to distinguish printed high-contrast patterns such as letters or line drawings while scanning them with the camera. This results in binary high-contrast *tactile* images, or tactile representations of printed material. Since then, other variations on this principle have been investigated such as the VITAL [5], which achieves an exceptionally slim form factor.

Another promising approach to portability is electro-stimulation by which the skin is electrically coupled to an array of electrodes in order to create small currents in the superficial skin layers. This principle was the basis of early portable tactile displays for aiding the visually impaired [148]. There are many possible variants in current waveforms, electrode geometries, and skin coupling methods. One of the latest designs is described in [79].

Today, three types of portable tactile stimulators are commercially available. The first is the Braille reader for the visually impaired, e.g. [158]. The second is the vibrator found in portable phones, pagers, and game controllers by which a motor spins an eccentric mass to shake the supporting structure. The third is the voice-coil transducer which applies vibrotactile stimulation to the skin as an aid for the deaf to lip read.[2]

While having been adopted with some success by the commercial market, these tactile devices suffer from many limitations. Braille readers, which have pins that

[2]See Audiological Engineering Corp. Tactaid™ - www.tactaid.com.

rise above a surface to produce Braille dots, are usable only for rendering refreshable Braille. Vibratory signals delivered by portable devices and voice-coils are spatially diffuse. They under-exploit the human ability to perceive, discern, and process intricate distributed skin deformation patterns. Although recent industrial developments have enabled a wider variety of waveform patterns,[3] the predominant sensation is still that of "buzzing".

4.1.2 Distributed Lateral Skin Deformation

Vibrotactile feedback will probably always have a role to play in information displays as a mechanically simple way to demand a user's attention. However, the sensations delivered by devices that vibrate evoke a distant and unnatural computer world. A likely explanation is that most ecologically familiar vibrotactile sensations arise from human-made objects, such as rotating machinery.

While it is unlikely that tactile displays will soon be able to convey the gamut of real-life tactile sensations, we believe that it is important to investigate tactile interaction techniques which feel natural and therefore, are better suited to the human reality. Touch feedback resulting from distributed skin deformations is potentially richer in content and form than vibrotactile stimulation.

Over the past few years we have developed a family of devices that operate according to a principle that departs from the conventional skin indentation approach [57, 116, 170]. Instead, tactile sensations are created by deforming the skin laterally, with negligible normal deflection. Because this display technique works by causing lateral skin deformations, we assign it the neologism *laterotactile.* A variety of sensations can be created this way, for example the sensation of sliding a finger over small shapes. We have recently succeeded in using this technique to display virtual Braille dots [93], as well as tactile graphics [173]. Because of its potential for miniaturization, laterotactile technology is a natural candidate for the integration of tactile displays in mobile devices.

[3] See Immersion Corp. VibeTonz^{TM} product literature.

4.1.3 Paper Overview

We explore the design, prototyping, and preliminary evaluation of a handheld device called the 'Tactile Handheld Miniature Bimodal device' (THMB, pronounced "thumb"). It features a tactile display mounted behind a liquid crystal display (LCD). This device enables the manipulation of information available visually on the screen as well as tactually on a transducer in contact with the user's left thumb (Fig. 4.1). The 'haptic interface' comprises a sliding active surface which is able to create a range of programmable tactile sensations. A miniaturized laterotactile display makes this possible.

In the remainder of this paper, we introduce the THMB and its interface, then describe its construction with sufficient detail to allow re-implementation. We close with an illustrative application example and a discussion of practical considerations. Several other scenario examples are discussed in [95].

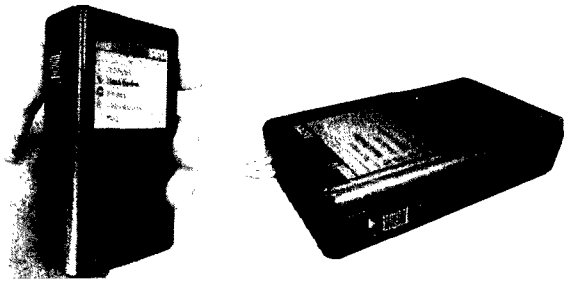

Fig. 4.1 The THMB device.

4.2 The THMB: Designed for Bimodality and Ergonomics

The overarching design goal was to support tactile feedback correlated with control movements, while being sensitive to ergonomic considerations. We achieved the synergistic association of sensory-motor functions in a small space and with few resources by integrating a laterotactile display within a thumb-operated slider on the left side of a hand-sized case, and locating it directly under the LCD screen.

Complementary modes of interaction can be supported this way, taking advantage of the human ability to process and integrate information in parallel [110].

Because of its hyper-agility, the thumb enjoys a unique status among fingers [176]. Its joints and musculature allow it to dexterously reach much of the inside of the hand. Many mobile devices rely on it for data input, for dialing, for writing text messages, or for browsing the web. The thumb also often operates scroll wheels positioned on the left side of a case, thereby freeing the right hand for other tasks. With the THMB device, the thumb operates the slider while the palm and remaining fingers hold the device at a convenient distance from the eyes (Fig. 4.2(a)). Several other hand postures may also be employed.

The tactile display is programmed to cause tactile sensations in the thumbpad as exemplified in Fig. 4.2(b). These sensations can be related to what is seen on the screen or they can convey signals on their own. The active surface of the tactile display protrudes slightly from the left side of the case, making it easy to slide the unit up and down by flexion of the thumb's distal phalanx. This mechanism also operates as a push-button triggered by light inward pressure.

4.3 Specifying distributed lateral skin deformation with the THMB

With distributed lateral skin deformation, a progressive wave of local tangential deformation in the fingertip skin gives the sensation of a small shape sliding on the finger [57]. Clearly, there are many ways to specify moving skin strain patterns, which vary in perceptual realism and distinctness. One technique that we have used in the past has patterns of deflection such that for one actuator i, the deflection d_i is:

$$d_i(y_s, t) = P(y_s + (i-1)\epsilon, t), \; i \in [1 \cdots N] \tag{4.1}$$

where, referring to Fig. 4.3, N is the total number of actuators, t is time, ϵ is a spatial period preferably equal to that of the display, about 0.9 mm for the THMB, and $P(.)$ is an admissible function. The variable y_s stands for the position of a slider

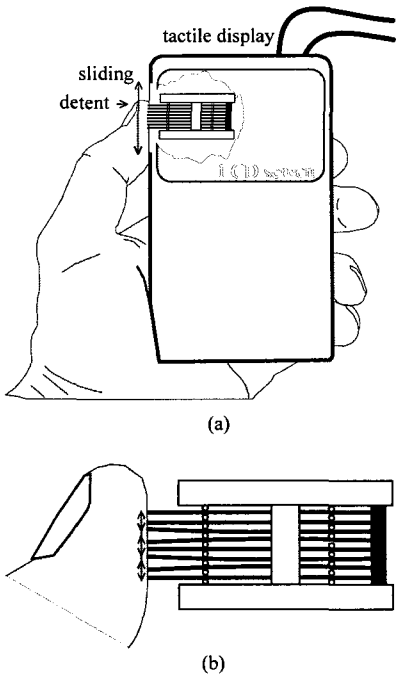

Fig. 4.2 (a) Overview of the THMB interface and (b) close-up of the tactile display stimulating the skin by lateral deformation.

on which the display is mounted; it can also serve as a simulated position if no slider is available.

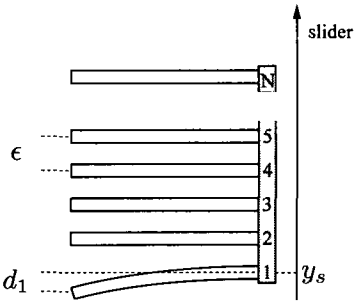

Fig. 4.3 Variables for the deflection function of Eq. (4.1).

Simple deflection patterns do not depend on time and must be explored actively.

One example is a sinusoidal grating $d_i(y_s) = A\sin([\pi(y_s + (i-1)\epsilon)]/2)$. Other important cases consist of features traveling automatically at constant speed. Fig. 4.4 shows five actuators undergoing a sequence of deflections and their consequences on the deformation of the skin for a feature moving upward.

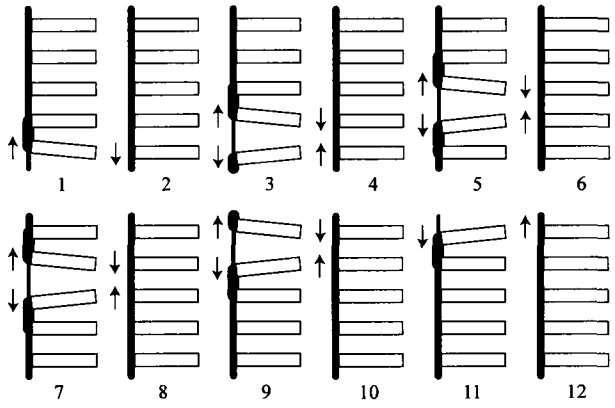

Fig. 4.4 Example of a moving feature with the slider fixed. Only key frames are shown. In practice transitions are smooth.

Generally, when the finger first comes in contact with a static undeflected display surface (e.g. frame 12 in Fig. 4.4), the skin is in a state of nearly neutral strain. The initial sensation is that of a flat surface because the actuators are sufficiently dense and the receptors embedded in the skin respond to deformation [70]. The sensation of a moving feature corresponds to traveling waves with specific deformation trajectories. In such instances, an individual patch of skin could be first in a relaxed state. It would then expand, relax, compress, and become relaxed again. In the meantime, the neighboring patch would undergo the same trajectory, but with a phase difference. To employ a physics terminology, if the phase velocity is equal to the group velocity, then a given fixed pattern will be felt to be moving at that velocity. If this velocity is equal and opposite to that of the slider, then the pattern is felt to be stationary with respect to the THMB enclosure. The THMB supports more complex functions which are left for future studies.

4.4 System Description

The transducer can easily be manufactured as a power-efficient, compact unit having a low component count following a technique inspired from earlier devices [93, 116]. Piezoelectric actuators can be inexpensive, light, power-efficient, and in fact are already in common use in many portable devices, including refreshable Braille displays and mobile phones. The technology is mature, well-understood, and steadily improving. Therefore, it is well suited for portable devices and allows for tactile displays with high spatial resolutions that require only one single part per actuator. On the downside, piezoelectric actuators are limited in the displacement they yield and require high activation voltages.

In this section, we seek to provide sufficient details for the replication of the THMB display, given access to modest fabrication facilities and, whenever possible, using off-the-shelf components with a view to demonstrate practical viability.

4.4.1 Piezoelectric Benders

The display makes use of piezoelectric bimorphs, also called benders, to generate lateral skin deformation. A bimorph bender has two layers of piezoelectric material bonded to a central electrode, plus electrodes coated on the outer faces. When an electric field is created in these layers with proper polarization, they experience opposing dimensional changes. The stress distribution created by the common boundary results in bending deformation of the entire slab [142]. The dynamic constitutive relationships are complex, but at frequencies below the first resonant mode, the electrical behavior of an actuator is well approximated by that of a capacitor. This means that when operating under a DC voltage, an actuator remains bent but consumes no power.

Benders are usually mounted in a cantilever configuration as shown in Fig. 4.5(a), but in the THMB, they are stacked between two rows of thin rods as illustrated in Fig. 4.5(b). This configuration yields a displacement-stiffness tradeoff that differs significantly from that of ordinary cantilevers. For a given free deflection, this structure will always be stiffer, but have a smaller deflection [170]. A high stiffness is

essential because skin is a tough medium to deform. The dual-pinning mount allows actuators to be smaller and to operate from lower voltages for an equivalent resulting tactile stimulation.

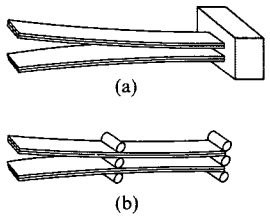

Fig. 4.5 Two possible configurations for the piezo-actuators: (a) cantilever mount and (b) stiffer dual-pinning mount used for the THMB.

4.4.2 Preparation of the Benders

Raw piezoelectric sheets obtained from the manufacturer are first prepared to provide for electrical connections compatible with the tight spacing requirements. The preparation is analogous to the fabrication of microelectronic circuits: a sequence of deposition and removal of material. Fabrication can be done in an ordinary mechanical workshop rather than in a cleanroom, and tools are a temperature-regulated soldering station, a precision knife, a pair of scissors, and model painting brushes. The materials are piezoelectric bender plates (eight units from Piezo Systems® T215-H4-203Y), conductive tape (3M Corporation, EMI Copper Foil Shielding Tape 1181), standard nonconductive adhesive tape, conductive paint (Loctite® Quick Grid™ found in car rear window defogger repair kits, product #15067), and polyurethane varnish.

As illustrated by Fig. 4.6(a), the slabs are first coated with tracks of conductive paint to improve their resistance to scratches. The two external electrodes are then electrically connected using strips of conductive tape wrapped around the plates. The conductive strip adheres directly to the external electrodes so patches of nonconductive tape must first be applied to prevent contact with the central electrode. To provide access to the central electrode, soldering pads are created by grinding off

channels in the external electrodes. These pads are connected to the central electrode using a small amount of conductive paint applied on the side faces to which a thin wire can then be soldered. The resulting solder joint must be thinner than the space between two neighboring actuators. Finally, the tips of the plates are coated with several varnish layers to isolate them from the user's skin.

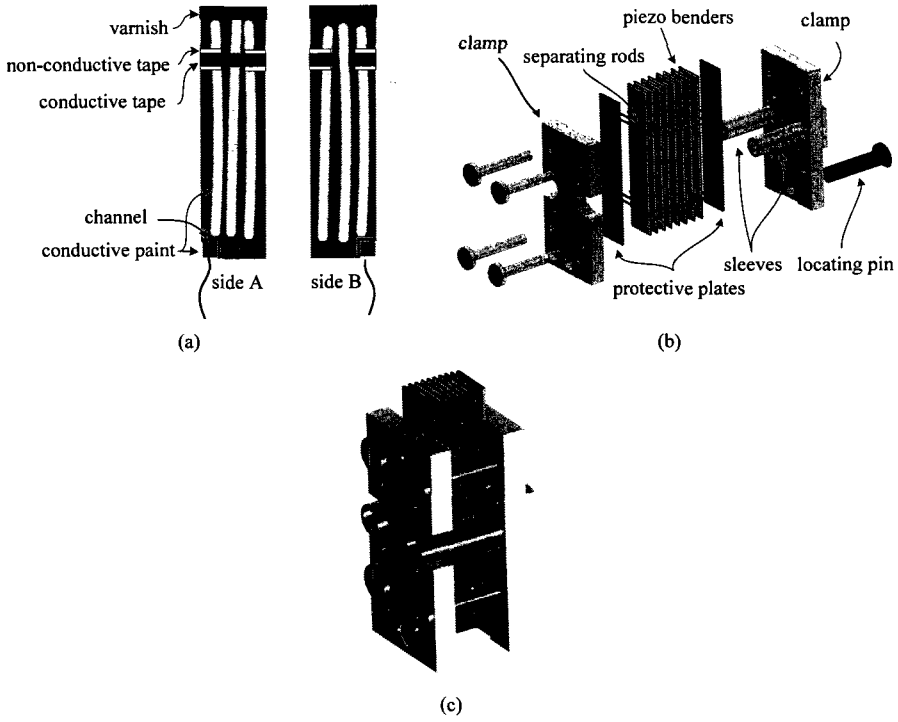

Fig. 4.6 Building the transducer: (a) preparation of the piezoelectric actuators, (b) assembly and (c) ready to be mounted to the slider.

4.4.3 Transducer Assembly

The tactile display prototype is made of eight 0.38-mm-thick piezoelectric benders (6.4-by-31.8 mm), sandwiched between 2 rows of 0.5-mm-diameter brass rods that act as hinges of a dual-pinning configuration (see section 4.4.1) and as ground connections. Two passive brass plates at the ends protect the benders from excessive

lateral stress applied by the user's thumb (Fig. 4.6(b)). The assembly is held together between two aluminum clamps secured with four screws that act as support for the rods. Once assembled, the transducer constitutes a single unit where only the actuator tips are free to move. In the present design, the rods are located at 5 mm and 23 mm respectively from the bottom of the actuator. The benders rest at the bottom on a plastic locating pin. Once the transducer is assembled, as shown on Fig. 4.6(c), the size of the exposed contact area is about 8.7 mm by 6.4 mm. During the assembly phase, locating pins (not shown) secure the piezoelectric benders alignment along the display axis.

4.4.4 Integration Inside the Case

The basic transducer is integrated in a ready-made case in order to provide for a sliding mechanism combined with a push-button. As seen in Fig. 4.7(a), two parallel guiding stainless steel shafts connect two parallel plates, and pass through the assembled transducer's two sleeves. Structural rigidity is provided by a strut that also supports a linear potentiometer. The axis-to-axis distances of the two steel shafts, the transducer's sleeves, and the potentiometer strut are all matched with exacting precision in order to guarantee smooth operation. Sliding travel is limited by spacers to a distance of 11 mm. This results in the sub-assembly shown in perspective view in Fig. 4.7(b).

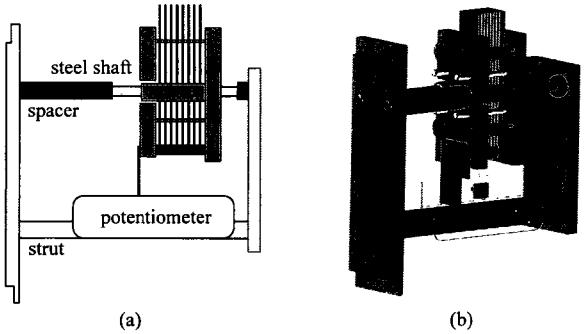

Fig. 4.7 Sliding mechanism: (a) 2D schematic of the sliding mechanism and (b) 3D view of the integrated system.

The sub-assembly is inserted in an off-the-shelf plastic case (Hammond® 1593 series instrument enclosure). Such prototyping boxes typically provide slots to secure printed circuit boards. As seen in Fig. 4.8, we use these to anchor a pivoting mechanism that transfers the detent behavior of a miniature, board-mountable switch to the tactile display surface itself. This way, the THMB device is given a slide-and-press-detent behavior similar to that of a scroll wheel with the difference that the tactile sensations associated with scrolling are not fixed but programmable in a number of ways.

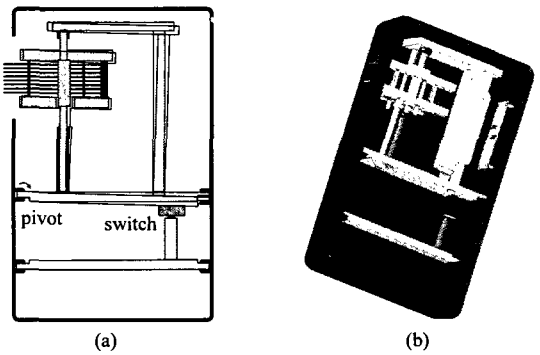

Fig. 4.8 (a) Integration of the sub-assembly in prototyping box and (b) 3D view of the integrated system.

The large number of connections packed in a small space makes for a significant wire management problem. A strategy similar to that used in electromechanical consumer devices such as printers and scanners is employed. Referring to Fig. 4.9, two miniature printed circuit boards, one moving (PC-1) and one fixed (PC-2), act as concentrators. A 0.5 mm flat flex cable (from Parlex®) allows for a smooth and reliable connection. Two round, miniature-gauge silicone cables (Cooner Wires® AS323REV1) route the 14 signals to external electronics.

4.4.5 Electronics

For the purpose of this prototype, the power and control electronics were adapted from earlier projects [93, 116], and were kept external to the casing. Briefly, the digital system comprised a personal computer (PC) and a field-programmable gate array

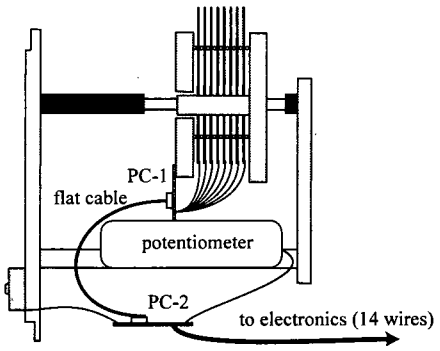

Fig. 4.9 Wire management.

(FPGA) development board (Constellation board™ 10K50E from Nova Engineering®) supporting universal serial bus 1.1 (USB) connectivity (Fig. 4.10).

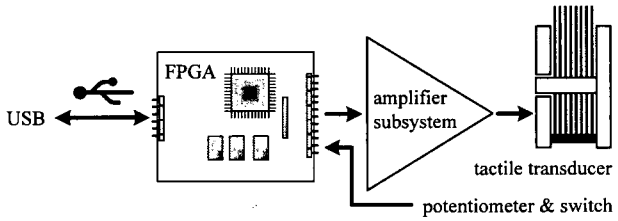

Fig. 4.10 Overview of THMB electronics.

The FPGA is programmed to convert the serial data coming from the personal computer to eight parallel 156.25-kHz pulse-width modulation (PWM) output signals. The eight digital level PWM signals are fed to eight analog comparators that step up the voltage to ±15 V. The resulting signals are low-passed and supplied to eight high-voltage operational amplifiers (Burr-Brown® OPA445) which in turn drive the eight benders at analog levels between ±45 V (Fig. 4.11). The FPGA also reads the output of an analog-to-digital converter detecting the position of the slider, and provides this information, along with the state of the switch, to the personal computer.

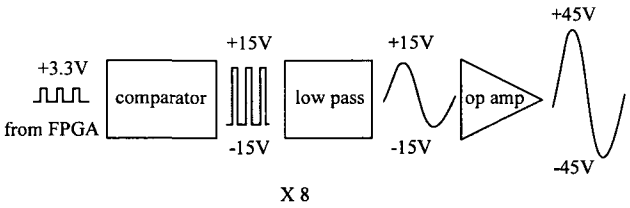

Fig. 4.11 Operation over one of the eight channels of the amplifier.

4.4.6 Liquid Crystal Display

The THMB is equipped with a 58.8 x 49.9 x 5.8 mm color LCD screen inserted in the casing, along with its control electronics. The LCD system accepts NTSC format signals and therefore, computers with a TV-out port can drive the LCD screen as an external monitor. A second miniature-gage flexible cable connects the THMB's screen to the computer's video-out port.

4.4.7 Software

The THMB's host software was written using the STReSS Software Development Kit (STReSSd, pronounced "stressed"). STReSSd consists of a cross-platform C++ object library for Linux and Microsoft Windows™. It provides a simple, common interface to most of our laterotactile displays (STReSS [116], VBD [93], and THMB). The library provides a layered interface allowing interaction with the displays at different levels of abstraction. Low-level interaction with the FPGA development board occurred through a separate Software Development Kit that we have released under an open-source license.[4]

4.5 Preliminary Evaluation: 'Scroll-And-Feel'

In the context of mobile interaction, users' cognitive resources are dedicated, in order of importance, to acting on the environment (e.g., walking, talking), to monitoring it

[4]Sidus - http://sidus.sourceforge.net.

(e.g., to scan the environment visually), and lastly to interacting with the mobile device [109]. For this reason, an undivided attention to the screen cannot be expected. During interactions such as browsing the web on a mobile phone, tactile feedback can provide continuous bridging information between glances at the screen.

With this in mind, a preliminary application was developed to exercise the device's capabilities, with the goal of implementing concurrent tactile and visual exploration of a graphical document that is too long to fit on the screen, and has to be scrolled up and down by using the slider.

4.5.1 Programming

User Control Under velocity control, it is possible to browse through an entire image by controlling speed and direction. The displacement range of the slider is divided into three zones. When the slider is inside a central neutral zone, the graphical image does not move. Moving the slider away from the center causes the image to scroll up or down at a speed proportional to the deflection from the center. Absolute positioning is applicable only within a small scope.

Tactile Mapping There are many ways to map a scrolling graphical image to a tactile sensation. We adopted a simple strategy that associated a tactile track to a graphical track. The two are scrolled simultaneously under slider command. The graphical track contains short paragraphs of text separated by headers and small figures (Fig. 4.12(b)). The tactile track is composed of a series of tactile haptic icons, see Fig. 4.12(a). A specific dynamic tactile sensation is felt on the thumb skin when a given tactile icon travels across the tactile window, a subset of the graphic track currently visible on the LCD (Fig. 4.12(c)).

Haptic icons were created for the three types of tags found in the document: *header*, *text* and *image*. The three types of tactile sensations are spatial functions of the position of the corresponding haptic icon in the tactile window. Referring to the three icons in Fig. 4.12(a), light regions represent a downward deflection of the corresponding actuator, while dark regions represent an upward deflection. When the page is scrolled, a given haptic icon activates different benders, thus creating

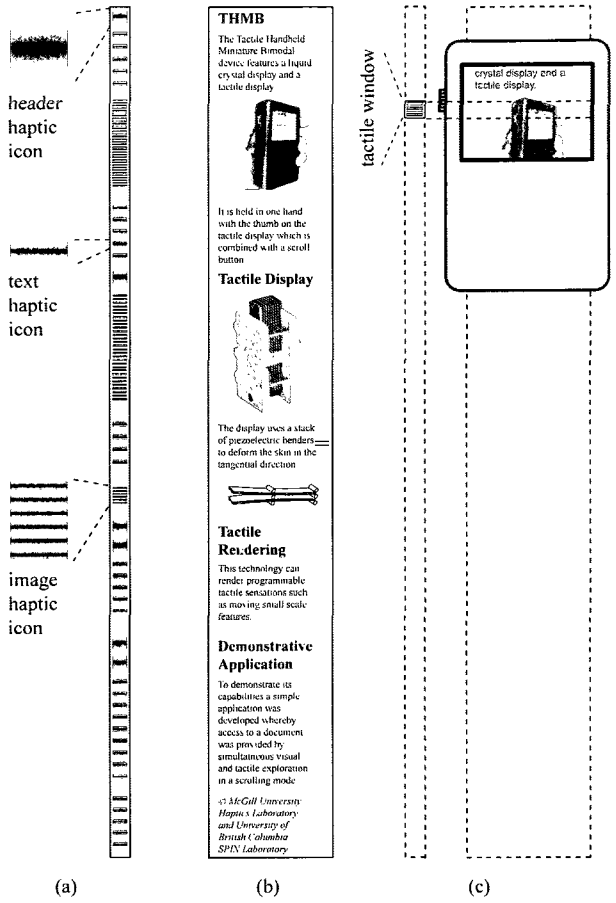

Fig. 4.12 Scrolling application for evaluation: (a) tactile track with the three types of haptic icons, (b) graphic track and (c) tactile window mapped to the center of the device's graphical window.

what can be described as *peristaltic* skin deformation such as the patterns that can be observed when a finger slides over a featured surface [91].

This mapping and its effect are illustrated in Fig. 4.13 for a *text* haptic icon. In this case, a single element travels upward in the tactile window. The positions of the eight actuators are shown in successive frames where patches of skin between two actuators, i.e. *taxels*, are stretched and compressed. Because the strain experienced by the skin is the result of the differential deflection of adjacent actuators, the strain

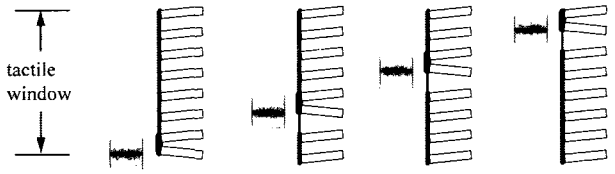

Fig. 4.13 Peristaltic deformation of the skin shown while a single line of text passes through successive frames of the tactile window (from bottom to top): position of the the *text* haptic icon within the tactile window (left column in each frame) and resulting skin deformation (right column).

is related to the *gradient* of the grey level of a haptic icon representation. Adjacent frames indicate the skin's deformation trajectory when the image scrolls by a distance equal to the display's spatial period.

A *text* icon spans about half of the height of a line. It causes a full but single back-and-forth bending motion of an actuator. The resulting sensation of a *text* icon moving from actuator to actuator is often felt as a tiny, sharp, and isolated traveling feature. *Header* icons are widened versions of the *text* icons and therefore, they are more diffuse and of lower contrast. They are also often perceived to be weaker and larger. Finally, *image* icons consist of a texture spanning an image's height. This generates a sensation rich in high frequencies. During exploration, the overall impression is that of tactile flow when multiple icons travel across the thumb.

4.5.2 Initial Reports

Six volunteers, all mobile phone users in their 20s with no previous experience with the THMB device were recruited. They were not compensated for their time. These participants freely explored the scroll-and-feel application and answered a few questions in 10-to-15-minute sessions. They were given no instructions other than the manner in which the device should be held.

The application represented by Fig. 4.12 was designed to be stylistically similar to news webpages found on the Internet. A red marker displayed at a fixed position on the left side of the graphical display delineated the smaller tactile window within

the graphical window. This way, the participants knew what part of the document they were exploring tactually. The graphical page was formatted to fit the width of a portable screen. It was displayed on the device's LCD for the two first participants and emulated on a PC monitor for the other four because of a hardware failure. The window sizes were the same in both configurations.

Half of the participants reported that the tactile sensations during scrolling were like touching a traveling train of ridges under the thumb. The illusion of motion was robust. Most of the participants (5/6), however, could not initially distinguish more than two different types of tactile sensations. The closely spaced ridges comprising the *image* icon and the loosely spaced *text* or *header* icons were so different that more subtle differences may have been obscured [97]. Upon further questioning, participants were eventually able to differentiate between them. Using this feedback, we were able to redesign the haptic icons to make all three clearly identifiable and differentiable. This development is beyond the scope of this paper and will be reported in the future.[5]

A few participants mentioned that "playing" with the device was pleasant, but could not verbalize exactly why. Not surprisingly, some (3/6) compared the *image* icon to vibration and/or texture, probably because of its high spatial frequency content. On the other hand, all seemed to agree that the *text* and *header* icons were compelling and completely different from anything related to vibrotactile stimulation. In general, all participants welcomed the velocity control; one of them even commented on its value. The subjective responses showed no differences between setups. We collected enthusiastic comments; some participants even envisioned environments and situations where our device could be useful. These and others are further developed in [95].[6]

Overall, the participants' ability to discriminate icons shows promise. It provides evidence that percepts richer than those currently available on handheld devices are possible in a mobile context.

