



VCU

Virginia Commonwealth University
VCU Scholars Compass

Theses and Dissertations

Graduate School

2012

Development of a Multiple Contact Haptic Display with Texture-Enhanced Graphics

David Burch
Virginia Commonwealth University

Follow this and additional works at: <https://scholarscompass.vcu.edu/etd>



Part of the [Biomedical Engineering and Bioengineering Commons](#)

© The Author

Downloaded from

<https://scholarscompass.vcu.edu/etd/2762>

This Dissertation is brought to you for free and open access by the Graduate School at VCU Scholars Compass. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of VCU Scholars Compass. For more information, please contact libcompass@vcu.edu.

Development of a Multiple Contact Haptic Display with Texture-Enhanced Graphics

This dissertation is submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at Virginia Commonwealth University.

David Spencer Burch
Bachelor of Science, 2006
Master of Science, 2008
Cranky Hermit, 2020
Department of Biomedical Engineering
Virginia Commonwealth University

Director: Dr. Dianne Pawluk
Associate Professor
Department of Biomedical Engineering

Virginia Commonwealth University
Richmond, Virginia
May 2012

Table of Contents

| | | |
|-------|---|-------------------------------------|
| 1 | Introduction | 1 |
| 2 | Background | 7 |
| 2.1 | The Haptic System | 12 |
| 2.1.1 | Neurobiology | 12 |
| 2.1.2 | Psychophysics | 16 |
| 2.2 | Dynamic Display Devices | 26 |
| 2.3 | Pattern Representation on Haptic Displays | 28 |
| 2.3.1 | Texture Patterns | 28 |
| 2.3.2 | Vibratory Patterns | 29 |
| 2.3.3 | Mixed Spatial and Vibrotactile Patterns | 29 |
| 3 | Display System Development | 31 |
| 3.1 | Specific Design Criteria | 32 |
| 3.1.1 | Design criteria based on tactile sensation | 35 |
| 3.1.2 | Other Design Criteria | 38 |
| 3.2 | Stylus Prototype | Error! Bookmark not defined. |
| 3.2.1 | Choice of System Components | Error! Bookmark not defined. |
| 3.2.2 | System Design and Method of Operation | Error! Bookmark not defined. |
| 3.2.3 | Discussion | Error! Bookmark not defined. |
| 3.3 | First Glove-based Prototype | Error! Bookmark not defined. |
| 3.3.1 | Device Testing | Error! Bookmark not defined. |
| 3.3.2 | Subject Testing | Error! Bookmark not defined. |
| 3.4 | Second Glove-based Prototype | Error! Bookmark not defined. |
| 3.4.1 | Device Hardware Components | 62 |
| 3.4.2 | Device Case Design | 70 |
| 3.4.3 | Control Program Design | 73 |
| 3.4.4 | Device Testing | 76 |
| 3.5 | Development of the “Texture” Feedback | 91 |
| 3.5.1 | Evaluation of Frequency, Waveform and Signal Modulation | 94 |
| 3.5.2 | “Texture” Patterns of Mixed Temporal and Spatial Dimensions | 101 |

| | | |
|----------|---|------------|
| 3.5.3 | Shift to Application-related Experiments | 112 |
| 4 | Evaluation of the Display System Performance | 133 |
| 4.1 | Image Creation and Representation | 135 |
| 4.1.1 | Image Set Groupings..... | 136 |
| 4.2 | Single-Finger Experiment..... | 139 |
| 4.2.1 | Method | 139 |
| 4.2.2 | Participant Population..... | 140 |
| 4.2.3 | Results..... | 140 |
| 4.2.4 | Discussion..... | 142 |
| 4.3 | Single- and Three-fingers Experiment..... | 142 |
| 4.3.1 | Method | 143 |
| 4.3.2 | Participants..... | 144 |
| 4.3.3 | Results..... | 144 |
| 4.3.4 | Discussion..... | 146 |
| 4.4 | Using Multiple fingers and Multiple hands Experiment..... | 148 |
| 4.4.1 | Pilot 1 | 149 |
| 4.4.2 | Main Experiment | 150 |
| 5 | General Discussion..... | 159 |
| 5.1 | Approach..... | 159 |
| 5.2 | System Design | 161 |
| 5.3 | Testing of Main Hypotheses and System Evaluation | 166 |
| 5.4 | Conclusion | 171 |
| 5.5 | Support..... | 172 |
| 6 | References | 173 |
| 7 | Appendix..... | 183 |
| 7.1 | Device Prototypes | 183 |
| 7.1.1 | Models..... | 183 |
| 7.1.2 | Device Circuits..... | 186 |
| 7.2 | Experimental Data | 193 |
| 7.2.1 | Device Testing | 193 |
| 7.3 | Developing Texture Feedback | 203 |

| | | |
|-------|------------------------------------|-----|
| 7.3.1 | MDS Experiment | 203 |
| 7.3.2 | Spatio-Temporal Patterns..... | 204 |
| 7.3.3 | Nonsense Objects..... | 205 |
| 7.4 | Experiments with Real Objects..... | 206 |
| 7.4.1 | Image Sets..... | 206 |

Abstract

DEVELOPMENT OF A MULTIPLE CONTACT HAPTIC DISPLAY WITH TEXTURE-ENHANCED GRAPHICS

By: David Burch

A thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at Virginia Commonwealth University.

Virginia Commonwealth University, 2012

Director: Dr. Dianne Pawluk, Associate Professor, Biomedical Engineering

This dissertation presents work towards the development of a multiple finger, worn, dynamic display device, which utilizes a method of texture encoded information to haptically render graphical images for individuals who are blind or visually impaired. The device interacts directly with the computer screen, using the colors and patterns displayed by the image as a means to encode complex patterns of vibrotactile output, generating the texture feedback to render the image. In turn, the texture feedback was methodically designed to enable parallel processing of certain coarse information, speeding up the exploration of the diagram and improving user performance. The design choices were validated when individuals who are blind or visually impaired, using the multi-fingered display system, performed three-times better using textured image representations versus outline representations. Furthermore, in an open-ended object

identification task, the display device saw on average two-times better performance accuracy than that previously observed for raised-line diagrams, the current standard for tactile diagrams.

1 Introduction

For the more than 3.4 million people [1] in the United States who are blind or visually-impaired, accessing visual information is often problematic. Many modern informative mediums, like the Internet, textbooks, and magazines, are increasingly dependent on visual images and graphics to present information. The paradigm that a picture is worth a thousand words is definitely an evident impetus in this transition, and undoubtedly visual imagery can convey information, especially new or unfamiliar information, with greater power than words alone. Kevin Carter's 1993 picture: *A vulture watches a starving child*, illustrated the gravity and enormity of the situation in South Sudan with tremendous pathos. For such images, a word description alone would struggle to evoke the same emotional response. Word descriptions are also an inadequate replacement for visual images that expose people to unfamiliar information or contain a lot of spatial information, both of which are difficult to describe fully in words..

In addition, rendering the information with words excludes the possibility of independent discovery or critical analysis of the information presented, as the image has already been interpreted by someone prior to the visually impaired person currently reading it. Precluding independent discovery and critical analysis of the information may put that person at a disadvantage in advancing in the workplace or educational environment compared to their peers with unimpaired vision, especially in disciplines like science and engineering, where these skills are fundamental.

Additionally, text descriptions can be inaccessible or impractical for some people and in some environments. Written descriptions traditionally require both a translation

and printing in Braille, and assume a knowledge of Braille by the person, which is surprisingly uncommon [2], but less so in education and work environments. Fortunately though, contemporary Text-to-Speech software enables greater accessibility by not requiring Braille literacy, but still is not appropriate in all environments. For example, in a classroom, the additional auditory noise generated may distract or disrupt others, while wearing headphones may isolate the student from class instruction, as well as socially from their peers. While text and oral descriptions are still perhaps the most pragmatic solution for the majority of time, these other situations may still potentially leave individuals who are blind or visually impaired at a disadvantage compared to their sighted peers. In addition, using auditory feedback precludes access by individuals who are deaf-blind.

For these reasons, alternative methods to render visual images and graphics are necessary. These methods all involve some form of sensory substitution, with the auditory and haptic systems being the most common, and relatively far more practical than our sense of taste, smell, pain, or balance, for reasons which are self-evident. The principle method of using auditory feedback, other than the speech descriptions previously mentioned, involves sonification, which provides nonspeech sounds to relay information to the user. Sonification, which regularly is used in areas where a high-demand is placed on the user's attention, has been applied to the field of diagram creation for individuals who are blind and visually impaired [3]. Unfortunately, sonification, suffers from the same disadvantage as the text-to-speech programs, as it isolates them from their regular environment.

This work centers on finding a solution to the problem of rendering 2-D graphics for individuals who are blind or otherwise visually impaired using the haptic system. Along with not having the disadvantages of sonification, as mentioned above, this is also expected to be a potential solution for those who may also have hearing impairments (deaf-blind). Towards this goal, a rough hierarchy of design factors was established to direct the development. The foremost factor is the idea of maximizing the use of the chosen sensory system; i.e., within the haptic system, how do we optimize usability?

The most important thing is for the reader to realize haptics is no more vision than are the auditory or olfactory senses. Both vision and auditory allow us to sense a wide or global field using a single physical dimension, electromagnetic or mechanical waves, respectively. Haptics has a much smaller sensory field, practically the size of a hand or two, but allows us to sense many physical properties, and it does so quickly. Just as familiar sights, sounds, smells, and tastes are immediately recognizable, so too are familiar materials. The hard coolness of iron, the roughness of sandpaper, the smoothness of silk—these properties are also immediately recognizable with our sense of touch. Further, it does not matter if a single finger is used or all them; identification is not slowed by the larger field of sensation. Nor does one have to attend to each point of contact separately.

However, haptics is not as well suited to certain tasks as vision. For example: reach into your bag or pocket without looking and find your keys. The haptic sense allows individuals to quite efficiently discern their keys from their wallet or mobile phone using the differences in gross shape, material properties, and weight between these objects. Had the task been to find a specific key among a group, say their house key from their

office key, then they would need to feel finer spatial details to make the identification, something which the haptic system would not find so easy, but vision would. The difference between the two is not that haptics cannot sense and integrate fine detail, but that unlike vision, haptics processes this information serially, whereas it can process the material property differences between their keys and their wallet in parallel.

This example highlights the idea that by using what is known about the neurobiology and psychophysics of the haptic system, one can establish what would be the natural response; however, no one has developed a single system with actuation that can fully utilize the skin's potential, or even come close. What is clear though, is that any system designed should be a dynamic display; the prevalence and accessibility to digital media via the internet demands a refreshable display. In order to truly provide equivalent accessibility to this visual information, individuals who are blind or visually impaired need a display that can also access this wealth of information without being tethered to unnecessary hardware like a printer.

In addition, there are several design considerations that must be taken into account which are based on the target group. Ultimately, anything built will potentially be purchased by someone, and they will weigh the purchase by its opportunity cost to them. This consideration was actually the second most important design consideration as most individuals who are blind or visually-impaired live in poverty in the United States [4]. Through the focus group sessions with the blind community, they indicated a rough cost limit of \$500, or about the cost of a tablet, laptop, or desktop computer, would be accessible to them financially. Pragmatically, this places considerable constraints on

what hardware options that could be used for the system, and in turn, how much of the haptic system could be utilized.

Other usability and accessibility factors that played a smaller role in the design process were safety and comfort, ease of use and intuitiveness, portability, compatibility, and maintainability. Safety and comfort are not only important for the acceptance of the device, but also for ethical reasons; any device made, especially one made for a vulnerable population, should have nothing more than a miniscule chance of causing harm. Furthermore, the device should not cause the user any unreasonable level of discomfort; they should never regret, or worse, fear using the device, because if they do, then the community at large will never embrace it. Likewise, ease of use and intuitiveness are important as this system will most likely see use in education, and a complex system will daunt children who might use it. Difficulty of use will not bode well for the system's adoption by older individuals as well, who might have sufficient difficulty adapting to use even a computer, much less a frustratingly intricate computer peripheral. Portability and compatibility increase the utility of the device by easing its transference to multiple environments with possibly multiple devices and/or operating systems. One of the major difficulties with any computer peripheral device during the 1990's through 2000's were they rarely compatible with both Windows and other operating systems—a difficulty that any new device should not revisit. Further, individuals who are blind or visually impaired should be expected to buy one system for work or school and another for home, simply because the system is too cumbersome to transport back and forth. Finally, maintainability is an important feature for any device; the lifespan of a device is dependent on the target market. While many people have

adopted a replace rather than repair mentality, this may not be true for individuals who are blind or visually impaired, due to their financial constraints [4]. Thus, the device should remain easily serviceable, so that relatively unskilled people are capable of making repairs, minimizing the need to have the system shipped off for long periods for repair work.

In the following sections, first the relevant background information on the neurobiology and psychophysics of the haptic system will be described, focusing on the research exploring how individuals interact with both real objects and 2-D information, and other devices that have been developed will be mentioned. After that, the development of the own device presented will be described, through its various iterations, and it will be shown how initial research, as well as the related psychophysical research, has shaped its design. Finally, the experiments that validate the system design will be dissertated, specifically, how the chosen method for haptically rendering visual images is better than current methods.

2 Background

The pervasiveness of visual imagery and media in the Twenty-first century is an inescapable fact, one that has even redefined what it means to be literate. The Association of College and Research Libraries declared that this “visual literacy empowers individuals to participate fully in a visual culture”. They define visual literacy as being the ability of an individual to “to effectively find, interpret, evaluate, use, and create images and visual media” [5]. Individuals who are blind or visually impaired have a difficult time achieving and maintaining visual literacy due to the intrinsic challenges of limited to total vision loss. For them, equivalent access to the visual imagery and media has largely been inadequate; although, some groups, such as the American Printing House for the Blind, the National Center for Accessible Media, and the Descriptions for the Blind, have worked diligently to provide increased levels of access.

As mentioned in the Introduction, the access to the information provided for them is not always equivalent to the visual image; often, the access is a language (written or speech) description of the imagery or media. In many cases, a description alone is adequate for the exchange of information, but such access is not equivalent, as it does not utilize, or contribute to the development of, an individual’s visual literacy. In other cases, providing a description precludes individuals from independently distilling key information from an image and learning new concepts for themselves. These two outcomes of this limited access have very negative effects: the first outcome isolates individuals who are blind and visually impaired from the visual component of their

society's culture, while the second one places them at a tremendous disadvantage educationally and job-wise to their peers without visual impairment.

Traditionally, those who have sought to provide equivalent access have used methods of creating a tactile diagram: a representation of the image that the user feels, rather than sees, to obtain information. Raised-line drawings, embossers, silk-screen, thermoformers, and multi-material tactile experience diagrams are all common methods used to create these tactile representations [6]. Each of these methods involve different machinery and processes, but have several common features: (1) they all require unique materials and a fair amount of time in order to be made; (2) they all produce a static, physical representation of the visual image; and (3) they all generally require the image to be presented in an intermediate form, simplified, with the object(s) of interest shown on a proper scale.

The previously mentioned methods all require a physical alteration of a material to produce the tactile image. This requires either semi-permanent deformation, as occurs with raised-line drawings, embossers, silk-screen, and thermoformers, or addition of new material, as with tactile experience diagrams. Standard printing paper does not maintain deformation well, or when it does, it tends to weaken the paper enough where it will quickly warp and tear. Thus, raised-line drawings and embossers require paper unique to their processing method. Raised-line paper consists of paper with embedded alcohol-containing microcapsules; the image is printed using an ink-jet printer, and then passed through a heating element, such as a Tactile Image Enhancer™, which causes the microcapsules painted with ink to swell and burst, creating a swelled, or puffed effect (hence the names, puff or swell paper). Embossers employ a simpler method: they use a

heavy cardstock paper and stamp small holes in it to create the deformed surface. Both the swell paper and the heavy cardstock are substantially more expensive than regular printing paper. Thermoformers and silk-screen both require a mold of the image, and either melt a plastic sheet, or stamp and aluminum sheet, respectively, to create their tactile image. While the latter two methods can more quickly create an image, once the mold has been made, the cost of the mold and the press, not to mention the space required for many molds, exclude its use to only printing house facilities. Lastly, tactile experience diagrams require a maker to manually cut or shape different materials, collage them together to render the parts or objects within visual image, and glue them all down to usually a cardboard backing. This can require considerable time and resources in terms of the different materials, but while it may have the highest per picture cost, it has the lowest startup cost. Unfortunately, all these methods only produce a static image, a tactile representation of a single visual image, which cannot be created quickly for dynamically changing environments and analyses, cannot be scanned and “reprinted”, cannot be attached via email, or stored in a digital format, and will deteriorate with time and need to be remade [6]. A dynamic display capable of working with digital media files, however, does not have these limitations, though its cost is dependent on the design choices made. Another difficulty is that the process requires a sighted person to make the diagrams for the user.

Currently, there really is not a good way to automatically convert most visual images into a form readily available for tactile representation; this is a problem not only with static tactile diagrams, but with the display device presented here. Vision has a much finer spatial resolution and much higher bandwidth for information than our haptic

system [7], so the amount of information presented in a tactile diagram must reflect this greatly reduced bandwidth. Some [8-10] have created programs to automatically translate a visual image into a tactile-ready one, called the TACTile Image Creation System (TACTICS) [8,9]; later, Krufka and his colleagues [10] created another program that used vector graphics to perform the same task. Both of these methods reduce a picture down to an outline representation, as is necessary for generating a raised-line or embossed drawing; however, this strategy has not shown much merit in other studies.

Many researchers found [11-14] that presenting visual information through outline representations (i.e., raised-line and embossed drawings) can be problematic: identification rates from 25% [13] to 33% [11] are typical, but the identification rate approaches 100% [15] for real objects. Bai and Kennedy [16] have mentioned that the poorer performance may be due to the ambiguity in between lines within an outlined representation, as lines representing the border for a part may occlude other portions, indicate part-projection into three-dimensional space, or have some other meaning. The viewing perspective taken by the two-dimension image projection itself may also alter the image, so that an individual' mental model of what is felt does not fit well with their *a priori* model of what the image should look like. This would indicate people have two difficulties with outline representations: (1) determining object/part identity, and (2) determining the orientation of parts relative to the image projection.

However, these two problems are not observed when using vision; the image set used to generate raised-line drawings for [11] came from a set of highly recognizable images [63] when using vision. As previously mentioned, the bandwidth for haptic sensory information is far less than it is for vision [15, 17], but also, the field of view for

vision is quite large compared to haptics, which for processing touch is practically limited to the individual's fingertips. The limited field of view for haptics, in turn, requires sequential exploration in order to process the whole image, whereas with vision, the entire picture can be processed in parallel.

Yet, when the field of view for vision is reduced to approximately the same sized window as it would be for touch, then the performance gap between the two sensory systems disappears [17]. However, when Loomis and his colleagues [17] increased the field of view for touch they found no effect on performance using raised-line drawings. Here, the performance gap cannot be explained entirely by the difference in the two fields of view. Instead, the sensory feedback obtained from outline image representations forces the haptic system to process the information serially, regardless of the available field size. However, Klatzky and her colleagues [12] did observe an improvement from using 1 finger versus 5 freely moving fingers for 2-D raised line drawings: both an approximate 13% increase in the percent correct, and a 17% decrease in the response time. They hypothesized that the difference in their results were due to the conditions being more extreme (i.e., 5 vs 2) and allowing the fingers to move *freely*. This stark difference between these cases illustrates the haptic system's complexity when processing information.

To better understand this complexity, the known neurobiology and psychophysics of the system will be briefly discussed, with emphasis placed on information processing by the haptic system and research performed on visual image representations. Then, the related devices and display methods will be discussed to draw comparisons and contrasts between these devices and the display system presented.

2.1 The Haptic System

2.1.1 Neurobiology

2.1.1.1 *Tactile Mechanoreceptors*

Haptics as a sensory system is perhaps most easily described as a functional combination of three sensory subsystems: our tactile, kinesthetic, and proprioceptive senses. Working in unison, these sensory systems allow us to obtain geometric and mechanical (i.e., material) properties and perform many motor tasks without additional visual input. Within each sensory subsystem there are one or more sensory receptors types that respond to either forces exerted externally on the body, or forces exerted internally within the body. These receptors vary in the range of spatial intensity (magnitude per unit area) or temporal intensity (magnitude per unit time) that they are most sensitive, their receptive field sizes, and their location in the body.

Tactile sensing in the human hand, the part of the body most relevant to my research, is most often attributed to four main mechanoreceptor endings: Merkel cells, Meissner's corpuscles, Ruffini endings, and Pacinian corpuscles. Merkel cells are located in the basal layer of the epidermal folds that project into the dermis. They are the terminal receptor of unmyelinated afferent nerve fibers that slowly adapts (SA) to sustained indentation, and have noted sensitivities to points, edges, and curvature. They are densely populated (100 per cm²) within the fingertip and individual afferent fibers, called SA1 fibers, which have receptive fields between 2 to 3mm [18]. However, their receptive field size belies their spatial acumen; they have a maximum spatial resolution of 0.5mm, though 1-2mm is more common [18, 32]. When the stimulus becomes smaller than the receptive field

size, a single terminal branch of the axon becomes dominant, accounting for the finer spatial resolution through sensitivity to strain energy density, rather than synaptic inhibition as in the visual system.

A couple of potentially important details about Merkel receptors that concern device interaction are 1) they respond linearly to skin deformation up to indentations of approximately 1.5mm [19], and 2) they are sensitive to angled orientation to approximately 5 degrees [20]. These two mechanisms of the Merkel receptors allow them to sense object curvature very well, and they are the only receptors that show the ability to discriminate curvature based on their evoked response [18]. Additionally, Merkel receptors are affected very little by changes in scanning velocities (velocity of the hand when tactually exploring an object) up to speeds of 80 mm s^{-1} , and are at least an order of magnitude more sensitive to dynamic than stationary stimuli [18].

Meissner's corpuscles are ovoid structures that are located at the epidermal-dermal junction in the dermal papillae that differ from Merkel receptors in several key ways. First, Meissner's corpuscles are insensitive to static to very-low ($<0.1 \text{ Hz}$) frequency indentation, and indentations beyond 0.4mm, which makes these afferent fibers fast adapting (FAI) [18]. Secondly, they are more sensitive to lower frequency vibrations ($<50 \text{ Hz}$) and light, dynamic pressure, perhaps in part due to their close proximity to the epidermis. Thirdly, they are also more sensitive to dynamic indentation than Merkel receptors and respond uniformly over their entire receptive field of 3 to 5mm, and hence they have much poorer spatial resolution than Merkel receptors. They are also more densely populated in the fingertips than Merkel receptors at 150 per cm^2 [18], so that many afferents will respond to a local stimulus.

Ruffini endings are spindle-shaped structures located within the subcutaneous connective tissue, and are most sensitive to skin stretch, rather than indentation like Merkel (SAI) afferents. They are less densely populated in the hand than either SAI or Meissner-type (FAI) afferents, and have receptive fields in the order of centimeters rather than millimeters. The Ruffini endings are believed to be responsible for the SAI afferents; however, Hopkins [21] suggested that long chain formations of Merkel receptors may form von Frey hair-type structures capable of the response seen in SAI afferents, based upon studies on the glabrous skin of the hand on monkeys, which have no observed Ruffini endings. SAI afferents have been classified as being large-field, slowly adapting fibers that are sensitive to stretch and also most sensitive to vibrotactile frequencies from 100 to 500 Hz [18].

Pacinian corpuscles (PC/FAII) are multilayered, oval-shaped structures located in the dermis and the subcutaneous fat layers beneath, and sense primarily vibration, but also pressure, and also rapidly adapts to stimuli. Unlike Meissner's corpuscles, FAII receptors are located much deeper, are far less densely populated, so that their receptive field sizes are very large (up to several cm) and have poor spatial resolution [18]. They are also sensitive to displacements up to 10 nm at 200 Hz [22], but sharply (~60 db/decade [18]) attenuate lower frequency vibrations. It is this sensitivity, which is about hundred times greater than that of the FAI receptors, that makes the FAII receptors largely responsible for sensing tactile stimuli transmitted through a probe [18], an important observation to keep in mind when designing any haptic interface.

As mentioned before, PC receptors are more sensitive to higher frequency than lower; in fact, they have u-shaped threshold curve as a function of displacement versus

frequency that peaks in the 200-250 Hz range, and then decreasing sensitivity as the frequency continues to increase, eventually becoming insensitive around 500 to 800 Hz [22, 23].

2.1.1.2 Kinesthetic Mechanoreceptors

The muscle-spindle organs, along with the previously discussed Ruffini endings (SAII afferents), contribute largely to our sense of kinesthesia [24], as it relates to haptic exploration. There are two types of muscle spindles: the primary-type is sensitive to both changes in muscle length and the velocity of change; the second type is sensitive only to changes in muscle length. Thus, position can be sensed from both types, whereas movement is only sensed by the first type. Muscle spindles can also encode movement history by altering their basal firing rates depending on previous movement and muscle contraction. However, muscle conditioning, which usually occurs with changes in muscle activity as in exercise, can lead to errors in the perception of limb position [25]. Another interesting note is that the primary muscle spindles are also sensitive to vibration; vibration presented on the elbow distorts the perception of movement of that joint [26]. Unlike with muscle conditioning where only the perception of position is effected, vibration can distort the perception of movement as well.

Ruffini endings are believed to contribute to the perception of both position and movement via their sensing of skin strain over a joint, as stretching the skin can create the illusion of movement [24]. Perception of skin strain as a mechanism for sensing movement is perhaps most important in the finger joints, where the controlling muscles are located in the forearm, with the connecting tendons crossing over several different joints [27].

2.1.2 Psychophysics

2.1.2.1 *Perceptual Channels*

Psychophysically, the contribution of the various mechanoreceptor afferents to the perception of touch is still not that clear. Physically, the FAI afferents are most effective at signaling sudden, small forces on the hand, and functionally provide the brain with information regarding slippage and grip control. The SAI afferents are responsible for obtaining information regarding shape or form, as well as many aspects of texture, and are sensitive to motion, with a near-flat response as a function of stimulus frequency [18]. The Pacinian corpuscles are the most sensitive to vibration, especially at higher frequencies, and thus, are responsible for conveying information transmitted through an intermediary object, such as a probe. Finally, the SAI afferents contribute to the perception of direction of a deforming force that produces skin stretch; however, FAI afferents can provide signaling regarding motion as well.

The sensory information from the mechanoreceptors from the early psychophysics studies led researchers to formulate a duplex model of mechanoreceptor [27, 28], which claimed there were two channels: one capable of temporal and spatial summation, and the other not. The first of these channels was identified as corresponding to the Pacinian receptors; however, the other channel was unable to be linked to any specific receptor, and was simply called the non-Pacinian, or NP channel [29]. Later, the response for the NP channel was divided into three channels: NPI, NP II, and NP III [30]. These frequency response curves were at first considered to be matched the four observed receptor-type afferents, as shown in **Figure 1**; however, this has not yet been verified [31]. Further, how these channels interact has also not been worked out; it is believed [31] that the

threshold response is determined by the channel with the greatest sensitivity, and that the suprathreshold response is dependent on the integration of the separate channels activated.

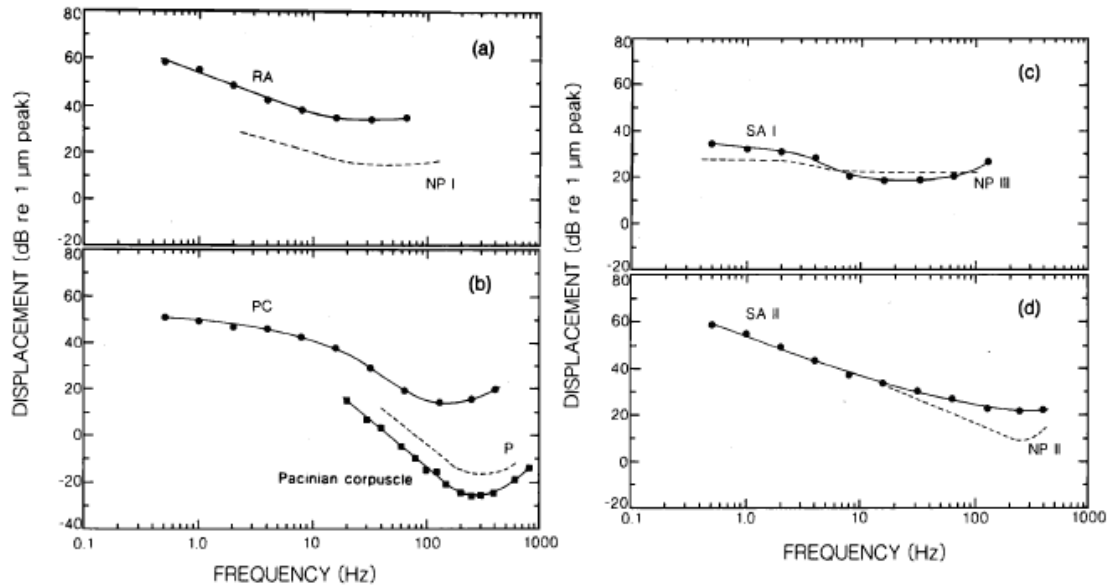


Figure 1: The comparative recorded physiological response (RA, PC, SA1, SA2) to the P and NP channels. Printed in [30].

2.1.2.2 Spatial Acuity

The spatial acuity of the tactile system has traditionally been found using a two-point discrimination test: for static deformation the acuities are reported as between 2-4mm [32]; however, using methods of a two-alternative force choice evaluation, a lower threshold of approximately 1-2mm is observed [18]. Two-point discrimination as a function of frequency is notably higher than that observed statically, ranging from 2.5mm to 5mm for frequencies between 50 Hz and 700 Hz (**Figure 2**) [33]. Other experiments ([34]) have shown that spatial acuity decreases with increasing temporal frequency, with the most likely explanation being that the integration of two or more channels, or receptor types, interferes with the perceptual primarily responsible for fine spatial resolution.




Figure 2: The two-point discrimination (in mm) as a function of pulse frequency. Printed in [33].

2.1.2.3 Pitch and Loudness

As the haptics system perceives touch along physical dimensions of both frequency and amplitude, to pitch and loudness in the auditory system. Some researchers (e.g., Hollins and his colleagues) refer to these parameters using the same terms. As found by [23], for a constant amplitude, the perceived intensity of a vibration increases with frequency, and the two dimensions interact perceptually [35]. The perception of pitch, however, as a function of amplitude and frequency remains subjectively variant [36, 37], so that no single method to describe perceived pitch by its ratio of frequency and amplitude would be effective for all people. Therefore, the concepts as loudness and pitch are not independent perceptual variables, similar to sound [114]. This may be due to vibrotactile stimuli being encoded at the cortical level by their energy, rather than by independent dimensions of amplitude and frequency [38].

Prolonged exposure to vibrotactile stimulation notably results in a decreased perceptual sensitivity, called vibrotactile adaptation. Adaptation on either the P channel or the NP channel(s) does not always result in adaptation on the other [39], as long as the

other channel is sufficiently insensitive to the stimulus. The extent, in terms of both lowering sensitivity and duration of the adaptation is proportional to the channel's initial sensitivity for that stimulus [39]. Additionally, the FAI afferents have been found to be more sensitive to increases in the adapting stimulus' amplitude than the FAII afferents [40], which has been suggested to be caused by differences in their ion-channel mechanisms.

2.1.2.4 Perception of Real Objects

When individuals touch something, they can encode, within as few as 200ms, many of the physical properties: compliance, coarse local structure/inertial and surface material properties for that object from what is known as a haptic glance [41]. When contacting a real object, people sense this information independently across multiple fingers, with the additional fingers causing no increase in response time [42], suggesting that parallel processing of information occurs during this glance. Among the properties that may be encoded, texture and hardness integrate together very quickly, and a combination of the two improves response time, when the information is redundant. Addition of planar contour information, however, is slower, and adding it redundantly offers no further improvement [43]. Local structure, which along with mass distribution, affects the object's inertial properties, can be encoded vary quickly, but finer spatial properties, such as contour geometry, are processed using serial processing [42], and thus, may not exhibit any added benefit from multiple fingers.

As I mentioned previously, people can quickly (around a second) and accurately (near 100%) identify many common objects [15]. Considering that this does not give someone enough time to make a detailed trace of contour geometry, objects, if they're

small enough to be hand-held, are identified primarily by factors that are encoded relatively early on, such as the compliance, surface texture, and coarse shape/structure, and this encoding can occur presumably in parallel [42, 43].

However, surface material properties, such as texture, are not always processed in parallel, depending on the task. Kappers and her colleagues found that in both a search-task [45] and a subitizing task [46], that parallel processing is observed only in cases where the different stimuli are highly salient from each other. Likewise, Lederman and her colleagues [42] also showed that less discriminable features were processed serially, and as a result, additional fingers provided no advantage in the task.

Adding constraints that inhibit perception of specific physical dimensions, such as a splint to disable hand-enclosure of the object and remove coarse spatial information, does not majorly reduce object identification, as long as spatial and kinesthetic feedback was unconstrained [44]. However, when restricted to using only a single end-effector with limited tactile feedback, either one finger in a rigid sheath or a probe, identification accuracy decreased by 50% and exploration time increased by more than 800% [44].

This suggested that the multi-modal three-dimensional information (i.e., compliance, surface texture, and coarse shape/structure) obtained from real objects is critical to object identification. In contrast, relying only on single-point contact vibration and contour information obtained through a rigid sheath or probe is grossly insufficient compared to the bare-finger response. While fine spatial detail, such as contour geometry, is often needed to make higher order, or supraliminal judgments about object identity (is it a house key or a car key?), integrating this information occurs more slowly [42], suggesting they are processed serially, rather than in parallel.

2.1.2.5 Perception of 2-D Information

In the beginning of this section, the issues concerning tactile diagrams and previous research performed using outline diagrams and real objects were discussed to show just how far short the former fall compared to the latter. In the section immediately proceeding this one, how the kinesthetic, proprioceptive, and tactile systems (i.e., surface material, and fine spatial information, as well as parallel versus serial processing of information), contributes to how individuals identify real objects were discussed. When three-dimensional information is removed, it is similar to placing constraints on kinesthetic and inertial information; this leaves surface material properties, and spatial contour as being the primary sources of information from which perceptual judgments are made.

Of the major surface material properties: texture, luminosity, reflectance, thermal and electrical conductivity, and hardness/compliance, only a few can be reliably actuated, and even fewer can be actuated portability and for little cost. The perception of texture does play an important role in haptic identification [42]; it can help classify an object or part of an object into dimensions of smoothness/roughness, softness/hardness, and slipperiness/stickiness [47], or some other comparable adjective [48]. In both of these experiments, the multi-dimensional scaling showed that the smooth/rough dimension explained the bulk of the perceived difference across materials.

2.1.2.5.1 Roughness and Vibrotactile Perception

Research has shown [49, 50] that the perception of roughness can occur through both passive (object touches you) and active (you touch the object) haptic sensing. The characteristics of roughness can be defined by two variables: interaction with the surface

and surface constitution. Surface roughness is dependent on many factors: the user's interaction with the surface; primarily, the force of contact and the relative speed between the surface and skin, groove width between indentations, and the size of the indentations [51, 52]. From this, surface roughness can be divided into two subcategories of features that define roughness overall: coarse and fine. For coarse roughness, groove width strongly, and ridge width weakly contributes to the overall perception roughness [51], as long as these features have a spatial period of 0.2mm or greater [52]. For roughness created by small texture features ($< 0.2\text{mm}$), vibration is the primary means of encoding, and perception of these textures are subject to adaptation, as are all vibrotactile feedback [53]. While sharp transitions, such as edges do contain higher frequency components that are encoded using "fine roughness", adaptation of these components do not significantly affect the perception of coarse roughness [53]; instead, coarse roughness is intensively coded through skin deformation. In some cases of indirect contact with the surface, such as when a subject uses a rigid probe to sense the surface, vibration does play a role in the perception of coarse roughness [53].

However, vibrotactile signals alone should not be considered "roughness" by themselves, merely part of it. Vibrotactile signals are encoding by temporal frequency of the stimulus, whereas roughness is more defined by groove width and applied force. For vibration, integrating complex vibrotactile signals compromised of multi-tone (or frequency summation) notes may increase errors in perception, whereas simpler notes with fewer components are more salient [54]. Depending on the conditions, this may be due to masking, where a more perceptually dominate signal interferes with a weaker one, adaptation, where sensitivity to one or more components is diminished, or response

competition, where the complexity of the signals, or the temporal congruency of their presentation strains the individual's attention.

Adaptation to stimuli for small textures ($< 0.2\text{mm}$), and presumably vibrotactile “textures”, will occur and prolong for a duration dependent on both stimuli intensity and channel (i.e., P or NP) sensitivity [39]. Furthermore, selective fatigue only occurs within the same channel and not across multiple channels. Selective enhancement due to temporal or spatial summation only occurs within the PC channel [56], as summation is not known to occur for the other channels. Temporal masking effects are also selective (i.e. only observed within channels, but not across them), suggesting some isolation, rather than integration, exists across the mechanoreceptor populations.

Another issue is vibrotactile masking, which is similar to adaptation in that it affects perception within a single channel. Temporal masking can occur when a stimulus with a lower threshold is temporally collocated near another signal (e.g. a 250Hz vibration is presented temporally near a 400Hz vibration); the closer the masking stimulus is presented to the signal stimulus, the greater the masking of the less-sensitive vibration. Tan and her colleagues have found [55] that the Weber fractions for signals presented along with maskers will decrease as the signal amplitude increases. They also found that amplitude discrimination thresholds are higher for lower amplitude signals presented with maskers. The memory of one spatial pattern may also persist and interfere with the perception of the next pattern felt. This temporal masking is greatest when the asynchrony between presentations is less than 100ms, but the memory can persist for up to 1200ms [58]. The reader should note that in the experiment, the patterns consisted of lines and letter-like shapes, displayed on an Optacon, vibrating at 230 Hz.

