Exploring the Effectiveness of Haptic Alarm Displays for Critical Care Environments

Parisa Alirezaee

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

January 2020

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

2020/01/29

ABSTRACT

Noise in critical care units, in particular, from patient monitor alarms, is harmful for the clinicians and patients alike. This has motivated research aimed at shifting the delivery of physiological vital sign information and announcement of alarm events from visual and auditory devices to haptic transducers. We compare performance in perceiving and identifying the specific type and level of a vital sign that has entered a high or low state, i.e., an alarm event, using several designs of a vibrotactile display, against that of the traditional audio alarm in conjunction with a graphical patient monitor. Multiple vibrotactile rendering conditions were presented, in addition to the baseline audio alarm with graphical patient monitor display. A distractor activity was used to simulate competing task demands in the clinical environment. Responses were assessed with respect to response time and accuracy. With sufficient anatomical separation of the actuators, certain vibrotactile information rendering strategies demonstrated performance that was not significantly different from that of the baseline condition, both in response time and accuracy. We conclude that vibrotactile delivery of vital sign information can support alarm-state vital sign identification competitive with graphical and auditory alarm display conditions, without significant impact on performance on a parallel attention-demanding activity. Our findings suggest the possibility of improving high-impact healthcare environments by replacing disturbing audio alarms with vibrotactile information delivery to clinicians.

RÉSUMÉ SCIENTIFIQUE

Le bruit dans les unités de soins intensifs, en particulier celui des alarmes du moniteur du patient, est nocif pour les cliniciens et les patients. Cela a motivé la recherche visant à remplacer des dispositifs visuels et auditifs par des transducteurs haptiques pour la transmission de l'informations sur les signes vitaux physiologiques et l'annonce d'événements d'alarme. Nous comparons les performances en termes de perception et d'identification du type et du niveau spécifiques d'un signe vital qui est entré dans un état haut ou bas, c'est-à-dire un événement d'alarme, utilisant plusieurs conceptions d'un affichage vibrotactile, avec celles de l'alarme sonore traditionnelle en conjonction avec un moniteur graphique du patient. Plusieurs conditions de transmission vibrotactile ont été présentées, en plus de l'alarme audio de base avec l'affichage graphique du moniteur du patient. Une activité de distraction a été utilisée pour simuler des tâches concurrentes dans l'environnement clinique. Les réponses ont été évaluées en fonction du temps de réponse et de la précision. Avec une séparation anatomique suffisante des actionneurs, certaines stratégies de transmission de l'informations vibrotactile ont montré des performances qui n'étaient pas significativement différentes de celles de la condition de base, à la fois en terme de temps de réponse et de précision. Nous concluons que la transmission vibrotactile de l'information sur des signes vitaux peut avoir une identification de signes vitaux en état d'alarme concurrencielle aux conditions d'affichage d'alarmes graphiques et sonores, sans impact significatif sur les performances pour une activité parallèle exigeant de l'attention. Nos résultats suggèrent la possibilité d'améliorer les environnements de soins de santé à fort impact en remplaçant les alarmes sonores perturbantes par la transmission d'informations vibrotactiles aux cliniciens.

ACKNOWLEDGEMENTS

I would like to thank my colleagues at the Shared Reality Lab, Pascal Fortin, Antoine Weill Duflos, Jeff Blum, Taeyong Kim, Roger Girgis, Preeti Vyes, for their helpful suggestions, their kind support and their friendship that made this journey interactive, instructive, and fun. I want to also thank my supervisor Prof. Jeremy Cooperstock for his help in preparing the conference and journal paper submissions.

I also acknowledge my dear friends for their emotional support throughout all the stages of studying at McGill. Without them to vent out the graduate school frustrations by spending quality time traveling, exploring, reading poems, jogging, talking, laughing and crying this journey would have been left unfinished.

I would like to thank Prof. Karyn Moffatt, the external examiner of my thesis, for their encouraging comments and valuable feedback that substantially enhanced the quality of my thesis. I also want to thank Mary Towers, my tutor at the McGill Writing Centre, who helped me through writing up this thesis manuscript with patience and a great knowledge on academic writing.

And at the end, I would like to thank my family for their support. My dad for being a role-model for hardworking and my mom for her unlimited kindness.

Contents

List of Figures

List of Tables

List of Acronyms

WHO	Worlds Health Organization
MAUDE	Manufacturer and User Facility Device Experience
OR	Operating Room
ICU	Intensive Care Unit
PICU	Pediatric Intensive Care Unit
BP	Blood Pressure
HR	Heart Rate
O_2	Blood Oxygen Saturation
TV	Tidal lung Volumes
etCO_2	end-tidal partial pressures of Carbon Dioxide
MAP	Mean Arterial blood Pressure
PoIE	Principle of Inverse Effectiveness
PEST	Parameter Estimation by Sequential Testing
TOST	Two-One-Sided T-test
IES	Inverse Efficiency Scores
CIRMMT	Centre for Interdisciplinary Research in Music Media and Technology
CIM	Centre for Intelligent Machines
SRL	Shared-Reality Lab
ICAD	International Conference on Auditory Displays
IRB	Institutional Review Board
dBA	A-weighted Decibels
SPL	Sound Pressure Level
HVAC	Heating, Ventilation, and Air Conditioning
F	Female
M	Male
CAD	Canadian Dollar

Unnecessary noise is the most cruel abuse of care which can be inflicted on either the sick or the well

—Florence Nightingale, *Notes on Nursing*

Chapter 1

Introduction

Hospital operating rooms (ORs) and intensive care units (ICUs) are extremely noisy environments. Whereas the World Health Organization (WHO) recommends sound levels do not exceed 40 dBA[1] in areas where patients are being treated or are recovering [1], one study found that the average sound level of patient rooms in three Veterans Affairs facilities measured 51 dBA [2]. Similarly, a study of five UK adult ICUs reported sound level peaks above 85 dBA during daytime periods of approximately 25 minutes every hour [3]. These studies indicate the importance of addressing high noise levels in hospital ORs and ICUs, which may be accomplished by focusing on the primary sources of such noise. These sources include variable-pitch pulse oximeter, speech, and alarms from clinical monitors. While patient monitoring alarms help clinicians maintain vigilance while providing care and are thus vital to patients' safety, they also contribute to higher levels of ambient noise.

High noise levels in critical care units may impede the quality of care by negatively affecting both caregivers and patients. One way these negative consequences may influence clinicians is by deteriorating their ability to receive important messages from relevant audio alarms. In what follows we discuss how a dangerously high number of monitor alarms and false alarms contribute to the noise in hospitals and cause a variety of challenges for clinicians working in ORs and ICUs. These challenges include alarm fatigue, auditory masking of alarms, and negative impact of broadcast auditory messages on clinicians, which will be discussed in greater detail in the following section. In addition to posing challenges

[1] A-weighted Descibels, or dBA, is an expression for the relative loudness of sounds in air as perceived by the human ear.

to hospital staff, noisy critical care units may also negatively affect patients, resulting in sleep deprivation and long-term cognitive impairment. In an effort to create an environment more conducive to patient care by reducing hospitals' high noise level, this thesis proposes vibration stimuli to deliver patients' vital information to medical staff in a wearable manner.

1.1 Preface

This chapter was written by myself while my supervisor, Jeremy Cooperstock, helped me with comments on the content flow as well as editorial suggestions.

1.2 The Problem

In recent years, the use of digital patient monitors, each integrating a loud auditory display, has become more prevalent, consequently exacerbating hospital noise levels.[2] Back in 1986, an average of 2.9 alarms per hour were recorded in a general-purpose ICU [5]. Although understandably sampled in different hospital environments, by 2004, this number increased to 4.7 alarms per hour in the emergency department of Sunnybrook Hospital, Canada [6], and by 2010, 6 alarms per hour were recorded in the medical ICU of the University of Regensburg Hospital, Germany [7]. This increase seems problematic, especially given the limited ability of clinicians to distinguish what each alarm signals. In a 214-bed Canadian teaching hospital, OR staff reported that out of an average total of 26 alarms per surgery, they were only able to identify roughly 10 to 15 alarms [8]. The same study reported ICU staff were able to identify an average of 9 to 14 alarms out of a total average of 23 alarms in the ICU. As these results demonstrate, clinicians experience difficulty in identifying alarms in critical care areas [9]. Although this difficulty is troubling, physiological alarms do not always indicate a condition warranting clinical intervention. These false and nuisance alarms, defined in the following, may contribute to the critical care areas' cacophony.

False alarms are signals activated due to electromechanical noise in the physiological sensor, occurring despite the absence of an aberration in physiology, i.e., not a true health

[2] For example, when exploring the noise measurement studies from the past 50 years, it was found that Sound Pressure Level (SPL) has increased by an average of 16 dB(A) between 1965 and 2005. As for the sources of this ambient noise, they reported HVAC (heating, ventilation, and air conditioning) for low frequency sounds and alarms and mobile medical equipment for high frequency noise [4].

event. Nuisance alarms are those that correctly activate when sensors detect a deviation in a physiological parameter, but the deviation is not related to a true health event, and therefore, an alarm is not warranted. Although these alarms do not reflect a clinically relevant condition, their occurrence in hospitals is significant. Lawless conducted a survey at a university affiliated children's hospital to categorize ICU alarms over a seven-day period, finding 1481 false alarms (68%) out of 2176 total alarms [10]. Using a different method for recording alarms, Siebig et al. employed a novel video surveillance method to extract alarm data. The results of their observation, similarly, showed that out of 5934 alarms that were annotated, only 14.9% of alarms were reported as clinically relevant [7]. Talley et al. found a false alarm rate of 85% to 99% in cardiopulmonary monitors at a Pediatric Intensive Care Unit (PICU) at Children's National Medical Center in Washington DC [11]. A similar finding was recorded at the PICU of the Children's Hospital of Philadelphia showing 87.1% of "non-actionable" alarms out of 5070 total alarms [12]. As is evident from these findings, the occurrence of false positive alarms in critical care areas is dangerously high. When exposed to a high false alarm rate, users' response rate, i.e., the percentage of alarms to which the user responds, decreases [13]. Consequently, alarms that were ignored keep sounding for longer and contribute to the auditory overload of hospital settings.

We now elaborate on the challenges facing medical staff and patients, stemming from the noisy hospital environment. One of these challenges affecting clinicians is "alarm fatigue", also known as "crying wolf" phenomenon [10]. This phenomenon occurs when auditory delivery of patients' physiological information overloads clinicians' audio capacity [10]. As a result, clinicians become desensitized to the sound of alarms [14] [15]. Other behaviors resulting from alarm fatigue may be ignoring the alarm, turning off the alarm, turning down the volume [16], or ignoring important incidents associated with true alarms [10]. All of these behaviours eventually result in harm to patients, as documented in the FDA Manufacturer and User Facility Device Experience (MAUDE) database, a collection of device-related safety issues [15]. Another challenge clinicians face in noisy hospitals is that the sound of some alarms are masked by the noises of other OR or ICU equipment; in particular, when multiple alarms annunciate at the same time. Therefore, the alarms become nearly indiscriminable and missed [17] [18] [19]. Finally, the current method of auditory alarm delivery may not practically include personalized messages. Since the alarms are broadcast, rather than delivered exclusively to the individual(s) responsible for each

specific alarm, they might increase the noise as well as associated stress levels. Dr Joseph Schlesinger, our anesthesiologist colleague, demonstrates this matter by narrating one of his experiences during a trauma surgery:

> "...the ventilator on the anesthesia machine begins notifying for high peak airway pressures, likely due to pulmonary contusion from blunt trauma. Turning my back on the patient to manipulate the ventilator settings, the rapid transfusion system starts notifying for air entrained into the system, as the anesthesia resident failed to fill the canister rapidly enough with blood products balanced with the high infusion rate. Feeling overwhelmed by the multitude of alarms, the electrocautery machine starts notifying the anesthesia team. The surgeon turns to look at it, confused why there is an alarm. Indeed, it is out of my reach, and not a machine with which anesthesia ever interacts – I don't know how to fix the problem or silence the alarm."

As is apparent from this scenario, the surgeon who was undergoing an extremely attention-demanding and high-risk task, was distracted by an alarm for which the anesthesiologist was responsible. This example demonstrated the negative impacts of broadcast auditory alarms. While the aforementioned issues affect clinicians in their performance, high levels of noise also adversely affects patients' recovery.

For patients who have to spend time in the ICU, the noisy environment creates various hazards. These hazards may affect patients not only by caregivers' alarm mismanagement, but also by imposing direct harm to their health and cognition. In the span of four years, more than 500 patient deaths have been associated with faulty alarm management [20] [16]. Similarly, between 2016 and 2018 alarm hazards were considered among the top four of the list of Top 10 Health Technology Hazards reported by ECRI Institute.[3] While this evidence demonstrates the effects of alarm mismanagement on patients' health, excessive noise levels, especially the sound of alarms, also create stress, fear, and traumatic experiences on patients. This last effect, known as ICU delirium [21], was recorded in 50% - 80% of patients [22]. Neuropsychological studies have shown that the duration of delirium directly impacts the duration of future cognitive impairments in survivors of critical illness [23] [24]. Additionally, constant exposure to high levels of noise may also be problematic.

[3] Top 10 Health Technology Hazards for 2015, 2016, 2017, 2018

In children aged 7-19, it was shown that constant exposure to high levels of noise causes long-term cognitive impairments and learning difficulties [25]. For patients healing in the ICU, sufficient sleep is an important factor in their recovery. However, previous studies report that noise from patient monitors is a significant disrupter of sleep [26] [25], responsible for 17% of the awakening at the University of Pennsylvania Medical Center [27]. Sleep loss destabilizes metabolism and slows down patient's recovery, which results in patients with cardiovascular issues to experiencing unstable heart rate and blood pressure [28].

To sum up, "Hospitals are horrible places to get better", according to Joel Beckerman, sound designer, composer and a pioneer in musical approaches for products, brands and environments.[4] Now that we have reviewed the causes and effects of the multifaceted issue of noisy critical care units, the remainder of this chapter first focuses on our proposed solution and anticipated outcome, and concludes by indicating the contributions of this research in improving hospital patient monitoring technology.

1.3 Objectives

The hospital critical care environment imposes a heavy burden on clinicians' auditory channels. Our objective is to reduce the ambient noise created by the sound of patient monitoring alarms while supporting clinicians' awareness of patients' vital physiological parameters. In order to do so, this research takes a novel approach by proposing the use of haptic feedback for delivering this information. Haptic feedback refers to feedback generated by stimulating individuals' sense of touch, e.g. vibrations, heat, or pressure. We specifically explored the effectiveness of vibration signals as an information delivery mechanism. If the proposed vibration-only display can be as effective as the current audio-visual information delivery system, it can result in obviating the need for the loud and distracting alarm signals. Even though this proposed alarm system may not entirely improve the alarm architecture and abolish the flaws described in the previous section (e.g. false alarms, nuisance alarms, etc) but the silent nature of vibrotactile signals as well as other benefits that are described in the next chapter, will noticeably reduce harms resulting from noisy hospital environment. This offers the potential to achieve a quieter and safer critical care environment, for clinicians and patients alike.

[4] `https://99percentinvisible.org/episode/sound-and-health-hospitals/`

1.4 Author's Contribution

In the following, we summarize our research contribution regarding two aspects of utilizing haptic wearable tools for clinical information delivery.

1. The possibility of improvement in the perception of the alarm sound when delivered in a multimodal manner in conjunction with a sub-threshold vibration signal.

2. The capability of haptic monitoring of patients' vital signs in improving clinical performance.

Our investigation of the first aspect mentioned above failed to demonstrate the anticipated results. However, in an effort to explore the second aspect, we found significance in the effectiveness of a two dimensional physiological alarm delivery (type of the vital sign and its level) through vibrotactile signals. In further detail, certain vibrotactile information rendering strategies demonstrated performance that was not significantly different from that of the baseline auditory graphical condition both in response time and accuracy. This strategy also showed no significant impact on performance on a parallel attention-demanding activity. These findings support the possibility of replacing disturbing audio alarms with haptic feedback for the sake of improving high-impact healthcare environments.

