Designing Audio-Haptic Systems to Render Charts for Blind Individuals

Rayan Isran

Department of Electrical & Computer Engineering
McGill University
Montréal, Québec, Canada

October 2023

A thesis submitted to McGill University in partial fulfillment of the requirements for the degree of Master of Science.

2023/10/04

Abstract

We examine the possibility of leveraging haptic feedback for communication in two areas of haptic perception: passively, through the body, and actively, through externally controlled haptic devices. In the first section of this work, we carry out a perceptual experiment to understand discrimination thresholds of vibrotactile stimuli presented along the forearm. Improving research surrounding tactile acuity at this site of the body could allow designers to deliver information encoded using vibrations more usefully. We report on the minimum detectable differences in vibrational frequency, and findings indicate that subjects exhibit greater sensitivity at locations proximally close to the wrist, suggesting that this location of the forearm should be used to reliably deliver information.

In the second section, which forms a majority of this research, we attempt to address challenges faced by blind and low vision (BLV) users in interpreting graphical content, specifically line charts. Although approaches that combine audio and haptic feedback can improve accessibility, it remains unclear how force-feedback compares to tactile feedback in conveying information to users in a multimodal setting with sound. As refreshable tactile displays (RTDs) are generally more expensive, recent commercially available force-feedback devices such as the Haply 2diy could act as substitutes for BLV users if they are able to interpret content to a similar level of accuracy. Accordingly, we conduct a study to compare sonifications alone against sonifications combined with three distinct haptic conditions: two force-feedback approaches, and a tactile representation. For each condition, we tasked users with describing trends and making inferences about the data, and then asked them about their experience with each approach through a semi-structured interview. Results show enhanced understanding of the data with the inclusion of haptics, with similar user performance among the multimodal conditions and a preference towards the RTD renderings. We analyzed the user interviews to identify four themes that can influence a user's preference for each rendering approach. We then discuss the implications of these themes on future multimodal audio-haptic designs.

Such types of multimodal feedback could help overcome access barriers faced by BLV students in classrooms, where traditional visual instruction necessitates additional intervention from teachers. Existing methods of content delivery, such as tactile diagrams, have shown to result in poor recognition performance of material. These limitations may then cause students to be perceived as unable to participate equally in joint classroom activities, impeding their prospects in scientific careers. Creating effective audio-haptic representations of classroom material could provide BLV students with more immersive and inclusive experiences to minimize this stigma. However, this task requires a technical understanding of audio-haptic rendering, programming expertise, and significant overhead, which is typically beyond the scope of educators' expertise. To this end, we present a prototyping tool that can be used to create a desired tactile image on an RTD and provide interactive audio cues through touch gestures. The end audio-tactile designs can be exported, edited, and imported, allowing for easy distribution of experiences. The tool also provides a plugin system, which allows designers to focus on automatically generating audio-haptic renderings for specific graphical content, and an example plugin that creates interactive experiences for single- and double-line charts is presented. We hypothesize that our system could enable educators to provide a higher quality learning experience for BLV students in educational environments over existing techniques.

Résumé Scientifique

Nous recherchons la possibilité d'exploiter le retour haptique pour effectuer de la communication dans deux domaines de la perception haptique : de manière passive, sur le corps, et de manière active, grâce aux dispositifs haptiques télécommandés. Dans la première partie de nos travaux, nous réalisons une expérience perceptuelle pour comprendre les seuils de discrimination des stimuli vibrotactiles délivrés le long de l'avant-bras. L'avancée des connaissances sur l'acuité tactile de cette partie du corps peut permettre aux concepteurs et conceptrices de transmettre plus efficacement des informations codées par des vibrations. Nous présentons nos resultats sur l'écart minimale perceptible entre les fréquences, et les résultats indiquent que les sujets présentent une plus grande sensibilité proche du poignet, ce qui suggère que cet endroit de l'avant-bras devrait être utilisé pour délivrer des informations de manière fiable. Nous rapportons les différences minimales détectables dans la fréquence vibratoire, et les résultats indiquent que les sujets présentent une plus grande sensibilité aux endroits proximaux proches du poignet, ce qui suggère que cet endroit de l'avant-bras devrait être utilisé pour délivrer des informations de manière fiable.

Dans la deuxième section, qui présente la partie la plus conséquente de nos recherches, nous tentons de relever les défis auxquels sont confrontés les personnes aveugles et malvoyantes pour accéder à l'information contenue dans des graphiques, et plus spécifiquement les graphiques linéaires. Bien que les approches combinant le feedback audio et haptique puissent améliorer l'accessibilité, la comparaison entre le retour de force et le feedback tactile pour la transmission d'informations aux utilisateurs dans un environnement multimodal avec du son n'est pas encore claire. Les écrans tactiles rafraîchissables (RTD) étant généralement plus coûteux, les dispositifs à retour d'effort récemment commercialisés, tels que le Haply 2DIY, pourraient remplacer les utilisateurs de BLV s'ils sont capables d'interpréter le contenu avec un niveau de précision similaire. Par conséquent, nous recherchons la différence entre l'utilisation des sonifications seules et l'utilisation des sonifications avec trois conditions haptiques, dont deux utilisent le retour d'effort et une utilise une représentation tactile. Pour chacune de ces conditions, nous avons demandé aux participants de décrire les tendances et de faire une inférence sur les données présentées. Ensuite, nous les avons interrogés sur leur expérience lors d'une interview semi-structurée. Nos résultats ont démontré une meilleure compréhension des données avec l'inclusion des stimuli haptique avec une préférence pour les rendus sur tablette

Braille dynamique. Nous avons identifié quatre thèmes qui pourraient influencer la préférence dans l'approche de rendu grâce à notre analyse des interviews. Nous discutons des implications de ces thèmes pour la conception des interactions multimodales audio-haptiques.

Ce type de retour d'information multimodal pourrait aider à surmonter les obstacles à l'accès auxquels sont confrontés les élèves malvoyants dans les salles de classe, où l'enseignement visuel traditionnel nécessite une intervention supplémentaire de la part des enseignants. Les méthodes actuelles pour transmettre du contenu, telles que les diagrammes tactiles, ont montré qu'elles ne permettaient pas une bonne reconnaissance du matériel. Ces limitations peuvent alors amener les élèves à être perçus comme incapables de participer sur un pied d'égalité aux activités communes de la classe, ce qui limite leurs perspectives de carrière scientifique. La création de représentations audio-haptiques efficaces du matériel scolaire pourrait offrir aux élèves aveugles et malvoyants des expériences plus immersives et inclusives afin de minimiser cette stigmatisation. Cependant, cette tâche nécessite une bonne connaissance technique des rendus audio-haptiques, une expertise en programmation et des frais importants généraux, ce qui dépasse généralement le champ d'expertise des éducateurs. À cette fin, nous présentons un outil de prototypage qui peut être utilisé pour créer une image tactile souhaitée sur une tablette Braille dynamique et fournir des signaux audio interactifs par le biais de gestes tactiles. Les conceptions audio-tactiles finales peuvent être exportées, éditées et importées, ce qui facilite la diffusion des expériences. L'outil propose également un système de plugins, qui permet aux concepteurs et conceptrices de se concentrer sur la génération automatique de rendus audio-haptiques pour un contenu graphique spécifique. Nous présentons un exemple de plugin qui crée des expériences interactives pour des graphiques à une ou deux lignes. L'outil pourrait permettre aux éducateurs de fournir une expérience d'apprentissage de meilleure qualité aux étudiants aveugles et malvoyants.

Acknowledgements

Jeremy Cooperstock gave me considerable freedom to pursue ideas for research, and provided constructive criticism throughout the stages of planning experiments and writing this thesis. Indeed, his feedback on my writing has been invaluable, and will continue to benefit me for years to come. Yongjae Yoo helped steer the initial direction of much of my research work, and was always quick to respond to my technical questions. Sri Gannavarapu and Linnea Kirby were my first friends at McGill, and working with them on a Human-Computer Interaction course project almost exclusively through Zoom formed the first of many great memories here. Jeff Blum was a fun and insightful colleague, and it was always a pleasure to watch and listen to his incredible presentations. Cyan Kuo brought a lot of energy to the lab's atmosphere, handled a lot of my paperwork, and never hesitated to respond to my questions at short notice despite her busy schedule. Everyone at the Shared Reality Lab made the past three years wonderful, and the friendships of Yongjae Yoo, Jeff Blum, Cyan Kuo, Émmanuel Wilson, Rohan Akut, Corentin Conan, Juliette Regimbal, Jano Fu, and Venissa Carol are especially memorable. Jaydeep Singh was a kind friend, and I'm glad to have met him in-person at the Shared Reality Lab's 25th anniversary. Thanks to Antoine Phan for the campus tours, as short-lived as they were.

I would like to express gratitude towards my external examiner, Vincent Lévesque, for taking the time to review my thesis.

Thanks to Justin-St Pierre for warmly welcoming me to Montréal and lending me some of his speedrunning equipment. I sit here reading my thesis from the very same monitor he bought for me when I arrived in the city.

A special thanks to Corentin Conan for translating the Abstract of this document to French, and Anonymous for proofreading parts of the thesis and offering suggestions to improve the writing.

I would like to express gratitude towards my amazing parents, my sister, and my brother, for their unconditional love and support throughout the course of my studies.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Overview

The sense of touch, or haptics, has recently become an active and diverse research topic largely driven by a growing interest in the development and use of consumer electronics with touch screens [1]. These types of devices can transmit signals in two ways: passively, in which sensations are delivered to the user without any conscious input, e.g., through vibrations on body-worn devices, or actively, in which sensations arise from actions by the user, e.g., through use of gaming controllers in simulated virtual reality environments. By numbers, the largest application area of haptics is basic communication [1]—while the most prominent example of this is a phone vibration, haptics has shown to be effective in demanding environments to offload the audio and visual channels; for example, haptic wearables allow users to interpret information without recurring visual inspection in critical care units [2], military zones [3], or social settings [4]. Additionally, haptic devices provide researchers with opportunities to discover novel ways of delivering information to improve accessibility for blind and low vision (BLV) individuals.

Effective use of the haptic channel for such applications, however, requires a deep understanding of both the capacity of the human touch system, as well as allocation of attentional processing across sensory channels. In this thesis, we expand upon the existing body of research on these topics and assess the potential of utilizing the haptic channel for communication through both passive and active haptic feedback. Chapter 3 describes a perceptual experiment conducted on the forearm, to understand how it can be more effectively used to present information passively to a user. In Chapters 4 and 5, which form the main focus of this work, we explore whether and how various types of haptic feedback can help BLV users in understanding information in line charts. We conclude by presenting a tool that can be used to prototype interactive audio-haptic renderings with minimal technical knowledge, which could be used in classrooms to replace traditional methods of presenting information to BLV students that are costly, time-consuming, and generally do not provide an adequate level of detail to these students. The tool can also be used to automatically generate audio-haptic experiences for specific types of graphics, thus providing designers with a quick way to distribute these experiences. This portion of work is part of a broader ongoing effort to develop and deploy accessible user experiences for graphical content online through a web extension [5].

We establish some brief background on the two types of haptic feedback in the subsequent section, and discuss the motivation for the thesis topics further in Sections 1.3 and 1.4.

1.2 Types of Haptics Feedback

In human-computer interaction (HCI), haptic feedback is an umbrella term that refers to sensations stimulated through both *tactile* and *force-feedback* devices. Each of these terms define perception solely on the location of sensory receptors: tactile perception relates to stimuli detected by low threshold receptors under the skin within the contact area, and allows users to feel textures, temperatures, and vibrations. *Vibrotactile* stimuli specifically refer to sensations delivered by vibrations. Force-feedback devices are typically used to impart forces to a user, but some, like the one we used for a study in this thesis, may also allow for manipulation by the user as input to the computer. Although the tactile and kinesthetic channels work together to provide humans with the means to perceive and act on their environment, haptic devices are primarily designed to only make use of one of these channels. For example, Braille is used for tactile interpretation through exploration with the fingers using pin array devices, whereas force-feedback could be used to obtain information about the shapes or textures of objects in a virtual environment. Because of these distinct types of interactions, an understanding of the strengths and weaknesses of rendering strategies designed with haptic interfaces is necessary if a designer's goal is to efficiently deliver information through touch.

1.3 Haptic Perception

As in audition, the vibrotactile sense is capable of discriminating between frequency, amplitude, and other aspects of a signal. Therefore, such signals can be used to communicate information to the skin. However, the location of stimulation plays an important role for this purpose. Despite having the greatest tactile acuity, the fingers might not be a suitable location due to their constant interaction with the external environment. Several previous studies have assessed perception on the hands, back, and foot, but limited research exists on a potentially more unobtrusive location—-the forearm, limiting designers from effectively using this site. An enhanced understanding of touch characteristics in this region would expand the range of possibilities for communicating information in fields such as navigation, gaming, virtual reality, or monitoring systems. In this work, we conduct and evaluate an experiment to determine the suitability of the forearm as a location for perception. We measure the minimal change in frequency that can reliably be distinguished by a user, known as the *just-noticeable difference* (JND). Furthermore, because practical applications often involve the use of multiple actuators spaced apart at a site, we determine the effect of distance between actuators along the forearm on the frequency JND. An improved understanding of these property could allow researchers to make proper use of the available bandwidth of the haptic communication channel in this area.

1.4 Accessibility

Another research area in haptic perception is the creation and distribution of audio-haptic representations of visual information. There is a good chance that you navigated to this document using either a mouse or a touchpad without being consciously aware of your interactions with the device. For sighted individuals, navigating to and reading digital content is an intuitive task that requires minimal effort. For many of the estimated two billion people worldwide who are visually impaired [6], accessing digital content remains a struggle that has only exacerbated with the COVID-19 pandemic [7]. This figure is also expected to double by 2030 [8]. Additionally, research shows that individuals in the United States fear losing vision more than the memory, hearing, or speech [9].

Combined with the influx of data and widespread use of technology in society, the skill to learn and understand graphical content will become increasingly necessary for those who work in STEM fields where visualizations are extensively present. For blind individuals, screen readers are the most popular assistive technology devices for accessing this information, but they often require sifting through irrelevant information, and provide little useful information when alternative-text tags that are intended to provide a text summary of a graphic are frequently missing [10, 11]. Even for sighted individuals, the increase in use of deceptive visualization practices, whether knowingly or unknowingly, can often lead to misinterpretation of data [12]. This has motivated research to explore alternative forms of presenting graphics, such as through sound or haptics.

1.4.1 Haptics for Line Charts

Despite its relatively limited bandwidth, the haptic channel has recently become an active research topic for communicating data from visualizations. Technological advances have allowed for richer expressivity through low-cost tactile and force-feedback devices; for example, tactile drawings can be used to represent bar charts, or force-feedback simulations can enable a user to physically navigate a line graph. However, despite these attempts to convey information, studies reveal that haptic mappings are often simple and lacking, and do not necessarily serve the needs of a user [13]. This emphasizes the need for effective designs that allow users to interpret the desired data with ease.

Many studies show that presenting information through a combination of audio and haptic modalities prove better than either modality alone [14, 15, 16]. But little prior work has focused on how these modes can be effectively combined to deliver content, especially with respect to the *type* of haptic feedback. For example, users may perceive information from an audio-tactile rendering differently to an audio-force feedback rendering. Secondly, humans have a remarkable ability to detect asynchrony in multimodal interactions, which can impact the user's understanding and perception of the rendering. A lack of crossmodal synchrony can lead to difficulty in interpreting the information, which often results in BLV individuals becoming wary of embracing new technologies [8]. Consequently, understanding the benefits and drawbacks of each haptic mode is essential for creating meaningful renderings. To this end, we conducted a user study with BLV participants to assess the effectiveness of audio-haptic renderings for line charts using both a refreshable tactile display (RTD) and a force-feedback device. We evaluate each multimodal approach to a baseline audio-only rendering, and subsequently compare the performance and preference between each of the multimodal approaches.

Additionally, creating multimodal renderings can be an arduous task. It requires a technical understanding of haptic rendering and interfaces, and programming expertise. This cannot be expected in settings where the technology can be employed, e.g., for educators in classrooms. Therefore, a simple way to quickly create these renderings should exist. To achieve this, we present a prototyping tool that can be used to create a desired tactile image with interactive audio experiences on a pin array device. We also show how this tool can be used to rapidly create consistent renderings for targeted types of graphics, such as line charts.

1.5 Thesis Overview

This chapter described the motivation behind pursuing areas of haptic perception that involve both passive and active communication. We discussed the importance of understanding perceptual properties on the body, and providing accessible methods to convey information for BLV individuals.

Chapter 2 provides a summary of the existing literature on these topics and is divided into two portions: in the first, a general summary of vibrotactile perception on the human body is presented, with a specific survey of the arm as a spot for communicating information. In the second, a literature review on haptic exploration for the BLV is presented, with specific scholarship on charts.

Chapter 3 explores the effect of distance on perceptual acuity of the glaborous (non-hairy) area across the arm through a psychophysical experiment. We found statistically insignificant but suggestive insights on the effects of distance and the encoded frequency of vibrations on a user's ability to discriminate between stimuli.

Chapter 4 discusses a study that sought to understand whether multimodal haptic techniques aid unimodal spatialized audio in more effectively presenting information in line charts. We found that each of the multimodal audio-haptic cases outperformed unimodal audio and similarly to the multimodal force-feedback case in terms of user performance, but outperformed unimodal audio and mulimodal force-feedback subjectively via user preference. We conclude by presenting an overview of the observed benefits and drawbacks to each modality and their implications.

Chapter 5 builds on insights gained from the findings of the study in Chapter 4. We create and demonstrate the usefulness of an interactive prototyping tool to rapidly create ideated multimodal experiences for the RTD, with an emphasis on charts.

Chapter 6 discusses the implications of the findings across this body of research, concluding with directions for future work.

1.6 Summary of Contributions

Our primary contributions are described in Chapter 4 and Chapter 5. To the best of the author's knowledge, the study described in Chapter 4 is the first of its kind that investigates tradeoffs between multimodal tactile and multimodal force-feedback rendering strategies. We describe blind users' experience with renderings that use each of these strategies, and compare them to an audio-only approach to understand how the inclusion of haptics can impact a user's ability to perceive information in line charts. Through qualitative analysis of recorded transcripts from interviews, we identify several factors that can impact usability, such as synchrony between the modalities, user autonomy, and initial impressions of the haptic devices used in the study. The results of our study can inform future audio-haptic designs, which could provide BLV individuals with improved experiences to navigate graphical content. Our experiment was also performed using readily available or expected-to-be-released devices on the market. As such, the results of our study are of practical relevance to would-be-adopters of the technology.

In Chapter 5, we demonstrate an authoring tool for the RTD used in a study described in Chapter 4. This allows a designer to create and edit any desired tactile image on the device. We built upon the hardware to incorporate touch interactivity, which can provide supplementary information cues through audio during real-time finger exploration. The tool can be used by educators to quickly create audio-haptic renderings, allowing for flexible use in learning environments. Although there are already tactile authoring displays on the market, e.g., Thinkable's TactileView[1] or ViewPlus' IVEO,[2] TactileView does not provide audio interactivity, and IVEO requires high running costs due to the repeated use of embossing paper. The Tangram

[1] https://www.tactileview.com/

[2] https://viewplus.com/product/iveo-3-hands-on-learning-system/

Workstation is a collaborative tool in which sighted individuals can compose and edit tactile graphics for blind individuals on an RTD [17]. However, this system also does not provide a multimodal interactive experience. We believe our tool would be simpler for designers to use, and employs a robust device that is expected to be on the commercial market soon.

