

Getting Your Hands Dirty Outside the Lab: A Practical Primer for Conducting Wearable Vibrotactile Haptics Research

Jeffrey R. Blum¹, Pascal E. Fortin¹, Feras Al Taha¹, Parisa Alirezaee¹, Marc Demers¹,
Antoine Weill-Duflos¹, and Jeremy R. Cooperstock¹, *Member, IEEE*

Abstract—As haptics have become an ingrained part of our wearable experience, particularly through phones, smartwatches, and fitness trackers, significant research effort has been conducted to find new ways of using wearable haptics to convey information, especially while we are on-the-go. In this paper, instead of focusing on aspects of haptic information design, such as tacton encoding methods, actuators, and technical fabrication of devices, we address the more general recurring issues and “gotchas” that arise when moving from core haptic perceptual studies and in-lab wearable experiments to real world testing of wearable vibrotactile haptic systems. We summarize key issues for practitioners to take into account when designing and carrying out in-the-wild wearable haptic user studies, as well as for user studies in a lab environment that seek to simulate real-world conditions. We include not only examples from published work and commercial sources, but also hard-won illustrative examples derived from issues and failures from our own haptic studies. By providing a broad-based, accessible overview of recurring issues, we expect that both novice and experienced haptic researchers will find suggestions that will improve their own mobile wearable haptic studies.

Index Terms—Wearable haptics, haptic assessment, tactile perception, vibrotactile, in-the-wild studies.

I. INTRODUCTION

RUNNING in-the-wild user studies is challenging due to the many considerations that go into making sure that the system being tested functions properly outside of a laboratory environment, and that the results are reliable. Mobile haptics raises many unique issues that are non-existent with other

modalities, such as the need to couple the actuators closely to the skin in order to better control how the haptic effects will be perceived. The present article is motivated by the desire to provide insight into the issues encountered while conducting wearable haptic research, in particular, for the challenging scenarios associated with in-the-wild studies. Although “in-the-wild” can be used to describe a variety of experiments outside the lab, we primarily focus on studies taking place in the context of participants’ daily lives, including for multiple days or weeks. This may also include using heterogeneous haptic hardware that participants already have, such as when conducting research in the large [1], as well as when taking into account not just strict performance measures, but also how the haptic device fits participants’ more subjective preferences, such as willingness to wear the device in public. We focus on vibrotactile actuators, arguably the most widespread class of haptic devices, found in a variety of consumer electronics devices and used in many research studies. Even though much of the material may apply to other haptic stimuli such as temperature, stroking, squeezing, or ultrasonics, we limit our discussion to only a few comments outside of our vibrotactile focus, such as when lessons concerning vibrotactile haptics may be particularly inapplicable.

In this article, we do not rehash all of the motivations and issues with running in-the-wild studies in general, since this has been well covered in prior literature. Briefly, running studies outside the lab, especially over extended periods of time, introduces issues of controlling confounding variables and gathering valid data from participants. Kjeldskov *et al.* question whether this is worth the effort, noting that there may be limited benefit in uncovering usability issues that cannot be found in a more controlled laboratory study [2]. However, Kjeldskov and Skov revisited this assessment a decade later, determining through a review of more recent literature on the subject that there are benefits to such studies beyond simply uncovering basic usability issues. In the end, they “suggest moving beyond usability evaluations, and to engage with field studies that are truly in-the-wild, and longitudinal” [3]. Such in-situ, often longer-term, field studies are precisely what we address in this work, with the challenges of integrating haptics into participants’ real-life constraints. By focusing on the issues unique to running *haptic* studies outside the laboratory environment, or in some cases, in-lab haptic studies that seek

Manuscript received November 21, 2018; revised May 22, 2019; accepted July 3, 2019. Date of publication July 23, 2019; date of current version September 16, 2019. This paper was recommended for publication by Associate Editors Domenico Prattichizzo, Miguel Á. Otaduy, Claudio Pacchierotti, and Hiroyuki Kajimoto upon evaluation of the reviewers comments. (*Corresponding author: Jeffrey R. Blum.*)

J. R. Blum, P. E. Fortin, P. Alirezaee, A. Weill-Duflos, and J. R. Cooperstock are with the Centre for Interdisciplinary Research in Music Media and Technology, McGill University, Montréal, QC H3A 1E3, Canada, and also with the Department of Electrical and Computer Engineering, McGill University, Montréal, QC H3A 1E3, Canada (e-mail: jeffrey.blum@mail.mcgill.ca; pefortin@cim.mcgill.ca; parisa.alirezaee@mail.mcgill.ca; antoinew@cim.mcgill.ca; jer@cim.mcgill.ca).

F. Al Taha and M. Demers are with the Department of Electrical and Computer Engineering, McGill University, Montréal, QC H3A 1E3, Canada (e-mail: feras.altaha@mail.mcgill.ca; marc@cim.mcgill.ca).

Digital Object Identifier 10.1109/TOH.2019.2930608

1939-1412 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

to maximize their applicability to real-world situations, we offer a primer for new and experienced haptic designers. We aim to create an article containing the more practical information that many haptics practitioners may have wished was conveyed to them in a single package, rather than learned largely through their own experience.

II. HAPTICS BACKGROUND

Although we expect that experienced haptic practitioners will be largely familiar with the content of this section, we provide pointers to accessible existing literature that forms the background to, and complements, the information in the rest of this article. While not a comprehensive review in and of itself, we touch on key aspects of general wearable haptics and provide a breakdown of wearable haptics goals that help frame the discussion. Before examining more specific works, some readers may also appreciate a broad historical and societal perspective on haptics, as found in Parisi's book "Archaeologies of Touch" [4].

A. General Wearable Haptics Background

Even though we focus primarily on experiments involving vibrotactile haptic transducers, to appreciate the broader context of haptic research in general, we refer the interested reader to a number of papers that cover key background topics applicable to wearable haptic applications. Topics include which actuators to use for stimulation, where on the body stimuli are best provided, and how to best design stimuli given haptic perception limitations.

1) *Haptics All Over the Body*: Skin covers the entirety of the human body, providing a broad surface for rendering haptic effects. Despite our ability to sense vibration anywhere through the skin, perceptual thresholds and discrimination abilities differ depending on the specific location. It should be noted that some of the most sensitive areas are generally impractical for a wearable system. The lips and tongue [5], for example, are significantly more sensitive than thighs and shoulders [6], but are rarely used for delivery of haptic stimuli.

Two articles provide an extensive review specifically of hand-based haptic solutions, one by Pacchierotti *et al.* [7], which focuses on wearable haptic systems for the fingertips and hand, and a second by Culbertson *et al.* [8], which offers pointers for the generation of touch sensations and an overview of haptic devices and actuators. Karuei *et al.* explored a broader range of body locations, running an in-lab study comparing the perception of vibrotactile stimuli at seven different body locations, concluding that "Wrists and spine are generally best for detecting vibrations, and are also the most preferred, with arms next in line. Feet and thighs are poor candidates for vibrotactile displays, exhibiting the worst detection performance of those we tested and ranking lowest in user esteem", and also finding that walking while receiving vibrations significantly reduces perception [9]. Although perception is a critical concern, user preference can also weigh

into how desirable participants may find a system being tested, as discussed in Section IV-B.

2) *Vibrotactile Haptics*: Vibrotactile devices are generally inexpensive, operate at safe low voltages, are robust to environmental noise, and produce perceptible sensations even at very low power, such that they can be embedded into small devices such as smartwatches. Coupling with the body simply requires pressing them against the skin, e.g., with a strap, although the consistency of this coupling, especially between participants and between studies, can be an issue (Section IV-A).

Since this article confines itself to vibrotactile haptics, we recommend that those unfamiliar with devices such as Eccentric Rotating Mass (ERM) and Linear Resonant Actuators (LRA) consult reviews by Choi and Kuchenbecker [10], and Jones and Sarter [11], which cover vibrotactile actuators, how they are perceived, and their applications. Van Erp authored a set of guidelines focused on using vibrotactile displays in Human-Computer Interaction (HCI) applications [12], which provides a psychophysical background for vibrotactile information delivery. For haptic communication design, MacLean provides guidelines for designing tactile messages with minimal attentional load on the user [13]. Wang *et al.* reviews vibrotactile information communication techniques, and surveys performance measurements of response time and accuracy across 24 studies [14].

3) *Non-Vibrotactile Haptics*: In addition to the vibrotactile devices on which this article focuses, other haptic approaches have been explored for wearable use, including temperature [15], electrotactile stimulation [16], Electrical Muscle Stimulation (EMS) [17], pneumatic and hydraulic based compression [18], [19], and even blowing wind on the user's skin [20]. As an accessible overview of common haptic devices that are easy to obtain and use, including those beyond vibrotactile, we recommend the "Do It Yourself Haptics" articles by Hayward and MacLean, which cover the types and construction of mechanical haptic devices (Part I) [21], and designing usable haptic interactions (Part II) [22]. We further recommend Hayward's perception-oriented article on tactile illusions and demonstrations [23].

4) *Multiple Actuators*: A fundamental design decision for haptics experiments is how many actuators are needed. For many tasks, a single haptic transducer suffices, but tasks involving a spatial haptic display generally require more. For in-the-wild studies, additional haptics hardware implies greater power consumption and wiring, with more possible points of failure. As an example of conveying complex vibrotactile messages with multiple actuators, we suggest Novich and Eagleman's paper reporting on their haptic vest [24], which also provides a solid background on related haptic information communication systems, largely using multiple actuators.

When building multi-actuator systems, the minimum spacing between actuators will vary depending on where the stimulus is delivered, since spatial acuity is not consistent across different body locations (Section II-A1). Actuator spacing can be particularly problematic for vibrotactile stimuli because if the actuators are too close together, and the vibration is strong,

the skin itself will conduct the vibrations to nearby mechanoreceptors (i.e., Pacinian corpuscles), making it difficult to localize which actuator in the set is being triggered [25]. It should be noted that simultaneous presentation has different requirements than successive temporal pattern presentation in terms of minimum actuator spacing [26]. For example, Panëels *et al.* found that recognizing static on/off actuation of closely spaced actuators around the wrist was more difficult than when they were sequentially activated to make dynamic patterns of movement, e.g., in one direction or another [27].

B. Goals of Wearable Haptics

We divide wearable haptic systems into three categories, depending on the goal. Roughly, these are ordered from requiring the least haptic fidelity to the most, with implications discussed in Section III-B:

- (1) *Attentional*: provide simple signals aimed at capturing a user's attention, with little to no intrinsic information, e.g., mobile phone notifications
- (2) *Semantic*: convey specific meaning, e.g., tactile icons or "tactons" [28] for purposes such as task progress indicators, directions [28], or numbers [29]
- (3) *Experiential*: provide synthetic haptic feedback to enhance actions such as tapping a screen, turning a dial, or else rendering a uni- or multimodal experience, e.g., simulating a ground surface [30] or enhancing affect [31].