In addition to temporal masking at a single site, masking across fingers can occur as well. Masking can blur the distinction between multiple signals across fingers, preventing individuals from attending to more than one finger at a time; however, research suggests [57] that vibrotactile signals can be attended to simultaneously across two fingers, but there is much less of an attention deficiency when using two fingers on opposite hands.

The interference in identifying target information may be due to factors of response competition, rather than masking; the theory is that patterns presented to same spatial location creates a burden on the individual's attention, rather than interfering with early processing of the tactile information [59]. However, both can occur; the suggestion is that the number, complexity, and time between presentation for the patterns impacts their competitiveness and persistence of a response.

2.1.2.5.2 Perception of Raised Line Diagrams

Interpretation of raised-line diagrams remains difficult at best, with previously reported accuracies in open-ended identification tasks being low, 20-30% [11, 13]. Performance does increase when the task is close-ended [10]; however, close-ended tasks do not evaluate free and independent use of the diagram. Magee and Kennedy [60] also found that active exploration of raised-line drawings to have very poor performance. Moreover, they and D'Anguilli and his colleagues found that active exploration of the diagram was not necessary; passive exploration, where a second individual traces the subject's finger along the outline of an image, performs just as well, if not better. This indicates that the identification of outline diagrams is based mostly on the kinesthetic sense of position, integrated serially as the hand and arm explore.

If shape from outline diagrams is processed only through serial integration, then multiple fingers would not provide any additional benefit. Indeed, some research indicates just that; Loomis and his colleagues [17] found no significant difference between one and two fingers bound together. Additionally, Jansson and Monaci [62] had participants track opposite sides of an outline drawing using two fingers and found no difference in performance. However, Klatzky and her colleagues [12] did find a difference between one and five fingers, but here the performance difference may be due to using 5 unbound fingers versus two bound.

Prompted by a goal similar to my own, Thompson and her colleagues redesigned traditional raised-line diagrams to include “texture” [64 & 65] based on some of the work of Kurze [66]. Their first method involved filling in the space between borders, and they saw modest, but significant improvement compared to outline representations [64]. Their later method, however, called TaxyForm [65], involved encoding 3-D object orientation into a 2-D representation of the object. With TaxyForm, diagrams were also redrawn, so that parts were not occluding others from “view”, as sighted individuals might draw them. A “texture” consisting of basic line or fill patterns was added to the interior of parts, to either signify the part’s unique identity or to indicate its three-dimensional orientation relative to the perspective of the image (example in **Figure 3**).

However, their “textures” are no more representative of natural textures as any other synthetic method, yet the simplicity and observed benefit of their method is nonetheless impressive. They found that using TaxyForm compared to visually realistic images, that early-blind subjects had 50% accuracy (versus 12.5% for visually realistic) and late-blind subjects had near 70% accuracy (versus 44%), with some subjects having 100%

accuracy. The smallest gain was with blindfolded sighted subjects, with a 56% accuracy (versus a 50% accuracy for visually realistic) [65]. Not only is this among the highest performance observed with “raised-line diagrams”, their experiment allowed for open-ended answering of object identities, unlike other experiments with similarly high performance [10].




Figure 3: Example of two TaxyForm images from [65], images originally modified from [63].

2.2 Dynamic Display Devices

Unlike static representations of visual images, dynamic displays offer the versatility of rendering many images without the laboriousness of having to physically produce every image on demand. There have been many dynamic haptic displays, both developed and proposed, as well as a few recommended guidelines for their design (for an excellent review/guideline, see [68, 69]).

The Optacon [70], is the most well-known such device, as it was one of the few commercial devices ever available. It uses a camera to obtain the visual information, and renders information on a single-finger distributed display (6x24 pins), vibrating at 230 Hz, and also has the capability of displaying text as well. While a certain portion of the blind and visually-impaired community achieved a practical level of usability with the Optacon, they unfortunately were the exceptions [71]. In addition, it was both expensive and difficult to obtain maintenance.

Since the Optacon, other optical to vibratory feedback devices have been developed for individuals who are visually impaired. The SmartFinger [72] and Fingersight [73] both use an optical sensor to detect some feature in the environment and provide vibratory feedback based on that feature. They have also used a trackpad [72] and a laser [73] to track position. However, both groups looked only at using a single point of contact (single finger), and although they generated virtual textures via waveform synthesis, did not develop purposefully salient feedback for multiple points of contact.

Other systems have been developed using a position sensor to sense the location on a graphic and provide tactile feedback to one [74], [75], [76], [77] or two [78], [79], [80] fingers. However, there are potential limitations in the devices designed for two fingers in investigating the effectiveness of multi-finger use: (1) the fingers cannot be moved independently but are a fixed distance apart and (2) the devices cannot be rotated or, if they can, significant errors due to the rotation are incurred. Both of these issues are potentially problematic as, in discussions with teachers of students who are visually impaired, the stated advantage of multiple fingers is to aid in tracking a line by having multiple contact points along it. However, this can only be done to a very limited extent without free movement of the fingers and hand. In addition, the devices in [79], [80] use passive touch with the location information being sensed with the other hand; how this affects processing of the information is unclear. As already documented in [76] and, to some extent, in [78] as well, the position inaccuracy of the device used in [78], the VT Player, causes significant problems.

The intended outcome of this research is to develop a simple, cost-effective dynamic display system that will allow individuals who are visually impaired to search

the graphic for objects as they typically do with physical (static) graphics. Namely, the device design will ensure that individuals have free movement of all the fingers of one or more hands. This is expected to facilitate both the ease and rate of identification of a picture (and potentially the correctness as well). The feedback generated by the device will be highly salient and texture-like (although, not true textures) to allow for parallel processing of information [42]. Also the device is intended to use the “texture” feedback to encode information about part-identity and part-orientation relative to the viewing perspective, similar to [65], but with a dynamic display device.

2.3 Pattern Representation on Haptic Displays

2.3.1 Texture Patterns

A diverse range of models has been proposed for providing texture feedback. For force-feedback devices, many researchers have used a constraint surface model for textures, allowing the underlying control system to determine the appropriate forces to feed back to the user. These models have consisted of simple 1-D periodic waveforms, such as sine waves and square waves (e.g., [81], [82]) to more complex elements (e.g., [83]). Other researchers have considered more directly controlling both the normal and lateral forces to simulate textures (e.g., [84]).

Several researchers have examined the use of a variety of tactile displays for conveying texture information ranging from single point contact displays using friction to display texture gratings (e.g., [82]) to distributed displays producing normal (e.g., [85]) or tangential displacements (e.g., [86]). The parameters examined for representing textures included spatial frequency, spatial waveform shape and spatial orientation. In addition, Kyung and his colleagues [87] have examined more complex patterns that included

texture gradients and primitive shapes (e.g., circles, squares, etc.). They found that spatial grating orientation is most effectively represented on all types of displays, followed by texture gradients and then patterns using primitive shapes.

2.3.2 Vibratory Patterns

Although the use of vibrotactile patterns has been primarily explored for single actuator vibration displays (e.g., [88]), some groups have considered it on distributed tactile displays (e.g., [74]). Both temporal sinusoidal signals and more complex signals, such as amplitude modulated sinusoidal and harmonic-synthesized (e.g., square-waves), signals have been used. Study results [88], [89] on the discriminability of vibratory signal parameters, such as frequency, amplitude, modulation, and waveform shape have been mixed.

Some studies have found waveform shape to be the most discriminable, followed by frequency and then modulation [88]. Also relevant to the design is research ([90], [88], [91]) that has used modulated vibratory signals. For example, Murray and her colleagues [90] found that variations in both frequency and amplitude modulation produced reliable changes in the perception of signal magnitude, but the differences between the two methods did not reach significance. Additionally, Massie and Salisbury [92] found that both the length of the modulating pulses and their unevenness contributed to the discrimination of the vibratory patterns.

2.3.3 Mixed Spatial and Vibrotactile Patterns

Lévesque and Hayward investigated [93] the use of spatial waveforms of different spatial frequency and direction, as well as dots, to fill regions of interest. They also considered the use of a single vibration frequency to represent edges in conjunction with

the spatial patterns. However, they did not consider a systematic method for using these variations. They also did not mix spatial and temporal patterns as one object or part “fill” pattern. In contrast, the intended design will employ a systematic method for describing mixed spatial, vibratory patterns for use in encoding information into a simplified representation of a visual graphic presented in a dynamic haptic environment.

3 Display System Development

The initial project was to develop a device that utilized haptic feedback to indicate features, such as lines or areas, on a two-dimensional visual diagram. The method of system operation would be that using the device, individuals would explore the diagram workspace; if the sensor detected, either directly or indirectly, a feature, the device would trigger the tactile feedback. A flowchart of this operation is shown in **Figure 4**.

Individuals would cognitively integrate the tactile feedback from the device with the kinesthetic feedback from their hand and arm, building a mental model of the image explored. The ultimate goal would be to produce a dynamic display device with performance better than, or at very least equal to, that seen using raised-line diagrams, as these are still the standard for tactile diagrams.

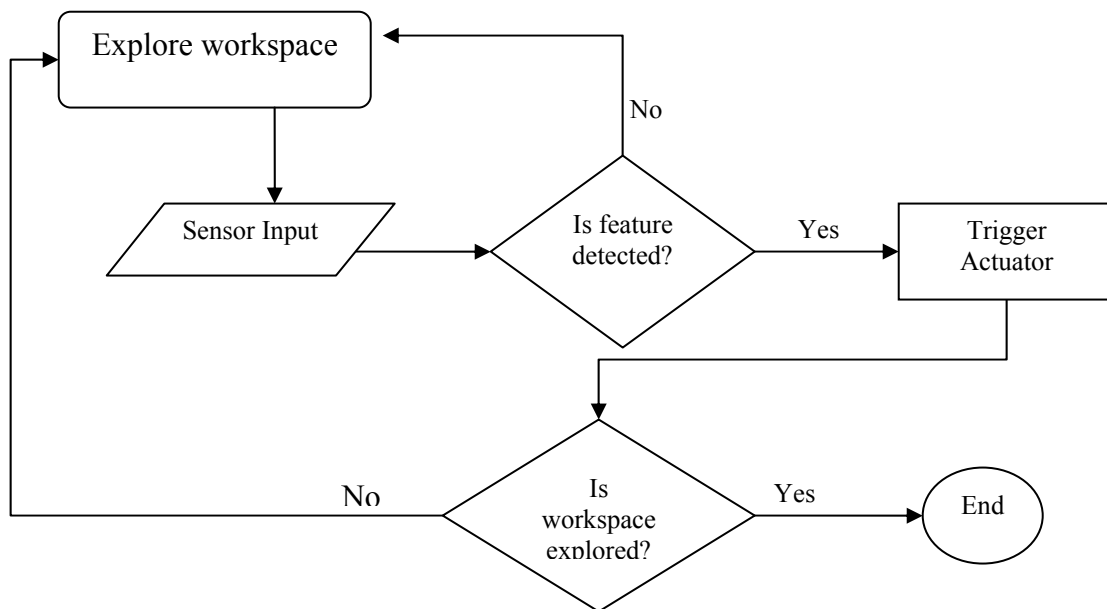


Figure 4: The basic operational flowchart for our system design.

3.1 Specific Design Criteria

In the Introduction, it was explained that the system was designed along several considerations, first of which was an understanding the abilities and limitations of the human component, in this case the haptic perception, because this has the most influence of the system usability.

This section (3.1) will discuss the following topics:

- The design criteria that are based on previous psychophysical literature, including:
 - Spatial resolution, which affects the ability to process fine details
 - Temporal resolution, which affects the ability to faithfully integrate the temporal and spatial aspects
 - Spatial concordance, which affects the how closely the device feedback mimics natural haptic feedback
 - The saliency of the feedback, which affects how well information can be encoded, and influences parallel processing
- The design criteria based on factors such as usability, including cost, safety, intuitiveness, compatibility, and portability.

In the sections 3.2 and 3.3, the engineering design process used to develop the current haptic display device will be described, specifically:

- The choice of system components, including the sensor, actuator, and method of operation
- Device testing, including the evaluation of design factors, such as:

- The psychophysical factors of spatial resolution, temporal latency, spatial concordance
- User safety
- Preliminary subject testing using the device on a few visual diagrams
- A discussion of the overall design, including the other usability topics of cost, compatibility, intuitiveness, and portability

After discussing the design of previous prototypes, the current haptic display design will be discussed in Section 3.4, which will include:

- The new hardware components selected to fix the design flaws with the previous prototypes
 - The new sensor that could be used to detect a wider range of features that would be used to encode the different texture feedback
 - The new actuator that would have a more diverse bandwidth of vibrotactile output for generating texture feedback
 - The redesign of the physical case to allow for greater user comfort
- The device control program to process the expanded range of input and output the new components would require, including a description of the temporal response performance
- Evaluation of device components, specifically:
 - The sensor characteristics and spatial resolution testing
 - The actuator characteristics and its magnitude response across vibrotactile frequency, which involves:

- Development of a compensation filter for the vibrotactile feedback and to obtain a constant subject magnitude for the feedback across frequency

Finally, in Section 3.5 the development of the feedback will be explained, from which a set of highly textures is then developed. This process includes:

- An evaluation of the temporal dimensions of the vibrotactile output, including:
 - Frequency, amplitude, waveform shape, and signal modulation
- Testing the saliency of mixed dimensions using MDS to reduce the number of dimensions from which potential textures would be created, specifically:
 - Temporal frequency and temporal duty cycle
 - Spatial period and duty cycle
- The shift to directly evaluating the performance of potential texture sets to identifiably encode part information. These texture sets were composed of mixed temporal and spatial features from the MDS experiment, and included:
 - Temporal frequency and spatial direction, which used an amplitude modulation of a single tone per part
 - Temporal frequency and spatial duty cycle, which also used amplitude modulation
 - Temporal Frequency and spatial direction, which used a frequency modulation of two tones per part

- A discussion of the experimental results and the chosen texture set

3.1.1 Design criteria based on tactile sensation

The first design considerations originating from the knowledge of the tactile sense are: (1) the spatial resolution that the device is capable of rendering, (2) the temporal resolution or signal latency (i.e., the coincidence of the tactile and kinesthetic information); and 3) the spatial concordance of the tactile and kinesthetic information. Then the importance of feedback saliency when using “texture” encoded information and multiple-finger designs will be discussed.

3.1.1.1 Spatial Resolution

The spatial resolution should match or be as close to human perception as possible to relay geometric information as accurately as possible. As mentioned in section 2.1.1.1, the lower limit on spatial resolution for two-point, forced-alternative discrimination task is 1-2mm [18]. However, the perceptual thresholds given in the literature are mean *threshold* values, with the total population having a distribution of spatial resolution thresholds [94]—dogmatically adhering to a threshold resolution may not yield any improvement in device performance. However, the spatial resolution the device is capable of rendering does affect another aspect: how big a display area needs to be to ensure that the whole picture is provided in the area. A reasonable upper boundary on the size of such an area is for it to be within an arm’s length of all individuals. However, for the sake of simplicity, the size of the display area chosen was the typical dimension of the media used, i.e., a standard piece of paper or a computer LCD screen. This size was not only practical, but previous scientific work on printed images also had similar size restrictions [11, 17, 65]. The reason behind having a size restriction from a human

factors perceptive is that too large of an image, even if comfortably within reach, takes too long to explore haptically, increasing the cognitive load placed on the user. In contrast, with too small of an area, features may be difficult to resolve, especially with a larger spatial resolution. When considering a display area the size of typical media (e.g., 8.5x11 piece of paper), from the perspective of image creation, a maximum line thickness of 5mm is practical: larger than this and some pictures become difficult to draw. This suggests an upper limit on spatial resolution of 5mm and a lower limit of 1mm.

3.1.1.2 Temporal Resolution

The temporal resolution, or time delay, is defined as the time between when the sensor detects a feature and when the actuator generates peak feedback. As the system currently discussed requires active user movement, significant delay can cause users to have a distorted perception of the object's position and shape. The temporal resolution has to be low enough as to not cause any noticeable spatial inconsistencies (or skips) within conditions of typical scanning velocities, which range from 20 to 50mm/s [17]. Unfortunately, every system will have some time delay. What matters however, is at what threshold do users start to notice the delay, and then what temporal resolution is required to meet the desired standards of spatial resolution given estimates for exploratory velocities.

Shogo and his colleagues [95] recommended the maximum allowable system latency to be 40ms when presenting textures, like those that will be used in more advanced device designs. However, at a scanning velocity of 50mm/s, a 40ms latency would create an image distortion of 2mm. Also, too large of latency may cause the device to skip actuation (false negative) for thin-features when users scan at faster velocities. This

means the acceptable limit of latency is also dependent on the device's spatial resolution and the maximum scanning velocity of the user. This limit can be approximated by setting a maximum scanning velocity of 50mm/s and 1mm as the thinnest line use, so then the maximum allowable latency becomes 20ms.

3.1.1.3 Spatial concordance

The spatial concordance between the location of the tactile feedback and kinesthetic information should also remain tightly coupled. When people explore an object with their bare hands, both the tactual information and the corresponding kinesthetic information integrate to give tactile feature information at a given point or area; the further these components are separated, the greater the user's difficulty will be trying to integrate the two. This is of particular importance if the device is to be expanded to multiple point contacts; otherwise, the mental effort involved in correlating the feedback could become very cognitively taxing. Therefore, the point of interaction between the user and the tactile feedback should be as close as possible to the location of the sensed feature.

3.1.1.4 Feedback Saliency

The saliency of the haptic feedback is another extremely important consideration as it is expected to ease the processing burden on the user, which is already very high for interpreting tactile diagrams. Several different results from Lederman and Klatzky's work [42-44] have found that texture, and other material properties, are much more salient than geometric properties for haptics. This suggests that using texture in a display method would increase a diagram's saliency. Further work [42, 96] also found that

processing was relatively constant even when textured information was applied to an increasing number of fingers, implying that it too is processed in parallel, in contrast to geometric information, which is processed in series. Lederman and Klatzky [42] and Kappers [45] also found that as the discrimination between two textures became more difficult, the processing time begins to increase with each additional finger, nullifying the parallel processing advantage of textures. Even with a single-point, user performance will be negatively impacted by poor discrimination, as textures used to describe different features end up feeling similar. Thus, it is expected that feedback saliency will be key in determining the success of the system. However, this does not have an absolute metric in the way spatial resolution or system latency does; instead, the different possible methods of generating feedback along the relative lines of human performance will need to be evaluated.

3.1.2 Other Design Criteria

Other, more generally relevant criteria: cost, safety and comfortable, ease of use and intuitiveness, portability, compatibility, and longevity, which have to deal with system accessibility are also important to define, especially cost, since it was the second most important point for our target population. As of 2005, less than half of the individuals who were blind or visually impaired had employment and the medium monthly income was less than \$1,500 per month [4]. The target “cost” for the device was set to be no more than \$500US, approximately the cost of a computer, as most individuals surveyed in a focus group study at the department for the blind and visually impaired in Virginia indicated that while still expensive, they would probably buy such a device if it worked well. The idea of “cost” however, is ambiguous: one point to keep in

mind is that the development cost of any prototype device will be naturally higher than a per unit manufacturing cost. A counterpoint however, is that any potential manufacturer will add in a marginal cost that includes profit.

An individual's intuition towards a device, and the device's ease of use, is a characteristic of the usability of the device, specifically, how easily learned are the device's functions. This can be measured using several metrics: the performance of users without any training, the time it takes to train individuals to a specified performance level or to automaticity, and through system usability surveys. To limit the time necessary to train individuals, the device should have a low cognitive demand. This can be measured using a system usability survey or task load index, such as the NASA TLX. The System Usability Scale [97] was employed, which uses a 5-point Likert scale to easily assess the system usability.

Another criterion for the device is that it has to be portable, as portability contributes greatly to the device's overall usability. The device's ability to be easily carried from home to the work environment or school is a crucial factor to contributing to the independence of the user, a major issue with individuals who are blind or visually impaired. It also contributes to the accessibility of the device: by being lightweight, it helps children adopt it as a tool they can easily use and carry with them. An added bonus would be if the device is, or could be made stand-alone, requiring no other equipment, such as a computer to work. Even as ubiquitous as computers are in this day and age, it is difficult to always have one in every situation where visual images are encountered.

The final criteria for the device are that it be safe and comfortable to use. The principle safety concern for most of these devices is eliminating the possibility of the user

developing HAVS (Hand-arm vibration syndrome) from using the device for prolonged periods of time. HAVS is brought about by time-dependent, frequency-dependent exposure to vibrations. The stronger the vibration amplitude, the shorter the safe time of exposure before HAVS is a problem. Certain frequencies also present an increased risk for triggering HAVS [98]. However, current research has shown a need to revise these standards, as the frequency weighting proposed does not match experimental data [99]. This will be an important component of selecting the “textures” to be used, not only with regard with their amplitude, but their frequency as well. The comfort of the device is another important consideration, as individuals are less likely to adopt a device that is not comfortable to use. Unfortunately, this design criterion is harder to accomplish when making initial prototypes, as the additional time needed to make rounded edges and polished surfaces can greatly increase the cost of manufacturing. Essentially, as long as the prototype device is judged to be of acceptable comfort to wear for testing, then it is accepted that further comfort related design issues will be addressed at a later point.

3.2 Stylus Prototype

The first prototype was limited in the scope of the design goals to just simply detect borders or outline figures. The reasons for this limitation were 1) the design was mainly to serve as a proof-of-concept that such a device could be cheaply developed, and 2) the device had to be designed within the time and cost constraints originally established for the project. The level of feedback necessary to render a border-image representation only requires binary output, as this is a Boolean operation (is there a border: true or false). As the design would only have two states of actuation (on or off), the sensor did not need to

be overly complicated; it only needed to detect highly contrasting differences (white versus black) in the graphics used.

The device contained: (1) an optical sensor, and a circuit that parsed widely distributed levels of grayscale contrast detection (that could either be on paper or a computer screen) into either a “white” background or “black” features (lines for the most part), and (2) a mechanically vibrating actuator that stimulated the same area of the hand used in sensing when the contrast detected was “black”. A push-button switch acted to interrupt the circuit between the battery, the actuator, and the sensor, to act as an on/off switch. These components were encased in a plastic stylus (an adult vibrator), so that the device operated similar to a pen: the user simply had to press the device against the surface and trace over the graphic, and the device would vibrate when it sensed a feature.

3.2.1.1 Choice of System Components

To maintain close proximity between the kinesthetic information, or the position of the optical sensor, and the tactual feedback from the actuator, all within a single-case, the size of both components were a concern. Several position-sensing systems were examined, such as using an absolute-position detecting RF transmitter/receiver pad, similar to what is used in graphics tablets, or by using dual-axis accelerometers. However, the Wacom Intuos5 medium Pen tablet, the latest graphics tablet as of March 2012, has a cost of \$350. Additionally, though it allows for multi-touch, this is restricted to a subset of gestures, not independent multiple points of interaction. These factors hardly meet two of the requirements for low cost (the tablet alone would cost 70% of allowable expenses before adding vibrotactile actuation), and expandability to multiple fingers.

Dual axis accelerators are far cheaper than a graphics tablet, starting around \$10 (as of March 2012), but would require new position calibration each time they were introduced into the workspace, which decreases intuitiveness if the user has to perform this calibration, or requires additional circuitry to automate the calibration process. Alternatively, an optical analog detector could sense the contrast between “features” and the background of a visual graphic, either on a computer screen or paper. This would be both intuitive to develop and cost effective. A photo-interrupter (GP2L26 by Sharp Electronics) was selected upon the recommendation of Dr. Paul A. Wetzel; the photo-interrupter worked over a wide range of voltages, was miniscule in size ($<3\text{mm}^3$), cheap ($<\$1$ per unit), and contained both a light-emitting diode and a phototransistor in a single unit, enabling the device to work without an additional lighting source. This feature meant that unlike photo-resistors, -diodes, or -transistors alone, the photo-interrupter did not require a photon source such as a light-emitting diode in order to generate a signal.

There are several advantages to using an optical sensor which directly senses the visual image. First, it means that for any given visual image, the position data of its features do not need to be stored in memory, and then compared to the user’s current position data using a sensor like an accelerometer or RF tablet, in order to render the kinesthetic part (spatial location) of the visual image. This conveys tremendous advantages by reducing cost and improving compatibility, by untethering the design from any specific computer processor, memory, or computer software. A computer can still be involved as the medium on which the visual image is presented, but there is no intrinsic need for a microprocessor for the design to work.

One additional benefit is that the visual outline diagram remains visible itself, allowing individuals with low vision to use any residual visual as well to aid in the interpretation the graphic. The visual presentation also allows a sighted observer to see the same information presented dynamically to the user and where the user is pointing. This is expected to allow individuals who are blind or visually impaired to communicate more effectively to someone who is sighted about the information. Using an optical sensor also means that the workspace is restricted by the size of the visual graphic, not the tactile device: an added advantage when examining large paper graphics, such as floor plans.

However, when using just an optical sensor without a software package requires that the visual graphic must be presented in a simplified, or intermediate form, for translation to haptics (i.e., with any conversion to an outline drawing, magnification and/or simplification already performed). This can be done by a software package, or manually by an additional person, but it does require an additional “step” in the process, and limits the independence of the individual who is blind or visually impaired.

Vibrotactile actuation was chosen as the method for producing haptic feedback due to its low cost, strong output, and the ease of its implementation (i.e. it did not require a microcontroller, timing circuitry, or software control) relative to the rest of the device circuitry. To provide the vibrotactile feedback, a solenoid pager motor was selected, as it was the cheapest available vibrotactile actuator, gave relatively high amplitude vibration feedback, and worked with very high repeatability using a binary output control. The main drawbacks of the pager motor were that it had a high power drain, the frequency of the feedback was dependent on the inertia properties of the counterweight, so it could not

be easily changed, and the amplitude of vibrations were too strong and uncomfortable for some people.

The overall device was housed in a stylus design; although a glove type mimics natural haptic exploration, the stylus device was far easier to implement as a proof of concept device, while maintaining temporal and spatial concordance of the tactile and kinesthetic information. A mouse-like device was rejected, as mice often convey a poor sense of spatial location to the user, as the point contact becomes the entire hand rather than a single finger or a stylus tip. **Figure 5** shows the stylus design for the first prototype, though the on/off switch was later moved to the front of the device.




Figure 5: The Stylus Prototype Design

The tip of the device is where the photo-interrupter was housed; behind it (not shown) was the push-button switch that acted as the power switch for the device. The photo-interrupter directly contacted the workspace; however, an acrylic lens (which had a 97% transmission of IR light at 900nm) was placed around the photo-interrupter to protect it from possible damage. The pager was placed motor behind the photo-interrupter, near the area where people would grip the device, to collate the tactile feedback with the

kinesthetic sense of position as best as possible. The circuitry that converted the analog sensor output to a binary control signal, which allowed current to flow through the power transistor and solenoid motor, was housed in the middle section of the stylus, and behind that, closest to where the case opened, was the battery.

3.2.2 System Design and Method of Operation

The GP2L26 by Sharp Electronics photo-interrupter consisted of an infrared photo-diode adjacent to a photo-transistor on a chip. It operated by emitting infrared light from the diode, which hits a nearby surface. The photo-interrupter used has an optimal 0.6-0.8mm range, but works well in the range of 0.4-1.4mm. The photo-transistor then detects the portion of infrared light reflected back off the surface, assuming no other emission of IR radiation by the surface. When enough photons strike the phototransistor surface, they provided a sufficient voltage potential at the transistor's base by the photoelectric effect to turn the transistor on, allowing for current to flow across the collector-emitter bridge. As the potential at the transistor base increased, the amount of current allowed through the collector-emitter increased, corresponding to a drop in the collector voltage. Thus, the drop in collector voltage produced an analog signal of changing voltage potentials that correspond to changes in reflected light hitting the transistor base. The amount of reflected IR light, in turn, was dependent on the absorption of IR light on the surface, which for our application, corresponded to the monochromatic shade underneath the sensor. Technically, there was some difference in reflectivity based on color hue, but saturation along a grayscale provided a more reliable means to base our system.




Figure 6: The control circuitry for the first prototype.

As mentioned before, the analog output signal from the photo-interrupter gets converted to a binary control signal, which then activates a power transistor that is tied to the solenoid motor. Specifically, the sensor output feeds through a series of four NP-diodes to clips the signal voltage and prevents negative feedback issues. The signal then passes to a Schmidt trigger circuit, whose hysteresis is tightly set to improve the spatial resolution of the device, turning on the vibrating actuator more effectively at finer resolution, without causing signal dithering for fine-resolution lines. The Schmidt trigger then fed into a modified Darlington-pair (the addition of resistor 6 shown on **Figure 6** is the modification) to boost the current allowed to the motor.

The single pager motor's counterweight caused the actuator to vibrate at approximately 203 Hz, stimulating mainly the PCs. Although this meant the vibration feedback onto the hand had poor spatial resolution, as the spatial resolution of the tip of the device, obtained by kinesthetic feedback, was the important issue (not that of the tactile resolution in the grasping finger), this was not a problem.

Finally, a 6V RadioShack Lithium battery (2CR-1/2N) was used to drive the circuits, allowing for 4.7 hours of operation in the on position (with the motor drawing 250 mA).

However, the life-time of the device is expected to be much greater than this due to the little amount of current drawn in the off position and the typically short intervals of motor operation during graphic exploration.

3.2.3 Discussion

The end result of the initial project was simply a proof of concept that a visual diagram, when presented in a simplified form, could be rendered solely through the haptic system using a low-cost device. The design satisfied many of its metrics: it was low cost, it was highly portable, highly intuitive, and somewhat compatible, though it did not work well with computer screens due to the IR radiation they emit. The design was made kinesthetically and temporally concordant by providing tactile feedback, based on the graphic, to not only the same hand, but the same part of the hand, as the kinesthetic information being obtained. Using the vibratory actuation of a pager motor had a very short temporal response, was cost-effective and relatively small. However, it did have some disadvantages in that (1) it was not extendable to multiple fingers (2) it could not actuate a large variety of outputs to produce “texture-like patterns” to take advantage of multiple fingers, although this was not a concern at this stage of the development, and (3) produced some discomfort relative to the amplitude of the vibratory feedback. Overall, the prototype achieved its short-term goals, and inspired us to move on to develop more complex, multi-fingered designs.

3.3 First Glove-based Prototype

In order to allow for an expansion to a multiple-finger device, the design was shifted from a stylus to a glove-like design. The selection of the photo-interrupter, the processing circuitry, and the motor remained the same for this prototype; only the device

casing changed. The glove version had three advantages over the stylus design: (1) by placing the sensing elements on the fingertips, it mimics natural haptic exploration; (2) it allows for independent device movement of multiple devices without requiring the user to have to hold them separately; and 3) it enables the possibility for parallel-processing of certain kinds of tactile feedback. The previous design was well-suited for extension to multiple sensor/actuator pairings that could operate simultaneously and independently from one another; however, the added circuitry, and power requirements, could be burdensome. To compensate, the circuitry was offloaded to a small box mounted on the back of the palm; theoretically, this box could also contain a larger lithium battery to power the multiple devices. **Figure 7** shows the glove-styled prototype: the circuitry and battery for the device are housed in a small project box from RadioShack (A), wires output to the pager motor (B), and input from the case housing the optical sensor and the push-button switch (C). Aside from the custom manufactured case for the optical sensor, all the components cost less (\$12) less than a pair of the work gloves they attach to (\$16).




Figure 7: The glove prototype model

In the previous prototype, the stresses placed on the lens during exploration would sometimes damage the photo-interrupter leads or the contact with the push-button switch. To eliminate this, a larger case was designed by Justin Owen to house the optical sensor and push-button switch, which would reduce the shear stresses that occurred during exploration, and which could also be mounted onto a glove. This design is shown in **Figure 8**. The advantages of the design are that it allows for easier access to the photo-interrupter and push-button switch for modification or repair. Also, since the case was made out of a transparent acrylic plastic, it did not require a separate lens as the first prototype did.




Figure 8: The sensor case designed by Justin Owen.

3.3.1 Device Testing

This device was evaluated along several criteria to determine how well it met the design criteria. The criteria of spatial resolution, temporal delay, and maximum vibration amplitude are intrinsic characteristics of the device, independent of any user, and therefore, are described first. Cost was another important criterion; however, it does not require “testing”, so it will be discussed briefly at the end in Section 3.3.6. The final criteria deal with device performance, and for this human participants were needed to properly evaluate the system. The first test performed with subjects dealt with whether

the device caused noticeable adaptation; this data is hard to prove non-invasively, and therefore is largely self-reported. The second test dealt with user performance in identifying simple geometric details using the device, and demonstrated a basic idea of how well a single-finger version might perform.

3.3.1.1 Spatial Resolution Testing

For evaluation purposes, the spatial resolution was defined for the devices as being: (1) the thinnest line for which the device would trigger for 50 out of 50 trials and (2) the smallest separation between two lines that the device could distinguish for 50 out of 50 trials. The spatial resolution characteristics were defined with these two metrics as they are analogous to those referenced perceptually, and are of practical importance to the design of the intermediate visual diagram. The photo-interrupter and system caused some limitations for this evaluation, namely (1) the sensor was rectangular in shape, producing a non-symmetric field of view, so the the horizontal and vertical axes would have to tested separately; and (2) the system was designed to deliver binary output for white and black surfaces, so testing resolution along a grayscale would be beyond the scope of this device.

To evaluate the thinnest-line performance, 10 very-long (>10cm) vertical lines with thicknesses of 0.33mm, 0.50mm, 0.66mm, 0.75mm, 1.00mm, 1.33mm, 1.50mm, and 1.66mm, were scanned using a single sweep at approximately 50mm/s with the device moving perpendicular to the orientation of the lines. The lines were printed at 1200 dpi on a HP LaserJet 4000N laser printer, and then each line's thickness was verified using calipers and magnifying glass to insure uniformity and accuracy. A successful trigger event was only counted if the motor fully revolved. The test was then repeated with the

lines rotated 90°. The minimum spatial resolution that the device would trigger for 100% of the time was found to be the 0.66mm line in both the horizontal and vertical orientations; the device would not trigger, in any of the trials, for the thinner lines.

To evaluate what the smallest separation between two lines that the device would trigger for, lines were made, but this time with varying separation distances and line thicknesses. The test consisted of a set of three lines of equal line-thickness for gap distances of 0.5mm, 0.66mm, 1mm, 1.3mm, 1.5mm, 1.66mm, 2mm, 2.5mm, 3mm, 3.5mm, 4mm, 4.5mm, or 5mm (13 sets with 3 lines in each set), as measured from the line edge to the next line edge. Three variations of line thickness, 0.5mm, 0.66mm, and 1mm, were also tested to see if this affected the spatial resolution (total sets = 39). As before, the device was tested using both horizontal and vertical scans, as this was expected to increase vertical field size affecting the necessary separation. In the test, the device was swept over each set of lines 20 times, again at a speed of approximately 50mm/s. Both the change in the base voltage on the motor's current amplifier (as monitored through an oscilloscope) and visual observation were used to determine whether successful triggering had occurred. The minimal separation needed for horizontal scanning was found to be 1.33mm for 0.5 mm (14 times out of 20 trials: 70%) and 0.66mm thick lines (20 times out of 20 trials: 100%), and 1.5mm for 1mm thick lines (20 times out of 20 trials: 100%). This suggests that there is interaction between line thickness and spatial resolution in the horizontal direction. In contrast, the spatial resolution in the vertical direction was found to be invariant of line thickness: the minimal separation needed for vertical scanning was found to be 2mm for the line thicknesses tested (60 times out of 60 tries). Spacing less than this 2mm separation

caused the motor to operate continuously and there was no reason to doubt that it would operate differently for thicker lines. As was hypothesized before testing, the reason for this behavior likely lies with the differences between the horizontal and vertical fields of view due to the rectangular shape of the phototransistor base, which is longer in length (the axis that lies in the sagittal plane with respect to the rest of the device) than in width (the axis that lies in the frontal plane with respect to the rest of the device).

This would, theoretically make the minimum for a two-point discrimination to be 2.6mm, which was deemed acceptable as compared to the 50% threshold for human perception which was previously stated to be approximately 1-2mm. [See **Appendix 7.2** for test sets]

3.3.1.2 Temporal Latency

The time delay was evaluated between the optical sensor detection of a feature and the corresponding trigger of the motor for it by directly measuring the asynchrony using an oscilloscope to observe the output of the photo-interrupter on one channel, and the output of a small dual-axis accelerometer glued directly onto the motor casing on a second channel. The motor casing was attached on the glove so that the shaft was allowed to rotate freely and uninhibited by any connecting wires or the accelerometer. The device was then scanned over a large black area to trigger the motor, for which the time delay could be measured by comparing the difference between the peak voltages (going from off to on) or baseline voltages (going from on to off). For the evaluation of time delay, the peak voltages were compared, as this directly corresponded to when the feature was fully detected, and when the motor reach full-strength vibration. **Figure 9** below shows the Agilent DSO1002A oscilloscope screen during a measurement of the

device triggering from an off-state to an on-state: the top waveform is the output of the accelerometer attached to the motor and the bottom waveform is the output of the photo-interrupter.

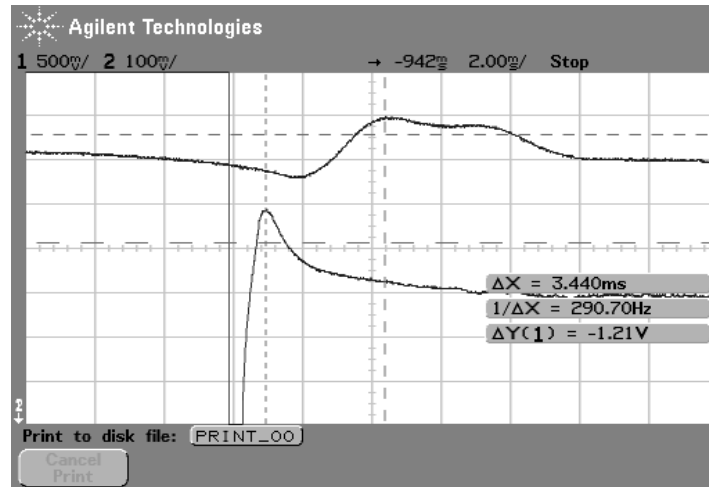


Figure 9: Measuring the temporal delay using the oscilloscope.