Chapter 2

Background

2.1 Preface

This chapter describes existing work on the topics discussed earlier. We divide this chapter into two parts: vibrotactile perception, expanding upon Section 1.3; and haptics rendering for chart information with an emphasis on line charts, expanding upon Section 1.4.

2.2 Vibrotactile Perception

The human skin is covered with mechanoreceptors that sense changes in temperature, pressure, texture, and vibration [18]. When stimulated by direct contact with external stimuli, these mechanoreceptors communicate information to the central nervous system (CNS), allowing individuals to interpret and respond to the environment. They are broadly of four types—(a) Merkel's disks, (b) Meissner's corpuscles, (c) Ruffini endings, and (d) Pacinian corpuscles—and are classified by depth from the skin surface, and how quickly they adapt to sensations. Pacinian and Meissner's corpuscles (PCs and MCs) rapidly adapt to tactile stimuli and only send signals to the central nervous system (CNS) when a change is detected, making them fast-adapting (FA) receptors. In contrast, Merkel disks and Ruffini endings (MDs and REs) continuously fire signals to the CNS and do not quickly adapt to tactile stimuli, making them slow-adapting (SA) receptors. The MCs and MDs are closer to the surface, whereas the PCs and REs are a few millimeters deep in the skin. These properties are summarized in Table 2.1.

Table 2.1: Receptor characteristics for different types of mechanoreceptors.

Receptor	**Receptive Field (mm)**	**Frequency Range (Hz)**	**Sensed Quantity**
FA-I Meissner corpuscles	3–4	5–60	Skin taps, grasp
SA-I Merkel disks	3–4	0–5	Form, textures, pressure
FA-II Pacini corpuscles	>20	40–400	Rapid vibrations
SA-II Ruffini endings	>10	0–10	Directional skin stretch, movement of joints

The PCs rapidly adapt to skin displacement, making them specialized for the detection of high-frequency vibrations (40–400 Hz, with a peak sensitivity of 200–300 Hz [19] throughout the body). This relatively wide bandwidth of detectable stimulation frequencies has been increasingly leveraged in the design of body-worn tactile devices to convey information [20, 21, 22, 23] or augment or replace sensory input from other modalities [24]. However, the amount of information that can be reliably conveyed to the user depends on the individual's ability to discriminate between stimuli of varying intensity.

Past studies have explored vibrotactile perception on the hands and fingertips as stimuli [25, 26, 27, 28, 29] because of their high mechanoreceptor density. However, the constant interaction of the fingertips in the external environment may make them unsuitable for placement of tactile devices.

Other tested locations include the lower back [30], the head [31], the torso [32, 33, 34], the upper leg [35], and the arm [34, 36]. Despite the forearm also being a commonly used and suggested site for the design of body-worn tactile devices [37, 38, 39, 40, 41, 42, 43, 44], fewer studies have examined the efficacy of stimuli delivered at this location.

Mahns et al. reported a marked increase in the just-noticeable frequency difference on the hairy skin of the forearm compared to the glabrous skin of the fingertip at vibrotactile frequencies of 200 Hz [45]. Similarly, Shah et al. investigated the two-point threshold discrimination both sequentially and simultaneously delivered across various dermatomes, or areas of the skin innervated by a single spinal nerve, on the arm and reported that the frequency discrimination ability is greater on the glabrous side of the forearm closer to the elbow than the center of the dorsal side [46].

In 1870, Vierordt stated via his *Law of Outward Mobility* that tactile acuity improves with increased mobility of a body part, with a change in slope of the function at joints such as the wrist or elbow [47]. This suggests that the two-point discrimination threshold decreases as one moves from the shoulder blade to the fingertips. This trend was later graphed by Boring (1942) [48] and has been replicated by recent research using more easily measurable methods [49]. Cholewiak et al. investigated the localization of stimuli at two varying frequencies ranging over the length of the volar forearm, determining that acuity is approximately lowest at the middle of the forearm and increases proximally, both towards the wrist, and the upper arm [50]. While this finding would appear to be in contrast to the results from the graph depicted by Boring, it does not violate Vierordt's principle as both the elbow and wrists are anatomical landmarks, i.e., each has increased mobility relative to nearby skin regions. This has also been observed in results by Morioka et al. who found that vibrotactile thresholds were lowest at the volar forearm among sites tested between this region, the fingertips, toes, and heel [28]. However, beyond the effects of glabrous vs. non-glabrous skin, no significant variation in vibrotactile perception has been noted over different regions of the arm. Our work aims to expand upon this research by determining discriminability threshold of this region by measuring the JND on the non-hairy side of the arm.

2.3 Accessibility Approaches for Charts

Tactile diagrams have been recommended as an effective way to convey graphical descriptions for BLV users [51]. These representations consist of printed content on Braille embossers or swell paper with raised material that can be felt by a user's fingers. These raised portions correspond to the different elements of the diagram, such as lines, curves, or points. Because tactile diagrams are quick and easy to produce, they are currently used in educational settings [52]. However, despite their effectiveness over verbal descriptions of tables [53], they have several limitations: the raised material wears out after frequent use, their size is often small owing to the high material cost and typical need for multiple copies of a diagram, and they must be reprinted if any changes to the diagram are required. Studies have shown that blindfolded, sighted participants have poor performance in recognizing common objects through tactile diagrams, largely due to an inaccurate perception of the global shape of these objects [54, 55].

Performance of congenitally blind participants on these representations was similar [55] or worse [56]. Therefore, alternative methods have been explored to convey visualizations.

2.3.1 Sonifications

In this rendering approach, data are assigned sounds to deliver information. An early study involving both sighted and blind participants found that sonification using pitch differences to indicate changes in a line chart yielded similar performance to tactile diagrams [57]. Further research over the past few decades has shown their benefit for BLV users in interpreting visualizations of one and two data series [58, 59], which has resulted in guidelines to present sound for tables and line graphs [60]. A more recent study comparing sonifications with tactile diagrams shows that sonifications are useful in providing general impressions of a line plot, whereas tactile diagrams allow users to explore points of interest at their own pace [61]. Because multimodal approaches have shown to result in improved experiences [62], efforts have been made to integrate these sounds to tactile diagrams [63, 64]. In this technique, a tactile diagram is placed on top of an audio tablet that stores information about the diagram. Audio representing the information can then be delivered when users touch an object of interest on the diagram. However, these tools, which are often unfamiliar to participants, also suffer from the same limitations as tactile-only diagrams. One study showed that users preferred a tactile only rendering to a tactile and audio approach to understand data from bar charts [53]. However, the authors note that this result may have been due to user unfamiliarity with the multimodal approach used, and unreliability in detecting touch gestures that provided audio feedback. To mitigate this problem, participants were asked to constrain their hand movements during exploration. Therefore, the results of their experiment could be attributed to the specific implementation of the participants approach and may not necessarily reflect multimodal principles.

2.3.2 Force-feedback Methods

Some approaches have used force-feedback haptics to deliver information. These are typically operated using a handle and used in two ways: *exploration*, in which the user can freely move the device handle to feel textures or effects in a virtual canvas, and *guidance*, in which the device guides the user's hand through a region on the canvas. These alternative interaction techniques can result in different ways information is perceived [65, 66]. Wai and Brewster conducted a study to assess the usability of exploration to present information in line and bar charts, using a high-fidelity SensAble PHANToM device [15]. Findings showed that a mix of blind and sighted participants were able to answer task questions correctly with high accuracy. However, due to the prohibitive cost of the device, they also compared its performance to a more affordable force-feedback mouse. While user accuracy with the mouse was lower, a multimodal approach that incorporated sounds to map chart values to pitch minimized the difference in user performance between the two devices. The study concluded that a multimodal approach, using both sonified audio and force-feedback exploration, was more effective in conveying bar chart trends than both tactile diagrams, and either modality alone. They also observed a similar pattern for line charts [16]. However, these approaches only focused on using an exploration approach. Limited research exists on understanding the comparative user performance between exploration and guidance using graphical content as the stimuli. Roberts et al. describe the potential strengths, drawbacks, and applications for each mode, but to the author's knowledge, no formal study has been conducted to evaluate these hypotheses [65]. Furthermore, it is unclear what the impact of audio is on user performance or preference between these force-feedback approaches.

2.3.3 Refreshable Pin Methods

While force-feedback devices are limited to a single point of contact, RTDs have been developed to allow for two-handed exploration techniques. Unlike Braille, which is arranged as a single line of characters, the grids in these devices are larger and can therefore support a higher resolution to display more complex shapes. One of the first such devices was the Optacon, which used a handheld scanner that could capture an image of printed material and convert it to a tactile representation [67]. Experiments conducted by the authors showed that it was useful in improving reading rates, and could enable engineers to interpret graphs and schematic diagrams.

Developed in the early 2000s, The HyperBraille[1] is a commercially available refreshable Braille display with one of the largest touch-sensitive pin-matrix displays on the market (7200 pins arranged in 120 columns and 60 rows). More recently developed in the early 2020s, the

[1] `http://www.hyperbraille.de/`

Graphiti[2] (Orbit Research, USA) and Dot Pad[3] (Dot Inc., South Korea) devices each contain a tactile grid of 60 (horizontal) × 40 (vertical) refreshable pins spaced comparably with Braille. TouchPad Pro Assistive Technology has developed a TouchPad Pro[4] that includes a color display and can raise pins to various heights.

Several applications have been developed using these types of devices, such as tools to improve accessibility for maps [68, 69, 70, 71], methods to explore geometric shapes with a zoom feature [72], tools for students with blindness to recall spatial images [73], and tactile animations [74]. However, few studies have focused on using them to convey visual content in scientific domains. Of these, Melfi et al. found that an audio-tactile reader outperformed tactile diagrams to convey line chart data from a PDF file [75]. Based on this work, Schwarz et al. developed a command line interface on the HyperBraille that displays quadratic graphs [76]. They conducted a study with five blind participants that showed high user accuracy in identifying shapes and regions. Participants indicated low workload, but cited some issues such as the low-resolution of the display, and the high learning curve to use a device with which they were unfamiliar. Albert describes a system using the Job Access With Speech (JAWS) screen reader software with a handheld refreshable pin array device that could be useful to understand mathematical content [77]; however, a small-sized tactile grid would be unsuitable for presenting graphical content that often contain a high volume of information, such as charts.

Despite a wealth of research contrasting multimodal audio-haptic renderings with unimodal approaches [16, 15, 53, 62, 78], to the best of the author's knowledge, no study has empirically investigated the tradeoffs between tactile and force-feedback techniques for representing graph data, nor compared between guided and user-controlled force-feedback rendering techniques. Additionally, it remains unclear how each of these approaches, when used in conjunction with audio, would compare among each other and a baseline unimodal audio approach in terms of user performance and preference. We consider such open-ended questions as important issues in the design of effective multimodal representations and accordingly, aim to answer some of them through the work in Chapter 4.

[2] http://www.orbitresearch.com/product/graphiti-plus/
[3] https://pad.dotincorp.com/
[4] https://touchpadpro.org/the-touchpad-pro/

Chapter 3

Evaluation of the Forearm for Haptic Perception

Preface

In this work, we investigate stimuli discrimination thresholds on the non-hairy side of the forearm. We calculate the just-noticeable change in vibrational frequency through a psychophysical experiment conducted with ten participants. Furthermore, the relatively long proximal length of the forearm provides designers with the opportunity to use multiple vibrotactile actuators to convey information. Thus, we also characterize the effect of distance on the JND. While results are statistically insignificant, they indicate a trend towards improved discriminability when actuators are placed closest to each other, suggesting the usage of closer distances between actuators in the design of wearable devices at this location if designers would like to support accurate frequency discrimination.

Author's Contribution

This work is the result of collaboration. Rayan Isran designed the experiment with guidance from Dr. Yongjae Yoo, implemented the hardware and software interfaces, recruited participants for a user study, and wrote this chapter. Dr. Yongjae Yoo verified the code that was written by the author for Section 3.2, and performed some of the statistical analysis. Prof. Cooperstock provided critical feedback and edited some of the contents.

3.1 Introduction

Touch-based technologies are prevalent in various aspects of our societies—from smartphones to gaming consoles to educational aids, amongst others. The advent and continued development of these technologies are of particular importance in the field of wearables, with a growing interest in using the forearm as a site for haptic interaction. To enable the design of better body-worn tactile devices, a good understanding of both the information transfer capacity and the somatosensory characteristics of vibrotactile stimuli is required. While many previous studies investigate vibrotactile perception on the hands, there exists limited research on perceptive characteristics of the forearm, preventing designers from effectively using the haptic channel for communication at this site. In this regard, we conducted a study that investigates frequency perception on the volar (non-hairy) forearm. Additionally, the relatively longer length of the forearm compared to many other sites on the body provides the opportunity to incorporate multiple actuators in designs. Thus, we also studied the effects of distance on the frequency discrimination threshold along the forearm.

Responding to the open questions identified in Section 2.2, our work described in this chapter investigates the effect of distance between two actuators on frequency discrimination ability. Specifically:

- **RQ1:** What is the minimum change in frequency that causes a perceptible change in sensation on the volar forearm?
- **RQ2:** How does the threshold from RQ1 change with varying proximal distance between the actuators (i.e., the interstimulus distance)?

An experiment was designed to test these hypotheses. Based on the approximate trend of tactile acuity across the arm, we expected that frequency discrimination will improve as the two-point threshold is smaller and closer to the wrist. The experiment was approved by the McGill Research Ethics Board (file #20-10-041) and participants were compensated CAD 10/hour for their time.

3.2 Methods

Ten participants (5 male, 5 female; $22-39$ years in age; mean = 26.9; SD = 5.4 years) with no known deficits of tactile perception were recruited from the McGill University community. Nine reported prior experience with participating in studies; three of them had experience with haptics, two specifically with psychophysics studies.

Participants were seated with their dominant arm outstretched over the apron of a wooden table with their palms facing up. Vibrotactile stimuli were delivered through two 5 mm linear resonant actuators (LRA) (Lofelt, L5), attached to the glabrous side of the participant's wrist using 3M double-sided VHB tape. To localize vibrations, foam pads surrounded the actuators and were strapped to the arm using Velcro such that the device was tightly coupled to the arm yet comfortable to wear, as shown in Fig. 3.1a. The "reference" actuator was affixed 3 cm on the volar forearm proximally from the wrist, and the second actuator was kept as the controlled distance (5 cm, 10 cm, or 15 cm) away from the reference actuator, as shown in Fig. 3.1b. The actuators were driven by an amplifier board connected to a laptop running a Python script that controlled the frequency and amplitude of the stimuli, as well as the presentation order for vibrations. Participants wore earplugs and noise-cancelling headphones playing pink noise, primarily to mask the sound from the vibrations throughout the experiment. A short test procedure was included to ensure that participants could perceive vibrations at both ends of the detectable range on the forearm.

3.2.1 Procedure

All stimuli were delivered as sinusoidal waves with an 800 ms duration, as any longer could result in information decay [79]. An initial familiarization stage was included in which participants could choose to deliver vibrotactile stimuli separately or in sequence (with an interstimulus interval of 800 ms), with varying frequencies and amplitudes selected through a slider. The range of the frequency was limited to those sufficiently perceivable on the forearm (40–400 Hz). Participants were told to take as much time as they needed to become comfortable with the vibrations, but typically only spent 5 minutes during the first session.

(a)

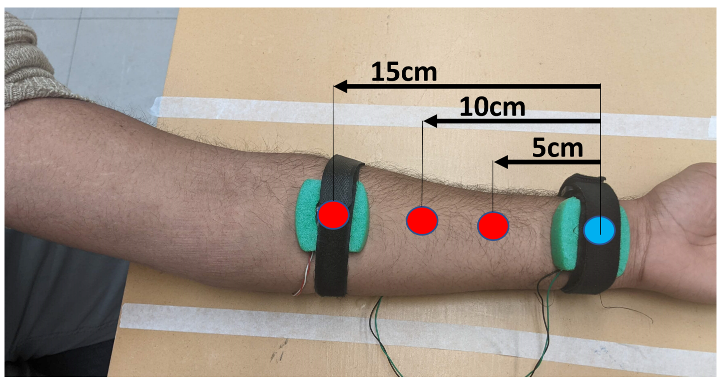

(b)

Fig. 3.1: Experiment setup showing: a) the delivery of vibrotactile stimuli, user input through the computer, and the use of noise-suppression headphones. b) the three inter-actuator distances used with the actuator closer to the wrist providing a reference point.

3.2.2 Calibration

To minimize the effect of frequency on perceived amplitude, a calibration procedure was performed to adjust amplitudes such that the vibrations were perceived as approximately equal in intensity across the entire range of frequencies used. This consisted of delivering two vibrations to the same actuator—a fixed reference frequency at 170 Hz, and a variable comparison frequency. The comparison frequencies ranged from 40 to 200 Hz in 10 Hz intervals, with two additional 5 Hz intervals between 150 and 170 Hz to include higher resolution testing near the reference frequency of 170 Hz, for a total of 24 pairs. For each comparison frequency, participants were told to adjust an amplitude slider until its intensity felt equal to that of the reference frequency.

Following the initial calibration for frequency response, an inter-actuator calibration was conducted, with stimuli covering the same set of frequencies used during the prior round. This was done to ensure that vibrations were perceived as equal strength when delivered by each of the two actuators, at equal frequencies, with an interstimulus interval of 800 ms. The intensity of the reference frequency was determined by the results of the initial calibration step. Participants adjusted an amplitude slider until the intensity of the reference actuator, presented first, felt equal to that of the "comparison" actuator. This process was repeated for each pair, with frequency values chosen pseudorandomly.

