

Evaluating Smartwatch-based Sound Feedback for Deaf and Hard-of-hearing Users Across Contexts

Steven Goodman, Susanne Kirchner, Rose Guttman, Dhruv Jain, Jon Froehlich, Leah Findlater

University of Washington

Seattle, WA, USA

{smgoodmn, sukirch, rguttman, Leahkf}@uw.edu, {djain, jonf}@cs.washington.edu

ABSTRACT

We present a qualitative study with 16 deaf and hard of hearing (DHH) participants examining reactions to smartwatch-based visual + haptic sound feedback designs. In Part 1, we conducted a Wizard-of-Oz (WoZ) evaluation of three smartwatch feedback techniques (*visual alone*, *visual + simple vibration*, and *visual + tacton*) and investigated vibrational patterns (tactons) to portray sound *loudness*, *direction*, and *identity*. In Part 2, we visited three public or semi-public locations where we demonstrated sound feedback on the smartwatch *in situ* to examine contextual influences and explore sound filtering options. Our findings characterize uses for vibration in multimodal sound awareness, both for push notification and for immediately actionable sound information displayed through vibrational patterns (tactons). *In situ* experiences caused participants to request sound filtering—particularly to limit haptic feedback—as a method for managing soundscape complexity. Additional concerns arose related to learnability, possibility of distraction, and system trust. Our findings have implications for future portable sound awareness systems.

Author Keywords

Deaf and hard of hearing; sound awareness; smartwatches.

CSS Concepts

• **Human-centered computing** → *Empirical studies in accessibility, Accessibility technologies.*

INTRODUCTION

Advances in wearables and audio processing provide new opportunities for portable sound awareness solutions [12,26,34]. A recent survey of 201 Deaf and hard of hearing (DHH) participants [9] found that, compared to smartphones and head-mounted displays, smartwatches were the most preferred portable device for non-speech sound awareness. Further, smartwatches were seen as useful, socially acceptable, glanceable (for all sound scenarios except captions), and advantageous because they can provide both haptic and visual feedback.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Association of Computing Machinery.

ACM ISBN 978-1-4503-6708-0/20/04...\$15.00.

DOI: <http://dx.doi.org/10.1145/3313831.3376406>



Figure 1. Participants used a smartwatch for sound awareness in three contexts during an *in-situ* exploration: a student lounge (top), bus stop (left), and café (right).

Most prior work, however, has focused on smartphones and older handheld devices [2,27,42], head-mounted displays [13,18,34], and custom wearable systems that provide limited information through a single modality (*e.g.*, [20,45]). For smartwatches specifically, Mielke and Brück [30,31] conducted a preliminary lab study with six DHH participants using a Wizard of Oz interface: when the wizard triggered feedback, a simple vibration occurred and a visual sound was displayed. User reactions were generally positive but given the limited nature of that study, many design questions remain. For example, how should a design most effectively combine visual and haptic feedback on the smartwatch? And, is there a role for haptic feedback that is more complex than simple vibration?

Further, smartwatches and other portable sound feedback systems will need to function in a variety of complex soundscapes. Constant vibrational sound notifications are not desirable [36], for example, and some projects have examined filtering sounds based on identity [2,28] or loudness [43]. Beyond these initial steps, there has been little investigation into how to design sound feedback for complex soundscapes. How should sound filtering be designed, and what are the implications for filtering when both visual and haptic feedback modalities are present?

To address these questions, we conducted a three-part study with 16 DHH participants: (1) a Wizard-of-Oz evaluation of a smartwatch prototype comparing three designs that offer

visual feedback plus different levels of vibrational feedback (*none, simple, tacton*), and exploring haptic and visual techniques to portray three sound characteristics (*loudness, direction, and identity*); (2) an *in situ* experience (Figure 1) where participants visited three locations (café, bus stop, student lounge) and used the watch to receive a preset sequence of sounds typical of the location; (3) a semi-structured interview covering the user's overall experience, the feedback design, sound filtering options, and possible issues of privacy and social acceptability.