The average time delay for ten trials was 3.4 milliseconds, with a variance of 0.7 ms: well below my criteria of 20ms; observationally, this latency did not generate any noticeable spatial inconsistencies for typical scanning velocities up to 100mm/s. Both the optical sensor and the accelerometer have negligible delays associated with them, with the larger of the two delays, the accelerometer, being on the order of a tenth of a millisecond according to its data sheet. [See **Appendix 7.2** for data]

3.3.1.3 User Safety

In terms of safety, the prototype was a low power electronic device with the active components encased in a plastic enclosure on the back of the hand. The device generated vibration feedback using a small, commercially available pager motor. However, there was a potential risk of hand-arm vibration syndrome (HAVS), so testing for user safety

remained an important concern. Symptoms of HAVS include: chronic tingling and numbness in the fingers, not being able to feel things properly, loss of strength in the hands, and the fingers going white (or paler). Some symptoms appear immediately, and others after a few months to a few years of chronic exposure. Not only does HAVS contribute to user safety, it also worsens system performance through loss of sensation.

The safety of the device prior was evaluated to using it in any preliminary subject testing. The evaluations were made in terms of the International ISO 5239-1 Standard [98] for determining the safe exposure to vibrations in the work place. The acceleration of the motor was recorded using an Analog Devices ADXL203 Dual Axis Accelerometer (the same one used to calculate time delay), which was mounted on the lateral side of the glove near the interphalangeal joint, approximately 1 cm away from the motor. The accelerometer was oriented so that the two axes lay in the plane of the motor's rotation, perpendicular to the motor's shaft. Formal data was not collected for the third axis, which would lie parallel to the motor's shaft, as acceleration along this axis was informally found to be very minor. The two channel output from the accelerometer was amplified with a signal gain of 20, sampled by a 16-bit PCI 6230 DAQ by National Instruments® at 1kHz, and read by LabView™, where both channels were filtered with band-limiting and frequency weighting filters, in accordance with the International Standard ISO 5349-1. Four trials with 10 samples each were recorded; the mean magnitude of acceleration for the vibratory motor was measured at 2.411 m/s^2 , with the maximum recorded acceleration at 2.420 m/s^2 . The single-axis average accelerations were 1.625 m/s^2 and 1.781 m/s^2 . **Figure 10** shows the pre-filtered Fast Fourier Transform obtained from LabView™, including the fundamental frequency of

approximately 203 Hz and two harmonics; the dc offset shown in the picture should be ignored, as it is a component added by the accelerometer’s zero-g base voltage. For 4 hours of total daily exposure (a high-end estimate of usage), the daily exposure value of our device is 1.697 m/s^2 , using the ISO 5349-1 guideline for calculating magnitude. This is well below the 2.5 m/s^2 acceleration level at which action needs to be taken. Since all of the acceleration frequencies were well beyond the 15.915 Hz frequency mentioned by the ISO as being known to exacerbate HAVS, and the total RMS acceleration magnitude was below the threshold for precaution at any exposure duration (including for the entire work day), the motor seemed to present minimal risk for the user.




Figure 10: The FFT of the accelerometer signal showing the frequency of vibration.

However, during the preliminary testing on actual graphics, 3 out of 7 of the subjects experienced a slight discomfort when using the device. They described the feeling as “a tingling sensation” in the finger and hand to “numbness” in the finger. They further described the numbness felt as a loss of tactile sensation, excluding static pressure. Subjects who experienced the discomfort said the onset came shortly after exploring the large, solid images. Some subjects further went on to say that exploring the outline images gave them no discomfort. None of the subjects experienced the discomfort long after ceasing to use the device (5 minutes or less). It is not clear whether the discomfort

resulted from adaptation to the vibratory stimulus, or from an associated effect of the vibration-related disorder, similar to an acute phase of HAVS. Although the acceleration test data indicated that the vibration magnitude was insufficient to cause long-term problems such as HAVS (using the ISO 5349-1 Standard), the Standard does not rule out the possibility of causing short term discomfort for some users. Some insensitivity to high frequency vibration (i.e., adaptation) was noted with the discomfort, but it was not easily apparent whether it was the main cause of the discomfort, or a side-effect of a more significant problem.

Also, while the International Standards ISO 5239-1 and ISO 5349-1 provide a guide for the maximum daily vibration exposure for an individual by proposing a frequency-weighted algorithm for measuring the acceleration magnitude of these vibrations, our tests have provided anecdotal evidence that they may not sufficiently address the problem. This hypothesis is supported by work of Dong and his colleagues [99] who found that the energy absorption distribution for vibration is dependent on the frequency, with higher frequency (>100Hz) energy being absorbed more locally. Further, they suggest that for frequencies higher than 16Hz, with some exceptions, “the relative weighting is higher than that of the ISO” [99]. This would explain these particular findings, but more testing would be needed to prove it to be the case. Regardless of whether the discomfort resulted from adaptation to the vibratory stimulus, or from a HAVS-like effect resulting from the vibration amplitude, discomfort associated with tactile stimulation was not acceptable for the device.

It was decided best to reexamine the choice of actuator, which led to the new choice of using piezoelectric buzzers. The buzzers have an advantage over the solenoid

motors in that they produce lower amplitude vibratory stimulation; however, this requires them to be placed directly against the skin to be felt, unlike the solenoid motors.

Fortunately, the thin casing for piezoelectric buzzers facilitates such placement, which improves the kinesthetic coupling of the optical sensor-actuator system. The piezoelectrics also enable a wider range of stimuli waveforms, allowing for potentially greater presentation of information. For these reasons, the piezoelectric buzzers were incorporated into the next version of the prototype.

3.3.2 Subject Testing

3.3.2.1 Adaptation

In addition to HAVS, the user's tactile perception was expected to become desensitized over time due to sensory adaptation [100], potentially making use of the device frustrating, or decreasingly effective. In a single subject test, a subject used the device for a period of two hours, with the input to the motor bypassed with a constant voltage so that the motor would stay on. Throughout the two hours, the subject showed no adaptation or discomfort to the vibratory stimulus. However, during a testing with diagrams, three out six of the subjects mentioned some discomfort, with loss of sensation being one of the mentioned effects. Currently, it is not clear whether this resulted from receptor desensitization or from an effect similar to, if not, HAVS. Adaptation remained difficult to assess and predict throughout the experiments. Therefore, to avoid this problem, subjects were required to take short, one-to-two minute rest breaks every ten minutes or so during experiments, and the break-frequency was increased if requested by a subject. Incidentally, the subject who tested for adaptation also participated in the study to identify objects and did not report feeling any HAVS-like symptoms.

3.3.2.2 Preliminary testing with 2-D Graphics

For this prototype, preliminary testing was started using basic geometric shapes and simple mathematical plots. These pilot experiments were intended to provide a more concrete proof-of-concept for the design, and get more application-based feedback to direct further development. Prior to this, no formal testing had been performed using a prototype device on actual graphics; rather, tests had been restricted to informal tests with a few subjects using the device on graphics presented on normal paper and raised line duplicates on swell paper for comparison.

3.3.2.2.1 First Experiment

The first pilot consisted of six subjects without vision loss, who were recruited from within the Department of Biomedical Engineering, age ranging from 20 to 25 years, all of whom volunteered for the test. Two of the subjects had some experience with the device, while the remaining four had no prior experience. The six subjects were taught how the device worked and were allowed to freely explore a 15 cm x 15 cm solid square on a 21.5 cm by 28 cm sheet of paper as “training”. The sparsity of this training was intentional, as it was desired to examine how well participants can performed with only a fundamental understanding of the device’s mechanical operation.

For the first part of the experiment, subjects were blindfolded using a sleep-shade and instructed to name or best describe the shape in picture in front of them. There were three pictures in all, with each picture presented separately. The images in the test included a solid triangle and a solid circle; an outlined square, triangle, and circle; and a parabola and a third-order polynomial curve. The dimensions for each of the outlined and solid images were approximately 15cm by 15cm, whereas the curves were

approximately 22cm in length. The paper they were printed on was standard letter sized (21.5cm by 28cm) and line thicknesses for all non-solid images were 5mm. The subjects were not informed of what pictures were included in the pool prior to the experiment, nor were subjects given feedback about their performance during the experiment. For the latter group of images, determining the overall shape of the curve substituted as a correct answer in place of naming the shape. Although subjects were not given time restrictions during their exploration of an image, no subject took more than three minutes for any single image in the task.

For the second part of the experiment, the subjects were asked to explore one of two x-y graphs (both extended the full length of a 21.cm by 28cm sheet of paper, and had a line thicknesses of 5mm). After a brief time exploring, subjects were then asked questions about which maxima was greatest, which minima was lowest, and which slope had a greatest magnitude. Subjects were not timed for this exercise and typically people took anywhere from three to six minutes to complete the task. After the completion of this task, subjects were prompted to provide any commentary that they wished to give about the device.

The result means are shown in **Table 1**; for the preliminary tests, the subjects had an aggregated average that was slightly higher than 50% accuracy (59.2%). Accuracy was low for the circle in solid and outline representations, as well as the spline-curve, possibly due to the difficulty in perceiving the contour of a curve with the device. A matched-pair t-test was performed ($df = 5$) to compare the two types of images, solid versus outlined, which yielding a p-value of 0.613, with an effect size and Cohen's d (-.16 and -.31,

respectively). Owing to the small sample size of the experiment, the significance of experiment remained dubious.

| <u>Image</u> | <u>N</u> | <u>% Accuracy</u> |
|------------------|----------|-------------------|
| Solid Triangle | 3 | 67 |
| Solid Circle | 3 | 33 |
| Outline Square | 1 | 100 |
| Outline Triangle | 2 | 100 |
| Outline Circle | 3 | 33 |
| Parabola | 3 | 67 |
| 4-Point Spline | 3 | 33 |
| Graph 1 | | |
| Maxima | 3 | 33 |
| Minima | 3 | 67 |
| Slope | 3 | 0.0 |
| Graph 2 | | |
| Maxima | 3 | 33 |
| Minima | 3 | 100 |
| Slope | 3 | 67 |

Table 1. Results of the First pilot on graphics.

3.3.2.3 Discussion

The evaluation of this prototype design showed two principle things: (1) individuals could use the device to render at least rudimentary visual diagrams, and (2) the multi-fingered glove design was at least feasible, in part due to point 1 being true, even if the multi-fingers portion of it was not yet implemented. However, the pager motor was not the best actuator choice. It lacked adequate ability to provide a variety of salient vibrotactile stimuli to help distinguish between different features, and it had the added disadvantage of producing noticeable discomfort among some of the test subjects. The optical sensor also presented issues as the function of the device was expanded to present multiple features through textures, which would mean encoding the different textures in multiple levels of the grey scale. However, the optical sensor was limited in the changes it could detect accurately and consistently. Therefore, the next prototype sought to

address these issues, and the added issue of new, diverse vibrotactile feedback, before being extended to allow for use of multiple fingers/devices.

3.4 The Current Display Device Design

The design describe here represents the third display prototype. The previous designs, had design flaws that prompted many of the current design choices. Among those, the optical sensor and actuator for the new prototype were chosen to enable the use of texture-like feedback. The intended advantage of using textures is based on the observed effectiveness of using tactile-experience pictures as compared to raised-line drawings. This is likely because tactile-experience pictures, through the use of different textures to represent objects/parts, better convey object/part identity or uniqueness over raised line forms.

However, more than this, it was hypothesized that the improved results in using tactile experience pictures stemmed from: (a) the high saliency of the textures, making the information easier to process and (b) the parallel processing of textures by multiple fingers. Work by Thompson and her colleagues [65] suggested that texture could be additionally used to encode information regarding orientation of objects and/or object parts. In contrast, it was believed that in using outline-representations without the texture-feedback, individuals would not perform as well with this less salient information and would not be provided with any benefit from using additional fingers/devices. This set of hypotheses formed the core of this dissertation research. The reader should note that incorporating texture-like feedback into the prototype system required not only changes in the hardware and programming, but also in how the images themselves were represented. In the rest of this section, the former changes will be discussed, and the

discussion of the texture-like feedback and creation of the intermediate “textured” image representation will be left for Section 3.5.

3.4.1 Device Hardware Components

3.4.1.1 The New Optical Sensor

When the system was redesigned to use “textured” feedback, versus the on/off vibrotactile feedback of the previous model, three issues had to be addressed together: (1) how to most effectively create tactile “textures”, (2) how to control the actuation, and finally (3) how to represent the texture information in the visual diagram for translation to tactile textures. The term “textures” is used in parentheses, as the vibrotactile feedback used only represents a small subset of the psychophysical characteristics sensed with real textures; therefore, what is generated by the device should not be confused as being high-fidelity representations of real textures.

When reconsidering the type of sensor to determine the local graphical information beneath the fingertip, optical sensing was kept. This was, most of all, because of the ease of extending the sensing from one finger to multiple fingers. This was problematic for other systems such as using RF transmitters with a tablet. Optical sensing also had the advantages of cost, compatibility, and sensor size, making it the better option. When considering between different optical sensors, a further consideration was the ease of transforming the sensor output into something that could be used to encode the localized graphical information using the “texture” feedback. In the first prototype, this was relatively simple: the photo-interrupter provided analog output easily convertible to a 1-bit output through the use of a single comparator. To provide clearly distinguishable encoding “values” a color sensor was considered.

One alternative was to use a sensor as complex as a CMOS digital camera chip, which has a much higher resolution and color resolution as compared to the previously used simple photo-interrupter. However, the inherent problems with CMOS camera sensors in this type of application versus simpler optical sensors are abundant: higher bit-rates than needed, low frames per second (introducing greater system latency), greater cost, and power. Furthermore, there are still additionally issues with camera optics, such as the lens and more importantly, the lighting and f-stop characteristics. Also, while optical mice use less advanced camera-type optical sensors and low-cost, they still suffer from many of the same disadvantages, plus are designed to detect velocity relative to position, not local color. In addition, as only a single point contact was considered per finger, a distributed sensor was not needed.

For these reasons, the focus was narrowed to optical sensor that directly sensed color, such as a photo-transistor, -resistor, or -diode. Since it was also intended to replace the actuator and a distributed actuator display was an option, other optical sensors, such as a CMOS or CCD photoarrays were also possibilities. Although, the selection of a single-point contact versus a distributed contact display will be discussed in Section 3.4.1.2, presently it will be said that a single-point contact device was selected, so a more complex sensor was not necessary.

All of the possible sensors work directly with the illumination provided by a computer or tablet-device screen. Previously, the optical sensor was designed to work with standard printed graphics, it could have been modified to read from a screen by changing the triggering voltages for the Schmidt trigger. Interaction with a computer was deemed necessary for the design to have the greatest level of access to visual diagrams.

By programming the device to encode a unique texture for every color it is set up to detect, the system is set up to easily render graphics haptically with the use of textures.

The potential sensors evaluated were the BCS series by TDK and the OPT101 series by Texas Instruments photodiodes, as well as the Agilent HDJD-S831-QT333 CMOS chip, the Hamamatsu S9706 Digital RGB Color sensor, and the Hamamatsu S9032-02 Analog Color sensor. The BCS series had two disadvantages: they were primarily sensitive to only one color, requiring an additional inverse transform to normalize the response, and had too large of a receptive field (~4mm). The OPT101 series had a smaller receptive field (~2.5mm), but was primarily sensitive to the infrared range. Since most monitors give off a fair amount of heat, this and the chip sensitive to the IR EM spectrum were very ineffective. The Agilent QT333 CMOS chip also had too large of a receptive field (>5mm), aside from being very heat sensitive, and thus, hard to work with, to meet the spatial resolution desired for the device. The Hamamatsu S9706 Digital CMOS sensor gave 12-bit output for red, green, and blue in a serial output array and had the advantage of two resolutions, low (0.36mm x 0.36mm) and high (1.2mm x 1.2mm). Unfortunately, the sensor required an integration time that is proportional to the log of the luminosity of the light and the log of the number of bits outputted on each channel. Typically, the luminosity for most LCD monitors is between 200 and 500 lumens; using the low value of 200 lumens, the low resolution would require a 25ms integration time to output only 4 of the 12 bits for each color, which would be too long in comparison to the time delay deemed adequate in the design criteria (20 msec).

The S9032-02 analog color sensor contains red, green, and blue color diodes with a total area of $2\pi \text{ mm}^2$, which puts it in the range of the spatial sensitivity desired for the

device. Initially, the device was rejected as having too poor of a spatial resolution; however, the sensor came with a lens that was not needed for our application; by removing this lens and adding a pin-hole aperture to the casing, the sensor obtained a more acceptable spatial resolution. Ultimately, the S9032-02 RGB color sensor from Hamamatsu was chosen as the sensor for the device, as it met the spatial resolution requirement, the analog signal output adds only negligible latency, it is easier to work with by not requiring an additional timing circuit, and costs less per chip (\$5 for the S9032-32 compared to \$19 for the S9706).

The S9032 sensor utilized a unique Labview program to sample the sensor input and then provide actuator control dependent on the local color sensed. This was not seen as a permanent solution; rather a temporary design choice that allowed me to quickly change many variables in software, rather than hardware, including the calibration for the optical sensor input. The intended design would use an embedded controller, or digital logic circuitry, contained in a case on the back of the hand. However, as this hardware was not necessary to test either the hardware or the human performance of the system, its development was postponed for some undetermined time in the future.

3.4.1.2 The Haptic Actuator

Due to the limitations and the possible safety issues concerning the use of the solenoid pager motor, other actuator types that provided more than a single monotonic vibration and had better control of the vibration amplitude were considered. Two classifications of tactile actuators were considered: actuators that had a single contact point localized over an area of skin (point-contact) and actuators that had multiple contact points distributed localized over an area of skin (distributed). The reader should note that

point-contact actuators do not necessarily prevent the display from incorporating multiple-fingers, only that each finger would have only a single point of feedback, as in the previous prototype. The main considerations for weighing the two options were (1) the haptic perception of the feedback, (2) the cost, and (3) the size.

The possible perceptual differences between distributed and point-contact actuation have not been widely looked at. Jansson and Monaci looked at using one or two fingers to explore tactile maps [62], and Loomis and his colleagues [17] to identify raised-line diagrams of common objects, with both finding no significant difference. Lederman and Klatzky [44], and later Jansson and Monaci [101], independently found constraints placed on bare fingers that restricted use of spatially distributed information worsen identification of real objects, as did reducing the number of fingers. However, contour information like that tested by Jansson and Monaci [62] and Loomis and his colleagues [17] is processed serially primarily using kinesthetic feedback [60], suggesting that either additional points or increased spatially distributed information would provide little advantage over the other. Furthermore, it remains very unclear if the feedback from a rigid sheath on real objects can legitimately compare to that from vibrotactually-generated textures. More equivalent comparisons made with force feedback devices have mixed results between those finding an advantage to using more points [102] and those who did not [103]. Other perceptual considerations are spatial resolution and masking effects. Spatial resolution does worsen as frequency increases [33], though the biggest decrease comes after 200 Hz. Vibrotactile masking can also play a role, and while Craig [57] found less masking with one finger than two on the same hand using a distributed display (the Optacon), he also found less masking using one finger on each hand.

Overall, little can currently be concluded, as no one method perceptually illustrates a clear advantage.

The factors of cost and size do, however, strongly favor point-contact actuation. A single electronic Braille cell with 8 distributed pins currently costs \$62 [104] compared to a few dollars for a small piezoelectric or voice-coil speaker (point-contact). The Braille cell also requires higher driving voltages (180-200V), and has a larger profile than most speakers, making case design more difficult. For these reasons, a point-contact display was chosen, but it was noted that the choice may change once a better comparative body of research exists for the two actuation methods.

Among the point contact actuators, speakers offered great flexibility for producing a wide-range of amplitudes and frequencies, with the two main types being a voice-coil or piezoelectric crystal. Voice-coils are electromagnetically driven linear actuators consisting of a cylindrical coil of wires surrounding a magnetic core. When current passes through the coil of wires, a magnetic field is generated that either attracts or repels the magnetic core, depending on the direction of the flow of current. Voice-coil speakers have the advantage of being able to be driven by lower voltages than piezoelectric materials, and can produce greater displacements than piezoelectric speakers, with the force produced being proportional to the number of coils, the diameter of those coils, and the current passing through the coils. The piezoelectric speaker is a disk-shaped piezoelectric device that deforms in response to an electric potential across it. Its displacement can be controlled by increasing the voltage across the material, by increasing the size of the disk, or by sandwiching multiple disks together to create different layers. It has the advantage of being driven by lower currents and having a

slimmer profile than voice-coil actuators. Piezoelectric speakers and voice-coils both have about equal cost, which ranges from around a few dollars for small, commercially available units to hundreds of dollars for custom made parts. Both voice coils and piezoelectric speakers have no DC response and very limited response for frequencies below a certain range (typically <10Hz).

Twenty different voice-coils and piezoelectric speakers were tested, primarily judging their characteristics of size and maximum output amplitude. A bilayer piezoelectric speaker (Taiyo Yuden) was chosen over the others, due to its thin, flat profile (0.8mm), which allowed the overall device to be smaller, and its stronger vibratory output compared to similarly-sized voice-coils and piezoelectrics. The piezoelectric speaker had an output frequency range from 1Hz to 20kHz, with a resonance frequency peak at 2500 Hz. It could generate sustained output for a large range of voltages, from $1V_{p-p}$ to over $40V_{p-p}$. The thin profile (0.8mm thick) allowed it to easily fit into more compact case designs, making for a less cumbersome and more easily wearable device.

Later, Taiyo Yuden discontinued production of their bilayer piezoelectric; this part was replaced with the Susumu multilayer piezospeaker, due to it being the same size (20mm diameter, >1mm thickness) and having similar frequency response characteristics. At the time of the replacement, several device cases had already been cut, which alone took many months, so a total redesign was a bit out of the question. However, if this had not been, the APS2714S-R multilayer speaker by PUI Audio would have been preferred, as it had greater low-frequency response, which is important for tactile feedback applications, and a stainless steel construction that most likely would have made it a more durable actuator.

The main problems with the actuators used (both models) was that: (1) sufficient pressure applied directly on the surface of the piezoelectric speaker will dampen the vibrations and may cause destruction of the actuator, and (2) the ceramic material does not hold up well against the skin's natural acidity and oils. To mitigate the first problem, the case was designed so that the user's finger rested primarily on the case, with the actuator only contacting the fleshy portion of the central distal phalange. Training on how to use the device also emphasized not pushing too hard on the actuator-contact site. However, this didn't prevent all people, especially those with smaller hands, from applying forces too high for the actuator. The second actuator from Susumu also tended to break more frequently than the Taiyo Yuden piezospeakers, which was possibly due to the lack of a metal-based construction. The use of stainless steel for the PUI Audio piezospeakers, however, may eliminate this problem.

The driving signal for the piezospeaker came from the output of the computer's D/A (max 10V, 5mA and additional amplification came from a circuit, which at first filtered the signal (described in detail in the next section), then amplified it with a gain of two.

The actuator's frequency response within the range below 800 Hz was not provided for by the manufacturer; therefore, the frequency response was tested under unloaded and loaded conditions to characterize both the native response and the application response for the piezoelectric. This was then used as a guideline to build a filter that not only compensated for the lack of flatness of the actuator's frequency response, but also along changing tactile perception of the signal magnitude across frequency.

3.4.2 Device Case Design

As mentioned before, the new design revolved around being more “glove-like” in shape, to allow for multiple interaction points on the fingers. However, the previous prototype, which used a construction glove to serve as a means to mount components, was not very comfortable to wear. For this prototype, individual wearable devices were developed. Two advantages this design had over a glove were (1) the user could wear as many or as few devices as needed without having to alter a glove, and (2) the device would not need different sized gloves, since the device would only be worn on the finger. Different finger sizes were accommodated using elastic band to mount the device case to the finger; for larger variations in finger sizes, several different sized elastic bands were cut and sewed.

The new design also continues to allow for independent device movement for multiple devices, one per finger, without the user having to hold them. This allowed the fingers to have more of a natural, independent exploration of the diagram.




Figure 11: The exploded view of the new prototype design casing.

Figure 11 shows the concept for the third version of the device. The case consists of three layers held together by four screws (only two screw holes are visible). The bottom part of the case has a pinhole aperture for the optical sensor. This acts to increase the spatial resolution for the optical sensor by restricting light transmitted towards the sensor at non-orthogonal angles. Behind it is an additional hole, where a small Dacron button rests; this button gets forced against the push-button switch when the device case is pressed against the screen.

Next, the PCB board containing the optical sensor rests between the top and bottom case parts and is held in place by the screws that pass through it. The PCB board also contained some buffering and filtering circuitry, as well as a push-button switch. The push-button switch activates when the Dacron button gets pressed against a surface; the weight of the device alone is not a sufficient enough force to do this, so it requires active pressure by the user. One benefit of this is that the device will only require power when it is being used. A second advantage is that the device does not falsely activate (because it will turn off) when it is pointed upwards towards another light source, such as a lamp.

The circuitry present on the chip provides line buffering, low-pass filtering, and adjustable gain for each of the three color channels. **Figure 12** shows the circuit diagram [See **Appendix** for larger copy]. The first set of operational amplifiers serve as a buffer for the photodiode and as a low-pass filter with a cut off frequency of 106Hz. The filtering component was necessary to remove high-frequency noise generated by the photodiode and to band-limit the signal such that the second-stage op-amps could properly increase the signal potential. The control program, which will be described in a later section, sampled the signal at 10kHz, so the 106Hz f_c was one-fiftieth of the

Nyquist frequency. These inverting amplifiers (the buffer inverted the original signal, so by having an inverted gain the signal was corrected) provided an adjustable gain for each of the channels, so that their outputs could be calibrated using hardware. The push-button switch provided an interrupt in the piezoelectric speaker line; however, when the device is driven off an internal power supply such as a battery, the push-button switch will serve as an interrupt between the power supply and the rest of the circuitry.




Figure 12: The pre-DAQ circuitry

The top part of the device case houses the piezoelectric speaker, the screw terminals, and serves as the contact surface for the user's finger. An elastic band hooks into a groove located between the bottom and top portions. It holds the finger in place and stretches for many different finger sizes. Additionally, the contact surface for the finger on the device is shaped such that the distal eminence of the fingertip slightly rests on the piezoelectric speaker, but so the remainder of the distal phalanx rests on an inner ledge of the casing. This: 1) provides a better seating of the finger on the device to prevent slippage of the finger in the device as it drags over the surface being explored and 2) prevents the pressure applied by the finger from dampening the vibrations of the piezoelectric speaker. **Figure 13** shows the real device with its separate parts labeled.



Figure 13: The real prototype device. Features: (1) pinhole aperture, (2) push-button opening, (3) RGB sensor, (4) push-button switch, (5) piezo-actuator, and (6) groove for the finger.

3.4.3 Control Program Design

The control program was developed in Labview on a computer. However, this was not seen as a permanent solution; rather a temporary design choice for exploratory purposes in using the device. The intended design would use an embedded controller, or digital logic circuitry, contained in a case on the back of the hand. However, as this hardware was not necessary to test either the hardware or the human performance of the system, its development was postponed to some undetermined time in the future.

The LabView control program was designed to input the signal from the optical sensor, determine the color by comparing the input to pre-determined value ranges for specific colors, and then output a driving signal for the actuator specific to that color. The design stemmed from two separate, earlier programs that were used to determine the optical sensor resolution and control the piezoelectric speaker during device testing. The input from the finger device gets sampled continuously using a NI ADC PCI-6221 16-bit DAQ at a rate of 10 kHz reading 10 samples at time; each of the three colors: red, green, and blue, are sampled separately on their own channel. Each device is also sampled separately; therefore, 5 devices require 15 separate channels, and more devices would

have required a DAQ with more input channels, such as the PCI-6225 DAQ. The signals for each channel are divided up and labeled by device and color, then multiplied by a constant, which helps provide a software calibration.

For each finger device, the three channels for color enter a Case loop within a While loop. Within the Case loop, the three values are compared to the preset color value-ranges; if the three inputs fall within the predefined ranges for a color, set through a separate subVI, the case loop triggers another waveform generator subVI to output the corresponding output signal to the DAC. If the three values do not match a pre-defined value, the Case loop will default to its null case, and halt the While loop until new values enter the loop, which given the sampling rate, occurs within 1ms. Execution of the While loop may also be stopped manually on the Front Panel. **Figure 14** shows the Labview block diagram of the control program.




Figure 14: The LabView block diagram for a single device.

The Labview program generates the output waveforms using the square-wave pulse generator subVI and some array manipulation so that full voltage range of the DAC could be used. Each case loop has an option to create digital output, which allows a 1-bit array output to be sent to the DAQ. This digital signal was used to control a separate border frequency that bypasses the other circuitry; the reason for this will be explained in the Section 4.1. The two outputs for each device are merged and sent to a NI PCI-6230 16-bit PCI DAQ. Initially, there were several problems with the output: (1) it had high temporal latency (approximately 40-50ms), (2) it did not output the correct frequency, and (3) it would interrupt signal generation when the output from one or more case loops changed. The proposed solution from the NI representative, to reduce sample rate and change to on-demand sampling, did not work, nor did using different subVIs to generate the waveforms at specific frequencies. Oddly enough, using a 200kHz sampling rate writing a fixed number of samples at a time, in this case 2000, fixed problems 1 and 3, but not 2. However, then it was noticed that the frequency of the output signal was not determined by the setting on any subVI, but by setting the size of the array to write per iteration of the loop. After some trial and error, it was found that frequency could be determined by the ratio of the number of samples written by the DAQ divided by the array size written by the subVI, and then multiplied by the sampling rate. This improved the total system latency to 20ms as measured on an Agilent DSO1002A oscilloscope.

After fixing the problems, the Labview control program was set to recognize 15 colors: red, green, blue, white, black, gray, purple, aqua, yellow, dark red, dark green, dark blue, dark purple, dark aqua, and dark yellow. Later on, this was narrowed down by removing the darker versions of the colors, as they were not needed to render the textures

we developed. However, the system could, in theory, recognize a wide range of colors, and with better noise reduction and amplification, may possess 8-bit per channel color resolution, similar to an RGB bitmap. This level of color detection, was however, far beyond the scope of the project, and the ability of individuals to discriminate, much less identify, the corresponding vibratory feedback.

This VI was based on previously designed hardware that consisted of three voltage comparators that looked at the output of a single photodiode. This photodiode was the BCS series photodiode by TDK that later was rejected in favor of the S9032-02 RGB color sensor. The three comparators acted as a parallel Analog to Digital converter, generating a 2 bit output, which was then passed to a bank of clock circuits operating at different frequencies. If the digital output was 11, then the logic circuit output one frequency, if 10 then another, if 01 then another, and if 00 there would be no output. Although it was never developed, a similar system for the S9032-02 RGB color sensors, this early project illustrated how digital control circuitry could replace the LabView VI.

3.4.4 Device Testing

3.4.4.1 The Optical Sensor Resolution

The purpose of evaluating the optical sensor resolution was to ensure that sensor performed within the guidelines for acceptable spatial resolutions of 1-5mm, developed from the literature.

To evaluate the spatial resolution for the new prototype, two factors were determined, the thinnest line (absolute resolution) and the smallest separation between two lines (relative discriminatory resolution) that the device can detect. However, this had to be performed not just for black lines, but every color that was to be used. To evaluate the

first part, the absolute resolution, sets of parallel lines were drawn with line thicknesses of approximately 0.3mm, 0.6mm, 0.9mm, 1.2mm, 1.5mm, 1.8mm, and 2.1mm (based on a 0.3mm pixel size), with each line set containing 56 lines, for a total of 392 lines. To evaluate the second part, the relative resolution, for each set of 56 lines, it was subdivided into sets of 8 lines with line separations of approximately 1mm, 1.5mm, 2.0mm, 2.5mm, 3.0mm, 3.5mm, and 4.0mm, with the separation distance measured from the end of one line to the beginning of the next. [See Appendix 8.4.1 for example test]. By scanning the device across these sets and recording the number of times the device triggered, the two resolutions could be evaluated in a single task. This was then repeated the process for select combinations of colors: white, red, blue, green, yellow, purple, and aquamarine lines, both at full and half saturation, against a black background; red and green lines against a yellow background, red and blue lines against a purple background, and green and blue lines against an aquamarine background; and finally red, green, blue, yellow, purple, and aquamarine lines of full saturation with their half saturation values as a background.

The colors used in testing were expected to give the poorest results in terms of spatial resolution due to the poor color contrast between them, not because these colors should be used adjacent to each other. All combinations were drawn using MSPaint; saturation values were manipulated by adjusting the RGB value for the color within the program. Full color saturation refers to a value of 255 for a particular color (e.g. yellow would have a RGB value of 255, 255, 0), while half saturation refers to a value of 128, or half the full saturation level. The ability to detect the correct color was assessed on a true or false basis: the device had to correctly identify the color for every line in a set to be

considered true. A constant scanning velocity was maintained of approximately 50mm/sec, but due to the sheer number repetitions, fatigue was an issue. If the device was able to correctly identify most of the lines (five or more of the eight in the set), but not all, then the set was repeated on the basis that human error in scanning the lines might have occurred. Upon repeating the trial, if the device still could not determine the color than a false statement was given for that test. [See **Appendix 7.2** for results].

It was found that against a black background, the full-saturation colors (red, green, blue, yellow, purple, and aquamarine) have an absolute spatial resolution of 1.2mm and a relative discriminatory resolution of 1.5mm. The corresponding half-saturation colors have a poorer absolute resolution at 1.8mm, but a better relative discriminatory resolution at 1mm. For color combinations with poor saturation contrast, i.e. full-saturation colors against a background of their half-saturation colors, the absolute resolution was 0.9mm and the relative discriminatory resolution was 2mm. For color combinations with poor hue contrast, i.e. colors of equal saturation values that would be adjacent on a color diagram, the absolute resolution is 1.8mm and the relative discriminatory resolution is 1.5mm for full-saturation combinations and 2mm for the half-saturation combinations. Overall, device, even in its worst case operating conditions, still produces reasonable spatial resolution of 2mm for both the absolute and the relative discriminatory resolutions.

3.4.4.2 The Actuator Frequency Response

The bandwidth capabilities and response characteristics were tested to assess whether any compensation would need to be performed for the feedback to be able to be produced with acceptable fidelity. The amplitude and frequency characteristics for the

piezoelectric actuator were measured using a LCL-113G Full Bridge Thin Beam load cell (Omega Engineering), recording the displacement of the actuator surface. One end of the load cell was mounted to a wooden platform using a metal washer and screw; the platform itself was clamped down onto the table. On the other end a small, metallic probe (10 g) was attached to the load cell using a metallic plate and cyanoacrylate adhesive; the probe was then fixed to the surface of the piezoelectric actuator using the same adhesive. The leads from the load cell bridge were connected to an AD521 Instrumental Amplifier with a gain of 1000, where any offset drift was manually adjusted for using 10k Ω trim-pot. The output from the OpAmp was then low-pass filtered using a 2nd order Butterworth with a cut-off at 1000 Hz to remove noise beyond the range of frequencies of interest. Then the signal was amplified again, this time using a single inverting AD544 OpAmp and a gain of 100. The displacement was calibrated by comparing the peak-to-peak voltages, obtained using an Agilent Oscilloscope, for ten 25.4 μ m displacements, which was measured using an Aerospace Dial Indicator. Using smaller displacements would have been preferred to calibrate the load cell readings, but could not be easily preformed with the equipment available. However, the accuracy of the calibration ratio was not very important, as the absolute displacement measurement was less important than the relative change in displacement across the range of amplitudes and frequencies the actuator might be used.

The Susumu actuator response was measured for sinusoidal frequencies from 10 to 500 Hz and input voltages 1, 5, 15, and 20V. A sinusoidal chirp was used to analyze the frequency characteristics. To simulate the various forces a fingertip might apply on the actuator, several weights were added, fixed to the top of the metallic probe, so that the

force transmitted similarly down against the actuator as a finger would. These weights consisted of modeling clay surrounding small sections cut from a metal rod, and were 43g, 77g, and 105g (approximately 1N) in weight. The weighted values were only recorded for frequencies around the range of the resonance peaking, for a commanded output voltage of 15V. The frequency response was also characterized for the Taiyo Yuden actuator. However, the actuator responses could not be honestly compared as improvements were made to the characterization method, specifically the load cell amplifiers and the weights used, when the actuator was switched to the Susumu actuator.

The frequency response was relatively flat in the frequency range of interest, between 10-200 Hz for all weights. For the unloaded case (**Figure 15**), a resonance peak occurred at 350 Hz. The actuator had a very close to linear response across amplitude for frequencies between 10 and 200 Hz (**Figure 16**), with some greater deviation seen at 200 Hz for 20V. The addition of weight had two effects on the piezo displacement: it lessened the response proportional to the weight applied and down-shifted the resonance peak (**Figure 17**). Finally, the mashed-finger case was tested, i.e., having an individual press hard on the actuator/load cell setup, and then find the resonance peak. Due to the individual not being able to apply steady pressure, not to mention the added signal noise generated, the resonance peak could only be roughly estimated to be about 230 Hz, and no further recorded measurements were made, as shortly afterwards, the actuator broke. Overall, it was expected that the applied weight on the actuator to relatively light and sufficiently steady for individuals due to the case design and training.