[5]See Chapter 5 for an analytic tool that can be used to identify easily distinguishable icons.
[6]Refer to Chapters 5 and 6 for some examples of potential user applications.

4.6 Implications for Feasibility and Future Work

We have built what is to our knowledge the first realistic prototype of a mobile handheld interface capable of direct-contact distributed tactile feedback. We hope that this tactile delivery technology will enable portable tactile displays in the same way that LCD technology has enabled portable visual displays. The key is lateral skin deformation, which utilizes piezoelectric technology in an efficient configuration for a compact, high-resolution tactile transducer.

The process of building the first generation THMB suggests that it will be feasible to incorporate the technology into consumer devices. In its present prototype form, the THMB is tethered, with remotized circuitry and power requirements. The dimensions and weight of the transducer still need to be reduced by a factor of two or three to meet the requirements for a practical mobile device. Nevertheless, nothing in the design makes these limitations terminal.

The participant's first impressions suggest that the THMB can convey compelling tactile sensations in a multimodal use environment, and these encouraging results motivate further investments. In [25], Chan et al. successfully identified perceptual groupings of haptic icons based on distinctiveness. Such a systematic technique should be adapted to future developments of the tactile sensations displayed on the THMB. We have also already embarked on application-specific studies of information manipulation in the small space available in a mobile device, and we will be further evaluating how direct-contact tactile feedback can provide added value in terms of both performance and comfort.

Acknowledgment

This work was performed in part with support of the Fonds Québécois de la recherche sur la nature et les technologies (FQRNT), the Natural Sciences and Engineering Council of Canada (NSERC) and PRECARN/IRIS. The authors would like to thank Don Pavlasek for insights and expertise, Joe Boka and Bob Thomson for their technical support, Andrew Gosline for helpful comments and all participants in the initial

reports for providing valuable feedback.

Chapter 5

Perceptual Analysis of Haptic Icons: an Investigation into the Validity of Cluster Sorted MDS

Preface to Chapter 5

The following chapter appeared in

> Pasquero, J., Luk, J., Little, S. MacLean, K. E., Perceptual Analysis of Haptic Icons: an Investigation into the Validity of Cluster Sorted MDS, *Proceedings of the 14th Symposium on Haptic Interfaces For Virtual Environment and Teleoperator Systems*, IEEE Virtual Reality, 2006, pp. 437-444.

Chapter 5 describes the validation of the cluster-sorted multidimensional technique that was used to partially characterize the THMB. This perceptual characterization of the device constitutes a first step in an effort to associate practical meaning to the tactile sensations that can be created. Taken together, the THMB platform described in Chapter 4 and the validation of the perceptual characterization described in this chapter provide a solid framework for the development of applications for handheld mobile devices that make use of tactile feedback. These were first envisioned in:

> Luk, J., Pasquero, J., Little, S., MacLean, K., Levesque, V. and Hayward, V., A Role for Haptics in Mobile Interaction: Initial Design Using a Handheld Tactile Display Prototype, *Proceedings of CHI 2006*, 171–180, 2006.

Contributions of Authors

Pasquero made the most substantial contribution to this next chapter by writing over 80% of the paper and generating all the figures and tables but one. He conducted the analysis of the data that is presented and used it to perform the analytical validation of the cluster-sorted MDS approach that is described. Luk and Little devised the software application used for the experiment and designed the set of tactile stimuli. Luk also ran the experiment and collected the user data. Prof. MacLean provided

supervision in the form of discussions and did some parts of the writing. She also edited the manuscript.

Perceptual Analysis of Haptic Icons: an Investigation into the Validity of Cluster Sorted MDS

Abstract

The design of usable haptic icons (brief informational signals delivered through the sense of touch) requires a tool for measuring perceptual distances between icons that will be used together as a set. Our experiences with one potentially powerful approach, Multidimensional Scaling (MDS) analysis of perceptual data acquired using an efficient cluster sorting technique, raised questions relating to the methodology for data collection. In this paper, we review key issues relating to perceptual data collection method, describe an example data set and present its initial MDS analysis, and then examine the impact of collection method on MDS outcome through a secondary analysis of the data and the inherent structure of the algorithm components. Our analysis suggests that an understanding of these issues is important for the method's effective use, but has not exposed any major flaws with the process.

5.1 Introduction

Using synthetic tactile feedback, we can communicate information about the state or content of a system, or the occurrence of an event, to a human user via cutaneous touch. Mobile devices which make use of tactile feedback (e.g., vibrating pagers and cell phones) have been available on the consumer market for several years. The tactile stimulation that they generate, however, is limited in its expressive capability; typically they make use of a single tactile element that vibrates at a unique frequency. More recently, mobile device manufacturers have been integrating tactile actuators with increased degrees of freedom into their products. This is possible, for instance, by using actuators for which the vibrating frequency can be controlled. The goal, inspired by the common use of audio icons in desktop interfaces and mobile telephony [16], is to allow for the design and construction of specific and short abstract tactile messages that can easily be interpreted by users with minimal cognitive effort [97].

At the center of our research lies the goal of understanding how to maximize the level of useful information content that tactile/haptic feedback devices can convey. To that end, we aim to develop a systematic analytic method to characterize the stimulus space of tactile feedback devices in order to inform the design of tactile icons.

In this paper, we present our progress to date in clarifying and validating the use of multidimensional scaling (MDS) for this purpose. While it seems clear from past efforts that MDS analysis can provide relevant and (in some sense) accurate information about the perceived relationships among a set of test stimuli and thus guide stimulus set design, we were not satisfied with our ability to interpret and quantify these relationships, nor with our understanding of the impact of perceptual data collection method on the MDS analysis outcome.

Here, we begin by reviewing key issues relating to data collection method, then proceed to examine the impact of collection method on the MDS result via a detailed analysis of an example data set. While an exhaustive critique will be a substantial project, the insights presented here should aid in the wise use and accurate inter-

pretation of results provided by a potentially powerful tool. Our analysis does not expose any major flaws with the process, and we thus conclude that it can be a valuable tool to evaluate the expressive capability of haptic devices.

5.2 Background

5.2.1 Multidimensional Scaling

MDS comprises a set of methods that display the structure of a dissimilarity matrix as a geometrical representation in a given number of dimensions [184].

A dissimilarity matrix contains $N * (N-1)/2$ individual perceived dissimilarity scores between N items of a set. MDS optimization algorithms try to map the perceived dissimilarity scores δ_{ij} between items i and j from a dissimilarity matrix onto an M-dimension geometrical space (usually Euclidean) within which distances d_{ij} between items represent relative levels of dissimilarity:

$$d_{ij} \approx f(\delta_{ij}) \quad i,j \in [1,2\cdots,N] \tag{5.1}$$

where $f()$ is a monotonic function. On the Euclidean space, distances are calculated with

$$d_{ij} = \sqrt{\sum_{a=1}^{M}(x_{ia} - x_{ja})^2} \tag{5.2}$$

where x_{ia} and x_{ja} are the item coordinates that the MDS algorithm is solving for.

Dissimilarity data are always mapped with a variable degree of error to the geometrical space, and several measures are used to evaluate the goodness of fit. For instance, one popular method introduced by Kruskal [86] and used to report goodness-of-fit factors in this paper, consists of minimizing the *stress* function:

$$S = \sqrt{\frac{\sum_{i=1}^{N}\sum_{j=i+1}^{N}(f(\delta_{ij}) - d_{ij})^2}{\sum_{i=1}^{N}\sum_{j=i+1}^{N} d_{ij}^2}} \tag{5.3}$$

Interpreting the resulting MDS plot and extracting the underlying perceptual

axes governing the signal space can be a frustrating task. The dimensionality of the space under investigation and the exact meaning of its prominent axes cannot be inferred easily from any goodness-of-fit factor or from simple observation of the MDS plot. Allowing the MDS algorithm to use a large number of dimensions will result in a geometrical representation that is matched closely to the data, but which can be very difficult to visualize and interpret. On the other hand, limiting the MDS algorithm to a small number of dimensions might not lead to the most exact representation of the data, but it could offer a more intuitive view of its primary underlying structure. This is of particular importance when dealing with the analysis of perceptual data, which is known to be both very noisy and the product of an intricate system. The mechanisms that govern human sensorial perception are complex and it is unlikely that any MDS analysis will reveal how they work. However, our premise, confirmed qualitatively by past work, is that MDS analysis can still expose the salient perceptual dimensions and provide guidelines to the design of new meaningful and easily distinguishable artificial stimuli.

5.2.2 Obtaining Data for The Dissimilarity Matrix

The conventional means of acquiring dissimilarity values is to directly compare all possible pairs of stimuli in the set of interest. Subjects are presented with each pair, then asked to rate their degree of similarity. Similarity scores are then converted to dissimilarity scores through a simple transformation. Unfortunately, this paired comparison method suffers from a lack of consistency as the number of pair-wise comparisons increases (e.g., a stimulus set of 25 requires 300 comparisons and a stimulus set of 35 requires 595 comparisons): subjects forget their grading scale, and experiment time lengthens dramatically.

Cluster-Sorting Method: To avoid the problems mentioned above, Ward developed a cluster-sorting technique for an experiment designed to study the salient properties of the physical environment as perceived by humans [174]. He presented a set 20 photographs representing real physical environments to a group of subjects, and asked them to categorize the photographs into different clusters according to a set

of rules. The sorting task was repeated a number of times with a different number of clusters for each trial. Dissimilarity matrices were constructed based on the number of times a particular pair of stimuli was sorted into the same cluster. Stimulus pairs that were often grouped together received a low dissimilarity score, whereas pairs that were never grouped together were assigned a maximum dissimilarity score.[1]

While Ward's method was originally developed for visual stimuli, MacLean, Enriquez, Chan and others applied a similar technique to sets of haptic stimuli. Brief computer-generated haptic signals, or haptic icons, were constructed by varying parameters such as frequency, magnitude and shape and were presented to subjects through a force feedback knob [97] or more recently, a vibrotactile mouse [25]. Users were asked to classify the haptic icons into different clusters; the sorting task was repeated five times varying the number of clusters. Results from an MDS analysis of this type of data has demonstrated a separation of the icons that follows intuition while lending extra structural detail. In the knob-generated data set, for example, while frequency seemed to be the salient dimension overall, other perceptual dimensions such as shape also emerged from the MDS plots when only signals pertaining to a certain range of frequencies were considered.

Hollins et al. have also made use of a group-sorting strategy to obtain dissimilarity data for the study of the underlying perceptual dimensions of tactile sensations [61]. However, in this case, the data for the MDS analysis was the result of a single sorting task during which subjects were asked to classify 17 real-life surface textures in a minimum of 3 but maximum of 7 groups. In a second phase of the experiment, subjects were asked to rate the textures according to different assumed properties (e.g, roughness, temperature, flatness) and the resulting scales were mapped to the 3D MDS space obtained in the first phase. Hollins et al. found that two dimensions of the MDS space corresponded closely to roughness-smoothness and hardness-softness scales respectively. The third dimension could not be related to any of the remaining scales. While valuable for understanding tactile perception, their perceptual data collection strategy differs from ours in its implementation. Moreover, we are more

[1]The reader is referred to [159] for a review of alternate methods that can be used to collect dissimilarity data for a MDS analysis.

concerned with the understanding of the perception of synthetic computer-generated tactile icons, rather than the existing tactile properties of real objects.

5.2.3 Objectives and Organization

The strength of the cluster-sorting method lies both in its greater speed of execution - reducing the period over which subjects must retain a calibration - and, we hypothesize, in the more global comparative nature of the sorting task as compared to pairwise comparisons. However, we have recently observed that the method's global nature may also be a weakness, because it creates a complex pattern of correlations among the elements of the dissimilarity matrix. It is the need to understand the impact of this restriction on independence on MDS outcome (which nevertheless seems to show plausible patterns) that motivated the investigation we report here.

Our analysis is illustrated using actual data obtained from a cluster-sorting task using a new tactile device that we have built. Because the device differs from other displays in the way it displays tactile stimuli to the thumb tip, the resultant sensations were novel and subjects should have no pre-existing categorizations.

In Section 5.3, we present our example data set, including an overview of tactile device and experimental method. Section 5.4 introduces our key critical questions, and in section 5.5 we explore each of the question in turn with secondary analyses and discussion. Finally, we close with conclusions and recommendations.

5.3 Data: Hardware, Stimuli and Procedure

5.3.1 Hardware

Tactile Stimulation by Lateral Skin Stretch

Applying a traveling wave of local tangential deformations against the fingertip skin can induce the sensation of a small-scale shape sliding on the finger [57]. During such interaction, no indentation normal to the skin surface occurs. The moving pattern of lateral skin stretch is perceived as a traveling feature. This illusion was

previously exploited to display truncated Braille characters to the visually impaired with a prototype device called the Virtual Braille Display (VBD) [93].

Tactile Handheld Miniature Bimodal Device

Following encouraging results obtained with the VBD, we built a miniature version of the tactile display (TD) and integrated it into a prototype for a mobile handheld device. The result, referred to as the Tactile Handheld Miniature Bimodal (THMB) device (Figure 5.1), comprises a miniature tactile display for the thumb tip and an LCD screen, both assembled inside a small plastic case. Because our focus here is on the data analysis, detailed implementation information is reserved for later publication.[2]

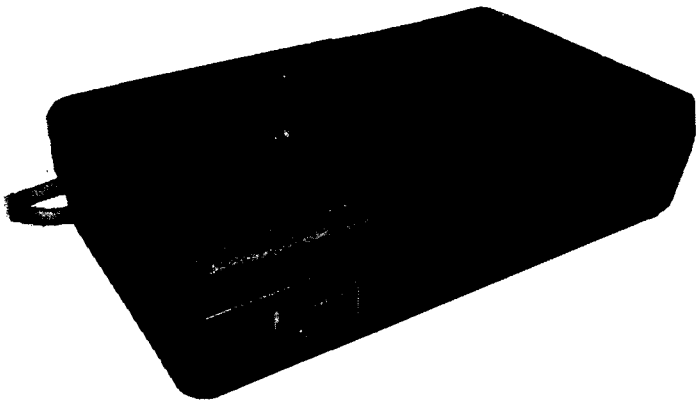

Fig. 5.1 THMB device.

Typically, the THMB device's casing is held in the left palm and is secured in place between the thumb and the four remaining fingers, similarly to how a PDA or cell phone is held (Figure 5.2). The tactile display consists of a stack of eight piezoelectric benders intercalated between brass rods, which protrudes slightly through a narrow slit on the left side face of the device's case. The user's thumb tip rests against

[2]A first version of the prototype was described in Chapter 4. Drawings of a second improved version can be found in Appendix B.

the summit of the the TD. When activated, the TD's piezoelectric actuators induce lateral skin deformation to the thumb tip by bending.

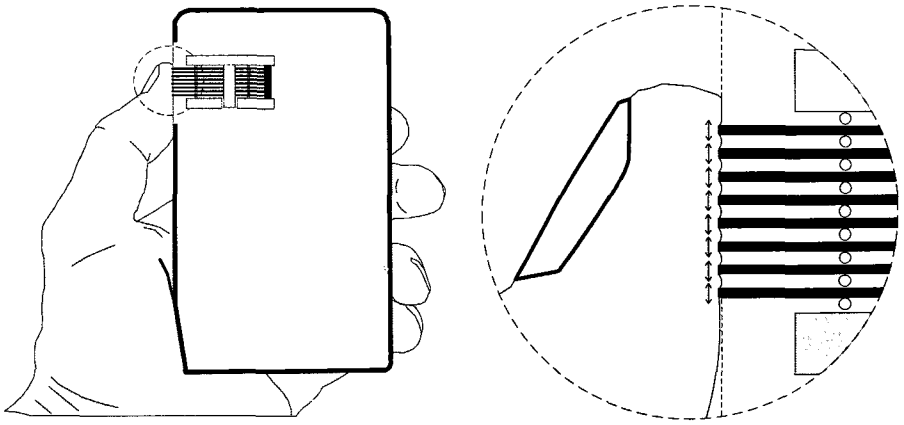

Fig. 5.2 Interaction with the device: (a) overview of the THMB interface and (b) close-up on the thumb against the tactile display.

The amount of bending of a particular piezoelectric actuator is controlled with a voltage applied across its electrodes. A PC host running Linux generates the 8 control signals, one for each piezo-actuator and sends them to interfacing electronics where they get filtered and amplified before being applied across the piezoelectric benders. The resulting control voltages range from $\pm 50V$ and are updated at 3125 samples/sec. They are encoded with a single byte and therefore can only take 256 different values.

5.3.2 Design of the Haptic Icon Stimuli

Because each of the 8 piezo actuators can be controlled independently, the device may be programmed to display 8-pixel animations, each pixel capable of 256 activation levels. Sets of 8 piezo activation states can be thought of as a frame; a time series of such frames define a given tactile stimulus, which thus varies across both time and space. With the THMB, we wanted to explore alternative shape configurations as well as ways to animate the stimuli over time.

Selection of Stimuli: Using a custom-built prototyping application to explore the stimulus space, we generated a wide variety of waveforms, defined as the displacement of one piezo element as a function of time. We selected five waveforms for further investigation based on the criteria that they produced qualitatively different tactile sensations while remaining as simple as possible. The intent was to begin with a candidate set of canonical basic waveforms referred to in this paper as TRI, ROLL, SAW, BUMP and EDGE (Figure 5.3). A waveform was animated across the display using a *tweening* method that consisted of playing the waveform on successive piezo elements with a phase delay of 1 sample. The resulting animation could then be slowed down by reducing the frame rate; or it could be played backwards, resulting in the sensation of the traveling tactile pattern moving in the opposite direction.

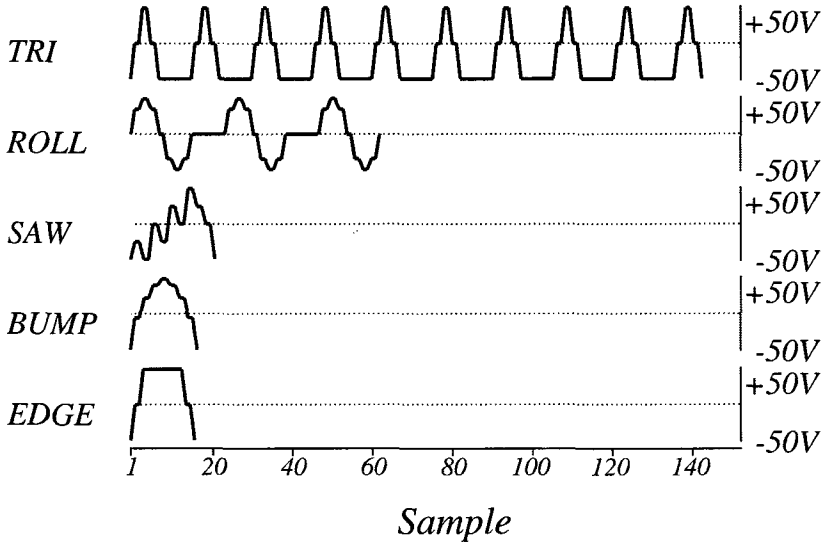

Fig. 5.3 Five waveforms used for the experiment.

Stimuli: The study used a $5 \times 3 \times 2$ combination of stimulus factors, resulting in 30 total tactile animations used as stimuli. The five waveforms mentioned above, three animation speeds, and two amplitudes were varied to produce the patterns. Animation speeds consisted of $\frac{1}{10}^{th}$, $\frac{1}{15}^{th}$ and $\frac{1}{20}^{th}$ of the maximum 3125 frames/sec rate imposed by the hardware. Amplitude could either be full scale (F) or half scale

(H). Because some waveforms were longer than others, the total duration of the animation (i.e., the amount of time a stimulus was present under the user's finger) was calculated from the waveform and animation speed. We kept track of this duration meta-parameter throughout the data analysis because it was a readily observable characteristic of the stimuli (Table 5.1). Stimuli varied in duration between 74ms and 960ms. For the purpose of the data analysis, the stimuli were labeled according to the following indexing: $<$ `waveform` $><$ *amplitude* $><$ *duration* $>$.

Table 5.1 Overall animation durations calculated from waveform and speed parameters.

waveform (Fig. 5.3)	**wavelength** (in samples)	**duration** (ms) at speed $\frac{1}{10}$	$\frac{1}{15}$	$\frac{1}{20}$
TRI	143	480	720	960
ROLL	62	221	331	442
SAW	20	86	130	173
BUMP	16	74	110	147
EDGE	16	74	110	147

5.3.3 Experiment

Subjects

Ten participants (7 male) aged 19 to 31 were recruited for the experiment; none had any prior experience with the THMB device or similar displays. The participants were told to hold the device in their left hand with their left thumb resting lightly on the tactile display, and were allowed to take breaks as necessary.

Method

The participants used the method described in [97] to conduct the cluster-sorting task, wherein a GUI displayed graphical tiles for each stimulus which, when clicked, played the respective stimulus, and which could be moved about and sorted into the cluster boxes. The software was modified to present the 30 tactile stimuli used in this experiment, and the onscreen tiles representing the stimuli were not marked.

Each participant performed five similar cluster-sorting trials. In the first trial, the stimuli were sorted into a user-selectable number (from 2 to 15) of clusters. In the subsequent four trials, subjects were required to sort the stimuli into 3, 6, 9, 12, or 15 clusters presented in random order, with the trial containing the number of clusters closest to the user-selected first trial clusters being eliminated. For example, if the user selected 8 groups for the first trial, their subsequent trials would consist of $\{3, 6, 12, 15\}$ groups in random presentation order.

Data

A similarity matrix was calculated and converted to a dissimilarity matrix in the same way described in [97]. The cluster-sorting method assigns similarity points to a pair of stimuli each time they are put together in the same cluster. The number of similarity points given to a pair of stimuli placed together in a given sorting trial is equal to the number of clusters in that trial. For instance, if two stimuli were present in the same cluster during the trials involving 3 and 9 clusters, the pair is assigned 12 points $(3 + 9)$. For the first trial for which subjects get to choose the number of clusters, the amount of points alloted is adjusted to fit the closest integer of the set $\{3, 6, 9, 12, 15\}$. The similarity points are summed over a subject's trials and a simple inverse linear transform is applied to obtain a matrix with dissimilarity scores that range from 0 to 1000.[3] A score of 0 indicates that the pair of stimuli always appeared in the same cluster across all trials, whereas a score of 1000 indicates that they were never paired.

Results

Dissimilarity matrices from all subjects were combined to create an average dissimilarity matrix (shown in Table 5.2) and the data was submitted to a 2D MDS analysis using SPSS™ 13.0 (Euclidean distance algorithm with ordinal/untied data).[4] Figure

[3] In our case, if s_{ij} is the sum of the similarity points obtained for a pair of icons i and j, the corresponding dissimilarity score d_{ij} is computed as $d_{ij} = -\frac{1000}{45} s_{ij} + 1000$.

[4] Ordinal/untied refers to data embodying rank information where elements with identical scores are free to be decoupled into successive rank positions.

5.4 depicts the stimuli's resulting spatial arrangement, upon which three groupings have been manually overlaid by the authors. Orientation of the graph and its axes returned by SPSS^{TM} are irrelevant since MDS analysis is only concerned with the relative positions of the stimuli in space, rather than their absolute positions. Therefore, the plot has been rotated for clarity. Five lab colleagues uninvolved in the project were also provided with the plot without stimulus labels and asked to group the stimuli into as many natural clusters as desired based on spatial layout; similar configurations resulted in all cases. MDS plots in 3D were also generated but no extra structural information could be extrapolated.

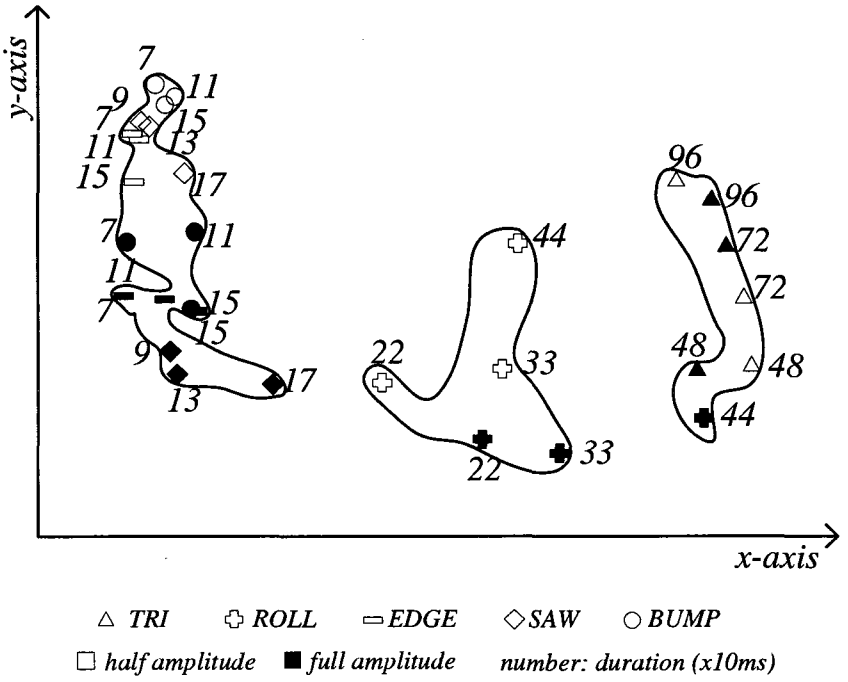

Fig. 5.4 Ordinal untied MDS plot from the average dissimilarity matrix obtained with the cluster-sorting experiment (*stress* = 0.083). Signal durations (from Table 5.1) have been rounded to the nearest multiple of 10ms for clarity.

To consider sensitivity of results to MDS algorithm, the individual subject matrices were also run through a replicated MDS (RMDS) algorithm. The resulting plots and

waveform	TRI						SAW						ROLL					
amplitude	Half (H)			Full (F)			Half (H)			Full (F)			Half (H)			Full (F)		
duration (ms)	480	720	960	480	720	960	86	130	173	86	130	173	221	331	442	221	331	442
TRIH720	640																	
TRIH960	793	693																
TRIF480	640	873	907															
TRIF720	707	747	800	780														
TRIF960	900	773	700	780	547													
SAWH86	1000	1000	993	1000	1000	1000												
SAWH130	1000	1000	993	1000	1000	1000	513											
SAWH173	1000	1000	993	1000	1000	1000	580	507										
SAWF86	1000	1000	1000	993	1000	1000	920	920	893									
SAWF130	1000	1000	1000	993	1000	1000	933	933	927	367								
SAWF173	993	993	1000	987	993	993	947	947	893	667	513							
ROLLH221	980	987	993	973	980	987	960	987	920	893	827	713						
ROLLH331	933	940	973	947	973	980	973	973	973	967	967	960	653					
ROLLH442	880	933	907	907	947	967	960	953	953	993	993	987	867	527				
ROLLF221	907	967	987	953	960	980	993	993	993	893	880	707	780	907	887			
ROLLF331	820	913	980	847	967	973	1000	1000	973	973	973	927	900	807	793	667		
ROLLF442	700	740	866	840	887	946	1000	1000	1000	1000	1000	987	960	873	780	813	633	
BUMPH74	1000	1000	993	1000	1000	1000	487	633	620	953	967	973	987	973	987	993	1000	1000
BUMPH110	1000	1000	993	1000	1000	1000	480	533	540	953	967	973	960	973	960	993	1000	1000
BUMPH147	1000	1000	993	1000	1000	1000	460	533	533	953	967	960	960	973	960	993	1000	1000
BUMPF74	1000	1000	1000	1000	1000	1000	840	713	753	707	820	867	947	993	993	993	973	1000
BUMPF110	993	993	993	993	993	993	700	707	753	820	800	833	873	987	960	967	993	993
BUMPF147	1000	1000	1000	993	1000	1000	860	887	767	753	807	753	653	953	967	860	953	1000
EDGEH74	1000	1000	993	1000	1000	1000	527	373	573	867	933	947	973	973	987	993	1000	1000
EDGEH110	1000	1000	1000	1000	1000	1000	540	487	573	907	893	940	987	980	960	993	1000	1000
EDGEH147	1000	1000	1000	1000	1000	1000	600	613	453	800	880	887	893	993	967	993	1000	1000
EDGEF74	1000	1000	1000	993	1000	1000	873	793	800	700	740	833	947	987	1000	980	973	993
EDGEF110	1000	1000	1000	993	1000	1000	800	853	833	747	707	693	780	960	993	893	980	1000
EDGEF147	1000	1000	993	1000	1000	1000	860	873	860	787	793	773	793	920	967	947	973	993

waveform	BUMP						EDGE				
amplitude	Half (H)			Full (F)			Half (H)			Full (F)	
duration (ms)	74	110	147	74	110	147	74	110	147	74	110
TRIH720											
TRIH960											
TRIF480											
TRIF720											
TRIF960											
SAWH86											
SAWH130											
SAWH173											
SAWF86											
SAWF130											
SAWF173											
ROLLH221											
ROLLH331											
ROLLH442											
ROLLF221											
ROLLF331											
ROLLF442											
BUMPH74											
BUMPH110	273										
BUMPH147	307	353									
BUMPF74	833	873	767								
BUMPF110	887	800	747	487							
BUMPF147	947	880	860	780	647						
EDGEH74	480	560	553	707	787	853					
EDGEH110	540	573	593	753	753	860	553				
EDGEH147	740	660	653	687	567	647	653	547			
EDGEF74	900	887	867	400	507	773	787	827	720		
EDGEF110	880	880	873	600	533	647	840	813	707	640	
EDGEF147	953	900	926	620	567	473	873	813	713	607	647

Table 5.2 Average dissimilarity matrix obtained from all 10 subjects' dissimilarity matrices.

goodness-of-fit stress factors for the non-metric RMDS analysis were similar (but not identical) to the ones reported in this paper.