Our findings confirm the importance of combining visual and haptic feedback for sound awareness [9,30,31], but extend past work by showing that vibration is particularly important for push notifications to draw attention to visual details, and that the haptics have the potential to support more discreet and immediate sound alerting. Participants saw utility in vibration patterns (tactons), emphasizing this to be a promising direction for future work. In terms of soundscape complexity, the *in situ* experiences caused all participants to request sound filtering—particularly to limit haptic feedback—with varied advantages seen for filtering by sound identity, direction, or loudness. We also report on important concerns that will need to be addressed in future designs, including learnability of the tactons, the possibility of distraction, and issues of trust in the system.

This paper contributes: (1) a deeper understanding of the complementary roles of visual and vibrational feedback for a wearable sound awareness device; (2) evidence of the potential for small sets of haptic patterns to convey sound information; and (3) characterization of initial subjective responses to soundscape complexity and potential means of managing that complexity based on three pre-set locations. We also close with a discussion of design considerations and directions for future work.

RELATED WORK

We review prior work on the sound awareness needs of DHH users, visual and haptic approaches to sound awareness, and smartwatch usage and interaction.

Sound Awareness Needs

Designing effective sound awareness technology requires understanding the wide-ranging abilities and preferences of the DHH community. Several studies have surveyed DHH people on sounds of interest, highlighting a strong desire for urgent and safety-related sounds (*e.g.*, alarms, sirens) and social interaction support (*e.g.*, name calls, door knocks) [2,9,19,28,32,42]; Mielke *et al.*'s [31] preliminary study on smartwatch-based sound awareness reflects these trends.

Sound interest is also influenced by cultural and contextual factors. For example, people who prefer communicating orally are more interested in sound awareness—both overall and specifically for captioning—than those who prefer sign language [9,17]. DHH respondents in Findlater *et al.*'s survey [9] also predicted that social context (*e.g.*, with friends *vs.* strangers) would impact the use of a sound

awareness tool, and a majority desired to have sound filtering rather than being informed of all sensed sounds. Relatedly, sound awareness needs may differ by physical location, with past work asking interview or survey respondents about being at home, work, or mobile [2,28].

Researchers have also studied what sound *characteristics* are most desired, finding that some (identity, location, urgency) are generally more important than others (volume, duration, pitch) [2,9]. However, relative utility may differ by location or how the information is conveyed. For example, in the home, sound identity and location may be adequate [19], while directional indicators are important when mobile [32].

Our study is informed by the above work. Further, in contrast to studies that hypothetically asked about contexts of use [2,9,28], participants in our study experienced sound feedback in different settings with multiple filtering options.

Sound Awareness Technologies

Sound awareness technologies can be categorized as stationary, handheld (*e.g.*, smartphone or PDA), or wearable. Early HCI research focused on stationary designs such as desktop displays [15,28,29,44]. More recently, however, attention has shifted to portable tools: smartphone apps for environmental sounds [2,21,32,42] or automatic captioning [16,25], and both head-mounted [13,18,22,34] and wrist-worn wearable solutions [20,30,31,45]. User evaluations of these portable tools, when present, have been limited to the lab or a single environment (*e.g.*, a classroom setting [42]) and have highlighted the tools' potential to provide communication support [18,20,34] and alert to urgent situations [2,31,32]. They have not, however, probed key practical issues of how to manage soundscape complexity via different filtering options and the potential implications of such filtering—which our study begins to do.

In terms of feedback modalities, several studies recommend combining visual and vibrational information for sound awareness [2,9,23,30], and the ability to do so is seen as a strength of smartwatches [9,19,30,31]. User evaluations of prototypes that combine visual and haptic feedback, however, have been limited to using vibration as a secondary modality to draw attention to the visual information, and have not compared different approaches for combining the two modalities [2,30–32,42]. In contrast, we assess different combinations of visual and vibration feedback, including the use of tactons [4] to convey richer feedback via vibration.

Outside of HCI, wearable vibrotactile approaches without visual displays have been studied, often for sensory substitution of auditory information [6,10,45,46]. For example, Yeung *et al.* [45] transformed pitch information to vibro-patterns via a 16-channel tactile forearm display. However, obtrusive form factors (*e.g.*, waist-mounted [6], neck-worn [10]) made many of these devices impractical for everyday use. While early wrist-worn vibrotactile sound aids showed promise (*e.g.*, *Tactaid 7* [10], *TAM* [43]), frequent vibrational feedback had a high attentional cost, especially in




Figure 2. The Wizard-of-Oz prototype used for our lab evaluation. The wizard used a smartphone app (right) to remotely trigger visual (left) and vibrational feedback after sound events, such as a phone ringing (center).

noisy environments [36]—a concern we return to in our study. As tactile approaches provide much lower information throughput than visual approaches [8], the extent of sound information to convey haptically and its combination with visual feedback remain open research questions.