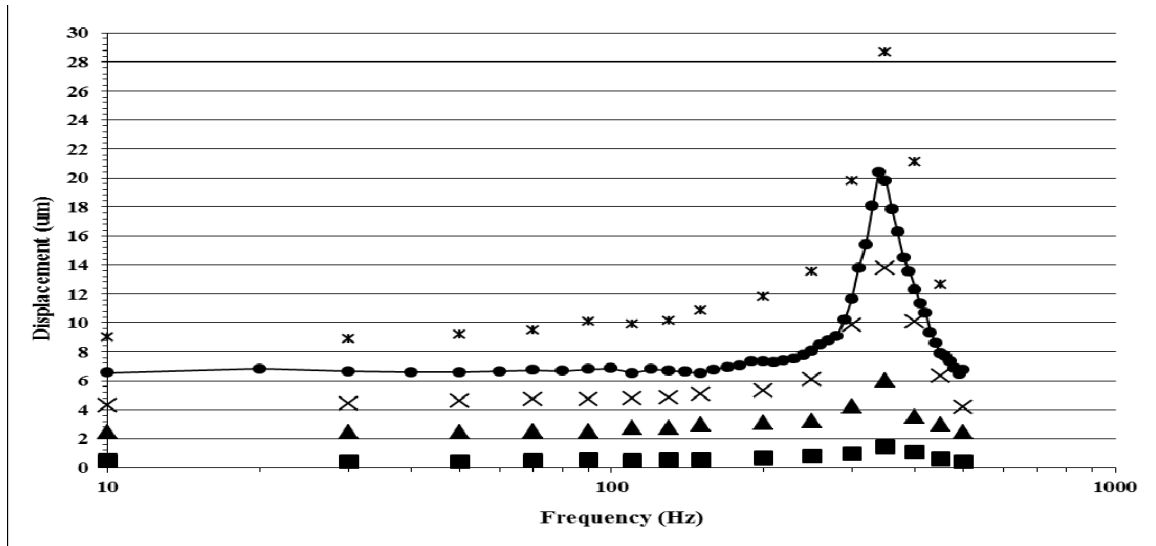


Figure 15: The unloaded response for 1V (square), 5V (triangle), 10V (X), 15V (circle) and 20V (stars). The frequency response was measured every 10Hz for 15V, while fewer readings were recorded at other amplitudes.

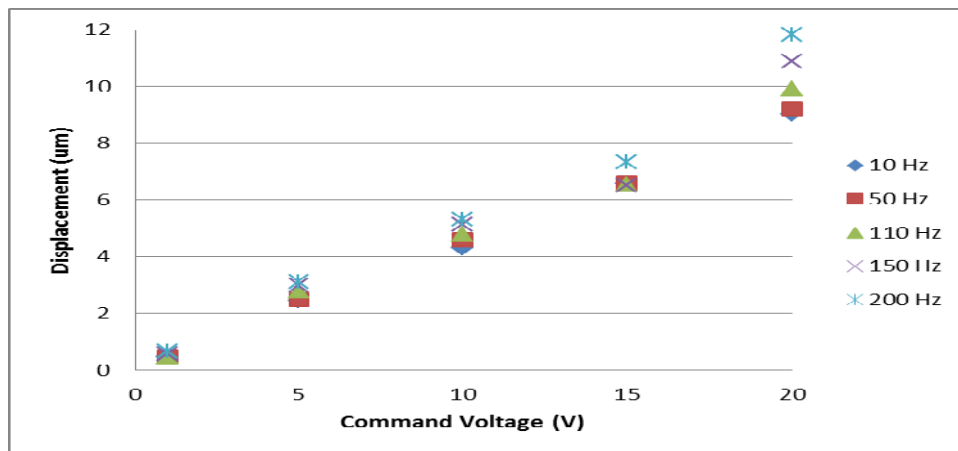


Figure 16: The actuator response across driving voltage. The closeness of the points for a given voltage shows the flatness of the response for this frequency range, while the linear rate of increase ($\sim 0.5\mu\text{m}/\text{V}$) shows that over this range, the actuator has a linear response.




Figure 17: The weighted frequency response for the actuator at 15V. As weight increased, the magnitude of displacement decreased as well, but not linearly. The resonance peak lowered in frequency as weight increased as well, even if the exact peak is not shown.

3.4.4.3 The Perceptual Filter

Since the actuator starts becoming non-flat past 200 Hz, in addition to having a resonance peak beyond 300 Hz, adjustment to the output signal would be necessary. The device was intended to use temporal square-waves to generate vibrotactile feedback due to people’s expressed preference for them over pure-tone signals; however, unintended amplification of certain harmonics would distort their perception, especially those affected by the resonance peak. The easiest solution was to filter the output, either to match resonance specifically, or just low-pass filter the signal beyond 200 Hz to attenuate the peak and all other higher frequencies.

However, there is a further consideration to be made: human perception of vibration amplitude of sinusoidal signals varies with frequency tactually, just as it does in the auditory system. The perceived (subjective) magnitudes of these vibrations vary drastically at different frequencies for the same signal amplitude [23]. This means that when presented at equal magnitudes, higher sinusoidal frequencies can be excessively

salient, while lower frequencies may not be detectable at all. Cholewiak and his colleagues [109] found that for higher frequencies, square-wave gratings were detectable as soon as their third-harmonic components were detectable. In pilot work, few frequencies above 40 Hz could be easily distinguished as anything but a harsh buzzing sensation, presumably due to the poor discrimination between the different harmonic frequencies. This indicated that actuator magnitude equalization would be required to achieve a greater usable range of frequencies.

In addition, adaptation, which is necessary to minimize so the device can be used as long as possible, is proportional to the intensity of the stimulus above the perceptual threshold, which varies as a function of frequency [105]. Without equalization, the excessive perceptual magnitudes at higher frequencies would lead to quicker adaptation for a given perceptual channel.

In adjusting the amplitudes along a curve of equal perceptual magnitude, could either involve (1) decomposing the signal or specifically scale only the component frequencies used, or (2) filtering the output signal along a curve corresponding to an equalized perceived magnitude. Equalization was chosen along a single isosensitivity contour, similar to that used in [100] and in sound equipment, i.e., creating a tactual equalizer. However, unlike Tan's group, it was chosen to use the 15 dB sensation level (15 dB_{SL}) for two reasons: (1) it is within a comfortable operating range [106] and (2) it attempts to minimize adaptation by only using sufficient, rather than excessive, stimuli magnitude. It should be noted that the 15 dB_{SL} isosensitivity contour is not simply the threshold frequency curve/contour, but 15 dB stronger in amplitude. Instead, the 15 dB_{SL} is the *isosensitivity contour* generated when a spectrum of frequencies are perceptually

equalized in magnitude compared to one or more standards, which are 15 dB above their threshold values. The 15 dB_{SL} contour is based on the perception for pure-tones, rather than multi-tonal waveforms, as the greater sensitivity for higher frequencies would distort the low frequency perceptual response.

In order to accurately equalize the subjective magnitude dimension over the desired frequency range (10 to 200 Hz), a pre-filter was used to perform two tasks: (1) low-pass filter to compensate for the actuator characteristics, and (2) perceptual magnitude equalization for the human user characteristics along the 15 dB_{SL}, as described by [23] for 10 to 200 Hz. When determining the perceptual magnitude response, rather than treat the actuator and human tactile system separately as done by [107], they were treated together. Although the resulting compensation system is specific to the device, the methodology used may be applied to other devices as well.

To reduce the complexity of compensation for the actuator, it was decided to treat this as a linear system, which as **Figure 16** in the previous section shows, it is fairly close to being linear. As described previously, the perceptual filter was considered at an isosensitivity contour at 15 dB_{SL} using several participants' data. In developing the perceptual filter, the absolute threshold was first determined across frequency for each participant using the device. Then, those amplitude values were used to determine the 15 dB comparison standards, i.e., the controls, as in [106]. From these perceived magnitude comparisons across the frequency spectrum, the 15 dB_{SL} isosensitivity curve was generated for each participant, and averaged this across participants. Using this curve like a Bode plot, a filter was designed with a matching attenuation, which equalized the signal to obtain a flat “perceptual” response—on average. Lastly, the filter response was

verified using both new and old (i.e., the same ones who aided in creating the contour) participants, to show that people do indeed perceive the different frequencies to be equal, with no more than a ± 3 dB deviation across the 10 Hz to 200 Hz frequency range.

3.4.4.3.1 Absolute Threshold Testing

First, in order to establish a baseline for participants to perform the 15 dB subjective magnitude task, the absolute thresholds of seven participants was determined, across frequencies. The frequencies used were 10, 20, 50, 75, 100, 150, 200, 230, 270, and 300 Hz. The method of adjustment was used [107] to determine the participant's threshold for each frequency. In this method, participants were provided with two knobs on a Labview front-panel, which they could use to adjust the physical amplitude of the actuator's signal. They used one knob to provide coarse amplitude control and the other fine amplitude control, similar to a microscope. Participants were required to adjust the signal to the level that they just noticed it. To minimize participant bias towards any particular knob position, the knobs had no indicator for magnitude. Additionally, at each frequency, a random number was subtracted from the amplitude adjustment control by the Labview program; thus, unknown to the participant, the start point for each frequency had a different, random offset.

Seven participants repeated the task ten times in different sessions, with the frequencies presented in a different randomized order each time. The threshold curves for each participant were averaged, where each threshold value is determined by the arithmetic mean of the 10 repetitions. The data plots for the participant thresholds are represented in **Figure 18**.

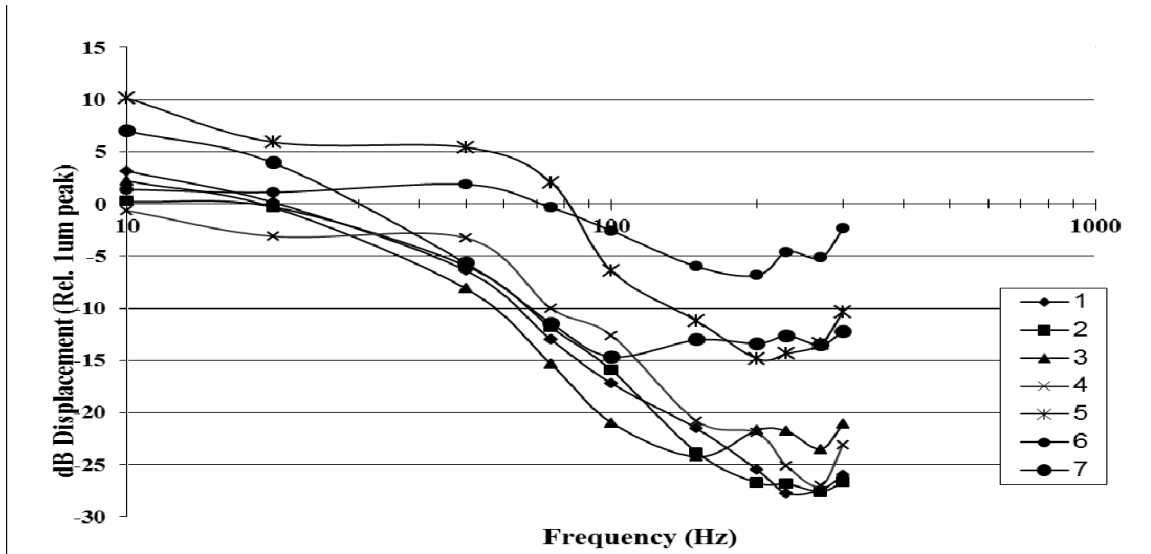


Figure 18: The participants' threshold curves as a function of frequency, relative to a $1\mu\text{m}$ displacement by the actuator.

3.4.4.3.2 15 dB_{SL} Determination

The method of matching by adjustment [23], [107] was used for determining the 15 dB_{SL} isosensitivity contour. This contour, as shown in [23], is not simply a matched contour to the threshold, so amplification of the threshold curve alone would not have sufficed for my purposes. Prior to the experiment, for each participant, their absolute threshold was determined and plotted (previous section). For both parts of the subjective magnitude experiment, the participants were presented with two alternately repeating sinusoidal bursts at specified frequencies. The first burst had fixed amplitude that was 15 dB greater than the participant's previously measured absolute threshold for that frequency; this was the standard that was mentioned previously. The second burst had an adjustable amplitude knob controlled by the participant. The two pulses were both 600ms in duration, and were separated by a 1400ms silent period.

Participants were asked to adjust two knobs, similar to those used on the absolute threshold task, to control the amplitude of the second burst. They were required to stop

when they felt that the bursts had matching amplitudes. Again, the knobs did not have any markings and a random value was added to the knob output to minimize participant bias.

In the first part of the experiment, participants made subjective magnitude adjustments for test frequencies of 10, 20, 50, 75, 100, 150, 200, 230, 270, and 300 Hz, using 50 Hz and 230 Hz as standard frequencies, i.e., all test frequencies were matched against the 50 Hz standard, then participants repeated the task for the 230 Hz standard. For this part of the experiment, the amplitudes of the test frequencies were adjusted in magnitude to match the fixed amplitude frequency standard. Participants repeated the task ten times, five times for each frequency standard; the frequencies were presented in a different randomized order for each repetition.

In the second part of the experiment, participants were asked to then complete the task by adjusting the amplitude of each burst of the frequency standard to match that of each test frequency burst, which was at a fixed amplitude 15 dB higher than the participant's threshold value for that frequency, similar to [23]. Participants then repeated this task ten times, five for each frequency standard.

Five of the seven participants who performed the absolute threshold testing completed the magnitude matching by adjustment task: one participant was omitted due to being an outlier in terms of their absolute threshold curve, while the other participant was unable to complete the second task. For each part of the experiment, the resultant matching amplitudes were averaged over the 10 trials to obtain an average 15 dB subjective magnitude contour (the 15 dB_{SL} isosensitivity contour). Resultant data contours (2 per participant) were normalized with respect to their individual response at

10 Hz, so that all the curves represented their response relative to the same point, the 10 Hz response (**Figure 19**). This data transformation was necessary to obtain an aggregated data set that was not dependent on participants' absolute threshold values. The geometric mean was then calculated, similar to [23], to replicate his process, across participants for the transformed data of both parts of the experiment. Finally, the geometric mean was found for the two resulting contours and then the overall group-mean isosensitivity contour was plotted (**Figure 20**).

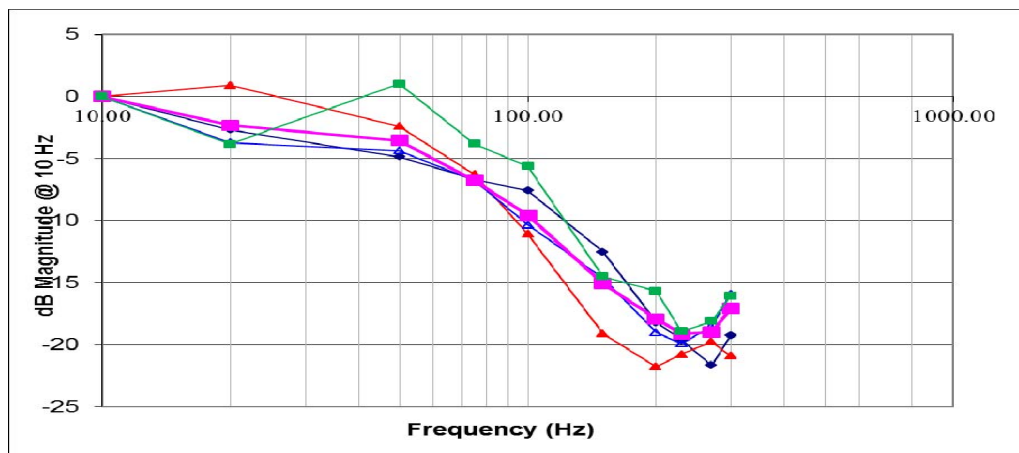


Figure 19: Individual 15 dB_{SL} isosensitivity contours plotted relative to their response at 10 Hz.

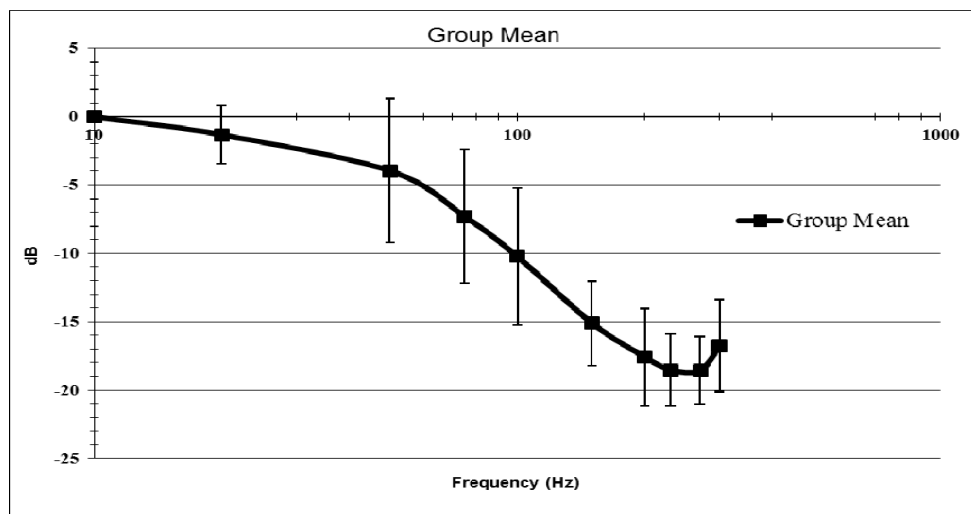


Figure 20: The Group-mean 15 dB_{SL} isosensitivity contour with standard deviation bars.

3.4.4.3.3 Filter Design and Validation

Using the group mean 15 dB_{SL} isosensitivity curve (**Figure 20**), a filter was designed that matched the perceptual response. This filter would have the same shape as the curve in **Figure 20**, but actually represents the opposite effect, since **Figure 20** shows the increased sensitivity, or gain, of the human perceptual system relative to frequency. This this, the signal was equalized to the group mean isosensitivity contour, scaled appropriately for the user to their 15 dB SL at 10Hz (the frequency at which the results were normalized).




Figure 21: Group Mean (with std. deviation) and Filter-response curves.

The resonance frequency of the actuator still had to be compensated for within the tactile perception range (see **Figure 15**), which had the issue of shifting lower in frequency as increasing force is applied to the actuator surface. While the case design attempts to limit the amount of force the user applies to the piezoelectric actuator during normal use, the resonance could potentially shift from 340-350 Hz down to 270-300 Hz. Therefore, it was decided to design the filter to (1) match the 15 dB isosensitivity contour for frequencies between 0 and ~200 Hz, and (2) low-pass filter frequencies above 200 Hz

at a rate of 40 dB per decade. The resultant filter curve, shown in **Figure 21**, had less than 3dB deviation from the 15 dB_{SL} isosensitivity contour previously found.

The filter was validated to ensure that it equalized the perceived feedback and actually produced a flat response curve. This test involved participants performing the subjective magnitude adjustment task again for 15 dB; however, this time the signals were passed through the filter before driving the actuator. It was expected that there would be some deviation due to the variance of participants' subjective magnitude curves from the group mean, so a 3 dB deviation tolerance from the 0 dB line was accepted, with the perceptual values were subtracted from the 10 Hz 15 dB_{SL} value.

The filter was validated using three of the original seven participants, and then again for several new participants in a subsequent study. The participants' threshold at 10 Hz was established, and then amplified that value to reach the 15 dB level, with the filter providing the equalization from there. This step of the experiment became the procedure for the setting the 15 dB amplitude for all participants throughout my experiments, and ensured that everyone performing the experiments felt approximately the same magnitude of vibration. For the filter validation, after determining their 10 Hz threshold and establishing their 15 dB amplitude, subjects were then asked to perform the matching by adjustment task, but this time only the first portion and only with two repetitions, not ten. All participants performed within 3 dB of the filter baseline, as shown in **Figure 22**, with a few exceptions of 4dB deviation at 200 Hz.

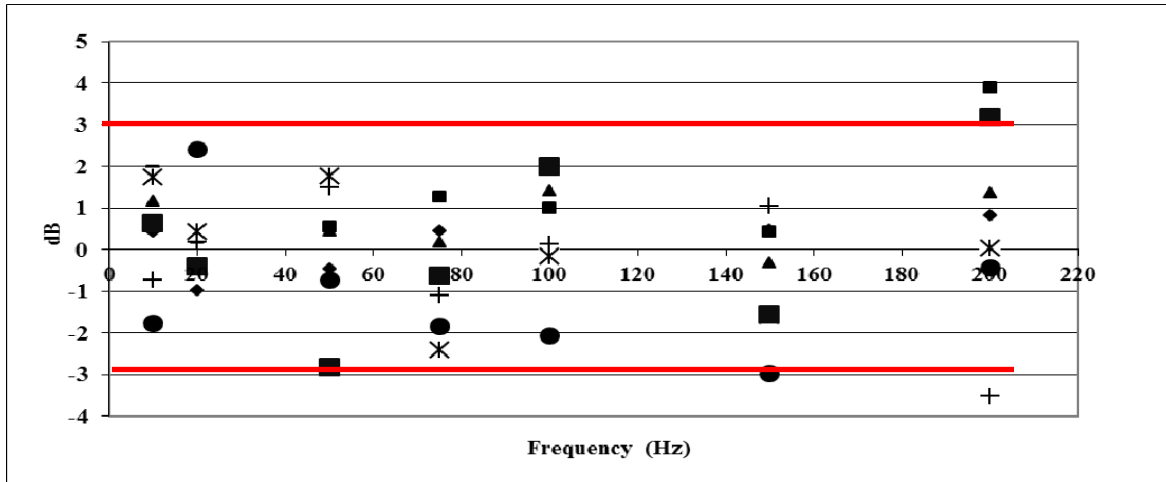


Figure 22: Validation of the filter by 7 participants. Only 2 data points at 200 Hz were above the acceptable limit, and no feedback was used at 200 Hz.

3.5 Development of the “Texture” Feedback

The purpose of redesigning the device was to enable the device to take advantage of the haptic system’s ability to quickly process material properties, such as texture, and do so in parallel across multiple fingers. If the haptic system can blindly identify real objects quickly and with near perfect accuracy [15], then why should individuals who are blind and visually impaired settle for the poor performance observed using raised-line drawings? However, in order for the “textured” feedback to have any chance of improving performance, then it was needed to not only redesign the design, but a proper method for generating “texture” as well. As mentioned before, the three factors that would need to be addressed were (1) how to most effectively actuate the “texture” feedback, (2) how to control the actuation, and finally (3) how and what to encode the texture information in the visual diagram. So far, how the chosen sensor could be used to detect color was discussed, implying that it would be using color to encode at least part of the textured information. Also, the range and abilities of the actuation system (consisting

of the piezoelectric, the LabView™ program, and the perceptual filter) were discussed. The points left to discuss are (1) what types, and how much information do should be conveyed to the user with the “texture”, (2) what mathematical dimensions of the vibratory signal work best to create highly salient “textures”, and (3) how will these dimensions combine into usable images.

The major information types to be encoded were (1) object and part identity, and (2) 3-D orientation relative to perspective. As mentioned in the background section, correctly determining part identity and part orientation were two common problems with most outline representations used in raised-line drawings. However, how to encode and how many unique textures would be needed to represent unique parts, and how to encode the part orientation separately from part identity, were the next big research questions that needed to be resolve.

The vibrotactile feedback generated by the actuator can vary along several temporal dimensions: amplitude, frequency, duty-cycle, harmonic synthesis, and/or amplitude- and frequency-modulation. The key motivation for a dimensions selection was how many salient variables each dimension could yield using the device. This work began when a replacement was sought for the solenoid motor actuator with something that generated more diverse feedback, so certain elements, specifically the filter and the device case, did not exist from the start. In fact, some of this work resulted in the identifying the problems that lead to the development of the filter, and while the actual device case (Figure 13) took some time to make, the design was made a long time prior, so mock-cases were developed for testing the actuator.

The temporal dimensions were first examined for the vibrotactile feedback by looking at how many discernable levels of amplitude and frequency could people notice. Using the Just Noticeable Differences (JNDs) for amplitude and frequency, however, was not found to be an appropriate method for generating textures, as it was extremely difficult to perceive these differences. This was a key point: feedback had to be identifiable, not just discriminable 50% or 75% of the time, especially when it came to using that feedback to encode information about part orientation. Previously, van Erp presented guidelines [106] of using 4 different amplitude levels and 9 different frequencies. This was validated in very preliminary tests with piezoactuator: it was found that most people could only discriminate about 3 to 4 amplitude levels, but around 7 different frequencies.

However, the perception of amplitude and frequency vibration is known to interact [35-37], and so using both dimensions as means of creating different “texture” variables may create confusion, reducing the saliency of the feedback. Since frequency was found to have the largest potential dimension for creating variation, it was decided to keep testing it, and to equalize for the perceived magnitude across frequency (the filter mentioned in the previous section) to avoid any interaction with frequency. The section of frequency over amplitude was also supported by research from [89], who found that amplitude, or vibration strength, had only very few (two-three) usable levels, compared to frequency. Other groups [90, 85], indicated that creating complex vibrotactile temporal waveforms through signal modulation (AM and FM), synthesis, and duty cycle could yield salient dimensions from which “textures” could be generated; therefore, it was decided to pilot some of these to look at possible interactions and problems with saliency.

3.5.1 Evaluation of Frequency, Waveform and Signal Modulation

3.5.1.1 Experiment 1

In this the first preliminary feedback experiment, the purpose was to examine the dimensions of frequency, waveform-shape (harmonic synthesis), and signal modulation, and compare their saliency via a confusion matrix. The specific aim was to see if any dimensions were particular salient, or if any were confused with one another. Hoggan and Brewster [88] had previously found waveform shape to be an effective parameter, followed by frequency, and lastly amplitude modulation, for generating Tactons, or vibrotactile tones that encode information similar to a visual icon. Kyung and his colleagues [85] tested identification rates for spatial waveform patterns, using stimuli of vibration with both no movement and active touch cases, static actuation with no movement and active touch, probe sensing, and moving waveform, using a distributed display of pins. While they found that the identification for active touch to be best overall, they also observed large variation in identification using spatial waveform shapes using both a distributed display and a probe. Signal modulation had also previously shown mixed performance, with some showing limited success [88], others finding better success [90], with no significant performance difference between amplitude and frequency modulation.

In this experiment 12 stimuli were tested, which consisted of three variables of frequency: 45, 80, and 140 Hz, and four variables from the other two dimensions: a sinewave (no synthesis or modulation), a saw-tooth wave, a square-wave, and a 15-Hz amplitude-modulating signal. The number of variables, three frequencies and four other, corresponded to a rough plan to encode for three part-identifiers, and for four

orientations: vertical, horizontal, rounded, or other. A thirteenth variable of a 400 Hz square-wave signal was also examined, as it was desired to test the saliency of a very-high frequency component for possible use as a possible means to encode “borders” between parts or surrounding objects.

Six subjects with no visual impairment between the ages of 22 and 32 years were tested using a single device worn on the index finger of their dominate hand to see how well they could identify the 13 different output modes of the device; additionally, the no-output case was tested to see if participants would confuse that case with any other. This experiment preceded the development of the perceptual filter (Section 3.4.4.3), and so amplitude was set to a constant value for all participants. Digitally filtered pink noise was presented through headphones during the experiment to eliminate any auditory feedback the device provided.

Participants went through a training period that consisted of an explanation of the device operation and description of each output choice, followed by a practice exercise. During the practice exercise, participants had to try and guess the frequency (low, medium, or high) and waveform type of the output signal (sine wave, square wave, sawtooth wave, and modulated wave), or whether the texture was the 400 Hz border texture or the no-output case; all participants received feedback on their choices. Each stimulus was presented in the determined order unless participants specifically requested to feel a stimulus again.

The experiment took place after a 5 minute rest period after training; the 14 stimuli were presented to the participants in a randomized order, and they were asked to identify the stimulus, with their answer being recorded. As with the practice exercise, participants

could request to go back and sample any two specific texture outputs, but were given no feedback or other help in their choices to select.

A confusion matrix was generated from the data to better visualize how participants perceived the various stimuli, rather than perform a comparison of means, as the latter, with so few subjects, would not show how specific stimuli were being perceived.

| G\C | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 |
|-----|---|---|---|---|---|---|---|---|---|---|----|----|----|----|
| 0 | 6 | | | | | | | | | | 2 | 1 | | |
| 1 | | 3 | | | | | | | 1 | | 1 | | | |
| 2 | | 1 | 3 | | | 1 | | 1 | | | | | | |
| 3 | | | 1 | 2 | | 1 | | | 2 | 1 | | | | |
| 4 | | | 1 | | 2 | | | 1 | | | | 1 | | |
| 5 | | | 1 | | | 2 | | | | | | | | |
| 6 | | | | 2 | | | 5 | 1 | | 2 | | | | |
| 7 | | | | | 1 | | | 2 | | | | | | |
| 8 | | 1 | | | | 2 | | | 3 | | | | | |
| 9 | | | | 2 | | | 1 | | | 3 | | | | |
| 10 | | 1 | | | 1 | | | 1 | | | 3 | 1 | 1 | |
| 11 | | | | | 2 | | | | | | | 2 | | |
| 12 | | | | | | | | | | | | 1 | 4 | |
| 13 | | | | | | | | | | | | | 1 | 6 |

Table 3: The confusion matrix for the texture sensing test

Table 3 shows the confusion matrix for the test; the columns are the correct responses and the rows are the guessed responses: numbers along the diagonal correspond to the number of correct responses. During the experiment, subjects were not informed that each texture gets presented only once; thus, responses along a row do not necessarily have to add up to six, whereas in each column they do. Case 0 corresponds to the 400 Hz square-wave stimulus, and case 13 to no-stimulus case. Cases 1, 2, and 3 correspond to the sine-wave frequencies 40, 80, and 140 Hz, respectively; likewise, cases 4-6 correspond to the sawtooth-waves of increasing frequency, cases 7-9 to the square-waves, and 10-12 to the amplitude-modulated waveforms.

The confusion matrix shows that participants confused the low and medium modulated textures with the 400 Hz border a total of three times [See Appendix 8.2 for data]. Participants also confused the modulated signals with each other, in no apparent specificity. Participants exhibited great confusion between identifying the sine-waves, sawtooth-waves, and square-waves of the same frequency; only the high frequency sawtooth wave had a high (5 out of 6) identification rate. Neither the sawtooth-waves, nor the square-waves were confused with the same waveform type of a different frequency. Subjectively, most participants identified the square-waves as feeling the most distinctive of the waveform types (sine-, sawtooth-, square-, or modulated), and had more of a pop-out effect.

From this pilot, several observations were made: (1) waveform type was not a sufficiently salient enough dimension to use in generating identifiable “textures”; (2) participants found the single very high “border” frequency to be highly salient; (3) participants preferred the feel of the square-waves, even if the data does not indicate that this method is any better, or any worse, than sawtooth-waves. However, this pilot had some series faults as well: too few subjects to properly counterbalance against learning bias and participants were allowed to “re-feel” up to two stimuli blurred the lines between having an identification task and an easier discrimination task. The first problem could not have been remedied with anything other than having eight more participants, an unfortunate circumstance of trying to quickly perform an experiment in a limited timeframe. The second problem, however, was actually a design choice; under circumstances of regular usage, it would be reasonable to include a key or legend, accessible to the user at all times, to explain which feedback stimulus corresponds to

what specific information. However, because this methodology means the experiment conducted, and the data collected, is not truly reflective of identification saliency, this was seen as poor design choice. While future experiments would not allow participants to go back and feel specific “textures”, the training portion was expanded, so that participants would be more gradually weaned off feedback, as they were in this experiment, meaning longer training periods and hopefully better learning of the task prior to testing.

3.5.1.2 Experiment 2

The purpose of the second experiment was to determine how many salient, identifiable points could be gleaned from the dimension of temporal frequency. To perform the experiment, eight participants were asked, all with no visual impairment, to identify different temporal frequencies. As mentioned in the Background Section, the human range of vibration perception ranges from 1-2 Hz to over 500 Hz, but utilizes different mechanisms to detect low frequencies (< 80 Hz) and high frequencies (> 80 Hz). Based on this, the experiment was broken up into three parts, each testing a single set of randomly chosen high frequencies: no output, 91, 147, 222, 300, 354, and 469 Hz; and another set of low frequencies: no output, 11, 17, 31, 47, 61, and 79 Hz. Each participant was told that the frequencies were called A-G, with A being the slowest, and G being the fastest in the set.

For training, participants were allowed to feel each vibrotactile frequency, and then followed this with a practice identification run through the frequencies, with feedback for each response. Finally, after a 2 minute break, participants went through a second practice run, where the number of feedback responses given was restricted to two.

Participants were blindfolded and played masking pink-noise for all participants as in Experiment 1. For both portions of the experiment, participants were presented a single frequency at a time and were asked to guess which frequency, A-G, it was, without any feedback. Each frequency in the set was presented once, with the presentation order counter-balanced between participants; frequency set presentation order was also randomized and counter-balanced (why 8 participants were used instead of 7), so that half completed the high frequency set first and the other half completed the low frequency set first. The confusion matrices for both sets are shown in **Tables 4 and 5**.

| | | | | | | | |
|-----------|---|----|-----|-----|-----|-----|-----|
| High freq | 0 | 91 | 147 | 222 | 300 | 354 | 469 |
| 0 | 8 | | | | | | |
| 90 | | 6 | 1 | | | | |
| 150 | | 2 | 6 | 1 | | | |
| 210 | | | 1 | 6 | 1 | 1 | |
| 300 | | | | 1 | 5 | 3 | |
| 360 | | | | | 2 | 4 | 1 |
| 450 | | | | | | | 7 |

| | | | | | | | |
|----------|---|----|----|----|----|----|----|
| Low freq | 0 | 11 | 17 | 31 | 47 | 61 | 79 |
| 0 | 8 | | | | | | |
| 11 | | 6 | 3 | | | | |
| 17 | | 2 | 5 | | | | |
| 31 | | | | 5 | 2 | | |
| 47 | | | | 2 | 4 | 2 | 1 |
| 61 | | | | 1 | 2 | 4 | 3 |

| | | | | | | | |
|----|--|--|--|--|--|---|---|
| 79 | | | | | | 2 | 4 |
|----|--|--|--|--|--|---|---|

Tables 4 (top) and 5 (bottom): Confusion matrices for the frequency test

From the confusion matrices, it was noticed that for both sets, confusion was greatest between adjacent frequencies within the set. Also, for the low frequency set, the upper range of the frequencies, those greater than 40 Hz, seemed to get confused more than those on the lower end. As the JND for frequency increases logarithmically with frequency, increasing the frequency in this manner might improve performance. To test this, the test was run again using frequencies that varied along a log-2 scale, starting at 10 Hz and going to 320 Hz. The same eight people participated again, using the same method, but now only letters A-F. The data are shown in the confusion matrix (**Table 6**).

| Log 2 | 10 | 20 | 40 | 80 | 160 | 320 |
|-------|----|----|----|----|-----|-----|
| 10 | 8 | | | | | |
| 20 | | 8 | | | | |
| 40 | | | 8 | | | |
| 80 | | | | 8 | | |
| 160 | | | | | 7 | |
| 320 | | | | | 1 | 8 |

Table 6: Confusion matrix for the frequency test using a log-2 scaling

The log-2 scale results showed remarkably better performance than either of the previous two frequency sets. While this may be due to the increased learning the participants had by this third experiment and the slightly reduced number of variables, most participants verbally stated that the wider spread between frequencies made the task a lot easier. To quickly confirm if these results were reproducible, four new participants performed an identification task using just four frequencies: 12, 25, 50, and 100 Hz after a brief period of practice, and found that 15 out of 16 times the frequency was correctly

identified, further supporting the idea that wider separation, such as using a log-2 scale, between selected frequencies improved their identification.

3.5.2 “Texture” Patterns of Mixed Temporal and Spatial Dimensions

To render visual images, the system would at least need to be able to render enough “textures” for at least three unique object parts, and four different orientations; in other words, 12 unique and highly salient textures. Rendering 12 unique textures represented a reasonable diversity from which to parse and encode object/part information. There seemed little advantage in rendering more textures, as each additional texture becomes harder for the user to identify and would only be needed for very complex diagrams; however, diagrams are normally simplified in order to be manageable by the haptic system and a very complex diagram unlikely.

At this point in the experimentation with textures, the perceptual filter was developed. After developing the perceptual filter (Section 3.4.4.3), work towards establishing dimensions from which highly salient “textures” could be generated continued. After the previous experiment (Section 3.5.1), the dimension of frequency remained an option; however, the dimensions of waveform shape: sine-, sawtooth-, and square-waves remained uncertain. The effect the filter would have on these variables were quickly tested using a few sighted individuals.

During testing, most individuals could no longer distinguish a sine-wave from a sawtooth-wave, but they indicated square-waves still were noticeably different. Again, most individuals said they preferable the feel of the square-wave versus a sine-wave of equal-magnitude; therefore, it was decided to rule out the use of harmonic synthesis as a possible dimension and decided to use square-waves rather than another shape when

generating temporal signals in LabView™. This left frequency, modulation, and duty cycle as possible temporal dimensions of the feedback signal that could be used to generate “textures”.

In addition to these temporal dimensions, using the spatial dimensions of the images themselves to create the “textures” was considered. Spatial dimensions are derived from the fact that the device was designed, via the LabView™ control program, to turn color on the computer screen into vibrotactile feedback. Therefore, variations of color across the image translate into variations in the vibrations the user feels, creating spatially-dependent “modulation” of the feedback signal as the user scans across the color variations. This also means that if a particular object or object part with a two-color striped pattern was painted, and those colors each had a unique frequency encoded for them, then when the user scans the device across this pattern, they will feel a spatially-dependent frequency modulated signal. Other groups (93, 112) have used spatial patterns combined with vibratory signals to generate texture-like patterns; however, while these groups looked at creating discriminable patterns, the goal was to use identifiable patterns to encode specific information relating to the object or object part.

Using spatial patterns within the visual image on the computer screen opened up several new dimensions from which “textures” could be created, besides the one given above. Using the simplest patterns—color stripes—were looked at first. Stripes that reached from one side of a part or object to the other were logically more detectable than alternatives, such as overlapping stripes (think plaid), or repeating small features. Variations due to striped patterns could be used to create the following dimensions: spatial modulation of frequency, spatial modulation of amplitude, spatial modulation of

duty-cycle, and spatial direction (i.e., the direction of the stripes relative to the absolute position of the screen/user interface). The purpose of the next set of experiments was to determine the most prominent of these perceptual dimensions, including temporal frequency and temporal duty-cycle, through MDS analysis.