Chapter 2

Literature Review

As discussed in detail in the previous chapter, there are numerous consequences associated with high noise levels in hospitals. In this chapter we review scholarship that explored the effectiveness of vibrotactile alarm-related information delivery mechanisms in a clinical scenario. First we explore previous work on the benefits of multi-modal presentation of information on individuals' perception. Then we investigate prior research efforts that examined the effectiveness of vibrotactile alarm/monitoring systems in improving clinicians' awareness of patients' physiological information. We conclude this chapter by highlighting the main research gaps concerning vibrotactile alarm systems.

2.1 Preface

This chapter was written by myself, and my supervisor, Jeremy Cooperstock, contributed by instructive suggestions and edits.

2.2 Vibrotactile and Wearable Technology

Multi-modality refers to the presentation of two or more stimuli simultaneously. There is an extensive literature demonstrating that stimuli are better perceived when presented in a multi-modal manner. When stimuli from two modalities are simultaneously presented, the Principle of Inverse Effectiveness (PoIE) [29] suggests that an enhanced neural response can be achieved. In addition, multi-modal feedback "humanizes interactions with computers" while maximizing human cognitive and physical abilities relating to working memory, attention and decision making [30]. For example, it was shown that the co-occurrence of a sound increases the accuracy and enhances the sensitivity for detection of near-threshold visual stimuli [30]. Similarly, when performing a task consisting of a series of "drag-and-drops", the bimodal effect of haptic and visual feedback resulted in decreased self-perceived mental demands as well as faster performance compared to those of unimodal presentations of the same feedback [31]. The encouraging effects of multimodal feedback led us to consider exploring this method in the design of hospitals'alarm system.

In designing our multimodal alarm system, we considered the vibration signal as one of the information delivery mechanisms. The literature elucidates various advantages for utilizing this modality in the complex environment of hospital critical care units. In the context of warnings, different stimuli have benefits and drawbacks. Visual displays benefit from a greater information capacity by employing various shapes, colors as well as the location on the display. However, when visual displays are employed to deliver warnings, the user must divert their attention away from the task at hand to a specific location for receiving necessary information. Conversely, auditory and haptic warnings direct attention to themselves; given that the sense of audition and touch can be perceived regardless of ones attention orientation. In the case of haptic feedback, this capability shows even more success than other signals, since vibrotactile cues exploit a comparatively underutilized sensory

channel[1] relative to vision and audition [32]. Haptic warnings must remain relatively simple given the "limited perceptual resolution and processing bandwidth of the tactile channel" as suggested by design experts [33]. Therefore they should not require heavy interpretation. In addition to this benefit, researchers have demonstrated that vibrations are relatively competent in communicating time-sensitive information in complex environments. In a panel discussion, Baldwin et al. explored the effectiveness of different modalities in high consequence environments where time-sensitive information is delivered to the user, [33]. They concluded that when tactile signals are effectively mapped to the desired meaning, haptic feedback can be a promising information display in complex environments, provided that concurrent tasks do not overload the user's cognitive resources. Additionally, tactile signals offer the capability of selective delivery of the alarm message to the clinician(s) for whom it is relevant [34]. This personalized information delivery can reduce stress levels of the other staff, patients, and visitors that are present in the room. We were equally interested in exploring the effectiveness of vibration signal when such a signal is delivered in a sub-threshold manner. In considering how to do so, we will review the relevant literature in the following.

Visell et al. showed that the addition of sub-threshold stimuli affected the participants' perception of compliance [35]. This effect provides encouraging evidence that even at a sub-threshold level, such stimuli may be sufficient to affect alarm perception. Therefore, we theorize that the co-occurrence of a sub-threshold haptic signal and an audio alarm would allow participants to perceive sounds below their auditory threshold of perception. In the OR and ICU setting, this would allow a reduction of the alarm intensity while maintaining the effectiveness in cueing the caregiver. The next section investigates the scholarship that employed a vibrotactile alarm/monitoring system in clinical scenarios with the purpose of improving clinical performance.

2.3 Previous Vibrotactile Solutions in Hospitals

This section explores prior work that proposed haptic alarm/monitoring systems as a potential improvement to the noise in critical care areas. We organize this review according to methodology of where researchers delivered vibration signals on the body. Locations

[1] except for the hands of the medical staff while grasping tools and instruments

that vary from arm and wrist to the waist as well as the torso were chosen in previous explorations. On the other hand, hospitals' sterilization constraints require specific considerations on placement of such hardware on medical staffs body. Therefore, we put extra attention on this aspect of methodology.

Ng et al. implemented a haptic arm band with two actuators to convey alarms related to three levels of change in patient heart rate, both increasing and decreasing. The alarm state was mapped to spatial location, i.e., the actuator closer to the wrist pulsed once (for 600 ms) to indicate a decreasing heart rate alarm, while the level of decrease was delivered to the actuator near the elbow by the number of pulses and the interval between them [36]. Results indicated statistically significant superior performance in alarm identification using the haptic arm band than with auditory alarms, but no difference in response time or learnability. Cobus et al. designed an armband with three vibration motors and compared perception accuracy, learnability, distinguishability, and perceived urgency for six different vibration patterns through interviews with nurses, but did not compare performance against that of the traditional auditory signal [37].

Since the arm is often not a suitable location for attaching such haptic hardware in a clinical context, other researchers have instead delivered patient information to other body locations. As an example of such an approach, Ford designed a haptic belt with four tactors, two worn in the front, and two at the posterior of the participant, to communicate changes in two physiological parameters: peak airway pressure indicated by the tactors on the right side of the body, and minute volume ventilation indicated by those on the left [38]. In either case, the anterior tactor of the pair vibrated to indicate increase, and the posterior tactor to indicate a decrease of the corresponding parameter. Results on a simulated anaphylaxis incident, with anesthesiologists and anesthesiology residents, indicated that participants were significantly faster in treating the case when receiving information from vibrotactile stimuli than the control condition of conventional auditory alarms ($p < 0.05$). However, the study found no significant differences in situational awareness of the participants between the two conditions.

Ferris and Sarter designed a garment holding sets of C2 tactors (`https://www.eaiinfo.com/tactor-info/`) at spatially relevant body locations: near the lungs to represent tidal lung volumes (TV), on the spine for end-tidal partial pressures of carbon dioxide ($etCO_2$), and on the upper left arm for mean arterial blood pressure (MAP). Patterns of actuation

followed what the authors termed a "metaphorically accurate natural mapping for each of the parameters", with increasing parameters associated with progression of actuation toward tactors placed higher on the clinician's body. The authors also tested a hybrid mapping in which, in addition to the spatial progression of actuation, as described above, the severity of the change in the parameter was associated with the salience of vibration [39]. Tactor placement was determined by the associated anatomical locations of these vitals. Clinical performance was then assessed on a simulated intubation task, with the use of the vibrotactile garment compared to the baseline auditory-visual display. The addition of tactile signals was found to result in a significant reduction in detection time of abnormal conditions and correction time for normalization of the physiological signs. However, no analysis was conducted as to which modality the participants relied upon to monitor the patient status. Moreover, the vest itself, with all of its constituent tactors, represents a fairly significant burden, both in terms of apparel and hardware requirements.

Although the prior literature described above investigated the benefits of adding vibrotactile display of patient state, some of these studies were limited to the communication of a single vital sign, while others did not consider suitable anatomical placement of the tactors that would be suitable for use in a clinical context. Moreover, the prior research has not considered the potential of this modality to *replace* or *reduce* the dependence on auditory alarms and graphical displays in the clinical context. Replacing or reducing such dependence result in less auditory alarms going off therefore quieter hospitals which is beneficial for patients and caregivers alike. Thus, inspired by the prior work, especially the studies conducted by Ferris and Sarter, we seek to demonstrate the benefits to clinical performance of a vibrotactile patient monitor as a potential alternative to the baseline condition of auditory and visual displays in the OR or ICU.

Chapter 3

Did you feel that? Developing Novel Multimodal Alarms for High Consequence Clinical Environments

As an attempt to attenuate the auditory sense overload, we propose the use of a multimodal alarm system in critical care units. Specifically, such a system that utilizes multisensory integration of haptic and auditory channels. We hypothesize that by combining these two channels in a synchronized fashion, auditory perception threshold of participants will be lowered, thus allowing for an overall reduction of noise volume in hospitals. In this chapter, once we elaborate on this hypothesis regarding audio-haptic multimodal effect, then we detail our methodology to verify such claim. We conclude this chapter by exploring future directions to improve the proposed multimodal alarm system with the purpose of reducing hospitals' sound exposure level.

3.1 Preface

This study was published at the Proceedings of the International Conference on Auditory Displays (ICAD) [40] and the text of Chapter 3 is taken from parts of that paper. I am the leader author on the paper and my colleagues, Roger Girgis and Taeyong Kim contributed to the paper as co-authors. The project's clinical collaborator, Joe Schlesinger, is the fourth author and Jeremy Cooperstock is the senior author on this paper. In the project, myself and my colleagues, Roger and Taeyong, were responsible for the experimental design and implementation, while I was solely responsible for data collection and analysis. The content of the Experiment and Methodology sections of the paper, which also shaped this chapter, were written by myself while suggestions and edits were provided by my colleagues, Roger and Taeyong, as well as Prof. Cooperstock and Dr. Schlesinger.

3.2 Methodology

As discussed in Section 2.2 there is evidence from the literature to support the hypothesis that multisensory integration may lead to participants perceiving sound at a lower threshold. We are equally interested in determining whether this effect may hold when the non-audio stimulus is delivered at a sub-threshold level. That is, can we reduce the level of auditory alarms in a clinical environment by delivering a complementary non-auditory stimulus, ideally, one that the clinician does not even perceive? In this section we illustrate our proposed approach for testing and answering this question.

In order to investigate this question, it was first necessary to determine the unimodal thresholds of perception for both the auditory and non-auditory stimuli. Our experiment therefore consists of three measurements:

1. haptic (vibration) perception threshold
2. auditory perception threshold
3. auditory perception threshold when combined with haptic stimulus

One of the most popular methods to map the relationship between physical stimuli and psychological response of the participant [41] is Parameter Estimation by Sequential

Testing (PEST) [42], [43], an adaptive staircase method that has shown its adaptability and robustness in obtaining a perceptual threshold value. Figure 3.1 represents a typical double staircase for measuring the auditory threshold for one of the participants.

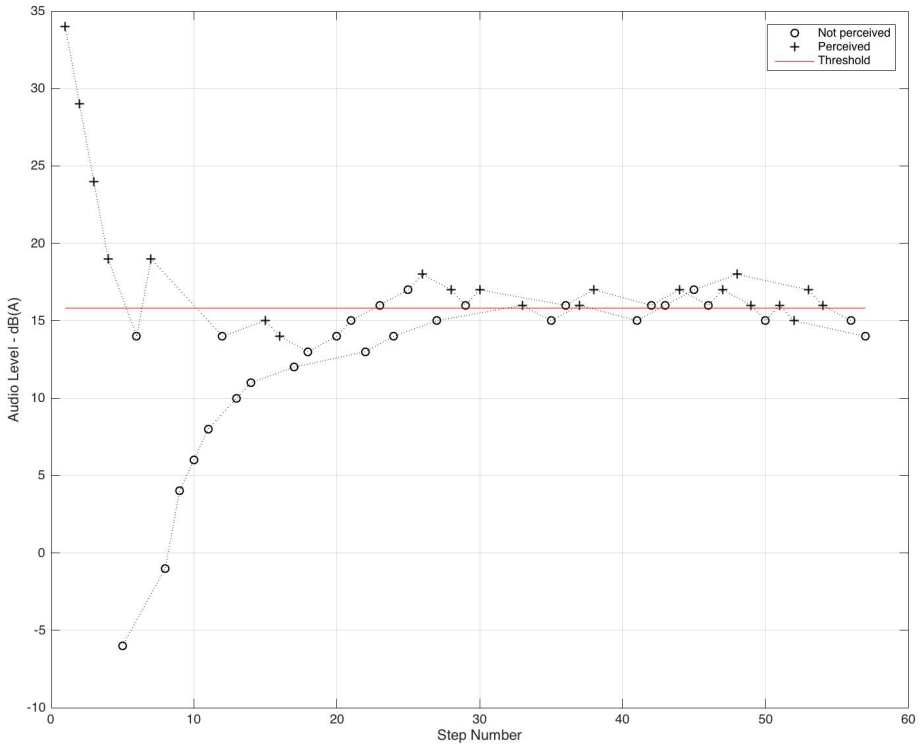

Fig. 3.1: An example of the double-staircase method to determine the auditory threshold of perception for one of the experiment participants.

An improvement is to double the step size in response to several identical responses, and halve the step size in response to a change in consecutive responses. This helps achieve faster convergence and improves participant focus, and was therefore adopted for our testing [41].

To reduce the effects of bias that arises after several identical responses to a given stimulus, Cornsweet suggested the use of the random double-staircase method [44]. The test participant is presented with two staircases, starting from values above and below

the assumed threshold, respectively. The step size begins relatively high to ensure fast convergence, and as the two staircases approach each other, the step size is reduced to ensure a smooth combination of the two. This results in a range of values bounding the threshold of perception.

Throughout our experiment, we employ the PEST procedure coupled with the use of the random double staircase to determine the threshold of perception.

3.3 User Experimentation

3.3.1 Environment

The experiment was performed in the Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT). The lab is acoustically insulated from the surrounding rooms. In addition, participants wore Beyerdynamic DT 770 Pro (Heilbronn, Germany) circumaural headphones during the experiment, and the ambient temperature was maintained between 21 and 25 throughout the tests, thereby ensuring a well-controlled environment.

3.3.2 Stimuli

To provide the vibrotactile stimulus, we used a Tactile Labs Haptuator Mark I (Montreal, Canada) [45], which allows for independent variation of the amplitude and the frequency. For our experiment, the haptuator was connected to a Sparkfun TP2005D1 audio amplifier (Boulder, CO, USA), and strapped snugly to the participants' leg, above the ankle, using a Velcro band as shown in Figure 3.2. One of the potential confound is that actuation of the vibrational device may be audible, contributing to the sound volume of the alarm stimulus. However, this was mitigated by the low intensity of vibration, the placement of the actuator on the participants' leg, and the use of closed headphones throughout the experiment.

The choice of placing the vibration device on the ankle rather than the wrist was motivated by our intended use case of delivering alarm signals in a medical environment, for which hygienic constraints preclude the wearing of devices on the hands or wrists. In order to support this decision, we conducted an additional experiment, employing the same staircase method as described in previous section, and compared the vibration perception

Fig. 3.2: Position placement of the actuator on the participants' ankle, from experiment instructions. The image of the foot is drawn by an anonymous student in our lab.

threshold of participants when the device is attached to wrist versus ankle [46]. We found that haptic perception properties are similar between these two body locations. Therefore, ankle was chosen as the location to wear the vibration band.

Delivery of stimulus during the experiment and logging of measurements for the double staircase was managed by a MATLAB script (MathWorks MATLAB R2016a, Natick, MA, USA). A one-second auditory stimulus was extracted from a recording of the Philips MP-70 (Amsterdam, Netherlands) patient monitor red/crisis alarm. The frequency spectrum of the alarm sound is shown in Figure 3.3. The choice of a one-second duration was deemed to be reasonably short to help eliminate guesses, and sufficiently long so as to include the salient auditory characteristics of the alarm signal.

The vibratory stimulus was generated using a sine wave at 175 Hz, output using the MATLAB sound() function for a duration of 1 second. For the combined auditory-haptic stimuli, the signals were output in unison, using a stereo audio splitter to separate the audio (left channel) and haptic (right channel) signals.