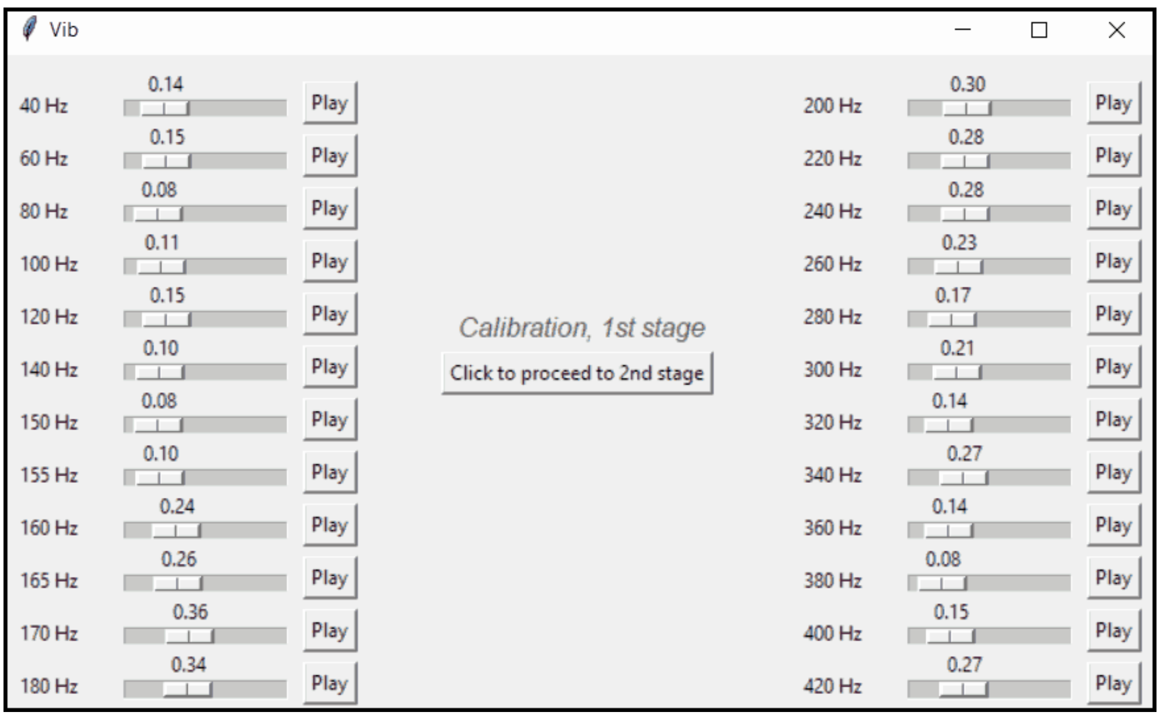

Fig. 3.2: A screenshot of the software's calibration interface.

3.2.3 Experiment

The main experiment was carried out using a 1-up 1-down adaptive staircase procedure, an established method for finding perceptual thresholds [80]. In this procedure, the intensity of the stimulus is adjusted based on the participant's responses to previous trials. The stimulus intensity is increased or decreased depending on the participant's response, with the aim of finding the intensity level at which the participant can just detect the stimulus. The "up-down" staircase procedure involves starting with a stimulus intensity level that is easily detectable and then decreasing the intensity after a correct response and increasing it after an incorrect response. This continues until the threshold is found, as explained below.

One vibration from each pair was the fixed reference frequency and the other was a comparison frequency, starting either from the lower or upper bound. The latter was computed by multiplying (upper bound) or dividing (lower bound) the reference frequency by 1.6. An additional multiplier between 0.8 and 1.2 was applied to randomize the comparison frequency value in the first pair of vibrations presented to each participant. The presentation of the two vibrations was randomly ordered. If participants indicated that the vibrations felt the same, the subsequent comparison frequency was set further away from the reference frequency by a factor of 20%. If the participant indicated that the comparison frequency felt different, the comparison frequency was moved closer to the reference frequency by a factor of 20%. During each reversal (i.e., when the participant's current response differed from the previous one), this multiplicative factor was reduced by 8%. Progressively reducing the step size in this manner is a common technique to efficiently converge to a reliable perceptual threshold in an adaptive staircase, and although there is no general rule on the size of the steps, these values are typically selected such that there are fewer than five responses between reversals [80]. Accordingly, pilot staircases were conducted on the coauthors of the experiment to determine the step size values used for the study. To account for frequencies outside of the calibration range, a linear interpolation method was used to obtain the intensity for the values using the closest lower and upper calibration values. Additional staircases were conducted with a reference frequency of 80 Hz to observe whether lower frequencies would impact stimuli discrimination performance.

A staircase ran until either 30 trials or 11 reversals were recorded. These values were chosen to produce a reliable threshold value in adaptive staircase methods [80, 81]. The same procedure was then repeated with another random pair from the set of reference frequencies and starting bounds. This was done until all pairs from the set were completed for the controlled distance between the actuators, i.e., a total of four staircases. To minimize fatigue, participants were instructed to take a two minute break between staircases.

For all participants, the sessions were spread across three consecutive days, taking approximately 1 hour for the first session and 45 minutes for each of the two subsequent sessions. Each day of the experiment involved completion of all four staircases at a particular inter-actuator distance, with these distances chosen in counter-balanced order across participants.

Participants were provided with an explanation of the purpose of the study after completing the experiment and filling in a post-experiment questionnaire.

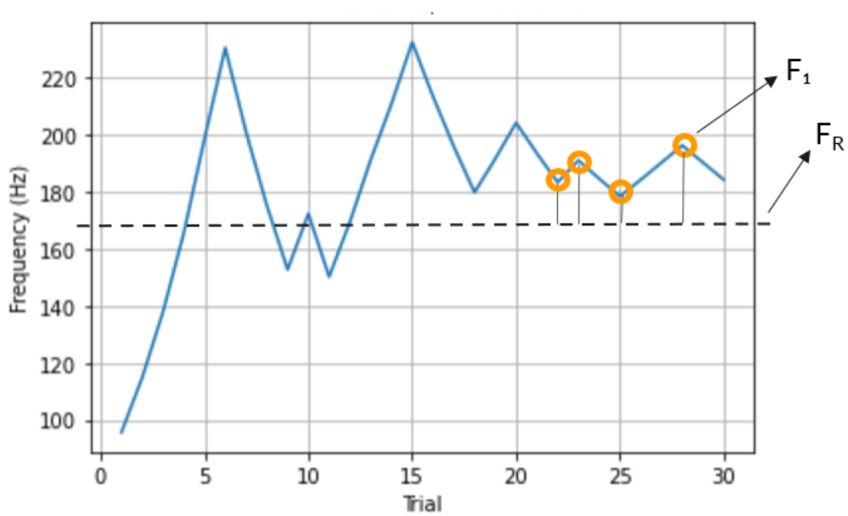

Fig. 3.3: An example of a staircase, showing the last four reversal points. The reference frequency is depicted as dotted line, with a value of 170 Hz.

3.2.4 Data Analysis

The difference between the reference frequency and each of the last four reversal points was calculated to determine the JND for each staircase, as shown in Fig. 3.3 and Equation 3.1. The average JNDs for each distance are shown in Fig. 3.4. A three-way repeated measures ANOVA with the JND as the dependent variable was conducted to ensure that there were no interaction effects between the conditions (distance, reference frequency, and direction) or order of presentation.

$$JND = \frac{\frac{1}{N}\sum_{i=1}^{N}(F_i - F_R)}{F_R} \times 100 \tag{3.1}$$

3.3 Results

JNDs were reported for each reference frequency and comparison frequency set, for each of the distance parameters recorded, for a total of 12 readings per participant. Effects of the reference frequency and starting point of the staircase (near the lower or higher end of the frequency range) were significant ($F(1,8) = 11.79$ and 9.022, $p = 0.009$ and 0.017, and $\eta_p^2 = 0.211$ and 0.273 respectively), while that of distance was not ($F(2,16) = 1.056$, $p = 0.371$, $\eta_p^2 = 0.0584$). The JNDs for the reference frequency = 80 Hz case was 33.1% (SD = 13.0), whereas it was 27.1% (SD = 11.4) for the 170 Hz case (p < 0.002). JNDs from upward and downward staircase directions were 32.8%, and 27.5% (SD = 12.2, 12.6%) respectively. Although the differences in JNDs associated with lower frequency were only slightly improved over that of higher frequencies, these results nevertheless differ from prior research suggesting the opposite for the fingertips [27]. We hypothesize that this may be due to stronger tactile masking effects [82, 83] and a higher density of mechanoreceptors on the fingertips than on the forearm [84].

In contrast, ANOVA test results revealed that the effects of distance were not statistically significant, despite six of the participants indicating that they found stimuli easier to distinguish when the inter-actuator distance was its maximum, i.e., 15 cm, versus only three who indicated the opposite. The effect size ($\eta_p^2 = 0.0584$) suggests that distance could have a considerable effect, but would need to be verified with a larger study. These observations could be related to a confounding factor of the mechanoreceptor density across the arm, affecting the ability to discriminate between stimuli at similarly encoded frequencies on the volar forearm. This would be consistent with research indicating that localization ability is greater closer to the wrist than the proximal center of the forearm [50].

However, the inclusion of a higher number of participants or reduced step size intervals between the selected distances is required to make a concrete comparison, given that a large number of deviations between individual results, particularly in the 10 cm case, potentially stem from the center of the volar forearm being comparatively less sensitive a location in terms of vibrotactile perception [84].

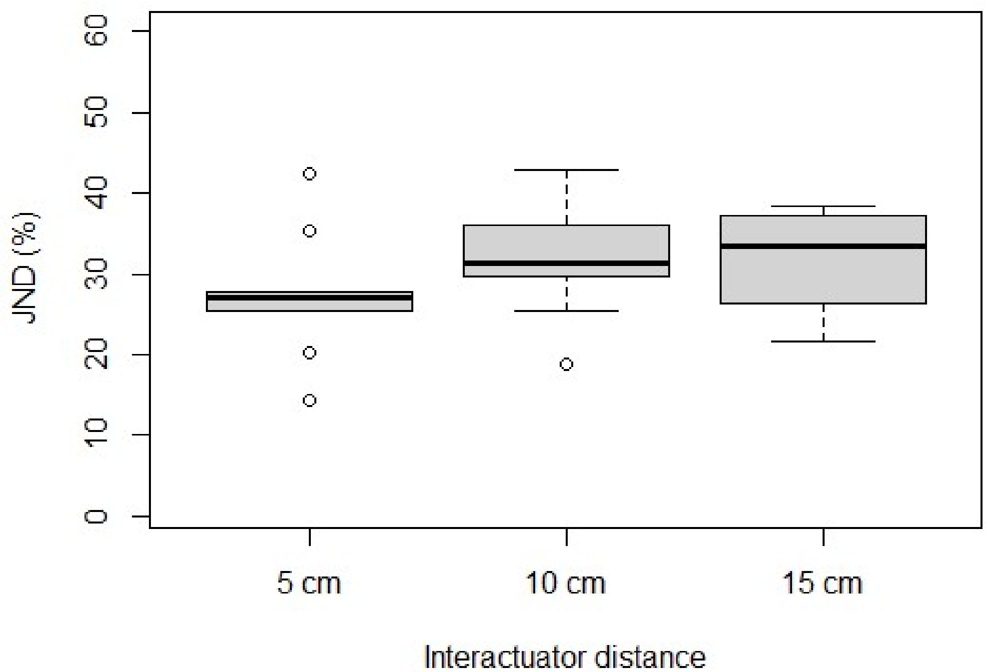

Fig. 3.4: Average JNDs for each distance condition. Outliers are denoted by a circle.

3.4 Conclusion

We examined the influence of distance on haptic sensation perception with actuators placed on the volar forearm. Although most of the trends observed did not yield statistical significance, they provide clues about the influence of distance on haptic perception that are consistent with current literature: notably, frequency discrimination appears best at smaller interstimulus distance for the sites tested.

Despite our use of a calibration process to minimize the non-linear frequency response of the L5 actuators, minor intensity differences between stimuli of similar frequencies were no doubt still present. Nonetheless, the hardware used in our studies is typical of that found in end-user products in virtual reality and wearables, so we believe the results are relevant to current consumer technologies. We note as well that evidence points to equal-subjective intensity curves being similar in shape to the threshold vs. frequency function [85], but it is not feasible for tactile designers to include an equal subjective magnitude procedure for each human operator.

Stimuli discrimination studies conducted on the fingertips indicate that JND is independent of amplitude [25], suggesting that calibration procedures are not necessary. However, it is unclear whether this is true for other locations on the body such as the forearm. It would also be interesting to observe whether tactile acuity varies with different haptic devices, for example, a custom linear tactile array device as used by Cholewiak and Collins [50]. The effects of interstimulus timing and age have also been shown to affect perception of vibrotactile stimulation [86, 50]. Lastly, participants have shown to adapt to vibrations over time, improving their stimuli discrimination ability [87, 88]. Further research is required to examine all of these issues while investigating the research questions described earlier in this chapter.

Chapter 4

Evaluation of Multimodal Line Graphs for Blind Individuals

Preface

This chapter discusses the motivation, design, and results of a user study that sought to understand whether and how multimodal audio-haptic approaches can help BLV users interpret line graphs more effectively than non-speech audio only experiences. Eight elderly participants took part in the study, six of whom were partially or legally blind. We characterized their performance in identifying trends, and asked them open-ended questions about their experience with the audio and multimodal audio-haptic modes presented. We then qualitatively analyzed the results to describe the benefits and drawbacks with each approach, and the implications for future designs.

Author's Contribution

Dr. Yongjae Yoo wrote code to reverse-engineer the RTD used in this study. Florian Grond created the sonified audio renderings. Cyan Kuo wrote the ethics application, recruited participants, and took notes during many user studies. Corentin Conan conducted the study for francophone participants under the guidance of Sri Gannavarapu and myself, who jointly designed the study and conducted the experiment for anglophone participants. Rayan Isran wrote the code to generate the haptic renderings described in Section 4.2.3. The thematic analysis procedure described in Section 4.4.3 was the result of collaboration between Sri Gannavarapu, Corentin Conan, and Rayan Isran. This chapter was mostly written by Rayan Isran, with the exception of Section 4.5, a majority of which was written by Sri Gannavarapu.

4.1 Introduction

The ability to interpret data through visualizations is a fundamental aspect of learning and communication in scientific fields [89, 90, 13]. Such visualizations include charts, which are one of the most common forms of recorded data on the web [89]. However, access to such information remains a struggle for BLV individuals, for whom current methods of content delivery are limited to textual representations of these graphs, e.g., synthesized speech or Braille [13, 91], or reduced visual representations, e.g., tactile diagrams [13, 92], neither of which adequately capture the dense and complex nature of graphics [92, 93, 54].

Another approach to convey data trends in graphs is to use non-speech audio, in which the visual environment is mapped to an auditory soundscape. Parameters such as the x- and y-axis values in a line chart can be mapped to perceivable audio parameters such as spatial panning, and pitch, respectively [94, 95]. BLV users tend to perform better than their sighted counterparts in discriminating frequencies and localizing sound [96, 97], motivating the use of sonifications to communicate spatial information for accessibility [98, 99]. To further augment understanding of the data, haptic effects, whether delivered through force feedback [15, 16, 100], or tactile devices, such as refreshable tactile displays (RTDs) [75, 101, 102, 103], can be delivered together with audio to provide more detail or reinforce the information conveyed. Notably, RTDs can provide a tactile experience that is considered the gold standard of rendering 2D information for BLV users in the classroom. However, there has been limited research evaluating such displays in conjunction with sonifications. Additionally, the high cost of RTDs presents an access barrier for blind communities, limiting their widespread use. Force-feedback devices have also been explored in education settings. However, as Baud-Bovy and Balzarotti note in their extensive 2017 review, while such devices were at one point affordable for the gaming market due to high demand, justifying the investment in mass production, past studies investigating applications for BLV education typically relied on devices that are expensive or no longer available [104]. To the best of the authors' knowledge, there has been little to no research since then that explored the use of force-feedback as a sensory substitution approach in BLV education.

Recently, more affordable force-feedback devices such as the Haply 2diy [105] have become available, which represent a potential alternative to RTDs for multimodal presentation of graphical content. However, little is known with regards to the tradeoffs between force-feedback and tactile rendering approaches, whether individually or in combination with sonifications, aside from the fact that the degree to which these sources of information integrate can affect the quality of interaction in multimodal interfaces [78, 106, 107]. Through a qualitative study, our work seeks to address this shortcoming of knowledge by answering the following research questions:

- **RQ1:** *Can the use of multimodality in rendering approaches improve the user experience for BLV users compared to an audio-only (sonified) experience?*
- **RQ2:** *What are the differences in user perception of line chart renderings that employ only audio or both modalities together via a force-feedback device or RTD?*

4.2 User Study

To evaluate our rendering strategies to represent graphical content, we conducted a study using line charts and maps as the stimuli. For reasons of space, this chapter only discusses charts.

4.2.1 Participants

Eight BLV users were recruited through one of three channels: the Regroupement des aveugles et amblyopes du Montréal métropolitain (RAAMM), the Canadian Council of the Blind, and word of mouth. Since participants recruited from the RAAMM included a mix of francophone and anglophone users, we conducted experiments in English and French. The McGill Research Ethics Board approved this experiment under file #21-10-031. Participants were compensated CAD 10/hour.

4.2.2 Apparatus

The devices used for our study are shown in Figure 4.1. Force-feedback was delivered using the Haply 2diy (Haply Robotics, Montreal) [105], a 2-degree-of-freedom grounded force-feedback device, connected to a laptop via USB connection. The device can be operated in two distinct ways. In the exploration approach, users can move the position of an end-effector over a workplate. This input controls the position of a graphical replica displayed on screen by a computer program, which can interact with simulated objects in a virtual environment to deliver force sensations to the user. These sensations include effects such as texture, damping, or elasticity. In the guidance approach, the end-effector can be programmed to automatically move to set locations on the workplate. Thus, when the user grasps the handle, their hand can be physically guided along a predefined path.

Tactile renderings were delivered using a Dot Pad connected to a laptop via Bluetooth for communication. On this device, individually addressable pins can be raised on the tactile grid to collectively represent a tactile image that the user can explore with their fingers. The Dot Pad is silent once the pins are raised, and the force-feedback device is quiet; none of the participants indicated that device sounds interfered with their rendering interpretation.

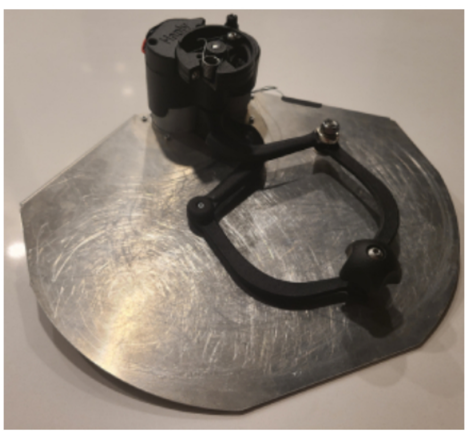

(a) The Haply 2diy device used for force-feedback, showing the end-effector extended towards the bottom of the plate.

(b) The Dot Pad device (image taken from `https://web.archive.org/web/20230410040328/https://pad.dotincorp.com/`).

Fig. 4.1: Images showing the haptic devices used in this experiment.

4.2.3 Stimuli & Modelling

We extracted data from five line charts from various websites that contained graphs rendered using the *HighCharts*[1] graphing framework. Two of these were COVID-19 trends, one represented the currency exchange rate between the US Dollar and Euro, and the remaining two

[1] `https://www.highcharts.com/`

graphs were stock chart trends of two separate companies. Graphs were chosen such that each line chart represented a distinct type of trend. Fig. 4.2 depicts two graphs that were used. Fig. 4.2a shows a stock trend with multiple local minima and maxima, whereas Fig. 4.2b shows three minor peaks and a large peak. Two graphs were presented for force-feedback active guidance and Dot Pad active exploration each, and one was presented for the force-feedback active exploration case. The rendering strategy for each condition is described in the next four subsections.

(a) A graph showing the trend of a company's stock value.

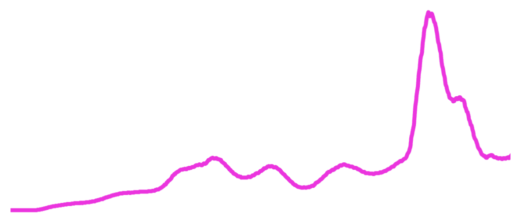

(b) A graph showing the trend of confirmed worldwide COVID-19 cases from 2020 to 2022.