Smartwatches

A key benefit of smartwatches for sound awareness is that they are *mainstream* devices [9,30,32]—which can reduce stigmatization compared to dedicated assistive devices [41]. In this general context, smartwatches are most often used for activity tracking and receiving notifications from a paired smartphone [33,39]. Smartwatch interactions typically consist of only brief glances [38], so visual designs need to emphasize glanceability and space efficiency (e.g., [1,5]). Smartwatch-based haptics are typically employed for simple notifications: the watch vibrates, and the user can either ignore the alert or check the watch screen for detail [38]. However, more complex haptic output also shows potential: the wrist has high perceptual sensitivity to vibrotactile patterns [24]. Recently, smartwatches have been used as supplementary haptic displays for video games [35] and for passive learning of Morse code [40]. The design recommendations above and potential for rich haptic feedback on smartwatches influenced our design choices.

METHOD

To elicit user preferences for watch-based sound awareness, we employed a design probe method with 16 DHH participants. The method included a Wizard-of-Oz prototype evaluation in the lab, a demonstration of how such a system could work in practice in three *in situ* settings (e.g., in a café), and a semi-structured interview. We investigated haptic feedback preferences, sound filtering, contextual factors, privacy, and social concerns.

Smartwatch Prototype

To provide a realistic smartwatch-based sound awareness user experience and to enable us to compare different types of visual and vibrational feedback, we designed a Wizard-of-Oz prototype that consisted of two parts: a “wizard” interface running on an Android-based smartphone (Honor 7X) and a participant interface running on an Android-based




Figure 3. The smartwatch display shows the *direction, loudness, and identity* of sounds, such as: (a) a loud door knock in front of the wearer, (b) a moderate phone ringing to the left, and (c) a quiet name called to the right.




Figure 4. Visual illustrations used to introduce participants to the three tacton sets: *direction, loudness, and identity*. Lines of varying length indicate the relative duration of each vibration within a tacton. Tactons for four directions (a) were based on PocketNavigator [37], while three tactons for loudness (b: low, medium, high) and identity (c: door knock, phone ring, name call) were our own design.

smartwatch (Mobvoi Ticwatch E). The wizard interface could trigger events on the watch via Bluetooth (Figure 2).

Visual feedback. Informed by [1,29], we designed our visual feedback with a minimalist, high-contrast, and glanceable aesthetic (Figure 3). The display conveys three properties of sound: *direction* as three ~90° arcs pointed towards the sound source; *loudness*, which fills the directional arcs depending on sound amplitude (three discretized levels), and *identity*, which shows the classified sound event as text in the screen’s center. For this prototype, we implemented direction relative to the wearer’s torso (in front, behind, to the right, or to the left), assumed perfect sound classification, and did not support co-occurring sounds (only the loudest sound was shown). We return to these design decisions in the Discussion.

Haptic feedback. Past work has paired visual feedback with a simple vibration for notification [2,30–32,42]. In addition, we explore vibrational patterns (tactons), including how to best convey sound characteristics with tactons. Based on the capability of our off-the-shelf smartwatch (e.g., vibration output at fixed frequency and amplitude), we designed our haptic feedback as follows:

Simple vibration: A single 500ms vibration occurs with each sound event to notify the wearer.

Tacton sets (vibration patterns): Informed by Brewster and Brown’s [4] study of tactile icons (“tactons”), we separate small sets of tactons to convey each of the following sound characteristics: direction, loudness, and identity. Each tacton consisted of a pattern of on/off vibrations at a constant intensity and between 200–1200ms long. We defined all tactons in this way based on prior work showing DHH users

ID	Age	Gender	Cultural Identity	Self-reported Hearing Loss
P1	19	NB	deaf	Profound
P2	62	M	deaf	Profound
P3	53	M	deaf	Profound
P4	54	W	deaf	Profound
P5	33	W	Deaf	Profound
P6	46	M	Deaf	Severe
P7	51	W	Deaf	Profound
P8	56	M	deaf	Severe
P9	61	M	deaf	Severe
P10	61	M	HH	Moderate
P11	69	W	HH	Moderately severe
P12	86	M	HH	Moderately severe
P13	74	W	Deaf and HH	Profound
P14	69	W	Deaf	Profound
P15	69	M	Deaf	Profound
P16	27	W	Deaf	Profound