Whereas *attentional* haptics are designed to be interruptive, grabbing the user's attention despite distractions to indicate that an important event has occurred, *experiential haptics* are specifically designed to be purely background support for other activities. *Semantic* haptics, however, can span both foreground and background applications. In the foreground, such systems convey information while they have the user's attention. However, MacLean describes background "ambient" haptics as the "supporting player" providing a low effort, unconscious affective communication channel, in which the user can gradually become aware of the context, maintain awareness of an ongoing situation, and follow up on any changes [32]. Note that some systems may have elements of multiple categories. For example, a tacton may be designed not only to convey a specific piece of information, but also to be jarring enough to grab attention even if the user is distracted, such as Immersion Corporation's "Instinctive Alerts", as mentioned in [4]. Thus, it would have both attentional and semantic elements. Similarly, a forceful screen tap may generate an experiential stronger haptic click, which in addition to telling the user that they have pushed with a certain force, also provides semantic information about a different selection being made.

III. HARDWARE

Laboratory studies can use fragile, custom-built haptic devices with little regard for portability or power constraints. Moving to in-the-wild use, however, requires hardware that remains expressive but is also portable and robust enough to provide reliable data throughout the experiment. The ideal of

solving this with small, high-fidelity, low-power, inexpensive haptic actuators built into robust, easy-to-wear consumer electronics, with all capabilities fully available to developers, has not yet been fully realized, but improvements are being made. Here we discuss examples of limitations in currently available commercial devices, as well as the tradeoffs of actuator expressivity, power, and robustness.

A. Commercial Versus Custom-Built Devices

While the actuator itself defines the core of the haptic experience, how that actuator is physically packaged, and how the developer can access its features, also impact its utility. Even though commercial haptic products can provide off-the-shelf robustness for in-the-wild use, they tend to limit the customizability of, and access to, the underlying haptic hardware. If the application is basic enough, simple control of ERM or LRA actuators in phones and smartwatches can be sufficient, and are an easy path to quickly prototyping and carrying out experiments. In the best case, this lets participants use their own existing phones or wearables during studies, without the need to learn how to use, carry, and maintain additional hardware specific to the experiment, subject to limitations discussed in Section VI-B.

Often, however, the tight constraints of commercial devices mean that exceeding their core intended purpose makes them unsuitable for novel applications. For example, the Taptic Engine in the Apple watch is essentially a very capable, fast response LRA-style vibrotactile actuator. At the time of writing, the soon-to-be-released Core Haptics API will provide rich access to the Taptic Engine on iPhones, but the Apple Watch will continue to be limited to a small set of predefined effects.¹ Even if it was fully accessible, battery life may suffer under heavy haptic use to a degree that would preclude all-day use, since the device was not designed primarily as a haptic tool. Thus, in order to maximize flexibility, it may be necessary to either heavily modify commercial hardware (e.g., by integrating a larger battery), or design and build experimental hardware despite the additional work and drawbacks.

That said, progress is being made. For example, newer Android devices (e.g., Samsung Galaxy S9, OnePlus 7pro, Google Pixel 3) include 256 levels of vibration amplitude control, which broadens the possibility for higher-fidelity haptics beyond simple on/off patterns.

B. Low Versus High Fidelity

Vibrotactile actuators are available with a wide variety of performance and control characteristics. For commercially deployed hardware, cost is often an overriding concern, so simpler, lower fidelity ERM devices are still most typical. Newer devices including Lofelt's L5 (Fig. 1) and Apple's Taptic Engine offer greater expressive capabilities, allowing for delivery of vibrotactile stimuli over a wider frequency range.

¹ <http://developer.apple.com/documentation/watchkit/wkhaptictype>

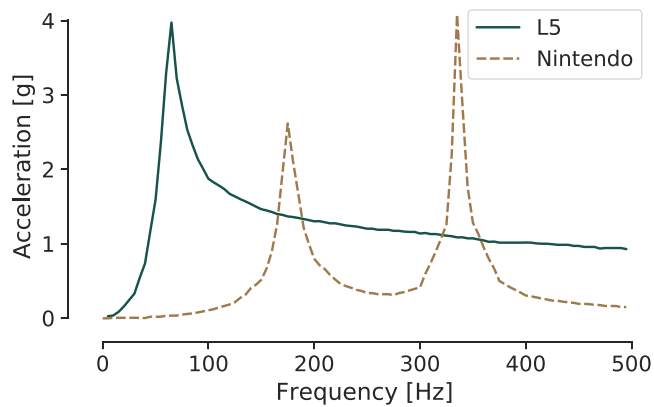


Fig. 1. Comparison of LRA (from Nintendo Switch gamepad) and Lofelt L5 (wideband actuator) acceleration as a function of frequency, at constant power. Figure adapted from Lofelt.² In contrast, ERM acceleration increases linearly with frequency, as applied power is increased.

High-fidelity haptic actuators such as the L5 are more suitable for covering a large part of the 0 Hz to 1000 Hz range of haptically perceptible frequencies [33], [34]. These HiFi devices are mostly being used for rendering textural information [35]. Examples include linear haptic taps while rotating the crown on a smartwatch [36], [37], music enhancement, and improved gaming immersion.³ Of course, once these actuators are included for rendering experiential effects, e.g., textures, they are also available for delivering richer tactons that render semantic information.

However, especially for in-the-wild studies, the sensory experience is often polluted by ambient vibrations, as discussed in Section V-A, such as those from moving vehicles in which the participant is riding, or self-induced motion [38]. In such cases, the richness of the effect can be less relevant than making sure that the vibrations are perceived at all.

This highlights the difference between focused use and ambient information delivery, two of the primary use cases for wearable haptic systems. On the one hand, when interacting directly with a small device like a smartwatch, the user typically needs to keep their arm still and focus on the experience, e.g., when twisting the crown on the Apple watch, and feeling the virtual “clicks” of detents as it scrolls through a list of items. The user is likely relatively motionless in order to perform the fine motor manipulation of the crown, minimizing motion that can confound the haptic effects, making higher-fidelity feedback more perceptible, and therefore more worthwhile. However, when receiving haptic information in the background, while not focused on interacting with the device, these benefits can be largely masked by overall motion.

As a result, for applications where ensuring perception is critical, in-the-wild experiments often use low fidelity vibrotactile actuators due to their combination of perceptual intensity, power efficiency, low cost, and compact size.

C. Power Consumption

A recent Immersion Corporation study found the impact of haptic actuators on everyday power consumption in smartphones to be minimal.⁴ However, smartphone battery capacity is on the order of ten times that of commercially available smartwatches (e.g., Apple Watch 279 mA · h and the Pebble model 301, 130 mA · h). Wang estimates that in a day of normal use, an ERM in a smartphone can consume 37 mA · h [39], which would represent 29 % of the Pebble battery. This implies that even though the presentation of haptic effects may represent negligible energy expenditures on a smartphone, they can quickly become a critical battery drain on a wearable device, especially if frequent or continuous vibrations are presented over a long period of time.

Anecdotally, in our studies with the Pebble smartwatch vibration motor triggering multiple times per minute throughout the day [40], the battery life drops from between four and seven days under normal use, to a little more than a single day, therefore requiring daily charging. Although the application uses Bluetooth communication and the accelerometer, which also consume power, the Pebble documentation is adamant in telling developers to, “use the vibration feature sparingly, because sustained use will rapidly deplete Pebble’s battery.”⁵ As a consequence, a participant who forgets to charge the device at night will almost inevitably run out of power the next day, resulting in the loss of data until they can charge. Further demonstrating the relationship between frequency of stimulation and battery consumption, the Basslet from Lofelt only enables 6 h of continuous music enhancement.

Supporting the significant impact of haptics on battery life, the Apple Watch developer documentation states, “Use of the haptic engine also consumes power, and using the engine too much may create a noticeable drain on the Apple Watch battery.”⁶ This is also reflected in the Android power save mode, which automatically disables haptic feedback, implying that the drain is non-trivial. Wearable devices are only becoming smaller, pushing all components to be more power efficient so that the battery can shrink while still delivering at least a full day of normal use.

Knowing that frequent presentation of haptic effects can impact battery life, hapticians have a few options that allow them to get the most out of their wearable systems. A first approach that minimizes power consumption is to choose a system that uses more efficient actuation mechanisms. Wang compared the energy consumption of three types of readily-available vibrotactile actuators: piezoelectric, LRA and ERM [39]. The conclusion was that piezoelectric consumes the least amount of energy for short effects (e.g., bump, click, pulse, alert), but drains a comparable amount of energy to an ERM for longer effects. On the other hand, the LRA can save

⁴ <https://www.immersion.com/wp-content/uploads/2017/08/haptic-technologies-consume-minimal-power-in-smart-phones.pdf>

⁵ https://developer.pebble.io/developer.pebble.com/docs/c/User_Interface/Vibes/index.html

⁶ <https://developer.apple.com/documentation/watchkit/wkinterfacedevice/1628128-play>

² <https://lofelt.com/blog/wideband-haptics-inside-nintendo-switch>

³ <https://lofelt.com/>

60–80% power in comparison to an ERM, and was found to be the most efficient actuation type in terms of force generated per unit of current.

A second approach is to obtain the greatest perceptual impact from the delivered stimulus under a given power budget. This requires knowledge of the human sensory system, since transmitting more raw force to the skin does not always translate into improved perception. Vibrotactile actuators are typically optimized to operate at a particular frequency that best stimulates the skin's receptors (details in Sections II-A2 & III-F). To this end, Jung and Choi tested the perceived magnitude of vibrotactile stimuli as a function of power consumption, finding that the difference between the vibration amplitude and the detection threshold is a good estimate for the perceived magnitude of the vibration [41]. Their results point to the possibility of predicting the perception of vibration by dividing the vibration amplitude at a given power (from the datasheet) by the detection threshold at the corresponding frequency.

At a higher level, haptic practitioners should consider taking advantage of opportunities to reduce power consumption when designing their applications. Promising techniques include using different pulsing strategies, minimizing the haptic signal duration according to the user's perception of the signal [42], or rendering only when appropriate and when chances of adequate perception are high [38]. Further work is required to determine what vibrotactile characteristics and tacton designs are the most "power efficient" for human perception.