Fig. 4.2: Examples of line charts used in the experiment.

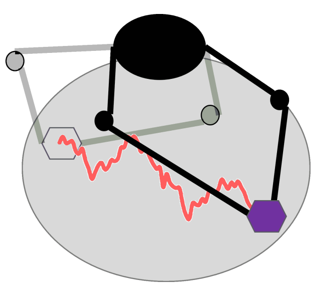

(a) A mockup of the 2diy device showing the path the end-effector would trace for a graph in the guided case.

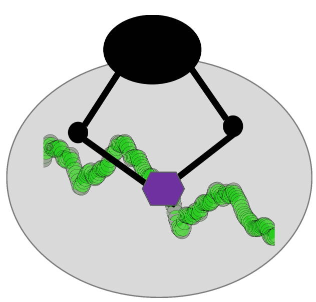

(b) A view of the 'tunnel' simulation for the same device in the exploration case.

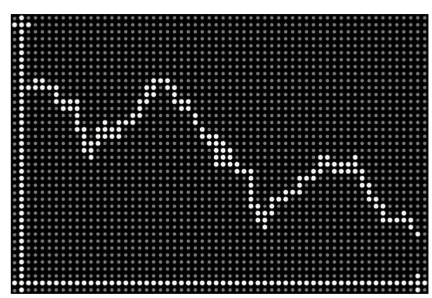

(c) A mockup of a graph as it would be shown on the Dot Pad.

Fig. 4.3: Mockups of the haptic rendering techniques used in the study.

4.2.3.1 Sonifications

Non-speech audio was synthesized and spatialized in SuperCollider[2] using higher-order ambisonics and a binaural decoder. To adopt a mapping choice considered natural for auditory line graphs, the y-axis was mapped to pitch, i.e., the higher the y-value of the data point, the higher the central frequency of the short sound representing a data point. The mapping range was expressed in Musical Instrument Digital Interface (MIDI) notes, from the range A♯4/B♭4 (70) to A♯8/B♭8 (118). Sounds representing data points were played sequentially from the left ear to right ear. For example, the first data point was played on the left ear, the last data point was played on the right ear, and intermediary values were played in between the ears. The reader can access sample renderings online.[3]

4.2.3.2 2diy Guidance

In this approach, the user is guided along a trend by the end-effector. The lines were each discretized into an array of coordinates, and then upsampled to address differences in the number of points among graphs. To simulate a continuous tracing effect, the end-effector guided the user's hand along these points, from left to right, at a fixed rate. The graphs were upsampled to provide approximately 700 points, through which the participant's hand was guided at a rate of 70 points per second, thus, covering the entire graph in 10 seconds. The force required to move to the next point was computed instantaneously in a 1 kHz simulation loop using a spring-damper model [108], and parameters that defined the force-feedback experience were fine-tuned empirically through pilot tests to minimize jitter in the guided movement.

4.2.3.3 2diy Exploration

In this approach, the user can move the end-effector, subject to elastic forces that help constrain the user's trajectory to stay on the line that represents the graph. We enabled active user control in this condition by using the hAPI Fisica physics library[4] in Processing to create a virtual environment that dynamically responds to users' input movements. Similar to the guidance condition, the line was discretized and then upsampled as a set of equidistant objects closely

[2] https://supercollider.github.io/
[3] https://rayanisran.github.io/multimodal/
[4] https://github.com/haphub/hAPI_Fisica

spaced together, such that they formed a virtual "tunnel" through which the graphical replica of the end effector could pass, as shown in Fig. 4.3b. If the replica touched the boundary of any object, the user would feel an elastic force pushing the end-effector back inside the tunnel, constraining their movement. To obtain information about the trajectory of the line, the user can move the end-effector along the tunnel.

4.2.3.4 Dot Pad Exploration

To fit the graph within the constrained resolution of the tactile grid, a moving average filter was applied to reduce the number of data points to a maximum of 60. An x- and y-axis were rendered to provide reference lines, as shown in Fig. 4.3c.

4.3 Methodology

4.3.1 Procedure

Participants were seated at a desk, provided with AKG K240 MKII over-ear headphones, and interacted with either the 2diy or Dot Pad for the associated haptic conditions. Headphones were worn during portions of the experiment when audio was rendered. The experimenter sat on the opposite side of the desk, presenting renderings through a laptop, and swapped to the second device once the participant had completed all tasks with the first device. Participants were encouraged to think out loud, with verbal expressions captured through a laptop microphone for later use in a thematic analysis. To provide context for later discussion, a camera captured their hand movements while participants were interacting with each device. An image of the setup is shown in Fig. 4.4.

The experimental procedure is summarized by the following pseudocode and explained in further detail below:

(a) All materials used in the experiment, including devices used for renderings.

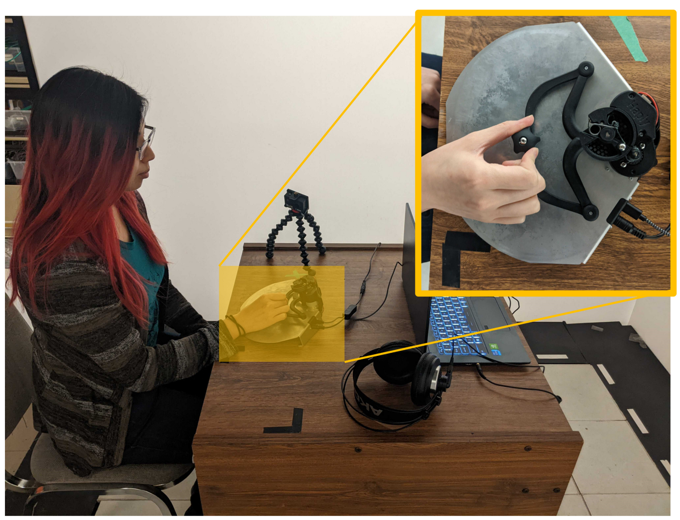

(b) A stand-in participant using the 2diy. Their hand grip is emphasized to illustrate the single-handed interaction typically used to operate the device, compared to the two-handed exploration technique for the Dot Pad.

Fig. 4.4: The experiment setup.

```
multimodal_rendering_ conditions ← {Dot Pad, 2diy with guidance, 2diy with exploration}
pre-test questionnaire
for condition ∈ multimodal_ rendering_conditions, pseudorandomly selected do
    training phase with the associated haptic device
    for each graphic used in condition do
        presentation of the sonified rendering of the graphic, followed by task questions
        presentation of the multimodal rendering of the graphic, followed by task questions
    end for
    semi-structured interview
end for
post-test questionnaire
```

Each participant was introduced to the experiment and asked them a series of pretest questions to collect data on their demographics and familiarity with the technology used in the experiment. Additional questions were asked about their experience with charts to understand current user behaviours and challenges.

The experimenters explained how each of the haptic devices worked, and how the participant could interact with it. At the start of each of the four experimental conditions, described in Section 4.2.3, a training phase was conducted to familiarize participants with examples similar to those included in the actual trials. Participants were then presented with renderings in each condition, starting with the sonified version of each graph, followed by the three multimodal conditions, presented in pseudo-random order. In each multimodal condition, audio was presented before haptic feedback.

For each condition and each graph presented, we asked the following two task questions, for consistency, always starting with the first:

1. Can you describe what the chart is doing? In what general direction are the values trending? *Rationale: evaluate participant's performance in detecting a trend*.

2. Can you identify the relative difference between the values of the starting and ending points of the graph? *Rationale: evaluate participant's ability to discern further details*.

For each condition, we conducted semi-structured interviews to capture qualitative feedback regarding user experience and rendering-specific metrics such as the detection of trends.

For example, if a participant exhibited any confusion with the trend being conveyed in the rendering, we focused on this source of confusion and tried to elicit a description of what they were experiencing. In instances where participants provided responses with insufficient detail, the interviewer proactively sought additional information by asking for elaboration or clarification to ensure a comprehensive understanding of their responses.

Participants were asked the same set of questions, in addition to whether what they felt was similar to what they heard. They were welcome to replay a rendering as many times as they liked to answer a question. Once they answered all questions, the next example was played for that condition, starting again with the audio-only case first. We presented the next block of conditions once every question was answered for all of the renderings presented for a condition. Once participants went through all examples with a particular device, they were asked which condition they found the quickest to obtain an overall "feel" for the general trend, and which they found provided the most "accurate" representation.

Following the experimental trials, participants were also asked whether they would prefer to interact regularly with audio or multimodally.

The entire experiment was conducted in a single two hour session.

4.3.2 Data Analysis

We conducted a thematic analysis to gain a deeper understanding of user interactions, using a reflexive approach informed by guidelines proposed by Braun et al. [109]. All three experimenters were involved in the analysis to mitigate interpretation bias. Each member independently reviewed the data and generated an initial set of codes using Taguette (an open-source qualitative analysis software) [110] to associate tags to the transcripts. The codes were discussed as a group, and then refined to ensure accuracy, comprehensiveness, and fidelity to the data. Once we had our final set of codes, we organized them into categories by grouping similar codes together and providing each group with a descriptive label. We reviewed the categories and identified overarching themes that emerged from the data, which provided a higher-level understanding of the patterns observed across users.

As a secondary analysis, we evaluated a participant's performance in each example by assigning binary scores to the task question. To understand their ability to infer trends, we asked them to provide a brief description of what they heard or felt in the renderings. A score of one was assigned if they could state the correct graph trend. This could be a brief description, such as “up, down, and then up again” for the graph in Fig. 4.2a. A score of zero was assigned if they stated that they could not understand the trend from the rendering, or if they stated a wrong trend. These scores were tallied and a ratio of correctly answered trends to total number of renderings presented was calculated.

4.4 Results

4.4.1 Pretest Questionnaires

The results of the first set of pre-experiment questions reflected a varied background among participants: four were congenitally blind; six were partially or legally blind. Four were proficient in Braille while the other four indicated having little to no experience with it. Three had familiarity using a force-feedback device, and three had experience listening to non-speech audio through a computer. All participants indicated a high-level of experience using computers.

From the results of the second set of questions, all but one participant stated they came across charts during their regular browsing on a computer or phone, to check weather updates, investments, politics, or scientific diagrams. Many stated their reliance on screen readers (e.g., JAWS), which can only recognize the type of graphic but not the information within the image. Consequently, users stated they were only able to obtain a glimpse of the information through related text context, such as an article that cited the graph. Some users indicated reliance on others to obtain the information they need, but most of them only used assistive tools on their phone or computer that were often lacking. As one participant shared their discontent:

> “So they will write it like the screen reader doesn’t know nor care, right? Like it’ll look at a table and it’ll say “a table with two columns and 155 gazillion rows” [. . .] you just arrow down, down, down... well, not that fast ’cause you have to hear the [table headers]. [. . .] it’s a poor man or a sad man’s chart. There is nothing that gets even close to the comfort you guys have of looking at a chart. Not only does [the graph] speak more [to the sighted user], but more importantly, it speaks faster.”

4.4.2 Task Performance

Every participant responded to each task question, although in some cases clarification was needed, but all of them were able to understand the questions asked. All three multimodal cases showed better performance than the audio-only case. However, there was no difference in performance among the three multimodal cases. No noticeable patterns could be obtained within groupings of congenitally and non-congenitally blind participants, or partially blind and completely blind participants, potentially due to a small sample size.

Participants were able to successfully deduce the graph trends approximately 74% of the time with unimodal audio, and subsequently 97% with the addition of haptic information. A breakdown of the success rates are shown in Table 4.1. The improved scores between unimodal and multimodal representations suggest that the addition of haptics, for each modality, improves understanding of line charts. Further investigation focusing on task performance is needed, but the results here give reason to be optimistic.

In addition, we performed preliminary non parametric analysis on the participants' answers. Results were inconclusive due to the limitations of the task design and small number of data points. However, although we cannot objectively deem renderings in one condition as being better than another, we hypothesize that each condition has different strengths when conveying information, based on our thematic analysis.

Despite high accuracy and similar results among each of the multimodal cases, all eight participants stated a preference for the multimodal Dot Pad rendering. Between the two multimodal force-feedback conditions, six participants indicated a preference for guidance, while two preferred exploration. These results, along with demographic details, are summarized in Table 4.2.

Table 4.1: Success rate of information retrieval for each multimodal condition against unimodal audio.

Interaction	Number of Trials	Unimodal	Multimodal
Dot Pad Exploration	16	81.3%	100%
2diy Exploration	8	62.5%	100%
2diy Guidance	16	68.8%	87.5%

Table 4.2: Demographic characteristics of participants, and preferred approach.

Participant	Visual Acuity	Age (Years)	Preferred Condition	Preferred Force-Feedback Condition
1	Congenitally Blind	70	Dot Pad	Guidance
2	Acquired blindness	69	Dot Pad	Guidance
3	Congenitally Blind	70	Dot Pad	Guidance
4	Acquired blindness	51	Dot Pad	Guidance
5	Congenitally Blind	65	Dot Pad	Exploration
6	Acquired blindness	61	Dot Pad	Exploration
7	Congenitally Blind	58	Dot Pad	Guidance
8	Acquired blindness	64	Dot Pad	Guidance

4.4.3 Qualitative Analysis

In general, we observed that haptics enabled participants to describe trends in greater detail. Participants also indicated that the haptic modalities provided them with a more accurate sense of vertical magnitude. Physically moving or being moved by the 2diy's handle, or guiding the fingers along the tactile graphic on the Dot Pad improved their understanding of changes in the vertical extent, which were seen as "compressed" in the audio renderings. For example, most participants were surprised that there were three minor waves before the peak in Fig. 4.2b, believing that section of the graph to be more "flat-line like" in the audio, and at how high the largest peak was. Finally, details present at the beginning or end of a graph that were sometimes not picked up by participants through audio were quickly perceived with haptics. Even after participants relistened to the audio in such cases, some had difficulty trying to reconcile the perceived "incongruency" between the audio and associated haptic modality.

Our thematic analysis identified four main themes, on which we elaborate below, that together answer our research questions. The first theme describes the initial and lasting impressions participants had with the rendering strategies, since such perceptions play an important role in their intention to adopt, and ability to use, a technology [111, 112]. The second theme considers the learning curves of each condition, assessing the difficulty participants experienced in adapting to each type of rendering, and determining whether there was any variation among the participants in their learning. The third theme describes how participants interacted with the renderings to obtain information, and the fourth theme relates to the mental workload associated with each condition.

4.4.3.1 Tactile exploration left a more positive impression than audio and force-feedback.

We observed that among audio rendering variables, tonality had the greatest impact on participants experience. Instead of a pure tone, we used a broadband sound that elicited mixed emotions. Some participants noted that the high audio quality made it difficult to follow sonification trends, perhaps because they associated the quality with real-life sounds (described on initial listening as “weather noise”, “shattering jet-engine”, “wind”, and “spirals”), which may have been distracting. However, the spatialization and timbre of the audio were well-received once participants understood the mappings, which was often after several listenings. One participant, who self-reported more experience with sonifications, remarked that the sound was pleasant to listen to for extended periods of time:

> “[Sound quality is] important if you’re going to do this all day [...] I used a synthesizer for 30 years and one day [...] I had to change [...] those sounds you’re using are good sounds, they are relaxing, they are almost cozy.”

For the 2diy renderings, the smoothness of the motion seemed to make the greatest impact on their impressions. Unpleasant interactions with the renderings were often due to rapid changes in directions or oscillations experienced at the end-effector for the guidance condition, or repeatedly encountering elastic forces for the exploration condition, which was observed in participants trying to navigate the graph quickly.

Participants were more at ease exploring on the Dot Pad, likely influenced by familiarity with the device’s affordances and its static nature. Experience with Braille likely facilitated understanding of the device’s characteristics for some participants, boosting their confidence in exploration. Additionally, participants were more relaxed with the Dot Pad display, not needing to grasp a physical cursor in order to explore the graphical components.

> “The top fingertip [Dot Pad], one was better in that it was something I wasn’t having to grapple onto it, being guided along or forced through. The information was there to explore so it was more interactive. It was more engaging rather than being followed along or trying to fumble your way through.”

4.4.3.2 Participants took some time to adapt to the sonifications and 2diy renderings.

Participants stated that audio renderings posed the greatest challenge for comprehension, likely attributed to the unfamiliarity with the sounds used and the fast pacing of the renderings. Although three participants were able to interpret low-level details through this modality to the same degree as they did with audio-haptic representations, the remaining five found it difficult to react to small and instantaneous changes in pitch, especially during sections of high volatility in the data trend. As one user remarked:

> "The premise is trying to thinking about the stock prices, it's like learning another language [. . .] I'm sure artists or musicians would probably do really well with this because their ears are fine tuned in a certain way. But learning the acoustic aspect of it, there's a learning curve [. . .], less so with the tactile device."

Participants expressed that they only needed a single example to understand our rendering strategy on the Dot Pad, and exploring on the Dot Pad required the least effort to obtain information. They also stated that this strategy provided them with the highest level of confidence in the data. When asked to identify the approximate difference between the first and last point in the graph, this condition provided the easiest approach, as all participants put one hand on each side of the graph and were able to mentally compare the difference very quickly. Such techniques were not possible on the 2diy or with audio representations, leading to participants feeling uncertain about their responses to the task. Additionally, an unexpected finding with the Dot Pad was how the tactile haptics improved their retrospective understanding of the audio representation. This suggests that training with multimodal tactile renderings could facilitate a user's interpretation with unimodal audio experiences.

Participants were less accustomed to using an end-effector, particularly in the force-feedback guidance condition. Applying excessive force on the handle unintentionally impedes its motion, while a lack of sufficient force hinders effective shape perception. Maintaining a fixed and uniform force on the end-effector for a smooth experience can be challenging. Participants took time to "optimize their grip" on the device, which led to inconsistencies in force sensations during their initial guided tour experiences. However, they found the device functionally useful when the rendering experience was steady, which was typically after they had replayed the guided tour two or three times. As compared to tactile rendering, participants require extra time and effort to familiarize themselves with the force-feedback renderings, likely due to the increased complexity of the device. In this respect, the learning curve was similar to that of audio-only renderings.

4.4.3.3 The 2diy has potential to convey volatility, but users preferred the Dot Pad's autonomy and clarity.

The 2diy device provides a single point of contact, which inherently reduces the information bandwidth available on the haptic channel. In the guidance condition, this precluded participants from obtaining a quick overview as they had to wait for the end-effector to go through take them through the graph sequentially from left to right. However, participants liked the guided tour, citing its ability to communicate volatility in the data effectively, and found it to be engaging because it guaranteed that the end-effector physically moved their hand along every point in the graph. This was in contrast to the exploration conditions, for which there was a possibility of missing out on information. Additionally, although the exploration condition facilitated fast graph navigation, participants had difficulty "correcting their path" if they had unknowingly strayed from its trajectory and encountered a boundary. We observed that this resulted in extended exploration time, equal to or surpassing the time spent in the guided tour condition.

The Dot Pad allowed for rapid two-handed exploration. Five participants primarily used one hand as a reference point while navigating the chart with the other, whereas three employed a two-handed sweeping motion along the grid. Despite these differences in the interaction, all participants naturally switched to the "reference point method" when tasked to compare values. One user stated that it was easy to follow along the tactile chart because their fingers covered pins representing both the current "finger location" and the next few points on the rendered graphs. This allowed them to anticipate patterns in the data more gradually compared to the audio and force-feedback conditions.