Subjects assigned labels to the different sorting boxes based on what we have reduced to four essential categories: duration (e.g., quick, very long, slow, short, fast), magnitude (e.g., moderate, weak, medium, strong), multiplicity (e.g., single, double, multiple, one, two) and description (e.g., rhythmic, beat(s), thud, heart beat, broken, vibes, sensation). While not immediately analyzable, these categories may lend insight into the perceptual dimensions subjects used in their classifications.

Preliminary Discussion

A few features are evident through inspection of Figure 5.4. The stimuli seem clustered naturally in three major groups ranged along the x-axis, with the shorter BUMP, SAW and EDGE signals to the left, ROLL in the middle and TRI to the right. There thus seems to be some influence of either pattern and/or overall duration on the major grouping; however, a duration argument may be at odds with the longer groups' internal structure. For the TRI and ROLL groups, duration (221-960ms) increases on the y-axis, while for the BUMP/SAW/EDGE cluster, overall duration seems roughly aligned horizontally. One of the longest ROLL stimuli is associated with the TRI cluster.

Neither amplitude nor duration visibly affect the TRI stimuli; whereas both the ROLL and BUMP/SAW/EDGE clusters seem strongly differentiated by amplitude (on the y-axis) and slightly by duration (on the x-axis).

The confounding of the relatively brief SAW, EDGE and BUMP stimuli is at least partially due to electromechanical filtering of the hardware; wave shape detail at these durations was later confirmed to be beyond the device's temporal display resolution.

While interesting and clearly meaningful at some level, the complex structure of these results immediately raised questions. For example: what is the actual perceptual *amplitude* of the differences found between these signals – can it be quantified? The y-axis seems to be interpreted as amplitude for the shorter signals, and as duration for another cluster which happens to contain the longest signals. Is this di-

mension perceptually contiguous, or is it used independently by each cluster? These uncertainties lead to closer examination of the data collection technique, and the more detailed analysis that follows.

5.4 Guiding Questions

The cluster-sorting method for collecting perceived dissimilarity data is a framework that drastically decreases the time required to obtain a dissimilarity matrix from a large set of stimuli. Rather than requiring subjects to gradually construct a personal dissimilarity scale through successive direct comparisons between all possible pairs of stimuli – to give an order of magnitude, the experiment reported in section 5.3.3 would have required 435 direct comparisons – it provides them with an implicit reference scale that facilitates consistency in their answers. On the other hand, the resulting dissimilarity matrix is the outcome of a process with strong structural constraints where, at the same time, there may exist multiple grouping strategies which would express the same degree of dissimilarity. Our curiosity about how the resulting interdependence in the dissimilarity matrix might impact the MDS outcome included these questions:

Q1 When applying the cluster-sorting technique, what assumptions are made concerning the structure of the collected data?

Q2 How does a dissimilarity matrix obtained from the cluster-sorting technique differ from a dissimilarity matrix constructed from a set of classic pairwise comparisons?

Q3 Does data obtained via the cluster-sorting technique preserve the structural information about the stimulus space at every level (for each sub-space)? Similarly, if one perceptual dimension is far more prominent than the others, will the other minor dimensions still be represented accurately in the dissimilarity matrix?

Q4 How is noise in the data-gathering process reflected in the dissimilarity matrix obtained with the cluster-sorting method?

In the following section, we examine each of these questions with further discussion and where possible, analysis of the cluster-sorting technique itself, data collected with it, or both.

5.5 Analysis of Cluster-Sorting Method

5.5.1 Assumptions on the Nature of the Dissimilarity Data Collected

MDS algorithms differ by assumptions made regarding input data distributions. Therefore, it is important to clarify relevant attributes of the data obtained with cluster-sorting technique. MDS algorithms are either metric or non-metric. Metric algorithms try to map the stimuli in a Euclidean space by keeping the distances between the stimuli as much like their corresponding dissimilarity scores as possible. Non-metric algorithms focus more on preserving the rank of the dissimilarity scores.

Therefore, by applying a metric algorithm to the data obtained from the cluster-sorting method, we formulate the assumption that the connection strength between two stimuli grouped together is linearly dependent on the number of clusters present. For instance, an association between two stimuli grouped in a sorting trial with 15 clusters is considered to be 5 times more powerful than a association between any other two stimuli grouped in a sorting trial with 3 clusters. While both Ward and MacLean obtained positive results when they applied a metric MDS analysis to their data, it is hard to demonstrate that a metric assumption is justified at all times.

On the other hand, the data in the dissimilarity matrix can also be considered to be non-metric. In this case, it is is assumed that summing the similarity points that accrue over successive sorting trials does not destroy the monotonicity of the data. This weaker assumption relaxes the constraints on the MDS analysis. However, an MDS analysis now results in a graphical representation that can only give an idea of the extent to which stimuli differ, without providing as much quantitative meaning to the graphical inter-stimuli distances.

Results reported in 5.3.3 are produced with a nonmetric analysis on the assumption that the data is ordinal. We also considered results generated with different metric MDS algorithms (including the INDSCAL algorithm). The MDS metric plots shared some similar features with their nonmetric counterparts, such as proximity in space of all the TRI stimuli. However, the plots did not exhibit any clear visual clustering of the stimuli into the 3 major groups described in the discussion section 5.3.3 and the stimuli appeared to be more continuously spread out across the perceptual space, making it harder to interpret the results. Without invalidating any metric analysis of the data, our results indicate that perceptual data obtained using the cluster-sorting technique seems best assumed to be ordinal.

5.5.2 Dissimilarity Matrix Properties

A matrix obtained from a set of classic pair-wise comparisons contains score-elements that are independent of each other. Each dissimilarity score reflects a dissimilarity intensity that is in relation to the others, but changing one score element won't affect the other elements of the matrix. On the other hand, a dissimilarity matrix obtained with the cluster-sorting method is structurally constricted. One cannot re-adjust a single score for a particular pair comparison, without having to readjust other elements of the matrix. This is a direct consequence of the requirement for a minimum of one-stimulus-per-cluster imposed by the experimental procedure and the fact that assigning a stimulus to a cluster also means that it is not put into any of the other clusters of the same trial.

Dissimilarity scores obtained with the cluster-sorting method have one of a set of discrete values. For the cluster sets used here, these vary from 0 to 1000 with a resolution of 66.6 (Figure 5.5). For 10 of the 16 possible cases, there exists more than one combinatorial way to obtain the same score. Moreover, the random distribution of the set of score values across the entire dissimilarity matrix is not a normal distribution since the matrix elements (i.e., all the different pairs) are not independent.

This was observed in post-analysis and shown in Table 5.3, which displays the

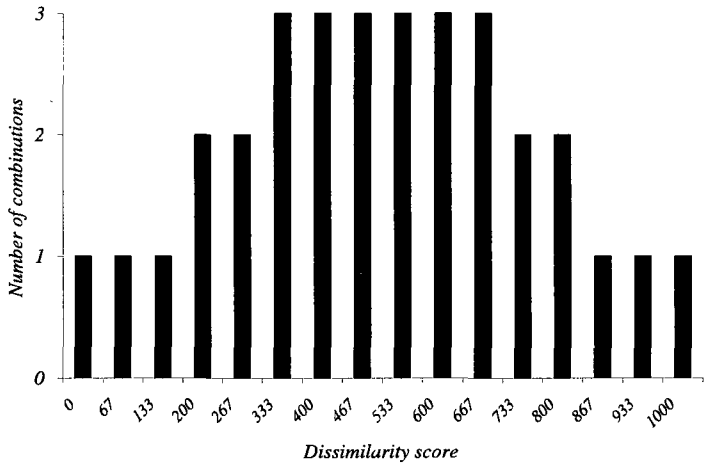

Fig. 5.5 Number of possible combinations to obtain a single dissimilarity score with a cluster-set of $\{3, 6, 9, 12, 15\}$.

number of clusters actually chosen by our subjects on the first sorting trial, as well as the number of maximum dissimilarity scores of 1000 and of 0. The latter provide some indication of the dissimilarity matrix structure: a score of 1000 indicates two stimuli that were never grouped together across the five trials whereas a score of 0 indicates that they were always paired. Across all 10 dissimilarity matrices, 20 associations (ranging from 2 to 4 stimuli in size) with scores of 0 were made and all but one involved SAW, EDGE and BUMP stimuli exclusively - for instance, subject #4 systematically put the stimuli BUMPF74, BUMPF110 and EDGEH147 in the same clusters across all trials. Moreover, only 2 out of these 20 0-dissimilarity associations included stimuli of different amplitudes (one of the exceptions is the example given above for subject #4). Conversely, the average dissimilarity matrix indicates that 107 of the 435 possible pairs that could have been formed across the five trials were never made (Table 5.4).

The high occurrence of dissimilarity scores of 1000 is easily explained. Over all five sorting trials, numerous pairs of stimuli are never grouped together and therefore are assigned a maximum dissimilarity score. Caution must be taken when interpreting a 1000 score. It cannot represent infinite dissimilarity because there is no absolute

Table 5.3 Key properties of the individual dissimilarity matrices. Percentages in parentheses indicate a proportion of the dissimilarity matrix's 435 elements. Scores of 1000 and 0 represent cases where a stimuli pair was always and never grouped together respectively.

Subject index	Nb. init. clusters	Nb. of 1000's	Nb. of 0's
1	6	210 (48%)	0 (0%)
2	5	241 (55%)	2 (0%)
3	8	245 (56%)	2 (0%)
4	5	238 (55%)	9 (2%)
5	12	253 (58%)	3 (1%)
6	8	233 (54%)	4 (1%)
7	2	203 (47%)	1 (0%)
8	7	269 (62%)	1 (0%)
9	6	258 (59%)	8 (2%)
10	6	207 (48%)	5 (1%)

Table 5.4 Key properties of the average dissimilarity matrix. The average dissimilarity matrix, shown in Table 5.2, is formed by averaging each score element of the 10 subjects' matrices. Percentages in parentheses indicate a proportion of the average dissimilarity matrix's 435 elements.

	Nb. of 1000's	Nb. of 0's
Avg. Matrix	107 (25%)	0 (0%)

zero of perceptual similarity. Humans are very good at finding associations between items, irrelevantly of the amplitude of the contrast between the items presented. However, it can either stand for the maximum non-infinite difference between two stimuli among the set of stimuli displayed, or simply represent a measure of uncertainty of the data. The latter occurs, for instance, if not enough cluster-sorting trials have been run. This dual nature of 1000 scores strengthens the argument that an MDS analysis should assume cluster-sorted data are ordinal/untied.

The actual distribution of scores obtained through a simulation of completely random sorting for this method is shown in Figure 5.6.[5] Also shown is the score distribution for a typical subject who chose 6 clusters for the initial trial, and whose individual MDS result exhibited a clear and plausible clustering.

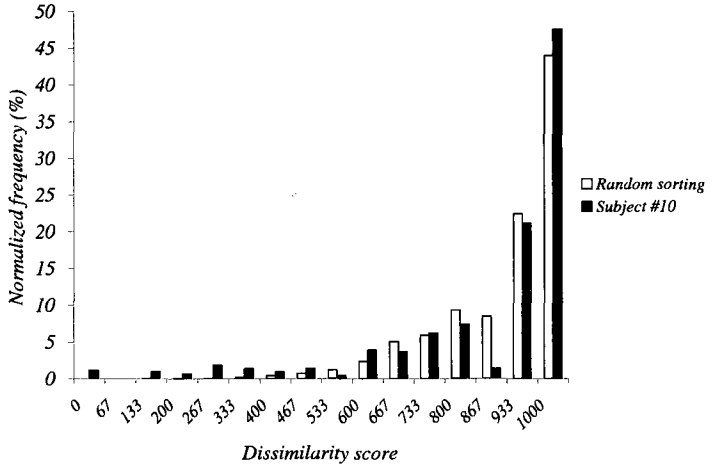

Fig. 5.6 Score distributions for a typical subject and random sorting. Frequency indicates a proportion of the 435 score-elements present in a 30-stimulus dissimilarity matrix.

In both cases, high dissimilarity scores are much more present than low dissimilarity scores, with a peak occurrence of the maximum dissimilarity score of 1000 (accounting for more than 40% of all the scores in the matrix). Interestingly, the

[5]The random distribution was generated by averaging the results of 10000 iterations of the replication of a random classification of the stimuli following the sorting rules described in 5.3.3.

random distribution is almost monotonic, except for a score of 867 that occurs less often than a score of 800, explained partially by the fact that there exists more combinatorial ways for a stimulus pair to obtain a score of 800 than a score of 867.

The first trial allowed subjects to choose a preferred number of sorting clusters. This extra degree of freedom in the experimental procedure is intended to expose what subjects feel is the natural number of clusters for the set of stimuli presented (and which shows substantial and interesting variety, as documented by Table 5.3). However, it adds yet another degree of complexity to the analysis of the dissimilarity data. Figure 5.7 illustrates how the choice of the initial number of clusters affects the random distribution of the 1000-scores across the matrix. We chose to show the effect on the 1000-score frequency since it is the highest, but it is easy to understand that the entire random score distribution is dependent on the initial number of clusters. Of all the sorting trials, the one with the least number of clusters affects the distribution of the 1000-scores across the matrix the most because it forces a higher average number of stimuli per cluster. Therefore, the dissimilarity matrix obtained from a subject who selected two clusters for the initial sort will contain a wider distribution of scores than that of a subject who selected four clusters.

These results indicate that a validation of the subjects' data – for instance obtaining p-values that denote with high probability that the subjects's clustering were not random – is not a trivial task because the random distribution is ill-defined. This also suggests that comparing the scores of the individual dissimilarity matrices through standard techniques such as standard deviation might not be relevant, and should probably not be used to discard subjects.

For the analysis reported here, we validated the individual dissimilarity matrices obtained by trying to find any flagrant inconsistencies in subjects' answers. For instance, we examined the individual MDS plots of each subject. We also compared the number of 1000's between the subjects individual matrices and the random distribution. These inspections did not reveal any inconsistency and therefore we made use all of our subjects' data for our analyses.

Further work is required to assess the significance of the dependencies among the matrix cell values, the discretization of the scores and the unusual shape of the score

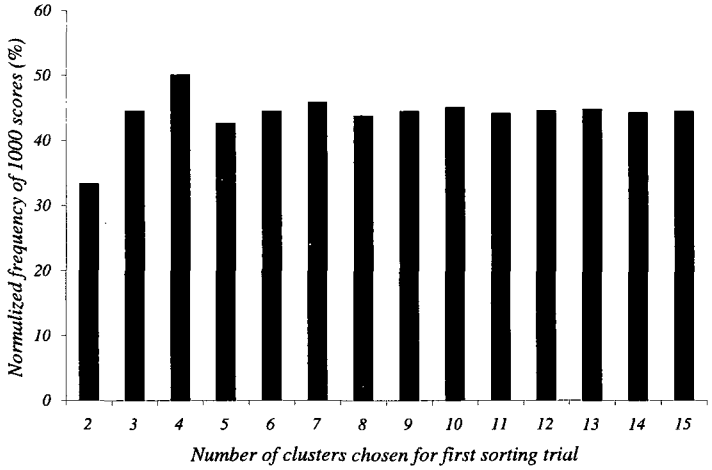

Fig. 5.7 Average normalized frequency of 1000-scores obtained in a randomized cluster sorting execution, as a function of the number of clusters chosen for the first trial. Frequency indicates a proportion of the 435 score-elements present in a 30-stimulus dissimilarity matrix.

distribution because they are likely to have implications on the results and validation of any analysis (MDS or other).

5.5.3 Capture of MDS Group Internal Structure

For our example results, the MDS plot's y-axis seems to absorb different stimulus design parameters for different large groupings (Section 5.3.3). This made us wonder about the extent to which the cluster-sorting method can capture and reveal sub-dominant stimulus distinctions. Clearly, a 2D MDS analysis cannot provide enough detail to resolve group internal structure (if it varies from group to group) because of the dimensionality restriction it imposes on the algorithm; but does a dissimilarity matrix acquired with the cluster-sorting method nevertheless contain the more subtle distinctions specific to a given major group found by a low-dimensional analysis? If present, can it be recovered without resort to a less-intuitive, higher-dimension MDS analysis? To address these questions, we ran a second experiment.

Method: The experiment protocol was similar to the one described in 5.3.3,

with the following differences and key parameters.

We tested five subjects, none of whom had taken part in the first experiment. For the first part of this experiment, subjects were again asked to classify the entire set of 30 stimuli, but this time with a cluster-set of $\{5, 10, 15\}$ rather than the set used for 5.3.3. In the second part, subjects proceeded to execute the same classification task three more times, on the 6 stimuli of three different particular waveform types (TRI, ROLL and EDGE, respectively) and for a cluster-set of $\{2, 3, 4\}$.

Data: From the first part of the experiment, we generated a control dissimilarity matrix (sized 30×30) from which we isolated the three 6×6 submatrices belonging to the subsets being examined for internal structure. From the second part, we produced three full 6×6 dissimilarity matrices, one for each waveform type. We then ran a 2D MDS analysis for the 30×30 control matrix and each of these six 6×6 matrices, and compared the resulting plots.

Results and Discussion: The resulting 2D MDS plot for the control dissimilarity matrix (Figure 5.8(a)) is very similar to our MDS plot obtained with the original experiment (Figure 5.4). Since these were derived from cluster sorts using different cluster-sets, this suggests that cluster-set does not strongly impact the outcome of the MDS analysis.

Moreover, MDS plots (not shown) generated from submatrices of the 30×30 control matrix strongly resemble their intra-waveform matrix counterparts ((b), (c) and (d) of Figure 5.8). Specifically, all graphs display a clear spatial demarcation between stimuli of half amplitude and stimuli of full amplitude. Furthermore, most of the gradation in duration as well as its orientation present on the intra-waveform plots is also visible on the submatrix plots. Thus, our data seems to offer evidence that the cluster-sorting method does capture the underlying structure of non-dominant distinctions.

A key observation is that these subtleties are not apparent from inspection of the global MDS plot based on the control 30×30 dissimilarity matrix. This suggests that re-running the MDS algorithm on sub-sets of the entire stimulus space (thus relaxing the *stress* on the algorithm when it is required to find consistent dimensions for an entire diverse set of stimuli) can offer valuable information on lower structural levels

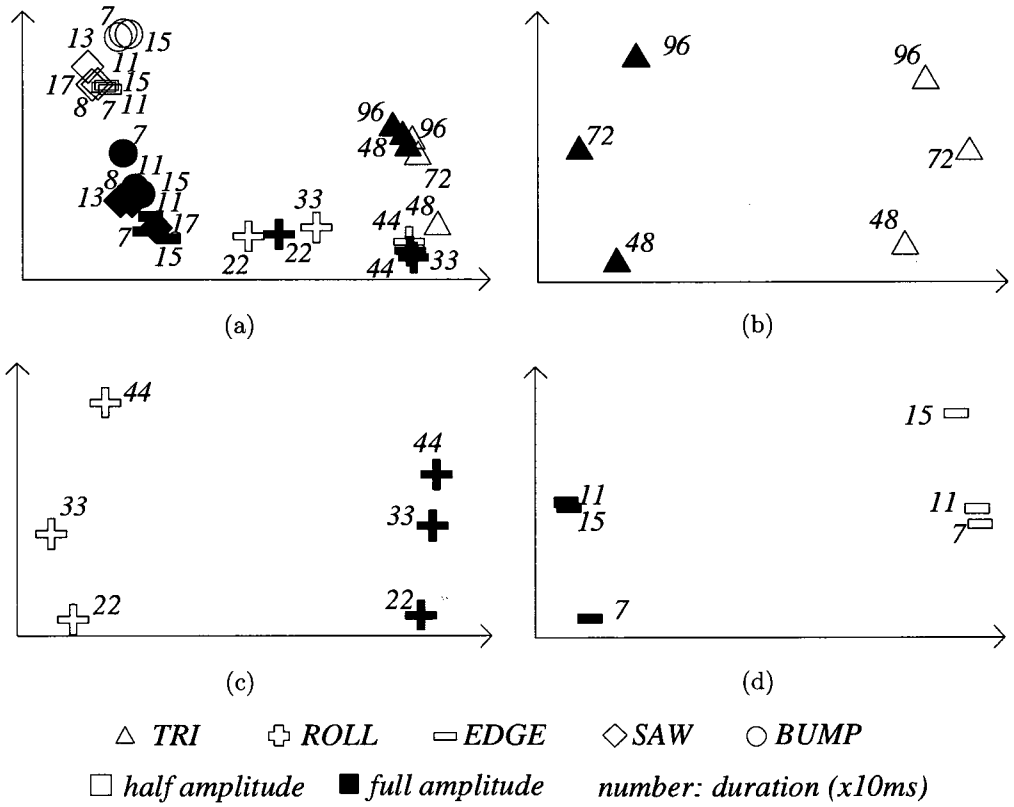

Fig. 5.8 Ordinal untied MDS plots obtained from grouping the different waveforms into cluster-sets of $\{2, 3, 4\}$: (a) control with all 30 stimuli and with isolated (b) TRI, (c) ROLL and (d) EDGE stimuli.

of the entire stimulus set that are not visible from the global MDS plot.

5.5.4 Effects of Noise

Perceptual experiments always suffer from noisy data due to subject fatigue and calibration drift. The effects of perceptual noise on data obtained from a cluster-sorting experiment are hard to quantify, even if assumed to be random, because a fixed level of randomness will weigh differently on the overall noise from one trial to another. The similarity point assignment process tends to amplify noise for trials with a large number of clusters. For instance, the noise-induced erroneous sorting of a stimulus when 15 clusters are present is more significant than with only 3 clusters because an improperly matched pair is assigned 15 similarity points in the former case and only 3 in the latter. To complicate things, it can be argued that a 15-cluster discrimination task will generate more uncertainty than a 3-cluster task because it imposes a heavier cognitive load on the subject. If this is true, it would mean that the most noisy part of the data is also precisely the one that accounts for the largest similarity scores (or lowest dissimilarity scores).

Despite the fact that we are currently unable to account for the level of noise in the data, we are not too alarmed by this potential problem since, as reported in 5.5.2, our results were reproducible. This suggests that, even though noise might get amplified for the trials with a large number of clusters, the initial amount of noise generated by applying the cluster-sorting technique is not important enough to significantly distort the dissimilarity data. However, if this was to become a problem in future experiments using the same technique, one possible solution would consist of applying a corrective factor that decreases the weight of trials containing a large number of clusters.

5.6 Conclusions and Future Work

The advantages of the cluster sorting algorithm for generating perceptual MDS input data are significant: in particular, it is superbly efficient in both time and consistency of subject response. The concerns we have identified are at minimum important to

be aware of. Based on the reported investigation, advantages seem to outweigh the concerns, and it appears justifiable to continue using this technique with well-informed caution. With respect to each of our Section 5.4 questions, we specifically conclude:

Q1 Perceptual data obtained using this cluster-sorting technique is best assumed to be ordinal, and analyzed with a non-metric MDS algorithm.

Q2 Key characteristics of a cluster-sorted dissimilarity matrix, in particular a restriction on element independence, may impact MDS output. Further investigation is required to quantify both the magnitude of this restriction, and the robustness of MDS algorithms to such inter-matrix correlations.

Q3 Hidden patterns in the resulting plot of a MDS analysis carried out on the entire stimuli set can become apparent when selected submatrices of the full dissimilarity matrix are submitted to the same MDS algorithm. This tends to indicate that the cluster-sorting technique also captures detailed information about sub-level of stimuli distinction that are not visible by sole inspection of the global MDS plot.

Q4 It is difficult to make sense of the distribution of noise across the dissimilarity matrix obtained from the cluster-sorting method (mainly due to noise amplification for trials with a large number of clusters). However, the fact that we were able to replicate our results over two independent experiments seems to indicate that the data-gathering technique is not generating enough noise to significantly distort the data.

This work represents the first step in an ongoing effort to develop a suitable mechanism for analyzing the perceived differences and usability of haptic icons. Our results for questions 2 and 4 in particular need further development. Finally, we are also considering alternatives to MDS for perceptual analysis of haptic icons.

Acknowledgment

The authors would like to thank the reviewers for their contribution to the clarity and completeness of this paper. We specifically thank Prof. Lawrence Ward from the Department of Psychology at the University of British Columbia for an informative discussion about the cluster-sorting method, and Don Pavlasek and Jozsef Boka from the Mechanical Workshop in the Electrical and Computer Engineering department at McGill University for their invaluable technical support in building the THMB device.

Chapter 6

Tactile Feedback Can Substitute for Visual Glance in Mobile Interaction

Preface to Chapter 6

Chapter 6 integrates the knowledge acquired through out the work described in the preceding three chapters and exercises the capabilities of the THMB device. After having implemented the concept of lateral skin stimulation, embodied the technology into a handheld prototype and validated its perceptual characterization, all the tools were available for the design of a mobile interaction for which a measurable added-value of tactile feedback could be exposed.

Contributions of Authors

Pasquero conducted the entirety of the work that is described in this chapter. Hayward supervised and edited the manuscript.

Tactile Feedback Can Substitute for Visual Glance in Mobile Interaction

Abstract

In this paper we describe two experiments using a hand held device equipped with a transducer that is capable of providing sensations much richer than vibrations and thumps. Our experimental hardware and software platform makes it possible to program sophisticated interaction metaphors where input and feedback are tightly coupled. In the first experiment, we modeled a scrolling task inside a long list, performed in concurrence with a distractor task. We found a marginal improvement in performance due to the tactile feedback but that came at a higher attentional cost dedicated to operating the device. This effect was manifested as a net decrease in performance in the distracting task. In the second experiment, we improved the interaction paradigm and the task model. To evaluate the effect of tactile feedback on visual allocation, instead of measuring performance in a secondary task, we measured the time consumed by subjects looking at the screen to perform the scrolling task. In this experiment, we demonstrated a 28% decrease in subject's reliance on vision when tactile feedback was enabled.

6.1 Introduction

Manufacturers of mobile devices are well-aware of the importance of providing the richest tactile experience possible to the end users. In the past, this meant designing keypad buttons that felt and behaved as expected; buttons could not be too large or too small, and they needed to return tactile cues, such as sharp detents, that communicated their state and their function clearly [140]. Failure to take these factors into account resulted in devices that were not adopted by the users. While these observations are still relevant today, manufacturers of cell phones and personal digital assistants are now also exploring different ways to design devices capable of active and programmable tactile stimulation. Their efforts reflect a desire to take advantage of the sense of touch in order to address the limitations introduced by the small size of mobile devices.

Mobile devices provide interaction affordances which inherently are more restricted than those possible with desktop solutions. This is particularly true of information visualization possibilities [28, 76]. By their very nature, mobile devices must sometimes be used in loud or poorly-lit environments where the users' auditory and visual channels are degraded, and where the user's attentional resources must be shared among several competing tasks. According to the context, mobile devices must contend with requirements that can range from producing alerting signals that must not be missed (e.g., during rescue operations) to avoiding distractions (e.g., during meetings), with many cases in between. Sometimes, users must be notified of important events or state changes, without impairing their capability to function otherwise (e.g., when driving [135]). In many cases, mobile devices endeavor to push large amounts of information through narrow technological channels and through perceptual channels that are already operating close to saturation.

In this context, the largely untapped tactile feedback channel is a natural candidate for addressing the difficulties in giving mobile devices true interactive capabilities. The tactile modality constitutes a rich and usable channel for human-computer interaction that has, so-far, resisted many attempts at a more systematic use.

The small size of the screens found in mobile devices poses a special challenge.

Because of the lack of display space, information in mobile devices is usually organized like a deep hierarchical structure, rather than a wide structure more typical of personal computers [28]. To access content, users of mobile devices are often confronted with the cognitively demanding process of navigating between multiples layers of menus in the interface. Interruptions from the environment complicate this task further. Providing tactile feedback to support document navigation has the potential to reduce user disorientation and the amount of attentional load that must be devoted to the device. This would free up cognitive resources that users could choose to allocate to other activities.

We explore this avenue using a hand held prototype that was specifically developed for this purpose. This prototype is capable of providing much richer tactile feedback than that possible using vibrotactile stimulation. The device serves as a flexible hardware and software platform for the design and evaluation of application concepts that take advantage of a combination of graphical and tactile feedback.

In this paper, we investigate two task models where we attempt to use tactile feedback to simultaneously enhance interactive performance and decrease reliance on vision. The results of a first experiment do not reveal any significant enhancement due to tactile feedback. Nevertheless, they inform the implementation of a second experiment by underscoring the importance of designing for a closer link between perception and input action. In the second experiment, tactile feedback allows subjects to complete a scrolling task more efficiently for a measurably lower visual cost.

6.2 Related Work

Research in multimodal interaction has long regarded tactile feedback as an indispensable channel. For instance, in the 1960's vibrotactile feedback was proposed as an aid to jet pilots [33]. Since then, it has been shown that when used in combination with vision in an interface, the tactile channel can provide added value in the form of an enhanced experience and/or a measurable performance increase [21]. A combination of recent advances in the miniaturization of actuator technology and the increasing understanding of the touch modality as a channel for artificial com-

munication have sparked an exploration of the possible ways to leverage touch in mobile interactions ([25,41]). In the last few years, the use of brief artificial tactile means to convey abstract information has been studied for mobile interactions with the following goals:

- to compensate for limited screen real-estate and cognitive overload [90],
- to communicate information rich in emotional content [133],
- to provide immediate tactile feedback that is closely linked to gestural input control [24,106,124], and
- to substitute for a visual interface [20,123].