D. Wear and Tear

Haptic experiments can quickly go awry when the device breaks down. This can be due to failure of the actuators themselves or damage to the connecting wires or other components. When designing a haptic device for long-term use, body position can make a large difference in longevity. Our lab, for example, has created haptic footwear, with actuators located under the foot itself [43]. During experiments, the wires connecting the various sensors and actuators would be repeatedly flexed, causing failures. Creating robust overall haptic devices is an engineering challenge requiring careful design, fabrication, and testing.

Wear and tear on the vibrotactile actuators themselves also bears comment. Since vibrotactile effects are usually generated by mechanical devices, typically small motors, they eventually wear out and fail. For common applications such as notifications throughout the day, or small haptic effects when touching the screen, the vibration motors are expected to outlast the lifetime of the overall device. Manufacturers typically characterize the lifetime of their ERM motors by the mean time to failure (MTTF). For example, according to the manufacturer data sheet, the Precision Microdrives Pico-Vibe™ 304-108 ERM motor has a MTTF of 836 h with a 50% duty cycle. Although this is sufficient for most typical applications that only trigger the motor occasionally for brief periods, it may preclude long-term use for studies that require potentially extended periods of continuous actuation [44]. To overcome

this limitation, the same company produces a Dura Vibe™ line of brushless motors that are specifically designed for lifetimes beyond 10000 h.⁷ Unfortunately, when using commercial devices such as smartphones and smartwatches that include a vibration motor, it can be difficult to determine exactly what type of actuator is inside, and therefore anticipate the lifetime.

There is, however, a more insidious problem. Ideally, vibrotactile actuator performance would be consistent across their lifetime, but in reality, for many such units, their characteristics change as they age and undergo mechanical wear and tear. For the aforementioned 304-108 ERM, Precision Microdrives reports that the vibration amplitude can drop significantly in the first few hundred hours of use, then be stable, then drop again before complete motor failure.⁸ Anecdotally, we have noticed similar effects in actual use. Pebble smartwatches worn for many months while testing a prototype system [40] seem to lose intensity before the motor completely fails. This has led us to better track which devices have been used most, so that we can give study participants "fresher" devices from our inventory.

Haptic studies sometimes report only the make and model of the actuators or end-user device (e.g., smartwatch) used in the experiments, and not necessarily their age or measured characteristics. When constructing multi-actuator devices, it is common to use actuators of the same type, for example, in a haptic belt [45]. There is an implicit assumption that multiple actuators from the same manufacturer will perform equivalently, and that using the same actuator across participants, or across studies in different labs, will make the results directly comparable. Unfortunately, this assumption is not necessarily justified, and may be problematic for some studies, especially when discrimination of amplitude levels is necessary. Even using only actuators of similar age may not be enough to eliminate variability, since manufacturing, type of use, or environmental differences may all mean that the breakdown differs across devices.

When actuation consistency is critical, avoiding vibrotactile actuators such as brushed motors that are known to degrade is recommended. When this is not possible, we recommend characterizing each device at the beginning of a study and reporting its rise time, fall time and steady state frequency response as the most transparent solution.

E. Communication and Synchronization

In a laboratory study, a controller can be directly wired to one or more actuators, allowing very accurate temporal triggering of complicated tactons (Fig. 2a). However, for in-the-wild studies, running wires is likely impractical, so a device such as a smartphone may need to wirelessly trigger actuators located at different points on the body (Fig. 2b). This raises issues of timing and synchronization, since common wireless protocols such as Bluetooth can have significant latency and jitter.

⁷ <https://www.precisionmicrodrives.com/content/ab-018-driving-brushless-long-life-vibration-motors/>

⁸ <https://www.precisionmicrodrives.com/content/lifetime-of-vibration-motors/>

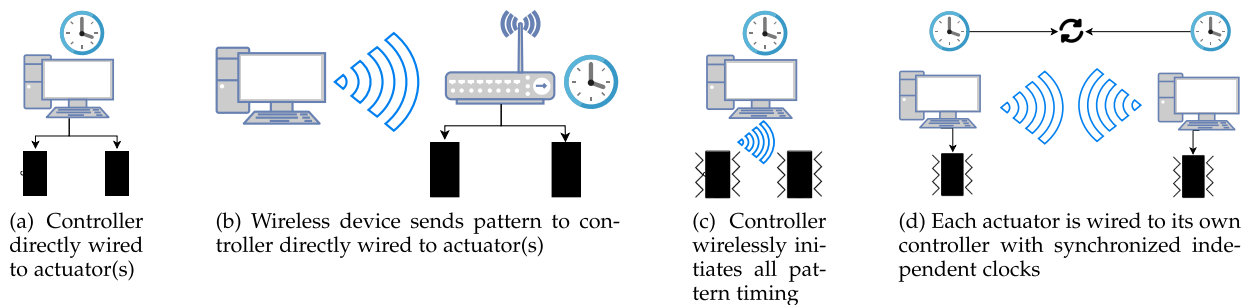


Fig. 2. Architectures for controlling the timing of haptic patterns, from fully wired (a), to completely wireless (d).

To illustrate these issues, we contrast the approaches of two commercial devices. Presumably to cut costs or simplify development, the Xiaomi Mi Band 1S⁹ uses a control system where the phone sends separate “motor on” and “motor off” commands to generate tactons, with the timing controlled from the connected phone (Fig. 2c).¹⁰ With this approach, the start and stop timings are subject to the latency and jitter of the underlying Bluetooth LE wireless communication. A developer can play some tricks to mitigate these issues, such as decreasing the connection interval timing before beginning a tacton, then restoring a slower connection interval when the tacton is complete, in order to save power. However, depending on the platform, the minimum connection interval can differ significantly, ranging from the BLE-defined minimum of 7.5 ms, to 30 ms.¹¹ Note that these are minimum connection intervals, and not necessarily the defaults, which can be around 50 ms, long enough to potentially make a difference in tacton generation, since the perceptible difference in duration is around this threshold for short vibrations [46].

Even with such tricks, however, such a system cannot finely control the motor at an interval down to a few milliseconds. Such a capability is useful, for example, to implement a Pulse Width Modulation (PWM) approach that controls an ERM’s perceived intensity by pulsing it rapidly on and off. This is possible on devices that have a local clock for timing the actuations, e.g., in a PWM library for the Pebble.¹² In this case, the motor may be driven for as little as 1 ms out of every 10, successfully reducing perceptibility when compared to continuously powering the motor [38]. The desired effect can be triggered from a Bluetooth-connected phone, but the tacton itself is timed locally on the Pebble (Fig. 2b). This model can be extended to the case of multiple actuators in one area of the body, such as a belt [45] or vest [24], where the actuators are close enough together to be wired to a single controller. This way, a Bluetooth-connected device can send a wireless message with all of the parameters, or an index into a library of tactons coded on the actuating device, and have it generated locally across all of the wired actuators. This does not avoid latency and jitter issues with timing the

start of the tacton as a whole, but for most applications, this is likely a less important consideration than the fidelity of the tacton itself.

However, such a strategy can break down when there are multiple actuators spread more broadly across the body, such that wiring them together is not possible, e.g., on each of the feet. In this case, the actuators would ideally be independently wireless, yet still able to render tightly synchronized tactons for effects such as directional sweeps. A low-latency non-Bluetooth wireless connection may sufficiently mitigate the timing issues, but is also likely to eliminate the possibility of using most commercial devices such as smartwatches and smartphones without modifications. Alternatively, synchronization issues can also be overcome if the individual devices have an accurate internal clock that can be independently synchronized to an external reference, e.g., a GPS clock signal or Network Time Protocol (NTP) server, or via a wireless clock sync protocol that synchronizes the clocks locally [47]. In these cases, the central controller can send a wireless command to all of the independent, but clock-synchronized, haptic nodes to begin their pattern at a specific time in the future (Fig. 2d). So long as the delay is longer than the maximum latency plus jitter of the wireless link to all nodes, reliable simultaneous triggering is possible. However, unless the required hardware and computing power for these synchronization methods is already in place for other purposes, adding it specifically for the clock synchronization will increase the cost and complexity of each actuator node.

F. Acoustic Considerations

Vibrotactile devices are prone to creating audible noise as a side effect of their mechanical actuation. If there is no isolation, the vibrations of the haptic-rendering mechanisms may be heard by the participant. In perceptual studies, this can lead to uncertainty as to whether the participants are reacting to the actual haptic sensation, or merely to the sound produced by the actuation. If the latter, any effects may disappear if quieter actuators are available in the future. To prevent this during lab-based experiments, the sound is typically masked with earplugs or pink noise through headphones, but this is obviously inappropriate for an in-the-wild study.

In order to better understand the problem, it is informative to compare the spectra of auditory and haptic vibrations, represented in Fig. 3.

⁹ <https://mobiletechtalk.co.uk/xiaomi-mi-band-1s-review/>

¹⁰ Although there is no formal API for the device, reference the <http://gadgetbridge.org> open source codebase for controlling the Mi Band.

¹¹ <https://punchthrough.com/blog/posts/maximizing-ble-throughput-on-ios-and-android>

¹² <https://github.com/jeffbl/pebble>

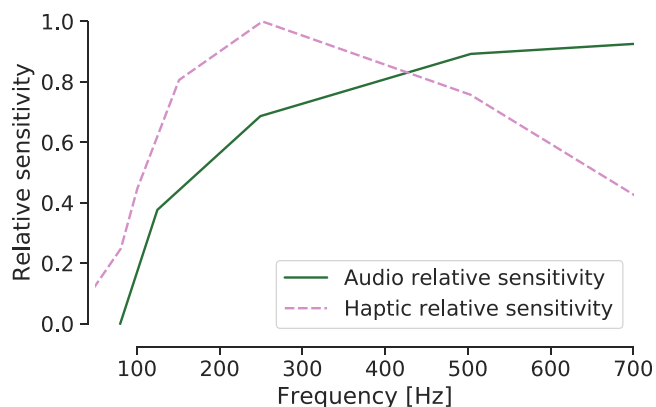


Fig. 3. Relative sensitivity to vibration [48] versus sound [49].

It is apparent that even though it is limited to low frequencies (0 Hz to 1000 Hz [33]), the haptic spectrum is superimposed over the range of human hearing, 20 Hz to 20 kHz [50]. In fact, hearing remains reasonably sensitive as low as 100 Hz [51], although frequencies up to 300 Hz will not be heard at sound pressure levels under 20 dB [49]. As such, it is generally recommended to use frequencies below 300 Hz for haptic rendering. Studies in multiple haptic research domains target the frequency of 250 Hz [52] as it corresponds to the peak sensitivity of Pacinian corpuscles [53]–[55] (Fig. 3). Indeed, one of the most commonly used research vibrotactile devices, the C-2 Tactor (Engineering Acoustics, Casselberry, FL, USA), is optimized for vibration at 200 Hz to 300 Hz. This does not preclude haptic designers from creating effects using other frequencies, as long as they are aware that perception will become more difficult.