A missing affordance stated by four participants was the lack of control over the rate at which audio played or the 2diy guidance rendering moved the end effector, contributed to negative emotions, such as frustration or confusion, in relation to these strategies. It became clear that autonomy—control over one's interactions—was an important feature in understanding the data. The ability to use multiple points of contact and control the speed of the interaction resulted in a unanimously positive experience for the Dot Pad condition.

> "You have to be a slave to the [2diy]. I much prefer the Braille display if there's a question to be asked."

> "And if [the speed] is user adjustable [. . .] then you could go several times over [the rendering] and decide what level you're getting more information. If [the speed is] a fixed value, then you don't have that control."

4.4.3.4 User-controlled strategies result in lower cognitive load.

We found that the "user-controlled" strategies, i.e., the Dot Pad and 2diy exploration conditions, were in general more efficient, with reduced demands on short-term memory. When participants were asked if they could locate a specific region on the graph, they found it easy to take control of the devices and query information as needed.

In contrast, participants requested to play back the audio and 2diy guidance tour multiple times to create and refine their mental mapping of the plot. This was the case both when trying to parse the graph contents to obtain an initial overview of the data, and when locating specific points on a graph. As a result, participants required more time to answer questions in these conditions. In addition, for the 2diy guidance condition, they had to focus on following the end-effector motion, which created an additional layer of cognitive load on top of parsing the graph contents. Moreover, they had to wait for the entire guided tour to complete, and then recall specific details in the graph, creating an additional memory burden. Another issue with this condition was inconsistency in the guided motion due to unintentional resistance by the user:

> "I'm not sure I'm honoring the movements that should happen. I think I'm blocking them unwillingly [. . .] I can see I'm lagging right? There is a delay between what happens and what my hands is realizing and just trying to compensate ... I got hacked by the lag [. . .] I was destroyed. Part of my concentration."

This effect could be analogized to sighted individuals watching a video encoded with a noticeable variation in frame rate, which would similarly distract from the experience of parsing the video contents [8]. One participant suggested that an option to pause and replay a desired selection of the guidance could partially alleviate this cognitive burden. Currently, if participants encountered difficulty understanding any part of the graph, they had to wait until the guided tour concluded before they could replay the rendering. This led to a time-consuming experience, as the distraction from the perceived discrepancy hindered their comprehension of the remainder of the guided tour. Navigating in segments could allow them to understand the information more clearly.

Some participants also expressed difficulty in building a mental model with the 2diy exploration condition, as encountering the elastic forces interfered with their memory of their previous position. We considered an alternative approach for this rendering strategy, which involved modelling the graph lines with strong damping textures that resist user motion. In this approach, participants could identify their position on a line if they had to exert more force on the end-effector to navigate along the path. However, we opted for the current approach because we believed the high variability in the shape of a typical time-series visualization could make it challenging for participants to stay on the line. Nonetheless, this deliberation highlighted the significance of considering multiple rendering strategies for the exploration condition, as certain approaches may offer more effective contextual cues than others.

Finally, proprioceptive skill is known to degenerate in elderly individuals due to their decreased dexterity and cognitive processing abilities [113, 114]. This disparity becomes more profound with larger distances in movement [113], suggesting that sudden and large spikes in data trends would be harder to keep up with in the 2diy guidance condition. We observed a high level of hand movement variability in one of the two participants who preferred the exploration condition, suggesting that guided movement was difficult for them to follow. In contrast, tactile acuity appears to be stronger in blind individuals than sighted participants and shows no age-related decline [115]. This may be one reason why our participants were able to navigate the graph contents more readily with the Dot Pad.

4.5 Discussion

A summary of the observed benefits and drawbacks to each condition is shown in Table 4.3. All three multimodal strategies allowed for better understanding of line graphs compared to sonification-only cases. This suggests that the haptic renderings were effective in delivering information about the graphs regardless of differences in their affordances. However, all participants preferred the multimodal Dot Pad strategy, suggesting that the lack of comfort and usability may have detracted from their appreciation of our sonifications and 2diy renderings. We believe that some of these issues are easily addressable, e.g., adding controls to navigate the rendering in piecewise fashion and adjust the pace of a non-user controlled rendering to a user's liking. Through analysis of recorded transcripts, we identified three factors that may impact user preference: potential improvements to the rendering strategies, targeted training, and ensuring greater temporal and spatial congruence to support user perception.

Feedback from our participants indicated a clear preference for having control over the rendering process, in particular for the pace of the rendering. Moreover, a heavy reliance on short-term memory was found to hinder the participant's learning abilities. To mitigate this issue, the incorporation of memory aids, such as tactile reference points, or control over audio speed or pitch range, should be employed. Our observations of the participants also revealed difficulties with understanding magnitude differences in charts, which was mapped using pitch. This is likely due to the fact that pitch discrimination deteriorates with age [116]. Moreover, recent statistics on the BLV community [117] suggest that this population tends to skew older, as many of the conditions leading to blindness occur later in life. Given that all our participants were above fifty years of age, the use of pitch as a means of conveying spatial information may pose challenges.

With the 2diy, ensuring stability is essential for creating positive initial and lasting impressions. This can be achieved through modifications to the control algorithm or through mechanical changes to create a smooth experience. In this respect, we have recently observed significant improvements in the most recent generation of 2diy devices, coupled with our implementation of a more advanced control strategy. For this force-feedback device, we also asked participants whether they preferred the exploration condition or the guided tour. Ultimately, it may be beneficial to combine these strategies, allowing users to explore the data after an initial audio-plus-guidance pass that provides the initial context about the graph. Providing information from multiple force-feedback perspectives in this way could enable a clearer understanding of the data than either strategy alone [65]. A combined rendering strategy could improve the usability of the 2diy device to the point where it becomes a viable alternative to the Dot Pad to explore time series data.

4.5.1 Importance of Familiarity and Targeted Training

Another major takeaway from the analysis was the need for more targeted training sessions or examples. The results indicate that the Dot Pad was more "intuitive" for participants. However, this may speak more to the participants' familiarity with the components of the system rather its design [118, 119]. This was most evident with the audio renderings, which many participants had to repeat multiple times to be able to understand them, but gained familiarity and comfort through practice as the study progressed. However, as revealed in Theme 4.4.3.2, our findings also indicate that there are participants who possess a high level of comfort with sonifications. Therefore, their training could be geared towards delving deeper into the technical aspects of the renderings to further refine their skills.

One observation from the study was that haptic feedback can not only reinforce information but also aid in the comprehension of audio renderings. With this in mind, incorporating the Dot Pad or a similar tactile device into future training could be beneficial. However, due to the cost of the device, it may be more feasible to host it in a public community space, such as a library, where users can have periodic access to improve their proficiency in using the audio renderings.

4.5.2 Temporal and Spatial Congruence as a Key Factor for Multimodal User Experiences

In this study, participants indicated that they felt a discrepancy in the duration of the audio and the haptic guidance. We presented audio and haptic feedback sequentially, which may have contributed to the perceived discrepancy. If both were presented simultaneously, the effect may have been less pronounced. However, synchronizing pre-made audio files with force guidance will likely be challenging, especially due to variations in grip and weight applied to the end-effector, which can alter the friction applied to it. Nevertheless, there is potential for improvement both from refined software control of the 2diy, and moving to a newer version of the device, which features a magnetic pen-shaped end-effector instead of a knob, as well as enhanced motors and encoders. These improvements could enhance position tracking by ensuring the pace of guidance aligns with that of the audio file—that is, positioning the user's hand by the 2diy at the corresponding location in the data series being described at that moment by the audio track. Another possible solution is to use real-time audio synthesis, taking advantage of the 2diy's existing position-tracking capabilities to provide dynamic audio in real-time to convey details.

We also observed a significant difference in how users perceive the magnitude of values in the vertical extent between audio and haptic renderings. This difference may be attributed to the vertical-horizontal illusion, a well-known phenomenon in haptic literature that suggests people tend to perceive changes in the vertical direction to be greater than the equivalent changes in the horizontal plane. Renier et al. [120] hypothesize that the origins of the phenomenon are from the visual modality, and therefore, should have no effect on congenitally blind individuals who have never been introduced to the visual illusion. Indeed, they found that the illusion did not have an effect on the congenitally blind. However, in our study, we observed a high occurrence of this discrepancy even in congenitally blind individuals, suggesting that our findings may be more related to our rendering strategy than to the illusion itself, specifically the note-to-pitch mapping used for sonifications. Although we used one of many such established mappings [121], the existing literature does not speak to their effectiveness when used in conjunction with another sensory modality. Considering our observations indicate a potential link between the mapping range and perceived discrepancy, additional research is needed to understand how the perception of this mapping is influenced by the introduction of a secondary modality.

4.5.3 Limitations

The generalizability of our findings to other populations may be limited due to a small participant pool, which only consisted of individuals over the age of fifty. As well, due to time constraints, we presented only two charts for the 2diy guidance and Dot Pad exploration modes, and one chart for the 2diy exploration condition. Because of the limited number of examples provided, users may not have been familiar with the rendering strategy used, which could have influenced their evaluations. User preference may have also been impacted by a novelty effect, as many users experienced renderings designed with these approaches or devices for the first time. Longitudinal studies could help mitigate some of these effects.

We did not incorporate synthesized speech into the audio for these renderings. Given the widespread use of screen readers to convey information, it is possible that the inclusion of speech to provide more targeted information may affect a user's preference or willingness to use a specific approach among the three haptic conditions.

Finally, although this study broadly sought to improve accessibility in understanding graphical content, the results and implications presuppose that users will be able to set up and operate the devices in a real-world environment. It is important to be acutely aware of and address the numerous technical issues associated with operating these haptic tools [8] in an individual home setting.

4.6 Conclusion

We developed and presented a number of audio and haptic renderings to convey information in line charts. Based on our results, users gained a better understanding of the data with the inclusion of haptics. Audio was generally stated to be useful in providing a quick overview of the trends, whereas the haptics allowed users to perceive the shape more accurately. As such, the multimodal interfaces convey line chart information more effectively to BLV users than the baseline audio sonification. To understand the differences among the three multimodal conditions, we conducted a qualitative analysis. We analyzed recorded transcripts from eight participants and devised six themes that can shape user preference: learnability, cognitive demands, audio-haptic congruence, user background, implications of the affordances with each modality, and users' initial and lasting impressions.

To our knowledge, this is the first study that compares user experience between force-feedback and refreshable tactile representations of graphs in combination with audio. Considering the higher cost of RTDs, this study provided valuable information in showing that when presented directly after sonifications, force-feedback can convey graph details to a similar level of accuracy, despite differences in device characteristics and rendering strategies. Although findings in this work revealed limitations and technical issues with the audio and force-feedback conditions, some benefits also emerged. These observations and design considerations can provide a foundation for further development of best practices in representing visualizations through multimodal audio-haptic interfaces.

Table 4.3 Summary of the observed benefits and drawbacks to each condition.

Condition	Observed Benefits	Observed Drawbacks
Audio	• spatialization and pitch mapping were well received • quality of audio was liked and sounds used were pleasant to listen to • generally quickest at providing overview of general trend of data	• requires a trained ear to understand smaller variations in data • non-user controlled strategy means more concentration is required • pitch mapping limits population • sounds used can make relative differences difficult to compare • users found the pacing too fast and had to playback the audio a number of times to make judgements when answering questions
2diy Guidance	• communicates shape and volatility of line best when experience is smooth • user preferred 2diy condition for accuracy	• non-user controlled strategy means more concentration is required • prone to technical issues, particularly with stability and lag • requires most psychomotor ability
2diy Exploration	• allows for exploring at own pace • allows for quicker exploration than guidance once trained • can navigate back and forth in graph, unlike guidance • flexible approach, as there are multiple ways to render the desired information to provide contextually relevant cues, such as virtual walls, elastic forces, or textures	• focus required to navigate graph path • virtual walls were found to be too stiff and unwieldy by some users
Dot Pad	• can anticipate patterns in the data at explored finger location • two-handed interaction allows for more complex exploratory procedures • ability to easily maintain interaction with device fosters a more comfortable exploration experience • allows for exploring at own pace	• None observed for the rendering strategy (but device cost is a factor)

Chapter 5

Interactive Multimodal Chart Experiences with the Dot Pad

Preface

Findings from the previous chapter suggest that BLV users prefer using an RTD to interpret line charts. However, one limitation of the delivered renderings was the lack of speech audio to provide targeted information about points of interest (POIs) on a graph. Here, we expand upon these renderings to provide additional information using an interactive touch experience. We present a new prototyping tool that can detect touch gestures on the tactile grid through an IR touch frame to deliver audio events. We then showcase richer line chart renderings created using this tool. We believe these renderings could potentially replace tactile diagrams in classroom settings to provide a higher quality learning experience.

Author's Contribution

This research was a collaboration effort between Dr. Yongjae Yoo, Juliette Regimbal, Cyan Kuo, Sri Gannavarapu, and Rayan Isran as work for the IMAGE project referenced in Chapter 4 and Section 1.1. Dr. Yongjae Yoo assembled the hardware and produced a proof-of-concept demo for an interactive audio-haptic experience, whereas Sri Gannavarapu and Rayan Isran conceptualized the framework for rapidly creating experiences and wrote the software. Juliette Regimbal, Cyan Kuo, Sri Gannavarapu, and Rayan Isran all brainstormed design ideas for

multimodal chart renderings, and Rayan Isran implemented them. Juliette Regimbal provided additional ideas for sonification renderings.

5.1 Motivation

The increasing use of haptic interfaces in computer systems [1] has created a need for software that allow users to create, modify, and distribute haptic experiences [122]. However, developing such systems is a time-consuming process that typically requires programming expertise. Additionally, designing effective haptic experiences can be complex and unintuitive, necessitating technical knowledge of haptic rendering and interfaces. Such challenges pose limitations for educators who wish to leverage these these haptic technologies to enhance the classroom experiences for BLV students. To address these obstacles, we present a prototyping tool that allows users to create multimodal audio-haptic experiences without requiring any technical programming knowledge. We developed this tool specifically for the Dot Pad as it was the most intuitive device for BLV individuals to use, based on findings from the previous chapter.

The tool was created with three major goals in mind. First, it should allow designers to actuate a desired set of pins such that any tactile image can be rendered on the device. Second, there should be an easy way to deliver targeted audio events upon touch inputs. Last, it should provide functionality to take in data of a specific format, and render a predefined audio-haptic experience. For example, line charts typically have two axes, and the data line trends upwards when the value increases. To accurately represent a line chart through the tactile interface would require the user to manually highlight pins that collectively represent the data. The line chart could be treated and rendered as a bitmap, but data in the bitmap would still have to be accurately mapped to the desired locations on the pixelmap, and any information in the graphic that might be undesirable to display on the tactile grid (e.g., chart title, axis values, or extraneous info in the graphic) would have to be removed manually. Each of these approaches is time-consuming and error-prone. Therefore, a framework to prototype a quick and consistent experience should exist.

5.2 Architecture

The proposed solution presents two modes: an *editing* interface and a *rendered* interface. The editing interface allows a designer to prototype their interactive experience. This includes both a quick and easy method to actuate pins on the tactile grid of the Dot Pad, and deliver audio according to a user's finger gestures along the actuated pins. The rendering mode would then generate this interactive experience for the end user. When creating the tool, four aspects were considered:

- A quick and simple method to actuate a desired set of pins, defined as a region, on the device.
- A framework to handle the interactivity between touch inputs and audio events at a region.
- A way to save the created interaction experience so that it can be loaded or quickly modified later.
- A framework that allows developers to create custom plugins that can rapidly generate prototypes for specific types of graphics.

This section discusses how the first three points are incorporated into our system, and the subsequent section describes the last point—i.e., the plugin framework, with examples to show how it can be used to produce audio-tactile line chart experiences.

5.2.1 Pixel Editor & Rendering

A 60×40 raster grid was created that corresponds to the tactile grid of the Dot Pad. Designers can use various editing tools to mark or erase tiles on this grid, such as a virtual pencil or eraser. A render command on the interface can actuate the pins on the Dot Pad that correspond to the tiles marked on the editor. This command transmits packets of data through BLE-5's GATT protocol. These packets contain device information, a checksum, and the status of each pin. The device firmware organizes these pins into sets of 4×2, which represent one byte of information. Each byte contains the the height status of each pin within this group ($0 =$ lowered, $1 =$ raised). As there are 300 groups of pins on the tactile grid, 300 bytes are sent to convey the tactile image.

To create audio events tailored to touch inputs on individual regions, there must first be a way to distinguish between these regions. Color coding is used for this purpose. The designer can select a region from a drop-down menu. Tiles drawn with that region will be indicated by a unique color. A mockup of the initial interface design is shown in Fig. 5.1.

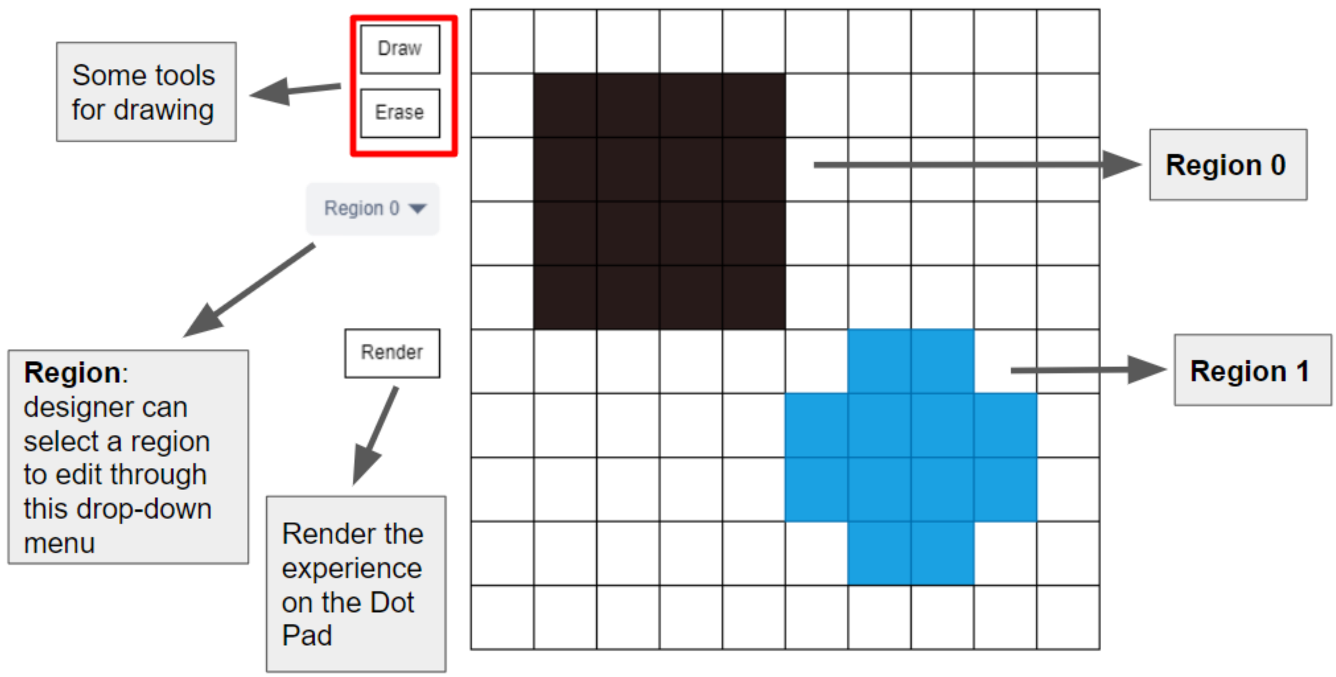

Fig. 5.1: Initial mockup sketch of the interface, on a scaled down 10 × 10 tactile grid.