3.3.3 Experimental Procedure

Participants first completed a pretest questionnaire to screen for possible health conditions that might exclude them from the experiment. They were then asked to read an instruction sheet, put on the headphones, and assisted with securing of the haptic band just above the ankle at a comfortable location.

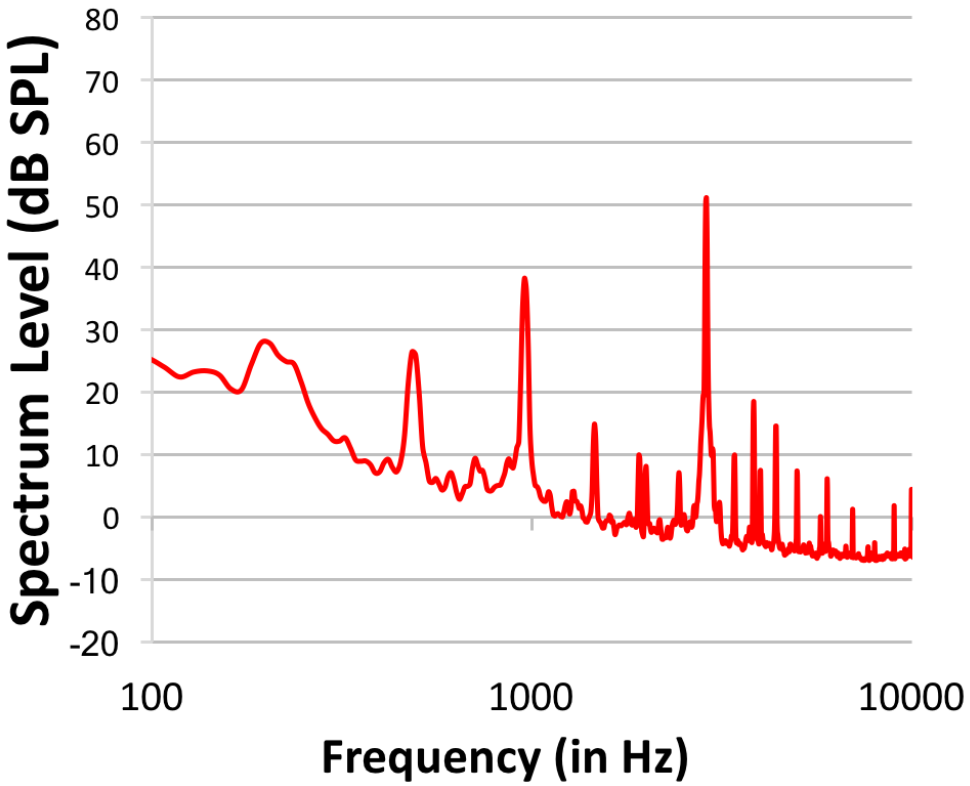

Fig. 3.3: Alarm waveform, measured with a class II Amprobe SM20A sound level meter provided a weighted output of 49 dB.

The strap was secured in a snug fashion for good coupling between the actuator and the skin, but not so tight that it caused discomfort. Participants were then asked to place their right foot in a comfortable position and to immobilize it for the remainder of the experiment.

For each experimental condition, participants initially carried out a training/calibration step to familiarize themselves with the experimental stimuli and adjust these to a level in which they were barely perceptible, using a coarse staircase method. This block took approximately 10 minutes to complete.

Participants then proceeded through the first block of the experiment, which determined their haptic perceptual threshold. This block of the experiment took around 15 minutes to complete. In the second block, we determined the auditory perceptual threshold, both

with and without a combined haptic stimulus. This was done by intermingling two tests, randomly selecting half of the trials for presentation of unimodal auditory stimuli, and half for presentation of combined audio-haptic stimuli with the haptic stimuli delivered at a fixed level. This intermingling was done to avoid potential habituation effects that may have biased the threshold estimates in either direction. The participants took around 25 minutes to complete this block.

To ensure that participants did not go too long between successive perceivable stimuli, which we observed during pilot testing as a significant source of fatigue, the system randomly delivered 20% of the stimuli at 3.5 standard deviations above the average of the intensities of the last six stimuli, a choice we determined empirically as adequate to ensure reasonable supra-threshold perception.

For our purposes, participants had to respond within 2 s from the onset of the stimulus, i.e., no more than 1 s following its presentation, by clicking on a button displayed in a simple graphical user interface; otherwise, it was assumed that they did not perceive the stimulus. The system then reduced or increased the subsequent stimulus intensity by a defined step size so as to maintain the intensity just at the edge of perceptibility. The step size was reduced as both staircases converged, i.e., as the difference in intensities between the upward and downward staircases decreased. This process continued until a minimum step size was reached, then six reversals were counted on each staircase for the threshold estimation. The perceptual threshold was estimated as the mean of the stimulus levels of the last six reversals from each staircase. For each participant, the number of audio or haptic stimuli presented until their perceptual threshold is measured depended on the recorded responses. Therefore, the total number of trials in each block of the experiment varied for each participant. The whole experiment took around one hour to complete.

The aforementioned design was employed for all of our experiments, with variations as described in the following sections.

3.4 Findings

3.4.1 Pilot

Initially, we fixed the sub-threshold amplitude of vibration at 2.5 standard deviations below the perceptual threshold determined in the first block.

Pilot testing was performed on 11 lab members (10 male, 1 female), over the span of three days. The mean age of the participants was 30.4 years with a standard deviation of 8.6 years. Data from one of these participants were excluded from the analysis for failing to respond to supra-threshold "wake-up stimuli", which suggested a lack of attention during the experiment. The test participants participated on a voluntary basis and did not receive monetary compensation for their time.

Five of the 10 participants whose data were retained for analysis exhibited a slightly lower auditory threshold when measured in the multimodal audio-haptic condition than in the unimodal audio-only condition. Although the results of these initial tests were only borderline in terms of statistical significance, we were encouraged to carry out a larger experiment with additional participants who were naive to the experimental hypothesis.

3.4.2 Full Experiment

The same experimental procedure was then applied to a new group of participants, naive as to the experimental hypothesis. These participants were not informed that sub-threshold vibration was delivered (in conjunction with half of the audio stimuli) during the second block. The test was conducted on 12 participants (10 male, 2 female), over the span of three days. The mean age of the participants was 28 years with standard deviation of 3.5 years. This study was approved by the Institutional Review Board (IRB) at McGill University. Participation in the study was voluntary and informed consent was obtained from each participant. A 10 CAD compensation was given to each participant for their time.

Data from one of these participants, exhibiting a difference between the audio-only and audio-haptic thresholds greater than six times, was excluded from the analysis as an outlier. The test participants participated on a voluntary basis and received monetary compensation for their time. The auditory threshold values measured across 12 participants for both the

audio-only condition and the audio-haptic condition is shown in Figure 3.4. The differences between the conditions are shown in the bar graph of Figure 3.5.

As can be seen, the data did not support our hypothesis that the perceptual threshold is reduced in the multimodal condition. Counter-intuitively, the trend suggested an opposite effect, although not significant ($p = 0.78$, ci = [-0.65, 0.83]). This led us to consider the possibility that our haptic stimuli was too far below the perceptual threshold, and was thus not contributing to the effect.

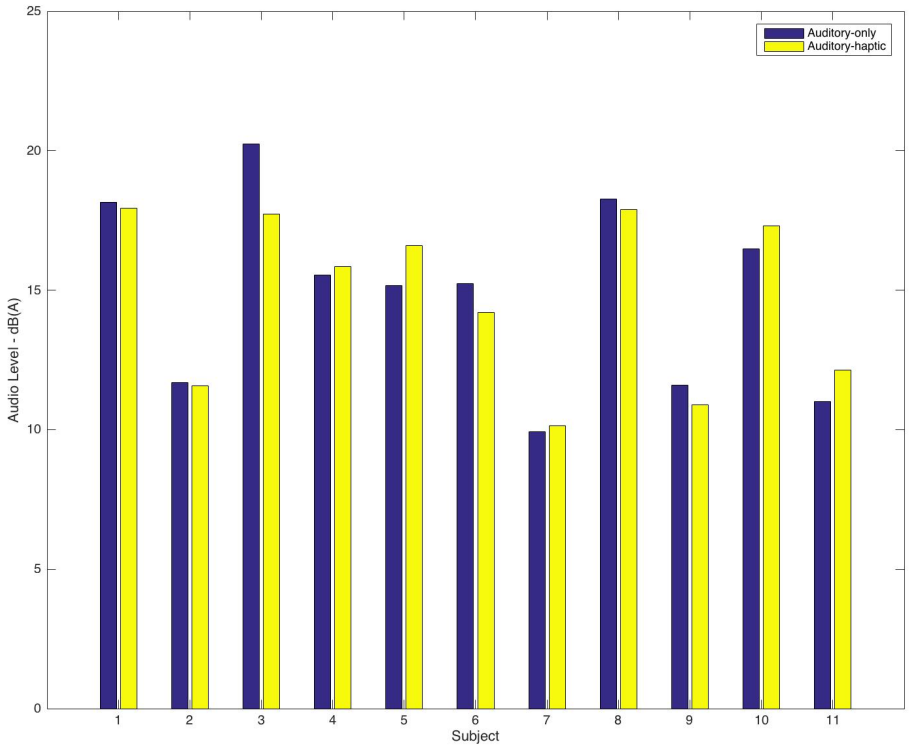

Fig. 3.4: Threshold data obtained from the audio-haptic and audio-only threshold measurement over the 11 participants whose data were retained.

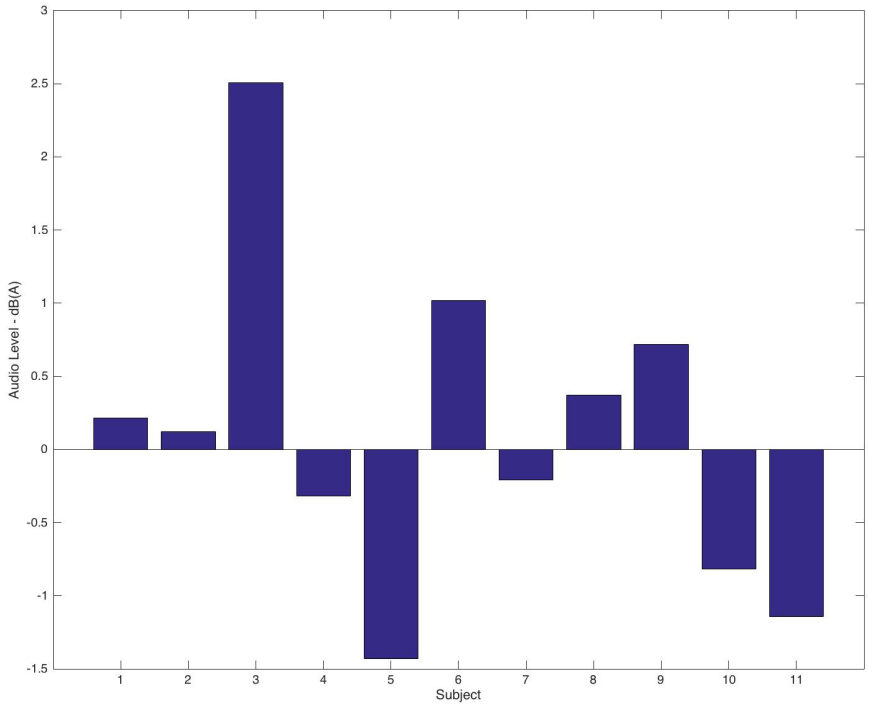

Fig. 3.5: Differences between audio-only and audio-haptic threshold values from Figure 3.4. A positive value indicates that the auditory perceptual threshold was reduced in the multimodal condition.

3.4.3 Increased Haptic Intensity Level

To address the possibility that the haptic stimuli were too far below threshold to have an impact, we then conducted a further experiment in which the amplitude of vibration was increased to 0.5 standard deviations below the threshold determined in block 1. As before, participants were not informed that they might feel vibrations during the third block.

Ten participants (4 male, 6 female) were recruited and the test took place during one day. The mean age of the participants was 27.3 years with standard deviation of 5.7 years. Data collection from one of the participants could not be completed on account of the participant changing the computer's output volume in the middle of the test. This participant was excluded from the analysis. The test participants participated in the study

on a voluntary basis, signed an informed consent form upon arriving at the experiment, and and received 10 CAD monetary compensation for their time.

The perception threshold in the auditory domain measured across 9 participants for both the unisensory and multisensory conditions is shown in Figure 3.6. The differences between the conditions are shown in the bar graph of Figure 3.7.

By way of response to a post-test questionnaire, 5 out of 9 participants indicated that on occasion, they perceived the haptic stimulus during the second block.

Despite the increase of the level of haptic stimulus, we found no statistically significant difference between the threshold of audio perception in the unisensory and multisensory conditions ($p = 0.21$, ci = [-0.90, 0.23]).

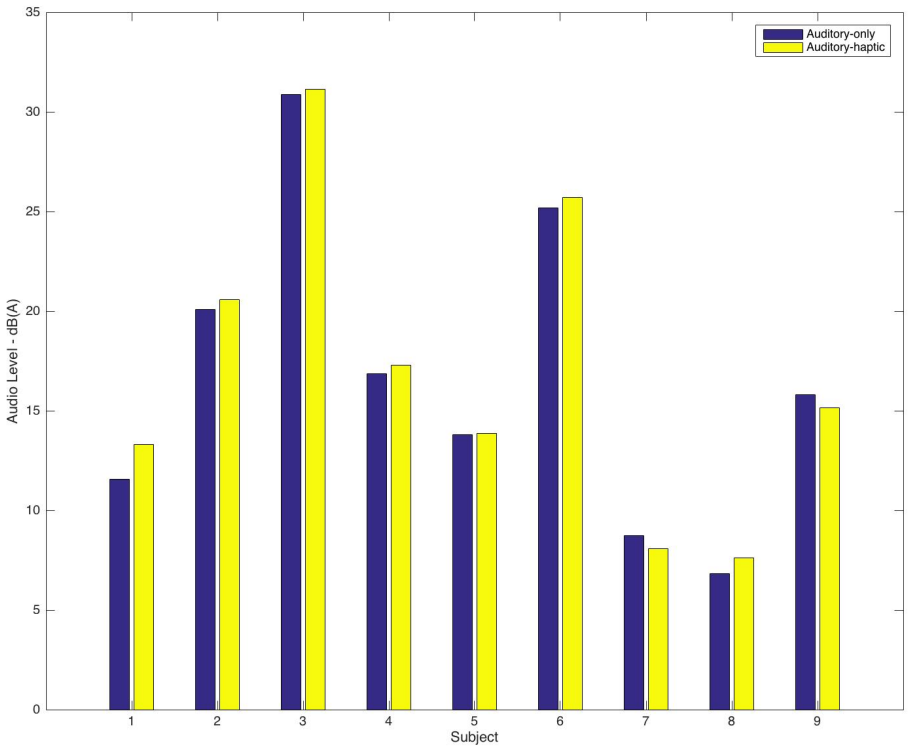

Fig. 3.6: Threshold data obtained from the audio-haptic and audio-only threshold measurement over the 9 participants from the increased haptic intensity level experiment.