A second recommendation is to avoid loose physical components that could create noise when moving. It is important to test not just actuators, but their integration into the hardware system as a whole, to make sure that their mechanical design avoids noise. Examples of higher frequency noise likely caused by loose component rattling and parasitic vibrations are shown in Fig. 4, in which a very audible ERM inside a Pebble smartwatch (model 301) is compared to the voicecoil in a Lofelt Basslet, and to an unidentified actuator, assumed to be an LRA, in a smartphone (Oneplus 5T). The frequencies above 250 Hz emitted by the Pebble ERM are the main cause of audible noise. Removing the higher harmonics results in a dramatically reduced acoustic signature, as can be heard in the examples at <http://srl.mcgill.ca/toh/Pebble>. In reality, however, the recommendations above are often difficult to apply to commercially available haptic actuators, for which the characteristics are typically immutable, or at best, arduous and time-consuming to modify.

If there is still a noise issue with the apparatus, vibration amplitude can be reduced in order to minimize acoustic noise when in a quiet environment, and increased when the participant is moving. This strategy can be particularly useful when conducting experiments on devices such as smartphones or smartwatches. In one of our ongoing studies [40], we are purposely lowering the intensity of the vibrations while the user

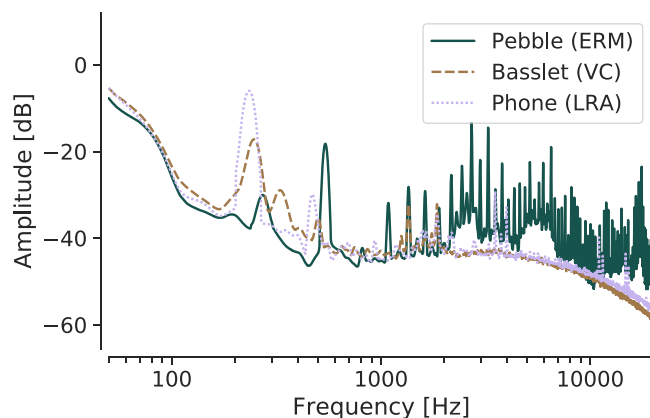


Fig. 4. Comparison of audible frequencies emitted by three different wearable haptic devices. All three actuators are targeting around 250 Hz, with the ERM and LRA by design, and the voicecoil of the Basslet driven specifically at that frequency. Importantly, note the extra audible frequencies above 250 Hz from the Pebble watch as compared to the lower levels from the Basslet. Unsurprisingly, the Pebble is much noisier when vibrating.

is relatively motionless, as measured by an accelerometer inside the smartwatch. This is motivated by the assumption that if the participant is not moving, they will still perceive the less intense stimuli [38]. However, by lowering the intensity of the haptic stimulus, it should be less audible, and therefore less annoying, both to the participant and to others nearby.

IV. INTERFACING TO THE BODY

Once the technical challenges of choosing and driving the actuators have been resolved, there remain important considerations around how to couple the actuators to the body. Additionally, an in-the-wild study, especially one lasting for days or weeks, will only be successful if it offers a system that participants are excited about, or at a bare minimum, willing to wear during their daily activities.

A. Haptic Coupling

Unlike with tightly controlled perceptual experiments, a number of issues make mechanical coupling a particularly difficult challenge for wearable haptics systems used for in-the-wild studies. Participants will engage in activities that cause physical displacement of actuators, as well as change the tightness of the coupling, due to muscle contraction that can cause devices to intermittently lose contact with the skin. Prior work has shown that a decrease in static force or pressure between the skin and the actuator is correlated with greater difficulty perceiving the delivered stimuli, and lower perceived stimuli intensities at various body locations [56]–[59].

While tightening a strap or using an elastic fabric to attach devices to the body can partially mitigate these problems, such solutions cannot guarantee consistent coupling characteristics during all daily activities, much less between multiple participants or across different studies. This is especially true of experimental protocols spanning multiple days that require participants to remove the device for charging and put it back on at the same location [29], [38]. In these cases, the

experimenter can, for example, count the number of holes in the clasp of a watch band after verifying appropriate coupling, and ask the participant to always use the same clasp position [38]. However, they cannot *guarantee* consistency as they rely, in most cases, on relative and subjective placement and tightness guidelines. Even assuming perfect participant compliance with the instructions, using the same clasp position will not guarantee the same coupling. For example, swelling from exercise injuries can cause a 9% increase in arm circumference [60]. To allow within and between participant, as well as cross-study comparison of results, coupling would ideally be determined using a calibrated reference or an absolute force or pressure measurement.

In haptic interfaces worn at highly active body locations such as the feet, the coupling situation is particularly problematic. Due to the nature of the gait cycle, pressure alternates between the toes and heels with each step, which can momentarily decouple the participant's foot from the actuators. Anlauff *et al.* demonstrated that the perception rate of tactons delivered under the feet dropped significantly when going from a resting (> 92%) to a walking state (> 66%) [43]. A comparison of the perception rates at different body locations under sitting and walking scenarios further demonstrated that the performance drop on the feet was almost double that of the wrist [9]. While the study did not specifically investigate the causes of that difference, it could be explained by the compounding effect of the suppression of haptic perception and attention generally experienced in active scenarios [61], added to the masking and decoupling effects that are specific to the feet and other densely jointed body locations [43].

Researchers have used load cells to measure the force under haptic actuators as part of laboratory studies, finding differences in perception due to coupling effects [58], [59]. Despite the supplemental planning during hardware design, employing such load cells as part of a dynamic calibration approach for mobile studies currently seems to be the most promising solution. The development and evaluation of such standardized methodology and means of reporting coupling characteristics would benefit the haptic research community as a whole.

B. Fashion and Participant Compliance

Unlike devices that render graphical or auditory content, haptic actuators must be “on body”, not just close to the body, in order to be effective. For an in-laboratory experiment, participants are unlikely to complain about wearing unwieldy, odd-looking devices. For an in-the-wild study, however, the size and aesthetics of a wearable haptic system can play a role in participant compliance or make recruitment more difficult, and therefore interfere with gathering reliable scientific data. This is especially true of prototype devices with which research groups often work. For example, the Movelet of Dobbstein *et al.* produces novel haptic effects, but would likely need significant streamlining to be accepted by users in everyday circumstances [62].

As noted above, coupling the actuator(s) to the skin is important for efficiency and reliability of haptic delivery.

However, the means of affixing them must also keep comfort in mind, especially if worn all day. Note that this must also take into account clothing worn in different weather and by different people. Anecdotally, we have had issues with hook and loop bands rubbing on the pants of participants, as well as wearing through stockings, causing participants to complain or move them to other locations on the body during pilot studies. We have therefore updated the straps to use primarily softer elastic material to avoid these issues. Others have suggested a magnetic clasp system, like that found on the Apple Watch, which can also be adjusted to any size without the issues caused by hook and loop fasteners. Nonetheless, some participants may be loathe to wear an apparatus in certain circumstances even despite careful design. In one of our studies, a participant flat out stated that they, “would not use [the leg-worn haptic device] in summer time” due to potential discomfort issues, even though they were willing to wear it in the cooler weather when the study happened to be running.

Aside from appearance and comfort, actuator noise can not only confound results (Section III-F), but also potentially reduce participant willingness to wear the device at all, at least when this noise might cause undue attention or embarrassment. Reducing the sound of the actuator(s) will help to make the system less obtrusive.

Lastly, it is worth noting that even with a well-designed system, there may still be resistance to wearing haptic devices, depending on the user population and the application. In some cases, haptics may not be a desirable solution to a problem, given other alternatives that potential users find preferable. This is an area where lab studies may be particularly misleading, since positive performance results may incorrectly be interpreted as indicating that a haptic approach would be successful with the intended user population. However, user acceptance may not follow for reasons unrelated to performance. For example, several groups have created haptic belts, and found that these perform well in lab environments for visually impaired guidance, e.g., Edwards *et al.* [45], which also cites a number of other vibrotactile belt projects. However, Golledge *et al.* surveyed visually impaired individuals about what types of assistive devices they would find most compelling for providing ongoing directional information in navigation tasks [63]. They found that vibrotactile devices placed around the neck or torso (i.e., a haptic belt) were ranked lowest in user preference, below a set of presented non-haptic options. Thus, despite scientifically verified performance benefits, such navigational haptic feedback systems have not yet seen significant uptake in commercial products.

V. ENVIRONMENTAL AND PERCEPTUAL EFFECTS

When conducting experiments in a laboratory, the environment and tasks can be tightly controlled. Upon moving out into the real-world, numerous issues and confounds appear. Although both individual differences and environmental factors can impact vibrotactile sensitivity [64], personal factors such as age, gender, and preexisting injuries are specific to the user. While they should be considered when running haptic studies,

they are invariant between in-lab and in-the-wild conditions, so are not discussed here. However, many environmental factors can be easily held constant during a lab study, but not during in-the-wild experiments. Examples include haptic noise from external vibrations, e.g., riding in a car, self-generated motion, such as walking, which reduces the perceived intensity of vibrotactile stimuli [65], and even weather conditions, for instance, cold weather, which can impact skin temperature and hydration, and in turn, alter haptic sensitivity [64].

This section addresses some of the unique problems raised for in-the-wild haptic studies, which spring from how users adapt over time in their perception of haptic stimuli, as well as from factors that are inherent to an individual, and from environmental differences that are uncontrollable when the user goes about their daily tasks.

A. Environmental Confounds

Just as we can become confused or overloaded by the constant stimuli received by our ears and eyes, the same is true for haptic stimuli. One of the major advantages of a laboratory study is the ability to control the environment, thereby minimizing the probability that such problems will occur. For an in-the-wild study, however, environmental confounds are likely to be one of the primary headaches, and can cause different issues. *Masking* [66] occurs when the haptic sense is overwhelmed by other haptic stimuli, such that the receptors are overcome with confounding haptic noise and the signal is imperceptible, e.g., not perceiving a phone vibration when riding a bike down a bumpy road. Note that non-vibrotactile haptic stimulation, e.g., temperature, may be masked by different confounds. *Gating* occurs when another stimulus, whether haptic or from another modality such as sight or sound, causes a haptic stimulus to be filtered out by the brain due to attention being allocated elsewhere [67]. For example, a painful event can overshadow a vibrotactile stimulus that would otherwise be perceived [68]. In addition to simply missing a stimulus, these effects can also result in *change blindness* [69], or failure to notice when a stimulus changes due to masking or gating. Pragmatically, this would mean a participant could be perceiving a haptic stimulus, but miss a change in the signal that was intended to convey important information. Although much of the existing literature is concerned with haptic detection thresholds, especially as related to signal amplitude or intensity, other in-the-wild studies also address tacton differentiability. For example, in the case of classifying numerical values delivered through vibration, Cauchard *et al.* showed a 36.4% decrease in the accuracy while the participant was running comparing to the stationary condition [29].