5.2.2 Hardware

As the Dot Pad does not contain any finger tracking capability, the tactile grid was overlayed with a 10.4" third-party touch screen IR frame connected via USB connection to a computer. The frame contains a set of IR transmitters on the top and right sides of the panel, and a set of IR receivers on the bottom and left sides. As the user touches a location, the path of the IR is disrupted and the rays do not reach the receiving end. The processor can use this disruption to pick up the coordinates of the touched location.

5.2.3 Touch Calibration

To detect gestures, the HammerJS[1] package was used. This package recognizes a variety of inputs, such as single-tap, double-tap, and panning motion events. Parameters can be set for each recognizer to provide additional control, for example, the hold time required to detect a

[1] https://hammerjs.github.io/

tap or minimal pan distance to detect panning motion.

Acquiring the (x, y) coordinate tapped on the Dot Pad through the IR frame requires a one-time calibration procedure as there is no communication between the two devices. For this, several equidistant pins were actuated on the tactile grid both horizontally and vertically using this prototyping tool. Each location was tapped multiple times, and the position values detected by the IR frame were recorded. An average position value was computed for each position. A linear mapping equation was then empirically derived using the sets of actual coordinates and corresponding IR frame position values.

5.2.4 Gesture Detection

A designer should be able to an assign an audio event to a region based on a detected touch gesture. Additionally, different gestures on the same region should be able trigger different events. For example, suppose a region is marked by a row of tiles moving from left to right on the grid, i.e., represented as a horizontal line. The user could tap a point on this line to read out its coordinate. The user could also swipe across this region to play a spatialized audio rendering of the coordinates that moves from the left ear to the right ear. To consider these possibilities, a prepopulated drop-down menu was added for touch gestures and a list box shows the event associated with that region and gesture. While functionality for multiple inputs and audio events was considered, only a single gesture and audio event was programmed for simplicity: a single-tap input, and a text-to-speech (TTS) audio output.

As the tool runs on a local web server, the Speech Synthesis Utterance API[2] is used to deliver speech. The user can input the text to tag through a text tool. This text is then added to the currently active region and gesture indicated by both drop-down menus respectively.

An example of this interaction process is shown in Fig. 5.2. We first check whether the tapped location corresponds to a region. Then, we determine whether this region has any TTS event associated with the single-tap gesture. If so, the TTS audio event is fired.

Because of imprecision in detecting the exact gesture location due to inherent noise from the IR sensors and imperfections with the calibration procedure, some leeway was provided in that a touched point close to a raised pin may be considered as the location of the pin itself. This was done by computing the Euclidean distance between the closest actuated pin on the

[2] `https://developer.mozilla.org/en-US/docs/Web/API/SpeechSynthesisUtterance`

tactile grid and the user's detected location. If this value was within a certain threshold, the detected location was equated as the closest point for gesture detection.

(a)

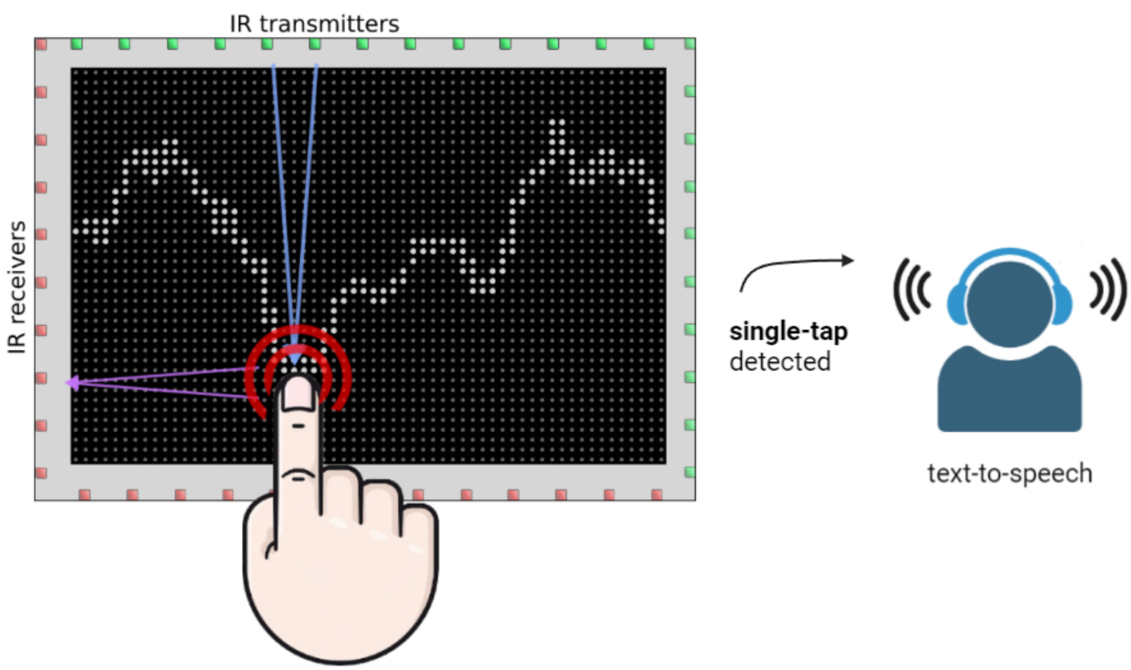

(b)

Fig. 5.2: Overview of the flow. a) Simplified flow of the process. b) A top-view mockup showing how this interaction scheme operates with a single-tap gesture indicated.

The data associated with the entire interaction experience, i.e., the tilemap information for marked regions, their associated gestures and audio events, is stored in a JSON file. This information can then be loaded by the tool for later usage.

5.2.5 Practical Example

Consider the scenario of conveying details of a map through this audio-tactile interface. A designer could recreate the map on the editor, using different regions to tag POIs with a TTS event using a single-tap gesture action. The generated experience could allow a user to explore the tactile map and tap on certain areas of the grid for additional information. This would allow them to obtain further details as needed, as opposed to listening to an audio rendering beforehand.

While such an interactive experience could provide a richer experience than provided by the haptic or audio modalities alone, the process of recreation can become tedious, especially if there are multiple maps. If a change to the experience is desired, each map rendering must be edited separately. To overcome this limitation, the framework also allows developers to create plugins that add predefined experiences for specific types of graphics. A developer can write code to detect data of a specific format and mark it as a specific type of graphic. Then, separate touch events can be called to run predetermined code to generate desired events if this type of graphic is detected. Therefore, if a change is desired, only the existing code that generates the experience must be edited rather than the interaction file itself.

In the following section, we demonstrate a chart plugin that was created and present examples of how it can be used.

5.3 Chart Renderings

To expand tool functionality, a plugin was developed that enables designers to rapidly generate multimodal renderings for line charts. This tool contains a user element that prompts a designer to load HighCharts JSON data and performs three major operations:

- It manipulates the graph data as necessary to fit the line using as much of the tactile grid space to allow for higher resolution exploration. For example, a graph containing 15 points is linear interpolated to utilize all 60 columns of the tactile grid. Conversely, a moving average filter consolidates a graph of 100 points to squeeze into the 60 column width.

- It renders this manipulated data visually to the bitmap editor and adds x and y axes so that a tactile graph can quickly be rendered to the Dot Pad. The entire line chart is assigned a region, but any edits can easily be made on the interface. Therefore, users do not need to recreate the graph on the device manually.

- It automatically adds a speech and non-speech audio experience based on certain user gestures. This is described in more detail in the subsequent section.

5.3.1 Gesture-Based Audio Experiences

The line chart plugin automatically adds two types of audio experiences to a chart: synthesized speech and non-speech audio.

The speech experience is separate from the TTS tagging designers can add to a region on the bitmap. It is activated upon a single tap gesture and narrates the title of the x-axis or y-axis, or coordinate information depending on the tapped location. The algorithm reads the closest point on the graph using the tapped location and reads the x-value, followed by the y-value.

A synthesizer was created using the Web Audio API[3] to deliver the non-speech audio experience. The synthesizer uses a sine wave oscillator to deliver continuous dynamic audio as the user pans their fingers across the rendered pins that form the line chart. The oscillator uses spatialization and pitch mapping to represent the change in values along the x and y-axes respectively. As the user moves along a segment of the line that trends upwards, they hear a gradual increase in pitch and vice versa. The audio would stop if they move their fingers off the tactile grid or away from the line. This was done by computing whether the difference between the user's detected (x, y) coordinate from the IR frame and closest rendered point on the line was within a predetermined threshold.

To see how this plugin could be used, we describe its use for single and double line charts.

5.3.2 Single-Line Charts

Here, the designer can load a specific single line chart example. This visually renders it to the bitmap and generates an audio-haptic experience. The rendering model is similar to that described in Section 4.2.3, but also includes dynamic and sonified audio to complement the cutaneous exploration. Users can tap POIs on the tactile grid to obtain information about the

[3] https://developer.mozilla.org/en-US/docs/Web/API/Web_Audio_API

coordinate through synthesized speech, or pan their fingers along any section of the tactile graph to perceive the trend of the data in that section through non-speech audio. These two techniques help them understand a gestalt overview of the data, as well as any targeted information.

However, while the single-line experience can provide a richer form of access to chart data, it is limited in that only one-dimensional data for one trend can be displayed on the tactile grid at a time. Comparing multiple trends of the same type of data, such as the COVID-19 count of a country against the worldwide count, would involve loading and rendering two separate charts. This would be time-consuming and require users to keep a mental mapping of each trend in their short-term memory. To overcome this limitation, we expanded the chart plugin to support an audio-tactile rendering for two-line charts. However, using the same single-line rendering approach would not be feasible in this case, as it would be difficult for users to distinguish between two lines using the same approach. Therefore, we discuss additional rendering techniques to overcome these challenges.

5.3.3 Multi-Line Charts

As with single line chart data, the plugin takes in real multi-line chart examples from the web that utilize the HighCharts charting framework. This data file contains information about each line in the plot. As shown in Fig. 5.3b, the plugin renders each line on the bitmap as a separate region, indicated visually by a separate color. The overlapping of values was not accounted for, so the second line trend in the extracted data takes priority on the bitmap display.

Functionality is similar to the single-line audio case: both lines are rendered to the tactile display, and non-speech and speech audio renderings are automatically added to the experience. The y-axis value mapping was retained, so users would hear a change in pitch as they navigated along up and down along a line. The frequency of the audio oscillator to generate the sine wave tone was preset to use a different frequency for each line, separated by one octave. In this example, the carrier frequency for one line was set to 100 Hz while the second was set to 200 Hz. These frequency values were chosen because the difference between them is significant enough for non-musicians to easily distinguish [123]. Therefore, users would hear a 100 Hz tone while exploring one line and a 200 Hz tone exploring the other. They would hear a change in tone if they navigated to the other line on the tactile grid. This would help them identify which line they were exploring, and they would only have to remember the line by the encoded tone frequency.

We also hypothesize that users would be able to detect points of intersection on the graph with greater confidence through this interactive audio-tactile approach than a tactile-only one. The system continuously detects the user's position while exploring a line. If the user reaches a point that intersects with another line, the system detects the line that was just explored and switches to the tone that represents the other one. Even if the user remained on the line that was explored before touching the overlapped point, there would be a brief, audible tone "swap". The system's capability to detect points and switch tones could allow the user to discern intersecting data points accurately and quickly.

However, users still have to retain a mapping of each line in their short-term memory, as they can only listen to one audio rendering at a time. Therefore, comparing the trends requires additional cognitive processing. For example, if a user wants to understand the y-axis value difference between the two lines at a given x-axis location, they would have to remember the value of one line at the current location, move their finger to find the location of the value of the other line on the tactile grid, mentally compare the two values, and then repeat this process for various locations on the graph to obtain a gestalt overview. Thus, switching between the two graphs using the multimodal approach requires high short-term memory needs.

To cope with this challenge, a secondary, *sonar* mode was created; this includes an additional feature to denote the instantaneous difference between both lines at a selected x-axis location, represented by intermittent pulses in the background audio. The user would hear the audio tone while exploring the current graph as before, as well as additional pulses in the background from a secondary audio oscillator at a lower volume. The rate of pulsation was set to be a function of the difference between the y-values at the current x-axis location, so the sound would pulsate faster when the values swing closer, and slower when they drift further apart. In this way, the user would be able to tell the approximate difference between the corresponding values on each graph while navigating on a single line. However, as this approach only denotes the difference, a user would not be able to discern whether the touched point is situated above or below the value of the other graph. To address this, we use the pitch of the pulse was to indicate the direction. If the current point is higher, the frequency is set to 450 Hz and the user would hear a higher-pitched tone; 100 Hz if it is lower. A two-step octave difference was used to allow for easier discrimination than the two separate tones while switching between graphs.

The aim of this approach is to allow users to explore at their own pace while relying on minimal extra cognitive processing compared to the single-line chart overview case. The organization and presentation of the various non-speech audio elements is intended to facilitate a gestalt overview of the information, while touch readouts can be used to scan for details. However, users may prefer to interact with only one or some of these audio feedback components. It is also likely that users would require a lot of training to effectively make use of these rendering approaches. Even so, the usefulness of these techniques can only be accurately determined through user studies, which we did not conduct.

5.4 Challenges & Future Work

In additional to the challenges of presenting multi-line chart experiences, several other obstacles must be overcome to achieve the broader goal of creating more accessible experiences with this tool:

- The interactivity is not robust. The Dot Pad device does not feature internal finger tracking capability, so an external IR touch frame was used for gesture detection. This detection was not always precise and relied heavily on software correction techniques. It can be difficult to detect gestures due to the different ways in which users orient their fingers to interact with the tactile pins. When we tested our multimodal chart experiences, the initial panning motion for sonification was sometimes inadvertently detected by the TTS readout algorithm; such an unexpected disruption can result in a jarring experience for end users, and, as stated in Chapter 4, bad initial impressions can have lasting effects.
- Secondly, while the showcased designs aim to provide users with a richer representation of the data, future work should determine how users would interact with these experiences and whether they are more helpful than other multimodal techniques of interacting with data, or existing strategies employed by users. For example, testing is required to determine whether users would find the previously described sonar mode useful in quickly understanding the variation between two lines, or whether this would require

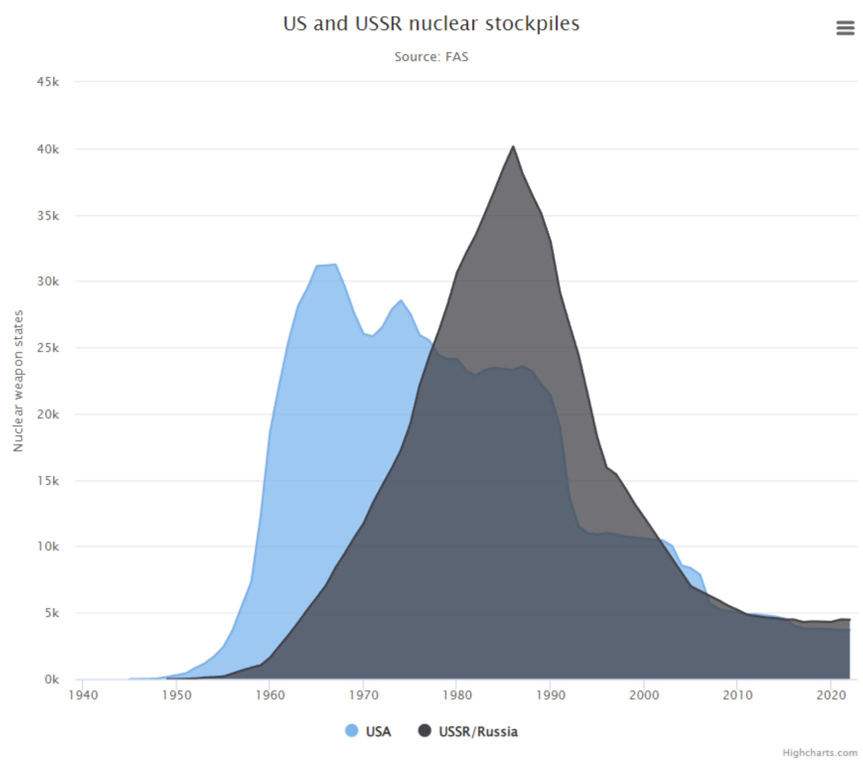

(a)

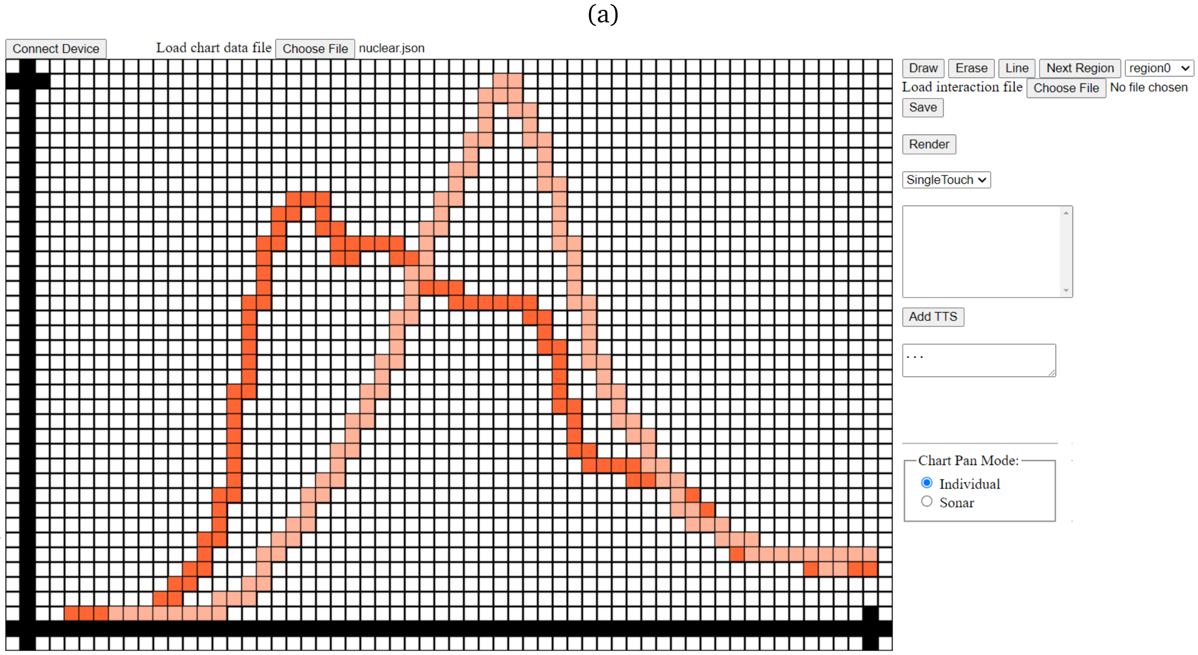

(b)

Fig. 5.3: Example of a double line chart rendering. a) The web graphic, taken from HighCharts. b) The loaded chart on the interface with the *Individual* audio rendering mode activated.

extra cognitive processing. Based on suggestive results from Chapter 4, it is possible that training would be required for them to effectively make use of this approach.