The method of Multi-dimensional scaling was chosen, which directly maps the perceived dissimilarity of physical factors onto n-dimensional space, to see along what dimensions people make their judgments. However, the number of physical parameters was very large for a human-subject experiment, with six treatments and three to four variables in each treatment.

One hypothesis was that the temporal frequency and temporal duty cycle may collapse into one dimension, similar to the spatial parameters of roughness that collapse into groove width [27]. Therefore, these two dimensions were examined directly in a smaller MDS design, with the hopes of reducing the dimensions needed to be considered later. Also examined was spatial modulation frequency and spatial modulation duty cycle, again, with the intention of reducing the dimensions to be considered later down to a more manageable number.

3.5.2.1 MDS Analysis of the Dimensions

In this experiment, MDS was used in hopes of reducing the number of potential temporal and spatial dimensions down by determining which among them were the most salient. Ten sighted, right-handed participants (age 22 to 28) took part in the experiments. None of the participants indicated that they had any malady known to affect tactile sensitivity. All participants wore the final prototype device on the index finger of their dominant hand. They also wore headphones playing pink noise to block any

auditory feedback from the device. Prior to the experiments, participants' absolute thresholds at 10Hz were determined using the method of adjustment as in Section 3.4.4.3.3. Then for the experiment, all stimuli output magnitudes were adjusted to 15 dB above the participant's threshold value. Participants received training on how to use the device, and for the second experiment, were additionally trained to move the device a constant speed of 30mm/sec back and forth in a constant horizontal direction. Participants were blindfolded for the duration of the experiment, but could remove the blindfold during requested breaks.

For each experiment, participants were asked to perform a pair-wise comparison of stimuli, ranking them directly based upon their perceived dissimilarity. The two experiments each contained a total of 9 individual stimuli and 45 unique stimuli pairs. Experiment one only contained stimuli that varied along temporal frequency and duty cycle, while experiment two only contained stimuli that varied along spatial modulation frequency and duty cycle. The order that the experiments were performed was counter-balanced between all participants, and all participants participated in both experiments. The 45 stimuli pairs were presented in randomized order during both experiments. Multidimensional scaling analysis was performed, using SPSS, separately on the aggregate dissimilarity data collected from each part of the experiment.

3.5.2.1.1 Experiment One: The Temporal Factors

In this part of the MDS experiment, the temporal dimensions of frequency and duty-cycle were crossed and ranked by their dissimilarity to determine which factor was perceptually the most salient.

3.5.2.1.1.1 Method

In this experiment, two stimuli were displayed on a computer LCD monitor placed in front of the subject, on a table, flat on its back. Each stimulus was a solid, single-color rectangle, 7.5cm by 12cm in size, with a 3cm gap between the two stimuli. There were 9 colors, each activating a temporal square-wave signal that varied along temporal frequency and duty cycle. The temporal frequencies chosen for the experiment were 30 Hz, 60 Hz, and 120 Hz; the three duty cycles chosen were 30%, 50%, and 70%. All commanded signals were passed through the filter to normalize the relative magnitude perception across frequencies.

3.5.2.1.1.2 Results of Experiment 1

Figures 23 and 24 show the two-dimensional and one-dimensional MDS solutions, respectively, for the aggregate data. These results show the consistent grouping of the variables of same temporal frequency along the x-dimension of the plot; however, in terms of temporal duty cycle, there is no consistent ordering among variables. In the figures, A, B, and C correspond to duty cycles of 30%, 50%, and 70%, respectively; and the number following the letter corresponds to the frequency of the stimulus.




Figure 23: The two-dimensional solution of the MDS analysis.

Both the two-dimensional (**Figure 23**) and one-dimensional (**Figure 24**) solutions showed that frequency is the most salient dimension, through the close grouping of equal frequencies along an-axis. The variation in temporal duty cycle causes very little dissimilarity among the stimuli, except for a single case at 60 Hz (A60), as shown in Figure 21. Overall, no consistent pattern was found in the dissimilarity mapping for the different duty cycles along the y-dimension (**Figure 23**), which supported the conclusion that the two dimensions collapsed into a single dimension corresponding to temporal frequency.




Figure 24: The one-dimensional solution of the MDS analysis.

3.5.2.1.1.3 Discussion of Experiment 1

From the results of the MDS analysis, it was concluded that the x-dimension is the most useful in describing the perceived dissimilarity, and the temporal frequency of the signal the best fit this parameter, making it the more salient of the two dimensions tested. However, the results also showed that temporal duty cycle did influence the dissimilarity, as there was still some dissimilarity amongst stimuli with this parameter. The lack of any pattern along the y-dimension as a function of duty cycle, however, suggests that this is not a second dimension.

A possible explanation of the results is that the data shown in **Figure 23** can actually be described by a curve of one-dimension, similar to a polar plot. This curve may actually represent a more complex parameter that is dependent on the frequency content of the signal, which resulted from the multiple harmonic frequencies making up the square wave, different duty cycles altering the frequency content of the square-wave, and the frequency attenuation of the filter. Whatever the case may be, it was clear that this dimension was definitely not dependent on the temporal equivalent of “groove width”. For the purposes of my device design, the results provided justification for my use of temporal frequency over temporal duty cycle as a “texture” parameter, as it is the more discriminable of the two.

3.5.2.1.2 Experiment Two: The Spatial Modulation Factors

3.5.2.1.2.1 Method

In this experiment, much of the methodology remained the same as in the previous experiment. The stimuli were presented as rectangular patches in the same placement and dimensions. Unique to this experiment were the visual components within these rectangles: spatial square-waves were created of varying wavelengths and duty cycles by making colored, striped patterns on the visual display. For example, one pattern may consist of gray and black stripes with the juxtaposition of a gray or “high” stripe and a black or “low” stripe forming a single spatial wavelength. (**Figure 25**) The ratio of the width of the “high” stripe to the “low” stripe created the duty cycle. The device, via the optical sensor, detected the colors and used them to modulate the temporal waveform (in this case a 60 Hz square wave). Based on pilot testing, the “high” stripe corresponded to

the 15 dB amplitude for the participant, while the “low” stripe corresponded to an amplitude 50% that for easy distinguishability from each other (**Figure 26**).




Figure 25: Example of pair-stimuli for the spatial factors experiment. The dimensions and scaling of the images has been altered to fit this paper.

The spatial modulation of the 60 Hz square wave varied for combinations of three spatial periods and three spatial duty cycles. For the experiment, spatial periods of 10 mm, 12 mm, and 15 mm, and spatial duty of 25%, 50%, and 75% were used. As in the previous experiment, all actuator-driving signals passed through the filter before being actuated.

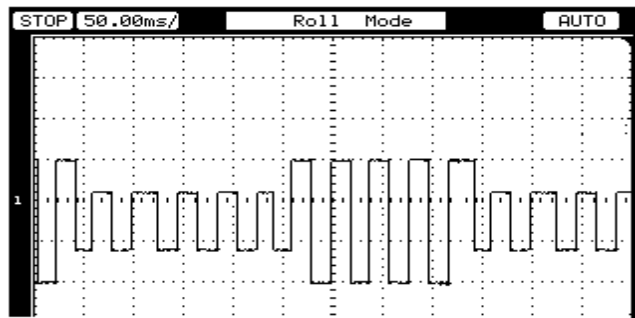


Figure 26: The spatially modulated, temporal output signal.

3.5.2.1.2.2 Results

Figures 27 and 28 showed the two-dimensional and one-dimensional MDS solutions. In these figures, A, B, and C corresponded to spatial periods of 10 mm, 12 mm, and 15 mm, respectively; the number following the letter corresponds to the spatial duty cycle.

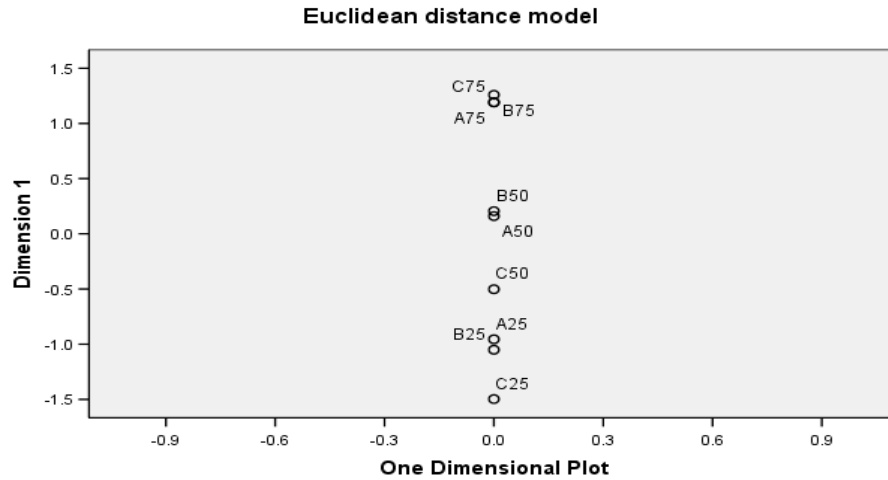


Figure 27: The two-dimensional solution of the MDS analysis.

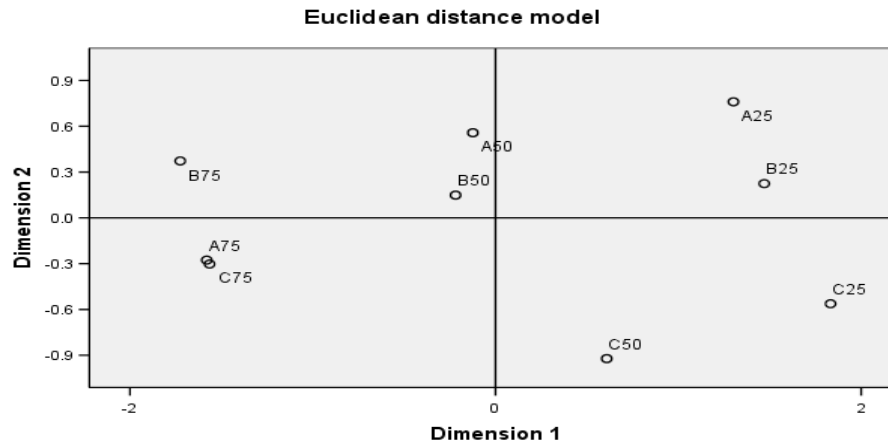


Figure 28: The results of the one-dimensional solution.

The distance data shown in Figure 24 can be explained along dimensions of spatial duty cycle (x-dimension) and spatial frequency (y-dimension), with the exception of point

A75. The data in **Figure 28** show that of the two spatial factors, spatial duty cycle is the most discriminable dimension. An interpretation of the plot shown in **Figure 27** is that twice as much dissimilarity is across the dimension best corresponding to spatial duty cycle as compared to spatial wavelength. The Scree plots for the two MDS experiments were inclusive as all Scree plots are for one-dimensional solutions; nevertheless, they are included in the **Appendix 7.3**.

3.5.2.1.2.3 Discussion

The second experiment showed that even over relatively large changes in spatial period, the most salient dimension could be explained along parameters of spatial duty cycle, rather than any interpretation of groove-width, the most dominate parameter in the perception of roughness by the bare hand [51]. This may be due to the spatially modulated-temporal pairings in the experiment being more comparable to the tactile-rhythms used by MacLean and Ternes [89], rather than sensations of groove-width, with spatial duty cycle more akin to long versus short notes and even versus uneven notes in their design. However, the dimensions did collapse into a single one corresponding to spatial duty cycle, so the experiment did advance the goal of reducing the number of dimensions that would need to be included in the main experiment.

Overall, it was concluded from these two experiments that for patterns generated from variables of temporal frequency and temporal duty-cycle, individuals rated the differences in the patterns along the dimension of temporal frequency. For patterns generated from variables of spatial period and spatial duty cycle, individuals rated the differences in the patterns along the dimension of spatial duty cycle. It was the plan to then combine these two dimensions, along with a third dimension of spatial direction,

into a single MDS experiment, with 27 unique patterns, three variables for each dimension, and 378 total pairs. However, it was not known how well the discriminability mapping (**Figures 23, 24, 27, and 28**) translated into user performance in identifying “textures”; just that one dimension was more discriminable than another. While insight was gained on what would mostly not work, little information was obtained about how well the two remaining dimensions would work in actual usage.

Therefore, it was decided to test the feasibility of these possible dimensions more directly, first through a few pilot studies that examined how well users could identify the different feedback, rather than discriminate between them. In this pilot, people's ability to identify spatial duty cycle and spatial direction for a single pattern was tested, but still as those cases varied with temporal frequency. Three spatial duty cycles (25%, 50%, & 75%), three spatial directions (Vertical, Horizontal, and Diagonal), and three temporal frequencies (12, 25, 50 Hz) were used as variables to create the patterns. The same method was used to create the spatial-duty cycle patterns as described in the second MDS experiment, but this time the patterns used three different temporal frequencies, and the stripes could be oriented/angled in one of three different orientations. In the experiment, three participants were asked to identify the duty cycle and the direction of the pattern. Having tested identification of temporal frequency three times, participants were not asked to provide any information about frequency. Instead, it was included only to see if temporal frequency noticeably interacted with the other two spatial dimensions.

It was found that participants could identify the spatial direction $83 \pm 4\%$ of the time, compared to $72 \pm 14\%$ of the time for spatial duty cycle. Most of the errors for identifying spatial direction were either false positives or negatives of the diagonal

direction, whereas errors made for spatial duty cycle showed no favoritism. While it was noticed that more mistakes were made on patterns at 12 Hz, this was not significant, perhaps due to the limited subject pool.

Overall, despite spatial duty cycle being mapped along similar-sized dimensions of discriminability (from the MDS experiment) as temporal frequency, the identification performance of this dimension, in this experiment, fell far short of that previously seen for temporal frequency. This indicated to me that my concerns about how applicable the MDS analysis would be in helping validate a set of highly salient “textures” were in part justified. Thus, rather than risk spending the time and money on a third, massive MDS experiment, it was decided to pursue experiments that would more directly evaluate user performance with possible “texture” sets.

3.5.3 Shift to Application-related Experiments

The purpose of these two experiments were to directly evaluate individuals’ ability to identify part identity and orientation information using the potential texture sets, with the hope that a single set would have superior performance in terms of identification rate.

Coming from the pilot studies on temporal dimensions (frequency, waveform shape and modulation), the MDS experiment (temporal frequency and duty cycle; spatial period and duty cycle) and the latest pilot with spatial duty cycle and direction, it was felt that while these experiments had helped to eliminate possible dimensions, so far, only temporal frequency had any supporting evidence for its use in generating highly salient “texture” feedback. The saliency of the chosen “textures” is crucial; as mentioned before, stimuli that are not saliency from each other are less likely to be as easily processed (single finger) or processed in parallel across multiple fingers. In turn, these

were two of the key advantages of using texture-like feedback in the first place, and without this potential benefit, it was not expected that any reduced exploration times and improved object recognition when using additional devices/fingers would be observed.

Furthermore, it was believed there was a distinction in the applicability between experiments that tasked individuals to identify encoded information from the presentation of a single pattern at a time, and how they would be presented during actual use: with multiple patterns of different shapes merged together, describing a single object. Therefore, it was decided the next pilot would focus on testing the two most seemingly salient dimensions from all prior experiments, temporal frequency and spatial direction, using a task more related to the device's intended application.

3.5.3.1 Experiment 1

As an evaluation of the feasibility for the method of using “textured” feedback to encode information pertaining to object/part identity and orientation relative to the viewing perspective was desired, in this experiment, participants would be required to identify part identity and orientation information from multiple “textures” per image. This experiment would primarily serve as a proof of concept for the method of encoding information using “textures” generated from temporal and spatial components of vibrotactile feedback.

Rather than test all possible texture sets at this exploratory stage, only a single texture set was tested on several individuals who are blind or visually impaired, as they are the target demographic for my system. It was hoped that if individuals who are blind or visually impaired could, in fact, use prototypes of the “texture” patterns to convey part identity and orientation, then this would verify a necessary step in my design method. It

was also important to test for any negative effects of interactions between the texture dimensions and between adjacent textures, adaptation of the textures and any other issues when users are allowed to freely explore pictures.

The initial texture set consisted of using temporal frequency to parse the individual object parts and spatial direction to indicate part orientation relative to the viewing perspective. As mentioned in the previous pilot experiment, the dimension of spatial direction refers to the angle or orientation of the color stripes within the pattern, relative to the position of the computer screen. Unlike the previous pilot, the dimension of spatial direction included not only vertically, horizontally, and diagonally stripes, but also a fourth pattern that contained no stripes—a solid colored “texture” pattern. This dimension of spatial direction was used to encode part orientation, as this seemed the most intuitive option; vertical stripes for a vertical orientation, horizontal stripes for a horizontal or flat orientation, diagonal stripes for a rounded or curved surfaces (the less intuitive one), and a solid pattern for anything else not covered. After that decision, using temporal frequency to encode part identity was the only option left; however, during the temporal-frequency pilot (Section 3.5.1.2), it was found that people made judgments of temporal frequency very quickly, usually within a few seconds. This exhibited high-saliency and low cognitive demand with temporal frequency makes it ideal to use when segmenting parts within an object. This was expect to be in line with visual performance, where determining part identity within an object is usually a very quick processing task.

The combinations used three temporal frequencies (12.5, 25, and 50 Hz) and the four spatial directions (vertical, horizontal, diagonal, and none) to generate the “texture” patterns. Additionally, in this pilot a different method of creating the striped modulation

of the signal was used. Before, during the MDS experiment and the quick pilot that followed (Section 3.5.2), the spatial modulation effect was generated by having a single frequency per pattern, with the two colors creating the stripes causing the actuator to vibrate at the 15 db participant-dependent sensation level (SL) for one color, and 50% this magnitude for the other color. This form of spatial modulation could be called spatial amplitude modulation, because the amplitude of temporal signal gets modulated by the spatial pattern. However, before the MDS experiment, due to a miscommunication, something very different was originally tested. Instead of using a single frequency with two different amplitudes, two different frequencies with the same 15 dB SL were used, creating a spatial frequency modulation of the signal (both signals are shown side-by-side in **Figure 29**). Out of curiosity, this method of creating the patterns was quickly piloted, using the same experimental method described for the pilot at the end of Section 3.5.2. The performance was found to be slightly improved with this new method: identification of spatial direction improved from $83 \pm 4\%$ to $88 \pm 5\%$, and spatial duty cycle improved from $72 \pm 14\%$ to $77 \pm 10\%$; and thus, it was decided that this method would be used for the next pilot.




Figure 29: The spatially-modulated signals: on left, the amplitude-modulated signal, and on right, the frequency-modulated signal.

The actual spatial modulation, represented on the visual pattern as a striped pattern (see **Figure 30**), has one color stripe of the visual pattern producing a temporal frequency of 100 Hz and the other color stripe producing a temporal frequency of either 12.5, 25 or 50 Hz (to indicate part identity). It was decided upon 100 Hz using the log-2 scale (Section 3.5.1.2), and as frequencies above 100 Hz seem to get encoded by energy transmission rather than through the temporal pattern, resulting in poorer discrimination [55]. A constant scanning velocity across the pattern will result in the temporal vibration pattern shown in **Figure 29** (right image).




Figure 30: The spatial frequency modulated patterns for three directions.

3.5.3.1.1 Method

Rather than the methodology of previous identification experiments, nonsense objects consisting of multiple parts, each with its own “texture” pattern were presented to participants. Nonsense objects with a large number of parts were chosen to ensure that the texture set had sufficient breadth for describing object scenes and part labeling would not be influenced by early identification of the object.

The participants were then tasked with exploring the entire image space, and 1) manually identify each part they felt and 2) describe what orientation, if any, they felt that part possessed. This design allowed me to validate the two constitutive factors; that of vibrotactile (temporal) frequency to parse object parts, and that of spatially-directed

modulated patterns to indicate part orientation relative to the viewing perspective.

Additionally, participants were informed during training that there would no questions regarding the shape of the object or parts, so that would not spend time trying to guess this information.

3.5.3.1.2 Participant Population

All 10 of the participants in the experiment had significant vision loss (legally blind) or were totally blind; exactly half of the participant population was legally blind and the other half near- or totally-blind. The participants had a mean age of 43.7 ± 16.8 years and of 7 participants who were not legally blind since birth, the average age for the onset of blindness was 27.1 ± 10.4 years. Additionally, 8 of the 10 participants were right hand dominant, 2 were left hand dominant, and all 10 participants either knew, or were currently learning Braille.

3.5.3.1.3 Object Design

The experiment consisted of twelve objects, each presented to the participant separately on the computer screen. Each of the objects contained 3 through 5 parts, pieced together like a collage. The parts were drawn with miscellaneous shapes of varying size, arranged within a 20x20cm space (two samples are shown in **Figure 31**; See Appendix for full experiment sample set). As discussed before, combinations of three part identity indicators (i.e., color/temporal frequency) and four spatial orientations (vertical, horizontal, diagonal, and none) were used to generate the parts. The 12 objects were evenly divided into objects with 3, 4, or 5 parts. The combinations of textures for an object were designed to represent cases much more difficult than what is reasonable

with real pictures. The object shapes were all “nonsense”, i.e., not representing anything close to a real object, as it was advantageous to challenge individuals to see how well they could separate and identify parts that may not be easily discriminated in either temporal frequency or spatial orientation.

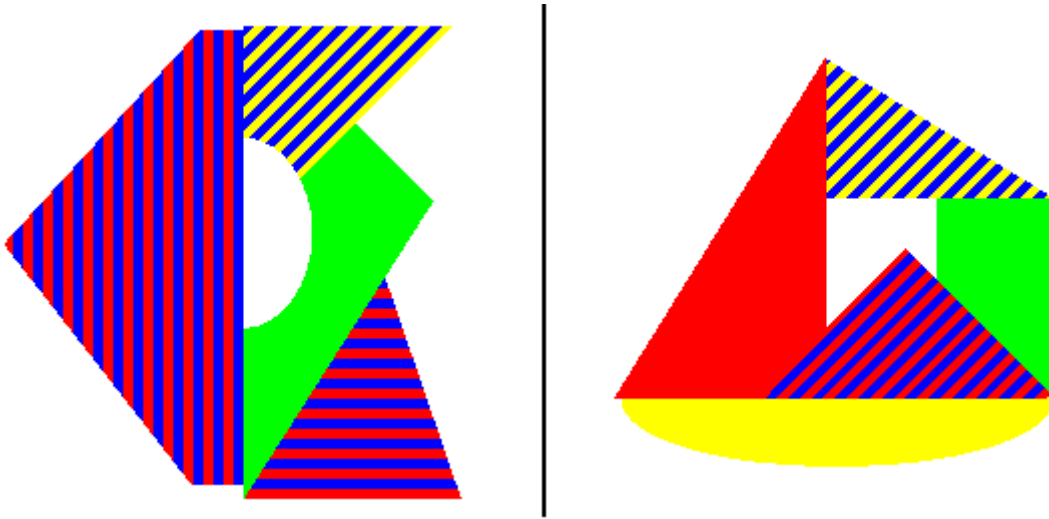


Figure 31: Sample objects from the experiment

3.5.3.1.4 Procedure

Prior to the experiment, participants were evaluated to determine their absolute threshold at 10Hz, from which their 15dB_{SL} scale factor (from the normalized version of the isosensitivity contour) was set. Next, they performed a quick validation of the equalization filter (Section 3.4.4.3.3); no participant had any significant deviation (more than 3 dB) from a flat response for frequencies used in the experiment (10-100Hz).

Then, participants underwent a 1 hour training period where they were introduced them to the device, practiced identifying “textures” within complex objects, similar to those presented in the experiment, and gave them feedback about their answers. The experiment itself consisted of performing the designated tasks for 12 objects not previously viewed, with object presentation order nearly counterbalanced among

participants. Participants were blindfolded during the experiment, and presented pink noise throughout the duration of the experiment to mask any auditory cues. Participants were asked to explore the entire screen using the device; when the participant found what they indicated was a unique part, they were asked to identify the part's orientation using the "texture" cues.

3.5.3.1.5 Scoring

Participant responses were recorded for when they identified, to their belief, a unique part (part identity task) and then, identified the part's orientation. Participant performance was evaluated by first calculating their score for the individual tasks for each object. The score was calculated by summing the correct number of parts, minus the number of parts incorrectly identified (note that the score could be negative if the number of parts incorrectly identified was more than the correct number of parts), then dividing by the total number of parts. For identifying unique parts within the object, there was no chance score, as participants were free to guess as many or as few parts as they felt.

Scoring for part-orientation was similar: number correct minus number incorrect, divided by the total number. The chance rate for the orientation identification task was 25%, since any part on an object could be in one of four orientations. During training the investigator informed participants that parts would only contain 4 possible orientations, thereby limiting each part orientation to a 4-alternative forced choice task, but participants were not informed as to the total number parts an object may have during the experiment.

3.5.3.1.6 Results

Figure 32 shows the scores from the various tasks in the experiment. The average total score for all participants was $86.3\% \pm 5\%$. The scores for the individual tasks were highest for counting the total number of parts ($94\% \pm 2.8\%$), followed by identifying parts without a pattern ($81.1\% \pm 11.7\%$) and lowest for identifying parts with a diagonally-orientation pattern ($61.2\% \pm 18.1\%$). The percent correct for all orientations are significantly different than chance ($p < 0.0001$). Using Fisher's LSD, the only comparison that showed no significant ($p > 0.05$) difference between cases were the horizontal and the non-pattern orientations ($p = 0.548$). Several participants were able to score perfectly in identifying horizontal and non-patterns when tasked; however, vertical and diagonal patterns were sometimes confused. In **Table 7**, the related signal detection probabilities are shown for the tasks; neither false alarm, nor correct rejection probabilities were calculated for the part counting task, as it is impossible to quantify these values in our particular experiment.

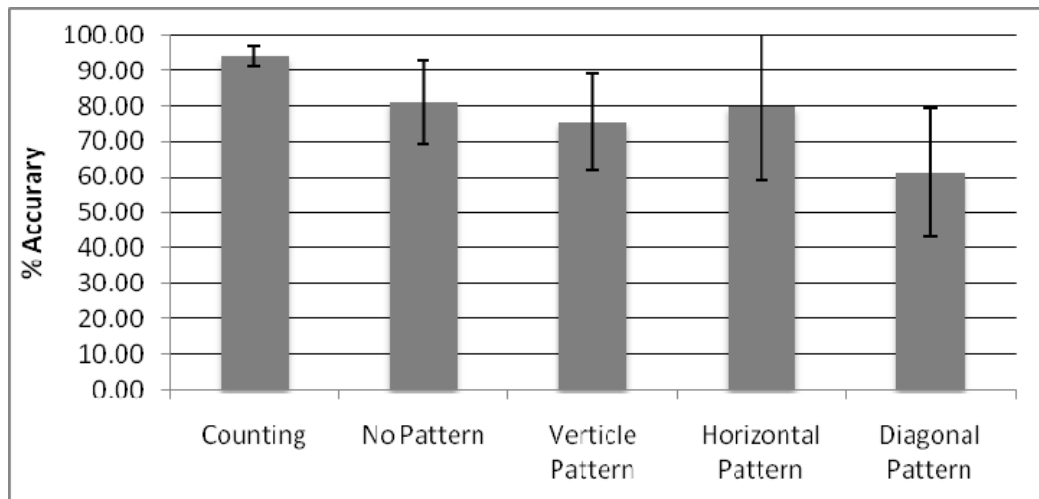


Figure 32: Percent scores from the experimental tasks.

| | Counting | None | Diagonal | Horizontal | Vertical |
|-------------|----------|-------|----------|------------|----------|
| <i>Hits</i> | 0.947 | 0.889 | 0.688 | 0.867 | 0.822 |

| | | | | | |
|---------------------------|------|-------|-------|-------|-------|
| <i>False Alarms</i> | -- | 0.047 | 0.050 | 0.012 | 0.031 |
| <i>Misses</i> | .053 | 0.122 | 0.312 | 0.133 | 0.178 |
| <i>Correct Rejections</i> | -- | 0.953 | 0.950 | 0.988 | 0.969 |

Table 7: Signal detection data obtained during the experiment

It was noted participant feedback, as well as my own observations, during the experiment:

- There were no substantial differences between the performance of participants with varying degrees of blindness, nor between participants whose vision loss was congenital versus adventitious. (*Observation*)
- Participants took roughly, on average, 2 minutes per object to find and describe all the parts, with the minimum and maximum times being around 30s to 7 minutes. (*Observation*)
- Having a border separating the individual parts would be of great benefit. (*Multiple participant comment*). This will be further investigated in the future.
- Accurate identification of spatially-oriented patterns gets worse the faster the individual moves while exploring. (*Observation*). This indicates a key requirement for successful training.
- Even without being asked, participants often spontaneously described the shape of a part, positively indicating the potential for doing the full task (object and object scene identification) when the texture set is optimized and other aspects, such as borders, are incorporated.

3.5.3.1.7 Discussion

The investigation of whether the chosen spatiotemporal pattern set could be used effectively to identify object part and part orientation in a complex picture proved to be a success. The experimental results showed that individuals who were blind or visually impaired could accurately (94%) use the device to parse the graphics into constitutive parts and also accurately (60-80%) determine the spatial-orientation of those parts, and these values were much greater than chance (25%). The relatively lower performance for the diagonal orientation was a concern; however, it was decided to wait until after all

possible texture sets had been tested in the subsequent experiment to see if this trend persisted. This gave support to the methodology used, and prompted the continuation of testing all of the remaining feasible dimensions in a similar, but larger experiment.

3.5.3.2 Experiment 2: The Main “Texture” Experiment

In this experiment, the three dimensions: temporal frequency, spatial duty-cycle, and spatial direction, along with the two methods of spatial modulation: amplitude and frequency, would be evaluated to determine if a single combination of dimensions created a texture set that performed better than the other options. Prompted by the results of the last experiment, this experiment would evaluate the performance of not just a single “texture” set, but all feasible “texture” sets left remaining. It was decided to keep the same basic methodological structure as the previous experiment, by having participants tasked with identifying parts and their orientation for nonsense objects. However, based on user feedback, some changes were made, notably training was longer, user breaks more frequent, and sessions were broken up over several days as not to fatigue the participants.

3.5.3.2.1 Texture Sets

The two dimensions from the MDS experiment were combined, temporal frequency and spatial duty cycle, along with spatial direction, into four possible sets, where one dimension indicated part identity and the other dimension indicated part orientation. Three sets were created using the same method of pattern creation described in Section 3.5.2 (i.e., spatial amplitude modulation or AM): temporal frequency with spatial duty cycle, temporal frequency with spatial direction, and spatial duty cycle with spatial

direction. However, based on the results from the previous experiment, the alternative method of spatial frequency modulation (FM) signals was also included for the case previously tested, with temporal frequency combined with spatial direction. If this latter set proved to have higher or comparable performance to set where AM temporal frequency and spatial direction were crossed, then it the intention to investigate whether FM produces better performance than AM when crossed with spatial duty cycle, should the AM temporal frequency and spatial duty cycle patterns also perform better than the AM temporal frequency and spatial direction patterns.

The number of sets in the experiment was narrowed down through piloting the texture set that was believed would have the worst performance: the combination of spatial duty cycle and spatial direction. It turned out that the suspicions were correct, as this texture set had abysmal performance (< 11% for any orientation) and the participants voiced their hatred for it. This left the experiment with only three test conditions, as listed in **Table 8**.

| | Set 1 | Set 2 | Set 3 |
|----------------------------------|---|---|---|
| Dimension 1, Part Identity | 12, 25, 50 Hz tones (AM) | 12, 25, 50 Hz tones (AM) | 12, 25, 50 Hz tones (FM) |
| Dimension 2, Part Orientation | Vertical, horizontal, diagonal, & no stripes | 0, 25%/75%, 50/50, 75%/25% Duty Ratios | Vertical, horizontal, diagonal, & no stripes |

Table 8: Test conditions for various “texture” sets.

The temporal frequencies chosen were 12, 25, 50, and 100 Hz, as previously had tested (Section 3.5.1.2) and used (Section 3.5.3.1) these ratios, which were more than a sufficient separation to satisfy the observed Weber fractions for vibrotactile perception [113]. For all the spatial patterns in the sets, the width of two stripes (1 spatial period) in a pattern was fixed to 10 mm across, as this dimension could be easily felt and explored (examples shown in **Figure 33 and 34**). The spatial duty cycles were set at 7.5/ 2.5 mm,

5.0/5.0 mm, or 2.5/7.5 mm in width, the ratios of which were selected because they were easy to discriminate when using the device (based on work from Section 3.4.4.1 and Section 3.5.2.1.2).

Finally, the spatial directions used were vertical, horizontal, diagonal, and no stripes, for the spatial orientations of vertical/upright, horizontal/flat, curved/spherical, and unspecified orientation, respectively. Samples of the patterns for each of the three sets are shown in **Figure 33**. In **Figure 34**, a sample picture from Set 3 of the experiment is shown.



Figure 33: AM Spatially Directed Patterns (Left), AM Spatial Duty Cycle Patterns (Middle), and FM Spatially Directed Patterns (Right). While each of the three blocks has a different orientation for the stripes, a test condition for Sets 1 and 3, the stripe orientation remained constant for Set 2.




Figure 34: Experimental Object using Texture Set 3.

3.5.3.2.2 Method

In order to test these three texture sets, an experiment was designed that evaluated the user's ability to determine part identity and part orientation, in complex and challenging combinations, without requiring individuals to identify real objects. As in the previous experiment, this was done to prevent any higher order processing judgments and its perceived representation from biasing the participants' responses as to what they felt.

For each of the three texture sets, twelve images, as in the prior experiment, were created and each image contained three to five separate parts, with an average of four parts per image (forty eight parts total). Since a texture set needed to encode three separate part identities and four part orientations (vertical, horizontal, spherical/curved, and unspecified orientation) for twelve possible combinations, the twelve images with forty eight total parts per texture set allowed me to have four repetitions for each of the combinations within the set.

Each "texture" set was presented in its own session, with no more than one session per day. The shape of the objects in each of the twelve object images did not differ between the three texture sets, so that the sets were of equal difficulty in that respect. As participants did not have to identify shape, and objects between sets were presented in a different randomized order, learning the objects to determine part identity (the only parameter tested that remained consistent between sets), was very unlikely to occur. All part orientations for all objects differed between sets, so even if participants did memorize this information, it would have been of no use and would have had a great, negative impact on performance. The presentation order was counterbalanced for both

the pictures in each set, and the sets themselves, between subjects to mitigate any related learning bias within the set.

Five of the twelve objects contained discontinuities between parts: holes and gaps where no vibration could be felt. These features were included since they are found in images of real objects, and it was believed these highly discriminable features (on-off vibration) could either positively or negatively influence exploration and recognition, depending on whether the participant saw it as a more salient break between two parts, or as the terminal boundary of the object.

All participants were blindfolded (unless totally blind) and subjected to hearing digitally-filtered pink noise to mask any auditory cues the vibrotactile feedback might provide. Before each experiment session, participants were trained on how to use the device and how to discriminate the different patterns within the texture set they would be tested on that day. Participant training lasted between 20 minutes to 1 hour per session, reaching a point where the participant personally felt confident and consistent in their answers. Images used in training were separate from those used in the experiment, with no image being presented more than once during a single session. During the experiment, participants were asked to explore the image, and provide answers as to when they felt a new part and what that part's orientation was. Participants could freely explore the image as long as they needed, although few chose to explore a single image longer than 5 minutes, though, an official time was not kept. Their responses were recorded for both part identity and part orientation, matching them to what was actually presented on the image displayed.

3.5.3.2.3 Participants

Nineteen participants, all of whom were blind or visually impaired took part in the experiment. **Table 9** shows participant information pertaining to age, their level of vision (total includes those who could see light and shadows, but nothing else), the age of onset for their vision loss (**Congenital**, **Early (0-18)**, and **Late (18+)**), their use of Braille (**Never/Learning**, **Sometimes**, **Regular**), and similarly, their use of tactile pictures. Sixteen participants were right-handed, three were left-handed, and only one subject self-reported neuropathy in his hands and fingertips, although this was not found to significantly affect his tactile threshold or his ability to perform the experiment. Participants were paid \$75 for their participation, exclusive of any compensation for travel expenses. The study received approval through the VCU Internal Review Board and all participants consented after having the experiment details explained to them, prior to the experiment taking place.

| Participant | Age | Vision | Onset | Braille | Tactile Pictures |
|-------------|-----|-------------|-------|---------|------------------|
| 1 | 54 | Tunnel (RP) | E | S | S |
| 2 | 49 | 20/400 | C | N | N |
| 3 | 51 | Total | E | R | S |
| 4 | 49 | Total | E | R | S |
| 5 | 52 | Total | L | R | S |
| 6 | 42 | Total | E | R | N |
| 7 | 40 | Total | C | R | N |
| 8 | 57 | Partial | E | S | N |
| 9 | 53 | Total | E | S | S |
| 10 | 43 | Partial | L | N | N |
| 11 | 49 | 20/200 | L | N | N |
| 12 | 55 | 20/200 | L | N | N |
| 13 | 45 | Total | C | R | N |
| 14 | 45 | Total | L | R | N |
| 15 | 42 | Total | L | N | N |
| 16 | 27 | Partial | E | N | N |
| 17 | 34 | 20/400 | E | N | N |
| 18 | 39 | Total | C | R | N |
| 19 | 70 | Partial | L | N | N |

Table 9: Participant Descriptions

3.5.3.2.4 Scoring

During the experiment, participants were asked to explore the object and identify each unique part by its orientation. As mentioned previously, each of the 12 objects in an experimental set contained between 3 to 5 separate parts, and each part had an orientation of vertical, horizontal, diagonal, or no orientation. No orientation, represented by a solid, single tone vibration, was easier for participants to recall than the terms unspecified or other orientation, the latter of which is how its use with real object images was practically intended.