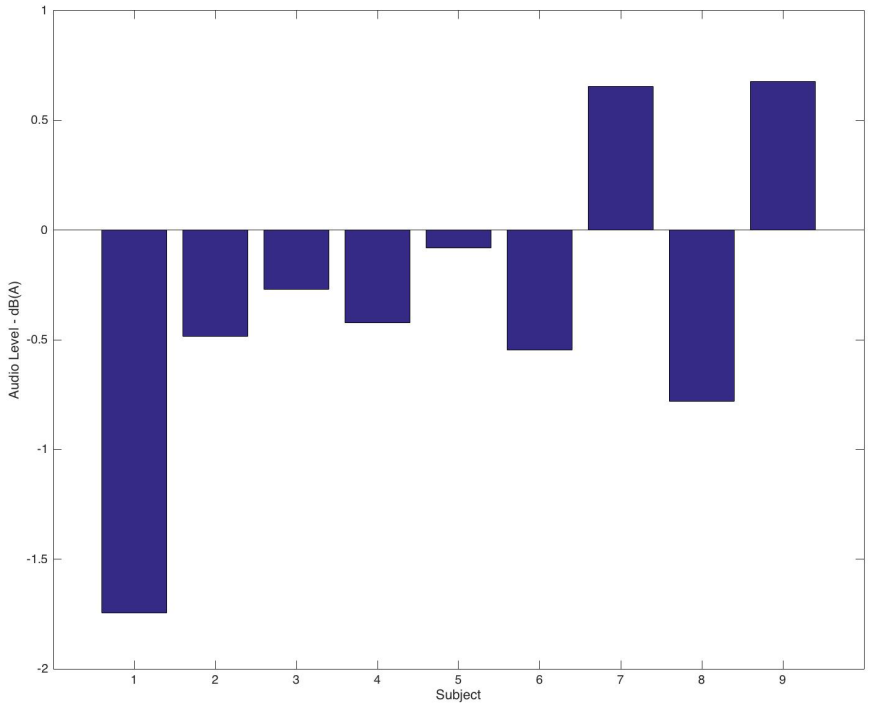

Fig. 3.7: Differences between audio-only and audio-haptic threshold values from Figure 3.6. A positive value indicates that the auditory perceptual threshold was reduced in the multimodal condition.

3.4.4 Circling Around the Thresholds

We considered several possibilities for the results of the previous section:

1. Either or both auditory and haptic perception thresholds varied throughout the experiment, e.g., due to fatigue or habituation, and so the thresholds measured in blocks 1 and 2 were unreliable.

2. The presentation of simultaneous sub- or near-threshold haptic stimuli interfered with auditory perception.

3. There is no multimodal integration benefit from haptic stimuli in conjunction with non-speech audio.

4. There is a multimodal integration benefit from haptic stimuli in conjunction with non-speech audio, but only for supra-threshold haptic stimuli.

To examine these possibilities, we conducted an additional exploratory experiment, in which we varied the level of auditory stimuli in a range of ± 2 standard deviations and varied the level of haptic stimuli in a range of [-2,+4] standard deviations around the unimodal thresholds determined in blocks 1 and 2.

The results of this exploration were revealing, although hardly conclusive: three of the five participants, demonstrated no discernible effect of haptic stimulus on audio stimulus detection, even at clearly supra-threshold levels of haptic stimuli. The remaining two participants demonstrated a possible effect, with slightly higher rates of audio alarm detection for sub-threshold audio when presented in conjunction with supra-threshold haptic stimuli (see Figure 3.8 for the results of one of these participants). However, it does not appear that these results are significant, and thus, we can neither confirm nor reject any of the possibilities described at the start of this section.

3.5 Conclusions

We postulated at the outset of this chapter of thesis that improved perception of auditory stimuli would result through multimodal integration with a complementary haptic signal, possibly even at sub-threshold levels. If so, we hypothesized that this could allow for attenuation of a and assist the practitioner in recognizing the alarm, thereby reducing the problems of stress and alarm fatigue in the clinical settings of the OR and ICU.

Through these experiments, we hoped to determine preliminary guidelines for the outcomes of implementing multimodal alarms, leading to a reduction in the demands on the audio channel. We hypothesize that a multimodal alarm system can attenuate alarm fatigue and assist the practitioner in recognizing the alarm.

While we have so far not been able to verify this hypothesis, we believe that the experimental protocol we developed to address the research questions here will prove valuable to the multisensory research community and can be applied to future experiments that seek to resolve some of the unanswered questions raised in Section 3.3. It is also possible that the PoIE, as observed in other experimental contexts, is only manifested in conjunction with

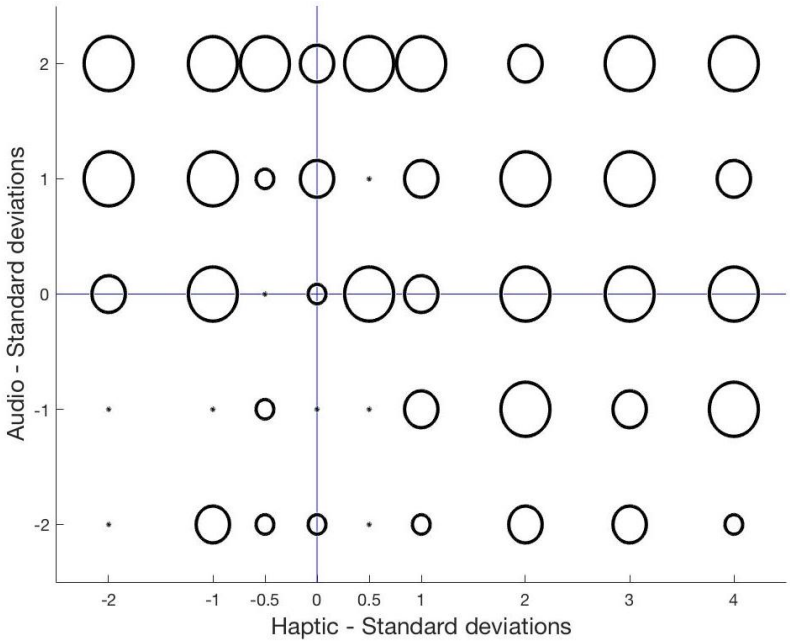

Fig. 3.8: Detection rate for one participant of audio stimuli as a function of audio and haptic stimulus level, relative to the thresholds measured in blocks 1 and 2. The size of each circle indicates the detection rate (out of three presentations) at each combination of parameters.

speech auditory stimuli, for which the neurophysiological responses are differently affected by the influence of a secondary stimulus modality. This would need to be determined through a separate experiment, employing a speech cue rather than an auditory alarm sound.

Future experiments can employ the PoIE by adding hospital background noise, auditory speech-in-noise tasks, and visual vigilance tasks to test if the hypotheses would still hold while testing clinicians during simulated emergencies requiring clinical pharmacologic intervention. These data will not only inform alarm design and improve patient safety, but have wide-ranging applications to other high consequence industries.

Chapter 4

Haptic Alarm Displays: Reducing Clinician Dependence on Patient Monitors and Auditory Alarms

The need to attend to patients' vital signs in high-consequence hospital environments poses significant demands on clinicians. Visual complexity and often problematic location of patient's monitoring displays and their integration with a loud, often irrelevant auditory alarm not only elevate stress level of clinicians but also result in potential health hazards for patients. These issues have motivated research aimed at shifting the delivery of physiological vital sign information and announcement of alarm events from visual and auditory devices to haptic transducers. Building on the work of Chapter 3, investigating the potential of multisensory audio-haptic integration to reduce the perceptual threshold for detecting alarms, in this chapter, our motivation is to introduce haptic information displays in hospitals and explore their effectiveness in improving clinician performance. If a haptics-only display can alert clinicians to abnormal physiological conditions it can result in obviating the need for the highly distracting auditory alarm signal to draw attention to the visual monitor. In this chapter we elaborate on our methodology for validating our hypothesis that vibrotactile delivery of vital sign information can support alarm-state vital sign identification competitive with graphical and auditory alarm display conditions, without significant impact on performance on a parallel attention-demanding activity.

4.1 Preface

This chapter is written based on our manuscript submitted to the ACM Transactions on Computing for Healthcare (HEALTH) journal which is currently under review. I am the leader author on the paper. A McGill postdoctoral fellow Antoine Weill-Duflos is the second author, our project's clinical collaborator Joe Schlesinger is the third author and my supervisor Jeremy Cooperstock is the senior author. In this project, I was responsible for the experiment design, implementation, and data collection, while Antoine carried out the statistical analyses and generated the plots. In terms of the written parts of this chapter, I reviewed the literature, and wrote the Methodology, Experiment, Discussion and Conclusion sections. Antoine was responsible for writing the result and analysis sections of each experiment. All the sections were thoroughly reviewed by Jeremy Cooperstock and Joe Schlesinger and suggestions and edits were provided.

4.2 Methodology

Our objective in this work is to demonstrate the feasibility of vibrotactile display of patient state as an alternative to auditory and visual displays. We expect this method to result in no significant impairment of performance, both in terms of response time to the alarm, and identification accuracy. If successful, such an approach would offer several benefits, namely, reducing the sound level in the OR and ICU environments, improving the ability of clinicians to retain their visual focus and attention on the patient, and offering the possibility of personalized alarms delivered to individual clinicians.

To investigate this possibility, we designed and conducted several experiments to determine how effectively participants could discriminate both the vital sign and its state (high or low) during alarm events, when this information was rendered haptically. In all of these experiments, participants were given a realistic context in which they had the role of a surgeon. We also required participants to perform a parallel distractor task that simulated the surgery they were performing. We thus explained to them that failing to attend to the patient's vital signs or poorly completing the distractor task will put their patient's life in jeopardy. Additionally, to reflect real-world conditions in which clinicians are actively attending to clinical activities in a noisy environment an audio recording from an operat-

ing room of Montreal's Sainte-Justine Hospital played continuously in the background, at a volume between 56.7 dB and 62.7 dB, as measured by an Scosche SPL1000 meter using C-curved measurement (`https://www.scosche.com/scosche-spl-meter-135db-max`).

Our participants were recruited from a generic population instead of clinical staff or medicine students, particularly since these studies were meant to build a proof of concept for further future research on the issue of noise in hospital. We believe, in the next steps of this research, recruiting participants that are directly involved in medical task set will bring a heightened perspective on the design of our proposed vibrotactile alarm system.

This research complied with the American Psychological Association Code of Ethics and was approved by the Institutional Review Board (IRB) at McGill University and Vanderbilt University. Informed consent was obtained from each participant. Participants signed a consent form at the beginning of the test. The remainder of this section describes common elements of the three experimental designs and apparatus, as shown in Figure 4.1, before we turn to the specifics of each.

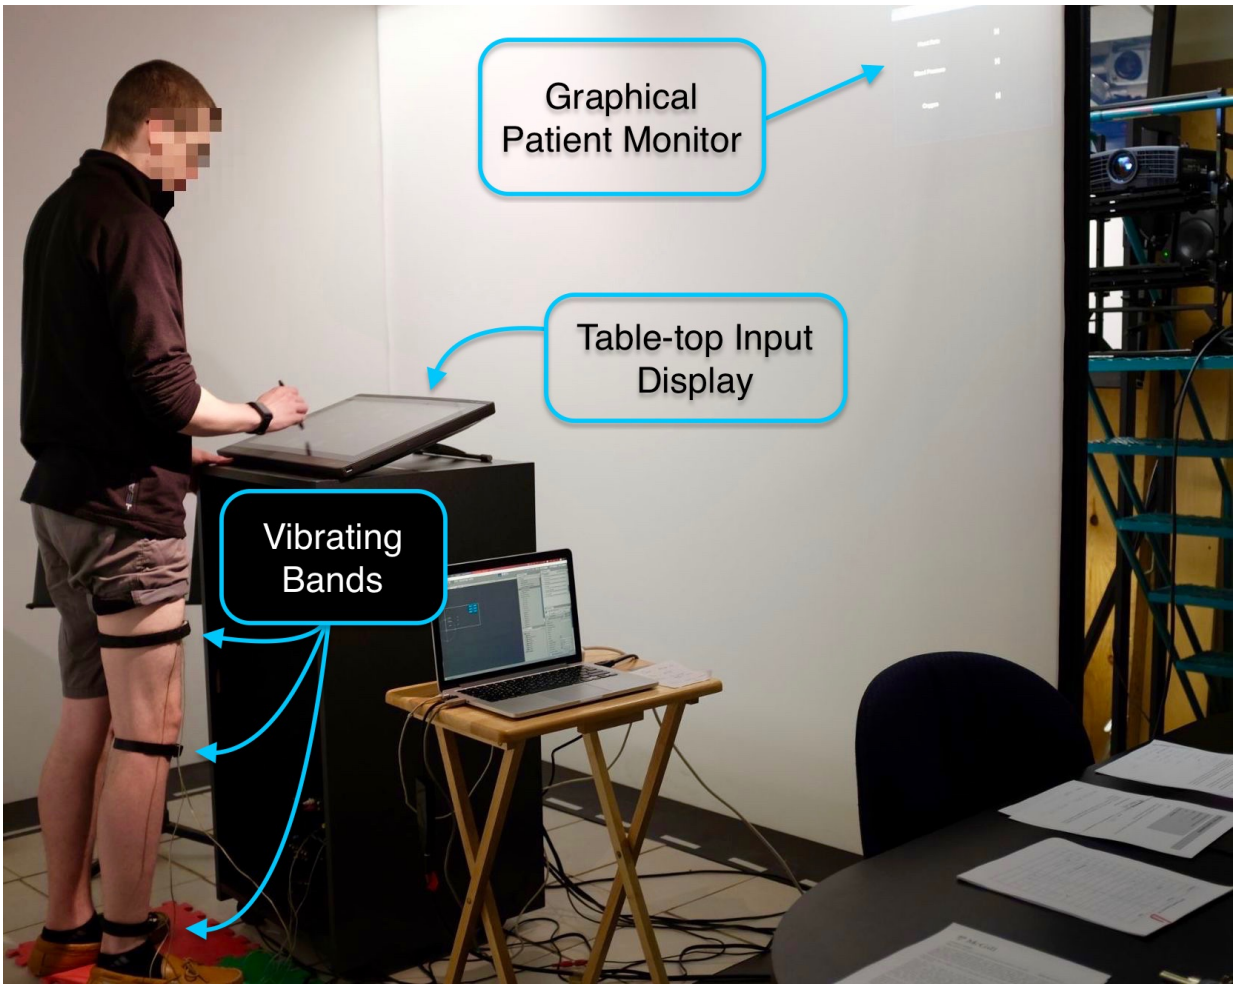

Fig. 4.1: Placement of the patient monitor, input display, and the vibrating bands, worn by the participant.

4.2.1 Input Interface

The graphical user interface that appeared on the input display is pictured in Figure 4.2. Participants detected alarm events by selecting one of the six buttons appearing in the upper right quadrant. In parallel, participants must attend to the competing distractor task, designed to occupy a portion of their attention and cognitive resources [47]. For this purpose, we employed a Fitts' law task [48], involving repetitive pointing to an on-screen target (red circle in Figure 4.2), that changes position after each selection. Importantly, this

distractor task does not interfere with the auditory or haptic modalities, which are occupied attending to the alarm signals. The Fitts' law task meets the ergonomic requirements of a medical procedure, in which the participant remains standing, with a downward visual focus, performing a precise task with the dominant hand, such as placing a right internal jugular central venous catheter. Moreover, the Fitts' law task had benefits over a memory or decision making task. The reason is that, since the nature of this experiment was cognitively demanding, especially given the short training time, a distractor task that is unfamiliar and requires training and cognitive resources would frustrate the participants and greatly deteriorate their general performance in the study. Lastly, even though there are many simple simulations of medical procedures available online, we could not find an open-source and freely available software to integrate with our study, therefore we implemented this simple task.