Table 5. Study participants. HH = hard of hearing.

prefer patterns over sustained vibration for attention-getting [14], and temporal patterns at a constant intensity level are easier to discern than patterns based on varying the intensity level [24]. For *direction*, we defined tactons for *left*, *right*, *front*, and *behind* based on Pielot *et al.*'s PocketNavigator [37], a tactile compass on a smartphone designed using the tactons framework [4]. We then created two more tacton sets of three patterns each: one set for loudness and one for sound (Figure 4); for example, a short-short-short vibration pattern indicated a door knock, while a short-long-short pattern indicated a name call.

Participants

We recruited 16 DHH participants through direct email, a hearing loss organization, and snowball sampling. Eight participants identified as men, seven as women, and one as non-binary. Participants were on average 55.6 years old ($SD=17.7$, range=19–84). Fourteen participants reported using hearing devices: eleven participants used hearing aids, five used cochlear implants, and two used both (Table 5).

Procedure

See supplementary materials for full detail on the procedure. Prior to the study session, participants completed an online questionnaire to collect demographics, current use of sound awareness technologies, important sounds in daily life, and initial reactions to new sound awareness solutions. The in-person procedure took place on a university campus and lasted ~90 minutes. Sessions were led by the first author, with one of three rotating members of the research team acting as a wizard. Participants were given the option to request communication support for the session: six opted for a sign language interpreter and two opted for a real-time captioner. Instructions and interview questions were presented visually on an iPad, while responses and follow-up discussion were spoken and translated to/from ASL.

The session consisted of three parts, the first and third of which took place in a quiet conference room. P11 had to leave unexpectedly at the conclusion of the lab-based design probe (Part 1) but returned to complete the remainder of the protocol 12 days later.

Part 1: Lab-based design probe (30 min): To give participants an idea of how a smartwatch-based sound awareness system could sense and convey different sounds, we presented a Wizard-of-Oz prototype in a lab setting. The participant sat at a conference table facing the door, with the facilitator on the opposite side and the wizard to the participant's left (Figure 2). After discussing current sound support strategies and soliciting reactions to the idea of smartwatch-based sound awareness, the participant placed the smartwatch on their preferred wrist (in contact with skin).

To gradually familiarize the participant with our Wizard-of-Oz prototype, we introduced three feedback designs in order of increasing complexity: *visual only*, *visual+simple vibration*, and *visual+tacton*. For the first two designs, we included a short description and three example sound events with different identity, loudness, and direction combinations (*e.g.*, Figure 3): (1) *door knock*, performed three times at high volume on a door in front of the participant; (2) *phone ring*, played at moderate volume to the left of the participant; and (3) *name call*, spoken at low volume while standing to the right of the participant. For each sound, the wizard remotely triggered the appropriate feedback on the watch.

For the third design (*visual+tacton*), we presented the three tacton sets (direction, loudness, identity) in counterbalanced order, with participants randomly assigned to an order (because there were 16 participants, two of the six orders only had two participants). For each tacton set, participants were given a visual reference sheet depicting the vibration patterns for the tactons (*e.g.*, Figure 4), as well as a demonstration of what each tacton felt like *without* any visual feedback. For this demonstration, the facilitator clapped at three different volumes in front of the participant (loudness), clapped at the same volume in four locations around the participant (direction), and created door knock, phone ring, and name calling sounds in front of the participant at constant volume (identity). Finally, the facilitator made the three example sounds described earlier (knock, ring, and name), with the wizard triggering both visual feedback *and* the appropriate tacton.

After being introduced to all three feedback designs (*visual only*, *visual+simple vibration*, *visual+tacton*), participants (1) rated the utility (*i.e.*, "useful in everyday life") of the three sound characteristics displayed on the watch in a set order (direction, loudness, identity); (2) rated the utility of each feedback design in random, counterbalanced order to minimize order effect from their demonstration; and (3) discussed each tacton set in the same order they were presented. Note: Because the study focused on the participant's subjective experience, we did not specify a hypothesis for this rating data and all quantitative results were considered secondary to the interview responses.