Attempts have been made to detect when masking or gating are likely to occur. At a coarse level, if a user is, for example, walking, this reduces the likelihood that a vibration at a given intensity will be perceived [9]. However, simply categorizing a participant as “walking” is a fairly coarse measure that encompasses a wide variety of possible motion intensities. At a finer level, measuring specific movement immediately prior to haptic delivery via accelerometer measurements, directly at

the site of stimulus, can also help predict whether it will be perceived [38]. This latter approach does not require any categorization of user activity, making it potentially less expensive to run on CPU and battery constrained hardware. Either way, coarse or fine motion information can potentially be used to determine when best to deliver a stimulus, or else be logged and factored into later analysis. In the future, such approaches may allow haptic feedback to be adjusted dynamically to be more perceptually equivalent regardless of confounding motion. Nonetheless, due to the difficulty of accurately predicting when such effects occur in the wild, a typical approach is to use more intense effects than would be necessary in a controlled environment, to make sure that they are perceptible even in less than ideal circumstances, subject to acoustic noise considerations (Section III-F).

The above environmental confounds typically reduce haptic perception. On the flip side, many people report hallucinating “phantom vibrations” [70], which although not real, cause participants to think they have received a haptic stimulus, typically a vibrotactile notification, even though none was administered. These phantom vibrations may interfere with perception rates or otherwise distort results of wearable haptic studies, especially for participants prone to these false signals, e.g., with more psychological dependency on cell phone communications [71], or a greater need for popularity [72].

While it may not be possible to mitigate environmental confounds in all applications, being aware of them helps haptic practitioners to adapt their designs to be more robust. For example, if a user’s motion is likely to cause masking (e.g., a user is most likely in motion if they are receiving navigation information), the system can render the haptic signal at multiple locations to increase the likelihood of it being perceived. This is strongly dependent on context. In an emergency room, for example, a haptic device rendering vital patient information should be designed to be robust to gating from competing audio and visual alarms by making the haptic signal much different from the audio signal (using different patterns) or by reinforcing it with an audio signal instead of competing directly with the other modalities.

B. Attention

In contrast to environmental confounds from external sources, a participant’s own internally guided attention also plays a role in haptic perception. Since haptics are especially well suited to conveying information in the background [32], issues of partial attention are particularly germane, but are difficult to measure.

Although an in-the-wild study provides valuable insight into how a wearable haptic system performs in actual use, when the system fails to produce the desired results, it can be difficult to determine whether it was due to the user’s attention being elsewhere. In-lab studies that attempt to introduce elements of real-world conditions often require careful design of an experimental activity to mirror the demands of the real task (s) in which the participants are assumed to be engaged. A common approach is to instruct participants to work on a

distractor task that occupies part of their attention and cognitive resources, and that ideally has a more continuous amount of mental demand than would be found in most real-world tasks, where attention can spike and recede as the task demands vary. Some that have been used in haptic studies include transcription and mouse-based data entry [73], repetitive pointing to a dynamic target on screen [74], visual tracking [75], puzzles [76], and attending to an audio stream and responding when a particular word is spoken [76]. These tasks may not be ecologically valid in and of themselves, but they can exercise similar faculties that would be engaged during real-world tasks, and allow conclusions to be drawn concerning how well a system will perform when the user is not paying full attention to the haptic information being presented.

Nonetheless, although these in-lab experiences can help improve the haptic system by pointing out flaws, they do not fully replace true in-the-wild testing, with real users performing real tasks.

C. Perceptual Variability Over Time

A participant's perception will change after repeated or continuous exposure to a haptic stimulus. Although this can occur in short laboratory studies, changes are likely to be more evident when participants are receiving haptic signals throughout the day, over an extended period of time. At first glance, this may appear to always be a negative effect, but it can be a crucial part of a haptic application, depending on the type of change.

Sensory adaptation [77], [78] is typically a decrease in sensitivity to a constant stimulus, occurring in the sensory system [79], [80]. An illustrative example is given by Langley *et al.*, "[One] is aware of the warmth of a fluid, but if the hand is kept immersed for a period of time the sensation of warmth disappears" [81].

Habituation [82], [83] is a form of learning in which a repeated stimulus ceases to produce a response. This is considered to be an attentional processing effect. It can be beneficial for some applications, and potentially detrimental for others. As Brewster points out, habituation can be the goal of an ambient system, since it means that the haptic stimuli fades into "the background of consciousness" [84]. Therefore, only dynamic, salient differences will be noticed, unless one turns their attention specifically to the stimulus, in which case it can be perceived. For notification applications, however, habituation can cause missed information, especially likely if most of the notifications are useless, thereby priming the participant to ignore them.

Adaptation is often confused with habituation [78]. The practical difference is that adaptation is physiological, such that the participant will be unable to notice changes in the signal even if they try, while habituation is behavioral, allowing the user to mostly ignore a stimulus, but still "tune in" to it by directing their attention. Habituation can thus assist a user in managing an ongoing signal, while with adaptation, any changes may be imperceptible.

Where habituation or adaptation are undesirable, providing resting periods can help to mitigate the effects. Vibrotactile

adaptation, for example, can be reversed in a matter of minutes [79], but this time can vary widely based on conditions [85]. Recent work by Kotowick and Shah also finds that switching between haptics and another modality, in this case visual, when performing a navigation task can reduce habituation and adaptation [85]. Their paper also contains a broad set of references to other dishabituation techniques and adaptation recovery times in various conditions.

Last, in perhaps the most extreme case of perceptual change over time, the delivery of strong, low-frequency vibration stimuli can produce significant harmful long term effects such as Hand-Arm Vibration Syndrome [86]. If working with applications that require this type of vibrotactile stimulus, due care should be taken to stay within safe limits.

VI. METHODOLOGICAL CONCERNS

Even after a haptic system has been developed and tested, the experiment and data collection procedures remain to be designed. This raises still more issues beyond those discussed in previous sections, including how to integrate into the participants' existing haptic landscape, since many people already carry and use vibrotactile devices, and how to best assess how well a system is working.

A. Deployment on Unknown Hardware

In-the-wild haptic studies are most often conducted using hardware that the researcher provides to the participants. This allows for a level of control that would be difficult if deploying a haptic application onto the participant's existing hardware. When dealing with GUI applications, differences between smartphones are largely abstracted away by standard APIs and functionality, allowing the researcher to deal with questions of screen size, quality and resolution. Such software infrastructure, which allows for an approximately common user experience, independently of hardware and operating system particularities, is largely non-existent in the world of haptics. Haptics on phones and smartwatches can vary considerably with respect to their fidelity, audibility, and coupling to the body. On the Android platform, this variety of devices is extreme, ranging from high-quality LRA actuators, to noisy ERM motors. One of the few studies to confront this specific issue is the HapTurk project by Schneider *et al.*, which proposed to translate high fidelity vibrotactile icons created using a high-end vibrotactile actuator, into a "low-fidelity vibration proxy" that preserved its key affective aspects when played on commodity Android smartphones [87]. However, these modified icons were created by hand; an automated process for creating complex haptic effects that translate to a wide variety of different haptic hardware remains an open problem.

B. Integrating With Existing Devices

Because of their ubiquity, many people already wear vibrotactile devices on their body, ranging from a cellphone in their pocket to smartwatches and fitness trackers. In a laboratory study, participants can be asked to remove these devices so

that they do not interfere with the haptic effects to be tested. However, when in the wild, participants are likely loathe to remove their existing haptic devices, since they rely on them for notifications. In addition, preventing users from accessing their usual technologies could influence their use of the experimental system. For users who count on their devices throughout the day, giving them up creates the possibility of missing an incoming notification, which can be sufficient to induce a significant amount of stress [88]. Especially for long studies, expecting participants to put aside their favorite mobile and wearable devices would be unreasonable, and would make recruiting difficult or reduce compliance during the experiment. In this case, using an alternative to vibrotactile haptic feedback, e.g. varying temperature, rubbing, or squeezing, may help participants distinguish haptic information from their normal vibrotactile notifications. This is especially the case if delivered at a body location near their already vibrating phone, smartwatch, or fitness tracker.

Another solution is to avoid areas of the body that are most often used by current commercial systems. However, even when separating vibrotactile effects by delivering them to different areas of the body, perceptual effects such as tactile gating and masking, as discussed earlier, can cause haptic stimuli to interfere with each other. In addition, there are other subtle issues to be considered even if the stimuli are not rendered at the same time. For example, Lakatos and Shepard demonstrated that humans are slower at perceiving haptic stimuli at a point on the body far from another body site to which they were recently attending [89]. This means that there may be a tradeoff between being able to discriminate vibrations at different body locations: the further apart the stimuli, the easier to discriminate, but the slower one can attend to them. Delving further into this topic, Gallace et al. review results of studies dealing with human limitations for processing of both unimodal (purely haptic) and multimodal stimuli, separated over the body [90]. Even in systems designed to coordinate the multiple actuators, this can be a difficult task. When running a study with a participant's existing, uncoordinated, haptic devices in place, which would effectively be triggering randomly unless otherwise controlled, one would expect the overlapping stimuli to interfere with each other. If so, this would motivate future work that would try and stage the different systems to avoid such interference, e.g., holding a text message vibration until after a fitness notification occurred, to avoid overlap and attentional issues.

C. Assessment

Traditional performance measures, such as response times and accuracy in perceiving stimuli, are also valid for in-the-wild haptic studies. Although such quantitative measures are valuable, for experiential haptic experiments, rigidly quantitative approaches may be less valuable than more subjective metrics, such as the perceived realism of a haptic simulation, its believability, or user engagement [91].

In either case, one promising future direction is using biophysiological measurements, such as heart rate or skin

conductance, to gain indirect insight into a user's perception and affect. This approach avoids self-evaluation biases, but is limited in what it can measure. For example, it may be possible to determine whether a user perceived a haptic signal [42], but there is currently no way to automatically gather information as to its perceived intensity, interpretation, or how much of it was perceived. As such, a combination of quantitative performance measures, biophysiological measurements, and self-evaluation is likely to yield a more robust assessment of haptic device usage and performance.