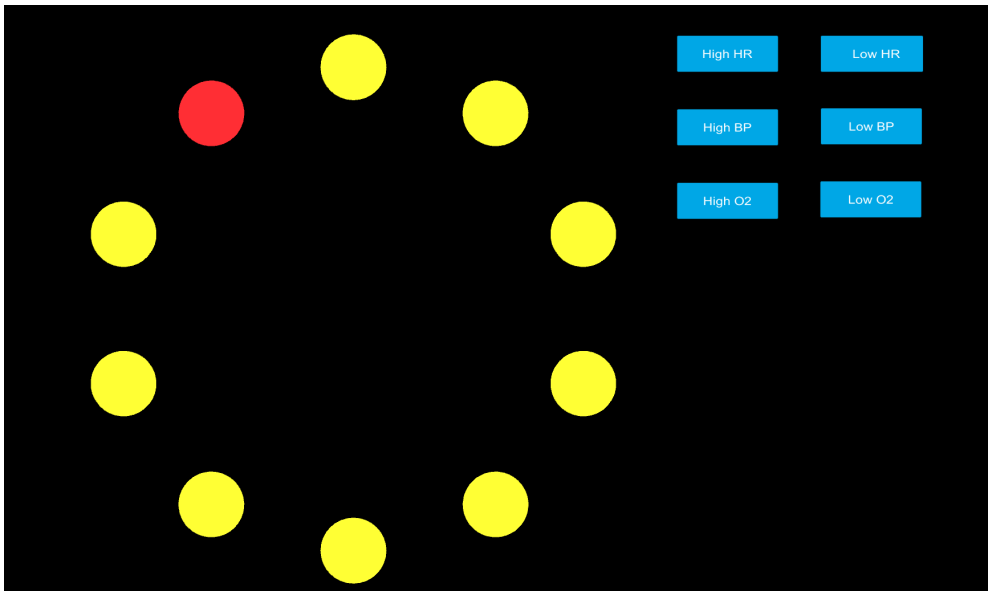

Fig. 4.2: Graphical User Interface (GUI) appearing in the table-top input display, used for the distractor task (yellow and red circles in left half of screen) and to record participant responses to the vibrotactile cues (blue boxes in upper right).

4.2.2 Patient model and vital sign display

Three vital signs were chosen and introduced to the participant: heart rate (HR), blood pressure (BP), and oxygen saturation (O_2). Each of three vital signs was represented internally as one of five levels: normal, nominal-low, nominal-high, low, or high. All vitals begin at a normal level, and only the low and high levels are used to trigger alarm events. At every simulation step, 2.5 s in duration, each vital sign has an 80 % probability of remaining in its current state, a 10 % probability of increasing, and a 10 % probability of decreasing, until reaching either of the extrema, i.e., low and high. This method is meant to randomize the time it takes for a parameter to reach alarm state from normal state.

A separate thread was responsible for driving the patient monitor, audio alarm, vibrotactile display, or a combination of these, as appropriate to the experimental condition. On each iteration, the three vital sign levels were checked, and the appropriate output is rendered, as described in the remainder of this section. For each experimental condition, response time is measured from the start of the display of the alarm event to the participant's response, indicating through a table-top pen display (Wacom DTU-2231 `https://www.wacom.com/en-cn/enterprise/business-solutions/hardware/pen-displays/dtu-2231`) to select which vital sign had gone out of bounds, and its level (low or high), thus simulating a clinical intervention to the alarm states.

4.2.3 Auditory Alarm

The audio alarm was based on a recording of the Philips MP-70 (Amsterdam, Netherlands) patient monitor red/crisis alarm, played at 74.5 dB, determined to be easily audible above the background noise. This was sounded simultaneously with the rendering of the associated vital sign through either the graphical patient monitor or the vibrotactile display, as per experimental condition, for a duration of 2.0 s, ensuring that response time measurements began from a common starting point.

4.2.4 Graphical Patient Monitor Display

A graphical patient monitor, situated at an azimuth of approximately 45° relative to the participant and elevation of approximately 1.7 m off the ground, displayed patient vital

signs. This configuration was based on typical arrangements found in hospital environments (see Figure 4.3). However, turning to glance at the patient monitor would involve a temporary disruption of attention to the distractor task running on the input display directly in front of the participant.

Fig. 4.3: Illustrative representation of monitor placement in an ICU, on the right side of the image. One clinician has her view of the monitors occluded, while the second clinician has his back to the monitors. Image titled "Clinicians" used under CC Attribution-Share Alike 3.0 license.

Each vital sign is displayed on a separate line by its name and the associated state. The display updates immediately in response to changes in vital signs of the patient model. For the purpose of highlighting states associated with alarm conditions, the graphical display indicates both low and high states by displaying a downward ↓ or upward ↑ arrow adjacent to the out-of-bounds vital sign, as illustrated for blood pressure in Figure 4.4. This simple informational layout was chosen so as not to require any interpretation of the displayed data on the part of our participants.

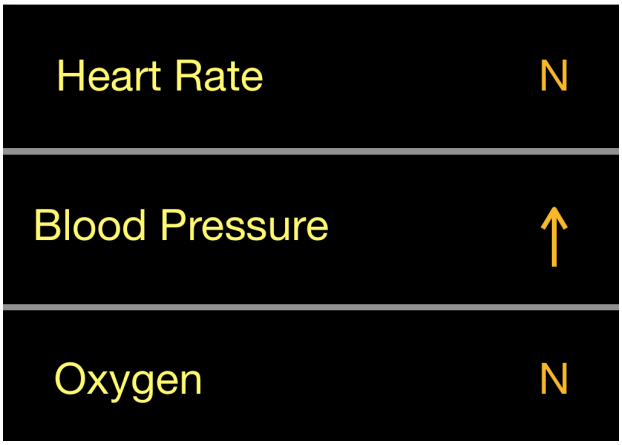

Fig. 4.4: Representation of an abnormal condition (high heart rate) delivered through the graphical display.

4.2.5 Vibrotactile Display

The vibration signal, encoding the patient vital signs, was delivered by Tactile Labs Haptuator Original tactors (`http://tactilelabs.com/products/haptics/haptuator/`), and were supplied with driving signals of 1.5 s duration at a frequency of 170 Hz. The duration and frequency were refined and chosen based on users feedback after multiple pilot tests.

Our design of haptic icons, or "Tactons" [49], to represent patient state was intended to convey through vibration both the type and level (high, nominal, or low) of three different vital signs: heart rate (HR), blood pressure (BP), and oxygen saturation (O_2). To do so, we drew from the guidance of previous research in this area. Brown et al. investigated the role of roughness, rhythm, and spatial location as vibrotactile parameters, reporting the highest recognition rate (99.8 %) when spatial location for information delivery was varied [50], [51].

In order to convey the *type* of abnormal vital sign, initially, we explored the possibility of using a vibration pattern based on the spoken (English) rhythm of the words representing each vital sign. Figure 4.5 shows the design of those tactons. However, feedback from a pilot test indicated that differentiating patterns with the same number of syllables ("blood pressure" and "oxygen") was difficult, and the patterns were too long, with the result that participants preferred instead to obtain the necessary information from the visual patient monitor. Therefore, we switched to considering spatial location of vibration bands as a

parameter to communicate the type of abnormal vital sign. This decision was supported by previous tacton studies. Brown et al. investigated the role of three vibrotactile parameters, roughness, rhythm, and spatial location, and reported the highest recognition rate (99.8 %) when spatial location for information delivery was varied [50], [51]. In a similar study, Barralon et al. proved that spatial location is the parameter that exhibited the best overall accuracy, comparing to rhythm and roughness, in delivering patients physiological information when prototyping their proposed tactile display [52]. Finally, for the studies described in the remainder of this chapter, we employed variant spatial location of multiple actuators as a means of representing the type and rhythm of vibrations as a means of conveying the level of abnormal vital sign. The rendering of patient information in each of the experimental conditions is shown in great details in Table 4.1.

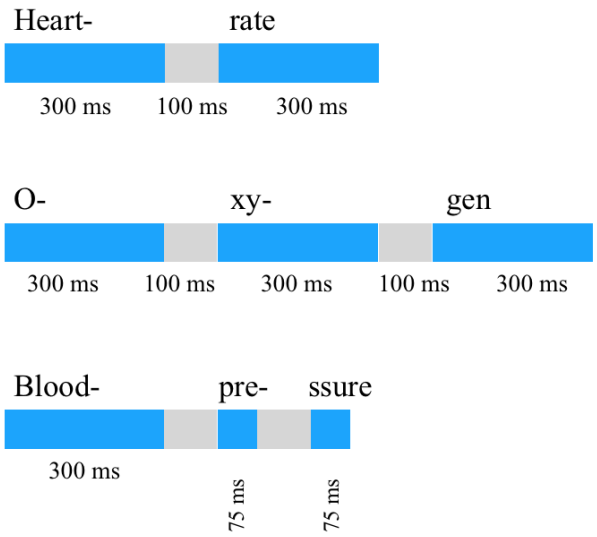

Fig. 4.5: Vibration patterns based on the spoken (English) rhythm of each vital sign's title; from top Heart Rate, Oxygen and Blood Pressure. Blue color represents vibration buzz while gray color shows silence.

Condition	Cues for		
	Normal Vital	Low Vital	High Vital
baseline	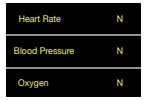	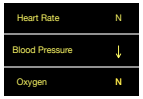 + Audio Alarm	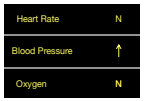 + Audio Alarm
swipe		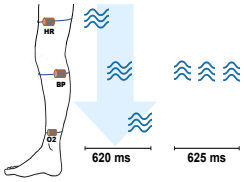	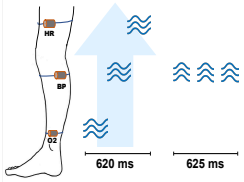
1-2-3		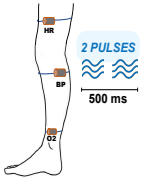	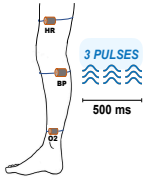
haptic alarm-only	[silent]	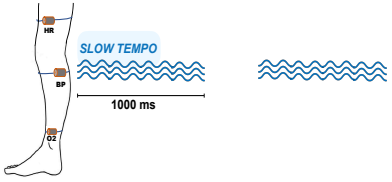	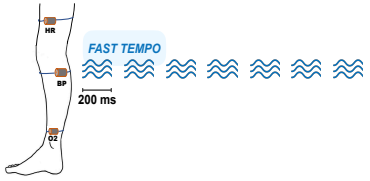

Table 4.1: Illustration of the vital sign feedback cues by condition, as used in the three experiments. The table shows a sample scenario in which BP enters an abnormal condition triggering an alarm event. The pattern remains similar for other vitals on different spatial locations on the leg. *Condition 1-2-3 is introduced in Experiment 2.

To familiarize participants with the experimental stimuli, they first carried out a training tutorial section. They moved to the actual experiment after, which was conducted in one of the three forms described in the next sections.

4.3 Experiment 1: Vibrotactile vs. Graphical

Our first experiment compared participant performance under haptic rendering conditions against those employing the graphical patient monitor, both with and without an accompanying auditory alarm. The latter was included as part of the experiment in order to assess whether haptic information display under an alarm condition could make the auditory alarm redundant.

We assigned each of the three actuator locations, on the participant's right leg, as shown in Figure 4.13 (left), to a specific vital sign. Normal vital levels were indicated by a short buzz (250 ms) of the corresponding actuator, with intervals of 2.5 s between each actuation. Abnormal low (or high) states were indicated by a "swipe" pattern of activation moving down (or up) the three actuators. The swipe consists of a vibration of 250 ms of each actuator, with a 65 ms overlap, for a total of 620 ms. This pattern is followed by a short gap of 130 ms, and then, a series of three buzzes of the actuator at the location corresponding to the vital sign being rendered. Each buzz lasts 125 ms with an inter-buzz gap of 125 ms, for a total duration of 625 ms, as illustrated in the "swipe" row of Table 4.1. The four experimental conditions were as follows:

- *baseline:* graphical display with auditory alarms
- *graphics only:* graphical display only; no auditory alarms
- *swipe+audio:* haptic patterns with auditory alarms
- *swipe:* haptic patterns only; no auditory alarms

In the *baseline* and *swipe+audio* conditions, auditory alarms sound simultaneously with the display of the alarm information, whether graphically or through vibrotactile haptic cues, directing the participants' attention to the alarm event. In the absence of alarm information in the *graphics only* condition, participants must glance periodically at the

graphical information display to notice alarm events. However, attending to the graphical display is implicitly discouraged, since this interferes with performance on the distractor task. The graphics only condition was included mainly to incorporate results from all the four possibilities of alarm modalities. Response time was measured from the onset of the alarm event until the participant entered by clicking on the corresponding selection on the pen tablet's graphical interface (see Fig. 4.2), which vital sign had gone out of bounds, and whether this vital was high or low.

The alarm event continues until the participant correctly acknowledges the type of alarm event, at which point, all vitals are reset to normal levels.

We conducted the test on 16 participants, 14M/2F, ages 21-35, $\bar{x}$=27. The order of experimental conditions was balanced across participants, with the constraint that the two haptic conditions were not presented consecutively. The experiments took place in a laboratory in the McConnell Engineering Building at McGill University. The recorded data and their interpretations are presented in the following.

4.3.1 Results

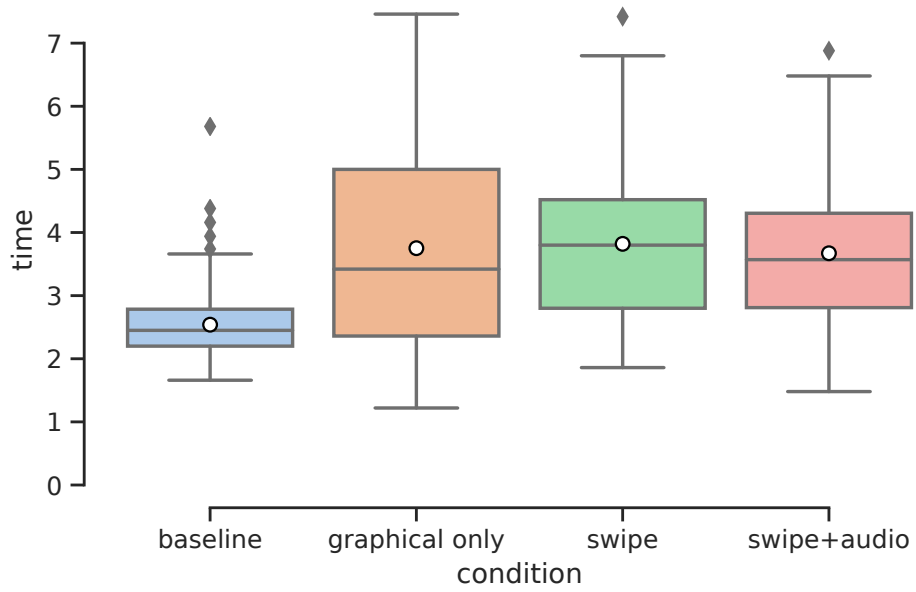

Fig. 4.6: Response time (in seconds) by condition for Experiment 1: vibrotactile vs. graphics.

Figure 4.6 shows the response time under the experimental conditions. The distribution of results does not follow a normal distribution, but is skewed, as confirmed by a Shapiro-Wilk test. Consequently, non-parametric tests are used to assess the statistical significance of the results. A Friedman test, as implemented in the python package [53], found a statistically significant difference between the experimental conditions, ($Q = 27.7$, $p_{value} < 0.001$). A pairwise comparison with a Holm correction confirmed the statistically significant difference between the *baseline* condition and the other three conditions ($p_{value} < 0.001$), with a large effect size (Hedges' g close to 1).

The overall scores for the distractor task are computed as the Fitts' index of performance [48] divided by the total time spent, per condition, with the results shown in Figure 4.7. A higher score is better.