Other contextual possibilities, all of which aim at promoting tactile feedback beyond its present role as an alerting mechanism, have also emerged.

Poupyrev et al., and then later Oakley et al., have studied the effects of providing vibrotactile feedback during a scrolling task that requires subjects to tilt a device for position and rate control [106,124]. These studies show improvements in scrolling performance in the form of lower error rates and faster time-to-target values resulting from the introduction of vibrotactile signals that confirm the results of gestural input.

Many approaches to stimulate the skin with miniature electromechanical devices exist [115]. To date, those approaches that are sufficiently practical to be integrated in a mobile device employ vibrations as the stimulation method, generated either by miniature DC motors spinning an eccentric mass, by voice-coils, or by piezoelectric actuators [88]. These miniature systems have the compact form factor, the robustness, and the low power consumption required by mobile devices. Recently, simple vibrations have been shown effective in communicating tactile messages that are suitable for human-machine communication [41]. However, the sensations evoked by vibrotactile stimulation can only convey a subset of all the expressive capabilities afforded by touch.

Over the last few years, we have been investigating a mode of skin stimulation that moves away from the conventional vibrotactile paradigm [93]. By using an array

of piezoelectric benders that exert distributed lateral deformation against the skin, it is possible to convey to users sensations of both fixed shapes (e.g. Braille dot and small-scale textures) and moving features under the fingerpad [92, 117]. We refer to this mode of interaction as *laterotactile* stimulation. Because the piezoelectric benders it employs can be operated over a large band of frequencies, laterotactile stimulation is capable of communicating a wide collection of signals that range from the sensation of a small dot to high-frequency textural sensations. The work we present here takes advantage of these capabilities.

6.3 Experimental Platform

We built a laterotactile transducer small enough to fit in a hand held object. The system, which serves as a research platform for multimodal interaction, is dubbed the THMB for Tactile Handheld Miniature Bimodal device [117]. It comprises a tactile display that is co-located with a liquid crystal display (LCD) to enable the coincidental display of graphical and tactile information. The present studies have been carried out with a second generation prototype of the THMB[1] that includes a number of enhancements suggested by the results of a perceptual characterization study [95], including an increase in the strength of the tactile sensations.

For experimental flexibility, the hand held enclosure is tethered to interfacing and power electronics. The system is driven in realtime from an external computer via a USB link that allows the actuator commands to be refreshed at 800 Hz. The enclosure is held in the left hand (Figure 6.1). The laterotactile transducer protrudes slightly from a rectangular opening located under the thumb. It is mounted on a slider so it can be moved up and down with a flexion of the thumb (Figure 6.1-c). The assembly also acts as a push-button when pressed in the direction normal to the sliding axis.

The slider is spring-loaded, as shown on Figure 6.2(a). This provides a passive resistive force against thumb movement that is helpful in locating the position of the transducer over the slider range, which is divided into three distinct zones. The top

[1]The reader is referred to the appendix for drawings of the new version of the device.

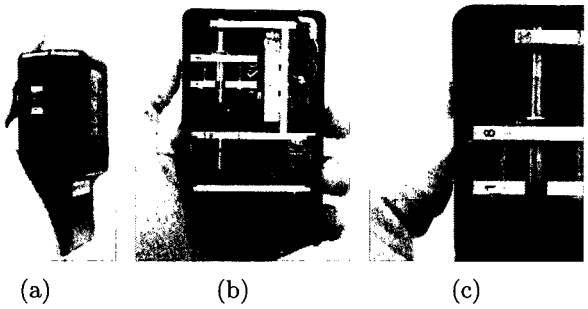

Fig. 6.1 The THMB device: (a) handheld case that combines a LCD and a tactile transducer, (b) sliding mechanism that also acts as an inward push-button and (c) close-up on the slider with range of sliding motion.

and bottom zones are spring actuated while the central neutral zone is unaffected, as illustrated in Figure 6.2(b).

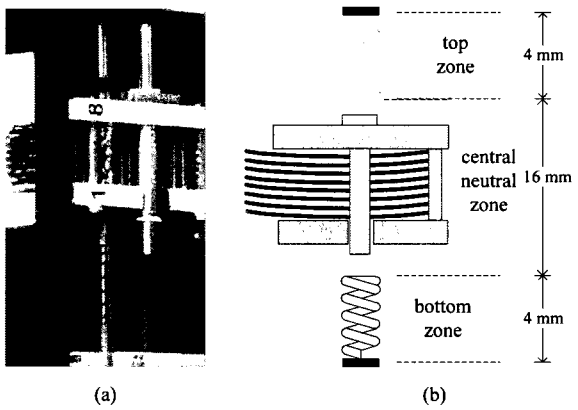

Fig. 6.2 Slider regions: (a) single spring that acts both in extension and in compression to bring the transducer back to a central neutral zone and (b) three distinctive sliding regions.

The laterotactile transducer has an array of eight 0.5 mm-thick piezoelectric benders activated by modulating the voltage across their electrodes. Stimulation is provided by programming the benders to cause tangential traction at the surface of the skin. The tips of the benders were beveled to increase their traction capacity. The display provides a 10 × 8 mm active tactile surface.

We also developed a graphical editor to support the design of laterotactile animations. The icon editor makes it possible to rapidly prototype and test tactile animations composed of traveling waves of different waveforms, trajectory profiles and amplitudes.

6.4 Experiment I

In the first experiment, subjects were asked to perform a scrolling task that displayed the motion-control behavior described in the next subsection. Because users are rarely scrolling in a mobile context uninterrupted, we also designed the experimental procedure to include a second phase that modeled interruptions from the environment. This was done by having subjects perform a second task in parallel to the scrolling. Subjects were required to monitor a distractor appearing on a separate screen. While scrolling performance was not expected to be enhanced significantly by tactile feedback, we nevertheless tested for that possibility. A more likely expectation was that tactile feedback would help subjects cope with interruptions in the multitasking situation.

We intended to model the situation facing mobile device users when locating an item in a very long list ($\approx$ 1000 items). Such tasks arise when one is searching for a correspondent in a list of contacts stored in a mobile telephone or looking for a particular song in a portable music player. Our model posited a conventional scrolling interface where a user is given control over scrolling movement up and down a list, by means of an input device.

6.4.1 Control Metaphor For Long Lists

Referring to Fig. 6.3(a), the subject could see a window of 7 items in a list of 1000 alphabetical items. There were markers to help structure it. A star symbol indicated a change in the first letter as in a dictionary and a vertical dash symbol was regularly placed every four items. Each of these symbols was associated with an easily recognizable tactile icon in an effort to assist navigation. Star icons had precedence over dash icons in case they would both fall on the same item. The

distractors were strings of digits displayed on an external screen that forced the subject to switch their gaze back and forth. The task, further described below, was to navigate to a target item in the list while spotting the occurrence of zeros in the distractors.

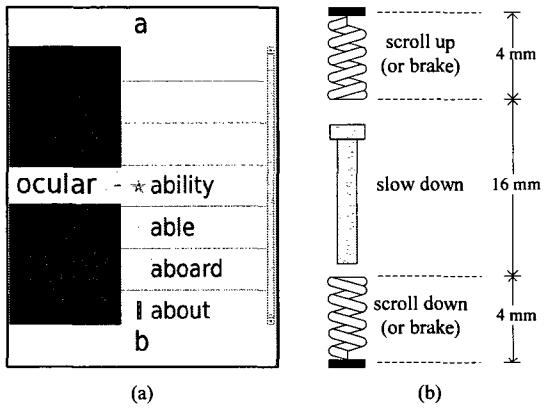

Fig. 6.3 Control metaphor for first experiment. (a) List of 1000 items with the target displayed to the left. A star next to a list item indicates a change in the first letter of the word. Vertical lines mark every four elements. (b) Driving control metaphor with three distinct regions.

A control metaphor based on input acceleration is well adapted to navigating inside long lists. This was implemented by setting the scrolling speed and direction of the scrolling motion according to the deflection of the button from its rest position. To limit the amount of attention required by this mode of input, velocity was prevented from winding up to infinity. The larger was the deflection, the higher was the terminal velocity. At maximum deflection, the acceleration and terminal velocity were set at 10 items/s^2 and 50 items/s respectively. If the subject partially released the button, diminishing deflection, the velocity would diminish too. Complete return to the neutral position would cause the scrolling velocity to decay exponentially to zero in 800 ms. It was also possible to bring the scrolling movement to rest by deflecting the button in the direction opposite to the scrolling velocity.

Tactile feedback was provided in the form of two types of highly distinguishable tactile icons: a short traveling wave was triggered each 4 list items while a longer

wave of higher spatial frequency indicated a change in the first letter of the sorted list (Figure 6.4(a)). Waves always traveled against the scrolling motion of the list to reinforce the sensation of moving over items as they passed by. When not activated, the piezoelectric actuators of the transducer were programmed to all bend towards the top (Figure 6.4(b)). As a tactile icon passed by, each actuator exhibited the same full back-and-forth bending motion to the bottom, but with a phase difference. To avoid any tactile icon interference at high scrolling speeds, the tactile animations could not be displayed too long on the tactile transducer. A traveling speed of 300 mm/s was selected because it combines the advantages of being fast and efficient at delivering directional cues [95]. The flow of tactile sensations under the thumb was similar to equally spaced markings on a road. As a result, it was possible to appreciate the speed of scrolling in the absence of visual information [24].

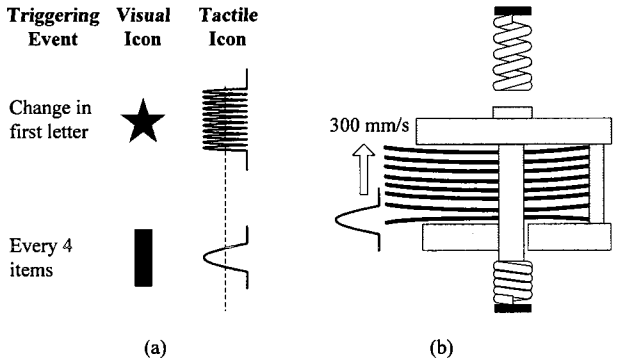

Fig. 6.4 Tactile animations. (a) Every four elements were marked by the sensation of a traveling pulse and a change in the first letter felt like a rapid vibration. (b) Tactile icons traveled across the tactile display against the scrolling motion.

6.4.2 Experiment Design

Sessions

Sessions were divided into two phases. In the first phase, subjects were uninterrupted (−INT). They were asked to perform a simple scroll-to-target task under two condi-

tions: with the tactile feedback enabled on the controller (+TF) and without (−TF). In the second phase, subjects were constantly interrupted (+INT). They were given the additional task to record the appearances of a visual distractor on an external screen.

In the first phase, subjects familiarized themselves with the scroll-to-target task during a short training session. Next, they executed the task for a set of 22 trials with the tactile feedback disabled and for another set of 22 trials with the tactile feedback enabled. The order was randomized.

During the second phase, subjects had to perform the distracting task concurrently with the primary scrolling task. They were given a training session to familiarize themselves with the simultaneous execution of the two tasks. The presentation order of the conditions with and without tactile feedback was also randomized across the subjects to minimize learning effects. The experimental phase with the distractor task was nevertheless conducted after the basic task to make sure that subjects were comfortable with the easier scrolling task before introducing the more strenuous distractor task. At the end of the session, subjects were given a Likert-type questionnaire to fill.

Subjects

Eleven subjects participated in the first study. Results from one subject were discarded. The subject in question was unable to focus on the task and gave erratic answers. The remaining subjects consisted of students (6 females and 4 males) between 18 and 27 years old (average: 20.9). All subjects were right-handed except one and most were regular mobile device users (8/10). Subjects were paid $25 to participate in the study, which lasted 1.5 hours.

Tasks

Subjects were instructed to navigate to a target item which was displayed to the left of the list. They were encouraged to reach the target as quickly as possible and with high accuracy. Because we were interested in the initial approach to the target rather

than its precise selection, targets were considered to be reached as soon as they were visible anywhere on the screen. Once a target was reached, subjects pressed on the device push-button to move to the next trial.

For the dual task condition, subjects were also instructed to record the presence of a visual distractor that appeared on an external monitor. The secondary task required subjects to press a key with the index finger of the hand not holding the device whenever the distractor '0' appeared on a second monitor. To prevent subjects from guessing the next appearance of the distractor, a 3-digit string was continuously refreshed on the monitor at irregular time intervals (1-3 seconds). Sometimes, one of the 3 digits displayed would be the distractor '0'. This procedure forced the subjects to switch gaze back-and-forth between the device screen and the second screen. In order to standardize the number of distractors for all subjects, the display of the distractor was triggered when specific index positions on the list were visited, rather than at random time intervals.

Trial Sets

The same set of trials was used under each experimental condition. It was composed of 22 distinct pairs of initial and target index positions on the scrolling list. Because the lists contained different words across the experimental conditions, the same index positions corresponded to different items on the list. Distances between the initial and target trial positions varied: six trials consisted of short distances of 50 items (D50), another six trials of medium distances of 250 items (D250) and the remaining ten of distances of 500 items (D500). Half of the trials required a bottom-to-top scrolling motion through the list and the other half a motion in the opposite direction. Trials were presented in pseudo-random order for each subject.

Corpus

To minimize learning effects, four lists of different items, one for each experimental condition, were used. A fifth list was also used for all training phases. The lists contained 1000 common 5-to-7-letter English words sorted in alphabetic order, with

the number of words starting with each letter selected according to Solso et al. [143].

Data Recorded

In both phases of the experiment, the scrolling trajectory profiles were recorded. The data collected included the time taken to reach the targets and the number of overshoots, defined as a scrolling motion that missed the target so far that it no longer was visible on the screen.

6.4.3 Results

The results are collected in Table 1. In all conditions, statistical significance indicated by figures in bold characters is the result of a Wilcoxon matched-pairs signed-rank test which does not make an assumption of normality for the data.

		−INT		+INT	
	Distance	−TF	+TF	−TF	+TF
avg. time to target (s)	D50	11.8	12.2	13.1	13.2
	D250	21.4	21.3	25.1	23.8
	D500	**24.4**	**26.5**	28.1	31.1
	all	20.2	21.2	23.20	24.25
number of overshoots	D50	2.8	2.4	1.7	1.9
	D250	3.3	3.7	**5.1**	**3.6**
	D500	**5.9**	**3.9**	6.7	7.2
	all	12.0	10.0	13.5	12.7
% distractors caught	D50	-	-	58.3	68.9
	D250	-	-	**77.0**	**65.0**
	D500	-	-	80.5	80.0
	all	-	-	71.9	71

Table 6.1 Raw data for the first experiment. Bold fonts indicate statistically significant differences in performance that can be attributed to the presence or absence of tactile feedback.

As was expected, we found scrolling performance variations between the conditions with and without the distractor task. On average, subjects took $\approx$ 3.2 s longer to reach the targets and did $\approx$ 2.1 more overshoots. On the other hand, we did not

find any overall statistically significant difference between cases with tactile feedback and without it. A closer look at the data broken up according to scrolling distances (Table 6.1) reveals some effects due to tactile feedback. These effects are summarized here:

- Under the undistracted condition, tactile feedback decreased the number of overshoots ($p < 0.05$) and increased the time to reach the target ($p < 0.05$) for distances of 500 items.
- Under the distracted condition, tactile feedback also decreased the number of overshoots ($p < 0.05$) but increased the proportion of distractors correctly detected ($p < 0.05$) for distances of 250 items.

The results from the Likert-type questionnaires are summarized in Figure 6.5. In general, subjects reported that both tasks were challenging. They found the tactile feedback to be helpful and pleasant but they did not believe it allowed them to reach the target faster or with better accuracy. They did not perceive the tactile feedback as being annoying and thought it felt different than the typical vibrotactile signal of a mobile phone. When tactile feedback was not present, subjects found that something was missing but could not verbalize what was missing. The overall size of the case enclosure was found to be too large for the average hand.

6.4.4 Discussion

In general, we did not find highly apparent performance improvement resulting from tactile feedback, except for the two cases where a decrease in the number of scrolling overshoots could be observed. This performance increase was at the cost of longer times to reach the target in one instance (D500 −INT) and at the cost of an increase in the number of distractors missed in another (D250 +INT). These findings reflect that different scrolling strategies were used by the subjects for the different scrolling distances. They also suggest that the tactile feedback acted more like an alerting signal rather than a navigation aid. Subjects might have reduced their scrolling overshoots because they were prompted to allocate more attention to the task of scrolling rather than to the task of detecting distractors.

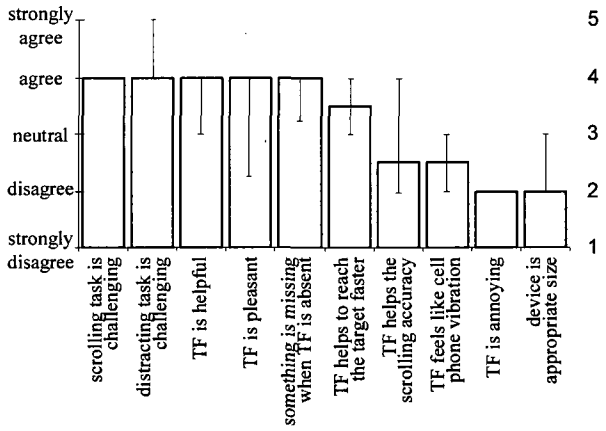

Fig. 6.5 Results from Likert-type questionnaire for first experiment. Mode results are reported. Error bars represent interquartile ranges (IQR).

A decrease in the number of overshoots due to tactile feedback could be viewed as showing an encouraging trend for improvement in overall scrolling performance. However, in the context of mobile interaction, this is not sufficient. Improved mobile interactions should liberate users from attending to the device and free up more time to attend to other tasks, such as monitoring the environment. Performance improvement achieved by drawing more attention to the device may be not viewed as a significant advance. Also, it is likely that despite the fact that our device gives sensations that are very different from that of conventional mobile phones, tactile stimulation may still produce the undesired behavioral consequence of grabbing attention, much like an auditory signal.

The results of this experiment also illustrate how tactile feedback can sometimes induce subjects to mistakenly believe that they perform better. For this experiment, performance measurements showed that tactile feedback was at best marginally useful and at worse detrimental, even though most subjects believed that it helped them.

6.5 Experiment II

The results of the first experiment motivated us to propose another task model intended to better represent the requirements of mobile interaction. The main design objective remained the same: multimodal mobile interaction, but with a focus on the synergetic integration of the control metaphor and the tactile feedback. Also the mode of testing for the subjects' behavior was modified. Instead of measuring speed and error rate, as is conventionally done, we attempted to probe subjects' behavior by recording the frequency at which they needed to look at the screen to perform the task.

To do this, we designed a very simple testing paradigm. The list was by default not visible on the screen. Each time the subjects felt that they needed visual information, they had to press and hold a key to make the list visible. The list was also modified from the first experiment to display a more obvious structure. It was made of the numbers 1 to 100 for easier navigation. This way, by recording the number of keystrokes and the manner in which subjects moved inside the list, we could paint an accurate picture of their behavior under different testing conditions. In addition, we re-designed the control metaphor to facilitate fine positioning. The overall task however remained that of navigating to a target item.

This paradigm attempts to model a type of mobile interaction where the human sensorimotor system is allowed to operate in closed loop with the device. This can be possible only if tactile feedback is mapped to motor behavior in such way that the loop is closed *without* consuming attentional resources. For this reason, we redesigned the controller to enable tighter coupling between user input and the resulting tactile sensations. The objective was to make the tactile feedback feel like the natural outcome of operating the controller, rather than like the decoupled sensation of a tactile track that is scrolling under the finger in parallel to the visual track. To achieve this, we associated a tactile icon to each item of the scroll list and designed a controller that supports fine positioning rather than an initial approach to a target from a long distance away. The design is rather involved so it is described in the next subsection.

6.5.1 Design of a Fine-Positioning Controller

A group of 8 subjects was first recruited to facilitate an iterative design process comprising a cycle of short informal evaluation sessions and calibration sessions. Six subjects had never tried the device prior to the evaluation session, and another two were graduate students who were very familiar with the project. The 6 naive subjects were paid $10 for their time. They also contributed their time to preliminary, pilot sessions prior to the full scale experiment. During calibration sessions, subjects operated the THMB device by trying different mappings during a scrolling task. They also described their experience to the designer. These comments were instantly addressed by adjusting the control parameters such as intensities and gains. The tool that made this process possible was an elaborate sensorimotor coupling interactive editor. At the conclusion of the design process, we were confident to have achieved an efficient and natural control metaphor that needed hardly any learning to be adopted by the subjects.

The resulting controller provides three scrolling modes that correspond to three different slider deflection regions: a stop mode, a discrete mode and a continuous mode (refer to Fig. 6.6). The deflection range associated to the stop mode roughly corresponds to the slider's central neutral zone that is unaffected by the spring return. When the transducer is in this region, the list is static. A small push against one of the boundary zones enables the discrete mode which causes the list to scroll by exactly one item. Pushing further into the spring zone switches the controller to a continuous scrolling mode where the list scrolls with an acceleration that is proportional to the amount of deflection, similarly to the controller of Experiment I, with the difference that the top acceleration and top terminal velocity were slower (a=3 items/s^2, v=12 items/s). Finally, when the button is released, the spring load returns the transducer back to the neutral zone where the list is brought to rest.

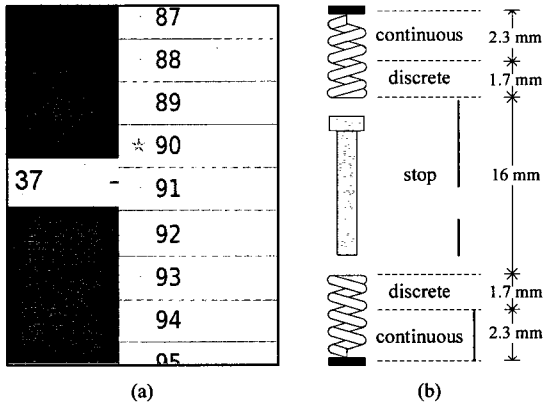

Fig. 6.6 Control metaphor for second experiment. (a) List of 100 numeric items where every 10 elements is marked with a high-frequency vibration. (b) Controller modes for the scrolling metaphor.

6.5.2 Experimental Design

Sessions

Experimental sessions were divided into three blocks of trials that corresponded to three different conditions: a training condition, a control (CTRL), and a condition with tactile feedback (TACT). The experimental task was the same for all three conditions. For the training and the control conditions, the tactile feedback generated by the controller was disabled. The training condition and the tactile condition were preceded by short introductory sessions where subjects were given time to get comfortable with the interaction and the experimental procedure. The training condition was always performed first. Next, in order to minimize learning effects, half of the subjects performed the control condition before the tactile condition, and the other half in reverse order.

The training condition had several aims. It allowed subjects to become familiar with the task and also provided them with a reference for the next two conditions. In the control and tactile conditions, subjects were reminded to minimize the number of key strokes and their duration. Finally, the training condition allowed the experimenter to ensure that subjects were performing the task at a pace compatible

with a natural range of operation. Abnormal behaviors, such as a subject trying to keep the screen invisible as long as possible, could be avoided this way.

Subjects

Ten graduate students between 25 and 34 years of age (mean 28.3) were recruited for this experiment. They were compensated $10 for their participation that lasted about 45 minutes. All subjects, but one, were right-handed. None of the subjects had participated in any of the previous phases of evaluation and calibration of the device.

Task

Subjects were instructed to scroll to a target and to select it. The list was made of 100 consecutive integers and was invisible by default. To make a viewing window, subjects had to press and hold down a key. They were instructed to minimize the number of key presses and their duration. The target number was displayed to the left of the list and was considered to be reached when it was aligned with a horizontal cursor located at the center of the screen (Fig. 6.6). Once the target was reached, subjects pressed the push-button to move to the next trial.

Tactile feedback was provided each time the scrolling list changed position. For every item that would move by the horizontal cursor on the screen, a traveling wave was felt under the thumb moving in the opposite direction to the scrolling motion. Tactile animations were very similar to the ones for Experiment I. Each list item were represented with sensation that could be best described as that of a moving pulse and items that were multiples of 10 triggered a tactile sensation closer to that of a vibration. The tactile feedback, therefore, carried information that combined relative motion cues and absolute position cues.

Trials

The block of trials under the control and tactile conditions were exactly the same but the presentation order of the trials were randomized. A trial was defined by an

initial index position and a target position that the subject needed to reach. Pairs of initial/target positions were separated by distances of 5, 8, 10, 30, 40 or 50 items. Each distance was represented 4 times, for a total of 24 trials per condition. Half of the trajectories required subjects to scroll upwards and the other half required them to scroll in the opposite direction.

Data Recorded

For each trial, we recorded the scrolling trajectory profiles and the times of occurrence of all events.

6.5.3 Results

We analyzed the difference between the control and tactile conditions for a number of dependent variables. Table 6.2 summarizes the average results over the ten subjects. It is interesting to note that the time windows for which subjects were pressing the key (average glance duration) were short (0.3 s). Hence, the experimental procedure modeled a multitasking situation where subjects are quickly looking at the device.

Statistical analysis using the Wilcox matched-pairs signed-rank test reveals a significant difference between the control and tactile conditions in the number of key strokes (associated to visual glances), the time between two key strokes (switch-back durations), and the number of overshoots. Closer analysis indicates a decrease of 28% in the number of key strokes required to reach the target when tactile feedback is enabled (Fig. 6.7) ($p < 0.01$). All ten subjects could reduce their dependance on the graphic screen when tactile feedback was provided. However, the decrease in the number of glances was not distributed evenly, ranging from a 2.5% drop for subject F to a 49% drop for subject A (Fig. 6.8).

The effect of tactile feedback can also be seen in an increase in the switch-back duration, defined as the average time spent attending to the environment, modeled here as an invisible viewing window (Fig. 6.9) ($p < 0.01$). With the exception of subject I, all switch-back durations collected correspond closely to those observed in a semi-natural field study during simple mobile interactions [109]. Subject I's

Independent Variable	CTRL	TACT
trial set duration (s)	373.2	394.8
visib. proportion (%)	7.8	6.0
nb. of push-button clicks	38.0	30.9
avg. glance duration (s)	0.3	0.3
nb. glances	**97.6**	**69.8**
switch-back duration (s)	**4.1**	**6.3**
nb. overshoots	**9.3**	**14.2**
overshoot length	4.4	3.0

Table 6.2 Summary results for 2nd experiment. Data represent average over the ten subjects. Statistically significant differences are shown in bold.

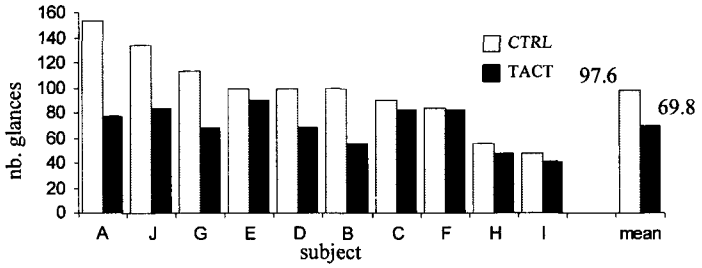

Fig. 6.7 Number of glances at the screen, or key strokes on the laptop.

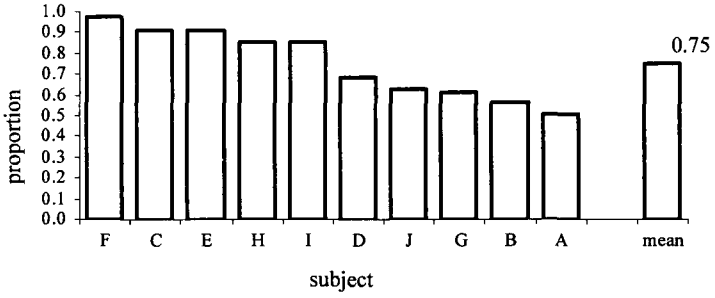

Fig. 6.8 Normalized proportion of the number of key strokes (TACT) with respect to the CTRL condition.

abnormally high switch-back durations when compared to the other subjects are probably a reflection of a strategy that consisted in counting exactly the number of

tactile stimulations felt. This strategy, however, resulted in a decrease of the number of key strokes in accordance with the results of the majority of subjects (Fig. 6.8).

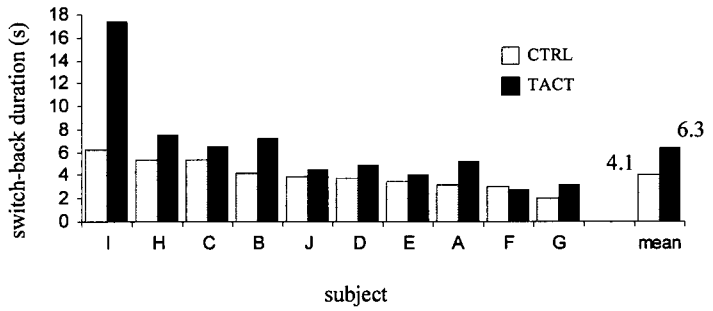

Fig. 6.9 Switch-back durations.

Results also show a statistically significant increase of 52% in the number of overshoots when tactile feedback is introduced (Fig. 6.10 ($p < 0.05$)). For eight out of the ten subjects, we observed an increase that ranges from one extra overshoot for subject G to a more than threefold increase for subject C. For the remaining two subjects, the number of overshoots dropped by one across the entire set of trials.