- The present process of setting up the tool is not easy. It requires users to run the tool on a local server, install an extension, and download the file data separately. There may be better ways to obtain the information to load onto the interface, such as directly integrating this tool into an extension.

- Lastly, although the work described in this chapter primarily focused on single-line charts, future work could include work on other types of charts, such as bar graphs or scatter plots. For example, plugins could be added to support multimodal interactivity for these types of graphics.

5.5 Conclusion

In this chapter, we presented a haptic authoring tool for an RTD. The core framework features a point-and-click bitmap editor that allows designers to actuate a desired set of pins. These pins can then be tagged with TTS audio, allowing designers to rapidly prototype audio-tactile images without requiring technical programming ability. Users can interact with these graphic representations through various types of finger gestures to hear synthesized speech that can provide additional or related information cues.

Developers can also use the framework to design plugins that focus on rapidly producing renderings for certain types of graphics. We demo a sample plugin for single line charts that loads data from real world examples from the web. The data is rendered visually on the bitmap, and both speech and non-speech audio experiences are automatically added to the interaction. With speech, users can tap points of interest on the rendered chart to hear readouts of desired coordinates. With non-speech audio, users can move their fingers along the chart to hear a sonified experience of the graph, in which x-axis values are spatialized and y-axis height values are mapped to pitch.

The plugin was then expanded to allow for two-line chart exploration, using a separate timbre to represent each line while retaining the same height to pitch mapping for the variation in y-axis values within each line. As a result, users can identify the line they are exploring based on the sound timbre, and detect intersection points through a change in tone. To further leverage audio capability in this plugin, we created an optional sonar mode. Here, the vertical difference between the two lines at any x-axis location on a line is represented by a set of pulses in the background audio, with the pulsating rate denoting the magnitude in difference. This feature was intended to help users understand how the lines vary without needing to lift their fingers off the tactile grid, thus minimizing exploration time and short-term memory load.

We also demonstrated that the tool has been designed in a way to allow developers to design modules specific to certain types of graphics, such as other types of charts, maps, or photographs. These modules can provide consistency in generated renderings within each type of graphic by applying the same mapping schemes to them.

The work described in this chapter sought to understand ways in which interactive experiences could be designed to exploit single-point-of-contact touch. From a practical perspective, users may find multi-finger interactions with the Dot Pad device more natural, as we saw in results from the user study described in Chapter 4. A multi-finger paradigm may provide several benefits to the user experience, such as enhanced spatial perception to gain a better understanding of the relationship between data points, and faster navigation by exploring multiple sections of the data simultaneously. However, this approach could add complexity to the interaction design, and lead to confusion if gesture-audio interactions are not well-designed; differences in the way stimuli are sensed among individuals will likely be more pronounced when multi-finger approaches are used [124, 125]. The use of multiple fingers to interpret audio cues could also increase cognitive load, potentially affecting the user experience. Follow-up work could compare single and two-finger exploration methods to understand how users could benefit most effectively from an audio-tactile experience through participatory design.

Although much future work remains to be done, both to improve the interface design and understand user experiences with this approach, we believe this tool can be used in classroom settings to replace traditional tactile diagrams. Despite a higher initial cost than Braille embossers, this tool could provide a significant long-term benefit for multiple reasons. Interactions can quickly be modified, saved, and stored electronically, thus eliminating the overhead of printing several diagrams that require a constant supply of costly material, physical space, maintenance, and time resources. Specialized plugins can be developed and used to rapidly prototype certain types of graphics that educators can easily share with one another.

Importantly, the interactions can potentially provide a richer and interactive form of access to learners, providing them with a higher quality educational experience.

Chapter 6

Conclusion

Haptic feedback has the potential to effectively convey information to a user, both passively and actively. A majority of the research described in this thesis focused specifically on delivering information in line charts through active exploration by a user. This understanding can be enhanced when non-speech audio is delivered prior to a haptic representation, as demonstrated in Chapter 4. Although unexplored in this thesis, a practical extension of the psychophysical work described in Chapter 3 could investigate whether or how accurately vibrotactile stimuli could convey the same information, given the significantly cheaper cost of vibrotactile devices over force-feedback and tactile interfaces. In this case, vibrotactile actuators could be placed along the forearm, and frequency or amplitude mappings could be used to represent changes in chart data. While some studies have focused on using vibrotactile feedback for data visualization [126, 127], there are practical factors to consider when delivering encoded data through this technique, such as decreased sensitivity to the stimuli over time, effects of environmental confounds, and user compliance to wearing body-worn devices [128]. Additionally, existing literature suggests that active exploration results in better memory retention [129]. Thus, we focused our efforts on using force-feedback and tactile devices that both require active user exploration to convey data. This formed the main focus of our thesis.

We designed and presented multiple audio-haptic renderings for line charts to BLV users through a user study. The work aimed to validate the effectiveness of multimodality for tactile and force-feedback approaches, and understand differences in the way information is perceived differently through them. While our results strongly suggest a benefit for all of the multimodal approaches, consistent with existing literature, users performed equally well at completing tasks among the three haptic conditions. In spite of this, all users stated a preference for the tactile case. To understand this seeming disparity, we qualitatively assessed their responses to open-ended questions about their experience with each condition, resulting in multiple themes that we believe can impact a user's preference: their background, the perceived affordances of a device, the cognitive resources required, the audio-haptic synchrony, initial impressions, and the steepness of the learning curve. These themes are linked, e.g., a steeper learning curve for a user likely correlates to a higher cognitive load. However, the observed benefits and drawbacks for each mode were based on limited examples that were presented to a user. It is possible, and likely, that additional training would mitigate any effects of novelty and consequently influence a user's experience. Many participants also noted technical usability issues with the force-feedback device; an upgraded version of the 2DIY was developed and acquired for the IMAGE web-accessibility project described in Section 1.1 shortly after this study was conducted, and further hardware improvements to the device over time are expected to result in improved force-feedback experiences. Lastly, the study was conducted with elderly participants. It is possible that younger individuals, who are generally more accustomed to technology and changes in technology, would adapt differently to the intricacies of a force-feedback device. In summary, there are several factors that may impact both user performance or preference.

Following this study, we also explored a further research direction with the RTD. While audio-haptic experiences can result in a more meaningful experience, they can be difficult to create. To address this issue, we developed and presented a tool that can be used to create and edit a desired tactile representation of an image in Chapter 5. We added a touch frame to the RTD and used it to provide interactive audio feedback, thus allowing for potentially immersive experiences. This included the use of touch gestures such as panning or tapping motions along the tactile image to deliver non-speech and speech audio renderings respectively. We hypothesize that this approach would improve the crossmodal congruence theme discussed in Chapter 4, and consequently allow users to obtain a higher level of trust in the data acquired through exploration. This method of generating renderings also has several advantages over tactile diagrams, such as a significantly reduced use of physical material, easy distribution of the audio-haptic media, and flexibility in quickly providing information to users.

As creating audio-haptic experiences for data of similar formats can be a time-consuming task, we developed a framework that allows designers to automatically generate a desired audio-haptic experience. With this approach, code can be written to automatically process the data, display the information on the RTD in an intended way, and provide various types of audio cues through gestures. We demonstrated this functionality with line charts in Chapter 5, but future work could explore the development of such audio-haptic experiences in many other STEM domains. For example, this could involve creating representations of atomic structures, structural formulas for organic compounds, maps, or electrical circuits. These potentially engaging and immersive experiences could make classrooms more inclusive to BLV students and allow educators to explore more ways of presenting data compared to traditional methods such as tactile diagrams.

In the scope of this thesis, we discussed haptic applications for BLV users. However, as mentioned in Section 1.1, haptic tools can also be used to offload the visual modality for sighted individuals. Earlier work by the author leveraged the 2DIY to allow users to quickly navigate to POIs in a densely populated scatter plot [130]. Here, the device was used in its guided mode to quickly navigate to a desired data point through text input, or take the user along a predefined set of salient points sequentially in a “story tour” mode of operation, analogous to the museum tour metaphor [65]. These points include statistical measures such as the interquartile values of the data series, which is information not directly retrievable through the visual modality. These techniques could therefore minimize the amount of time users spend inspecting a graph, and consequently, their required cognitive resources.

Our research findings suggest potential for using multimodal audio-haptic approaches to enhance content accessibility. The work also holds practical significance, as we used devices that are readily available or are expected to be on the market soon. Although there still remains significant challenges in deploying the technology [8], we believe continued efforts on iterating and testing designs of audio-haptic experiences that can later be deployed, such as in the IMAGE project [5], may provide improved experiences over existing commonly used assistive technologies, such as screen readers.

References

[1] T. A. Kern, C. Hatzfeld, and A. Abbasimoshaei, *Engineering Haptic Devices*. Springer, 2023.

[2] P. Alirezaee, A. Weill-Duflos, J. J. Schlesinger, and J. R. Cooperstock, "Exploring the effectiveness of haptic alarm displays for critical care environments," in *2020 IEEE Haptics Symposium (HAPTICS)*, pp. 948–954, IEEE, 2020.

[3] R. D. Gilson, E. S. Redden, and L. R. Elliott, "Remote tactile displays for future soldiers," tech. rep., University of Central Florida Orlando, 2007.

[4] J. Baldwin, *Affective Haptic Vest for Facilitating Interpersonal Communication*. PhD thesis, Rice University, 2022.

[5] J. Regimbal, J. R. Blum, and J. R. Cooperstock, "Image: a deployment framework for creating multimodal experiences of web graphics," in *Proceedings of the 19th International Web for All Conference*, pp. 1–5, 2022.

[6] World Health Organization, "Visual Impairment and Blindness." [Online; accessed 17-April-2023].

[7] H. Zhang, P. Nurius, Y. Sefidgar, M. Morris, S. Balasubramanian, J. Brown, A. K. Dey, K. Kuehn, E. Riskin, X. Xu, *et al.*, "How does covid-19 impact students with disabilities/health concerns?," *arXiv preprint arXiv:2005.05438*, 2020.

[8] V. Lévesque, "Blindness, technology and haptics," *Center for Intelligent Machines*, pp. 19–21, 2005.

[9] A. D. Flaxman, J. S. Wittenborn, T. Robalik, R. Gulia, R. B. Gerzoff, E. A. Lundeen, J. Saaddine, D. B. Rein, K. N. Baldonado, C. Davidson, *et al.*, "Prevalence of visual acuity loss or blindness in the us: a bayesian meta-analysis," *JAMA Ophthalmology*, vol. 139, no. 7, pp. 717–723, 2021.

[10] B. Leporini, M. C. Buzzi, and M. Buzzi, "Interacting with mobile devices via voiceover: Usability and accessibility issues," in *Proceedings of the 24th Australian Computer-Human Interaction Conference*, OzCHI '12, (New York, NY, USA), p. 339–348, Association for Computing Machinery, 2012.

[11] R. Kuber, A. Hastings, and M. Tretter, "Determining the accessibility of mobile screen readers for blind users," *UMBC Faculty Collection*, 2020.

[12] A. Fan, Y. Ma, M. Mancenido, and R. Maciejewski, "Annotating line charts for addressing deception," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2022.

[13] S. Paneels and J. C. Roberts, "Review of designs for haptic data visualization," *IEEE Transactions on Haptics*, vol. 3, no. 2, pp. 119–137, 2009.

[14] R. Sigrist, G. Rauter, R. Riener, and P. Wolf, "Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review," *Psychonomic Bulletin & Review*, vol. 20, pp. 21–53, 2013.

[15] W. Yu and S. Brewster, "Evaluation of multimodal graphs for blind people," *Universal access in the information society*, vol. 2, pp. 105–124, 2003.

[16] W. Yu, K. Cheung, and S. Brewster, "Automatic online haptic graph construction," in *Proceedings of EuroHaptics*, vol. 2002, pp. 128–133, 2002.

[17] J. Bornschein, D. Prescher, and G. Weber, "Collaborative creation of digital tactile graphics," in *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*, pp. 117–126, 2015.

[18] S. J. Bolanowski, G. A. Gescheider, R. T. Verrillo, and C. M. Checkosky, "Four channels mediate the mechanical aspects of touch," *The Journal of the Acoustical Society of America*, vol. 84, pp. 1680–1694, Nov. 1988.

[19] K. Johnson, "Neural basis of haptic perception," in *Steven's Handbook of Experimental Psychology: Sensation and Perception, Vol. 1, 3rd ed*, pp. 537–583, Hoboken, NJ, US: John Wiley & Sons Inc, 2002.

[20] J. M. Weisenberger and M. E. Percy, "The transmission of phoneme-level information by multichannel tactile speech perception aids," *Ear and Hearing*, vol. 16, no. 4, pp. 392–406, 1995.

[21] S. R. Williams and A. M. Okamura, "Body-mounted vibrotactile stimuli: simultaneous display of taps on the fingertips and forearm," *IEEE Transactions on Haptics*, vol. 14, no. 2, pp. 432–444, 2020.

[22] C. M. Reed, H. Z. Tan, Z. D. Perez, E. C. Wilson, F. M. Severgnini, J. Jung, J. S. Martinez, Y. Jiao, A. Israr, F. Lau, *et al.*, "A phonemic-based tactile display for speech communication," *IEEE Transactions on Haptics*, vol. 12, no. 1, pp. 2–17, 2018.

[23] S. Choi and K. J. Kuchenbecker, "Vibrotactile display: Perception, technology, and applications," *Proceedings of the IEEE*, vol. 101, no. 9, pp. 2093–2104, 2012.

[24] P. B. Shull and D. D. Damian, "Haptic wearables as sensory replacement, sensory augmentation and trainer–a review," *Journal of NeuroEngineering and Rehabilitation*, vol. 12, pp. 1–13, 2015.

[25] H. Pongrac, "Vibrotactile perception: examining the coding of vibrations and the just noticeable difference under various conditions," *Multimedia Systems*, vol. 13, no. 4, pp. 297–307, 2008.

[26] C. Hatzfeld and R. Werthschützky, "Vibrotactile force perception thresholds at the fingertip," in *Haptics: Generating and Perceiving Tangible Sensations: International Conference, EuroHaptics 2010, Amsterdam, July 8-10, 2010. Proceedings, Part I*, pp. 99–104, Springer, 2010.

[27] S. Kuroki, J. Watanabe, and S. Nishida, "Dissociation of vibrotactile frequency discrimination performances for supra-threshold and near-threshold vibrations," in *Haptics: Perception, Devices, Mobility, and Communication: International Conference, EuroHaptics 2012, Tampere, Finland, June 13-15, 2012 Proceedings, Part II*, pp. 79–84, Springer, 2012.

[28] M. Morioka, D. J. Whitehouse, and M. J. Griffin, "Vibrotactile thresholds at the fingertip, volar forearm, large toe, and heel," *Somatosensory & Motor Research*, vol. 25, no. 2, pp. 101–112, 2008.

[29] G. D. Goff, "Differential discrimination of frequency of cutaneous mechanical vibration," *Journal of Experimental Psychology*, vol. 74, pp. 294–299, June 1967.

[30] H. C. Stronks, D. J. Parker, and N. Barnes, "Vibrotactile spatial acuity and intensity discrimination on the lower back using coin motors," *IEEE Transactions on Haptics*, vol. 9, no. 4, pp. 446–454, 2016.

[31] K. Myles and J. T. Kalb, "Vibrotactile sensitivity of the head," tech. rep., Army Research Lab Aberdeen Ground MD Human Research and Engineering Directorate, 2009.

[32] J. B. Van Erp, "Presenting directions with a vibrotactile torso display," *Ergonomics*, vol. 48, no. 3, pp. 302–313, 2005.

[33] Ó. I. Jóhannesson, R. Hoffmann, V. V. Valgeirsdóttir, R. Unnþórsson, A. Moldoveanu, and Á. Kristjánsson, "Relative vibrotactile spatial acuity of the torso," *Experimental Brain Research*, vol. 235, pp. 3505–3515, 2017.

[34] E. Piateski and L. Jones, "Vibrotactile pattern recognition on the arm and torso," in *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics Conference*, pp. 90–95, IEEE, 2005.

[35] E. Wentink, A. Mulder, J. S. Rietman, and P. H. Veltink, "Vibrotactile stimulation of the upper leg: Effects of location, stimulation method and habituation," in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1668–1671, IEEE, 2011.

[36] J. M. Fontana, R. O'Brien, E. Laciar, L. S. Maglione, and L. Molisani, "Vibrotactile stimulation in the upper-arm for restoring individual finger sensations in hand prosthesis," *Journal of Medical and Biological Engineering*, vol. 38, pp. 782–789, 2018.

[37] I. Oakley, Y. Kim, J. Lee, and J. Ryu, "Determining the feasibility of forearm mounted vibrotactile displays," in *2006 14th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 27–34, IEEE, 2006.

[38] G. Luzhnica, S. Stein, E. Veas, V. Pammer, J. Williamson, and R. M. Smith, "Personalising vibrotactile displays through perceptual sensitivity adjustment," in *Proceedings of the 2017 ACM International Symposium on Wearable Computers*, pp. 66–73, 2017.

[39] A. V. Reinschluessel, S. C. Cebulla, M. Herrlich, T. Döring, and R. Malaka, "Vibro-band: supporting needle placement for physicians with vibrations," in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–6, 2018.

[40] G. Luzhnica and E. Veas, "Optimising encoding for vibrotactile skin reading," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2019.

[41] D. Tam, K. E. MacLean, J. McGrenere, and K. J. Kuchenbecker, "The design and field observation of a haptic notification system for timing awareness during oral presentations," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1689–1698, 2013.

[42] E. Pezent, A. Israr, M. Samad, S. Robinson, P. Agarwal, H. Benko, and N. Colonnese, "Tasbi: Multisensory squeeze and vibrotactile wrist haptics for augmented and virtual reality," in *2019 IEEE World Haptics Conference (WHC)*, pp. 1–6, IEEE, 2019.

[43] M. Aggravi, F. Pausé, P. R. Giordano, and C. Pacchierotti, "Design and evaluation of a wearable haptic device for skin stretch, pressure, and vibrotactile stimuli," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2166–2173, 2018.

[44] T. Moriyama and H. Kajimoto, "Wearable haptic device presenting sensations of fingertips to the forearm," *IEEE Transactions on Haptics*, vol. 15, no. 1, pp. 91–96, 2022.

[45] D. A. Mahns, N. Perkins, V. Sahai, L. Robinson, and M. Rowe, "Vibrotactile frequency discrimination in human hairy skin," *Journal of Neurophysiology*, vol. 95, no. 3, pp. 1442–1450, 2006.

[46] V. A. Shah, M. Casadio, R. A. Scheidt, and L. A. Mrotek, "Spatial and temporal influences on discrimination of vibrotactile stimuli on the arm," *Experimental Brain Research*, vol. 237, no. 8, pp. 2075–2086, 2019.