Although not unique to haptic studies, participants are likely occupied with other tasks that demand their attention when they are active in the wild. From an ethical and safety perspective, participants should be instructed to delay responses to a stimulus until it is safe to do so, and this must be taken into account in the analysis. As a case in point, for a haptic study requiring the user to press a smartwatch button each time they perceived a vibration throughout the day, we opted to use two separate buttons, one for when the user was able to press it right away, and a second button they were instructed to use if they could only respond after a delay [38]. This provides at least some idea of when response times may be longer than expected because of a delay in being able to mark the event safely, rather than a delay in perception.

VII. DISCUSSION

Some readers may be feeling discouraged due to the numerous practical issues that arise when conducting wearable, in-the-wild haptic user studies. Nonetheless, the value of such studies for gathering data on real-world use means that haptic practitioners will necessarily be confronting these issues for years to come. Despite this reality, we believe optimism is warranted, since perfection is not necessary, technical improvements are ongoing, and the remaining challenges can be overcome.

A. Do Not Let the Perfect Be the Enemy of the Good

In-the-wild studies inherently have more uncontrolled confounds due to the unpredictability of the environment. They can be painful to run, and may lead to more frequent negative results, but are valuable in discovering the practical limits of haptic systems in real-world use. It is a rare in-the-wild user study that does not deliver useful insights simply unobtainable in a laboratory environment. Even if some of the obstacles to in-the-wild studies are hard to overcome, being aware of them allows researchers to be conscious of what possible confounds may occur during an experiment. Accordingly, experimenters can minimize the impact of these confounds and/or better interpret the results of their studies based on the information in this article. For example, in one of our studies [38], while a higher quality, lower noise actuator and a more consistent between-participant coupling mechanism may have been beneficial, it was still possible to determine useful perception curves in the face of real-world motion confounds. Furthermore, the use of a real commercial actuator can better reflect the experience that users may achieve using their own

commodity devices, and opens the door to larger scale in-the-wild studies using participants' own devices, greatly extending the reach of such studies to broader user populations.

B. Technical Improvements Are Ongoing

For many years, consumer haptic devices were largely limited to ERMs with basic on/off control. With high-volume phones, handheld gaming systems, fitness trackers, and smart watches driving demand for improved haptic effects, actuators and the software to control them are now rapidly improving. This opens up new possibilities for conducting in-the-wild mobile studies, since researchers can increasingly count on the improved hardware, as well as control software, being available to more of the general population. For example, early versions of the Android API only supported vibrotactile on/off patterns. Querying to determine if the haptic hardware supports amplitude control was only added in Android API 26, released in August 2017, a capability that is critical for determining what effects can be rendered on a given device. Both expanding access to, as well as providing more information about, the actuators in a device, e.g., actuator type, axis, resonant frequency, and force measurements, would allow hapticians to better determine what is possible, and tailor effects to a specific device. Although Apple is to be commended for their upcoming release of the Core Haptics API for their iPhone devices, we can only hope that they and others expand their SDKs to allow access to the increasingly robust and expressive haptic hardware they are shipping to millions of users on more wearable devices such as the Apple Watch.

C. Standardization of Haptic Studies

While this article identifies a number of areas for future improvement, we argue that the most urgent issue in the haptic community at this moment is the lack of generalizability and comparability of studies' findings. We therefore invite researchers to pay particular attention to developing:

- Standardized assessment methods
- Standardized coupling methods and reporting, such as measurement procedures (e.g., load cell placement).

Contributions to these two aspects would allow the generation of more cohesive, consistent, repeatable, and comparable results that would strengthen the haptics community, bringing it closer to that of its audio and visual counterparts.

VIII. CONCLUSION

We have discussed a wide range of issues and suggestions for conducting in-the-wild, wearable, vibrotactile haptic studies, including both general background information useful to a novice haptic designer, as well as detailed specific lessons learned while designing our own haptic systems for real-world use. By pulling relevant information not just from academic literature, but also from commercial material, we hope to help bridge the gap between what is theoretically needed to conduct such research, and what is actually available in the marketplace, despite its shortcomings. By assembling this information into

one accessible article, we expect that haptics practitioners will be able to improve their chances of running meaningful and successful in-the-wild experiments with wearable haptic devices.

ACKNOWLEDGEMENT

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), RGPIN-2017-05013.

REFERENCES

- [1] N. Henze, M. Pielot, B. Poppinga, T. Schinke, and S. Boll, "My app is an experiment: Experience from user studies in mobile app stores," *Int. J. Mobile Human Comput. Interact.*, vol. 3, no. 4, pp. 71–91, Oct. 2011. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2440778.2440783>
- [2] J. Kjeldskov, M. B. Skov, B. S. Als, and R. T. Høegh, "Is it worth the hassle? Exploring the added value of evaluating the usability of context-aware mobile systems in the field," in *Proc. Int. Conf. Human-Comput. Interact. Mobile Devices Serv.*, 2004, pp. 61–73. [Online]. Available: http://link.springer.com/10.1007/978-3-540-28637-0_6
- [3] J. Kjeldskov and M. B. Skov, "Was it worth the hassle? Ten years of mobile HCI research discussions on lab and field evaluations," in *Proc. 16th Int. Conf. Human-Comput. Interact. Mobile Devices Serv.*, New York, NY, USA: ACM Press, 2014, pp. 43–52. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2628363.2628398>
- [4] D. Parisi, *Archaeologies of Touch*. Minneapolis, MN, USA: University of Minnesota Press, Feb. 2018. [Online]. Available: <https://www.jstor.org/stable/10.5749/j.ctt20mvgvz>
- [5] R. W. Van Boven and K. O. Johnson, "The limit of tactile spatial resolution in humans: Grating orientation discrimination at the lip, tongue, and finger," *Neurology*, vol. 44, no. 12, pp. 2361–2361, Dec. 1994. [Online]. Available: <http://www.neurology.org/cgi/doi/10.1212/WNL.44.12.2361>
- [6] F. Mancini *et al.*, "Whole-body mapping of spatial acuity for pain and touch," *Ann. Neurol.*, vol. 75, no. 6, pp. 917–924, 2014.
- [7] C. Pacchierotti, S. Sinclair, M. Solazzi, A. Frisoli, V. Hayward, and D. Prattichizzo, "Wearable haptic systems for the fingertip and the hand: Taxonomy, review, and perspectives," *IEEE Trans. Haptics*, vol. 10, no. 4, pp. 580–600, Oct.–Dec. 2017.
- [8] H. Culbertson, S. B. Schorr, and A. M. Okamura, "Haptics: The present and future of artificial touch sensation," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 1, pp. 385–409, 2018.
- [9] I. Karuei, K. E. MacLean, Z. Foley-Fisher, R. MacKenzie, S. Koch, and M. El-Zohairy, "Detecting vibrations across the body in mobile contexts," in *Proc. Conf. Human Factors Comput. Syst.*, New York, NY, USA: ACM Press, 2011, pp. 3267–3276. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=1978942.1979426>
- [10] S. Choi and K. J. Kuchenbecker, "Vibrotactile display: Perception, technology, and applications," *Proc. IEEE*, vol. 101, no. 9, pp. 2093–2104, Sep. 2013. [Online]. Available: <http://ieeexplore.ieee.org/document/6353870/>
- [11] L. A. Jones and N. B. Sarter, "Tactile displays: Guidance for their design and application," *Human Factors J. Human Factors Ergonom. Soc.*, vol. 50, no. 1, pp. 90–111, Feb. 2008. [Online]. Available: <http://journals.sagepub.com/doi/10.1518/001872008X250638>
- [12] J. B. F. van Erp, "Guidelines for the use of vibro-tactile displays in human computer interaction," in *Proc Eurohaptics*, 2002, pp. 18–22.
- [13] K. MacLean, "Foundations of transparency in tactile information design," *IEEE Trans. Haptics*, vol. 1, no. 2, pp. 84–95, Jul. 2008. [Online]. Available: <http://ieeexplore.ieee.org/document/4674348/>
- [14] Y. Wang, B. Millet, and J. L. Smith, "Informing the use of vibrotactile feedback for information communication: An analysis of user performance across different vibrotactile designs," *Proc. Human Factors Ergonom. Soc.*, vol. 58, no. 1, pp. 1859–1863, 2014. [Online]. Available: <http://pro.sagepub.com/lookup/doi/10.1177/1541931214581389>
- [15] G. Wilson, M. Halvey, S. A. Brewster, and S. A. Hughes, "Some like it hot: Thermal feedback for mobile devices," in *Proc. Special Interest Group Comput.–Human Interact. Conf. Human Factors Comput. Syst.*, ACM, 2011, pp. 2555–2564.
- [16] H. Pohl and K. Hornbæk, "Electricitch: Skin irritation as a feedback modality," in *Proc. 31st Annu. Symp. User Interface Softw. Technol.*, New York, NY, USA: ACM Press, 2018, pp. 765–778. [Online]. Available: <http://doi.acm.org/10.1145/3242587.3242647>