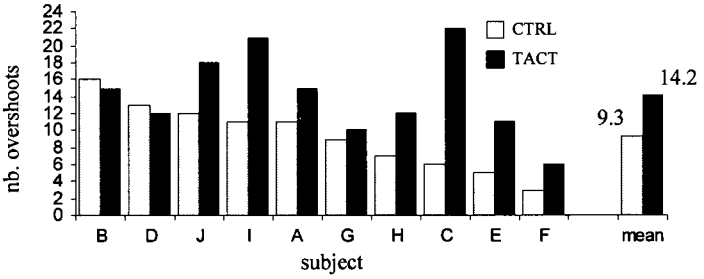

Fig. 6.10 Number of scrolling overshoots.

6.5.4 Discussion

We investigated the different scrolling strategies elicited by tactile feedback. When tactile feedback is disabled, subjects rely between glances on a mental model they

have constructed during the training phases or during previous phases of the experiment. They navigate by "dead reckoning" between glances. However, with the introduction of directional tactile cues that were well mapped to the control inputs, subjects are able to adjust their scrolling strategies and reduce the frequency at which they feel the need to look at the screen.

A inspection of the trajectory profiles gives extra information regarding the different scrolling strategies employed by the subjects. For instance, Figure 6.11 shows samples of trajectory that clarify why the overall number of glances is reduced and why there is an increase in the number of overshoots. In these four cases, the subjects' reliance on visual input was reduced. Note that not all trials are that explicit.

Figure 6.11(a) indicates that it was hard for the subjects to control precisely the movements of the list when there is no tactile confirmation of the consequences of their motor output. Subject C mistakenly estimated that the target was reached when it was still three items away. This prompted the subject to use an extra glance at the screen.

Figure 6.11(b) shows that tactile feedback can also be useful for long jumps. In the control case, subject D monitors a stop-and-go scrolling action with repetitive glances. Under the tactile condition, subject D initiates an initial approach to the target in continuous mode, eventually takes a look at the screen, and finally switches to the discrete mode to bridge the small distance left to the target. The combination of a continuous-mode approach with a discrete-mode fine positioning results in a smooth landing on the target that is reminiscent of optimal control strategies in feedback control theory [68].

Figure 6.11(c) gives another example of a strategy that utilizes different scrolling modes. The case with tactile feedback shows an example of an accidental push-button command at the start of a continuous scrolling motion. This typically happens whether tactile feedback is displayed or not and is an indication that the device ergonomics is not well matched to users' hands. When subjects engage the slider in continuous mode, they often accidentally press the button.

Finally, Figure 6.11(d) depicts an example of an overshoot-plus-recovery under the tactile feedback condition. In this case, subject D initially underestimates the

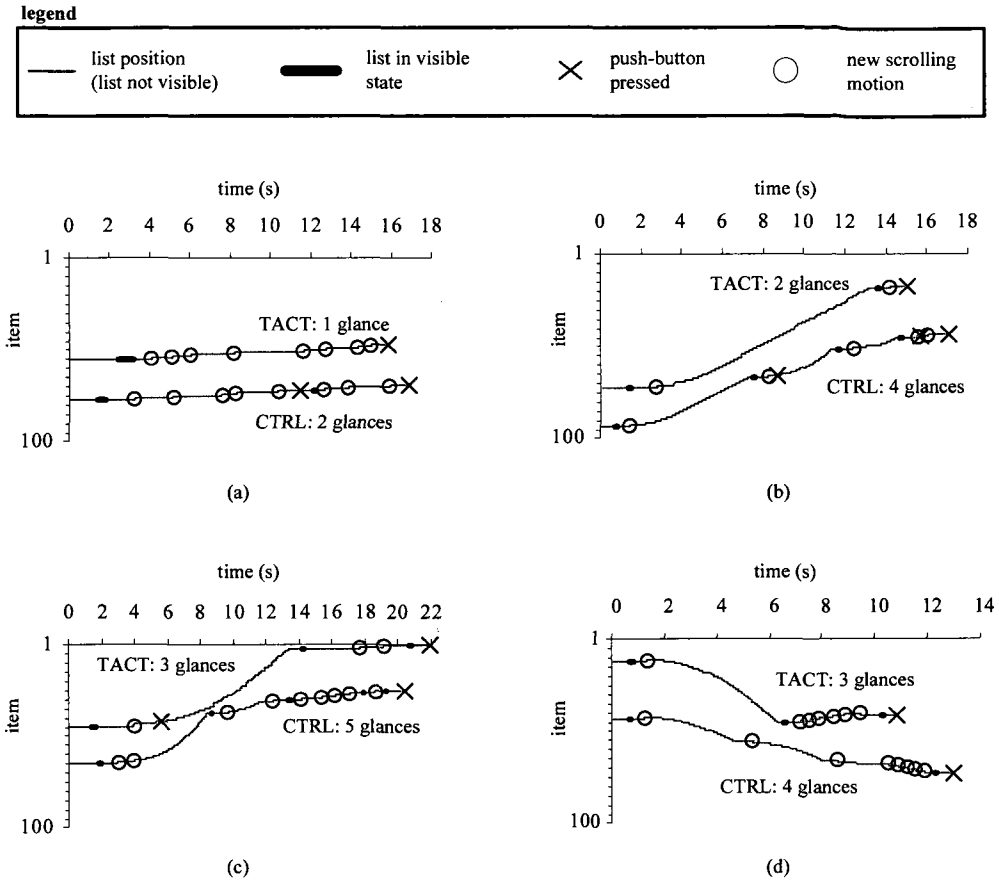

Fig. 6.11 Sample trajectory profiles illustrating examples of where the number of screen glances are reduced. The TACT condition is always shown on top of the CTRL which represents true initial and target positions. (a) Discrete mode for short distance (Subject C). (b) Continuous mode for long distance (Subject D). (c) Mixed mode for long distance (Subject J). (d) Overshoot-plus-recovery trajectory (Subject G).

scrolling speed and ends up passed the target. The error is promptly rectified with a switch to the discrete scrolling mode with very little consequences on the overall time needed to reach the target.

Figure 6.11(d) also reveals a small inconsistency with the controller that is repeated with each trial involving a downwards scrolling motion. Instead of scrolling

by one item towards the bottom of the list, the system initially scrolls up by a single item. This inconsistency does not seem to have effected the results obtained since all statistical results reported in Table 6.2 still hold when the cases of upwards and downwards scrolling motions are analyzed separately (not shown). In fact, the results suggest that subjects were not explicitly keeping count of the number of items scrolled by, but rather integrated the flow of tactile stimulation to locate their position.

At the conclusion of the experiment, subjects were asked to briefly describe the scrolling strategies they adopted when tactile feedback was present. Many reported using the discrete mode afforded by the controller whenever they were close to the target. A common strategy that was described consisted of first trying to reach a position close to the target in continuous scrolling mode. Then subjects would quickly glance at the screen before a final discrete adjustment in discrete mode. Most subjects also mentioned that the tactile icons positioned every 10 items and characterized by a higher frequency, were very useful in estimating current position.

6.6 Final Discussion and Conclusion

Taken together, the results from the two experiments highlight the possibility for a new role for tactile feedback in mobile interactions. They show that it is possible to make use of the tactile modality not only to assist, but to substitute for the visual modality in a task performed with a mobile device. This enhancement was enabled by a careful design of the tactile feedback provided by the device, not only from the view point of the tactile transducer design, but also from the view point of the feedback provided in response to input commands. Perhaps, the most important aspect of the information provided in the design used in Experiment II was how it enabled subjects to *anticipate* the sensory consequences of their motor output.

This principle is represented in the differences between the implementations of the tactile feedback for the two experiments. The scrolling interaction that we developed for Experiment I provided tactile feedback that did not couple input and output tightly. Instead, tactile signals were provided as a mere informational confirmation

to the interaction. The tactile cues were rich, pleasant and had information content regarding the direction and speed of movement, yet, they failed to provide feedback that made it easy to anticipate an event because they were too sparse in time and in space. This prevented subjects from taking advantage of anticipatory control over the scrolling motion.

The results of the first experiment illustrate a second point. Conventional metrics alone (error-rates or time measurements) are unlikely to uncover the full benefits of artificial tactile feedback. These performance measures are applicable to the visual modality which is anticipatory in nature, whereas tactile feedback employed in the conventional manner can only alert the user of events that have already occurred. In fact, tactile feedback may even be detrimental to the main task when the level of attention it demands outweighs the amount of useful information it provides. In Experiment I, the tactile feedback was probably perceived as an alarm that channeled subjects' attentional resources on the device rather than on the distractors. As a result, we observed cases where subjects had fewer scrolling overshoots but performed significantly worse at detecting the distractors. In a mobile context, this result can hardly be seen as a benefit, since the goal is to free up the user's attention.

To address these issues, in Experiment II, we proposed a control metaphor that differed significantly from that of Experiment I. We focused on a controller where action and perception are tightly coupled so that subjects could anticipate the results of their input sufficiently accurately to promptly re-adjust the scrolling trajectory.

Another contribution that is pertinent to mobile environment research is the testing paradigm we introduced. It enables the estimation of the amount of visual allocation by simply requiring them to have manual control over it. This method is practical and it proved to be efficient at simulating a mobile environment.

The results of the evaluation of the motion-controller used in Experiment II show that the tactile feedback could yield an average reduction of 28% in the number of glances subjects felt they needed to look at the screen. All ten subjects were affected positively by the tactile feedback. We consider these results to be very encouraging since they suggest new approaches to employ tactile feedback in mobile interactions.

6.7 Acknowledgements

This work was supported by the Natural Sciences and Engineering Council of Canada, in the form of a Collaborative Research and Development Grant. The authors would like to thank the industrial partners, the Nokia company and Immersion Corporation, for their support. They are also grateful to Prasun Lala and Vincent Levesque for invaluable comments.

Chapter 7

Summary and Future Work

Thesis Summary

This thesis has presented contributions to the design of fingertip tactile displays and to their use in mobile interfaces. Motivation for this work came from a need to compensate for some of the shortcomings of mobile devices for which typical desktop metaphor solutions are not applicable.

Early in this thesis, the use of tactile feedback was recognized as offering a potential solution to some of the challenges of mobile interaction. This is not a new idea. The sense of touch has long been identified as an underutilized channel for communication in HCI, especially when compared to audition and vision. However, examples of commercial devices that are capable of artificial haptics are not common. By contrast, touch plays an important role in our everyday-life activities. The slow deployment of tactile technology in the consumer world can be explained by three main factors: a limitation of the technology, our poor knowledge about touch and a lack of a clear vision about the true benefits of artificial tactile feedback.

Limitations of the Technology

Technological constraints continue to limit the development of electromechanical means that are capable of stimulating the skin. The requirements imposed by con-

sumer technology on tactile displays share many similarities with the ones imposed on their visual counterparts. Practical displays, whether tactile or visual, require that large arrays of low-power mechanisms be packed in a small space. In the case of visual displays, the essential transducer consists of a means to transform electrical energy into light. For their part, tactile displays make use of miniature moving parts in direct contact with the skin and the conventional actuation methods are fragile. Designing for small and robust actuators that must be capable of significant output forces against the skin remains a challenge.

Poor Knowledge About Touch

Our knowledge on touch is limited. The haptic sense has been studied for centuries, and more recently, by pioneers like Geldard [51]. However, a broad and distributed effort to study how the skin surface can be used as a successful communication device has only started to emerge in the last few years [48]. Our poor knowledge concerning the different mechanisms that relate to touch is well illustrated by a long list of unknowns. One example is our lack of certitude about the exact roles of the different mechanoreceptors embedded in the skin. Uncertainties about touch are not limited to the periphery. The different cognitive processes that take place in the higher levels of processing are also still open for debate. New tactile illusions are being uncovered on a regular basis, and occasionally an explanation for the phenomenon is offered. Nevertheless, to be truly practical as a communication channel, the haptic sense must also be understood when it is used in conjunction with other senses. Because there exists multiple ways by which different modalities can interact together (in a constructive or a destructive way), multimodal interactions are more complex in nature. The mechanisms by which touch competes with other modalities for the limited central cognitive resources have not been studied extensively. A better understanding of these mechanisms is crucial to the development of any useful tactile language because tactile communication will likely never replace visual and audio communication in consumer electronics. Rather, it will play a role that is matched to its unique capabilities.

Lack of Clear Benefits

Finding the advantages and means of leveraging the touch channel in human-computer interactions is not a trivial challenge. Visual communication occupies the vast majority of our interactions with computers. The visual system affords a broadband communication channel where large amounts of information can be displayed to users in a natural way. The auditory channel is also a natural candidate, but to a lesser extent. On the other hand, touch is, for the most part, a personal medium that is limited in its capabilities to transfer abstract information. For these reasons, it remains difficult to articulate how and to which extent haptic feedback can add value to consumer technology, even though it is ever-present in our everyday lives. Artificial tactile feedback has proven to be useful in certain contexts, as illustrated by the now omnipresence of vibrators in mobile devices. However, more often than otherwise, it is perceived as being a "gadget" that rarely succeeds at providing sustained value. This unfortunate reality can be explained by a lack of vision about the exact role of haptic technology.

The work presented in this thesis addressed some of the issues covered above. First, the development of the VBD device, described in Chapter 3, illustrated how making use of tactile illusions could simplify the manufacturing of tactile displays. This followed an approach that consists of taking advantage of tactile phenomena to compensate for limitations of the technology. In the case of the VBD, the illusion is one that results in a sensation best described as a 3D tactile feature that is moving on the user's skin. The phenomenon, which takes place without any significant mechanical normal indentation of the skin, was reproduced electromechanically with an array of piezoelectric benders in direct contact with the skin. The VBD was originally developed as a proof-of-concept prototype to explore the feasibility of displaying Braille dots. However, the transducer's expressive capabilities and its potential for miniaturization rapidly turned it into a natural candidate for exploring different tactile feedback solutions to issues with mobile interfaces [95].

Chapter 3 described how the transducer was integrated in a handheld mobile device prototype. Embedding the transducer in a small case that fits in the palm of

the hand helped picture what a real mobile device capable of rich tactile feedback would feel like. Users who tried the device, including participants in the experiments, would naturally visualize how a practical version would work. The miniaturization of the transducer had two functions. First, the engineering that it required moved one step closer to building a realistic tactile display for mobile interaction. Second, the small form factor of the resulting handheld device unlocked an entire space of possibilities for the envisioning of multimodal mobile applications that take advantage of tactile feedback. Because the device affordances suggested its use in a mobile context, the user experience was rendered more realistic. The move from an unsophisticated device (VBD) to a tangible prototype allowed the users to formulate a shared and descriptive language around the tactile sensations that could be communicated. Most users would "get it right away" and would often genuinely launch, without being prompted, into detailed descriptions of what they felt and how they saw it as being useful. In this context, it became possible to design experiments with the THMB that were more than psychophysical studies that measure user thresholds of tactile perception.

Specific applications with a potential to enhance mobile interaction were developed and tested. To make this possible, it was first necessary to assess the display capabilities of the THMB from a perceptual point of view. This step, which was achieved through a set of perceptual characterization studies first described in [95], raised questions about one of the techniques used to identify the underlying perceptual parameters of the THMB tactile stimuli. Cluster-sorted MDS, a method adapted by MacLean and Enriquez from Ward to efficiently collect and analyze data, had never been formerly validated [97, 174]. More specifically, the effects on the results of a MDS analysis on data that was collected with the cluster-sorted approach had never been examined. Chapter 5 described an empirical examination of the technique, illustrated with sample stimulus data directly taken from the THMB device. The analysis suggests that the dissimilarity matrix obtained from cluster-sorted methodology conserves important perceptual properties of the stimuli space. This is good news as it provides partial evidence that cluster-sorted MDS is a valid and efficient method to quickly identify groups and sub-groups of stimuli that can easily be distinguished

by users.

The last part of this thesis, Chapter 6, reported on the design and evaluation of a scrolling interaction with the THMB for which tactile feedback plays a central role. The goal was to design for a mobile interaction with a measurable improvement that could be attributed to the introduction of tactile feedback. The approach taken consisted of associating sensory functions to motor functions to achieve a holistic experience when scrolling a long list of items. A look at a measure of participants' reliance on vision, expressed by how many times they pressed a key that controlled the visibility state of the list being scrolled, suggested that tactile feedback could diminish the number of required glances at the screen of a mobile device.

Future Directions

The scope of the work described in this thesis is broad. Opportunities for improvement can take one or many of several potential directions. First and foremost, the design of laterotactile displays would benefit from a better understanding of the laterotactile illusion. This work involves a study of what really happens at the mechanoreceptors level when the skin is mechanically excited. In turn, this can inform the design of laterotactile displays by suggesting parameters for their fabrication, such as an optimal distance between the bending actuators in terms of ease of manufacturing and realism of the resulting sensation. To this extent, Wang and Hayward [171], as well as Kikuuwe et al. [82], have recently suggested promising but partial solutions based on contact mechanics and mathematical modeling.

Another interesting avenue for inquiry concerns the exploration of alternative means to generate the distributed deformation of the skin surface. Piezoelectric benders have proven to be quite effective at applying programmable patterns of skin deformation, but the study of other mechanisms, such as microelectromechanical system (MEMS), should also be pursued. Recent years have been marked by an explosion in the number of attempts at building tactile display prototypes. New promising actuation mechanisms have been built and different ways of exciting the skin are continuously revisited. These efforts will eventually pay off. With the

upcoming improvements in tactile display technology will emerge a need for more refined tactile rendering algorithms. Over time, the topic of tactile rendering is likely to approach a level of maturity that is comparable to that of current graphics rendering. Levesque and Hayward have already embarked in a methodical study of the different types of fundamental 2D tactile sensations that laterotactile stimulation can generate. One of the objectives of this research is to develop tactile maps that are formed of combinations of laterotactile primitives; for instance, borders on a map of the world could be represented tactually with patterns similar to the Braille dots of the VBD, while oceans could be symbolized by dynamic undulations [92].

This thesis has underscored the importance of making use of appropriate evaluation techniques to assess the perceptual capabilities of novel tactile displays. More work is needed on this front to reach a set of efficient and reliable tools that allow a complete mapping of the sensations that tactile displays can generate. The technique of cluster-sorted MDS is part of this toolset, but other methods must also be devised.

In this thesis, it was also argued that new evaluation methodologies must be developed in order to uncover the complete benefits of haptics. Applying typical HCI evaluation metrics that measure performance, such as task-completion time and error rates, are important but they risk overlooking the true advantages of using tactile feedback in consumer electronics. Future work should be concerned with the study and validation of evaluation tools that are better matched to the properties of touch. As illustrated in Chapter 6, this can be achieved by developing alternative metrics such as the number of times the screen of a mobile device is glanced at. Other methods should focus on revealing the affective qualities of tactile feedback. This, in turn, will open the door to new application specific scenarios. Future work promises to be an exciting combination of fundamental perceptual science mixed with applied engineering and careful HCI design.

Appendices

Appendix A

State-of-the-Art Tactile Displays

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
Electrostatic	Capacitor with Polyimide (PI) insulator	Capacitor formed of the conducting fluids in the fingertip acting as the first plate and external electrodes acting as the other plate. A voltage induced across the capacitor creates attraction between the skin surface and the external electrode surface.	49	Friction	- Reproduction of shear forces at the surface of the skin - Active touch device (sliding of finger)	- Large voltage (200-600 V) - Complex fabrication process of the PI layer - Fixed operating freq. (100 Hz) - Sensitive to humidity of skin	[153]
	Capacitor with polymeric elastic dielectric	Stimulator tip mounted on a stack of capacitors with polymeric elastic dielectrics. When a control voltage is applied across the capacitors, the dielectric material contracts and the position of the stimulator tip is set.	No proto.	Normal indentation	- Low cost, lightweight and flexible material - Potential for large strain (up to a few mm)	- Large operating voltage (100-1000 V) - Little current knowledge of the material properties and manufacturing process	[77]
Rheological	Electro-rheological (ER) fluid	ER fluid cell resisting the motion of the fingertip. The ER fluid changes from a liquid state to a solid state when exposed to an electric field. Altering the ER fluids state induces horizontal and vertical reactive forces during finger scanning.	25	Resistance to finger motion	- Low energy consumption - Simple mechanical design - Active touch	- Problems, such as liquid accumulation, related to the use of an ER fluid - Tradeoff between the resolution of the array and the force of the response (due to the hazard of having large control voltages close to each other)	[156]

continued on next page

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
	Magneto-rheological (MR) fluid	MR fluid placed in a Plexiglas box surrounded by solenoids. Inducing a current in a solenoid creates a magnetic field that changes the fluid to a near-solid in the vicinity of the solenoid.	16	Shape Softness	- Active exploration from the user - Both a kinesthetic device and a tactile device	- Low actuator spatial resolution - Need to wear latex glove - Large power dissipation (overheating)	[8]
Electromechanical	Piezoelectric	Bending bimorph carrying an L-shaped wire acting as the skin contactor.	100	Vibration	- Large working bandwidth (20-400 Hz) - Large spatial resolution of actuators ($1/mm^2$)	- Large control voltage (85 V) - Complexity of manufacturing	[147]
		Piezoelectric ceramic plates are assembled next to each other in a staggered pattern to form a 1D array of contactors.	88	Normal indentation	- Large bandwidth (0-1000 Hz) - Simple design - Controllable actuator amplitude	- Limited to displaying tactile signals in a single dimension - Mechanical coupling between the actuator plates - Weak maximum displacements of the actuators (11 μm)	[161]
		Mechanically amplified piezoelectric actuator driving a vibratory pin.	50	Vibration	- Controllable actuator amplitude (5-57 μm) - Simple circuitry	- Fixed operating frequency (250 Hz) - Limited to simple sensations of vibration	[62]
		Vertical movement of a contactor induced by a pair of piezoelectric levers.	48 (4000 virtual)	Normal indentation	- Vertical movement of up to 0.7 mm - The device is mounted on a sliding apparatus that permits the exploration of a large surface area without the need for an extensive number of actuators	- Small spatial resolution - Large control voltage (200 V) coming out of a power supply card - Low bandwidth (20 Hz)	[100]
		Skin contactors glued on a membrane that is deformed by a matrix of piezoelectric actuators.	64 (112 contactors)	Laterotactile	- Large spatial resolution of contactors - Portability of device (e.g., can be put on a computer mouse) - New mode of interaction (lateral stretch)	- Limited actuator displacement and force - Large control voltage ($\pm$200 V) - Indirect control of the positions of the contactors	[57]
		Array of piezoelectric benders that are vertically mounted.	64	Vibration (tangential to skin)	- Early example of tangential mechanical deformation of the skin - Designed to fit the palm of the hand - Capable of 63 levels of intensity	- Fixed output frequency (250 Hz) - Large spatial actuator resolution (8 mm centers)	[31]
		Similar mechanism to Braille cells. Horizontal bimorphs push up against contact pins when they bend.	30	Normal indentation	- Large bandwidth - Displacement of 0.7 mm - Blocking force of 0.6 N - Light and portable	- Large control voltage (150 V) - Moving parts	[87]
		Structural configuration made of 6 piezo-extenders. The vertical displacements of all the individual actuators add up.	4	Normal indentation	- High bandwidth - Use of extenders rather than benders - Displacement comparable to benders (90 μm)	- High control voltage (100 V) - Complex structure - Low spatial resolution	[107]

continued on next page

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
		The friction coefficient of a glass surface is modulated by making it vibrate at ultrasonic frequencies by a SINGLE piezoelectric actuator.	1	Friction	- Single actuator - Capable of a set of interesting and diverse texture sensations	- Limited surface - Requires active finger exploration	[177]
		The friction coefficient of a touch surface is modulated by making it vibrate at ultrasonic frequencies by an ARRAY of piezo-benders. Similar to [177] above.	> 30	Friction	- Can be extended to a large exploratory surface - No moving skin contactor pins - Able to induce sensation of texture - Low-voltage for piezo-actuators (15 V)	- Requires finger movement - Limited in expressive capabilities	[10]
	Motor	RC servomotor slightly rotating a lever arm on which a skin contactor is fixed. The small rotation of the lever arm results in a vertical motion of the contactor. A sheet of rubber covers all the contactors to create a spatial low pass filter.	36	Normal indentation	- Large vertical displacement (up to 2 mm) - Poor actuators spatial resolution (2 mm) compensated by a rubber sheet acting as a spatial low-pass filter - A mouse attached to the display permits active exploration	- Complex control system - Fairly big and cumbersome device - Low bandwidth (25 Hz)	[169]
		2 DOF mechanism consisting of servomotors that pull and push laterally on metal pins in contact with the fingertip skin.	4	Lateral stretch	- The pins/contactors have 2 DOF - Considerable displacement and force exhibited by the pins/contactors	- Fairly complex and large mechanical structure - Limited to 4 actuators	[38]
		Miniature DC motors that remotely pull on spring-loaded pins through a transmission pulley system made of nylon tendons. The tactile interaction occurs at the fingertip, but the actuators are located on the user's wrist.	16	Normal indentation	- Portable fingertip display - Low mass - Large displacement (25 mm)	- Low bandwidth (in the tens of Hz) - Friction in the transmission system	[136]
		The rotation of a step motor is transformed into vertical movement of a skin contactor by using a lead-screw mechanism.	4096	Normal indentation	- Large contactor force - Large contactor displacement (up to a few mm) - Considerable surface area (200 mm x 170 mm) and large number of actuators	- Very slow refresh rate (15 s) - Limited to display shape (i.e., no texture) because of high contactor spatial resolution (3 mm) - Complex and expensive control system	[139]
		Mechanism that converts the rotating motion of a miniature brushless motor into the vertical movement of a tactile element (tactel).	32	Normal indentation	- Large force exhibited (1.7 N) - Large tactel displacement (3 mm) - Compact	- Limited bandwidth - Large spatial resolution of tactels (2.7 mm)	[108]

continued on next page

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
		Stepping motor. Screw mechanism similar to [108] and [139] above.	256	Normal indentation	- Large pin displacement (6 mm) - Skin contactors also act as switches for input	- Large pins (diameter 5 mm) - Very large spatial resolution (10 mm) - Bulky and complex system	[81]
	SMA	A shape memory alloy (SMA) wire pulls a lever that lifts a skin contactor. The contactor indents the fingertip skin.	24	Normal indentation	- Large vertical extension of the contactors - Large contact force	- Hysteretic behavior of the SMA material - Low control bandwidth (10 Hz) - Large power dissipation	[84]
		Vertical pin fixed to the middle of a SMA wire like an arrow is mounted on the wire of a bow. Controlling the length of the SMA wire with an electric current moves the pin up and down. A latex rubber membrane acting as a seal is laid on top of the pins.	Line of 10	Normal indentation	- Interesting bandwidth for a SMA device (30 Hz) - Low strain of the SMA material amplified by an ingenious mechanical arrangement	- Hysteretic properties of the SMA material - Complex cooling system - Uses a line of skin contactors instead of a matrix distribution (as a consequence, the line edges are felt)	[175]
		SMA NiTi wire attached to a sprung pin in contact with the skin. An electric current induced in the SMA makes it contract and pulls the pin down.	64	Normal indentation	- Fairly large controllable strains of the SMA wires (up to 5% or 5 mm)	- Low operating frequency (1-3 Hz) - Large heat generation	[157]
		A SMA spring and linear counterspring are wrapped around a ceramic pin held in mechanical equilibrium. Changing the temperature in the SMA moves the pin up and down.	64	Normal indentation	- Large pin displacement - Simple and low-cost design	- Requires cooling and pins cannot be held in position for long sustained lapses of time - Thermal interaction between the actuators	[42]
		Very similar to design [42] above but with another SMA as the 2nd spring.	64	Normal indentation	- Low cost - Compact and light - Simple control electronics	- Binary state (up or down) - Low spatial resolution (2.6 mm) - Very low bandwidth for tactile graphics (1.5 Hz)	[165]
		Heat pulses through a membrane control its curvature state by taking advantage of the difference between the thermal expansion coefficients of the two SMA materials that compose it. When deformed, the membrane pushes on a skin contactor.	*not known*	Normal indentation	- Bistable mechanism that does not require energy consumption between state changes - Thin film technology allows mass manufacturing at low cost	*Not enough information available*	[178]

continued on next page

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
	Coil	Two fixed coils and a moving magnet suspended by two helical springs act as a motor controlling the displacement of a long stainless steel probe.	400	Normal indentation	- Large contact force (up to a few Newtons) - Large displacement of the actuators (up to 25 mm) - Good actuator resolution	- Very large control system - Very complex and expensive device	[119]
Pneumatic	Pneumatic valve & dimple	Array of pressurized silicone tubing. By changing the pressure in the chambers, the displacement of vertical contactors in the tubes is controlled.	25	Normal indentation	- Constant contact with the finger - No leakage and no pin friction - Controllable pin displacement (up to 0.7 mm) - Highly portable	- Very low bandwidth (5 Hz) - Low spatial resolution (actuators are 25 mm apart) - Undesired operating vibration resulting from the PWM control signal	[103]
		Pneumatic inflow controlling the pressure and vibration of stainless steel pins. Pneumatic muscle generating lateral forces to simulate friction.	16	Normal indentation Vibration Shear	- Compact and integrated package capable of three different types of stimulations - Large normal force (2 N) and displacement (35 mm) - Large vibratory bandwidth (20-300 Hz)	- Complex system - Few contact pins with fairly high spacing separation (175 mm)	[23]
		The pressure inside a small cavity is increased by heating up a gas near the boiling point. The change in pressure displaces a tactile element.	16	Normal indentation	Offers potential for: - simple miniature actuation - significant displacement and force	In current state: - low spatial resolution and bandwidth - high power	[168]
Electrocutaneous	Electrostimulation	Visual images are captured by an optical sensor mounted on the display before being translated into electrical tactile stimulation on the fingertip.	16	Electric current felt as: 1- vague pressure 2- acute vibration	- Can generate 2 distinct sensations (vague pressure and acute vibration) - Mounted on a force sensor to regulate the sensation magnitude and to decrease discomfort - Sensor directly mounted on tactile display	- Given the technology, fairly low spatial resolution (actuators are at least 2 mm apart)	[79]
		Active electrode becomes electrically connected to ground through the fingertip when the user touches it. The current passing through the finger creates a tactile sensation of vibration and pressure.	49	Electric current	- Simple method - Possibility of large tactile resolution - Flexibility (e.g., can be put in a glove)	- Can cause pain - Adaptation to the stimulus occurs very quickly	[78]
Thermal	Peltier Element	Peltier Element	1	Heat	- Very simple - Capable of simulating real sensations of the quality of materials under passive touch	- Only one single actuator - Incapable of presenting dynamic information such as pressure or strain	[64]