The task required participants to explore the image, local individual parts, and identify descriptive information using the vibrotactile feedback of the device, while keeping track of what they previously explored by building a mental model of the object. If participants identified two or more separate parts that were in fact a single part, only the first one would count as correct, and the additional parts would be counted individually as false positives for the part identification portion of their task. If they missed a part in their exploration, this would be counted as an incorrect for part identification and orientation. If they correctly identified a part as being new, but incorrectly identified its orientation, then this would only count as an incorrect description of part orientation. Participants were not informed of what the limits were for number of parts per object, as in the prior experiment, so there was no limit to the number of false positives possible for part identification (and no % chance of randomly getting the correct answer), unlike with orientation, were again explained during training that any given part would only have one of four possible orientations.

3.5.3.2.5 Results

Using a Generalized Linear Mixed Model with Logit Normal to fit the data, it was found that for identifying a correct part, there was no significant difference ($p = 0.81$) between sets, and the number correct was above 90% for all sets (See **Table 10** for ANOVA results). The part orientation did not significantly affect the part identification. However, the number of parts and the temporary frequency were significant factors which affected performance in the part identification task. Comparison of the odds ratios for identifying the correct part as temporal frequency varied showed that individuals were most likely to correctly identify a unique part when the part was encoded using 50 Hz, then 25 Hz, and least likely at 12 Hz (**Table 11**). The comparison of odds ratios as number of parts varied showed that individuals were more likely to correctly identify part when the object contained fewer parts (**Table 12**).

| Effect | F | (DF1, DF2) | p-value |
|------------------|------|------------|---------|
| Set | 0.21 | (2, 2666) | 0.8101 |
| Number of Parts | 8.42 | (2, 2666) | 0.0002 |
| Frequency | 8.42 | (2, 2666) | 0.0002 |
| Part Orientation | 0.80 | (3, 2666) | 0.4926 |

Table 10: Correct Part Identification: F-tests.

| Comparison | Odds Ratio | 95% CI |
|------------|------------|------------------|
| 50 vs. 12 | 2.1800 | (1.3999, 5.4318) |
| 50 vs. 25 | 2.4382 | (1.5709, 3.7843) |
| 12 vs. 25 | 1.1184 | (0.7822, 1.5991) |

Table 11: Correct Part: Odds Ratios as Temporal Frequency varied.

| Comparison | Odds Ratio | 95% CI |
|------------|------------|-------------------|
| 3 vs. 4 | 3.4112 | (1.2749, 9.1273) |
| 3 vs. 5 | 8.3962 | (3.0775, 22.9065) |
| 4 vs. 5 | 2.4614 | (1.1153, 5.4328) |

Table 12: Correct Part: Odds Ratios as Number of Parts varied.

Using the same Logit-Normal GLMM for part orientation, we found a significant effect between the different sets. Part orientation also had a significant effect for all sets combined. Number of parts had a near significant effect on identification, whereas temporal frequency was not found to be significant (**Table 13**). Participants were able to correctly identify part orientation in Set 3, FM temporal frequency and spatial direction, 82% of the time, compared to 67% for Set 1 and 57% of the time for Set 2 (**Tables 14**). However, when looking at only Set 3, lower odds ratios were found overall, with only the no orientation (represented numerically as -1) versus diagonal (45) and horizontal (0) versus diagonal having odds ratios above 1.5 (**Table 15**).

| Effect | F | (DF1, DF2) | p-value |
|------------------|-------|------------|----------|
| Set | 63.20 | (2, 2666) | < 0.0001 |
| Number of Parts | 2.73 | (2, 2666) | 0.0653 |
| Frequency | 1.43 | (2, 2666) | 0.2404 |
| Part Orientation | 19.60 | (3, 2666) | < 0.0001 |

Table 13: Correct Orientation: F-tests.

| Set | Probability | SE | 95% CI |
|-----|-------------|-------|------------------|
| 1 | 0.6674 | 0.027 | (0.6128, 0.7178) |
| 2 | 0.5678 | 0.029 | (0.5096, 0.6241) |
| 3 | 0.8150 | 0.019 | (0.7739, 0.8500) |

Table 14: Correct Orientation: Descriptive Statistics.

| Comparison | Odds Ratio | 95% CI |
|------------|------------|------------------|
| -1 vs 0 | 1.1503 | (0.7027, 1.8830) |
| -1 vs 45 | 1.9756 | (1.2227, 3.1921) |
| -1 vs 90 | 1.4027 | (0.8492, 2.3169) |
| 90 vs 0 | 0.8201 | (0.5167, 1.3017) |
| 90 vs 45 | 1.4085 | (0.9066, 2.1880) |
| 0 vs 45 | 1.7175 | (1.1181, 2.6382) |

Table 15: Correct Orientation: Odds Ratios as Orientation varied for Set 3.

3.5.3.2.6 Discussion

These results show that regardless of how part-orientation was encoded or what type of spatial modulation was used; the differences between the temporal frequencies remained highly salient, validating its usage to generate “textures”. The data also affirmed the conclusions of Overvliet, et al. [110], that a haptic search task is best when it contains fewer items, explaining why the number of parts in a diagram had a significant effect on whether the parts were parsed correctly. It can also be seen in **Table 11** that the odds of identifying a part are best when the part is a higher frequency; for the two lower frequencies the odds are nearly equal. These results could stem from the different mechanoreceptors most responsive to the different frequencies [18].

From the second part of the data dealing with part orientation, it was concluded that the different methods, as seen in Sets 1-3, did have a significant effect on performance. Set 3 (FM of spatially directed patterns) proved to be the most effective method tested to convey orientation information, and even significantly better than the AM case, which had the exact same spatial parameters as the FM set. Again, similar to separating parts, the number of parts had a near, but not quite, significant effect on the determination of orientation, again affecting the task. There was no observed significant affect for temporal frequency, but a significant effect for part orientation, specifically the diagonal orientation, on interpreting orientation information for the sets combined. However, when only the best performing set (Set 3) is considered, this effect is no longer significant. Whether this is an exmple of the oblique effect [111], that when presented alongside vertical and horizontal gratings, diagonal gratings often get confused with both, remains unclear. However, since Set 3 won the gauntlet to carry on into the next

experiments, this observation is a bit moot as far as further device application is concerned.

One reason why it was believe the FM set performed better than the AM sets was due to the previously seen differences between amplitude and frequency discrimination. It was repeatedly found that frequency was much more discriminable than amplitude using the selected actuator, and while the 50% amplitude was noticeable, it failed to produce the same pop-out effect seen as in the difference between the frequencies. It could be that amplitude, like frequency, needed a larger range along logarithmic scale in order to have a similar pop-out effect; however, this simply was not possible using the device actuator. It should be noted that the differences found between AM and FM does not conflict with the findings of [90] or [91], as spatial modulation was used rather than temporal modulation, to generate feedback.

Overall, this experiment finally settled upon a single, highly salient set of “texture” patterns to use in rendering visual objects. The next experiments will involve testing the performance of individuals who are blind and visually impaired using the device to identify real objects, in single- and multi-device setups.

4 Evaluation of the Display System Performance

After building a device that could be expanded for multiple-fingers, and developing a method for generating highly salient “textures”, the final portion of this work was to evaluate the display system using images of real objects. This evaluation would involve several stages: (1), an initial proof-of-concept test, using only a single-finger version of the device using a limited number of images; then (2) comparison between single-finger and multi-finger versions of the device, and finally (3) comparison of using multiple fingers on either one- or two-hands.

From these experiments it was hoped to test the central hypothesis, which is that rendering objects with information encoded “textured” feedback will provide significant improvement, measured by identification rate, over outline or raised-line representations. Furthermore, it was hypothesized that using multiple devices will provide an even greater benefit over a single device, in terms of improved identification rates and decreased exploration time, when using the textured feedback, but not for the outlined object representations. These hypotheses were based on the fact that the haptic perception of surface materials, including texture, is one of its main strengths, that highly salient, texture-like information can be processed in parallel, and that a display system truly capable of generating highly salient “texture” feedback over multiple independent devices in real time was created.

To validate the device, a series of experiments were performed to test the main hypotheses. First, however, the objects used had to be created, and then they had to be

organized into sets of approximate equal difficulty, so that no within set effects would significantly bias the results. This involved:

- Taking a larger set of commonly identifiable outline object representations and converting them first into TaxyForm, and then into the textured representations using the method concluded from Section 3.5.
- Having individuals who were blind and visually impaired rank the perceive difficulty of the objects, and using these rankings group the objects into sets of equal difficulty.

After this was performed, the central hypothesis: that rendering objects with information encoded “textured” feedback will provide significant improvement, measured by identification rate, over outline or raised-line representations, was tested.

This involved breaking up the hypothesis into four specific hypotheses, and testing them:

- Hypothesis/Experiment 1: That using a single device/finger, texture representations would perform better than outline representations.
- Hypothesis/Experiment 2: That using multiple devices/fingers would yield no performance improvement when using outline representations over a single device/finger, but would when using textured representations.
- Hypothesis/Experiment 3: That using multiple devices and multiple hands would improve performance over using a single-hand with the textured representations.

4.1 Image Creation and Representation

For the proposed comparisons it was necessary to produce many images of real objects for use with the device. Even though the display system was designed to convert the colors on any visual image displayed a computer screen into vibrotactile feedback, the user would have a very difficult time trying to make sense out of images that are very complex, too large or too small, are represented using unusual viewing perspectives, and/or contain different objects that have drastically different sizes, such as a mouse next to elephant [6].

Converting the original visual image into an intermediate form was necessary before presenting it on the computer screen, so that the device could render the final form with the vibrotactile feedback in way that is usable. It was expected this would involve several steps: (1) simplifying the image to provide only enough detail where the object could be identified on basic categorical levels; (2) changing the perspective, if necessary, to a frontal view, and making sure all major parts are shown; (3) resizing the image and applying borders between parts; and (4) applying the “textured” patterns to the parts in order to properly encode identity and orientation relative to the perspective.

There are many methods to transcribe electronic tactile pictures available, and others [76] are currently working an algorithm that would automate much of this process, so this side-project was seen as being outside the scope of this work. Instead, the images used were obtained from [63], and manually modified. Then, an expert of many years in the field of educational diagram creation for K-12 students who are blind or visually impaired, Janice Johnson, gave feedback on all the diagrams made.

The outline representations were visually identical to the pictures provided by [63]; when the device crossed over any line present in the representation, it would trigger a

high frequency (230 Hz square-wave) border tone, which was designed to bypass the perceptual filter so that it would pop-out from the other vibrotactile feedback. The “textured” object representations were made using the method described in Section 3.5.3.2.1 for Set 3, the spatial frequency modulation version. A border, which triggered the same vibrotactile output as it does for the outline representations, was used around the whole object, and within the object to help segment the parts. The image perspective was simplified and changed similar to how [65] did for her own images when testing TaxyForm; however, the “textures” were then applied to the image parts. A sample showing the progression from an outline-representation, to a TaxyForm representation, to the “texture” representation, is shown below in **Figure 35**.




Figure 35: An outlined, TaxyForm, and “textured” representation of an apple.

4.1.1 Image Set Groupings

The process of creating the images remained consistent for the experiments discussed in this section; however, the individual experiments each required a different number of pictures. As mentioned previously, images were used from [63], and with the aid of Mrs. Johnson, images were selected and drawn based on their familiarity and relative uniqueness, i.e., if two or more images were very similar, only one would be chosen.

The experiments would also require the images to be grouped into different sized sets, but each set needed to be of equal difficulty for the test to maintain homogeneity. To ensure that the sets were of equal difficulty, a small experiment was developed, where mostly individuals who are blind or visually impaired rated the pictures based on their perceived difficulty, then using this data, the images were grouped into sets of approximate equal difficulty and performed a t-test on the sets' difficulty ranks to validate the equality of their mean difficulty.

This was done twice; the first time with fewer images, and since the second time reused many of the participants, due to the limited population in this area, the images were regrouped so that the participants would not be familiar with the set. However, their experience identifying the individual pictures was of less concern, as many remained the same between experiments, this learning bias would be the same if they were using the display system repeatedly out in the real world. Nevertheless, all participants in the last difficulty ranking had previous experience with the images, so if they found them to be easily recognized, that would have been reflected in their difficulty ranking.

4.1.1.1 Image Difficulty Ranking Experiment

As mentioned previously in the last section, the images were tested twice to produce two different groups. The images were produced using TaxyForm's methods and embossed using a Tactile Image Enhancer. TaxyForm was chosen, rather than the method of object presentation developed for this project, because the images were visually equivalent (see **Figure 35**), and also so that individuals who may go on to

participate in the experiment were not biased ahead of time by any proto-image form, which may be altered before the main experiment.

For the first grouping, 32 images [See Appendix for list and images] were tested on two separate participant groups. The first group were asked (3 participants: 2 visually impaired, 1 sighted; each blindfolded) to identify and rate the objects on their difficulty, and after the test was completed, they could check their answers and provide any feedback to how they felt the image should be drawn. From the results, several objects were replaced or redrawn. Participants were allowed to use any difficulty scaling they wanted, and their scores were converted into a Student-T score to remove any scaling-bias between participant ratings. These objects were grouped into 4 sets (one being from [65]) of approximately equal difficulty. The second group (4 sighted and blindfolded) was asked to complete the same task with the new images grouped into the four sets. Again, each participant's rankings were standardized using a t-statistic; one-way ANOVA performed on the normalized rankings did not find any significance between sets ($df = 3$, $F = 0.656$, $p = 0.581$).

For the second grouping task, 54 images [See Appendix for list and images] were tested, out of which it was hoped to have 48 usable ones. Six participants, all of who were blind or visually impaired, rated the images based on their difficulty using any scale they wanted. Again, each participant's rankings were standardized using a t-statistic, and then averaged the results. Six of the most difficult pictures were excluded, and only minor changes were made to the other 48 based on participant and expert feedback; therefore, a second group rerate the images was not performed. The t-scores were averaged across participants, and grouped the images into sets based on their difficulty.

One-way ANOVA was then performed on the four sets' t-statistics, and again, no significance between sets was found ($df=3$, $F = .013$, $p = 0.998$).

4.2 Single-Finger Experiment

In the first experiment, to test the first hypothesis: that individual performance would be higher when using the textured object representations than it would be for the outline representations of real objects, which are currently the primary representation of tactile diagrams. The single finger/device performance using the “texture” representations would also be compared to that found with TaxyForm, which had the benefit of multiple bare-fingers [65]. This was intended to serve as a practical validation for the method and a launching point for future work using multiple point-contact devices across several fingers.

4.2.1 Method

For this experiment, participants' performance in identifying common objects was evaluated with no significant limits to what the object set could encompass. It was reasoned that outside of a clinical setting, people using the device may not have benefit of knowing what superlative category the object fits into, nor would they typically be trying to answer some multiple-choice question. Participants received brief (approximately 5 to 10 minutes) training on how the “texture” patterns encoded the separation of parts and part orientation (which remained the same from the previous experiment), and were told that all the images would be of common, everyday objects. Participants were evaluated on whether they correctly identified the object in question, and the time it took them to explore the object. Afterwards, participants were asked to complete a limited version of a system usability survey to compare their opinions between the two representations.

4.2.2 Participant Population

Of the 8 individuals who participated in this experiment, half had significant vision loss (legally blind) or were totally blind, and the other half had no significant vision loss. All participants had experience with the device in the 6 months prior to the experiment; this is largely due to difficulty of finding potential subjects who had not yet participated in our studies. The participants had a mean age of 38 ± 12 and half (2) of the visually-impaired/blind group had some experience with Braille.

4.2.3 Results

The group means for picture identification were 22% correct for outlined representations and 44% correct for textured representations. Analyses (McNemar) revealed that the difference was significant ($\chi^2 = 5.633$, $p = 0.018$) between the two methods in identification. Analysis using a generalized linear model, assuming a Poisson distribution and using a logit linking function, found no significance ($\chi^2 = 0.419$, $p = 0.518$) in the difference between participants with normal sight and those with visual impairment, nor any significant interaction effects between vision level and representation method ($\chi^2 = 0.014$, $p = 0.906$). Thus, the results were aggregated between the blind and sighted participants.

In addition, a matched-pair t-test was performed on the time data, but found no significant difference ($p = 0.136$) between the times taken for the outline representations (273 ± 147 s) and the textured representations (242 ± 124 s). The large variances seen with the time data were not only between participants, but within as well. Most participants took a wide range of times before making a guess at an object; an extreme case was one participant who took only 6 sec to correctly identify a chair, but took 8 min to incorrectly identify a lamp (said jar). These results are shown in **Figure 36**. The odds

ratio showed that participants were 2.52 times more likely to correctly identify the object when using the “texture” representations versus using the outline representations.




Figure 36: The ID rates and exploration times for the single-finger test.

Results from the System Usability Survey (SUS) [97] showed a preference for the textured method, with it receiving a score of 48 compared to a score of 35 for visually realistic representations, out of a possible score of 80. A matched-pair t-test of these results showed a significant ($p = 0.009$) difference between the two representation methods.

Additional participant comments that were noted during the experiment:

- Most all of the participants believed that having more contacts (i.e., more devices/fingers) would be beneficial.
- All participants noted that having the parts “filled” with an interior texture helped them keep track of the object’s interior versus exterior borders.

4.2.4 Discussion

In this first experiment, the first hypothesis, that object identification task using only a single device/finger would perform better, in terms of the identification rate, was supported by the results. The range of possible objects was purposefully kept as open-ended as possible, as in [64, 65], to get a true feel for the baseline real-world performance of the method. The object identification accuracies were low, as we expected, with 22% for the visually realistic representation and 44% for the textured representation. However, these were not too different from identification rates seen with the same objects using TaxyForm, which saw 12.5% - 50% [65] accuracies for visually realistic objects and 50% - 68.7% for textured objects, using static representations and multiple bare-fingers. However, exploration time did differ dramatically between our single-finger device and TaxyForm, with average time for our method taking around 4.5 min, compared to 1 min for TaxyForm. This was not surprising, as exploration with only a single point to provide information is a laborious task, and using devices on multiple fingers was expect to improve upon this. The results of the SUS also illustrated a difference between the two methods, with participants clearly preferring object representations with texture over outline drawings.

4.3 Single- and Three-fingers Experiment

In the second experiment, the main objectives were to evaluate whether using multiple devices/fingers would improve performance over using a single device/finger for both outline representations and textured representations. Stemming from the previous results (Section 4.2) with blind and sighted subjects, as well as analysis from previous literature, the two hypotheses: (1) there would be an improvement in accuracy from one

finger with the raised line drawings to one finger with the texture drawings due to the enriched representation used; although a shorter exploration time since would not occur since, in both cases, information would need to be processed serially; and (2) there would be an additional improvement between one finger and three fingers with the textured diagrams due to parallel processing, and here it was expected that both a higher accuracy and a faster response time would be observed.

The reasoning for the first hypothesis was that the results of this current experiment to agree with those of the previous one (Section 4.2). For the second hypothesis, with the addition of more devices working independently and simultaneously, it was anticipated participants could utilize parallel processing of texture information, similar to that observed by [42], facilitating the exploration of the object. Users would simultaneously (i.e., in parallel to the tracer finger) use the feedback provided by the other fingers/devices to provide coarse, vibrotactile information about any nearby parts that they felt. This was expected to lead to faster exploration times, resulting in a reduction in memory load, and improved object identification.

4.3.1 Method

There were two factors in this experiment, (a) the use of one device or three (single hand) and (b) objects with either outline or textured representations, which were crossed to give four conditions in a factorial design. The amplitude was set specifically for each participant at their 15 dB_{SL} prior to the start of the experiment using the method described in Section 3.4.4.3.2, in which their threshold was found for 10 Hz, then adjusted the amplitude to 15 dB higher, and the perceptual filter provided the equalization.

Participants then received brief (approximately 5 min) training on how to use the

device(s), and how to recognize the different spatiotemporal patterns and their associated encoded information. For each of the four test conditions, participants would explore and try to identify 8 different common objects, and both their response and their exploration time were recorded. The presentation order was counterbalanced for both the test conditions between subjects and for nearly all the objects within each set in a randomized order.

4.3.2 Participants

Seven participants took part in the experiment, all of which were either totally blind (3) or visually impaired (4). The mean age for the group was 51 ± 7 years. Participants were recruited from the local Rehabilitation Center for the Blind and Visually Impaired and were paid for their participation. Seven of the eight participants were right handed, and all of them reported some experience with Braille.

4.3.3 Results

From the data, it was observed that both for one- and three-finger device setups with outline representations, participants had 20% accuracy in object identification. The one-finger device setup with the textured representations had 32% accuracy, which is lower than the previously found 44% accuracy (Section 4.2.3). Meanwhile, the three-finger device setup with the textured representations had 52% accuracy. Repeated-measures GLM ANOVA with two within factors (number of fingers, type of representation) found significance for the representation method [$F(1,49) = 15.46, p < .001, p. \eta = .49, \text{obs. power} = .971$], while number of fingers and method*finger interaction both showed nominal insignificance [$F(1,49) = 3.27, p = .077, p. \eta = .25, \text{obs. power} = .426$] and [$F(1,49) = 3.49, p = .068, p. \eta = .26, \text{obs. power} = .207$], respectively. These low, but

not significant p-values prompted me to use matched-pair t-tests ($\alpha = .05$, $df = 55$) to compare the two outline cases, the two textures cases, the two three-finger cases, and the two one-finger cases, directly. The matched-paired t-test results showed no significance between identifications for the two outline conditions ($p = 1.0$, $r = .99$), but did show significance for the two textures cases ($p = 0.02$, $r = .306$). Identification rates between the two three-finger cases (outline versus “texture”) were found to be significant ($p < .001$, $r = .45$), but the one-finger cases did not show a significant difference ($p = .07$, $r = .24$).

The observed odds ratios for method and fingers showed that users were 1.95 times more likely to correctly identify the object using the textured method over the outline method, regardless of number of fingers used, and 0.59 times more likely to correctly identify the object using three fingers over one, regardless of the method used. Breaking these ratios down into their individual comparisons, when using three fingers with textures, participants were 3.39 times more likely to correctly identify the object over either outline condition, and 1.27 times more likely than using one finger with textures. When using just a single finger with textures, participants were 0.94 times more likely to make a correct identification than with either outline condition.

The mean exploration times, with mean standard errors, for the four conditions were: 269 s \pm 21 s for one-finger, outline; 249 s \pm 22 s for three-finger, outline; 224 s \pm 17 s for one-finger, textured; and 154 s \pm 12 s for three-finger, textured. Repeated measures GLM ANOVA with two within factors (number of fingers, type of representation) showed significance for both number of fingers and method [$F(1,49) = 14.07$, $p < .001$, $p. \eta^2 = .47$, obs. power = .96] and [$F(1,49) = 28.28$, $p < .001$, $p. \eta^2 = .61$, obs. power =

.99], respectively. Matched-pair t-tests ($\alpha = .05$, $df = 55$), showed that the time differences between the three-fingers with textures was significant to all other conditions ($p < .001$, $r > .5$), while the one-finger textures case was significant from one finger outline case ($p = .016$, $r = .58$), but not the three-finger outline case ($p = .149$, $r = .64$). The comparison between the two outline cases was again, not significant ($p = .278$, $r = .66$). These results are shown in **Figure 37**.

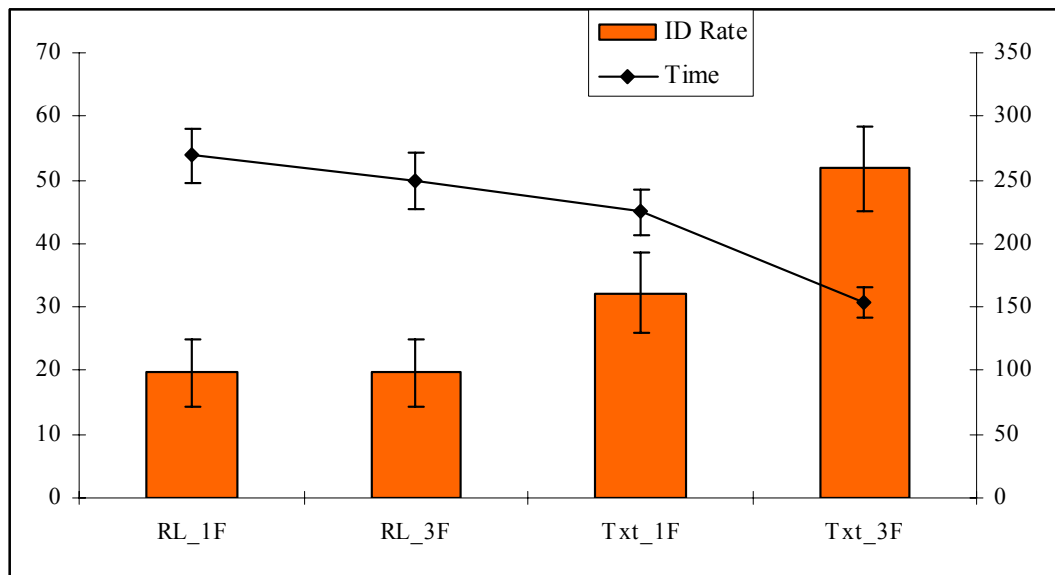


Figure 37: Mean Accuracies and Exploration times, with std error mean bars, for the 4 test conditions

4.3.4 Discussion

The results showed a clear performance trend of improved performance both for object identification and exploration time using additional fingers/devices when objects are represented with our “texture” method, supporting the second hypothesis. While both [12] and [102] similarly saw a performance increase when using multiple fingers [12] or points of contact [102], both used contour information only, and [102] did not look at a one-finger condition. Similar to [12], this performance increase could result from either parallel integration of the texture-encoded feedback or guided exploration; however,

provided the use of “texture” feedback by my system, this alternative hypothesis remains unlikely.

The case for parallel integration is supported by the fact that a similar increase in outline performance was not seen when the number of available fingers increased. If kinesthetic cues are alone sufficient to guide exploration, then the additional fingers should have provided some benefit, as they did in [12]. Instead, the results from this experiment show that the performance change for the outline cases best matched those of [62] or [103], supporting the hypothesis that spatiotemporal information in the “textures” used by my system are processed in parallel, similar to real textures [42, 43], whereas purely spatial information, like that available in raised-line drawings, or in [62], does not. However, the role that the saliency of information may play in guided exploration is unknown, or even how results of constraining feedback for 3-D object recognition [44, 101] are applicable when looking at 2-D representations that use complex vibrotactile feedback to render images. If guided exploration were to play a role, it would have expected to see a significant decrease in time between the two outline cases; however, this was not the case, again suggesting that no mechanism of guided exploration is at play here.

The influences of using sighted versus individuals who are blind on the results remain unclear; [12] and [62, 101] both used sighted subjects, whereas only individuals who are blind or visually impaired were used. Overall, there was not enough supporting evidence from this experiment to support the hypothesis of guided exploration, though, this may be something to consider in the future.

The results failed to confirm the first hypothesis, as a statistically significant improvement in object identification between raised-line and our textured representations was not observed when using the one-finger setup. Additionally, the observed accuracy using a single-finger device on textured representations was notably lower than previously found: 32% versus 44% [Section 4.2.3]. The most likely explanation for the discrepancy is that one or more individuals in the prior experiment had more skill with either tactile diagrams or shape recognition, although this remains unclear.

When compared to the results found in [65], the three-finger device using textured representations did perform within their accuracy range (52% compared to 50-68%), without the benefit of bare fingers or multiple hands. As several of the participants commented that they felt using two hands would be better, this will most likely be the next direction taken with the haptic display system.

4.4 Using Multiple fingers and Multiple hands Experiment

In the final experiment, the final hypothesis, which was that multiple hands and multiple fingers would have improved performance over using a single hand with the textured representations, was examined. In this experiment, the system performance for two factors: number of hands, and number of fingers per hand, was evaluated. Previously single- and three-finger device performance was examined for two methods of object representation: the traditional method of using an outlined representation of a figure, or the method developed in this project, which used a “textured” representation of the figure, where part identity and part orientation relative to its perspective are encoded using different, highly salient “texture” patterns. Participants were much more likely to identify the object using three fingers/devices with the “texture” representations than they

were using just a single-finger, or when using either a single- or three-fingers with the outline representations. The only question left remaining was how to optimize the design: what number of devices works best, and how should the user wear them?

First of all, to examine these questions, it was no longer felt that testing the outline object representations had any further experimental value. In none of my experiments using depictions of real objects did participants perform well using the outline representations, even when using additional fingers/devices. Trying to optimize the system design with the outline representations, even just to show again its failings, would be a waste of time; therefore, it was decided to only use the “texture” object representations for this final experiment.

Next, there was the general question of how many fingers/device per hand would be used. Obviously, the factor of number of hands is limited to either one or two, but the number of devices that could be worn on the hand may range from one to five. A 5x2 factorial experiment would either require a lot of within-condition repetitions (a very long experiment), or a lot of subjects. Considering the limited regional subject pool for individuals who are blind and visually impaired, and also that they are humans with limits on their available time, neither of these experimental designs were very attractive. Instead, a limit to the number of variables within the factor of fingers tested in the main experiment through pilot experimentation was sought.

4.4.1 Pilot 1

Six participants who were blind or visually impaired, all of whom had previous experience in one or more of my experiments, took part in the experiment. The first three participants were piloted with using four-, three-, two-, and one- device on a single hand,

first on identifying the texture patterns, second with tracing a few lines, third identifying a few geometric primitive objects, and finally, identifying two real object textured representations. While training them, all three participants asked not to use all four devices, stating reasons of comfort: all four devices limited their ability to move fingers independently; and saliency: most felt that they just ignored the feedback on the last digit or raised the finger off the screen to turn it off. Rather than continue with using four fingers per hand for the rest of the pilot, the experiment was changed, and instead looked at using just three-, two-, and one-finger on a single hand, and both hands with one-finger each. All participants were retrained using the same method previously described, and then tested using 2 objects per condition. The objects were drawn from the same pool of 48 objects described Section 4.1.1. The results showed no difference between the identification performance for two- or three-fingers using a single hand (42%), compared to 25% for one-finger and a single-hand and 50% for one-finger on each hand. A matched-pair t-test was performed for each set, but only found that the one-finger on one hand versus one-finger on both hand cases were significantly different ($df = 11, p = .039, r = .448$). Despite the small correlation and low degrees of freedom, it was concluded that there was not enough difference shown between using two or three fingers per hand to justifying using the increased number of devices.

4.4.2 Main Experiment

4.4.2.1 Method

For the main experiment, two factors each with two variables were tested: use of one or two hands, and the use of one or two fingers per hand. The results of the previously described (Section 4.4.1) pilot experiment showed no advantage in using more than two

fingers on a single-hand; however, the use of just one-finger on one hand, while also of no performance advantage, served as a control and a measure of relative performance across the three main experiments. Using the results of the experiment detailed in Section 4.3, it was estimated that using 12 objects per condition, it would necessary to have at least 12 subjects to have an experimental power of 0.80. The objects were made, selected, and organized into four sets for the experiment using the process described in Section 4.1. A total of 48 objects were used in the experiment [See Appendix for list].

Training for the experiment followed the example of the previous pilot. Participants were trained first on how to recognize the individual texture patterns and taught how the pattern encoded information. Second, they practiced tracing a few lines, and outline patterns, to practice following a contour. Then, participants identified a few geometric primitive objects, such as squares, circles, and higher-sided polygons. Finally, they were tasked with identifying a few real object textured representations, which were not presented again in the experiment, and gave them feedback.

During the experiment, participants were seated at a table with a computer screen placed horizontally on the table in front of them. Throughout the experiment, the participants were required to wear a blindfold unless they had a total loss of vision; they could request a break at any time, during which they could remove the blindfold. Participants were required to take a break about every 15 minutes, to prevent tactile adaptation from dulling their perception of the “textures”.

In the experiment, the participants were tasked with wearing devices according to each of the four test conditions: one-hand/one-finger, one-hand/two-finger, two-hand/one-finger, or two-hand/two-fingers. Conditions and objects were presented in a

random, completely counterbalanced order. Then, an object was presented on the screen, and the participant was instructed to explore the screen, using the tactile feedback to build a mental image of the object, and to guess what they felt the object was most like. They were told the objects were all common, everyday objects that would not be an overly superlative nomenclature (knife versus butcher’s cleaver, et cetera), that they were free to guess anything, and take as much time as they felt they needed. After they made their guess, their responses and the time spent exploring the object were recorded.

4.4.2.2 Participant Population

Twelve individuals who are blind or visually impaired took part in the experiment, eight of which had previous experience with the display device. **Table 16** shows participant information pertaining to age, their level of vision (total includes those who could see light and shadows, but nothing else), the age of onset for their vision loss (Congenital, Early (0-18), and Late (18+)), their use of Braille (Never/Learning, Sometimes, Regular), and similarly, their use of tactile pictures. Ten participants were right-handed and two were left-handed. Participants were paid \$10 per hour for their participation, exclusive of any compensation for travel expenses. The study received approval through the VCU Internal Review Board and all participants consented after having the experiment details explained to them, prior to the experiment taking place.

| Participant | Age | Vision | Onset | Braille | Tactile Pictures |
|-------------|-----|-------------|-------|---------|------------------|
| 1 | 54 | Tunnel (RP) | E | S | S |
| 2 | 51 | Total | E | R | S |
| 3 | 49 | Total | E | R | S |
| 4 | 52 | Total | L | R | S |
| 5 | 42 | Total | E | R | S |
| 6 | 40 | Total | C | R | S |
| 7 | 57 | Partial | E | S | S |

| | | | | | |
|----|----|---------|---|---|---|
| 8 | 53 | Total | E | S | S |
| 9 | 43 | Partial | L | N | N |
| 10 | 43 | 20/200 | E | N | N |
| 11 | 55 | Partial | E | N | N |
| 12 | 18 | Partial | C | N | N |

Table 16: Participant population descriptions.

4.4.2.3 Results

The results for both identification rate and exploration time are shown in **Figure 38**; H1F1 refers to one-hand/one-finger, H2F2 to two-hand/two-finger, and so forth. The identification accuracies were: one-hand/one-finger $36.8 \pm 4.5\%$; one-hand/two-finger $48.6 \pm 4.8\%$; two-hand/one-finger $54.9 \pm 4.5\%$; and two-hand/two-fingers $50.7 \pm 5.2\%$. GLM with repeated measures ANOVA found a significant effect for number of hands ($F(1,143) = 4.47$; sig = .045, eta = 0.173, obs. power = .556) and the interaction effect for hands*fingers ($F(1,143) = 4.09$; sig = .275, eta = 0.167, obs. power = .519), but not for the effect of fingers ($F(1,143) = 2.02$; sig = .275, eta = 0.089, obs. power = .193).

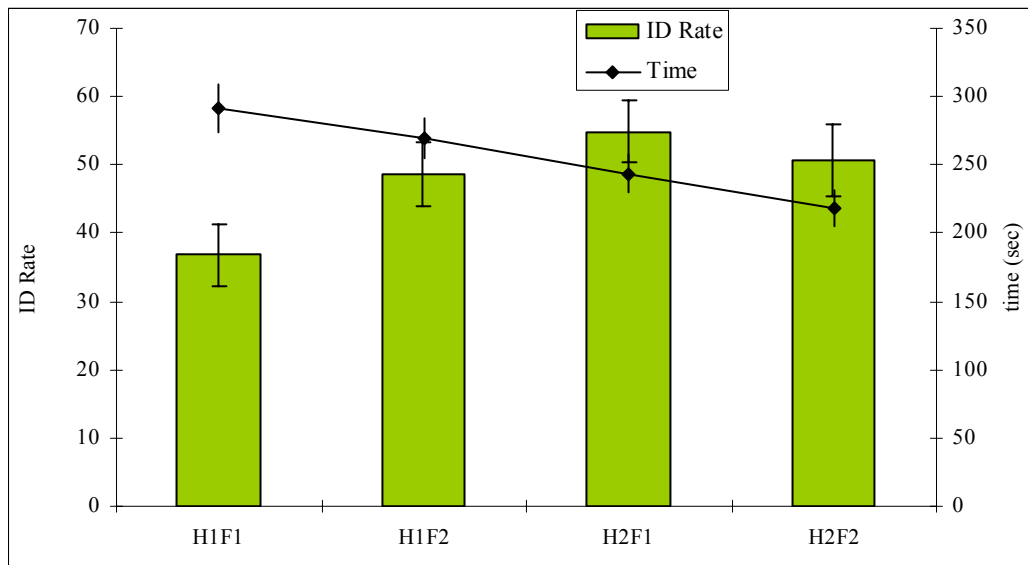


Figure 38: Results of the multiple-hand experiment for both ID rate and exploration time. Error bars indicate standard mean error.

Matched-paired t-test (**Table 17**) found significance between using a single finger on one hand, and all other cases, but found no other significance between the other condition comparisons for rate of object identification. The strong effect size and large Cohen's d values for the one-hand/one-finger versus multiple-finger cases shows that the experiment had enough power to make this assessment. The comparison between the one-hand/two-finger and two-hand/one-finger cases had near ($p = 0.07$) significance, which with an increased number of subjects or images may have significance. However, the mean performance difference between these two was only 6.3%, which shows only a very weak advantage difference.