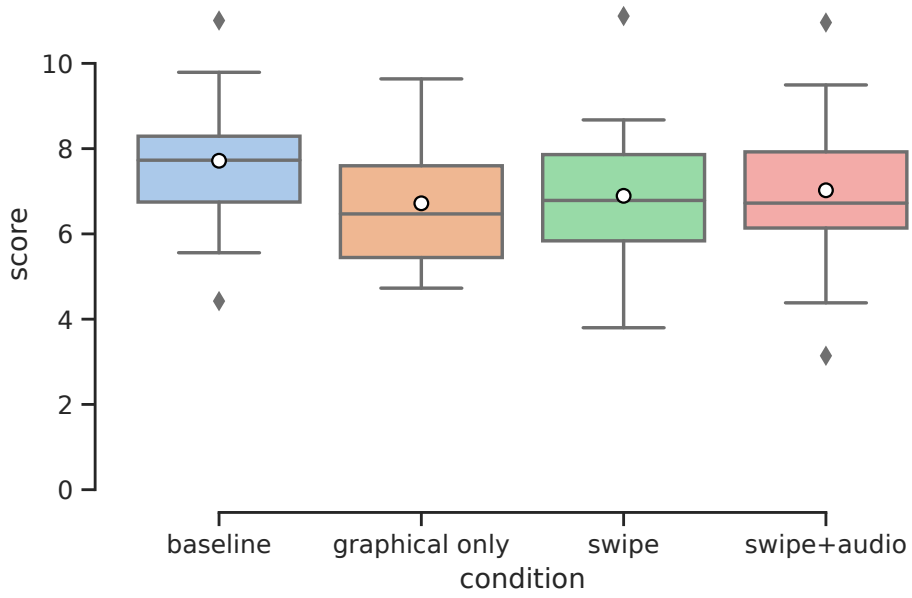

Fig. 4.7: Performance on the Fitts' law distractor task by condition for Experiment 1: vibrotactile vs. graphics.

For all conditions, the distribution of the scores likely follows a normal distribution, as checked with both QQplots and a Shapiro-Wilk test. A one-way repeated measures ANOVA confirmed a statistically significant difference between the experimental conditions, ($F(3, 45) = 3.8$, $p = 0.016$, achieved power of 0.92). A pairwise T-test confirmed a statistically significant difference between baseline and the *graphical only* and *swipe* conditions ($p < 0.05$) with a medium effect size (Hedges' g close to 0.5).

4.3.2 Analysis

Neither of the haptics-only conditions performed as well as the baseline. However, there was also no obvious difference in performance between the graphical information display and the haptic conditions. As anticipated, performance with the graphical information display was superior in conjunction with the auditory alarm, since participants were informed instantaneously when an alarm condition occurred.

These results raised the question of whether the inferior performance under the haptic

swipe conditions was endemic to haptic information display, isolated to the leg, or indicative of a limitation specific to our design.

4.4 Experiment 2: Distal Actuator Separation

As an alternative to confining actuator placement to the leg, our next experiment therefore investigated distribution of the actuators over distal body positions. In our prior work [40], we avoided placement of actuators on the arms, since there may be sterility constraints for any hardware below the elbow. However, other research found the arm to provide high accuracy of recognition, sensitivity, and comfort, compared to other body locations [54]. We introduced a new condition using both arms and a leg as the tactor locations, as illustrated in Figure 4.8, as a possible means of overcoming the limitations of haptic signal discrimination over three positions on the leg.

Distributing the actuators in this manner, with those on the arm kept near or above the elbow, avoids the sterility concerns noted above. However, with this arrangement, the hardware begins to resemble more closely the potentially cumbersome vibrotactile garment of [39]. Furthermore, in a more real-world setting, arm movement may lead to some masking of the vibration signal [55].

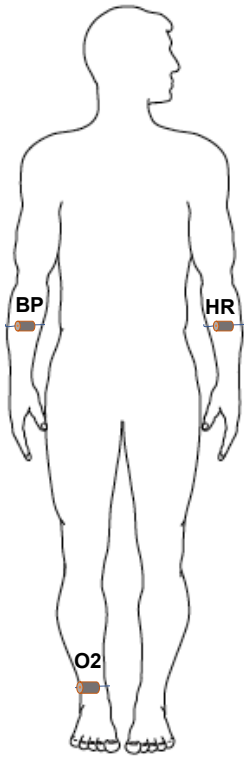

Fig. 4.8: Distal placement of actuator at the arms and one leg. The image of the body is drawn by an anonymous McGill student.

Each actuator was mapped to a different vital sign: right arm for BP, left arm for HR, and right ankle for O2. A short buzz of the actuator, 250 ms in duration, indicated a normal level of the corresponding vital sign. Given the non-linear arrangement of actuators, the representation of high or low vital sign levels could no longer benefit from the swiping up or down metaphor used in the previous experiment. Instead, we used a single pulse to indicate a normal level, double pulse, each of 200 ms duration separated by a gap of 100 ms (500 ms total), to indicate a low level and a triple pulse, each of 100 ms duration separated by a gap of 100 ms (also 500 ms total), to indicate a high level, at the corresponding actuator location.

For reference, we compared performance with this distal actuator arrangement against the baseline and leg-based swipe conditions used in the previous experiment. Thus, the conditions included:

- *baseline:* graphical display with auditory alarms
- *swipe:* haptic leg swipe patterns as used in previous experiment
- *1-2-3:* 1 pulse for normal, 2 pulses for low, and 3 pulses for high level of vital signs, delivered to distal arms and leg locations

The experiment was conducted with 13 participants, 6M/7F, ages 18-35, $\bar{x}$=24, none of whom participated in the previous experiment. Conditions were presented in counter-balanced order. One of the participants was excluded due to a hardware failure during the test.

4.4.1 Results

Response time performance for the three conditions is plotted in Figure 4.9.

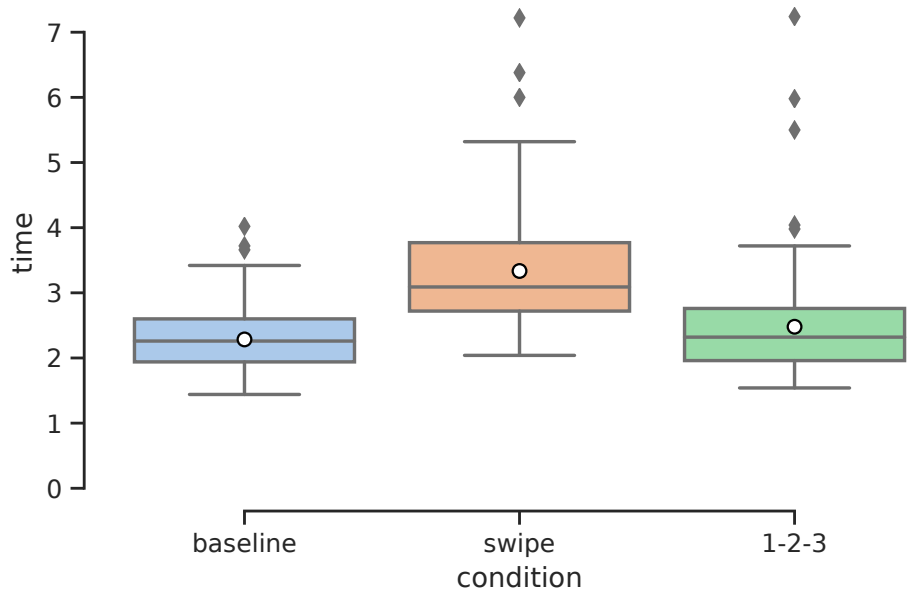

Fig. 4.9: Response time (in seconds) by condition for Experiment 2: distal actuator separation.

Same as the previous experiment, non-parametric tests are used to assess the statistical significance of the results (due to abnormality of their distribution). A statistically signif-

icant difference was found between the experimental conditions resulting from a Friedman test (Table 4.2). A pairwise Wilcoxon test confirmed that the differences between the *swipe* condition and the other conditions were statistically significant ($p < 0.001$) with a medium effect size (Hedges' g close to 0.5 for *1-2-3* and *swipe*, closer to 0.8 for *baseline* and *swipe*)

While response time under the *1-2-3* condition was similar to that of the baseline, response accuracy, as seen in Figure 4.10, suffered in comparison.

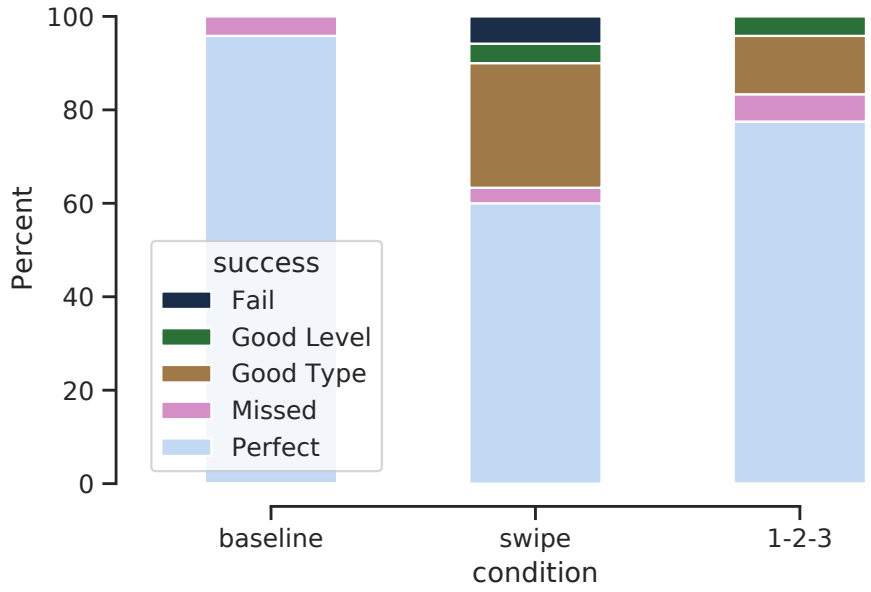

Fig. 4.10: Accuracy of first response by condition for Experiment 2: distal actuator separation.

Using this distal haptic rendering method, participants correctly recognized an average of 78 % of alarms on their first attempt, missing 6 %, and making an error on type and/or vital sign level on the remaining 16 %. Missed responses are considered to be those for which the participant does not respond before the second representation of the alarm, whereas failures are those for which the participant errs in identification of *both* the type and level of the vital sign. A non-parametric Friedman test of differences among the success rate (percentage of correct answer) was conducted (Table 4.2). A pairwise Wilcoxon test confirmed statistical difference between the baseline and the two other conditions ($p < 0.05$)

with large effect size (Hedges' g of 1.6 between baseline and *1-2-3*, 2 between *swipe* and baseline).

As a means of assessing the combination of response time (speed) and correctness of response (accuracy), we plot the Inverse Efficiency Scores (IES), per participant, in Figure 4.11, as proposed by [56]. Lower IES scores correspond to improved efficiency.

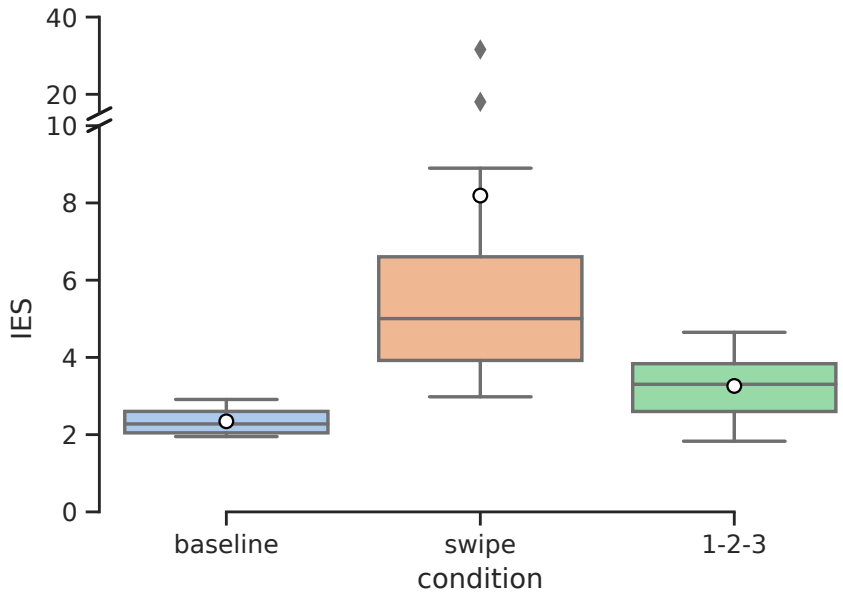

Fig. 4.11: IES scores by condition for Experiment 2: distal actuator separation. Note that for two participants with high error rates under the swipe condition, the calculated IES values were well above the 1.5 interquartile range.

We calculated the IES using median, rather than mean values of response time and accuracy, to minimize the possible effects of outliers. However, the results were similar in either case.

It bears mention that the IES is intended for scenarios in which a strong positive, linear correlation between response time and the proportion of errors is evidenced [57], which, from comparison of Figures 4.9 and 4.10, is not necessarily the case. Furthermore, the IES has an exaggerated, non-linear effect as error rate increases. Thus, the results shown here should be seen as another means of visualization, rather than one that serves as a scientifically sound basis for comparison.

	Source	ddof1	Q	p-unc
First RT	condition	2	15.5	0.000431
Success rate	condition	2	14.9	0.00058
IES	condition	2	15.2	0.000509

Table 4.2: Friedman test for Experiment 2: distal actuator separation.

Analysis of performance on the distractor task did not show any statistically significant effect of condition. However, the trend suggests a smaller mean score for the *swipe* condition (Figure 4.12), consistent with the results observed for the previous experiment.

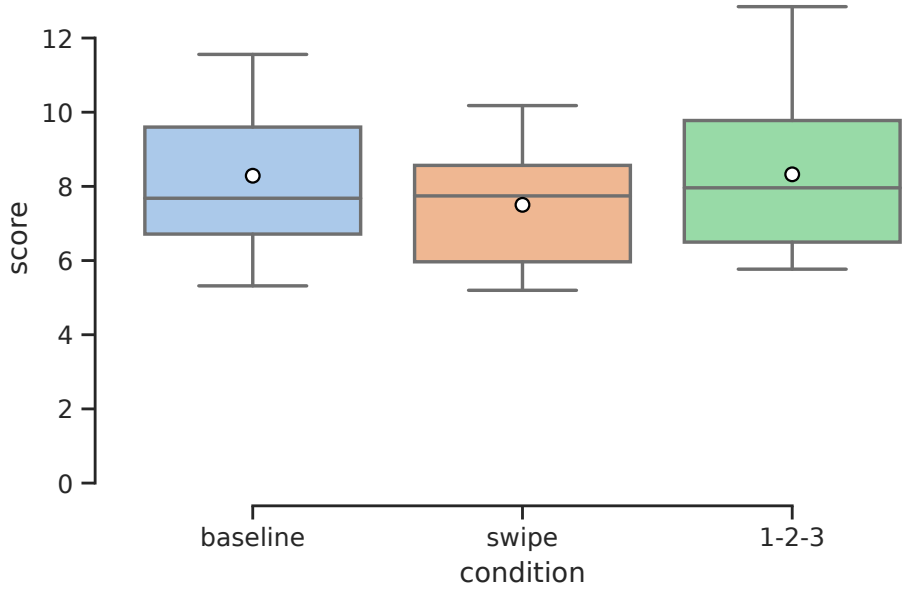

Fig. 4.12: Performance on the Fitts' law distractor task by condition for Experiment 2: distal actuator separation.

4.4.2 Analysis

This demonstrated that adequate performance with a haptics-only information display could be obtained. Despite the increase in error rate, these results were encouragingly competitive with that of a conventional hospital arrangement, which employs auditory alarms in conjunction with a graphical information display. However, the distributed placement of actuators over the arms and one leg is problematic from a clinical application perspective.