- [17] T. Duinte, J. Schulte, M. Pfeiffer, and M. Rohs, "Muscleio: Muscle-based input and output for casual notifications," *Proc. ACM Interact. Mobile Wearable Ubiquitous Technol.*, vol. 2, no. 2, pp. 64:1–64:21, Jul. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3214267>
- [18] H. Pohl, P. Brandes, H. Ngo Quang, and M. Rohs, "Squeezeback: Pneumatic compression for notifications," in *Proc. Conf. Human Factors Comput. Syst.*, New York, NY, USA: ACM Press, 2017, pp. 5318–5330. [Online]. Available: <http://doi.acm.org/10.1145/3025453.3025526>
- [19] T. Han, F. Anderson, P. Irani, and T. Grossman, "HydRoring: Supporting mixed reality haptics using liquid flow," in *Proc. 31st Annu. ACM Symp. User Interface Softw. Technol.*, New York, NY, USA: ACM Press, 2018, pp. 913–925. [Online]. Available: <http://doi.acm.org/10.1145/3242587.3242667>
- [20] Y. A. Shim, J. Lee, and G. Lee, "Exploring multimodal watch-back tactile display using wind and vibration," in *Proc. Conf. Human Factors Comput. Syst.*, New York, NY, USA: ACM Press, 2018, pp. 132:1–132:12. [Online]. Available: <http://doi.acm.org/10.1145/3173574.3173706>
- [21] V. Hayward and K. E. Maclean, "Do it yourself haptics: Part I," *IEEE Robot. Automat. Mag.*, vol. 14, no. 4, pp. 88–104, Dec. 2007. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4437756>
- [22] K. MacLean and V. Hayward, "Do it yourself haptics: Part II [Tutorial]," *IEEE Robot. Automat. Mag.*, vol. 15, no. 1, pp. 104–119, Mar. 2008.
- [23] V. Hayward, "A brief taxonomy of tactile illusions and demonstrations that can be done in a hardware store," *Brain Res. Bull.*, vol. 75, no. 6, pp. 742–752, 2008.
- [24] S. D. Novich and D. M. Eagleman, "Using space and time to encode vibrotactile information: toward an estimate of the skin's achievable throughput," *Exp. Brain Res.*, vol. 233, no. 10, pp. 2777–2788, 2015.
- [25] K. O. Sofia and L. Jones, "Mechanical and psychophysical studies of surface wave propagation during vibrotactile stimulation," *IEEE Trans. Haptics*, vol. 6, no. 3, pp. 320–329, Jul.–Sep. 2013.
- [26] J. B. F. van Erp, "Vibrotactile spatial acuity on the torso: Effects of location and timing parameters," in *Proc. First Joint Eurohaptics Conf. Symp. Haptic Interfaces Virtual Environ. Teleoperator Syst. World Haptics Conf.*, 2005, pp. 80–85.
- [27] S. Paneels, M. Anastassova, S. Strachan, S. P. Van, S. Sivacoumarane, and C. Bolzmacher, "What's around me? Multi-actuator haptic feedback on the wrist," in *Proc. World Haptics Conf.*, 2013, pp. 407–412.
- [28] S. Brewster and L. M. Brown, "Tactons: Structured tactile messages for non-visual information display," in *Proc. Fifth Conf. Australasian User Interface-Volume 28*, Australian Computer Society, Inc., 2004, pp. 15–23.
- [29] J. R. Cauchard, J. L. Cheng, T. Pietrzak, and J. A. Landay, "ActiVibe: Design and evaluation of vibrations for progress monitoring," in *Proc. Conf. Human Factors Comput. Syst.*, New York, NY, USA: ACM Press, 2016, pp. 3261–3271. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2858036.2858046>
- [30] 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, 2011.
- [31] G. Changeon, D. Graeff, M. Anastassova, and J. Lozada, "Tactile emotions: A vibrotactile tactile gamepad for transmitting emotional messages to children with autism," in *Int. Conf. Haptics: Perception, Devices, Mobility, Commun.*, P. Isokoski and J. Springare, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 79–90.
- [32] K. MacLean, "Putting Haptics into the ambience," *IEEE Trans. Haptics*, vol. 2, no. 3, pp. 123–135, Jul.–Sep. 2009. [Online]. Available: <http://ieeexplore.ieee.org/document/5184837/>
- [33] F. Reynier and V. Hayward, "Summary of the kinesthetic and tactile function of the human upper extremities," *McGill Res. Cent. Intell. Mach.*, Tech. Rep., CIM-93-4, pp. 1–50, 1993. [Online]. Available: <http://www.cim.mcgill.ca/~haptic/pub/FR-VH-92.pdf>
- [34] T. T. Diller, "Frequency response of human skin in vivo to mechanical stimulation," Master of Science, Mechanical Engineering, Massachusetts Institute of Technology, 2001. [Online]. Available: <https://dspace.mit.edu/bitstream/handle/1721.1/41116/RLE-TR-648-48839050.pdf>
- [35] J. M. Romano and K. J. Kuchenbecker, "Creating realistic virtual textures from contact acceleration data," *IEEE Trans. Haptics*, vol. 5, no. 2, pp. 109–119, Apr. 2012.
- [36] Apple Inc. Digital crown, user interaction, watchOS, human interface guidelines, apple developer, 2018. [Online]. Available: <https://developer.apple.com/design/human-interface-guidelines/watchos/user-interaction/digital-crown/>
- [37] Introducing iPhone 6s and iPhone 6s Plus with 3D Touch, 2016. [Online]. Available: <https://www.youtube.com/watch?v=kuHqfzcJYA>
- [38] J. R. Blum, I. Frissen, and J. R. Cooperstock, "Improving haptic feedback on wearable devices through accelerometer measurements," in *Proc. 28th Annu. Assoc. Comput. Machinery Symp. User Interface Softw. Technol.*, 2015, pp. 31–36. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2807442.2807474>
- [39] F. Wang, Haptic energy consumption, Texas Instruments Application Report SLOA194, 2014. [Online]. Available: <http://www.ti.com/lit/an/sloa194/sloa194.pdf>
- [40] J. R. Blum and J. R. Cooperstock, "Expressing human state via parameterized haptic feedback for mobile remote implicit communication," in *Proc. 7th Augment. Human Int. Conf.*, New York, NY, USA: ACM Press, Feb. 2016, pp. 1–2. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2875194.2875225>
- [41] J. Jung and S. Choi, "Perceived magnitude and power consumption of vibration feedback in mobile devices," in *Proc. Human-Comput. Interact.*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2007, vol. 4551, pp. 354–363.
- [42] P. E. Fortin, E. Sulmont, and J. R. Cooperstock, "Detecting perception of smartphone notifications using skin conductance responses," in *Proc. Conf. Human Factors Comput. Syst.*, New Orleans, LA, USA: ACM, 2019, Paper 190.
- [43] J. Anlauff, J. Fung, and J. R. Cooperstock, "Vibewalk: Foot-based tactons during walking and quiet stance," in *Proc. IEEE World Haptics Conf.*, 2017, pp. 647–652.
- [44] C. Leichsenring, R. Tünnermann, and T. Hermann, "Feelabuzz: Direct tactile communication with mobile phones," *Int. J. Mob. Human Comput. Interact.*, vol. 3, no. 2, pp. 65–74, Jan. 2011. [Online]. Available: <http://www.igi-global.com/article/feelabuzz-direct-tactile-communication-mobile/53217>
- [45] N. Edwards *et al.*, "A pragmatic approach to the design and implementation of a vibrotactile belt and its applications," in *Proc. IEEE Int. Workshop Haptic Audio Vis. Environ. Games*, IEEE, Nov. 2009, pp. 13–18. [Online]. Available: <http://ieeexplore.ieee.org/document/5356126/>
- [46] F. A. Geldard, "Some neglected possibilities of communication," *Science*, vol. 131, no. 3413, pp. 1583–1588, 1960. [Online]. Available: <http://www.jstor.org/stable/1705360>
- [47] B. Sundararaman, U. Buy, and A. D. Kshemkalyani, "Clock synchronization for wireless sensor networks: A survey," *Ad Hoc Netw.*, vol. 3, no. 3, pp. 281–323, 2005.
- [48] R. T. Verrillo, "Vibrotactile thresholds measured at the finger," *Percept. Psychophys.*, vol. 9, no. 4, pp. 329–330, 1971.
- [49] R. S. Dadson and J. H. King, "A determination of the normal threshold of hearing and its relation to the standardization of audiometers," *J. Laryngol. Otol.*, vol. 66, no. 8, pp. 366–378, 1952.
- [50] S. Rosen and P. Howell, *Signals and Systems for Speech and Hearing*, vol. 29. Brill, Leiden, The Netherlands, 2011.
- [51] S. A. Gelfand, *Essentials of Audiology*. Stuttgart, Germany: Thieme, 2016.
- [52] K. E. MacLean, "Haptic interaction design for everyday interfaces," *Rev. Human Factors Ergonom.*, vol. 4, no. 1, pp. 149–194, 2008.
- [53] M. Sato, "Response of Pacinian corpuscles to sinusoidal vibration," *J. Physiol.*, vol. 159, no. 3, pp. 391–409, Dec. 1961. [Online]. Available: <http://doi.wiley.com/10.1113/jphysiol.1961.sp006817>
- [54] R. T. Verrillo, "Effect of contactor area on the vibrotactile threshold," *J. Acoust. Soc. Amer.*, vol. 35, no. 12, pp. 1962–1966, Dec. 1963. [Online]. Available: <http://asa.scitation.org/doi/10.1121/1.1918868>
- [55] K. Shimoga, "Finger force and touch feedback issues in dexterous tele-manipulation," in *Proc. Fourth Annu. Conf. Intell. Robot. Syst. Space Exploration*, Troy, NY, USA: IEEE, 1992, pp. 159–178. [Online]. Available: <http://ieeexplore.ieee.org/document/671841/>
- [56] J. C. Craig and C. E. Sherrick, "The role of skin coupling in the determination of vibrotactile spatial summation," *Percept. Psychophys.*, vol. 6, no. 2, pp. 97–101, 1969. [Online]. Available: <https://doi.org/10.3758/BF03210689>
- [57] B. G. Green and J. C. Craig, "The roles of vibration amplitude and static force," *Percept. Psychophys.*, vol. 16, no. 3, pp. 503–507, 1974. [Online]. Available: <https://doi.org/10.3758/BF03198578>
- [58] M. Azadi and L. A. Jones, "Vibrotactile actuators: Effect of load and body site on performance," in *Proc. IEEE Haptics Symp.*, 2014, pp. 351–356.
- [59] S. Papetti, H. Järveläinen, B. L. Giordano, S. Schiesser, and M. Fröhlich, "Vibrotactile sensitivity in active touch: Effect of pressing force," *IEEE Trans. Haptics*, vol. 10, no. 1, pp. 113–122, Jan.–Mar. 2017.
- [60] J. N. Howell, G. Chleboun, and R. Conatser, "Muscle stiffness, strength loss, swelling and soreness following exercise-induced injury in humans," *J. Physiol.*, vol. 464, no. 1, pp. 183–196, May 1993. [Online]. Available: <http://doi.wiley.com/10.1113/jphysiol.1993.sp019629>