[47] K. V. Vierordt, "Die abhangigkeit der ausbildung des raumsinnes der haut von den beweglichkeit der kerpertheile," *Zeitschrift fur Biologie*, vol. 6, pp. 53–72, 1870.

[48] E. G. Boring, "Sensation and perception in the history of experimental psychology.," 1942.

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

[50] R. W. Cholewiak and A. A. Collins, "Vibrotactile localization on the arm: Effects of place, space, and age," *Perception & Psychophysics*, vol. 65, no. 7, pp. 1058–1077, 2003.

[51] Braille Authority of North America (BANA), "Guidelines and standards for tactile graphics." `http://www.brailleauthority.org/tg/`, 2010.

[52] M. Phutane, J. Wright, B. V. Castro, L. Shi, S. R. Stern, H. M. Lawson, and S. Azenkot, "Tactile materials in practice: Understanding the experiences of teachers of the visually impaired," *ACM Transactions on Accessible Computing (TACCESS)*, vol. 15, no. 3, pp. 1–34, 2022.

[53] C. Goncu, K. Marriott, and J. Hurst, "Usability of accessible bar charts," in *Diagrammatic Representation and Inference: 6th International Conference, Diagrams 2010, Portland, OR, USA, August 9-11, 2010. Proceedings*, pp. 167–181, Springer, 2010.

[54] A. Kalia, R. Hopkins, D. Jin, L. Yazzolino, S. Verma, L. Merabet, F. Phillips, and P. Sinha, “Perception of tactile graphics: Embossings versus cutouts,” *Multisensory Research*, vol. 27, no. 2, pp. 111–125, 2014.

[55] M. A. Heller, J. A. Calcaterra, L. L. Burson, and L. A. Tyler, “Tactual picture identification by blind and sighted people: Effects of providing categorical information,” *Perception & Psychophysics*, vol. 58, pp. 310–323, 1996.

[56] S. J. Lederman, R. L. Klatzky, C. Chataway, and C. D. Summers, “Visual mediation and the haptic recognition of two-dimensional pictures of common objects,” *Perception & Psychophysics*, vol. 47, no. 1, pp. 54–64, 1990.

[57] D. L. Mansur, M. M. Blattner, and K. I. Joy, “Sound graphs: A numerical data analysis method for the blind,” *Journal of Medical Systems*, vol. 9, pp. 163–174, 1985.

[58] S. Brewster, “Visualization tools for blind people using multiple modalities,” *Disability and Rehabilitation*, vol. 24, no. 11-12, pp. 613–621, 2002.

[59] L. Brown, S. Brewster, R. Ramloll, W. Yu, and B. Riedel, “Browsing modes for exploring sonified line graphs,” *Proceedings of BCS HCI 2002*, vol. 2, pp. 6–9, 2002.

[60] L. M. Brown, S. A. Brewster, S. Ramloll, R. Burton, and B. Riedel, “Design guidelines for audio presentation of graphs and tables,” International Conference on Auditory Display, 2003.

[61] K. Vines, C. Hughes, H. Holmes, V. Pearson, C. Kotecki, L. Alexander, C. Colwell, and K. Parks, “Sonification of depictions of numerical data,” *eSTEeM Final Report (Tech. Rep.)*, 2016.

[62] R. Sharma, V. I. Pavlovic, and T. S. Huang, “Toward multimodal human-computer interface,” *Proceedings of the IEEE*, vol. 86, no. 5, pp. 853–869, 1998.

[63] P. Blenkhorn and D. G. Evans, “Using speech and touch to enable blind people to access schematic diagrams,” *Journal of Network and Computer Applications*, vol. 21, no. 1, pp. 17–29, 1998.

[64] P. Roth, D. Richoz, L. Petrucci, and T. Pun, “An audio-haptic tool for non-visual image representation,” in *Proceedings of the Sixth International Symposium on Signal Processing and Its Applications (Cat. No. 01EX467)*, vol. 1, pp. 64–67, IEEE, 2001.

[65] J. C. Roberts, K. M. Franklin, and J. Cullinane, “Virtual haptic exploratory visualization of line graphs and charts,” in *Stereoscopic Displays and Virtual Reality Systems IX*, vol. 4660, pp. 401–410, SPIE, 2002.

[66] S. A. Panëels, J. C. Roberts, and P. J. Rodgers, "Haptic interaction techniques for exploring chart data," in *Haptic and Audio Interaction Design: 4th International Conference, HAID 2009 Dresden, Germany, September 10-11, 2009 Proceedings 4*, pp. 31–40, Springer, 2009.

[67] L. H. Goldish and H. E. Taylor, "The optacon: A valuable device for blind persons," *Journal of Visual Impairment & Blindness*, vol. 68, no. 2, pp. 49–56, 1974.

[68] L. Zeng and G. Weber, "Exploration of location-aware you-are-here maps on a pin-matrix display," *IEEE transactions on human-machine systems*, vol. 46, no. 1, pp. 88–100, 2015.

[69] B. Schmitz and T. Ertl, "Interactively displaying maps on a tactile graphics display," *SKALID 2012–Spatial Knowledge Acquisition with Limited Information Displays (2012)*, pp. 13–18, 2012.

[70] M. Ivanchev, F. Zinke, and U. Lucke, "Pre-journey visualization of travel routes for the blind on refreshable interactive tactile displays," in *Computers Helping People with Special Needs: 14th International Conference, ICCHP 2014, Paris, France, July 9-11, 2014, Proceedings, Part II 14*, pp. 81–88, Springer, 2014.

[71] L. Brayda, F. Leo, C. Baccelliere, E. Ferrari, and C. Vigini, "Updated tactile feedback with a pin array matrix helps blind people to reduce self-location errors," *Micromachines*, vol. 9, no. 7, p. 351, 2018.

[72] D. Prescher and G. Weber, "Comparing two approaches of tactile zooming on a large pin-matrix device," in *Human-Computer Interaction-INTERACT 2017: 16th IFIP TC 13 International Conference, Mumbai, India, September 25–29, 2017, Proceedings, Part I 16*, pp. 173–186, Springer, 2017.

[73] F. Leo, C. Tinti, S. Chiesa, R. Cavaglià, S. Schmidt, E. Cocchi, and L. Brayda, "Improving spatial working memory in blind and sighted youngsters using programmable tactile displays," *SAGE Open Medicine*, vol. 6, p. 2050312118820028, 2018.

[74] L. Holloway, S. Ananthanarayan, M. Butler, M. T. De Silva, K. Ellis, C. Goncu, K. Stephens, and K. Marriott, "Animations at your fingertips: Using a refreshable tactile display to convey motion graphics for people who are blind or have low vision," in *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, pp. 1–16, 2022.

[75] G. Melfi, L. Schölch, T. Schwarz, K. Müller, and R. Stiefelhagen, "Audio-tactile reader (atr): Interaction concepts for students with blindness to explore digital stem documents on a 2d haptic device," in *2022 IEEE Haptics Symposium (HAPTICS)*, pp. 1–6, IEEE, 2022.

[76] T. Schwarz, G. Melfi, S. Scheiffele, and R. Stiefelhagen, “Interface for automatic tactile display of data plots,” in *Computers Helping People with Special Needs: 18th International Conference, ICCHP-AAATE 2022, Lecco, Italy, July 11–15, 2022, Proceedings, Part I*, pp. 73–81, Springer, 2022.

[77] P. Albert, “Math class: An application for dynamic tactile graphics,” in *Computers Helping People with Special Needs: 10th International Conference, ICCHP 2006, Linz, Austria, July 11-13, 2006. Proceedings 10*, pp. 1118–1121, Springer, 2006.

[78] M. R. Mc Gee, P. Gray, and S. Brewster, “The effective combination of haptic and auditory textural information,” in *Haptic Human-Computer Interaction: First International Workshop Glasgow, UK, August 31—September 1, 2000 Proceedings*, pp. 118–126, Springer, 2001.

[79] B. Berglund, U. Berglund, and G. Ekman, “Temporal integration of vibrotactile stimulation,” *Perceptual and Motor Skills*, vol. 25, no. 2, pp. 549–560, 1967.

[80] T. N. Cornsweet, “The staircase-method in psychophysics,” *The American Journal of Psychology*, vol. 75, no. 3, pp. 485–491, 1962.

[81] W. H. Ehrenstein and A. Ehrenstein, “Psychophysical methods,” in *Modern Techniques in Neuroscience Research*, pp. 1211–1241, Springer, 1999.

[82] M. Tommerdahl, K. Hester, E. Felix, M. Hollins, O. Favorov, P. Quibrera, and B. Whitsel, “Human vibrotactile frequency discriminative capacity after adaptation to 25 Hz or 200 Hz stimulation,” *Brain Research*, vol. 1057, no. 1-2, pp. 1–9, 2005.

[83] R. D. Hamer, R. T. Verrillo, and J. J. Zwislocki, “Vibrotacile masking of Pacinian and non-Pacinian channels,” *The Journal of the Acoustical Society of America*, vol. 73, no. 4, pp. 1293–1303, 1983.

[84] W. Taube Navaraj, C. García Núñez, D. Shakthivel, V. Vinciguerra, F. Labeau, D. H. Gregory, and R. Dahiya, “Nanowire fet based neural element for robotic tactile sensing skin,” *Frontiers in Neuroscience*, vol. 11, p. 501, 2017.

[85] R. T. Verrillo, A. J. Fraioli, and R. L. Smith, “Sensation magnitude of vibrotactile stimuli,” *Perception & Psychophysics*, vol. 6, pp. 366–372, Nov. 1969.

[86] R. W. Cholewiak, “The perception of tactile distance: Influences of body site, space, and time,” *Perception*, vol. 28, no. 7, pp. 851–875, 1999.

[87] A. K. Goble and M. Hollins, “Vibrotactile adaptation enhances frequency discrimination,” *The Journal of the Acoustical Society of America*, vol. 96, no. 2, pp. 771–780, 1994.

[88] H. C. Stronks, J. Walker, D. J. Parker, and N. Barnes, "Training improves vibrotactile spatial acuity and intensity discrimination on the lower back using coin motors," *Artificial Organs*, vol. 41, no. 11, pp. 1059–1070, 2017.

[89] J. Heer, M. Bostock, and V. Ogievetsky, "A tour through the visualization zoo," *Communications of the ACM*, vol. 53, no. 6, pp. 59–67, 2010.

[90] K. T. Zebehazy and A. P. Wilton, "Quality, importance, and instruction: The perspectives of teachers of students with visual impairments on graphics use by students," *Journal of Visual Impairment & Blindness*, vol. 108, no. 1, pp. 5–16, 2014.

[91] B. Leporini, M. C. Buzzi, and M. Buzzi, "Interacting with mobile devices via voiceover: usability and accessibility issues," in *Proceedings of the 24th Australian computer-human interaction conference*, pp. 339–348, 2012.

[92] E. C. Bell and A. M. Silverman, "Access to math and science content for youth who are blind or visually impaired," 2019.

[93] F. K. Aldrich and L. Sheppard, "Tactile graphics in school education: perspectives from pupils," *British Journal of Visual Impairment*, vol. 19, no. 2, pp. 69–73, 2001.

[94] W. W. Gaver, "What in the world do we hear?: An ecological approach to auditory event perception," *Ecological Psychology*, vol. 5, no. 1, pp. 1–29, 1993.

[95] T. Hermann, A. Hunt, J. G. Neuhoff, *et al.*, *The sonification handbook*, vol. 1. Logos Verlag Berlin, 2011.

[96] A. Dufour, O. Després, and V. Candas, "Enhanced sensitivity to echo cues in blind subjects," *Experimental Brain Research*, vol. 165, no. 4, pp. 515–519, 2005.

[97] F. Gougoux, F. Lepore, M. Lassonde, P. Voss, R. J. Zatorre, and P. Belin, "Pitch discrimination in the early blind," *Nature*, vol. 430, no. 6997, pp. 309–309, 2004.

[98] L. M. Holloway, C. Goncu, A. Ilsar, M. Butler, and K. Marriott, "Infosonics: Accessible infographics for people who are blind using sonification and voice," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2022.

[99] X. Hu, A. Song, H. Zeng, and D. Chen, "Intuitive environmental perception assistance for blind amputees using spatial audio rendering," *IEEE Transactions on Medical Robotics and Bionics*, vol. 4, no. 1, pp. 274–284, 2022.

[100] D. K. McGookin and S. A. Brewster, "Graph builder: Constructing non-visual visualizations," in *People and Computers XX—Engage: Proceedings of HCI 2006*, pp. 263–278, Springer, 2007.

[101] G. Petit, A. Dufresne, V. Levesque, V. Hayward, and N. Trudeau, "Refreshable tactile graphics applied to schoolbook illustrations for students with visual impairment," in *Proceedings of the 10th international ACM SIGACCESS Conference on Computers and Accessibility*, pp. 89–96, 2008.

[102] S. O'Modhrain, N. A. Giudice, J. A. Gardner, and G. E. Legge, "Designing media for visually-impaired users of refreshable touch displays: Possibilities and pitfalls," *IEEE Transactions on Haptics*, vol. 8, no. 3, pp. 248–257, 2015.

[103] H. V. Rao and S. O'Modhrain, "2Across: A comparison of audio-tactile and screen-reader based representations of a crossword puzzle," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2020.

[104] G. Baud-Bovy and N. Balzarotti, "Using force-feedback devices in educational settings: a short review," in *Proceedings of the 1st ACM SIGCHI international Workshop on Multimodal Interaction for Education*, pp. 14–21, 2017.

[105] A. Weill-Duflos, N. Ong, F. Desourdy, B. Delbos, S. Ding, and C. Gallacher, "Haply 2diy: an accessible haptic platform suitable for remote learning," in *Proceedings of the 2021 International Conference on Multimodal Interaction*, pp. 839–840, 2021.

[106] S. J. Kerr and C. Siau, "Multi-modal user interfaces: Effects of redundant and complementary approaches in high stimuli environments," *EDITORIAL BOARD MEMBERS*, p. 16, 2016.

[107] J. Hartcher-O'Brien, *Multisensory integration of redundant and complementary cues*. PhD thesis, Oxford University, UK, 2012.

[108] K. Salisbury, F. Conti, and F. Barbagli, "Haptic rendering: introductory concepts," *IEEE Computer Graphics and Applications*, vol. 24, no. 2, pp. 24–32, 2004.

[109] V. Braun and V. Clarke, *Thematic Analysis*. American Psychological Association, 2012.

[110] R. Rampin and V. Rampin, "Taguette: open-source qualitative data analysis," *Journal of Open Source Software*, vol. 6, no. 68, p. 3522, 2021.

[111] N. Kim, B. Koo, J. Yoon, and K. Cho, "Understanding the formation of user's first impression on an interface design from a neurophysiological perspective-eeg pilot study," in *Proceedings of HCI Korea*, pp. 139–145, 2016.

[112] A. Cafaro, H. H. Vilhjálmsson, and T. Bickmore, "First impressions in human–agent virtual encounters," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 23, no. 4, pp. 1–40, 2016.

[113] Y. Cheong, R. L. Shehab, and C. Ling, "Effects of age and psychomotor ability on kinematics of mouse-mediated aiming movement," *Ergonomics*, vol. 56, no. 6, pp. 1006–1020, 2013.

[114] I. J. Deary, J. Corley, A. J. Gow, S. E. Harris, L. M. Houlihan, R. E. Marioni, L. Penke, S. B. Rafnsson, and J. M. Starr, "Age-associated cognitive decline," *British Medical Bulletin*, vol. 92, no. 1, pp. 135–152, 2009.

[115] G. E. Legge, C. Madison, B. N. Vaughn, A. M. Y. Cheong, and J. C. Miller, "Retention of high tactile acuity throughout the life span in blindness," *Perception & Psychophysics*, vol. 70, pp. 1471–1488, Nov. 2008.

[116] J. C. Steinberg, H. C. Montgomery, and M. B. Gardner, "Results of the world's fair hearing tests," *The Bell System Technical Journal*, vol. 19, no. 4, pp. 533–562, 1940.

[117] N. F. of the Blind, "Blindness statistics." `https://nfb.org/resources/blindness-statistics`, 2017. Accessed on August 25, 2023.

[118] J. Raskin, "Intuitive equals familiar," *Communications of the ACM*, vol. 37, no. 9, pp. 17–19, 1994.

[119] D. A. Norman, "The 'problem'with automation: inappropriate feedback and interaction, not 'over-automation'," *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, vol. 327, no. 1241, pp. 585–593, 1990.

[120] L. Renier, R. Bruyer, and A. G. De Volder, "Vertical-horizontal illusion present for sighted but not early blind humans using auditory substitution of vision," *Perception & Psychophysics*, vol. 68, pp. 535–542, 2006.

[121] J. Fitzpatrick and F. Neff, "Perceptually congruent sonification of auditory line charts," *Journal on Multimodal User Interfaces*, pp. 1–16, 2023.

[122] M. Eid, M. Orozco, and A. El Saddik, "A guided tour in haptic audio visual environments and applications," *International Journal of Advanced Media and Communication*, vol. 1, no. 3, pp. 265–297, 2007.

[123] B. Kollmeier, T. Brand, and B. Meyer, "Perception of speech and sound," *Springer Handbook of Speech Processing*, p. 65, 2008.

[124] F. Kalantari, D. Gueorguiev, E. Lank, N. Bremard, and L. Grisoni, "Exploring fingers' limitation of texture density perception on ultrasonic haptic displays," in *Haptics: Science, Technology, and Applications: 11th International Conference, EuroHaptics 2018, Pisa, Italy, June 13-16, 2018, Proceedings, Part I 11*, pp. 354–365, Springer, 2018.

[125] C. Zhang, D. R. Sahoo, J. Pearson, S. Robinson, M. D. Holton, P. Hopkins, and M. Jones, "Active PinScreen: Exploring spatio-temporal tactile feedback for multi-finger interaction," in *22nd International Conference on Human-Computer Interaction with Mobile Devices and Services*, pp. 1–11, 2020.

[126] W. Han and H.-J. Schulz, "Exploring vibrotactile cues for interactive guidance in data visualization," in *Proceedings of the 13th International Symposium on Visual Information Communication and Interaction*, pp. 1–10, 2020.

[127] F. G. Bastidas Cuya, R. L. Martins Guarese, C. G. Johansson, M. Giambastiani, Y. Iquiapaza, V. A. de Jesus Oliveira, L. Nedel, and A. Maciel, "Vibrotactile data physicalization: Exploratory insights for haptization of low-resolution images," in *Symposium on Virtual and Augmented Reality*, pp. 84–91, 2021.

[128] J. R. Blum, P. E. Fortin, F. Al Taha, P. Alirezaee, M. Demers, A. Weill-Duflos, and J. R. Cooperstock, "Getting your hands dirty outside the lab: A practical primer for conducting wearable vibrotactile haptics research," *IEEE Transactions on Haptics*, vol. 12, no. 3, pp. 232–246, 2019.

[129] R. T. Stone, "Actively guided practice enhances kinesthetic memory development," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 57, pp. 640–644, SAGE Publications Sage CA: Los Angeles, CA, 2013.

[130] R. Isran, K. Sepehri, K. Theivendran, and A. Anwar, "Towards more effective data visualization methods using haptics," in *2021 IEEE World Haptics Conference (WHC)*, pp. 590–590, IEEE, 2021.