4.5 Experiment 3: Increased Actuator Separation on Leg

We now sought to determine whether the limitations of the initial "single-leg" actuator placement could be improved. Specifically, we considered the possibility that further spatial separation of the actuators over the leg would improve discrimination ability, allowing for comparable performance to be achieved to the less palatable distal separation of actuators over the arms and leg. In this regard, previous research found that localization of the tactile stimulus is improved when delivered to locations close to joints [58]. This motivated an increased separation between the actuators on the leg, to new positions as shown in Figure 4.13 (right), each near a different point of mobility: ankle, knee and hip.

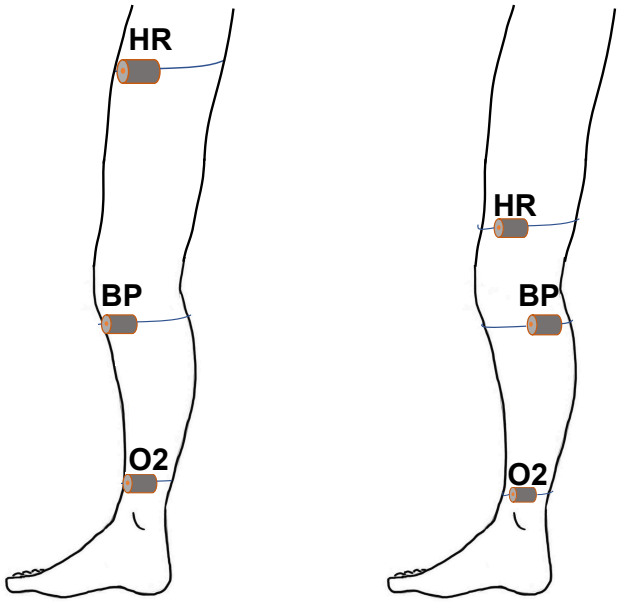

Fig. 4.13: Initial spatial separation of the vibrating bands on the leg as used in Experiment 1 (left) and revised placement as used in Experiment 3 (right). The image of the leg is drawn by an anonymous McGill student.

This experiment was also used to consider performance with a simpler alarm representation, consisting of a series of short pulses of 200 ms separated by gaps of 100 ms for a high vital sign or long pulses of 1 s separated by a gap of 500 ms for a low vital sign. The vibration continued pulsing until the participant responded with the correct selection to acknowledge the alarm event.

We expected that this would require less cognitive effort to parse than our previous representations, involving spatio-temporal swipe patterns or sequences of two or three pulses, to represent low and high states, respectively. At the same time, we considered dispensing with a representation for the "normal" vital sign level, restricting the rendering exclusively to alarm events (*haptic alarm-only*), more consistent with other alarm systems. Although there are strong arguments to be made for the benefit of periodic delivery of a signal indicating that each vital sign is within an acceptable range, we opted to focus the final experiment purely on perception and response performance to alarms. We thus included the baseline condition and three different haptic conditions, with actuators placed only on

the leg, for comparison:

- *baseline:* graphical display with auditory alarms
- *swipe:* haptic leg swipe patterns as used in previous experiments
- *1-2-3:* pulses, as previous, but delivered to the leg
- *haptic alarm-only:* rate-based patterns indicating low or high vital signs, as shown in Table 4.1

The test was conducted with 14 participants, 5M/9F, ages 18-30, $\bar{x}$=24, who had not participated in either of the previous experiments. Conditions were presented using Latin squares ordering. Two of the participants were excluded from analysis due to a failure in logging the data as well as non-compliance with the experimental instructions

4.5.1 Results

Response time performance for the three conditions is plotted in Figure 4.14.

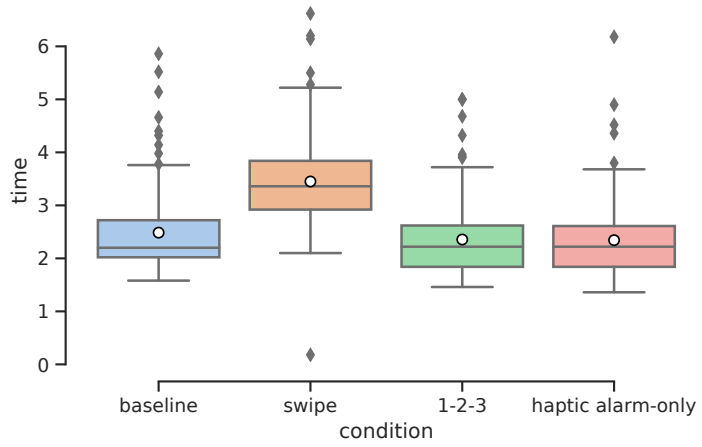

Fig. 4.14: Response time (in seconds) by condition for Experiment 3: increased actuator separation on leg.

Similar to the two previous experiments, we used non-parametric tests to assess the statistical significance since the results do not follow a normal distribution. A Friedman test found a statistically significant difference between the experimental conditions (table 4.3).

A pairwise Wilcoxon test, with Holm correction, confirmed a significant difference between the *swipe* condition and the other three conditions ($p < 0.001$, $W = 1569$ for *baseline*, $W = 765$ for *1-2-3*, $W = 545$ for *haptic alarm-only*). A significant difference was also found between the baseline and the *haptic alarm-only* condition ($p < 0.05$, $W = 2387$) with a small effect size (Hedges' g of 0.4). Equally important, the improved response time for these last two conditions did not come at the expense of a high error rate, as seen in Figure 4.15.

Fig. 4.15: Accuracy of first response by condition. Missed responses are considered to be those for which the participant does not respond before the second representation of the alarm, whereas failures are those for which the participant errs in identification of *both* the type and level of the vital sign.

A Friedman test of the response accuracy (Table 4.3) and a following pairwise Wilcoxon test, with Holm correction, confirmed that the *swipe* condition was significantly different from the baseline, with $p < 0.05$ and $W = 2$. However, the test found differences between the baseline and the other conditions of *1-2-3* and *haptic alarm-only*, with $p > 0.05$, for which we cannot draw conclusions.

A two-one-sided t-test (TOST) procedure applied to the response accuracy results, confirmed a within 15 % equivalence between the baseline and *1-2-3* condition ($p < 0.05$,

$dof = 11$). The equivalence test between *haptic alarm-only* condition and the baseline was not conclusive.

Results of combining response time and accuracy using the IES are shown in Figure 4.16.

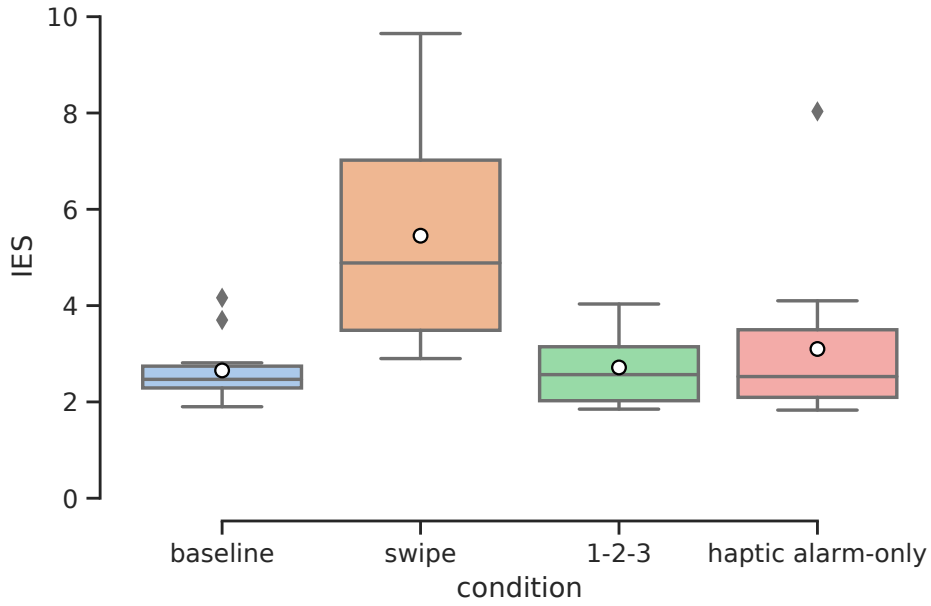

Fig. 4.16: IES scores by condition for Experiment 3: increased actuator separation on leg.

Lower IES scores correspond to improved efficiency. The results of a Friedman test of the response accuracy, shown in Table 4.3, followed by a pairwise Wilcoxon test, indicated that the IES values of the *swipe* condition were significantly different from the other conditions, with $p < 0.05$.

A TOST procedure on the IES scores with equivalence bounds of 0.5 s suggests equivalence between the baseline and the *1-2-3* and *haptic alarm-only* conditions ($p < 0.05$, $dof = 11$). The overall scores for the distractor task are computed as previously, with the results shown in Figure 4.17.

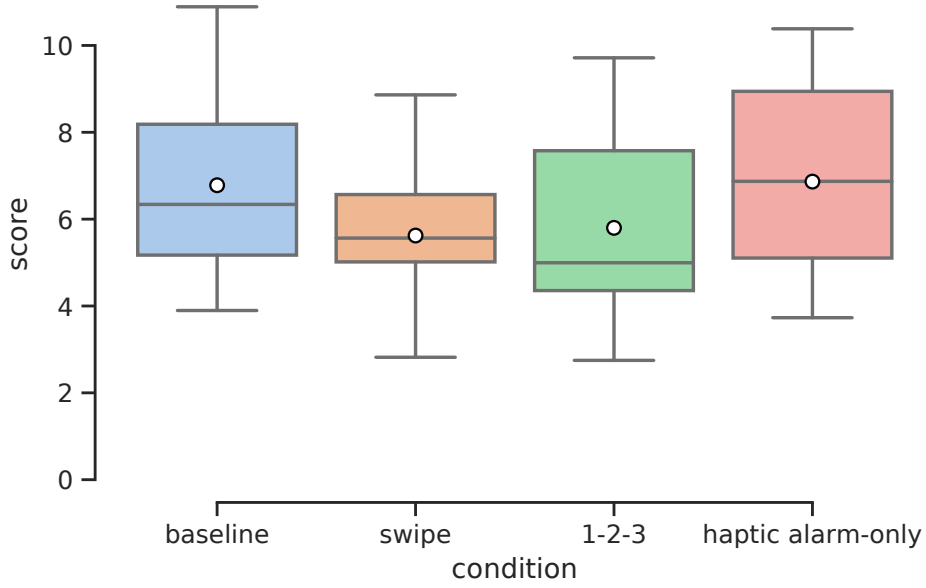

Fig. 4.17: Performance on the Fitts' law distractor task by condition for Experiment 3: increased actuator separation on leg.

For all conditions, the distribution of the scores likely follows a normal distribution, as checked with both QQplots and a Shapiro-Wilk test. A one-way repeated measures ANOVA found a statistically significant difference among the experimental conditions, ($F(3, 33) = 3.8$, $p = 0.02$, with an achieved power of 0.92). A pairwise T-test shows a statistically significant difference between the *haptic alarm-only* and *swipe* conditions ($p < 0.05$) with a medium effect size (Hedges' g close to 0.5) but no significant difference between the other conditions.

4.5.2 Analysis

The improved response time for the *1-2-3* and *haptic alarm-only* conditions, relative both to the baseline, and those obtained in the initial experiment, appears to have been the result of the greater separation of actuators on the leg. Interestingly, with regard to the IES results, these demonstrate superiority of the *1-2-3* condition over the baseline, suggesting not only that vibrotactile rendering of alarm events is feasible, but that it may even support improved clinical performance in recognizing and responding to these events. Moreover,

performance on the distractor task suggests, encouragingly, that the vibrotactile rendering did not impose a greater workload on participants than the baseline condition.

	Source	ddof1	Q	p-unc
RT	condition	3	26.4	7.86e-06
Success rate	condition	3	10.3	0.0159
IES	condition	3	17.6	0.000532

Table 4.3: Friedman test results for Experiment 3: increased actuator separation on leg.

We were also interested in understanding the correlation between the speed and accuracy in participants' response to the alarms. Specially since, based on many discussions we had with our clinical collaborator, we realized taking some seconds in attending to patient's alarm has less hazardous implications than getting the alarm wrong. A cursory analysis suggests that there is a negative correlation between the speed and accuracy of responses in all conditions apart from haptic-alarm-only. This shows that that participants who were more confident in their perception of the stimuli responded faster. Further research and analysis is needed to balance between these two parameters while removing the effects of bias on participants if guided in advantage of either parameter.

4.6 Discussion

The results of our experiments, in particular, Experiment 3, offer strong support for the possibility of reducing clinical dependence on auditory alarms and graphical displays. However, several important caveats regarding the experimental results are in order. First, given that our participant pool was drawn from the general university population rather than trained clinicians, their level of experience attending to patient monitors and medical alarms is not necessarily reflective of realistic hospital conditions. Accordingly, their skills in monitoring of vital sign levels in parallel with other workload is likely to be inferior, which may have impacted either or both response time and recognition accuracy, in particular across conditions. These factors serve as an advisory note for our intended future experimentation with clinician participants.

On the other hand, we biased our experiment against the vibrotactile condition by the choice of a visual information display that was trivial to interpret. The information display on actual hospital monitors is dramatically more complex and demanding of clinician visual search and interpretation. We should note, however, that these monitors could help clinicians resolve any uncertainty of parsing the information provided via the vibrotactile display. Had the graphical monitor been included in the vibrotactile conditions, as would be more reflective of an actual hospital environment, this could well have reduced the incidence of incorrect vital sign or level identification observed in our experiments.

From a practical deployment perspective, another contribution of our study is in demonstrating the efficacy of a haptic display comprising a reduced number (3) of actuators for conveying similar information as used in previous studies (17) for example implemented in a haptic vest [39].

While our intent is to deliver the necessary information via haptics only, thereby reducing demands on visual attention, the co-presence of these monitors would be likely to further reduce the error rates seen, for example, in Figure 4.15. It should also be mentioned that since training was minimal, the results are likely to improve with time.

For the present stage of this research, we made a modest compromise by distributing the hardware over different points on the body. While it is likely that this could be improved over time, we believe that the results obtained thus far will motivate further efforts to develop an optimized hardware package that is better suited to the intended use case.

4.7 Conclusions

The results of our studies demonstrate strong promise for reducing the highly disturbing noise of audible alarms in hospital OR and ICU environments. We found that the delivery of alarm information through vibrotactile actuators, instead, can be similarly effective as audio alarms in conjunction with graphical patient monitors in terms of both response time and accuracy, and without significant impact on performance on a parallel, attention-demanding task. It should be emphasized that this comparison was performed against a very simple graphical patient monitor, which required essentially no interpretation of displayed information. We anticipate that our approach may have attentional-sparing benefits across sensory modalities in cognitively demanding environments.

Further studies should be conducted to determine the practical effects of such haptic alarm information delivery mechanisms in conjunction with patient monitors, and under competing task loads more closely replicating actual hospital conditions. Nevertheless, it should be emphasized that the consequences of alarm fatigue will remain until the disturbingly high rate of false positive alarms, and not only the sheer number of alarms, is also reduced.

Future work should continue not only to tackle this serious problem, but also, to improve the robustness of vital sign information delivery via vibrotactile cues, and to reduce the hardware requirements for doing so with the goals of improving patient safety and monitoring.

Chapter 5

Conclusions and future work

5.1 Conclusion

Just like work in other complex data-rich domains, the task set of medical staff is characterized by high mental workload. This cognitive burden, resulting from information overload presented primarily through auditory channels, saturates the hospitals' soundspace. The noise in hospital settings, particularly from patient monitor alarms, is harmful both to the health of clinicians and patients. These detrimental effects, as detailed in the introduction of this thesis, motivated our proposal to improve hospital environments, specifically patient monitoring methods.