continued on next page

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
Air Jet	Piston	Air jet produced by controlling the pressure through a tube with a piston.	1	Normal indentation (by air jet)	- Stimulation of superficial receptors creating a sensation not reproducible with other TDs ("bug creeping under the skin") - No direct mechanical contact with the skin	- Impossibility to get a high-resolution array due to the size of the jet actuators	[1]
Hybrid	Piezo + Peltier element	Cantilevered piezo bimorphs pushing on vertical pins + Peltier element	1 Peltier + 30 pins	Heat and normal indentation	- Displays heat on thumb - Compact and mature technology - All the advantages of piezo technology - Multifinger display	- Requires water cooling system - Limited bandwidth of Peltier element	[180]
	Braille cells + linear motor	Custom adaptor that converts 3 off-the-shelf refreshable Braille cells into a higher spatial resolution display. The display is mounted on a tangential motor.	1 linear motor + 24 pins	Normal indentation and slippage	- Uses off-the-shelf components - High spatial resolution (2mm)	- All the same problems associated with Braille cells (high-voltage, low force, etc.)	[185]
	Coil + Peltier element	Small electromagnetic actuators with microcoils actuate flexible membranes at a specific frequency (also Peltier elements)	64	Vibration Heat	- Low cost fabrication technology - Relatively high density matrix (2 mm interspace) - High temporal resolution - Coupling between thermal feedback and vibrotactile interactions	- Low static force - Limited to vibrational and thermal interactions (i.e., no direct stimulation of slowly adapting skin mechanoreceptors)	[4]
Others	Pressure Valve	Drawing air from a suction hole contacting with the palm creates the illusion that the skin is pushed by a "muddler".	20	Suction pressure	- No interference between neighboring stimulators - Two kinds of basic patterns of stimulation (large holes and small holes)	- Very low spatial resolution (only appropriate for the palm of the hand) - Need for regulation of air pressure	[99]
	Surface Acoustic Wave	Burst of surface acoustic waves (SAWS) are used to modulate the amount of surface friction applied to a slider on which the user's finger rests. This allows the control of the shear stress applied on the finger's skin by the slider while moving. The SAWS are created by interdigital transducers.	n/a	Shear stress	- Original and unexplored method	- Not a direct-contact tactile display	[104]

continued on next page

Type	Actuator	Mechanism	Nb. Act.	Interaction	Interesting Properties	Drawbacks	Ref.
	PZT transducer (producing ultrasound)	Elastic gel is covered with an ultrasound reflector and is radiated with ultrasound. The net effect is one of induced pressure on the fingertip lying on the reflector.	10 and 30	Acoustic radiation pressure Vibration	- High spatial resolution (1 mm) - High refresh rate - Free from contact problems	- Bulky system - Weak continuous pressure force	[66]
	Ionic Conducting Polymer gel Film (ICPF)	ICPF cilium-shaped actuator submerged in water. Applying an electric field between the surfaces of the actuator makes it bend.	10	Vibration (at high freq.) Shear (at low freq.) Brushing	- The softness of the ICPF material allows for very delicate pressure - Low driving voltage (under 15 V) - Fairly high frequency operation (up to more than 100 Hz)	- ICPF actuators require to be submerged in water in order to bend - Low actuator resolution	[85]
	Dielectric elastomer	Braille cell composed of multiple levels of a dielectric elastomer actuator that bends when submitted to a voltage. The bending motion pushes on a Braille dot.	6 Braille dots	Normal indentation	- Offers potential to substitute for piezoelectric Braille cells to be manufactured in batch at low cost	- Huge control voltage (in kVolts) - Difficult control of the exact Braille dot height - Low spatial resolution not suitable for tactile graphics yet	[29]

Table A.1 State-of-the-art distributed tactile displays for the fingertip (updated from [114])

Appendix B

CAD Drawings for the THMB

This appendix holds the computer-aided design (CAD) drawings for the main parts of the 2^{nd} generation THMB device. The design and manufacturing of the THMB was the result of a close collaboration with Don Pavlasek, machine shop supervisor in the Electrical and Computer Engineering (ECE) department at McGill University.

All detailed drawings presented here were put together on Autodesk's AutoCAD® software application by Don Pavlasek. Most final drawings are based on conceptual drawings that I prepared prior to meeting with him. The final CAD drawings are the result of multiple design sessions where Don and I would discuss ideas and refine the design around a sketchpad until a satisfactory solution was reached. Manufacturing was conducted on a computer numerical control (CNC) machine in the ECE machine shop. I performed the cutting and preparation of the piezoelectric actuators in another machine shop that belongs to the McGill Haptics lab. I also took care of assembling the device.

I am ever grateful to Don for his sustained patience and for sharing his mechanical design and manufacturing expertise with me so generously.

SPRING LENGTH DETERMINATION, MAY 15, 2007
BASED ON 9.75 STACK HEIGHT AS MEASURED, MAY 2007

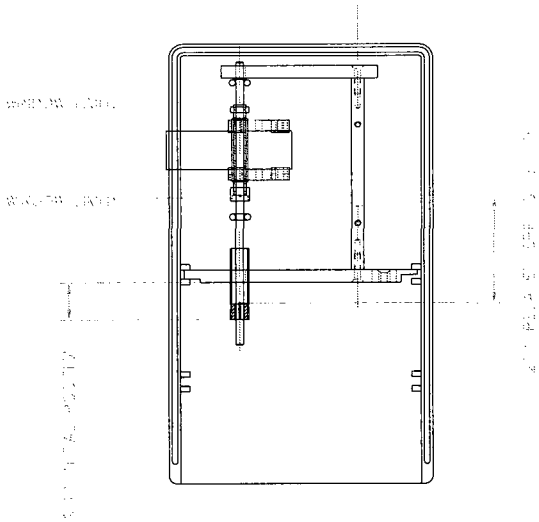

SHAFT AND BUMPER LENGTH DTERMINATION, OCT. 24, 2006

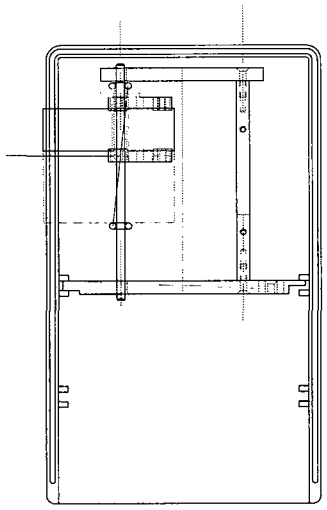

'TUBE 1'

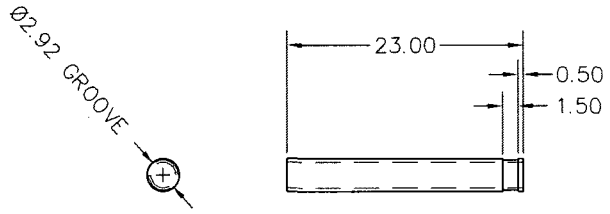

MAT'L: 3.22 OD x 2.43 ID BRASS TUBING

1 PC REQ'D

'TUBE 2'

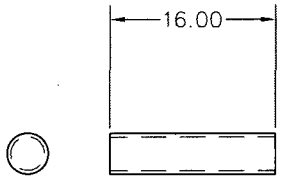

MAT'L: 3.97 OD x 3.27 ID BRASS TUBING

1 PC REQ'D

SPRING HOUSING CLIP

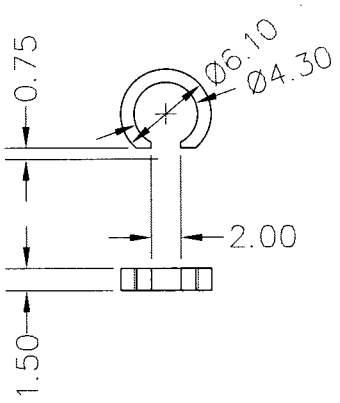

NZC SPRING HOUSING

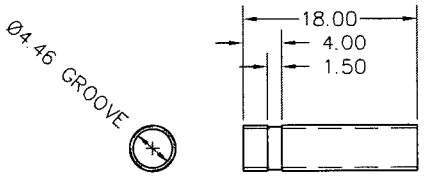

MAT'L: 4.76 OD x 4.05 ID BRASS TUBING

1 PC REQ'D

NZC SPRING HOUSING BUSHING

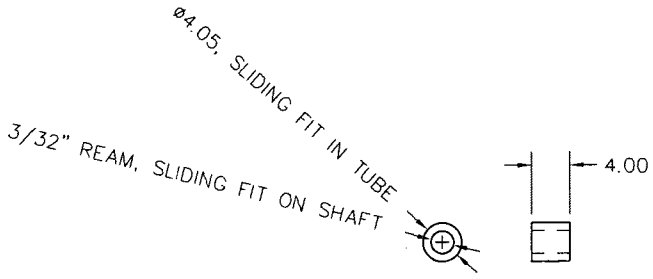

MAT'L: BRASS ROD

1 PC REQ'D

NZC SPRING HOUSING ASSEMBLY

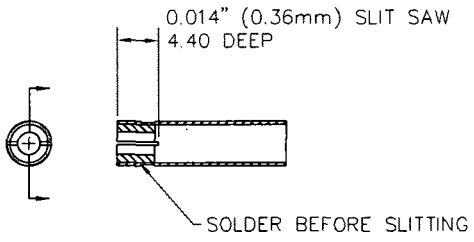

1 PC REQ'D

SPRING

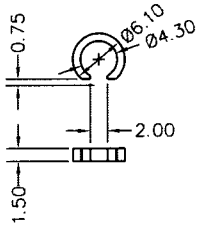

MAT'L: STEEL EXTENSION SPRING
0.014" WIRE, 0.142" OD. STRETCHED AS SHOWN

1 PC REQ'D

BLOCK END PLATES, MARCH 2007

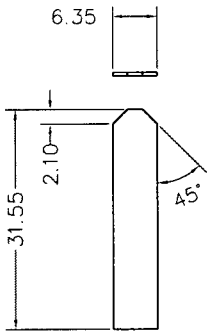

MAT'L: 0.020" (0.51mm) BRASS SHIM STOCK

2 PCS REQ'D

BLOCK END PLATE RELIEF DETERMINATION

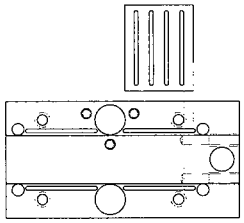

BLOCK END PLATES WITH RELIEF
APRIL 2007

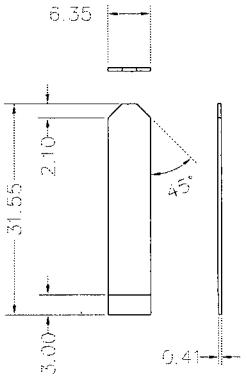

MAT'L: 0.020" (0.51mm) BRASS SHIM STOCK

2 PCS REQ'D

LEXAN 'DUAL PINNING' WIRE GUIDES

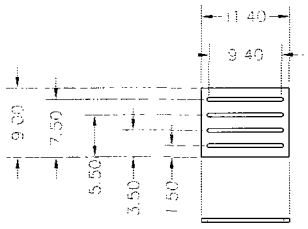

MAT'L: 0.020" (0.51mm) CLEAR LEXAN SHEET

4 PCS REQ'D

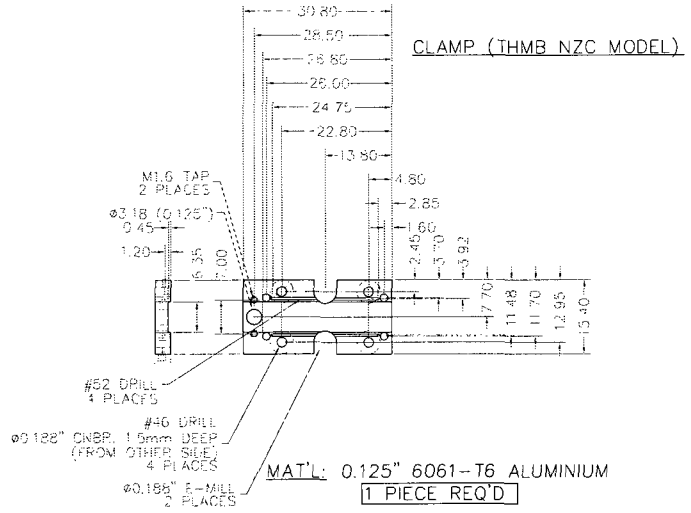

NOTE: M1.6 HOLES MOVED FROM 26.8mm TO 28.5, NOV. 2, '06

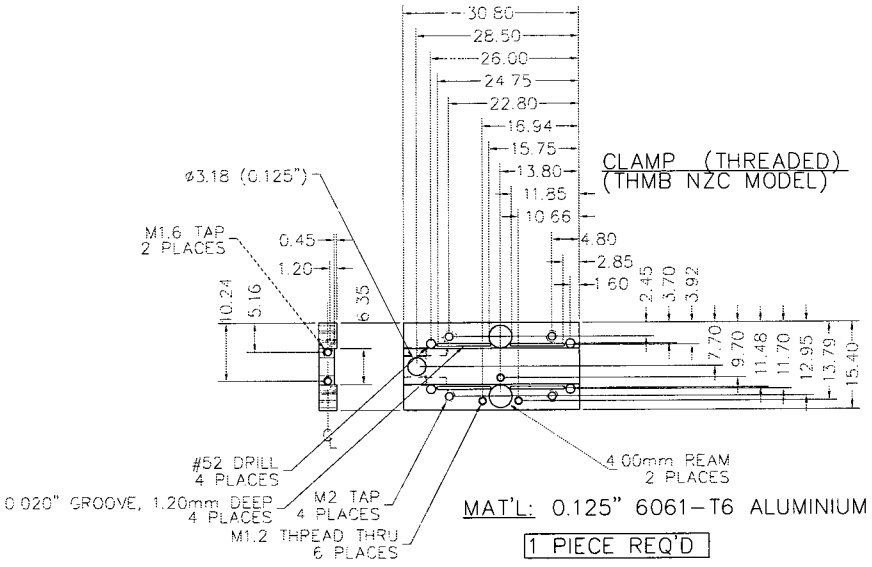
30.80
28.50
26.00
24.75
22.80
16.94
15.75
13.80
11.85
10.66
4.80
2.85
1.60
2.45
3.70
3.92
7.70
9.70
11.48
11.70
12.95
13.79
15.40
CLAMP (THREADED)
(THMB NZC MODEL)
ø3.18 (0.125")
M1.6 TAP
2 PLACES
0.45
1.20
10.24
5.16
6.35
#52 DRILL
4 PLACES
0.020" GROOVE, 1.20mm DEEP
4 PLACES
M2 TAP
4 PLACES
M1.2 THREAD THRU
6 PLACES
4.00mm REAM
2 PLACES
MAT'L: 0.125" 6061-T6 ALUMINIUM
1 PIECE REQ'D

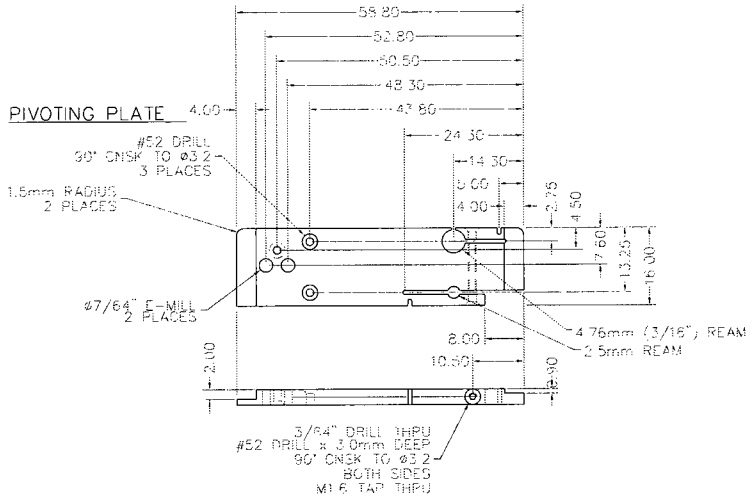
58.80
52.80
50.50
48.30
43.80
24.30
14.30
6.00
4.00
4.00
2.75
4.50
7.60
13.25
16.00
PIVOTING PLATE
#52 DRILL
90° CNSK TO ø3.2
3 PLACES
1.5mm RADIUS
2 PLACES
ø7/64" E-MILL
2 PLACES
4.76mm (3/16") REAM
2.5mm REAM
8.00
10.50
2.00
3.50
3/64" DRILL THRU
#52 DRILL x 3.0mm DEEP
90° CNSK TO ø3.2
BOTH SIDES
M1.6 TAP THRU
MAT'L: 0.125" 6061-T6 ALUMINIUM
1 PIECE REQ'D