From the calculation of the odds ratios for the conditions it was found that participants were 67% more likely to correctly guess the object for one-hand/two-fingers than one-hand/one-finger, 115% more for two-hands/one-finger than one-hand/one-finger, and 72% more for two-hands/two-fingers than one-hand/one-finger. The other comparisons reflected the observed lack of preference for any particular setup of multiple fingers; participants were only 28% more likely to correctly guess the object for two-hands/one-finger than one-hand/two-fingers, and 24% more than two-hands/two-fingers, while they were only 3% more likely for two-hands/two-fingers than one-hand/two-fingers.

| Pair | t | df | sig. 2tail | Cohen's d | r-value |
|-------------|--------|--------|------------|-----------|---------|
| H1F1 - H1F2 | -4.529 | 11.000 | 0.001 | -2.731 | 0.807 |
| H1F1 - H2F1 | -5.118 | 11.000 | 0.000 | -3.086 | 0.839 |
| H1F1 - H2F2 | -3.458 | 11.000 | 0.005 | -2.086 | 0.722 |
| H1F2 - H2F1 | -2.018 | 11.000 | 0.069 | -1.217 | 0.520 |
| H1F2 - H2F2 | -0.639 | 11.000 | 0.536 | -0.385 | 0.189 |
| H2F1 - H2F2 | 1.198 | 11.000 | 0.256 | 0.722 | 0.340 |

Table 17: Matched-pair t-tests for all mean ID Rate comparisons.

The results for exploration time are also shown in **Figure 38**. The identification accuracies were: one-hand/one-finger $296 \pm 18s$; one-hand/two-finger $269 \pm 15s$; two-hand/one-finger $243 \pm 14s$; and two-hand/two-fingers $219 \pm 13s$. GLM with repeated measures ANOVA found a significant effect for number of hands ($F(1,143) = 28.4$; sig = .000, eta = 0.407, obs. power = .999) and for the effect of fingers ($F(1,143) = 6.65$; sig = .011, eta = 0.21, obs. power = .726), but not for the interaction effect for hands*fingers ($F(1,143) = .022$; sig = .883, eta = 0.00, obs. power = .052). Matched pair t-tests found all comparisons aside from one-hand/one-finger with one-hand/two-finger to be significant; these results are shown in **Table 18**.

| Pair | t | df | sig 2-tail | cohen's d | r |
|-------------|-------|---------|------------|-----------|-------|
| H1F1 - H1F2 | 1.680 | 143.000 | 0.095 | 0.281 | 0.139 |
| H1F1 - H2F1 | 3.645 | 143.000 | 0.000 | 0.610 | 0.292 |
| H1F1 - H2F2 | 5.205 | 143.000 | 0.000 | 0.871 | 0.399 |
| H1F2 - H2F1 | 2.201 | 143.000 | 0.029 | 0.368 | 0.181 |
| H1F2 - H2F2 | 4.044 | 143.000 | 0.000 | 0.676 | 0.320 |
| H2F1 - H2F2 | 2.002 | 143.000 | 0.047 | 0.335 | 0.165 |

Table 18: Matched-paired t-tests for all comparisons of exploration times

4.4.2.4 Discussion

With this experiment, it was hoped that the fourth hypothesis, which was that multiple hands and fingers would perform better than a single hand would be supported, and that the device could be optimized for user performance, but the results of the matched-pair t-tests showed participants had no significant performance difference for identification rate between each of the multiple-finger conditions. This conclusion is further supported by the relatively low odds that a user is more likely to correctly guess the object using one of the multi-finger conditions over another. Perhaps a significant difference for the two-hand/one-finger condition over the other multi-finger conditions

may be found if more subjects were tested, but it would not affect the odds ratio given the same identification rates, and practically, it would mean users would only gain a slightly better chance over the other methods. There was an observed significant effect between the multiple-finger conditions for exploration time, but again, the practical value of designing the system in a particular manner just for a 50 s average advantage may be without merit.

The single-finger condition did show improved performance for ID rate in this experiment (37%) versus the last one (32%), but not quite as good as the first experiment (44%). The range of performance (49-55%) for the multiple-finger conditions in this experiment match that found in the previous one (52%), supporting the findings that three-fingers does not perform any better than two. The multi-finger identification performance also again fell into the range observed by [65] for printed diagrams using TaxyForm.

Exploration times were notably longer in this experiment, due to the data of two participants; it was decided not to exclude their data because they did follow directions, even if they still took much longer than other participants. The two-finger/two-hand case did have the fastest exploration times; however, it had only a 19% improvement over the two-finger/one-hand case and a 10% improvement over the one-finger/two-hand case. Despite being statistically significant, it was doubted that this advantage justified doubling the number of devices in the system, as well as the associated increase in cost.

From these results, it was concluded that using multiple-fingers does perform better than using a single-finger, a reaffirmation of my previous results and similar to the findings of [12]. The advantage seen in this and the last experiment when using multiple

fingers/devices indicates that the highly salient “textures” used to render visual diagrams are processed in parallel, similar to how material properties are processed in parallel [42]. However, it was found that there is little performance gain using more than two points of contact, similar to the single-hand bare-finger(s) and finger(s) in a glove responses observed in [101]. There also was no observed detriment to using more fingers/devices, although this was only tested up to three-fingers on a single hand and two-fingers on each hand. Thus, using multiple fingers/devices does clearly show a benefit, but beyond that, how the user wears the multiple devices may be a choice of personal preference, rather than one of performance optimization.

Over the course of the main experiments, some trends were observed that may direct future work in this area. The first observation made was that experience and performance did not have as much correlation as expected. People who had repeatedly participated in my experiments, and had experience with many of the object representations, rarely performed better than average. Considering everyone received the same training for each of the experiments, had familiarity and learning with the device been a significant contributing factor, those with additional experience have above average in their performance, but this was not the case.

In this experiment, the highest identification rates and fastest times were achieved by first time participants, both of who had little to no prior experience with tactile images. It was not believed, however, that this observation suggests a negative correlation between experience and performance, but rather another factor not previously consider may have a significant influence on performance. This factor may be the individual’s ability to build a visual model of what is felt. The visual/spatial intelligence of the individual may play

an important role in the performance of this, or any, display system. Assessment of the participants' visual intelligence has never been made as part of any of my studies, so there was no supporting evidence for this hypothesis. However, this may become an important blocking variable for future studies to better explain subject variance.

5 General Discussion

In this dissertation, the work towards developing novel, dynamic display device to haptically render two-dimensional visual diagrams for individuals who are blind and visually impaired is presented. This was necessary to help provide individuals who are blind and visually impaired more effective access to visual information than compared to traditional methods, in a world where information is increasingly provided solely through visual media. Towards this goal, first a prototype was developed to validate the idea that a cheap, simple, and dynamic display could haptically render visual graphics. Next, that display was further developed to allow for multiple-fingers, and then “texture” feedback to harness the strength of the haptic system. Then the use of “textures” was validated to encode information lacking on traditional tactile diagrams that could improve user performance. Subsequently, it was showed that the low-cost, multi-finger display developed, with its “texture” encoded graphics, performed better than raised-line diagrams. And finally, different system configurations of multiple fingers and hands were evaluated to optimize the performance of my design.

5.1 Approach

Traditionally, to render a visual image tactually meant making a raised-line diagram by first reducing a visual image down to an outline representation, printing that diagram out on special swell-paper, then heating up the paper to create the raised-line diagram. Despite this work, these diagrams had poor performance both anecdotally from educators interviewed, as well as experimentally [11-14]. In open-ended identification tasks, identification rates of 20-30% of common objects were typical [11, 13].

However, if the haptic system can identify real objects quickly and accurately [12], why does it perform so poorly for outline 2-D representations? Magee and Kennedy had found [60] that the identification of 2-D outline diagrams was the same if the users had only the sense of position corresponding to the contours of the image, or if they felt the object depictions themselves. This result indicated that serial processing of contour information via the kinesthetic sense played more of a role in the task than did the tactile information received. Serial processing, in contrast to parallel processing, is very time consuming, memory intensive, and cognitively demanding [42]. Compared to the performance with real objects, where it was found objects could be accurately identified with little to no time for detailed contour tracing [15], the haptic mechanisms that are used to perceive raised-line drawings are grossly inadequate.

However, Thompson and her colleagues [65] saw 2-4 times better performance using very simple “textures”, which better segmented parts and labeled their orientation relative to viewing perspective, despite essentially using raised-line diagrams. Even before Thompson and TaxyForm, tactile experience pictures had used similar methods of part segmentation and information encoding to great success, albeit with a great time cost in fabrication [6]. These results suggested that object identification could vastly be improved if only “textures”—specifically those that encoded object/part information—were made part of the tactile diagram.

Indeed, in examining the psychophysical function of haptics, Klatzky and Lederman also found that perception of coarse intensive discriminations, which include material properties, were processed early on and in parallel across multiple fingers [42]. In contrast, spatial information, such as geometric contours, is processed serially. In

addition, providing less ambiguous means to segment an object into parts, a problem with raised-line drawings [16], and providing a means to encode object/part orientation relative to a more intuitive viewing perspective [65], had already been shown to grant a performance boost. However, despite the apparent advantages using texture-encoded information had over current methods, this concept had never been applied to a dynamic display. The advantage of dynamic display over TaxyForm or a tactile experience is rather straightforward: why print, or manually create a diagram, when you can access any number of graphics using a refreshable device? In media enriched and highly interactive environments, remaining tethered to tactile printer that is neither portable nor relatively cheap is not a means of providing acceptable access.

5.2 System Design

The inclusion of texture-encoded information as a means to represent 2-D visual imagery spurred the development of the display system presented here; in fact, the time difference between Thompson and her colleagues publishing their results in [65] and my own development was only weeks. However, for the purpose of this project, the system development was not simply taking TaxyForm and directly using it to create digital images; rather, in addition to addressing the limitations of TaxyForm given above, the neurobiological and psychophysical factors that govern haptic perception in the first place were considered. Namely, Thompson's subjects and those using tactile experience pictures had the advantage of bare skin, multiple fingers, and both hands, but any device will place constraints on perception of stimuli, depending on the method of actuation [44].

In order to be successful, the design needed to not only preserve the quality of textured information, but also needed to accommodate for the use of multiple fingers. The feedback saliency was the biggest concern with designing for multiple fingers, as feedback that was less discriminable is processed serially rather than in parallel [42, 43, 45, and 46]. Independent movement between fingers was also seen as being potentially beneficial [12]. Finally, the device itself would have to allow for independent, simultaneous operation of multiple devices.

The first prototype developed did not focus on texture-like output or multiple fingers. Instead, it validated the basic idea that a haptic display device rendering 2-D visual information could be made affordable, portable, and intuitive. While this may seem like a self-evident research question, at this point, few devices outside those developed for research purposes had even attempted to render 2-D visual information for individuals who are blind and visually impaired, and among the main ones, the Optacon and NIST's display [67], neither cost less than \$5000 at the time (the Optacon cost several thousand back in the 1970's, and the NIST's display was never mass produced). Developing a device that could render visual images and cost only \$30 off-the-shelf to produce, while fundamentally simple, still illustrated the relatively unexplored potential in this area.

Although novel, low-cost, portable, and intuitive, the first and second prototypes lacked the capacity to output "texture" feedback, so to meet the project goal, a new design was necessary. For the final design, the optical sensor was changed from a photointerrupter capable of only sensing grayscale to one that could sense RGB color, and the actuator from a solenoid motor to a piezoelectric speaker. These decisions were made so that (1) the device could use color to encode the more complex vibrotactile

patterns that would be used to generate “texture” sets, and (2) the actuator could generate the more complex vibrotactile patterns.

The piezoelectric speaker was chosen over a distributed display like a Braille cell, because while the Braille could possibly provide more information [101], although there has been no test performed comparing a distributed display and a point-contact display using vibrotactile feedback. One problem with the Braille cell is that its dimensions (84mm x 6mm x 19mm) make it difficult to extend to a wearable design that allows for multiple independently moving fingers. Additionally, a Braille cell cost an order of magnitude more than the single-point contact actuators considered. In the future, this may be something that is looked at further: does the independent movement of between the fingers contribute much to performance and does distributed feedback provide enough of a performance improvement over a point-contact to justify the cost difference between the two?

The final display design developed [Section 3.4] met all the fundamental cost, size, and intuitive characteristics. Aside from the custom built plastic case, each device cost around \$40 in parts and weighed around 100g. However, this prototype required a good deal of additional circuitry, especially with the perceptual correction filter, and was still tied to a computer running LabView and an external power supply. These were not seen as permanent design choices, rather ones of temporary convenience, as the LabView program could be replaced with a microprocessor and the power supply with a battery.

The design remained highly intuitive, as had the previously model been: just put the device on, place your hand on the screen, and start exploring. The push-button switch controlled the power, turning on when the user contacted the screen, so the user never

had to hunt for an on or off switch. The device could also be calibrated to a monitor's particular color settings, and then would continue to work as long as those settings did not change. Repeatedly, the device scored well in the System Usability Survey, and no participant ever indicated that the device was too technically challenging for them to use.

In Section 3.4 described how the device satisfied the established criteria for spatial (<5mm) and temporal resolution (<20ms). It could detect a wider range of colors than necessary for my use, and could actuate over the necessary bandwidth for both amplitude and frequency. Owing to the perceptual interaction between these two dimensions [35-37], a perceptual correction filter was developed, similar to [100], but using the guidelines established by [106] and the methods from [23] and [107]. This filter acts very similar to an equalization filter for an audio system, as it matches the attenuation based on a perceptual magnitude curve, in this case for a 15 dB SL, creating a perceived flat magnitude response across frequency. Additionally, this filter compensated the actuator's resonance frequency by low-pass filtering the output past the bandwidth was intended to use to generate the "texture" feedback.

In section 3.5, it was described how a set of highly salient textures was developed, starting with fundamental dimensions of vibrotactile signals all the way to device specific cases of color/spatial-pattern generation. As previously said before, any "textures" for use with the system, particularly with multiple fingers, had to be sufficiently salient; otherwise, they get processed serially and would not see any advantage from using the additional fingers[42, 45, & 46]. Therefore, it was important not only to design textures so that they could be singularly interpreted by the information encoded, but to have the whole set of textures be highly salient amongst each other.

Testing of the possible dimensions from which a set of highly salient “textures” could be generated focusing on factors of evaluating confusion and discriminability between the dimensions. Of the temporal dimensions, it was continually found that temporal frequency was the most salient dimension, followed by waveform modulation, and waveform shape. These results agreed with some previous research [81, 90, 108], but not with others [88]. As phase was not looked at as a dimension, it should be noted that Cholewiak and his colleagues found the haptic system, at least for gratings, was insensitive to changes in phase [109].

Using multi-dimensional scaling, it was found that among temporal (frequency and duty cycle) and spatial (period and duty cycle) dimensions, temporal frequency and spatial duty cycle were the most discriminable. Since the connection between the comparative discriminability of different dimensions and their ability to encode information was ambiguous, a more directly applicable testing was performed. Using nonsense objects constructed from possible mixed dimensions, it was found that frequency modulated, spatially directed patterns worked best. The results supported previous research [108] that temporal frequency was the most salient dimension after correcting for amplitude, as seen by the very high performance for part identification. Although it was observed that some performance difference between amplitude and frequency modulated patterns existed, this did not necessarily conflict with [90] and [91], as they both used temporal modulation, whereas the patterns used here were modulated along mixed spatial and temporal dimensions.

The results also supported the previous conclusion from [110], which was that performance improves when an object or task consists of fewer items. An oblique effect

[111] was seen for the diagonally direct spatial pattern, leading to confusion between it and the other two spatially directed patterns. However, what effect this would have on the depictions of real objects was unclear, as real objects often would not frequently mix all of the different orientations together.

5.3 Testing of Main Hypotheses and System Evaluation

In Section 4, the experiments performed on depictions of real objects to sequentially test my four main hypotheses: (1) “texture” representations of objects would perform better than outline representations using even a single device/finger, (2) that there would be no performance advantage to using more than a single device/finger for the outline representations; (3) there would be a significant performance advantage to using multiple devices/fingers for the “texture” object representations; (4) that multiple devices distributed on each hand would perform better than a multiple devices on a single hand, were described in detail.

Evidence of parallel processing for real objects [15, 44] and for certain tactile properties [42-46], gave support to the idea that the performance discrepancies seen between the textured images and outline images was rooted in how the different features were perceptually integrated. The success of tactile experience pictures [6], and TaxyForm [65], further suggested that how the objects diagrams were represented with regard to part identity and perspective may play a role as well. Identification of raised-line drawings were shown to rely just as much on kinesthetic mapping of the contour as anything else [11, 60], and this geometric contour information had been shown to be processed serially [42].

From these directions, a design paradigm was established that “texture” representations, which more clearly segmented object parts and labeled their relative orientation, would perform better than their outline counterparts. Furthermore, users would see an additional performance boost when using multiple fingers/devices due to the parallel processing of coarse information with the “texture” representations, but not the outline representations. After design a system capable of generating the “texture” feedback and extending to multiple fingers/devices, the veracity of this theory was evaluated through the first three of my main hypotheses.

The final hypothesis stems more from a standpoint of system optimization than application of psychophysical research. In fact, little in the way of applicable research exists: Craig [57] observed that identification using one finger on each hand was better versus two fingers on the same hand when identifying a spatial pattern, and Jansson and Monaci [62] found no difference using two hands versus one when identifying a contour map image, but both of these tasks are different from each other, as they are from what was intended to do with the device. Therefore, it was concluded that looking at two-hand performance crossed with multiple fingers would not only represent an important system design evaluation, but also would possibly contribute to the broader field of haptics research.

For the first hypothesis, it was found object identification rates of approximately 30-40% for “texture” compared to 20% for outline representations, though no significant reduction in exploration time was observed. This result was expected, as while the “texture” representations were more easily interpreted, by using only a single finger/device, only serial processing of information was possible. The second hypothesis

was also confirmed, as one-finger and three-finger performance for the outline representations were the exact same. Contrastingly, the “texture” representations using multiple fingers/devices saw an approximately 50-60% compared to 30-40% for a single finger/device. (The ranges represent performance rates over several experiments).

The final identification rates (50-60%) for the multi-finger display using texture encoded information were not only much higher than the performance seen with raised-line diagrams in open-ended identification tasks [11, 13, 14, 17, 65], but also fell within the performance range (50-70%) for TaxyForm [65]. However, no performance difference of practical application was seen between the multiple hands/multiple fingers cases. While there was no strongly supporting evidence for the fourth hypothesis, this has an overall positive affect on the system design optimization: only two devices (versus more, at added cost) are really needed, and the user can choose how to wear them without it significantly affecting their performance.

These results not only show the success of my proposed system design, and the method of using “textures” to more clearly represent 2-D diagrams tactually compared to raised-line drawings, but also further support several paradigms about haptic devices in general. First is evidence in support of the theory that contour information gets processed serially [42], and most likely through kinesthetic information [60] rather than tactually, and therefore, does not benefit from more than a single finger or contact, as has been seen for raised-line diagrams [17, 62].

Secondly, there has not been a strongly supported paradigm for why tactile experience pictures [6], “textured” images [64], or TaxyForm [65] perform better than outline-representations. It was believed that texture (tactile experience pictures), or

texture-like properties (my system, [64], and [65]) distinguish themselves for two reasons: (1) they allow certain information to get integrated in parallel across multiple fingers, and (2) they provide a more clear projection of the visual image than outline representations. The results further support the conclusion that the texture information is indeed processed in parallel, as it was observed that both an increase in identification rate and a decrease in exploration time when using multiple fingers/devices over a single finger/device.

A possible alternative conclusion for the observed performance gap between the single and multiple finger/device(s) is that participants were able to use one device to guide the exploration of the other, as suggested by Klatzky and her colleagues [12]. This mechanism of guided exploration could provide a more efficient means of performing a serial integration task, but this idea has not been developed further. The biggest argument against this alternative conclusion is that no difference between the single and multiple finger/device cases was observed for the outline representations, whereas [10] did see an improvement when going from 1 to 5 fingers.

Additionally, it was believed that the clearer image projection used by the texture object representations contributed to the performance increase over the outline representations. Most outline representations of objects were represented in two-dimensions using an isometric projection, which allowed some preservation of a three-dimensional perspective, at the cost of potentially occluding some of the object or its parts. Thompson and her colleagues [64, 65], as well as Erdman [6] and Janice Johnson, an expert in tactile diagrams for K-12 education, all recommended against this practice.

Instead, these resources concluded that a frontal view of an object was preferable for tactile diagrams, with occluded parts shown extending out from their points of connection, e.g., the four legs of a table shown splayed from the table top. This allowed participants to feel all the major parts of the object and build a better mental model for the object they felt, making identification that much easier. By segmenting the parts and indicating their three-dimensional orientation relative to the frontal prospective, individuals were provided with additional information to reduce the ambiguities that plague outline representations [16]. However, the improved object representation would not explain any performance increase seen with the use of multiple-fingers; the only explanation for that is the more efficient processing of the “texture” feedback.

Lastly, the data and results throughout the experiments using depictions of real objects showed one other important characteristic, seen through the observed high variance, and observationally: some participants performed very poorly using tactile diagrams and some performed very well. As discussed in Section 4.4.2.4, often the participants who performed the best did not have the most experience with the device, or tactile pictures in general, and some of the participants with the most experience typically performed below average. Since during training all of the participant’s ability to sense and use the vibrotactile feedback was assessed, it was not believe that the perceptual difference explain the difference observed. Instead, it was believed that cognitive differences, specifically visual literacy (or intelligence) of the individual may explain why some individuals perform exceedingly well using the display system with little or no prior experience, while others, even after participating six or more times, still struggle. In any of the experimental cases, participants are required to build a mental model through

integrating all the feedback obtained during exploration, and then compare this model with their memory to see if there is anything that matches. However, if the individual has trouble building this mental model or lacks memory of a similar shape, then their performance deficiency is not owing to any inadequacy of the display system necessarily, but rather their cognitive limitations.

It may be that some of the poor participant performance occurred simply because those individuals had fewer opportunities to develop their visual literacy during their developmental phase growing up, a situation that the display system presented here may one day help correct. Anecdotally, when exploring a TaxyForm image of a boat made [Section 4.1.1.1], after wrongly guessing the object, the participant asked, “what was that.” A boat. “Oh, I’ve never felt or seen one those before. Is that how they look?”

5.4 Conclusion

The Introduction Section began with the example of Kevin Carter’s 1993 picture: *A vulture watches a starving child*, and my Background section discussing visual literacy, and how it connects us to an important part of our culture. The implication of not providing adequate access to visual information is that individuals who are blind or visually impaired may get isolated from culture and society. The haptic display system presented here not only attempts to bridge that gulf in access, it also helps expand the knowledge base for haptic devices using texture encoded feedback and multiple fingers. It was shown how a low-cost and portable haptic display can not only generate highly salient, texture enriched feedback for the interpretation of 2-D visual graphics, but that it can perform on par with some of the best existent methods, and two-to-three times better than the standard method of raised-line diagrams.

Future questions that remain are: (1) does the use of a distributed display versus a point-contact confer any advantage, (2) is participant performance affected by having independently moving or fixed fingers, (3) whether parallel processing or guided exploration are the perceptual mechanism involved when using more than a single finger/device, and (4) does participant visual literacy/intelligence influence performance enough that it should be treated as a blocking variable in future experiments?

5.5 Support

This work was supported in part by NSF IIS grant #0712936 and an A.D. Williams grant.

6 References

- [1] http://www.cdc.gov/visionhealth/basic_information/vision_loss.htm. Accessed September 2011.
- [2] http://www.nfb.org/nfb/Braille_Initiative.asp Accessed Sept. 2011.
- [3] B. N. Walker & L. Mauney (2010). “Universal Design of Auditory Graphs: A Comparison of Sonification Mappings for Visually Impaired and Sighted Listeners.” *ACM Transactions on Accessible Computing* 2, 3. Article 12 March 2010, 16 pages. DOI = 10.1145/1714458.1714459. <http://doi.acm.org/10.1145/1714458.1714459>
- [4] U.S. Census Bureau. Americans with Disabilities. <http://www.census.gov/prod/2008pubs/p70-117.pdf> Accessed September 2011.
- [5] ACRL Visual Literacy Competency Standards for Higher Education. <http://www.ala.org/acrl/standards/visualliteracy>. Accessed March 2012.
- [6] P.K. Erdman. *Tactile Graphics*. American Foundation for the Blind, New York, 1992, pp. 41-119.
- [7] D. Katz (1989) *The world of touch [Der aufbau der tastwelt]*. (L. Krueger, Trans.). Mahwah, NJ: Erlbaum. (Originally published in 1925)
- [8] T. P. Way and K. E. Barner (1997). “Automatic Visual to Tactile Translation—Part I: Human Factors, Access Methods, and Image Manipulation.” *IEEE TRANSACTIONS ON REHABILITATION ENGINEERING*, VOL. 5, NO. 1, pp. 81-94.
- [9] T. P. Way and K. E. Barner (1997) “Automatic Visual to Tactile Translation—Part II: Evaluation of the TACTile Image Creation System.” *IEEE TRANSACTIONS ON REHABILITATION ENGINEERING*, Vol. 5, No. 1, pp. 95-105.
- [10] S. E. Krufka, K. E. Barner, and T. C. Aysal. (2007) “Visual to tactile conversion of vector graphics,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, Vol. 15, pp. 310–321.
- [11] S. Lederman, R. Klatzky, C. Chataway and C. Summers (1990). “Visual Mediation and the Haptic Recognition of Two-Dimensional Pictures of Common Objects.” *Perception and Psychophysics*, 47 (1), Austin, Texas, pp. 54-64.
- [12] R. Klatzky, J. Loomis, S. Lederman, H. Wake, and N. Fujita (1993). “Haptic identification of objects and their depictions.” *Perception and Psychophysics*, 54 (2), pp. 170-8.

- [13] M. Heller, J. Calcaterra, L. Tyler, and L. Burson (1996). "Production and interpretation of perspective drawings by blind and sighted people." *Perception*, 25, pp. 321-34.
- [14] R. Hopkins. (2000) "Touching pictures." *British Journal of Aesthetics*, Vol 40, pp. 149-167.
- [15] R. Klatzky, S. Lederman, and A. Metzger (1985). "Identifying objects by touch: An 'expert system.'" *Perception & Psychophysics*, 37, 299-302.
- [16] J. Kennedy and J. Bai (2002). "Haptic Pictures: Fit Judgments Predict Identification, Recognition Memory and Confidence." *Perception*, 31, 1013-1026.
- [17] J. Loomis, R. Klatzky, and S. Lederman (1991). "Similarity of tactual and visual picture-recognition with limited field of view." *Perception*, 20(2), pp. 167-77.
- [18] K. Johnson. (2001) "The roles and functions cutaneous mechanoreceptors." *Current Opinion in Neurobiology*. 11, pp. 455-61.
- [19] F. Vega-Bermudez, K. Johnson (1999). "SA1 and RA receptive fields, response variability, and population responses mapped with a probe array." *Jo. of Neurophysiology*, 81, pp. 2701-10.
- [20] M. Dodson, A. Goodwin, A. Browning, and H. Gehring (1998). "Peripheral neural mechanisms determining the orientation of cylinders grasped by the digits." *Jo. of Neuroscience*, 18, pp. 521-30.
- [21] M. Paré, A. Smith, and F. Rice (2002). "Distribution and Terminal Arborizations of Cutaneous Mechanoreceptors in the Glabrous Finger Pads of the Monkey." *Jo. of Comparative Neurology*, 445, pp. 347-59.
- [22] A. Brisben, S. Hsiao, and K. Johnson (1999). "Detection of vibration transmitted through an object grasped in the hand." *Jo. of Neurophysiology*, 81, pp. 1548-58.
- [23] R. Verrillo, A. Fraioli, and R. Smith (1969). "Sensation magnitude of vibrotactile stimuli." *Perception and Psychophysics*, 6 (6A), pp. 366-72.
- [24] U. Prokx and S. Gandevia (2009). "The kinaesthetic senses." *Jo. of Physiology*. 587, pp. 4139-46.
- [25] J. Gregory, D. Morgan, and U. Proske (1988). "Aftereffects in the responses of cat muscle spindles and errors of limb position sense in man." *Jo Neurophysiology*, 59, pp. 1220-30.
- [26] G. Goodwin, D. McCloskey, and P. Matthews (1972). "The contribution of muscle afferents to kinaesthesia shown by vibration induced illusions of movement and by the effects of paralysing joint afferents." *Brain: A Jo. of Neurology*, 95, pp. 705-748.

- [27] D. Collins D, K. Refshauge, G. Todd, and S. Gandevia (2005). "Cutaneous receptors contribute to kinesthesia at the index finger, elbow and knee." *Jo. of Neurophysiology*, 94, pp.1699–1706.
- [28] R. T. Verrillo (1963). "Effect of contactor area on vibrotactile threshold." *Jo Acoustical Society of America*. 35, pp. 1962–6.
- [29] R. T. Verrillo (1968). "A duplex mechanism of mechanoreception." In: *The skin senses*. D. R. Kenshalo, editor. Springfield, IL: Thomas, pp 135–59.
- [30] S. Bolanowski, G. Gescheider, R. Verrillo, C. Checkosky (1988). "Four channels mediate the mechanical aspects of touch." *Jo Acoustical Society of America*, 84, pp. 1680–94.
- [31] G. Gescheider, S. Bolanowski, R. Verrillo (2004). "Some characteristics of tactile channels." *Behavioural Brain Research*, 148, pp. 35-40.
- [32] S. Weinstein (1968). "Intensive and extensive aspects of tactile sensitivity as a function of body part, sex, and laterality." *The skin senses*. D. R. Kenshalo, editor. Springfield, IL: Thomas, pp 195-222.
- [33] C. A. Perez, C. A. Holzmann, and E. Sandoval (1998). "Two point vibrotactile spatial resolution as a function of pulse frequency and pulse width." *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Vol. 20, No 5, pp. 2542-6.
- [34] S. J. Bensmaïa, J. C. Craig, K. O. Johnson (2006). "Temporal factors in spatial acuity: evidence for RA interference in fine spatial processing." *Jo. of Neurophysiology*, Vol 95, pp. 1783-91.
- [35] B. Taylor (1977). "Dimensional interactions in vibrotactile information processing". *Perception and Psychophysics*. 21 (5), pp. 477-81.
- [36] J. W. Morley and M. J. Rowe (1990). "Perceived pitch of vibrotactile stimuli: effects of vibration amplitude, and implications for vibration frequency coding." *Jo. of Physiology*, Vol 431. pp 403-16.
- [37] E. A. Roy and M. Hollins (1998). "A ratio code for vibrotactile pitch." *Somatosensory and Motor Research*, 15(2). pp. 134-45.
- [38] J. A. Harris (2006). "Psychophysical investigations into cortical encoding of vibrotactile stimuli." *Percept, Decision, Action: Bridging the Gaps: Novartis Foundation Symposium 270*. Vol 270.
<http://onlinelibrary.wiley.com/doi/10.1002/9780470034989.ch19/pdf>
- [39] M. Hollins, A. K. Goble, B. L. Whitsel, and M. Tommerdahl (1990). "Time course and action spectrum of vibrotactile adaptation." *Somatosensory and Motor Research*, 7, 205-221.

- [40] S. J. Bensmaïa, Y. Y. Leung, S. S. Hsiao, and K. O. Johnson (2005) “Vibratory Adaptation of Cutaneous Mechanoreceptive Afferents.” *Jo Neurophysiology*, 94, pp. 3023-36.
- [41] R. Klatzky and S. Lederman (1995). “Identifying Objects from a Haptic Glance.” *Perception and Psychophysics*. 57(8), pp. 1111-23.
- [42] S. Lederman and R. Klatzky (1997). “Relative Availability of Surface and Object Properties During Early Haptic Processing.” *Jo Experimental Psychology: Human Perception and Performance*, Vol 23 (6), pp. 1680-1707.
- [43] R. Klatzky, S. Lederman, and C. Reed. (1989). “Haptic integration of object properties: texture, hardness and planar contour.” *Journal of Experimental Psychology: Human Perception and Performance*, Vol 15, pp. 45-57.
- [44] S. Lederman and R. Klatzky (2004). “Haptic identification of common objects: effects of constraining the manual exploration process.” *Perception and Psychophysics*, 66(4), pp. 618-28.
- [45] M. Plaiser, W. Bergmann Tiest, A. M. Kappers (2008). “Haptic pop-out in a hand sweep.” *Acta Psychologica*. 128, pp. 368-77.
- [46] M. Plaiser, W. Bergmann Tiest, A. M. Kappers (2011). “The effect of feature saliency on haptic subitizing.” *Exp. Brain Research*. 209(1), pp. 29-34.
- [47] M. Hollins, S. J. Bensmaïa, K. Karloff, and F. Young (2000). “Individual differences in perceptual space for tactile textures: evidence from multidimensional scaling.” *Perception and Psychophysics*. 62(8), pp. 1534-44.
- [48] D. Picard, C. Dacremont, D. Valentin, and A. Giboreau. (2003) “Perceptual dimensions of tactile textures.” *Acta Psychologica*. 114, pp. 165-84.
- [49] S. J. Lederman, J. M. Loomis, and D. A. Williams (1982). “The Role of Vibration in the Tactual Perception of Roughness”. *Perception and Psychophysics*. 32(20), pp.109-16.
- [50] G. D. Lamb, (1983) “Tactile discrimination of textured surfaces: Psychophysical performance measurements in humans.” In *Human Hand Function*. Eds L. Jones and S. Lederman. Oxford University Press, New York.
- [51] S. J. Lederman and M. M. Taylor. (1972) “Fingertip force, surface geometry, and the perception of roughness by active touch.” Vol 12(5), pp. 401-8.
- [52] M. Hollins (2007). “The coding of roughness.” *Canadian Jo. of Experimental Psychology*.

- [53] M. Hollins, F. Lorenz, and D. Harper (2006). "Somatosensory coding of roughness: the effect of texture adaptation in direct and indirect touch" *Jo Neuroscience*. 26(10), pp. 5582-88.
- [54] H. Tan, C. Reed, L. Delhorne, N. Durlach, and N. Wan (2003). "Temporal masking of multidimensional tactual stimuli." *Jo of the Acoustical Society of America*, 114 (6), pp. 3295-3308.
- [55] A. Israr, H. Tan, and C. Reed (2006). "Frequency and amplitude discrimination along kinesthetic-cutaneous continuum in the presence of masking stimuli." *Jo of the Acoustical Society of America*, 120 (5), pp. 2789-2800.
- [56] R. Verrillo and G. A. Gescheider (1975). "Enhancement and Summation in the Perception of Two Successive Vibrotactile Stimuli. *Perception and Psychophysics*, 18, pp 128-136.
- [57] J. C. Craig (1985). "Attending to two fingers: Two hands are better than one." *Perception and Psychophysics*, 38(6), pp 496-511.
- [58] J. C. Craig and P. M. Evans (1987). "Vibrotactile masking and the persistence of tactual features." *Perception and Psychophysics*, 42(4), pp. 309-17.
- [59] J. C. Craig and P. M. Evans (1995). "Vibrotactile masking: the role of response competition." *Perception and Psychophysics*, 57(8), pp. 1190-1200.
- [60] L. E. Magee and J. M. Kennedy (1980). "Exploring Pictures Tactually." *Nature*, *Nature*, 283, pp 287-288.
- [61] A. D'Anguilli, J. Kennedy, and M. Heller (1998). "Blind children recognizing tactile pictures respond like sighted children given guidance in exploration." *Scandinavian Journal of Psychology*, 39(3), pp 187-190.
- [62] G. Jansson and L. Monaci (2003). "Exploring Tactile Maps with One or Two Fingers." *The Cartographic Journal*, 40 (3), pp 269-71.
- [63] J. G. Snodgrass and M. Vanderwart (1980). "A Standardized Set of 260 Pictures: Norms for Name Agreement, Image Agreement, Familiarity, and Visual Complexity". *Jo. of Experimental Psychology*. 6(2), pp. 174-215.
- [64] L. J. Thompson, E. P. Chronicle, and A. F. Collins (2003). "The Role of Pictorial Convention in Haptic Picture Perception". *Perception*, 32(7), pp. 887-893.
- [65] L. J. Thompson, E. P. Chronicle, and A. F. Collins (2006). "Enhancing 2-D Tactile Picture Design from Knowledge of 3-D Object Recognition". *European Psychologist*, 11(2), Hogrefe and Huber Publishers, pp 1-9.

- [66] M. Kurze (1997). "Rendering drawings for interactive haptic perception." *Conference on Human Factors in Computing Systems 1997*, Atlanta, USA. CHI 97 Electronic Publications.
- [67] "Tactile Graphic Display Device." http://www.nist.gov/public_affairs/factsheet/visualdisplay.htm. Accessed 2006, Site no longer accessible as of 2012.
- [68] L. A. Jones and N. B. Sarter (2008). "Tactile Displays: Guidance for their design and application". *Human Factors*, 50 (1), Santa Monica, CA, pp. 90-111.
- [69] V. Hayward and K. E. MacLean (2007). "Do it yourself haptics: parts 1 & 2." *IEEE Robotics and Automation Magazine*." Vol 14(4), pp. 88-119.
- [70] J. Bliss, M. Katcher, C. Rogers, and R. Shepard (1970). Optical-to-tactile conversion for the blind. *IEEE Transactions on Man-Machine Systems*. 11 (1), pp. 58-65.
- [71] J. C Craig (1977). "Vibrotactile Pattern Perception: Extraordinary Observers." *Science*. 196(4288), pp. 450-2.
- [72] H. ANDO, E. KUSACHI, and J. WATANABE (2007). "Nail-mounted tactile display for boundary/texture augmentation". *ACE '07*, ACM, Salzburg Austria, pp. 292-3, (June 2007).
- [73] G. STETTAN, et. al. (2007). "Fingersight: Fingertip Visual Haptic Sensing and Control." *IEEE International Workshop on Haptic Audio Visual Environments (HAVE)*, Ontario, Canada, 12-14 (October 2007).
- [74] K. Kyung and D.S. Kwon (2006). "Multi-sensory Perception of Roughness: Empirical Study on Effects of Vibrotactile Feedback and Auditory Feedback in Texture Perception". *ICAT*, pp. 406-415.
- [75] G. PETIT, A. DUFRESNE, V. LEVESQUE, V. HAYWARD and N. TRUDEAU. 2008. Refreshable Tactile Graphics Applied To Schoolbook Illustrations for Students with Visual Impairments. *ACM ASSETS 2008*, pp. 89-96.
- [76] R. RASTOGI, D. PAWLUK and J.M. KETCHUM (2010). "Issues of Using Tactile Mice by Individuals who are Blind and Visually Impaired". *Trans. on Neural Systems and Rehabilitation Engineering*.
- [77] J. OWEN, J. PETRO, S. D'SOUZA and D. PAWLUK (2009). "An Improved, Low-cost Tactile Mouse for Use by Individuals who are Blind and Visually Impaired." *ACM ASSETS 2009*, October 25-28, Pittsburgh, Pennsylvania, USA.
- [78] G. JANSSON ET AL. (1999). "Haptic Virtual Environments for Blind People" Exploratory Experiments with Two Devices." *International Journal of Virtual Reality*, pp. 8-17.