One way to approach the issue of noise in hospitals is by implementing a vibrotactile alarm system for communicating patient physiological information. The research described in this thesis assessed the effectiveness of multiple vibrotactile rendering conditions by comparing them to a baseline audio alarm with graphical patient monitor display, discussed in detail in Chapter 4. A distractor activity was used to simulate competing task demands in the clinical environment. We found that, with sufficient anatomical separation of the actuators, certain vibrotactile information rendering strategies demonstrated performance that was not significantly different from that of the baseline condition, both in response time and accuracy. Our findings suggest the possibility of improving high-impact healthcare environments by replacing disturbing audio alarms with vibrotactile information delivery to clinicians.

Although our results were promising, we recognize that our study may be improved by greater participant training and testing on real users, e.g., medical staff population. Our participants had little or no previous experience in receiving information from vibrotactile feedback in a wearable manner. As a result, they often felt overwhelmed and somewhat uncomfortable with the novel haptic signals delivered to their leg. This could be mitigated by including longer training times, possibly over the course of several days or more before conducting the test. Even though unfamiliarity with vibration signals might remain an issue for the majority of our intended clinician population, it would be feasible to fully familiarize them with this feedback as part of their rigorous clinical training. Additionally, extended training with the vibration patterns could help determine whether greater familiarity with the vibration displays improves one's ability to habituate to the signal, allowing it to more effectively support peripheral awareness. Another suggested improvement to ensure ecological validity of the haptic alarm system is to test on anesthesiologists and surgeons while they perform a simulated medical task. An experimental design that reflects a clinical task may generate results that are more predictive of performance in the real clinical setting.

5.2 Future Work

Recognizing the potential of implementing haptic alarms in hospitals, we now offer suggestions for future research. In our study, we observed that when changes to the location of vibrating bands and their pattern were made, participants' performance in receiving information from vibrotactile cues improved. Thus, we suggest further research may explore the effects of altering the vibrotactile pattern designs while investigating different vibration band placements to enhance our proposed haptic alarm design. Such work may focus on coordinating physiological parameters in real-time with vibration signal characteristics in a meaningful and intuitive manner (e.g., an alarm that monitors a patient's heart and feels like their heartbeat). A second suggestion for future research is to consider a single actuator for the delivery of patients' vital information. This direction of exploration is currently ongoing in our group. If future research were to further explore vibration pattern design as well as employing less bulky hardware, an advanced physiological display design can be accomplished. This can offer the potential of reducing clinicians' reliance on loud auditory alarms, therefore achieving quieter and safer critical care environments.

The efforts described here, and future efforts in the development of haptic and audio/haptic multimodal displays, show promise for improving clinical performance in both monitoring alarm events as well as conducting the primary medical procedure. Through continued efforts such as these, we can realize an advancement in critical care environments, that benefits clinicians in their awareness of patients and, more importantly, improves patient health outcomes.

References

[1] B. Berglund, T. Lindvall, and D. H. Schwela, *Guidelines for community noise*. World Health Organization, 1995.

[2] M. H. Otenio, E. Cremer, and E. M. T. Claro, "Noise level in a 222 bed hospital in the 18th health region-pr," *Revista Brasileira de Otorrinolaringologia*, vol. 73, no. 2, pp. 245–250, 2007.

[3] J. L. Darbyshire and J. D. Young, "An investigation of sound levels on intensive care units with reference to the who guidelines," *Critical Care*, vol. 17, no. 5, p. R187, 2013.

[4] I. J. Busch-Vishniac, J. E. West, C. Barnhill, T. Hunter, D. Orellana, and R. Chivukula, "Noise levels in johns hopkins hospital," *The Journal of the Acoustical Society of America*, vol. 118, no. 6, pp. 3629–3645, 2005.

[5] T. O'carroll, "Survey of alarms in an intensive therapy unit," *Anaesthesia*, vol. 41, no. 7, pp. 742–744, 1986.

[6] C. Atzema, M. J. Schull, B. Borgundvaag, G. R. Slaughter, and C. K. Lee, "Alarmed: adverse events in low-risk patients with chest pain receiving continuous electrocardiographic monitoring in the emergency department. a pilot study," *The American journal of emergency medicine*, vol. 24, no. 1, pp. 62–67, 2006.

[7] S. Siebig, S. Kuhls, M. Imhoff, U. Gather, J. Schölmerich, and C. E. Wrede, "Intensive care unit alarmshow many do we need?," *Critical care medicine*, vol. 38, no. 2, pp. 451–456, 2010.

[8] K. Momtahan, R. Hetu, and B. Tansley, "Audibility and identification of auditory alarms in the operating room and intensive care unit," *Ergonomics*, vol. 36, no. 10, pp. 1159–1176, 1993.

[9] C. Force, "Impact of clinical alarms on patient safety: a report from the american college of clinical engineering healthcare technology foundation," *Journal of Clinical Engineering*, vol. 32, pp. 22–33, 2007.

[10] S. T. Lawless, "Crying wolf: false alarms in a pediatric intensive care unit.," *Critical care medicine*, vol. 22, no. 6, pp. 981–985, 1994.

[11] L. B. Talley, J. Hooper, B. Jacobs, C. Guzzetta, R. McCarter, A. Sill, S. Cain, and S. L. Wilson, "Cardiopulmonary monitors and clinically significant events in critically ill children," *Biomedical Instrumentation & Technology*, vol. 45, no. s1, pp. 38–45, 2011.

[12] C. P. Bonafide, R. Lin, M. Zander, C. S. Graham, C. W. Paine, W. Rock, A. Rich, K. E. Roberts, M. Fortino, V. M. Nadkarni, *et al.*, "Association between exposure to nonactionable physiologic monitor alarms and response time in a children's hospital," *Journal of hospital medicine*, vol. 10, no. 6, pp. 345–351, 2015.

[13] J. P. BLISS, R. D. GILSON, and J. E. DEATON, "Human probability matching behaviour in response to alarms of varying reliability," *Ergonomics*, vol. 38, no. 11, pp. 2300–2312, 1995. PMID: 7498189.

[14] S. Sendelbach, "Alarm fatigue," *Nursing Clinics*, vol. 47, no. 3, pp. 375–382, 2012.

[15] M. Cvach, "Monitor alarm fatigue: an integrative review," *Biomedical instrumentation & technology*, vol. 46, no. 4, pp. 268–277, 2012.

[16] M. S. Kristensen, J. Edworthy, and S. Denham, "Alarm fatigue in the perception of medical soundscapes," *European Congress and Exposition on Noise Control Engineering*, pp. 745–750, 2015.

[17] P. Lacherez, E. Limin Seah, and P. Sanderson, "Overlapping melodic alarms are almost indiscriminable," *Human factors*, vol. 49, no. 4, pp. 637–645, 2007.

[18] B. Hasanain, A. D. Boyd, J. Edworthy, and M. L. Bolton, "A formal approach to discovering simultaneous additive masking between auditory medical alarms," *Applied Ergonomics*, vol. 58, pp. 500–514, 2017.

[19] L. M. Stanford, J. W. R. McIntyre, and J. T. Hogan, "Audible alarm signals for anaesthesia monitoring equipment," *International journal of clinical monitoring and computing*, vol. 1, p. 251, Dec 1985.

[20] S. Sendelbach and M. Funk, "Alarm fatigue a patient safety concern," *AACN advanced critical care*, vol. 24, no. 4, pp. 378–386, 2013.

[21] S. Ouimet, B. P. Kavanagh, S. B. Gottfried, and Y. Skrobik, "Incidence, risk factors and consequences of icu delirium," *Intensive care medicine*, vol. 33, no. 1, pp. 66–73, 2007.

[22] E. W. Ely, A. Shintani, B. Truman, T. Speroff, S. M. Gordon, F. E. Harrell Jr, S. K. Inouye, G. R. Bernard, and R. S. Dittus, "Delirium as a predictor of mortality in mechanically ventilated patients in the intensive care unit," *Jama*, vol. 291, no. 14, pp. 1753–1762, 2004.

[23] J. C. Jackson, R. P. Hart, S. M. Gordon, A. Shintani, B. Truman, L. May, and E. W. Ely, "Six-month neuropsychological outcome of medical intensive care unit patients," *Critical care medicine*, vol. 31, no. 4, pp. 1226–1234, 2003.

[24] R. O. Hopkins and J. C. Jackson, "Long-term neurocognitive function after critical illness," *Chest*, vol. 130, no. 3, pp. 869–878, 2006.

[25] M. Basner, W. Babisch, A. Davis, M. Brink, C. Clark, S. Janssen, and S. Stansfeld, "Auditory and non-auditory effects of noise on health," *The lancet*, vol. 383, no. 9925, pp. 1325–1332, 2014.

[26] M. Thomas, H. Sing, G. Belenky, H. Holcomb, H. Mayberg, R. Dannals, H. Wagner JR, D. Thorne, K. Popp, L. Rowland, *et al.*, "Neural basis of alertness and cognitive performance impairments during sleepiness. i. effects of 24 h of sleep deprivation on waking human regional brain activity," *Journal of sleep research*, vol. 9, no. 4, pp. 335–352, 2000.

[27] N. S. Freedman, J. Gazendam, L. Levan, A. I. Pack, and R. J. Schwab, "Abnormal sleep/wake cycles and the effect of environmental noise on sleep disruption in the intensive care unit," *American journal of respiratory and critical care medicine*, vol. 163, no. 2, pp. 451–457, 2001.

[28] B. B. Kamdar, D. M. Needham, and N. A. Collop, "Sleep deprivation in critical illness: its role in physical and psychological recovery," *Journal of intensive care medicine*, vol. 27, no. 2, pp. 97–111, 2012.

[29] M. A. Meredith and B. E. Stein, "Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration," *Journal of neurophysiology*, vol. 56, no. 3, pp. 640–662, 1986.

[30] T. Noesselt, S. Tyll, C. N. Boehler, E. Budinger, H.-J. Heinze, and J. Driver, "Sound-induced enhancement of low-intensity vision: multisensory influences on human sensory-specific cortices and thalamic bodies relate to perceptual enhancement of visual detection sensitivity," *Journal of Neuroscience*, vol. 30, no. 41, pp. 13609–13623, 2010.

[31] H. S. Vitense, J. A. Jacko, and V. K. Emery, "Multimodal feedback: an assessment of performance and mental workload," *Ergonomics*, vol. 46, no. 1-3, pp. 68–87, 2003.

[32] J. Ng and J. Man, "Vibro-monitor: A vibrotactile display for physiological data monitoring," in *Human Interface Technologies Conference*, 2004.

[33] C. L. Baldwin, C. Spence, J. P. Bliss, J. C. Brill, M. S. Wogalter, C. B. Mayhorn, and T. K. Ferris, "Multimodal cueing: The relative benefits of the auditory, visual, and tactile channels in complex environments," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 56, pp. 1431–1435, SAGE Publications Sage CA: Los Angeles, CA, 2012.

[34] P. Sanderson, A. Wee, and P. Lacherez, "Learnability and discriminability of melodic medical equipment alarms," *Anaesthesia*, vol. 61, no. 2, pp. 142–147, 2006.

[35] Y. Visell, B. L. Giordano, G. Millet, and J. R. Cooperstock, "Vibration influences haptic perception of surface compliance during walking," *PLoS one*, vol. 6, no. 3, p. e17697, 2011.

[36] J. Y. Ng, J. C. Man, S. Fels, G. Dumont, and J. M. Ansermino, "An evaluation of a vibro-tactile display prototype for physiological monitoring," *Anesthesia & Analgesia*, vol. 101, no. 6, pp. 1719–1724, 2005.

[37] V. Cobus, B. Ehrhardt, S. Boll, and W. Heuten, "Vibrotactile alarm display for critical care," in *Proceedings of the 7th ACM International Symposium on Pervasive Displays*, p. 11, ACM, 2018.

[38] S. Ford, J. Daniels, J. Lim, V. Koval, G. Dumont, S. K. Schwarz, and J. M. Ansermino, "A novel vibrotactile display to improve the performance of anesthesiologists in a simulated critical incident," *Anesthesia & Analgesia*, vol. 106, no. 4, pp. 1182–1188, 2008.

[39] T. K. Ferris, A. Texas, C. Station, N. Sarter, and A. Arbor, "Continuously Informing Vibrotactile Displays in Support of Attention Management and Multitasking in Anesthesiology," 2011.

[40] P. Alirezaee, R. Girgis, T. Kim, J. J. Schlesinger, and J. R. Cooperstock, "Did you feel that? developing novel multimodal alarms for high consequence clinical environments," Georgia Institute of Technology, 2017.

[41] S. A. Gelfand and H. Levitt, *Hearing: An introduction to psychological and physiological acoustics*, vol. 4. Marcel Dekker New York, 1998.

[42] M. Taylor and C. D. Creelman, "Pest: Efficient estimates on probability functions," *The Journal of the Acoustical Society of America*, vol. 41, no. 4A, pp. 782–787, 1967.

[43] M. Taylor, S. Forbes, and C. D. Creelman, "Pest reduces bias in forced choice psychophysics," *The Journal of the Acoustical Society of America*, vol. 74, no. 5, pp. 1367–1374, 1983.

[44] T. N. Cornsweet, "The staircase-method in psychophysics," *The American journal of psychology*, vol. 75, no. 3, pp. 485–491, 1962.

[45] H.-Y. Yao and V. Hayward, "Design and analysis of a recoil-type vibrotactile transducer," *The Journal of the Acoustical Society of America*, vol. 128, no. 2, pp. 619–627, 2010.

[46] D. M. Gay-Betton, P. Alirezaee, J. R. Cooperstock, and J. J. Schlesinger, "Haptic: Haptic anatomical positioning to improve clinical monitoring," in *Proceedings of the Conference on Design and Semantics of Form and Movement-Sense and Sensitivity, DeSForM 2017*, InTech, 2017.

[47] A. Chan, K. MacLean, and J. McGrenere, "Learning and identifying haptic icons under workload," in *Eurohaptics Conference, 2005 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2005. World Haptics 2005. First Joint*, pp. 432–439, IEEE, 2005.

[48] P. M. Fitts, "The information capacity of the human motor system in controlling the amplitude of movement.," *Journal of experimental psychology*, vol. 47, no. 6, p. 381, 1954.

[49] S. Brewster and L. M. Brown, "Tactons: structured tactile messages for non-visual information display," in *Proceedings of the fifth conference on Australasian user interface*, vol. 28, pp. 15–23, Australian Computer Society, Inc., 2004.

[50] L. M. Brown, S. A. Brewster, and H. C. Purchase, "A first investigation into the effectiveness of tactons," 2005.

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

[52] P. Barralon, G. Ng, G. Dumont, S. K. Schwarz, and M. Ansermino, "Development and evaluation of multidimensional tactons for a wearable tactile display," in *Proceedings of the 9th international conference on Human computer interaction with mobile devices and services*, pp. 186–189, ACM, 2007.

[53] R. Vallat, "Pingouin: statistics in python," *The Journal of Open Source Software*, vol. 3, p. 1026, Nov. 2018.

[54] F. Pece, J. J. Zarate, V. Vechev, N. Besse, O. Gudozhnik, H. Shea, and O. Hilliges, "Magtics: Flexible and thin form factor magnetic actuators for dynamic and wearable haptic feedback," in *Proceedings of the 30th annual ACM symposium on User interface software and technology*, pp. 143–154, ACM, 2017.

[55] L. Post, I. Zompa, and C. Chapman, "Perception of vibrotactile stimuli during motor activity in human subjects," *Experimental brain research*, vol. 100, no. 1, pp. 107–120, 1994.

[56] J. Townsend and F. Ashby, *Methods of modeling capacity in simple processing systems*, vol. 3, pp. 199–239. Jan. 1978.

[57] R. Bruyer and M. Brysbaert, "Combining speed and accuracy in cognitive psychology: Is the inverse efficiency score (ies) a better dependent variable than the mean reaction time (rt) and the percentage of errors (pe)?," *Psychologica Belgica*, vol. 51, no. 1, pp. 5–13, 2011.

[58] 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.