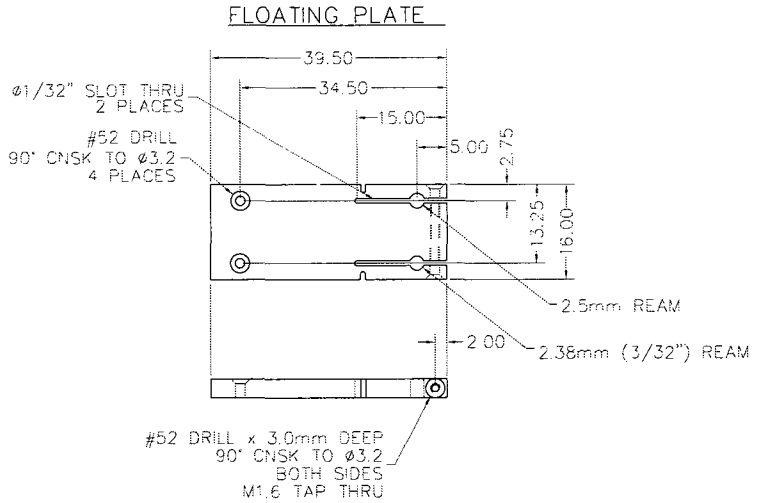

MAT'L: 0.125" 6061-T6 ALUMINIUM

1 PIECE REQ'D

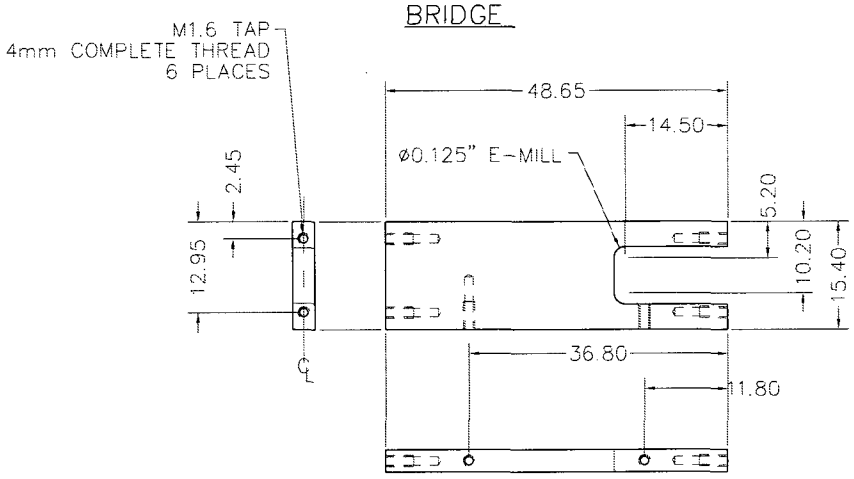

MAT'L: 0.125" 6061-T6 ALUMINIUM

1 PIECE REQ'D

BEARING SHAFT

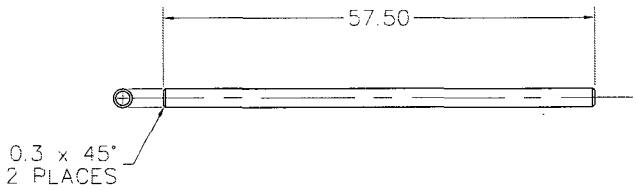

MAT'L: ø2.50mm 303 STAINLESS GROUND SHAFTING
1 PC REQ'D

NZC SHAFT
MAY 15, 2007 LENGTH REVISION

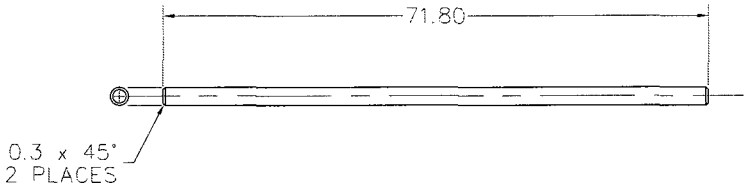

MAT'L: ø3/32" (ø2.38mm) 303 STAINLESS GROUND SHAFTING
1 PC REQ'D

GUIDE TUBE: TUBE BLANK

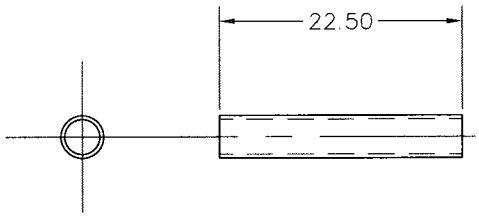

GUIDE TUBE: FLANGE BLANK

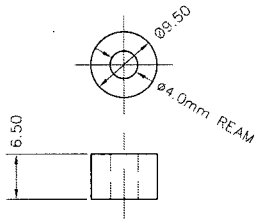

MAT'L: BRASS
NOTE: 0.1 x 45° CHAMFER, ALL CORNERS

1 PC REQ'D

GUIDE TUBE: SOLDERING

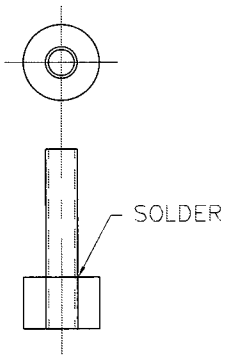

GUIDE TUBE: MACHINING

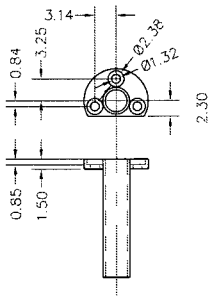

NOTE: 0.2 x 45° CHAMFER, ALL CORNERS

GUIDE TUBE: BUSHING

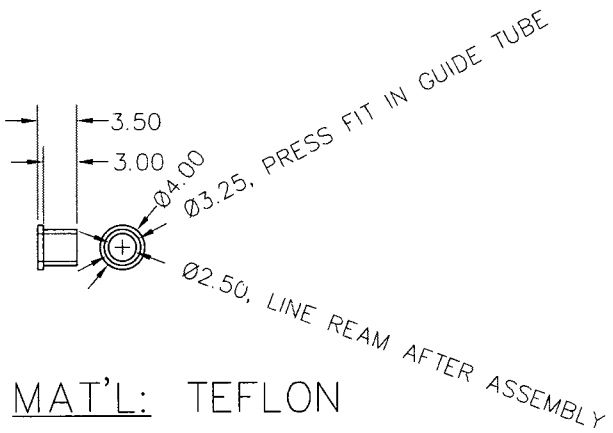

MAT'L: TEFLON

2 PCS REQ'D

GUIDE TUBE: BUSHING INSTALLATION

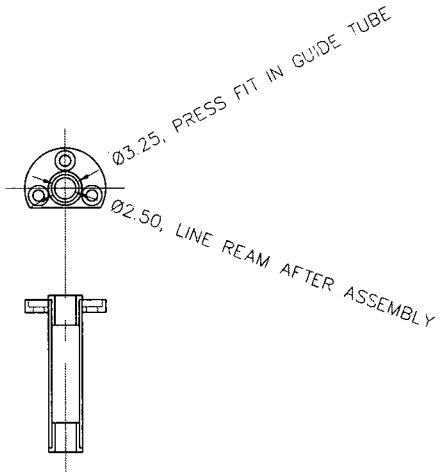

GUIDE TUBE: SCREWS

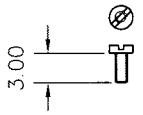

M1.2 X 3.0 CHEESE HEAD, ALLOY STEEL

3 PCS. REQ'D

STOP TUBE: NEW LENGTH (MAY 15, '07) AS MEASURED

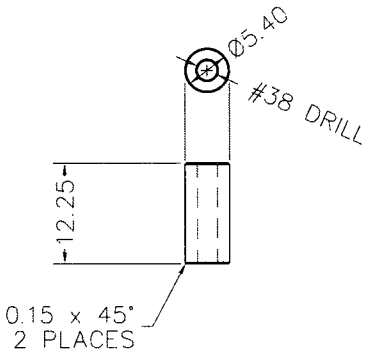

MAT'L: BLACK DELRIN

1 PC REQ'D

PIEZO BLOCK MODEL

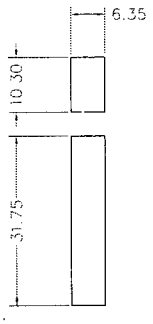

MAT'L: BLACK DELRIN

1 PC REQ'D

PIEZO BOTTOM STOP

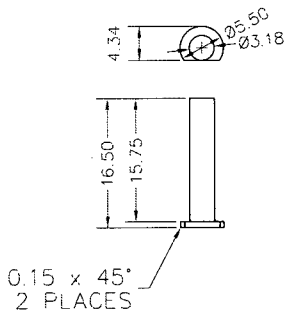

MAT'L: BLACK DELRIN

1 PC REQ'D

SWITCH BUMPER PLATE

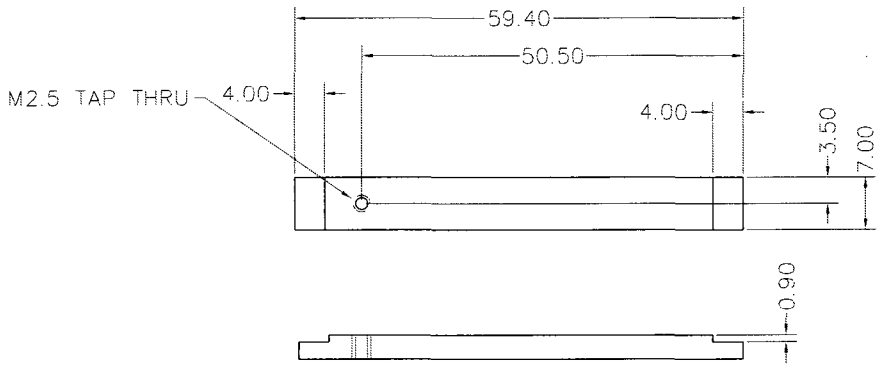

MAT'L: 0.125" 6061-T6 ALUMINIUM

1 PIECE REQ'D

SWITCH MOUNT

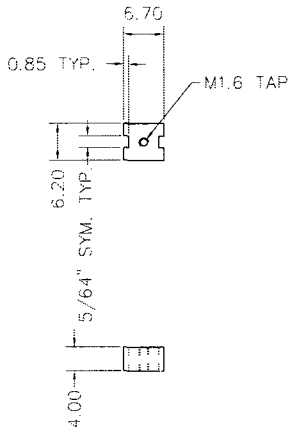

MAT'L: BLACK DELRIN

1 PC REQ'D

SWITCH BUMPER

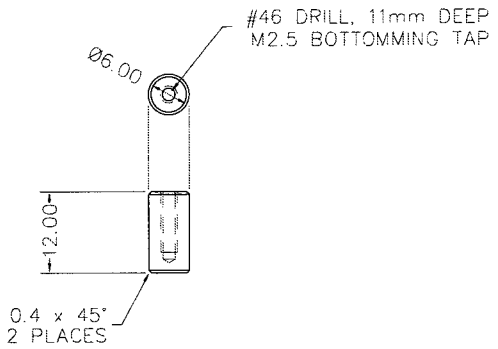

MAT'L: BLACK DELRIN

1 PC REQ'D

SLIDER WINDOW: NZC PRODUCTION
MAY 2007
NOTE: CASE DEPTH/GROOVES CORRECTED, MAY 29, 2007

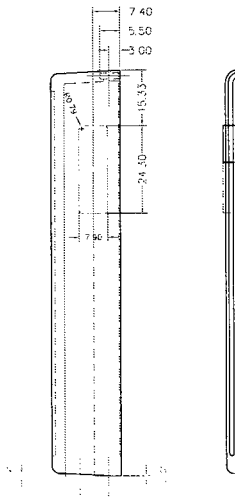

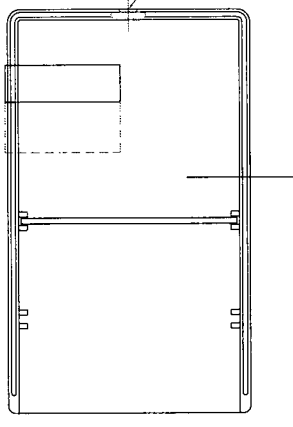

SLIDER CAVITY: NZC PRODUCTION
MAY 2007

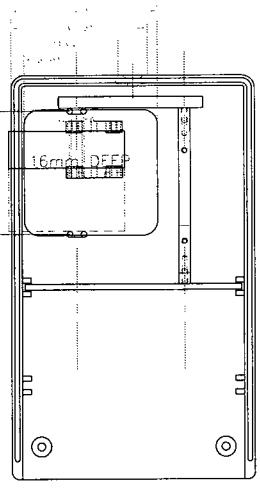

SWITCH BUMPER

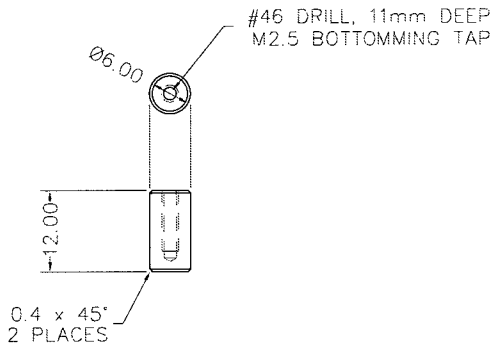

MAT'L: BLACK DELRIN

1 PC REQ'D

SLIDER WINDOW: NZC PRODUCTION
MAY 2007
NOTE: CASE DEPTH/GROOVES CORRECTED, MAY 29, 2007

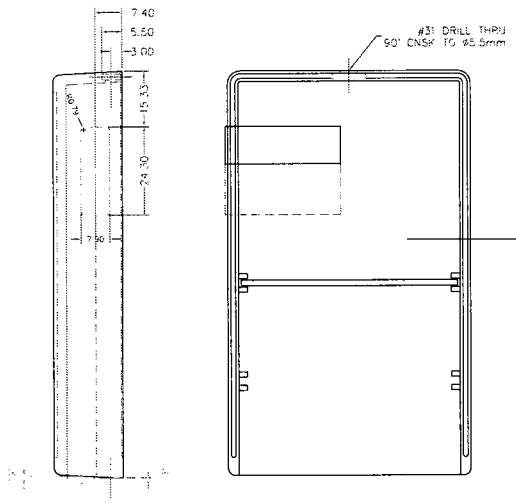

SLIDER CAVITY: NZC PRODUCTION
MAY 2007

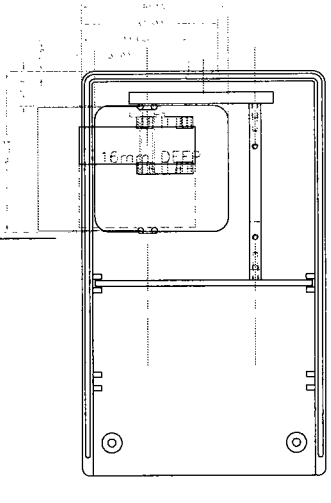

BACK/SCREEN ENCLOSURE ATTACHMENT TAB

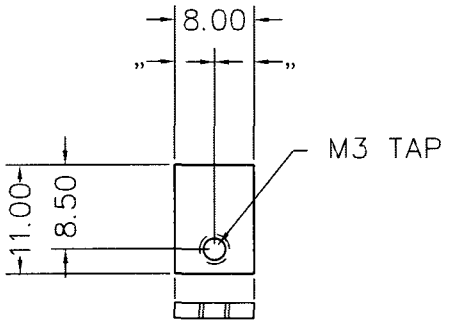

MAT'L: 1/16" 6061-T6 ALUMINIUM

1 PC REQ'D

SCREEN ENCLOSURE: CASE SECTION
MAY 2007
NOTE: CASE DEPTH/GROOVES CORRECTED, MAY 29, 2007

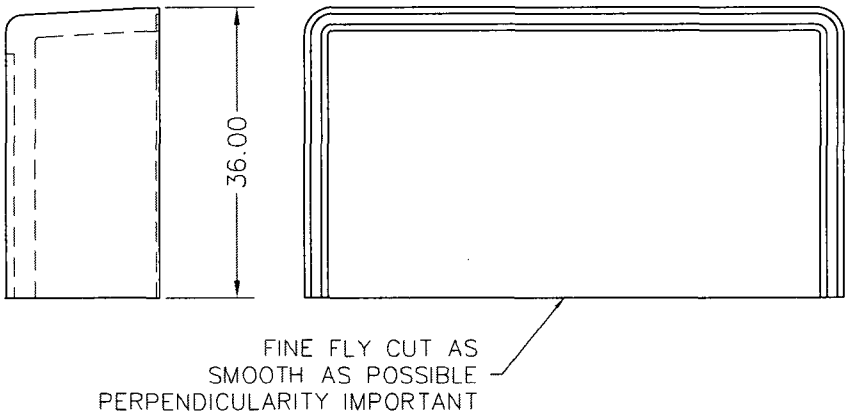

SCREEN ENCLOSURE
MAY 2007

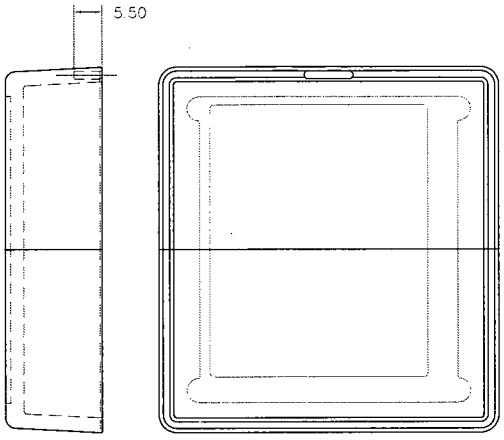

CONNECTOR CIRCUIT BOARD STAND-OFFS

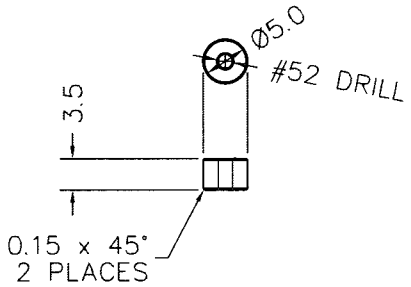

MAT'L: BLACK DELRIN

2 PCS REQ'D

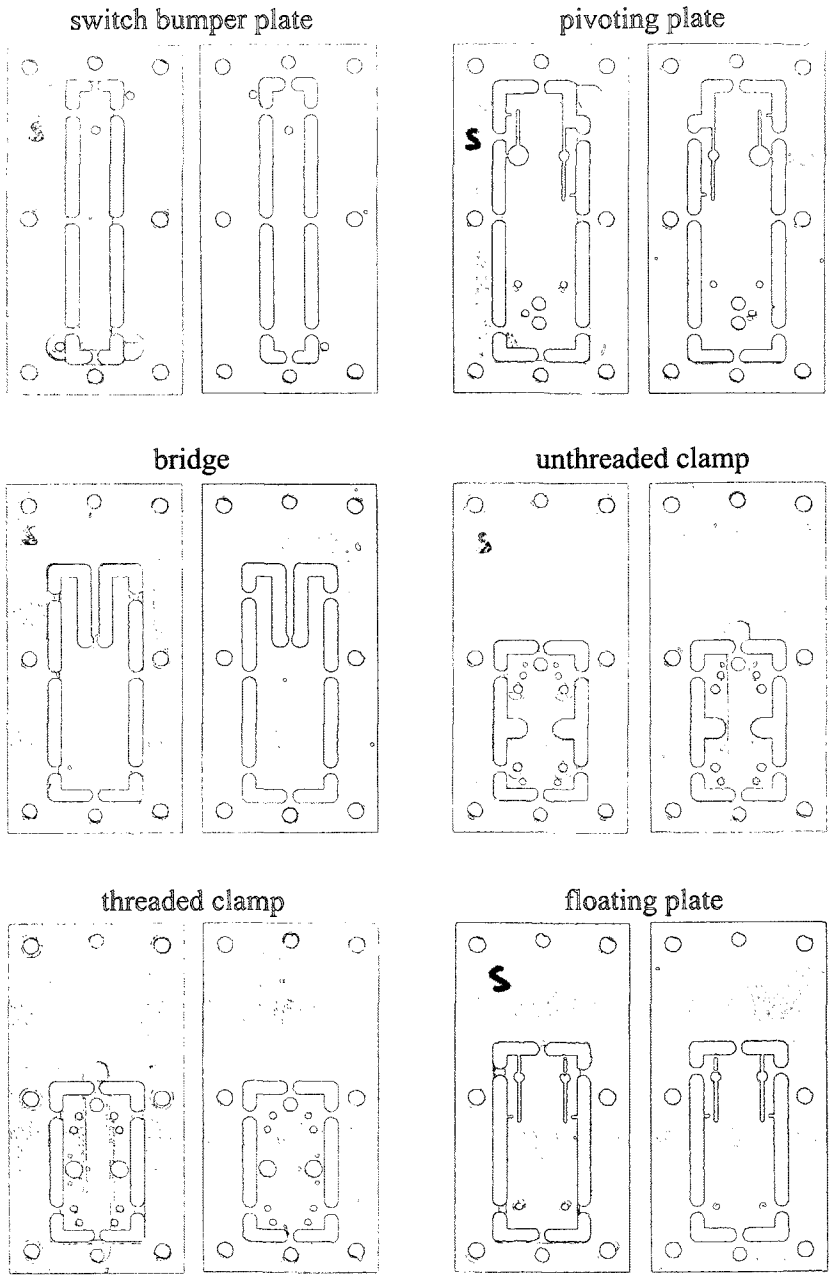

Fig. B.1 Main aluminum parts of the THMB device as they come out from the CNC machine.

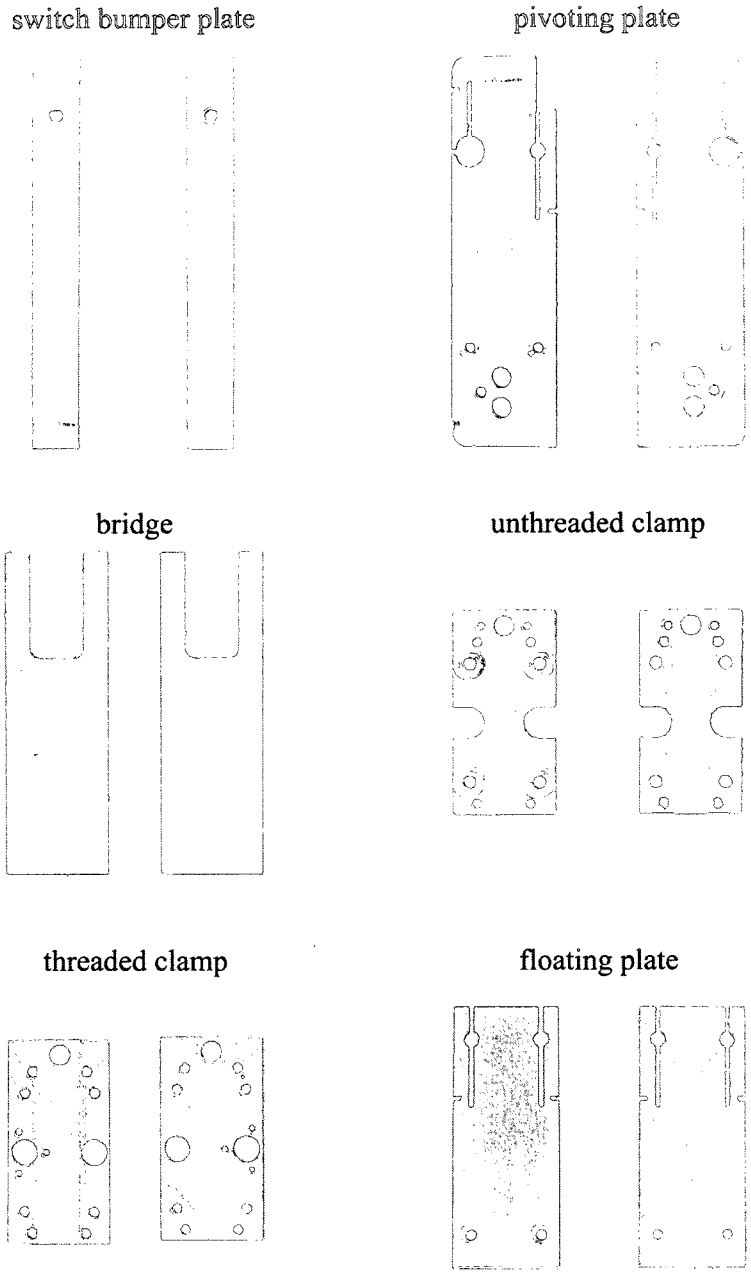

Fig. B.2 Main aluminum parts of the THMB device once cut-out from their respective fixture plates.

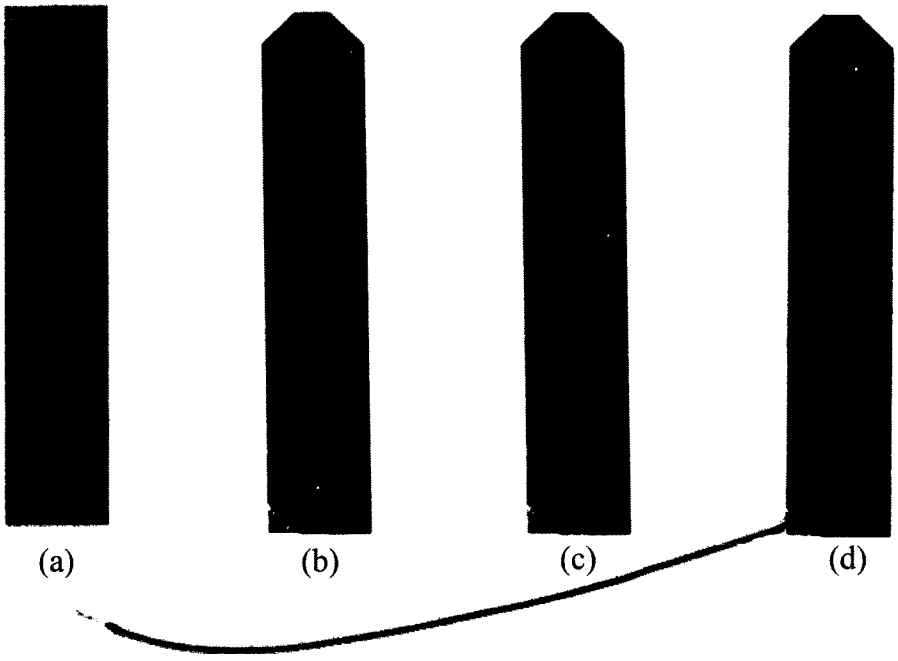

Fig. B.3 Preparation of the piezoelectric actuators for the 2^{nd} version of the THMB. (a) The actuator arrives from the manufacturer as a slab. (b) The tip is beveled and the bottom left is scraped off to allow easy access to the middle electrode. (c) A coat of electrically-conductive paint is applied to protect the thin outside electrodes. (d) A thin wire is soldered to the middle electrode.

Appendix C

Electronics that Drive the THMB Device

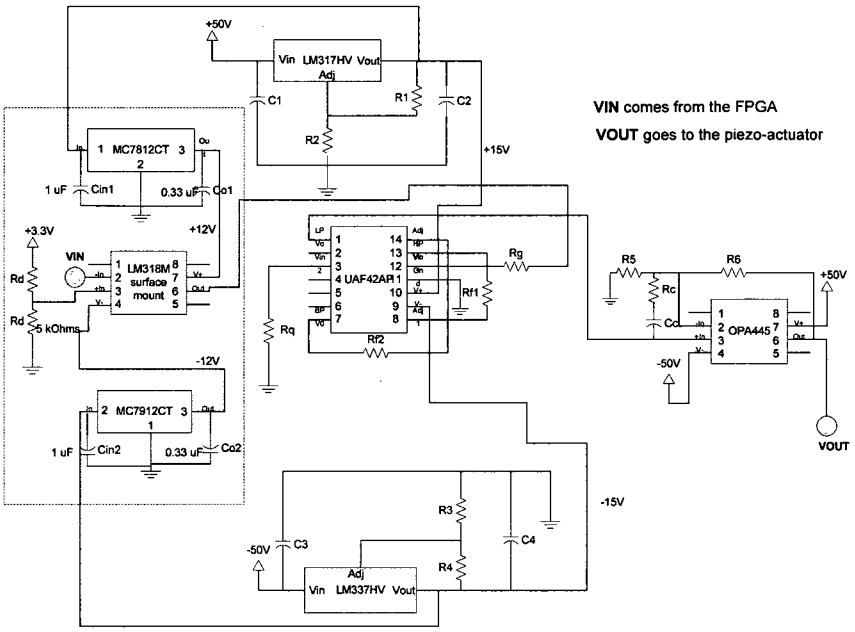

Fig. C.1 Interfacing electronics that control the position of a single piezo-actuator. The 3.3-V driving voltage comes from the Nova Engineering Constellation-10KE^{TM} development board which operates an Altera FLEX 10KE^{TM} chip. There are 8 of these channels in the control box, one for each actuator.

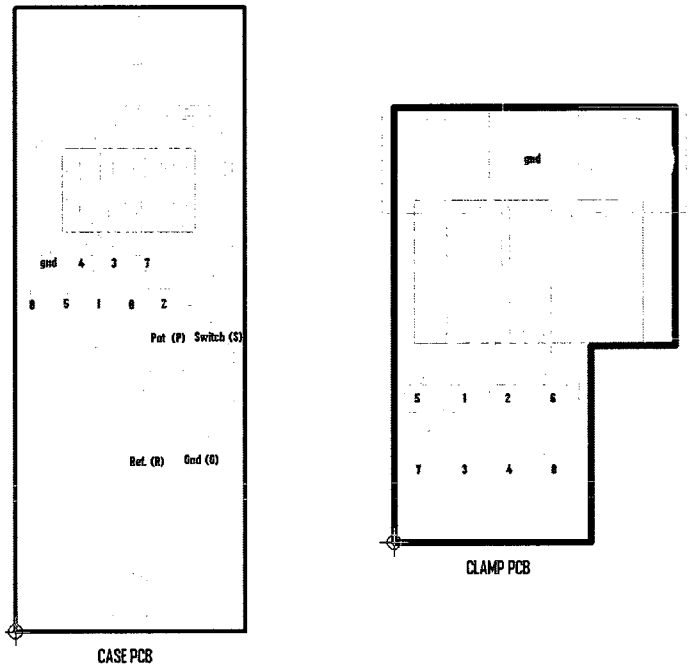

Fig. C.2 PCBs located inside the handheld case of the THMB.

Appendix D

THMB Components

Unless explicitly indicated, part numbers are from Digikey.

Screws	4 screws to secure bridge (M1.6 x 6mm flat head) 4 screws to clamp 2 NZC shafts (M1.6 x 6mm flat head) 4 clamp screws (M2.5 x16mm pan head - NB: used cheese head instead) 1 switch bumper screw (M2.5 x 12mm pan head) + nuts (M2.5) 1 switch mount screw (M1.6 x 12mm flat head) 2 case PCB mount screws (M1.6 x 12mm flat head) + glued nuts (M1.6) 2 clamp PCB mount screws (M1.6 x 4mm cheese head) 2 potentiometer mount screws (M1.6 x 4mm cheese head) 3 screws to mount guide tube on threaded clamp (M1.2 x 3mm slotted cheese head machine screws)
O-rings	2 o-rings to stop sliding motion (#004) 4 clamping o-rings (#004) 1 o-ring to fasten potentiometer to clamp PCB (#003)
Power Supply	1 (Hammond) box for new power supply (20cm x 30cm x 10cm) 1 On/Off Switch 2 LEDs

Manufacturing	2 Brasseler saw (.10mm) 911HEF-220 (or 180) 1 soldering flux for piezo-actuators (7 ml - #67 liquid flux from www.piezo.com) 2 miniature carbide 2-flute end mill square-end, Regular Flute, .020" Mill Dia, .060" L Cut (McMaster 8832A62) for the grooves in the lexan rod holder 2 miniature carbide two-flute end mill square-end, .031 mill dia, .1550" L-cut (McMaster 8915A34) 1 Permatex Quick Grid rear window defogger repair kit - conductive pain (Canadian tire - Permatex #15067)
Assembling	1 Loctite 290 (threadlocker) - very low viscosity (McMaster Loctite 290) 1 Loctite 271 (threadlocker with higher viscocity) (McMaster Loctite 271) 1 Diamond compound for polishing sliding rods and NZC tubes (McMaster 4776A37 or 4428A19) 1 Loctite 454 (Prism super glue surface-insensitive) (McMaster Loctite 454) 1 Protective coating for PC boards to avoid short-circuits (MG Chemicals 422A-340G)
Wires	2 Cooner Wire: (AS323Rev1) 16 .4mm flexible wire to access piezo-actuators 12 .5mm flexible wire for connections inside the handheld case. need multiple colors
Hand Held Case	1 Hammond case 1593 Q (HM359-ND) 2 10-pos conn. .5mm vertical (HFL110CT-ND) 1 CABLE FLAT FLEX 10POS .5MM 2" (HFF-10U-02-ND) 1 Push Button: SWITCH LT RT ANGLE 160GF 3.15MM (P12208S-ND) 1 Push Button: LT SWITCH RT ANGLE L3.15MM 130G (P10881S-ND) 1 Potentiometer: SENSOR POSITION LINEAR SLIDE 10KOhms (P12338-ND) 1 .20" Brass rod (stock #1596 - Engineering ing. Chicago Il - from McMaster)

Amplification Box	1 Pin Headers with latches (20 conn) (with latch AHE20G-ND or without latch AHE20H-ND) 1 Pin Headers with latches (10 conn) (with latch AHE10G-ND or without latch AHE10H-ND) 1 socket connector (10 connections) with polarizing key (CKC10G-ND) 1 socket connector (20 connections) with polarizing key (CKC20G-ND) 1 CONN HOUSING 4POS .100 DUA (WM2519-ND) 1 CONN TERM FEMALE 24-30AWG GOLD - pck of 100 (WM2513-ND) 1 Connector Housing - with Mounting Ears and Detent, Panel Mount (plug and receptacle) (WM1207-ND and WM1206ND) 1 .062" terminals (WM1002-ND and WM1003-ND) 8 Comparators: IC OP AMP HIGH SLEW RATE 8-SOIC (LM318M-ND) 8 Amp: IC HIGH VOLTAGE OP AMP 8-SOIC (OPA445AU-ND) 8 Filter: n IC UNIV ACTIVE FILTER 16-SOIC (UAF42AU-ND) 1 AD Converter: IC ADC SRL 10BIT 2.7-5.25V 8SOIC (MAX1243BCSA-ND) 1 HV Pos Regulator: IC REG POSITIVE ADJ TO-220 (LM317HVT-ND) 1 HV Neg Regulator: IC REGULATOR NEG ADJ TO-39 (LM337HVH-ND) 1 Heat sink for pos HV reg 1 Heat sink for neg HV reg 1 Fan 1 LV neg reg: IC REG VOLT NEG 1A 12V TO220 (MC7912CTGOS-ND) 1 LV pos reg: IC REG VOLT POS 1A 12V TO220 (MC7812CTGOS-ND) 1 Power socket + power connector 6 Pins: CONN PIN IDC 26AWG TIN - pck 100 (H3839-ND) 32 Metric Standoffs: STANDOFF M/F MET 5MM HEX 10MM L (4313K-ND)
Other	1 .014" spring Protective Headers (.100 x .100) 4 straight 50-connection - A26282-ND 4 right-angle 50-connection - A26300-ND socket connectors (.100 x .100) 7 16-connection - CKC16G-ND 2 CSC34G-ND

Appendix E

Ethics Certificate

This appendix holds the approved ethics certificate for research involving human subjects. The certificates cover the studies described in Chapter 3, Chapter 5 and Chapter 6 respectively. The studies described in Chapter 5 were conducted by Joseph Luk who was at the time a master's student at the University of British Columbia under the supervision of Prof. Karon MacLean.

References

[1] N. Asamura, N. Yokoyama, and H. Shinoda. Selectively stimulating skin receptors for tactile display. *IEEE Computer Graphics and Applications*, 18(6):32–37, 1998.

[2] P. Bach-y-Rita and C. C Collins. Sensory substitution systems using the skin for the input to the brain. *Journal of the Audio Engineering Society*, 19(5):427–429, 1971.

[3] P. Bach-y-Rita, K. A. Kaczmarek, M. E. Tyler, and M. Garcia-Lara. Form perception with a 49-point electrotactile stimulus array on the tongue: A technical note. *Journal of Rehabilitation Research and Development*, 35:427–430, 1998.

[4] M. Benali-Khoudja, M. Hafez, J. M. Alexandre, and A. Kheddar. Tactile interfaces: a state of the art survey. In *Proceedings of International Symposium on Robotics*, March 23–26 2004.

[5] M. Benali-Khoudja, M. Hafez, J.-M. Alexandre, A. Kheddar, and V. Moreau. VITAL: a new low-cost vibro-tactile display system. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2004 (ICRA '04)*, volume 1, pages 721–726, April 2004.

[6] S. J. Bensmaïa, Y. Y. Leung, S. S. Hsiao, and K. O. Johnson. Vibratory adaptation of cutaneous mechanoreceptive afferents. *Journal of Neurophysiology*, 94:3023–3036, 2005.

[7] P. Bertelson, P. Mousty, and G. D'Alimonte. A study of Braille reading: 2. patterns of hand activity in one-handed and two-handed reading. *Quarterly Journal of Experimental Psychology*, 37A(2):235–56, 1985.

[8] A. Bicchi, E. P. Scilingo, N. Sgambelluri, and D. De Rossi. Haptic interfaces based on magnetorheological fluids. In *Proceedings of Eurohaptics 2002*, pages 6–11, 2002.

[9] D. K. Biegelsen, W. B. Jackson, L.-E. Swartz, A. A. Berlin, and P. C. Cheung. Pneumatic actuator with elastomeric membrane and low-power electrostatic flap valve arrangement. US Patent No. 6807892, 2004.

[10] M. Biet, F. Giraud, and B. Lemaire-Semail. Squeeze film effect for the design of an ultrasonic tactile plate. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 54(12):2678–2688, 2007.

[11] J. Biggs and M. A. Srinivasan. Tangential versus normal displacement of skin: Biomechanics and perceived intensity. In *Proceedings of the International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 121–128, 2002.

[12] M. Blattner, D. Sumikawa, and R. Greenberg. Earcons and icons: Their structure and common design principles. *Human Computer Interaction*, 4(1):11–44, 1989.

[13] D. Blazie. Refreshable braille now and in the years ahead. Transcript of lecture available online at `http://nfb.org/legacy/bm/bm00/bm0001/bm000110.htm`, 1998.

[14] J. Bliss, M. Katcher, C. Rogers, and R. Shepard. Optical-to-tactile image conversion for the blind. *IEEE Transactions on Man-Machine Systems*, 11(1):58–65, 1970.

[15] S. A. Brewster and L. M. Brown. Non-visual information display using tactons. In *Proceedings of CHI '04: extended abstracts on Human factors in computing systems*, pages 787–788. ACM Press, 2004.

[16] S. A. Brewster and L. M. Brown. Tactons: Structured tactile messages for non-visual information display. In *Proceedings of the 5th Australasian User Interface Conference (AUIC2004)*, pages 15–23, 2004.

[17] P. L. Brooks and B. J. Frost. The development and evaluation of a tactile vocoder for the profoundly deaf. *Canadian Journal of Public Health*, 77 Suppl 1:108–113, 1986.

[18] P. L. Brooks, B. J. Frost, J. L. Mason, and K. Chung. Acquisition of a 250-word vocabulary through a tactile vocoder. *Journal of the Acoustical Society of America*, 77(4):1576–1579, 1985.

[19] P. L. Brooks, B. J. Frost, J. L. Mason, and D. M. Gibson. Word and feature identification by profoundly deaf teenagers using the Queen's university tactile vocoder. *Journal of Speech and Hearing Research*, 30:137–141, 1987.

[20] L. M. Brown, S. A. Brewster, and H. C. Purchase. Multidimensional tactons for non-visual information presentation in mobile devices. In *MobileHCI '06: Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, pages 231–238. ACM, 2006.

[21] J. L. Burke, M. S. Prewett, A. A. Gray, L. Yang, F. R. B. Stilson, M. D. Coovert, L. R. Elliot, and E. Redden. Comparing the effects of visual-auditory and visual-tactile feedback on user performance: a meta-analysis. In *ICMI '06: Proceedings of the 8th international conference on Multimodal interfaces*, pages 108–117. ACM, 2006.

[22] W. Buxton. Chunking and phrasing and the design of human-computer dialogues. In *Proceedings of the IFIP World Computer Congress*, pages 475–480, 1986.

[23] D. G. Caldwell, N. Tsagarakis, and C. Giesler. An integrated tactile/shear feedback array for stimulation of finger mechanoreceptor. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 287–292, 1999.

[24] C. Campbell, S. Zhai, K. May, and P. Maglo. What you feel must be what you see: Adding tactile feedback to the trackpoint. In *Proceedings of IFIP Interact'99, Edinburgh, UK, IOS Press*, pages 383–390, 1999.

[25] A. Chan, K. E. MacLean, and J. McGrenere. Learning and identifying haptic icons under workload. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 432–439, 2005.

[26] A. Chang, S. O'Modhrain, J. R. Gunther, E. Gunther, and H. Ishii. Comtouch: design of a vibrotactile communication device. In *Proceedings of the 4th conference on Designing Interactive Systems*, pages 312–320. ACM, 2002.

[27] C. E. Chapman. Active versus passive touch: factors influencing the transmission of somatosensory signals to primary somatosensory cortex. *Canadian Journal of Physiology and Pharmacology*, 72(5):558–570, 1994.