- [61] L. J. Post, I. C. Zompa, and C. E. Chapman, "Perception of vibrotactile stimuli during motor activity in human subjects," *Experimental Brain Res.*, vol. 100, no. 1, pp. 107–120, 1994. [Online]. Available: <https://doi.org/10.1007/BF00227283>
- [62] D. Döbelstein, E. Stemasov, D. Besserer, I. Stenske, and E. Rukzio, "Movelet: A self-actuated movable bracelet for positional haptic feedback on the user's forearm," in *Proc. Assoc. Comput. Machinery Int. Symp. Wearable Comput.*, New York, NY, USA: ACM Press, 2018, pp. 33–39. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=3267242.3267249>
- [63] R. Gollidge, R. Klatzky, J. Loomis, and J. Marston, "Stated preferences for components of a personal guidance system for nonvisual navigation," *J. Vis. Impair. Blind.*, vol. 98, no. 03, 2004. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.851.3955&rep=rep1&type=pdf>
- [64] M. S. Gandhi, R. Seseek, R. Tuckett, and S. J. M. Bamberg, "Progress in vibrotactile threshold evaluation techniques: A review," *J. Hand Ther.*, vol. 24, no. 3, pp. 240–256, 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.jht.2011.01.001>
- [65] R. Milne, A. Aniss, N. Kay, and S. Gandevia, "Reduction in perceived intensity of cutaneous stimuli during movement: A quantitative study," *Exp. Brain Res.*, vol. 70, no. 3, pp. 569–576, May 1988. [Online]. Available: <http://link.springer.com/10.1007/BF00247604>
- [66] M. Enriquez and K. E. MacLean, "Backward and common-onset masking of vibrotactile stimuli," *Brain Res. Bull.*, vol. 75, no. 6, pp. 761–769, Apr. 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0361923008000105>
- [67] C. E. Chapman, I. C. Zompa, S. R. Williams, J. Shenasa, and W. Jiang, "Factors influencing the perception of tactile stimuli during movement," in *Somesthesia and the Neurobiology of the Somatosensory Cortex*. New York, NY, USA: Springer, 1996, pp. 307–320.
- [68] S. J. Bolanowski, L. M. Maxfield, G. A. Gescheider, and A. V. Apkarian, "The effects of stimulus location on the gating of touch by heat- and cold-induced pain," in *Somatosens. Mot. Res.*, vol. 17, no. 2, pp. 195–204, 2000. [Online]. Available: <http://www.tandfonline.com/doi/full/10.1080/08990220050020607>
- [69] M. Auvray, A. Gallace, J. Hartcher-O'Brien, H. Z. Tan, and C. Spence, "Tactile and visual distractors induce change blindness for tactile stimuli presented on the fingertips," *Brain Res.*, vol. 1213, pp. 111–119, 2008.
- [70] M. Drouin, D. H. Kaiser, and D. A. Miller, "Phantom vibrations among undergraduates: Prevalence and associated psychological characteristics," *Comput. Human Behav.*, vol. 28, no. 4, pp. 1490–1496, 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.chb.2012.03.013>
- [71] D. J. Kruger and J. M. Djerf, "Bad vibrations? Cell phone dependency predicts phantom communication experiences," *Comput. Human Behav.*, vol. 70, pp. 360–364, 2017. [Online]. Available: <http://dx.doi.org/10.1016/j.chb.2017.01.017>
- [72] M. Tanis, C. J. Beukeboom, T. Hartmann, and I. E. Vermeulen, "Phantom phone signals: An investigation into the prevalence and predictors of imagined cell phone signals," *Comput. Human Behav.*, vol. 51, no. PA, pp. 356–362, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.chb.2015.04.039>
- [73] I. Oakley and J. Park, "Did you feel something? Distracter tasks and the recognition of vibrotactile cues," *Interact. Comput.*, vol. 20, no. 3, pp. 354–363, 2008.
- [74] S. Wall and W. Harwin, "Quantification of the effects of haptic feedback during motor skills task in a simulated environment," in *Proc. Second PHANToM Users Res. Symp.*, 2000, pp. 61–69.
- [75] A. Tang, P. McLachlan, K. Lowe, C. R. Saka, and K. MacLean, "Perceiving ordinal data haptically under workload," in *Proc. 7th Int. Conf. Multimodal Interfaces*, New York, NY, USA: ACM Press, 2005, pp. 317–324. [Online]. Available: <http://portal.acm.org/citation.cfm?doi=1088463.1088517>
- [76] A. Chan, K. MacLean, and J. McGrenere, "Learning and identifying haptic icons under workload," in *Proc. Eurohaptics Conf., Symp. Haptic Interfaces Virtual Environ. Teleoperator Syst., World Haptics. First Joint, IEEE*, 2005, pp. 432–439.
- [77] B. Wark, B. N. Lundstrom, and A. Fairhall, "Sensory adaptation," *Curr. Opin. Neurobiol.*, vol. 17, no. 4, pp. 423–429, 2007.
- [78] Various, "Neural adaptation," *Wikipedia*, Sep. 2018, Page Version ID: 859794034. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Neural_adaptation&oldid=859794034
- [79] G. A. Gescheider and J. H. Wright, "Effects of sensory adaptation on the form of the psychophysical magnitude function for cutaneous vibration," *J. Exp. Psychol.*, vol. 77, no. 2, pp. 308–313, 1968. [Online]. Available: <http://doi.apa.org/getdoi.cfm?doi=10.1037/h0025746>
- [80] J. F. Hahn, "Vibrotactile adaptation and recovery measured by two methods," *J. Exp. Psychol.*, vol. 71, no. 5, pp. 655–658, 1966.
- [81] L. L. Langley, E. Cheraskin, and R. Sleeper, *Dynamic Anatomy and Physiology*, vol. 38. New York, NY, USA: McGraw-Hill, 1963.
- [82] T. J. Tighe and R. N. Leaton, *Habituation: Perspectives from Child Development, Animal Behavior, and Neurophysiology*. London, U.K.: Routledge, Jul. 2016.
- [83] Various, "Habituation," *Wikipedia*, Nov. 2018, Page Version ID: 868638885. [Online]. Available: <https://en.wikipedia.org/w/index.php?title=Habituation&oldid=868638885>
- [84] S. Brewster and A. King, "An investigation into the use of tactons to present progress information," in *Proc. Int. Fed. Inf. Process. Conf. Human-Comput. Interact.*, Springer, Berlin, Heidelberg, NY, USA 2005, pp. 6–17.
- [85] K. Kotowick and J. Shah, "Modality switching for mitigation of sensory adaptation and habituation in personal navigation systems," in *Proc. Conf. Human Inf. Interact.*, New York, NY, USA: ACM Press, 2018, pp. 115–127. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=3172944.3172980>
- [86] N. J. Mansfield, *Human Response to Vibration*. Boca Raton, FL, USA: CRC Press, Oct. 2004. [Online]. Available: <https://www.taylorfrancis.com/books/9781134459032>
- [87] O. S. Schneider, H. Seifi, S. Kashani, M. Chun, and K. E. MacLean, "HapTurk: Crowdsourcing affective ratings for vibrotactile icons," in *Proc. Conf. Human Factors Comput. Syst.*, New York, NY, USA: ACM Press, 2016, pp. 3248–3260. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2858036.2858279>
- [88] N. A. Cheever, L. D. Rosen, L. M. Carrier, and A. Chavez, "Out of sight is not out of mind: The impact of restricting wireless mobile device use on anxiety levels among low, moderate and high users," *Comput. Human Behav.*, vol. 37, no. C, pp. 290–297, Aug. 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.chb.2014.05.002>
- [89] S. Lakatos and R. N. Shepard, "Time—distance relations in shifting attention between locations on one's body," *Percept. Psychophys.*, vol. 59, no. 4, pp. 557–566, Jun. 1997. [Online]. Available: <http://www.springerlink.com/index/10.3758/BF03211864>
- [90] A. Gallace, H. Z. Tan, and C. Spence, "The body surface as a communication system: The state of the art after 50 years," *Presence Teleoperators Virtual Environ.*, vol. 16, no. 6, pp. 655–676, 2007.
- [91] G. Cirio, M. Marchal, A. Lécuyer, and J. R. Cooperstock, "Vibrotactile rendering of splashing fluids," *IEEE Trans. Haptics*, vol. 6, no. 1, pp. 117–122, Jan.-Mar. 2013. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6226398



Jeffrey R. Blum received the B.S.E. degree in computer science from Princeton University, Princeton, NJ, USA, in 1994. He is a Ph.D. candidate in the Shared Reality Laboratory, Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada. His research interests include mobile software design and implementation, which he has been involved with both in industry, at Microsoft and two startup companies, and now in academia. Mr. Blum is a student member of McGill's Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT), McGill University, Montreal, QC, Canada, and the Association for Computing Machinery (ACM), New York, NY, USA.



Pascal E. Fortin received the B.Eng. degree in electrical engineering from the University du Quebec a Chicoutimi, Chicoutimi, QC, Canada, in 2015. He is currently working toward the Ph.D. degree in Electrical and Computer Engineering with McGill University's Shared Reality Lab, Montreal, QC, Canada. His research interests include physiologically adaptive human-computer interfaces and affective computing. Mr. Fortin is a student member of the Centre for Interdisciplinary Research in Music Media and Technology and the Association for Computing Machinery (ACM), New York, NY, USA.



Feras Al Taha is currently working toward the B.Eng. degree in honours electrical engineering with a software engineering minor at McGill University, Montreal, QC, Canada. He is an Undergraduate Researcher with McGill University's Shared Reality Lab, Montreal, QC, Canada. His research interests include haptic information communication and haptic interfaces.



Parisa Alirezaee received the B.Sc. degree in software engineering from the University of Tehran, Tehran, Iran, in 2015, and is currently working toward the M. Eng. degree in Electrical and Computer Engineering with McGill University's Shared Reality Lab, Montreal, QC, Canada. Her research interests include human-computer interactions in healthcare and assistive technology. She is a student member of the Centre for Interdisciplinary Research in Music Media and Technology.



Marc Demers received the B.Eng. degree in mechanical and mechatronics engineering from the École Polytechnique de Montréal, Montreal, QC, Canada, in 2018 and is currently working toward the M. Eng. degree with McGill University's Shared Reality Lab, Montreal, QC, Canada. His current research interests lie at the intersection of human-computer interaction, haptics, and artificial intelligence.



Antoine Weill-Duflos received the M.Sc. and Ph.D. degrees in advanced system and robotics, in 2013 and 2017, respectively, from University Pierre and Marie Curie, Sorbonne Universities, Paris, France, in the Institut des Systèmes Intelligents et de Robotique. He also received the Mechanical Engineering degree (M.Eng.) from the top-ranked French Engineer School Arts et Métiers. He is currently working toward the Postdoctoral Fellowship with McGill University's Shared Reality Lab, Montreal, QC, Canada. His current research interests include haptic device design, robotics, haptic perception, and virtual reality. Mr. Weill-Duflos is a student member of the Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT), McGill University, Montreal, QC, Canada.



Jeremy R. Cooperstock received the B.A.Sc., M.Eng., and Ph.D. degrees in 1990, 1992, and 1996, respectively. He is a Professor with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada. His research interests focus on multimodal immersive systems and distributed computer-mediated communication. Dr. Cooperstock is an IEEE member and founding member of the Centre for Interdisciplinary Research in Music, Media and Technology. He is an Associate Editor for the *Journal of the Audio Engineering Society*, recipient of the award for Most Innovative Use of New Technology from ACM/IEEE Supercomputing, a Distinction Award from the Audio Engineering Society, and best paper awards from IEEE TRANSACTIONS ON HAPTICS, International Conference on Mobile and Ubiquitous Systems, and IEEE Haptics Symposium.