- [79] K. ROVIRA and O. GAPENNE (2009). "Tactile Classification of Traditional and Computerized Media in Three Adolescents Who Are Blind." *Journal of Visual Impairment and Blindness*. July, pp. 430-435.
- [80] S. A. WALL, and S. BREWSTER (2006). "Non-Visual Feedback for Pen-Based Interaction with Digital Graphs." *International Journal of Disability and Human Development*. 5 (3), pp 179-185.
- [81] J. M. WEISENBERGER, M.J. KRIER, M.A. RINKER (2000). "Judging the Orientation of Sinusoidal and Square-Wave Gratings Presented via 2-DOF and 3-DOF Haptic Interfaces." *Haptics-e*, 1 (4).
- [82] M. BIET, et al. (2008). "Discrimination of Virtual Square Gratings by Dynamic Touch on Friction Based Tactile Displays." *Symposium on Virtual Environments and Teleoperator Systems*. Reno, Nevada.
- [83] B. UNGER, R. HOLLIS and R. KLATZKY (2007). "JND Analysis of Texture Roughness Perception using a Magnetic Levitation Haptic Device." *Proceedings of World Haptics*, Tsukuba, Japan.
- [84] M. A. OTADUY and M.C. LIN (2004). "A Perceptually-inspired Force Model for Haptic Texture Rendering." *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization*, 1, pp.123-126.
- [85] K. Kyung, et. al. (2005). "How to Effectively Display Surface Properties Using an Integrated Tactile Display System." *Proc. IEEE, Intl Conf. on Robotics and Automation*. Barcelona, Spain, (April 2005).
- [86] R. KLATZKY and S. J. LEDERMAN (2006). "The Perceived Roughness of Resistive Virtual Textures: I. Rendering by a Force-Feedback Mouse." *ACM Transactions on Applied Perception*, Vol 3, No. 1, pages 1-14.
- [87] K. KYUNG, J. LEE, and J. PARK (2008). "Haptic Stylus and Empirical Studies on Braille, Button, and Texture Display." Hindawi Publishing Corporation. *Journal of Biomedicine and Biotechnology*, Vol 2008.
- [88] E. Hoggan and S. Brewster. "New Parameters for Tacton Design." *ACM. CHI 2007*. San Jose, California, 2007.
- [89] D. Ternes and K. Maclean (2008). "Designing Large Sets of Haptic Icons with Rhythm." *EuroHaptics 2008*. LNCS 5024, Pages 199-208.
- [90] A. MURRAY, R. L. KLATZKY, and P. KHOSLA (2003). "Psychophysical characterization and testbed validation of a wearable vibrotactile glove for telemanipulation." *PRESENCE: Teleoperators and Virtual Environments*, 12, pp. 156-182.

- [91] I. SUMMERS, J. WHYBROW, D. GRATTON, P. MILNES, B. BROWN and J. STEVENS (2005). "Tactile Information Transfer: A Comparison of Two Stimulation Types." *Journal of the Acoustical Society of America*, 118, pp. 2527-2534.
- [92] T.M. MASSIE, and J.K. SALISBURY (1994). "The phantom haptic interface: A device for probing virtual objects". *Proc. of ASME Haptic Interfaces for Virtual Environment and Teleoperator Systems* 1, pp. 295—301.
- [93] V. LÉVESQUE and V. HAYWARD (2008). "Tactile Graphics Rendering Using Three Laterotactile Drawing Primitives." *Symposium on Virtual Environments and Teleoperator Systems*. Reno, Nevada.
- [94] K. O. Johnson and J. R. Philips (1981). "Tactile Spatial Resolution. I. Two-point discrimination, gap detection, grating resolution, and letter recognition." *Jo. of Neurophysiology*. Vol 46(6), pp. 1177-91.
- [95] S. Okamoto, M. Konyo, S. Saga, and S. Tadakoro (2009). "Detectability and Perceptual Consequences of Delayed Feedback in a Vibrotactile Texture Display." *IEEE Transactions on Haptics*. Vol 2, No 2, pp. 73-84.
- [96] T. Moore, M. Broekhoven, S. Lederman, and S. Ulug (1991). Q'HAND: A fully automated apparatus for studying haptic processing of spatially distributed inputs. *Behavior Research Methods, Instruments and Computers*, 23(1), pp. 27-35.
- [97] J. Brooke (1996). "[SUS: a "quick and dirty" usability scale](http://www.usabilitynet.org/trump/documents/Suschapt.doc)". In P. W. Jordan, B. Thomas, B. A. Weerdmeester, & A. L. McClelland. *Usability Evaluation in Industry*. London: Taylor and Francis.
<http://www.usabilitynet.org/trump/documents/Suschapt.doc>.
- [98] International Standard ISO 5349-1. (2001). "Mechanical vibration – Measurement and evaluation of human exposure to hand-transmitted vibration. Part 1: General Requirements". Switzerland.
- [99] R.G. Dong, J.Z. Wu, and D.E. Welcome. (2005) "Recent Advances in Biodynamics of Human Hand-Arm System." *Industrial Health*, **43**, 449-471.
- [100] A. Israr, P. Merkl, and H. Z. Tan (2004). "A Two DOF Controller for a Multi-finger Tactual Display Using a Loop-Shaping Technique." *2004 ASME Intl Mechanical Engineering Congress and RD&D Expo*. Anaheim, California.
- [101] G. Jansson and L. Monaci. (2006). "Identification of real objects under conditions similar to those in haptic displays: providing spatially distributed information at the contact areas is more important than increasing the number of areas." *Virtual Reality*, Vol 9, pp. 243-9.
- [102] Y. Kohno, S. Walairacht, S. Hasegawa, Y. Koike, M. Sasta. (2001). "Evaluation of Two-handed Multi-Finger Haptic Device." *SPIDAR-8. ICAT*, Tokyo, Japan, Dec 5-7.

- [103] A. Frisoli, M. Bergamasco, S. Wu, and E. Ruffaldi. (2005). "Evaluation of Multipoint Contact Interfaces in Haptic Perception of Shapes." *Multi-point interaction with real and Virtual Objects*. (eds) F Barbagli, D. Prattichizzo, K. Salisbury,. Series: Springer Tracts in Advanced Robotics, Vol. 18. Accessed December 2010.
- [104] D. Kendrick. (2009) "Lowering the Price of Braille: A review of the Seika Braille Display." *AccessWorld*, vol 10(4).
<http://www.afb.org/afbpress/pub.asp?DocID=aw100405>
- [105] G. GESCHEIDER, WRIGHT and R. VERRILLO (2008). "Information-Processing Channels in the Tactile Sensory System: A Psychophysical and Physiological Analysis". *Psychology Press*, USA.
- [106] J. R. P. van Erp (2007). "Guidelines for the Use of Vibro-tactile Displays in Human Computer Interaction." *World Haptics 2007*. Tsukuba, Japan.
- [107] G. A. GESCHEIDER (1997). *Psychophysics, the Fundamentals*. 3rd Ed. Lawrence Erlbaum Assoc. Inc. Mahwah, New Jersey.
- [108] K. MacLean and M. Enriquez (2003). "Perceptual Design of Haptic Icons." *Proc. Of EuroHaptics*. Dublin, Ireland, pp. 351-63.
- [109] S. Cholewiak, K. Kim, H. Tan, and B. Adelstein (2010) "A frequency-domain analysis of haptic gratings." *IEEE Transactions on Haptics*. Vol 3(1), pp. 3-14.
- [110] K. E. Overvliet, K. Mayer, J. Smeets, and E. Brenner (2008). "Haptic search is more efficient when the stimulus can be interpreted as consisting of fewer items." *Acta. Psychological*. Amsterdam, Vol 127(1), pp. 51-6.
- [111] E. C. Lechelt (1988). "Spatial asymmetries in tactile discrimination of line orientation: a comparison of the sighted, visually impaired, and blind." *Perception*. Vol 17(5), pp. 579-85.
- [112] Y. Ikei, K. Wakamatsu, and S. Fukuda (1997). "Vibratory Tactile Display of Image-Based Textures". *IEEE Computer Graphics and Applications*, Vol. 17(6), pp. 53-61.
- [113] J. C. Craig (1974). "Vibrotactile difference thresholds for intensity and the effect of a masking stimulus." *Perception and Psychophysics*, Vol 11, pp. 150-2.
- [114] J. A. Grau and D. G. Kemler-Nelson (1988). "The distinction between integral and separable dimensions: Evidence for the integrality of pitch and loudness." *Journal of Experimental Psychology: General*, 117, pp. 347-370.

7 Appendix

7.1 Device Prototypes

7.1.1 Models




Figure 39: Earlier Solidworks model for the stylus-design in an exploded view. Parts are: lens holder, lens, photointerrupter, cap, spring/switch mechanism, middle casing for circuitry and actuator, and the end portion of casing to hold a battery.




Figure 40: Fingertip/minimal casing design. This model was included for the initial patent filing.




Figure 41: Another fingertip design, this time with the sensor on the forefinger and a circuit on back.




Figure 42: Design sheet for a more detailed case design: the actuator would rest in the circular-cut out on top, the control circuitry and optical sensor on the middle circle (yellow), and there would

be an opening on the bottom so the sensor could read the output from the screen. The handles on the side are for an elastic band, so the device could be worn.




Figure 43: The final design: the biggest difference is the top case piece is thicker, and restricts contact to the piezo to reduce the force placed on it.



Figure 44: The physical device.




Figure 45: Three devices on a single hand being used on an image of a car.

7.1.2 Device Circuits




Figure 46: The first control circuit: the left portion describes the photointerrupter. 1, 2, & 3 are a 5V DC power source. The circuit contains a buffer, an inverting amplifier, then a comparator, tied to the motor. The motor would draw too much current from the source, and fry the chip, so output from the comparator was then feed to a darlington-pair transistor.

| Component # | Value (and type) |
|-------------|----------------------|
| 1,2, & 3 | Voltage Source (+5V) |
| 4 | Resistor: 150Ω |

| | |
|----|-------------------|
| 5 | Resistor: 22Ω |
| 6 | Resistor: 100Ω |
| 7 | Capacitor: 10μF |
| 8 | Resistor: 1kΩ |
| 9 | Resistor: 1.5kΩ |
| 10 | Resistor: 1kΩ |
| 11 | Variable Resistor |
| 12 | Resistor: 1kΩ |
| 13 | Resistor: 1.5kΩ |
| 14 | Motor |




Figure 47: The second control circuit got rid of the op-amps all together, and instead used 4 transistors.

| Component # | Type: Value |
|-------------|-------------------|
| 1 | 6V Li Battery |
| 2 | 5V Volt. Reg |
| 3 | Resistor: 100Ω |
| 4 | Resistor: 150Ω |
| 5 | Resistor: 22Ω |
| 6 | Diode 4148, n = 4 |

| | |
|--------|-----------------|
| 7 | Resistor: 1.5kΩ |
| 8 | Resistor: 100kΩ |
| 9, 10 | Resistor: 10kΩ |
| 11, 12 | Resistor: 1kΩ |
| M | Motor |

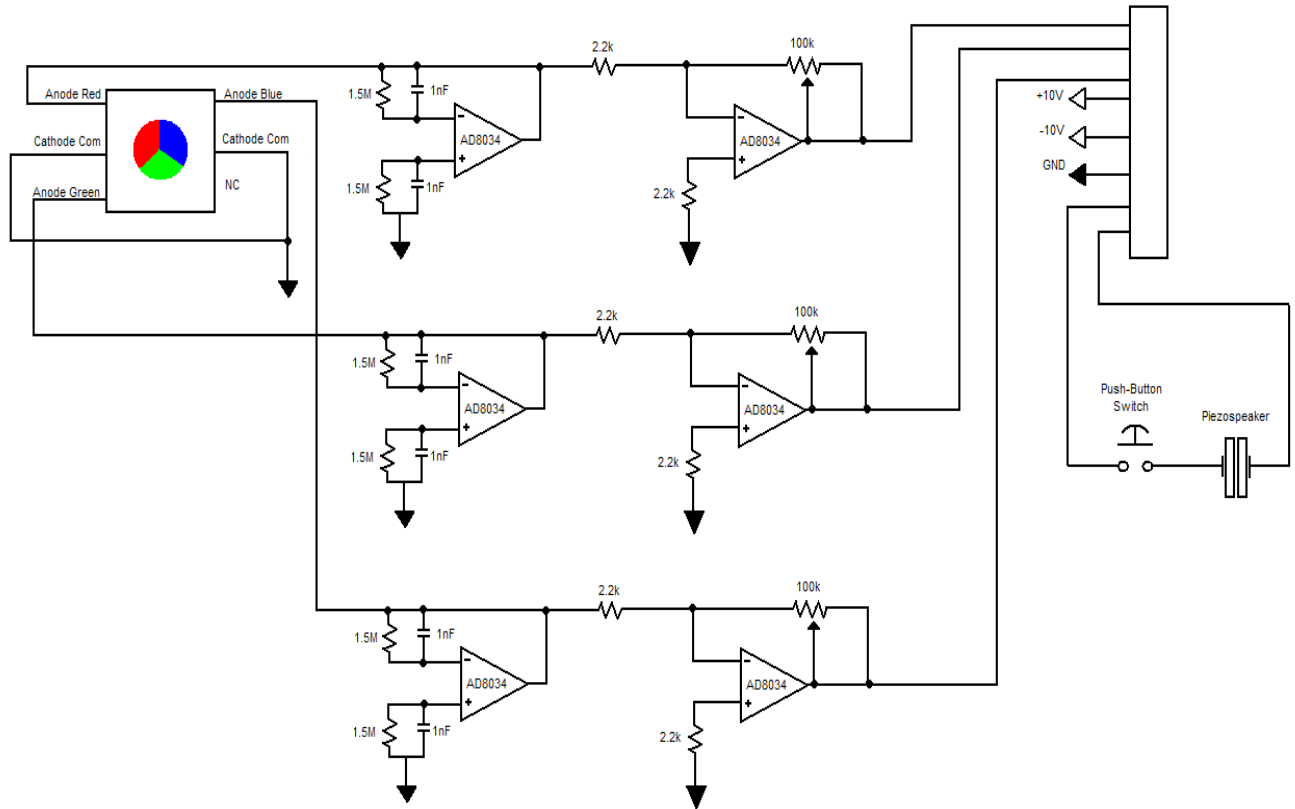



Figure 48: The pre-sampling circuit for the device.



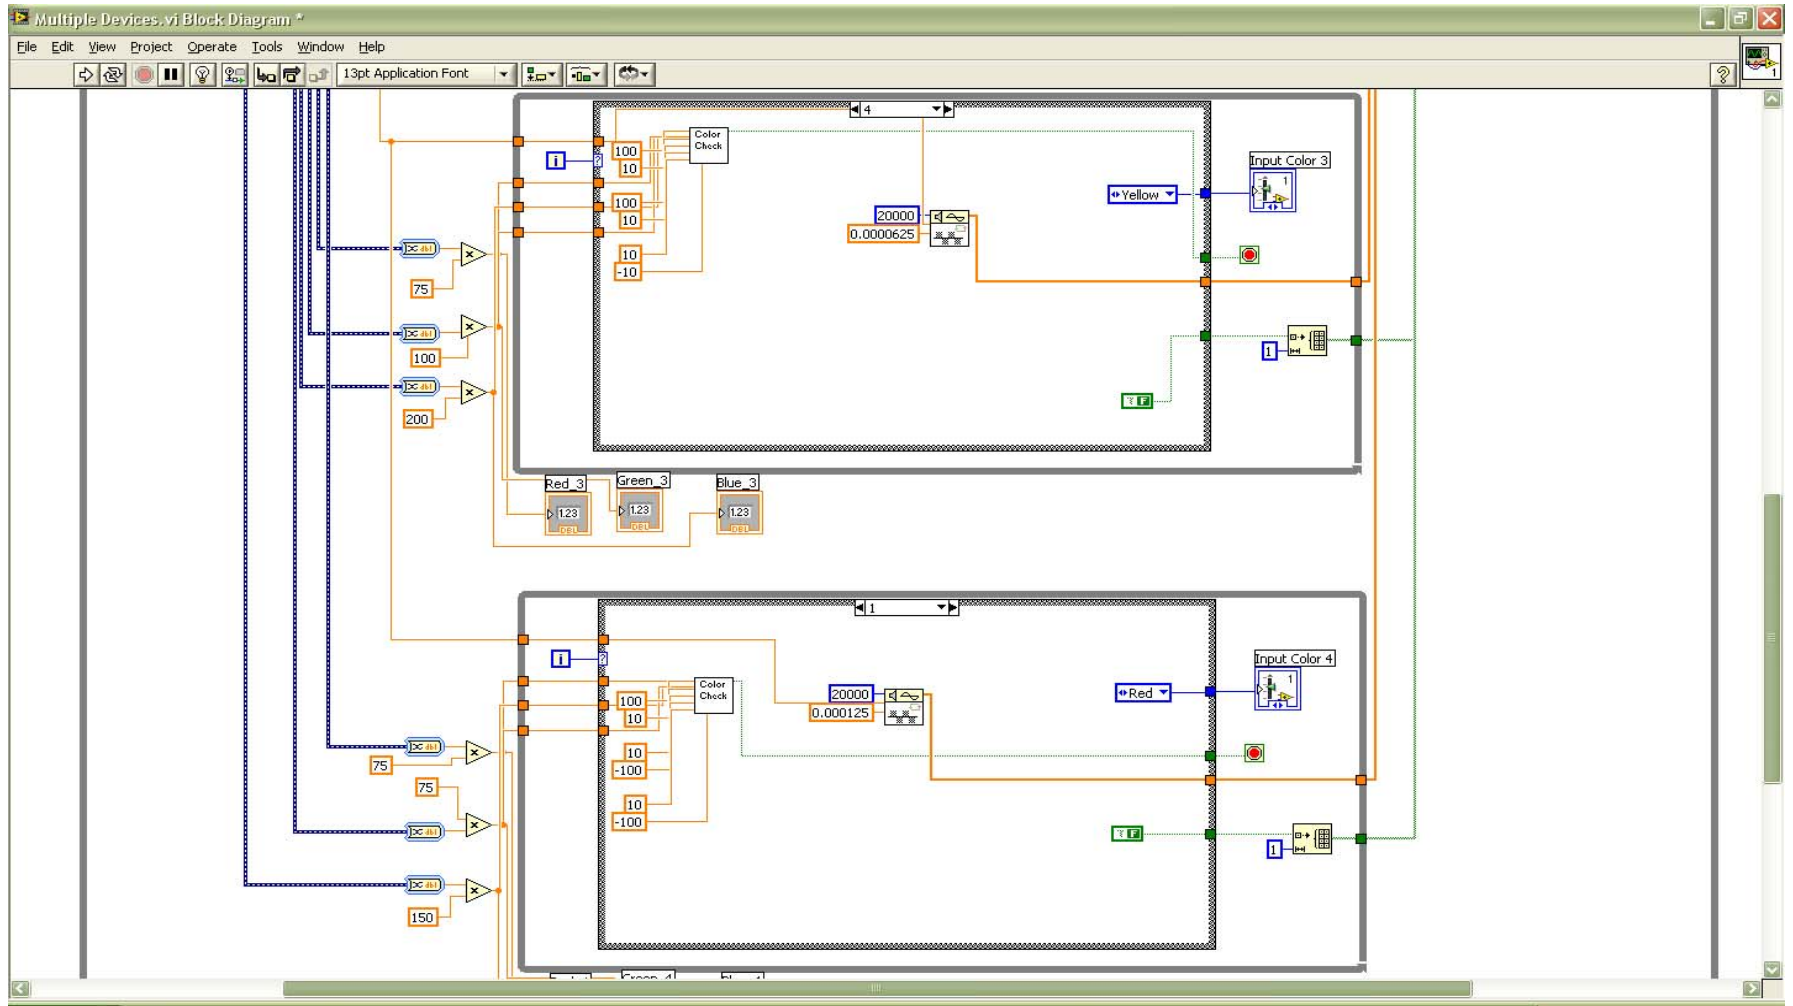



Figure 49: The Labview Control program spread over two pages; note a few of the numeric outputs for the channel colors are cut off.



Piezo Correction Filter




Figure 50: The first perceptual corrections filter design. This design sought to preserve more of the higher frequency bandwidth by filtering out the resonance peak. After the above filter was built, it was found that the resonance frequency was dependent on the applied pressure, so a bandstop + lowpass filter design (above) was replaced with just a lowpass design.

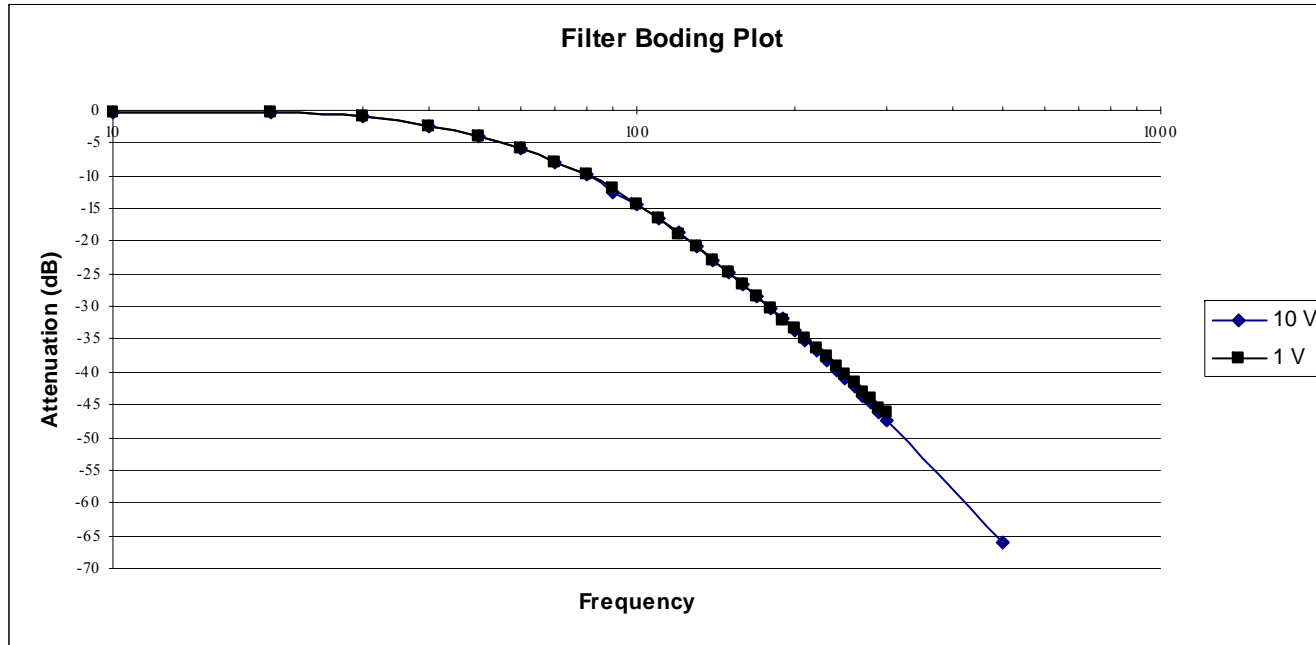



Figure 51: The perceptual filter as it currently is designed. The design is customized using Filter Solutions 2009™; it's based on a 2nd order Bessel lowpass with f_c at 100Hz, and a 4th order Bessel lowpass with f_c at 230Hz.

7.2 Experimental Data

7.2.1 Device Testing

7.2.1.1 Color Input

Here are the sampled inputs read by the Labview program for colors of specific RGB values. This was used as a guideline for creating the color detection subVI within Labview, but I designed the program to still allow for adjustment of these variables.

100 Sample Integration (fs=2000)

| Color Value | | | Input Value | | |
|-------------|-------|------|-------------|---------|----------|
| Red | Green | Blue | Red | Green | Blue |
| 255 | 0 | 0 | 8.16101 | 1.86998 | 0.970072 |
| | | | 8.05582 | 1.87148 | 0.972168 |
| | | | 8.03236 | 1.87681 | 0.974264 |
| 0 | 255 | 0 | 2.31247 | 11.0245 | 2.65641 |
| | | | 2.30021 | 11.0712 | 2.6714 |
| | | | 2.30152 | 11.0499 | 2.66021 |
| 0 | 0 | 255 | 1.51597 | 2.2089 | 9.37105 |
| | | | 1.53632 | 2.23295 | 9.38116 |
| | | | 1.57478 | 2.24533 | 9.40189 |
| 255 | 255 | 255 | 8.96398 | 11.9043 | 10.7234 |
| | | | 8.92193 | 11.8621 | 10.7132 |
| | | | 8.93259 | 11.8715 | 10.6937 |
| 0 | 0 | 0 | 1.43303 | 1.20543 | 0.719956 |
| | | | 1.4209 | 1.19389 | 0.71319 |
| | | | 1.44508 | 1.20819 | 0.719907 |
| 255 | 255 | 0 | 8.66946 | 11.2232 | 2.84582 |
| | | | 8.56739 | 11.152 | 2.81758 |
| | | | 8.54631 | 11.0567 | 2.80322 |
| 0 | 255 | 255 | 2.1886 | 11.5654 | 10.5924 |

| | | | | | |
|-----|-----|-----|---------|---------|----------|
| | | | 2.25732 | 11.5209 | 10.4762 |
| | | | 2.25524 | 11.5985 | 10.5556 |
| 255 | 0 | 255 | 7.90524 | 2.80627 | 9.40394 |
| | | | 7.97929 | 2.8682 | 9.56457 |
| | | | 8.17752 | 2.87385 | 9.6365 |
| 192 | 0 | 0 | 5.04929 | 1.49933 | 0.836802 |
| | | | 5.00509 | 1.5013 | 0.828905 |
| | | | 4.9927 | 1.46274 | 0.817603 |
| 0 | 192 | 0 | 1.7321 | 7.16525 | 1.75582 |
| | | | 1.78278 | 7.12898 | 1.76351 |
| | | | 1.78131 | 7.20656 | 1.77375 |
| 0 | 0 | 192 | 1.38893 | 1.75421 | 6.10299 |
| | | | 1.3997 | 1.75335 | 6.06048 |
| | | | 1.44697 | 1.75524 | 6.07823 |
| 192 | 192 | 192 | 5.62808 | 8.26634 | 7.40469 |
| | | | 5.31355 | 7.80635 | 6.91526 |
| | | | 5.48087 | 8.08326 | 7.10631 |
| 192 | 192 | 0 | 5.65144 | 7.75468 | 1.96172 |
| | | | 5.68228 | 7.74861 | 1.96143 |
| | | | 5.63906 | 7.77884 | 1.96097 |
| 0 | 192 | 192 | 1.93237 | 8.02141 | 7.37109 |
| | | | 1.8868 | 8.0027 | 7.37224 |

| | | | | | |
|-----|-----|-----|---------|---------|----------|
| | | | 1.89887 | 7.97638 | 7.32069 |
| 192 | 0 | 192 | 5.24633 | 2.18635 | 6.45832 |
| | | | 5.20097 | 2.17254 | 6.37756 |
| | | | 5.21047 | 2.17449 | 6.34033 |
| 128 | 0 | 0 | 2.96401 | 1.21527 | 0.692969 |
| | | | 3.16123 | 1.346 | 0.765067 |
| | | | 3.09082 | 1.31098 | 0.744572 |
| 0 | 128 | 0 | 1.42032 | 3.98749 | 1.15137 |
| | | | 1.41725 | 3.96151 | 1.14144 |
| | | | 1.57203 | 3.99959 | 1.20413 |
| 0 | 0 | 128 | 1.35921 | 1.44495 | 3.31237 |
| | | | 1.41393 | 1.45707 | 3.27589 |
| | | | 1.4 | 1.46114 | 3.31422 |

7.2.1.2 Spatial Resolution, Line Test

Testing the spatial resolution; 1pt = 0.3mm. TRUE/FALSE statements result from a 100% accuracy of 10 trials.

All colors compared against a black background

| WHITE | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | TRUE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |

| RED | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | TRUE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| GREEN | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| BLUE | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | TRUE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| Yellow | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|---------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| | | | | | | | |
|-----|-------|-------|-------|------|------|------|------|
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
|-----|-------|-------|-------|------|------|------|------|

| Purple | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|---------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| Aqua | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |

| Dark Red | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-----------------|-------|-------|-------|-------|-------|------|------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

| Dark Grn | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-----------------|-------|-------|-------|-------|-------|------|------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

| Dark Blue | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|------------------|-------|-------|-------|-------|-------|------|------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

| D Yellow | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-----------------|-------|-------|-------|-------|-------|------|------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

| D Purple | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-----------------|-------|-------|-------|-------|-------|------|------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

| D Aqua | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|---------------|-------|-------|-------|-------|-------|------|------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

Line color/background color: these trials represent the lowest saturation cases.

| R/DR | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |

| G/DG | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |

| B/DB | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |

| Y/DY | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |

| P/DP | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |

| A/DA | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|-------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE | TRUE |

Line color/background color: these trials represent the lowest contrast between hues.

| G/A | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | TRUE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |

| R/Y | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|------------|-------|-------|-------|------|------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | TRUE | TRUE | FALSE | FALSE |

| | | | | | | | |
|-------|-------|-------|-------|-------|------|------|------|
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |

| | | | | | | | |
|------------|-------|-------|-------|-------|-------|-------|-------|
| R/P | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| | | | | | | | |
|------------|-------|-------|-------|------|------|-------|-------|
| B/P | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
| 1mm | FALSE | FALSE | TRUE | TRUE | TRUE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| | | | | | | | |
|------------|-------|-------|-------|-------|------|------|------|
| B/A | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
| 1mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |

| | | | | | | | |
|------------|-------|-------|-------|------|------|------|------|
| G/Y | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
| 1mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 1.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |

| | | | | | | | |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| DG/DA | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2mm | FALSE | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE |

| DR/DY | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |

| DR/DP | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |

| DB/DP | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |

| DB/DA | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | FALSE | TRUE | FALSE | FALSE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | FALSE | FALSE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | FALSE | FALSE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | FALSE | FALSE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | FALSE | FALSE | TRUE | TRUE |

| DG/DY | 1pt | 2pt | 3pt | 4pt | 5pt | 6pt | 7pt |
|--------------|-------|-------|-------|-------|-------|-------|-------|
| 1mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 1.5mm | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE | FALSE |
| 2mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 2.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 3.5mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |
| 4mm | FALSE | FALSE | TRUE | TRUE | TRUE | TRUE | TRUE |




Figure 52: Example of a spatial resolution test using vary line thicknesses and vary line separations.

7.2.1.3 Temporal Resolution




Figure 53: Time latency measurement

| Trail | Time (ms) |
|---------|-----------|
| 1 | 3.3 |
| 2 | 2.3 |
| 3 | 3.8 |
| 4 | 4.3 |
| 5 | 4.2 |
| 6 | 4.1 |
| 7 | 2.6 |
| 8 | 3.4 |
| 9 | 3.2 |
| 10 | 2.6 |
| average | 3.38 |
| stdev | 0.71771 |

Time latency recordings (first and second prototypes)

7.2.1.4 HAVS and Accelerometer Data (solenoid motor)

| RMS Acceleration (m/s) | | | |
|------------------------|--------|--------|-------|
| Trial | Axis 1 | Axis 2 | Acc |
| 1 | 1.624 | 1.773 | 2.404 |
| 2 | 1.660 | 1.727 | 2.395 |
| 3 | 1.679 | 1.919 | 2.550 |

| | | | |
|--------------|-------|-------|-------|
| 4 | 1.577 | 1.788 | 2.384 |
| 5 | 1.643 | 1.736 | 2.390 |
| 6 | 1.616 | 1.775 | 2.401 |
| 7 | 1.632 | 1.787 | 2.420 |
| 8 | 1.652 | 1.760 | 2.414 |
| 9 | 1.552 | 1.746 | 2.336 |
| 10 | 1.615 | 1.800 | 2.418 |
| STDEV | 0.038 | 0.054 | 0.054 |
| AVG | 1.625 | 1.781 | 2.411 |


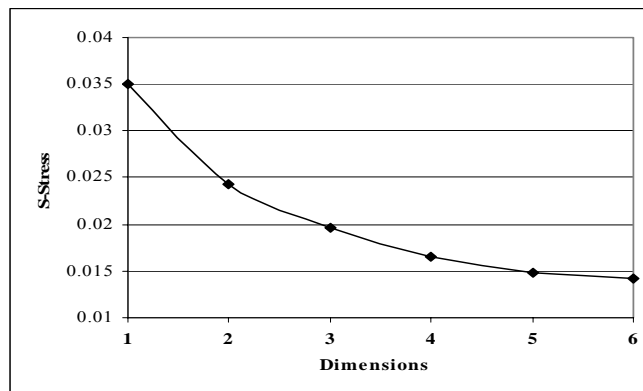


Figure 54: FFT of the solenoid motor feedback.


7.3 Developing Texture Feedback

7.3.1 MDS Experiment

Experiment 1 Scree Plot



Experiment 2 Scree Plot



7.3.2 Spatio-Temporal Patterns

All Sets: 3 Temporal Frequencies




Set 1: 3 Spatial Directions:

Round



Vertical




Horizontal




Set 2: 3 Spatial Directions


Round



Horizontal




Vertical




Set 3: 3 Spatial Directions


Round



Vertical

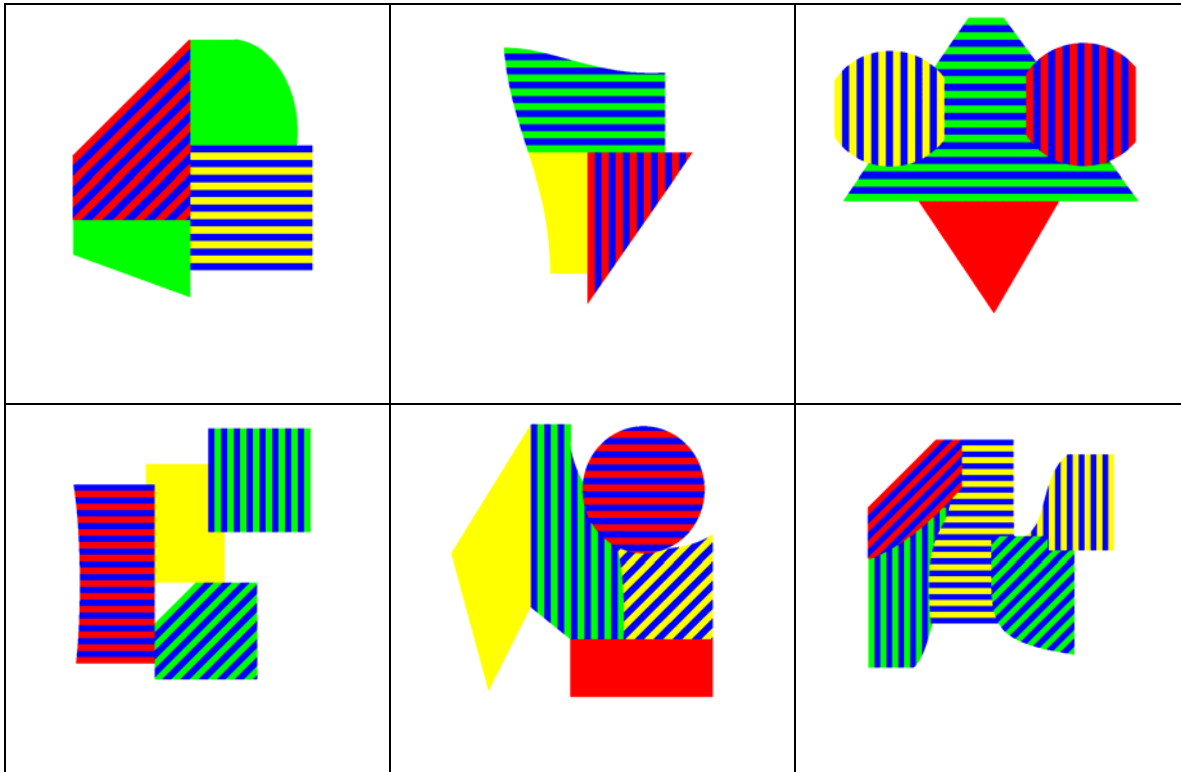



Horizontal



7.3.3 Nonsense Objects

Using Frequency Modulated, Spatially Directed Patterns (Set 3)







7.4 Experiments with Real Objects

7.4.1 Image Sets



7.4.1.1 Thompson's Set [from 65]

Contains: Umbrella, phone, glass, pan, bowl, bread, chair, & lamp.



Outline Representation


















TextyForm Representation





Texture Representation

































7.4.1.2 Complete Image Set










Set Contains: Airplane, apple, axe, bell, boat, bowl, bread, camera, candle, cane, car, chair, dresser (chest), clock, comb, computer, eyeglasses, fire hydrant, glass, hammer, handsaw, ironing board, iron, knife, lamp, lawnmower, mixer, mug, pan, pear, phone, plant, semi-truck, ring, scissors, skillet, pot, spatula, stool, stroller, table, tea kettle, teacup, toilet, trashcan, umbrella, vase, wine bottle, microwave, wrench.








Note: some texyform images were lost due to me saving over the file when I created the texture version.








































[image accidentally saved over]









[image accidentally saved over]













[image accidentally saved over]


















[image accidentally saved over]














[image accidentally saved over]


















[image accidentally saved over]





[image accidentally saved over]