[28] L. Chittaro. Visualizing information on mobile devices. *IEEE Computer*, 39(3):40–45, 2006.

[29] H. R. Choi, S. W. Lee, K. M. Jung, J. C. Koo., S. I. Lee, H. G. Choi, J. W. Jeon, and J. D. Nam. Tactile display as a braille display for the visually disabled. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2004)*, volume 2, pages 1985–1990, 2004.

[30] R. W. Cholewiak and J. C. Craig. Vibrotactile pattern recognition and discrimination at several body sites. *Perception & Psychophysics*, 35:503–514, 1984.

[31] R. W. Cholewiak and C. E. Sherrick. A computer-controlled matrix system for presentation to the skin of complex spatiotemporal patterns. *Behavior Research Methods and Instrumentation*, 13(5):667–673, 1981.

[32] M. A. Clements, L. D. Braida, and N. I. Durlach. Tactile communication of speech : Comparison of two computer-based displays. *Journal of Rehabilitation Research and Development*, 25(4):25–44, 1988.

[33] R. H. Cook. An automatic stall prevention control for supersonic fighter aircraft. *Journal of Aircraft*, 2(3):171–175, 1965.

[34] J. C. Craig. Modes of vibrotactile pattern generation. *Journal of Experimental Psychology: Human Perception and Performance*, 6(1):151–166, 1980.

[35] J. C. Craig. Some factors affecting tactile pattern recognition. *International Journal of Neuroscience*, 19:47–58, 1983.

[36] J. C. Craig. Tactile pattern perception and its perturbations. *Journal of Accoustical Society of America*, 77(1):238–246, 1985.

[37] J. C. Craig. Identification of scanned and static tactile patterns. *Perception and Psychophysics*, 64(1):107–120, 2002.

[38] K. Drewing, M. Fritschi, R. Zopf, M.O. Ernst, and M. Buss. First evaluation of a novel tactile display exerting shear force via lateral displacement. *ACM Transactions on Applied Perception*, 2(2):118–131, 2005.

[39] M. Dunlop and S. Brewster. The challenge of mobile devices for human computer interaction. *Personal and Ubiquitous Computing*, 6(4):235–236, 2002.

[40] B. B. Edin and N. Johansson. Skin strain patterns provide kinaesthetic information to the human central nervous system. *Journal of Physiology*, 487:243–251, 1995.

[41] M. Enriquez and K. MacLean. The role of choice in longitudinal recall of meaningful tactile signals. In *Proceedings of the IEEE Symposium on Haptic*

Interfaces for Virtual Environments and Teleoperator Systems, pages 49–56, March 2008.

[42] H. Fischer, R. Trapp, L. Schüle, and B. Hoffmann. Actuator array for use in minimally invasive surgery. *Journal de Physique IV*, 07(C5):609–614, 1997.

[43] E. Foulke. Reading braille. In W. Schiff and E. Foulke, editors, *Tactual Perception: A Sourcebook*, chapter Reading Braille, pages 168–208. Cambridge Univ Press, 1982.

[44] J. Fricke. Substituting friction by weak vibration on a tactile pin array. In *Proceedings of the IEE Colloquium on Developments in Tactile Displays*, volume Digest No. 1997/012, pages 3/1–3/3, 1997.

[45] M. Fukumoto and T. Sugimura. Active click: Tactile feedback for touch panels. In *Proceedings of CHI 2001*, pages 121–122, 2001.

[46] A. Gallace, M. Auvray, H. Z. Tan, and C. Spence. When visual transients impair tactile change detection: a novel case of crossmodal change blindness? *Neuroscience Letters*, 398(3):280–285, 2006.

[47] A. Gallace, H. Z. Tan, and C. Spence. Tactile change detection. In *Proceedings First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems (WHC'05)*, pages 12–16, 2005.

[48] A. Gallace, H. Z. Tan, and C. Spence. The body surface as a communication system: The state of the art after 50 years. *Presence: Teleoper. Virtual Environ.*, 16(6):655–676, 2007.

[49] R. H. Gault. Recent developments in vibro-tactile research. *Journal of the Franklin Institute*, 221:703–719, 1936.

[50] F. A. Geldard. Adventures in tactile literacy. *The American Psychologist*, 12:115–124, 1957.

[51] F. A. Geldard. Some neglected possibilities of communication. *Science*, 131:1583–1588, 1960.

[52] F. A. Geldard. The mutability of time and space on the skin. *Journal of Accoustical Society of America*, 77(1):233–237, 1985.

[53] G. A. Gescheider, R. T. Verrillo, A. J. Capraro, and R. D. Hamer. Enhancement of vibrotactile sensation magnitude and predictions from the duplex model of mechanoreception. *Sensory Processes*, 1(3):187–203, 1977.

[54] J. J. Gibson. Observations on active touch. *Psychological Review*, 69:477–491, 1962.

[55] A. W. Goodwin and H. E. Wheat. Magnitude estimation of contact force when objects with different shapes are applied passively to the fingerpad. *Journal of Somatosensory and Motor Research*, 9(4):339–344, 1992.

[56] V. Hayward, O. R. Astley, M. Cruz-Hernández, D. Grant, and G. Robles-De-La-Torre. Haptic interfaces and devices. *Sensor Review*, 24(1):16–29, 2004.

[57] V. Hayward and M. Cruz-Hernández. Tactile display device using distributed lateral skin stretch. In *Proceedings of the Haptic Interfaces for Virtual Environment and Teleoperator Systems Symposium*, volume DSC-69-2, pages 1309–1314. ASME, 2000.

[58] M. A. Heller. Active and passive touch: the influence of exploration time on form recognition. *Journal of General Psychology*, 110(2):243–249, 1984.

[59] H.-N. Ho and L. A. Jones. Development and evaluation of a thermal display for material identification and discrimination. *ACM Transactions Applied Perception*, 4(2):13, 2007.

[60] J. Ho and S. S. Intille. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. In *Proceedings of CHI'05*, pages 909–918, New York, NY, USA, 2005. ACM.

[61] M. Hollins, R. Faldowski, S. Rao, and F. Young. Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis. *Perception & Psychophysics*, 54:697–705, 1993.

[62] Y. Ikei, K. Wakamatsu, and S. Fukuda. Texture presentation by vibratory tactile display. In *Proceedings of the IEEE Annual Virtual Reality International Symposium*, pages 199–205, 1997.

[63] Immersion®. VibeTonz™ mobile player: An embedded solution for playing VibeTonz touch sensations on mobile handsets, corporate literature. www.immersion.com.

[64] S. Ino, S. Shimizu, T. Odagawa, M. Sato, M. Takahashi, T. Izumi, and T. Ifukube. A tactile display for presenting quality of materials by changing the temperature of skin surface. In *Proceedings of the IEEE International Workshop on Robot and Human Communication*, pages 220–224, 1993.

[65] H. Ishii and B. Ulmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of CHI'97*, pages 234–241. ACM, March 1997.

[66] T. Iwamoto, D. Akaho, and H. Shinoda. High resolution tactile display using acoustic radiation pressure. In *Proceedings of SICE Annual Conference*, pages 1239–1244, August 2004.

[67] T. Iwamoto and H. Shinoda. Ultrasound tactile display for stress field reproduction - examination of non-vibratory tactile apparent movement. *Proceedings of First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'05*, pages 220–228, March 2005.

[68] R. J. Jagacinsk and J. M. Flach. *Control Theory for Humans: Quantitative Approaches to Modeling Performance.* Lawrence Erlbaum Associates, 2003.

[69] R. S. Johansson and I. Birznieks. First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience*, 7(2):170–177, 2004.

[70] R. S. Johansson and A. B. Vallbo. Tactile sensory coding in the glabrous skin of the human hand. *Trends Neuroscience*, 6:27–31, 1983.

[71] K. O. Johnson. The roles and functions of cutaneous mechanoreceptors. *Current Opinion in Neurobiology*, 11(4):455–461, 2001.

[72] K. O. Johnson, S. S. Hsiao, and I. A. Twombly. Neural mechanisms of tactile form recognition. In M. S. Gazzaniga, editor, *The Cognitive Neurosciences*, pages 235–268. The MIT press, Cambridge, MA, USA, 1995.

[73] L. A. Jones, J. Kunkel, and E. Torres. Tactile vocabulary for tactile displays. In *Proceedings of Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, pages 574–575, March 2007.

[74] L. A. Jones, B. Lockyer, and E. Piateski. Tactile display and vibrotactile pattern recognition on the torso. *Advanced Robotics*, 20(12):1359–1374(16), 2006.

[75] L. A. Jones and N. B. Sarter. Tactile displays: Guidance for their design and application. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(1):90–111, February 2008.

[76] S. Jones, M. Jones, G. Marsden, D. Patel, and A. Cockburn. An evaluation of integrated zooming and scrolling on small screens. *Int. J. Hum.-Comput. Stud.*, 63(3):271–303, 2005.

[77] M. Jungmann and H. F. Schlaak. Miniaturized electrostatic tactile display with high structural compliance. In *Proceedings of Eurohaptics 2002*, pages 12–17, 2002.

[78] K. A. Kaczmarek, M. E. Tyler, and P. Bach-y-Rita. Pattern identification on a fingertip-scanned electrotactile display. In *Proceedings of the 19th International IEEE/EMBS Conference*, pages 1694–1697, 1997.

[79] H. Kajimoto, M. Inami, N. Kawakami, and S. Tachi. Smarttouch: Augmentation of skin sensation with electrocutaneous display. In *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 40–46, March 2003.

[80] T. Kamba, S. Elson, T. Harpold, T. Stamper, and P. Sukaviriya. Using small screen space more efficiently. In *Proceedings of CHI 96*, pages 383–390, 1996.

[81] Y. Kawai and F. Tomita. Interactive tactile display system: a support system for the visually disabled to recognize 3d objects. In *Proceedings of the second annual ACM conference on Assistive technologies (Assets '96)*, pages 45–50. ACM, 1996.

[82] R. Kikuuwe, A. Sano, H. Mochiyama, N. Takesue, and H. Fujimoto. Enhancing haptic detection of surface undulation. *ACM Transactions on Applied Perception*, 2(1):46–67, 2005.

[83] K. Kim, J. E. Colgate, and M. A. Peshkin. On the design of a thermal display for upper extremity prosthetics. In *Proceedings of Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 413–419, March 2008.

[84] D. A. Kontarinis, J. S. Son, W. Peine, and R. D. Howe. A tactile shape sensing and display system for teleoperated manipulation. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 641–646, 1995.

[85] M. Konyo, S. Tadokoro, and T. Takamori. Artificial tactile feel display using soft gel actuators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3416–3421, 2000.

[86] J. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1–27, 1964.

[87] K.-U. Kyung, M. Ahn, D.S. Kwon, and M. A. Srinivasan. A compact planar distributed tactile display and effects of frequency on texture judgment. *Advanced Robotics*, 20(5):563–580(18), 2006. Tactile display.

[88] P. Laitinen and J. Mawnpaa. Enabling mobile haptic design: Piezoelectric actuator technology properties in hand held devices. In *Proceedings of IEEE International Workshop on Haptic Audio Visual Environments and their Applications (HAVE 2006)*, pages 40–43, 2006.

[89] G. E. Legge, C. M. Madison, and S. J. Mansfield. Measuring braille reading speed with the MNREAD test. *Visual Impairment Research*, 1(3):131–145, 1999.

[90] R. Leung, K. MacLean, M. B. Bertelsen, and M. Saubhasik. Evaluation of haptically augmented touchscreen GUI elements under cognitive load. In *Proceedings of the 9th international conference on Multimodal interfaces (ICMI '07)*, pages 374–381. ACM, 2007.

[91] V. Levesque and V. Hayward. Experimental evidence of lateral skin strain during tactile exploration. In *Proceedings of Eurohaptics 2003*, pages 261–275, 2003.

[92] V. Levesque and V. Hayward. Tactile graphics rendering using three laterotactile drawing primitives. In *Proceedings of the 16th Symposium on Haptic Interfaces For Virtual Environment and Teleoperator Systems*, pages 429–436, 2008.

[93] V. Levesque, J. Pasquero, V. Hayward, and M. Legault. Display of virtual Braille dots by lateral skin deformation: Feasibility study. *ACM Transactions on Applied Perception*, 2(2):132–149, 2005.

[94] J. Linjama, J. Häkkilä, and S. Ronkainen. Gesture interfaces for mobile devices - minimalist approach for haptic interaction. In *Proceedings of CHI 2005*, 2005.

[95] J. Luk, J. Pasquero, S. Little, K. MacLean, V. Levesque, and V. Hayward. A role for haptics in mobile interaction: Initial design using a handheld tactile display prototype. In *Proceedings of CHI 2006*, pages 171–180, 2006.

[96] K. MacLean. Haptics in the wild: Interaction design for everyday interfaces. *To appear in Reviews of Human Factors and Ergonomics (HFES)*, 2008.

[97] K. MacLean and M. Enriquez. Perceptual design of haptic icons. In *Proceedings of Eurohaptics 2003*, pages 351–363, 2003.

[98] L. E. Magee and J. M. Kennedy. Exploring pictures tactually. *Nature*, 283:287–288, 1980.

[99] Y. Makino, N. Asamura, and H. Shinoda. A cutaneous feeling display using suction pressure. In *Proceedings of the SICE 2003 Annual Conference*, volume 3, pages 2931–2934, 2003.

[100] T. Maucher, K. Meier, and J. Schemmel. An interactive tactile graphics display. In *Proceedings of the Sixth International Symposium on Signal Processing and its Applications*, pages 190–193, 2001.

[101] S. Millar. *Reading by Touch*, chapter Hand-movements in Reading: Measures, Functions and Proficiencies, pages 56–97. Routledge, New York, London, June 1997.

[102] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63:81–97, 1956.

[103] G. Moy, C. Wagner, and R. S. Fearing. A compliant tactile display for teletaction. In *Proceedings of the IEEE International Conference on Robotics and Automation USA*, volume 4, pages 3409–3415, 2000.

[104] T. Nara, M. Takasaki, T. Maeda, T. Higuchi, S. Ando, and S. Tachi. Surface acoustic wave tactile display. *IEEE Computer Graphics and Applications*, 21(6):56–63, 2001.

[105] M. Nolano, V. Provitera, and C. Crisci. Quantification of myelinated endings and mechanoreceptors in human digital skin. *Annals of Neurology*, 54:197–205, 2003.

[106] I. Oakley and S. O'Modhrain. 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)*, pages 40–49. IEEE Computer Society, 2005.

[107] M. Ohka and Y. Muramatsu. Fine texture presentation system for tactile virtual reality. In *Proceedings of 2nd World Manufacturing Congress (WMC'99)*, pages 57–62, 1999.

[108] M. V. Ottermo. *Virtual Palpation Gripper.* PhD thesis, Norwegian University of Science and Technology Faculty of Information Technology, Mathematics and Electrical Engineering Department of Engineering Cybernetics, June 2006.

[109] A. Oulasvirta, S. Tamminen, V. Roto, and J. Kuorelahti. 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*, 2005.

[110] S. Oviatt. Ten myths of multimodal interaction. *Communications of the ACM*, 42(11):74–81, 1999.

[111] M. Paré, C. Behets, and O. Cornu. Paucity of presumed Ruffini corpuscles in the index fingerpad of humans. *Journal of Comparative Neurology*, 356:260–266, 2003.

[112] M. Paré, H. Carnahan, and A. M. Smith. Magnitude estimation of tangential force applied to the fingerpad. *Experimental Brain Research*, 142(3):342–348, 2002.

[113] M. Paré, A. M. Smith, and F. L. Rice. Distribution and terminal arborizations of cutaneous mechanoreceptors in the glabrous finger pads of the monkey. *Journal of Comparative Neurology*, 445:347–359, 2002.

[114] J. Pasquero. STReSS: A tactile display using lateral skin stretch. Master's thesis, McGill University, 2003.

[115] J. Pasquero. Survey on communication through touch. Technical Report TR-CIM 06.04, Center for Intelligent Machines - McGill University, 2006.

[116] J. Pasquero and V. Hayward. STReSS: A practical tactile display system with one millimeter spatial resolution and 700 hz refresh rate. In *Proceedings of Eurohaptics*, pages 94–110, 2003.

[117] J. Pasquero, J. Luk, V. Levesque, Q. Wang, V. Hayward, and K. E. MacLean. Haptically enabled handheld information display with distributed tactile transducer. *IEEE Transactions on Multimedia*, 9(4):746–753, June 2007.

[118] J. Pasquero, J. Luk, S. Little, and K. MacLean. Perceptual analysis of haptic icons: an investigation into the validity of cluster sorted MDS. In *Proceedings of the Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 437–444, 2006.

[119] D. T. V. Pawluk, C. P. van Buskirk, J. H. Killebrew, S. Hsiao, and K. O. Johnson. Control and pattern specification for high density tactile display. In *Proceedings of the Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, volume DSC-64, pages 97–102. ASME, 1998.

[120] J. R. Phillips, R. S. Johansson, and K. O. Johnson. Representation of Braille characters in human nerve fibres. *Experimental Brain Research*, 81:589–592, 1990.

[121] J. R. Phillips and K. O. Johnson. Tactile spatial resolution. III. A continuum mechanics model of skin predicting mechanoreceptor responses to bars, edges, and gratings. *Journal of Neurophysiology*, 46(6):1204–1225, 1981.

[122] J. R. Phillips and K. O. Johnson. Neural mechanisms of scanned and stationary touch. *Journal of Accoustical Society of America*, 77(1):220–224, 1985.

[123] A. Pirhonen, S. Brewster, and C. Holguin. Gestural and audio metaphors as a means of control for mobile devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'02)*, pages 291–298. ACM, 2002.

[124] I. Poupyrev, S. Maruyama, and J. Rekimoto. Ambient touch: designing tactile interfaces for handheld devices. In *Proceedings of the 15th annual ACM symposium on User Interface Software and Technology (UIST '02)*, pages 51–60. ACM Press, 2002.

[125] R. C. R. C. Petersen. Tactile display system. US Patent No. 6,734,785, 2004.

[126] C. Ramstein. Combining haptic and Braille technologies: Design issues and pilot study. In *Proceedings of the second annual ACM conference on Assistive Technologies*, pages 37–44. ACM SIGCAPH Conference on Assistive Technologies, 1996.

[127] C. M. Reed, W. M. Rabinowitz, N. I. Durlach, L. D. Braida, S. Conway-Fithian, and M. C. Schultz. Research on the Tadoma method of speech communication. *Journal of the Acoustical Society of America*, 77(1):247–257, 1985.

[128] B. L. Richardson, M. Symmons, and R. Accardi. The TDS: a new device for comparing active and passive-guided touch. *IEEE Transactions on Rehabilitation Engineering*, 8(3):414–417, 2000.

[129] J. Roberts, O. Slattery, and D. Kardos. Rotating-wheel Braille display for continuous refreshable Braille. In *Digest of Technical Papers, Society for Information Display International Symposium*, volume XXXI, pages 1130–1133, 2000.

[130] G. Robles-De-La-Torre. The importance of the sense of touch in virtual and real environments. *IEEE Multimedia*, 13(3):24–30, 2006.

[131] V. Roto and A. Oulasvirta. Need for non-visual feedback with long response times in mobile HCI. In *Proceedings of International World Wide Web Conference*, pages 775–781, 2005.

[132] J. Rovan and V. Hayward. Typology of tactile sounds and their synthesis in gesture-driven computer music performance. In *Trends in Gestural Control of Music*, pages 297–320. IRCAM, 2000.

[133] L. Rovers and H. van Essen. Design and evaluation of hapticons for enriched instant messaging. In *Proceedings of Eurohaptics 2004*, June 2004.

[134] A. H. Rupert. An instrumentation solution for reducing spatial disorientation mishaps. *IEEE Engineering in Medicine and Biology Magazine*, 19(2):71–80, 2000.

[135] D. D. Salvucci, D. Markley, M. Zuber, and D. P. Brumby. iPod distraction: Effects of portable music-player use on driver performance. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'07)*, pages 243–250. ACM, 2007.

[136] I. Sarakoglou, N. Tsagarakis, and D. G. Caldwell. A portable fingertip tactile feedback array? Transmission system reliability and modelling. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'05)*, pages 547–548, 2005.

[137] C. E. Sherrick. Cutaneous communication. In W. D. Neff, editor, *Contributions to Sensory Physiology*, volume VI, pages 1–43. Academic Press, 1982.

[138] C. E. Sherrick. Basic and applied research on tactile aids for deaf people: Progress and prospects. *The Journal of the Acoustical Society of America*, 75(5):1325–1342, 1984.

[139] M. Shinohara, Y. Shimizu, and A. Mochizuki. Three-dimensional tactile display for the blind. *IEEE Transactions on Rehabilitation Engineering*, 6(3):249–256, 1998.

[140] M. Silfverberg. Using mobile keypads with limited visual feedback: Implications to handheld and wearable devices. *Proceedings of Mobile HCI 2003*, 2795:76–90, 2003.

[141] A. M. Smith, C. E. Chapman, M. Deslandes, J. S. Langlais, and M. P. Thibodeau. Role of friction and tangential force variation in the subjective scaling of tactile roughness. *Experimental Brain Research*, 144(2):211–223, 2002.

[142] J. G. Smits, S. I. Dalke, and T. K. Cooney. The constituent equations of piezoelectric bimorphs. *Sensors and Actuators A*, 28:41–61, 1991.

[143] R. L. Solso, C. Juel, and D. C. Rubin. The frequency and versatility of initial and terminal letters in English words. *Journal of Verbal Learning and Verbal Behavior*, 21:220–235, 1982.

[144] M. A. Srinivasan and K. Dandekar. An investigation of the mechanics of tactile sense using two-dimensional models of the primate fingertip. *Journal of Biomechanical Engineering*, 118(1):48–55, 1996.

[145] M. A. Srinivasan, J. M. Whitehouse, and R. H. LaMotte. Tactile detection of slip: Surface microgeometry and peripheral neural codes. *Journal of Neurophysiology*, 63(6):1323–1332, 1990.

[146] B. Stöger and K. Miesenberger. The conventional Braille display - state of the art and future perspectives. In *Proceedings of the 4th international conference on Computers for handicapped persons*, pages 447–454. Springer-Verlag New York, Inc., 1994.

[147] I. R. Summers and C. M. Chanter. A broadband tactile array on the fingertip. *Journal of the Acoustical Society of America*, 112:2118–2126, 2002.

[148] S. Tachi, K. Tanie, K. Komoriya, and M. Abe. Electrocutaneous communication in a guide dog robot (meldog). *IEEE Transactions on Biomedical Engineering*, 32(7):461–469, 1985.

[149] M. Takasaki, H. Kotani, T. Nara, and T. Mizuno. Transparent surface acoustic wave tactile display. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1115–1120, 2005.

[150] H. Z. Tan, N. I. Durlach, C. M. Reed, and W. M. Rabinowitz. Information transmission with a multifinger tactual display. *Perception & Psychophysics*, 61(6):993–1008, 1999.

[151] H. Z. Tan and A. Pentland. Tactual displays for wearable computing. In *Proceedings of the International Symposium on Wearable Computers*, pages 84–89, 1997.

[152] A. Tang, P. McLachlan, K. Lowe, R. S. Chalapati, and K. E. MacLean. Perceiving ordinal data haptically under workload. In *Proceedings of the 7th International Conference on Multimodal Interfaces, ICMI'05*, pages 317–324, 2005.

[153] H. Tang and D. J. Beebe. A microfabricated electrostatic haptic display for persons with visual impairments. *IEEE Transactions on rehabilitation engineering*, 6(3):241–248, 1998.

[154] H. Tang and D. J. Beebe. Design and microfabrication of a flexible oral electrotactile display. *IEEE Journal of Microelectromechanical Systems*, 12(1):29–36, 2003.

[155] M. A. Taylor, A. R. Ferber, and J. E. Colgate. Assessing the efficacy of variable compliance tactile displays. In *Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, pages 427–432. IEEE Computer Society, 2007.

[156] P. M. Taylor, A. Hosseini-Sianaki, and C. J. Varley. An electrorheological fluid-based tactile array for virtual environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 18–22, 1996.

[157] P. M. Taylor, A. Mose, and A. Creed. The design and control of a tactile display based on shape memory alloys. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1317–1323, 1997.

[158] O. Tretiakoff and A. Tretiakoff. Electromechanical transducer for relief display panel. US Patent No. 4,044,350, 1977.

[159] L. Tsogo, M. H. Masson, and A. Bardot. Multidimensional scaling methods for many-object sets: A review. *Multivariate Behavioral Research*, 35(3):307–319, 2000.

[160] C. L. van Doren. A model of spatiotemporal tactile sensitivity linking psychophysics to tissue mechanics. *Journal of Accoustical Society of America*, 85(5):2065–2080, 1989.

[161] C. L. van Doren, D. G. Pelli, and R.T. Verillo. A device for measuring tactile spatiotemporal sensitivity. *Journal of Accoustical Society of America*, 81(6):1906–1916, 1987.

[162] J. B. F. van Erp. Guidelines for the use of vibro-tactile displays in human computer interaction. In *Proceedings of Eurohaptics*, pages 18–22, 2002.

[163] J. B. F. van Erp. Presenting directions with a vibrotactile torso display. *Ergonomics*, 48(3):302–313, 2005.

[164] F. Vega-Bermudez, K. O. Johnson, and S. S. Hsiao. Human tactile pattern recognition: Active versus passive touch, velocity effects, and patterns of confusion. *Journal of Neurophysiology*, 6(3):531–546, 1991.

[165] R. Velázquez, E. Pissaloux, M. Hafez, and J. Szewczyk. A low-cost highly-portable tactile display based on shape memory alloy micro-actuators. *Pro-*

ceedings of the 2005 IEEE International Conference on Virtual Environments, Human-Computer Interfaces and Measurement Systems (VECIMS 2005), pages 121–126, July 2005.

[166] R. T. Verillo. Psychophysics of vibrotactile stimulation. *Journal of the Acoustical Society of America*, 77(1):225–232, 1985.

[167] R. T. Verillo and G. A. Gescheider. Effect of prior stimulation on vibrotactile thresholds. *Sensory Processes*, 1:292–300, 1977.

[168] F. Vidal-Verdú and M. Hafez. Graphical tactile displays for visually-impaired people. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on [see also IEEE Trans. on Rehabilitation Engineering]*, 15(1):119–130, March 2007.

[169] C. R. Wagner, S. J. Lederman, and R. D. Howe. A tactile shape display using RC servomotors. In *Proceedings of the Symposium on Haptic Interfaces For Virtual Environment And Teleoperator Systems*, pages 354–356, 2002.

[170] Q. Wang and V. Hayward. In vivo biomechanics of the fingerpad skin under local tangential traction. *Journal of Biomechanics*, 40:851–860, 2006.

[171] Q. Wang and V. Hayward. Tactile synthesis and perceptual inverse problems seen from the viewpoint of contact mechanics. *ACM Trans. Appl. Percept.*, 5(2):1–19, 2008.

[172] Q. Wang, V. Hayward, and A. M. Smith. A new technique for the controlled stimulation of the skin. In *Proceedings of the Canadian Medical and Biological Engineering Society Conference*, September 2004.

[173] Q. Wang, V. Levesque, J. Pasquero, and V. Hayward. A haptic memory game using the STReSS^2 tactile display. In *Proceedings of CHI 2006*, pages 271–274, 2006.

[174] L. Ward. Multidimensional scaling of the molar physical environment. *Multivariate Behavioral Research*, 12:23–42, 1977.

[175] P. Wellman, W. Peine, G. Favalora, and R. Howe. Mechanical design and control of a high-bandwidth shape memory alloy tactile display. In A. Casals and A. T. de Almeida, editors, *Experimental Robotics V*, volume 232 of *Lecture Notes in Control and Information Science*, pages 56–66. Springer Verlag, 1998.

[176] F. R. Wilson. *The hand: How its use shapes the brain, language and human culture.* Random House, 1998.

[177] L. Winfield, J. Glassmire, J.E. Colgate, and M. Peshkin. T-PaD: Tactile pattern display through variable friction reduction. *Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, pages 421–426, March 2007.

[178] B. Winzek, R. Vitushinsky, and S. Schmitz. Tactile graphical displays by thin film composites with shape memory alloys. *Proceedings of First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'05)*, March 2005.

[179] J. O. Wobbrock, J. Forlizzi, S. E. Hudson, and B. A. Myers. Webthumb: Interaction techniques for small-screen browsers. In *Proceedings of ACM UIST 2002*, pages 205–208, 2002.

[180] G.-H. Yang, K.-U. Kyung, M. A. Srinivasan, and D.-S. Kwon. Quantitative tactile display device with pin-array type tactile feedback and thermal feedback. *Proceedings of IEEE International Conference on Robotics and Automation (ICRA'06)*, pages 3917–3922, May 2006.

[181] P. Yang. Electroactive polymer actuator Braille cell and Braille display. US Patent No. 2005/6881063, 2004.

[182] L. Yobas, D. M. Durand, G. G. Skebe, F. J. Lisy, and M. A. Huff. A novel integrable microvalve for refreshable Braille display system. *Journal of Microelectromechanical Systems*, 12(3):252–263, 2003.

[183] T. Yoshioka, B. Gibb, A. K. Dorsch, S. S. Hsiao, and K. O. Johnson. Neural coding mechanisms underlying perceived roughness of finely textured surfaces. *The Journal of Neuroscience*, 21(17):69050–6916, 2001.

[184] F. W. Young. *Encyclopedia of Statistical Sciences*, volume 5, chapter Multidimensional Scaling. John Wiley & Sons Inc., 1985.

[185] Y. Zhou, O. Masahiro, and O. Miyaoka. A tactile display presenting pressure distribution and slippage force. In *Proceedings of the Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 281–285, 2008.