Part 2: Contextual design probe (25 min): To probe participants' responses to the effects of context on sound awareness, we visited three campus locations (student lounge, café, bus stop) and presented a preset sound scene at




Figure 6. Sound map used for the bus stop location in the contextual probe. These maps oriented the participant to the physical space and prepared them for the ten preset sounds that would be conveyed on the smartwatch. The facilitator guided the participant to the location and to face in the direction indicated by the purple face, then pointed out physical landmarks before initiating the smartwatch feedback.

each one. Locations were visited in a set order, following this scenario:

Imagine you are on your way home but forgot your water bottle in the student lounge upstairs. After you pick it up, you go to [the café] in the building next door to pick up some coffee and then go to the bus stop to catch the next bus home.

In each location, participants were shown a map on the iPad to orient themselves to the preset sound scene (e.g., bus stop map in Figure 6). Each map included ten numbered sounds typical of the area, with circles around the number to indicate loudness. After the participant had reviewed the map, the wizard triggered the watch to display the list of sounds in sequence, with three-second pauses between sounds. Because open-ended sound identification is an active area of research [11], we chose to have the watch visually convey loudness and direction but not identity. Instead, we instructed participants to view feedback from the watch and connect it to each potential real-world sound source as a holistic experience to ground discussion after returning to the conference room. Due to background noise, participants were asked to hold in-depth discussion until after the visits.

Additionally, to spur participants to consider different sound filtering options, we employed simple vibration feedback but varied how it worked across the three locations. Instead of having vibrations occur for all ten sounds, which could be overwhelming in practice, the vibration notification only occurred for either the top three loudest sounds, the three sounds occurring behind the participant, or three of the more important sounds identified by the watch. For the latter condition, we imagine that a personalized sound system such as that proposed by Bragg *et al.* [2] would support sound feedback by allowing the user to specify a small set of high priority sounds to identify. The pairing of contextual location (lounge, café, bus stop) with vibration for loudness,

direction, or identity was presented in counterbalanced order, with participants randomly assigned to an order.

Part 3: Semi-structured interview (20 min): Finally, we asked semi-structured questions on participants' overall experience with the system, exploring contextual factors, filtering options, social acceptability, and privacy issues surrounding smartwatch-based sound awareness.

Data and Analysis

Session transcripts were analyzed using an iterative coding approach [3]; see Supplementary Materials for codebook. Two researchers independently read the first four transcripts and identified a small set of potential codes to form high-level themes. The researchers then met and developed a mutually agreeable codebook with a two-level hierarchy to apply holistically to the data. As additional transcripts came in, the two researchers split the data by odd and even participant numbers and independently coded every other transcript. Upon receipt of the final two transcripts (P15 and P16), the two researchers agreed we had reached thematic saturation. Both researchers then randomly coded one of the other's transcripts to check for inter-rater reliability. Codes with low Cohen's *kappa* scores were used to identify areas of disagreement and were updated, merged, or removed from the codebook until 71 codes remained (10 first-level, 61 second-level). Each researcher applied the updated codebook to the other eight transcripts they had not yet analyzed. Following another inter-rater assessment, the codebook was considered final, with an average *kappa* of 0.72 ($SD = 0.14$) and raw agreement of 0.91 ($SD = 0.07$). Again, all code applications were reviewed, and disagreements were resolved through consensus.

For ordinal rating scale data, we used R to perform Friedman tests, a non-parametric alternative to repeated measures one-way ANOVAs. In cases of significance, Wilcoxon signed-rank tests with a Bonferroni correction were used for post-hoc pairwise comparisons.

FINDINGS

Participants reported using their hearing aids and/or cochlear implants as well as smartphones for sound awareness—the latter primarily for automatic captioning ($N=8$), such as via Android's *Live Transcribe*. When discussing hearing aids and cochlear implants, participants described well-known limitations [7], including discriminating between sounds ($N=10$), inadequate background filtering (10), and poor speech comprehension (6). Sounds of interest reflected past work (e.g., [2,9,29]): social sounds and alerts were important, while indoor sounds (e.g., doors, typing) and background noise (e.g., birds, traffic) were less desired.

Before presenting our prototype, we asked participants to share their thoughts on using a wearable device for sound awareness. Similar to findings by Findlater *et al.* [9], most participants responded positively: nine were *very* or *extremely* interested, while the remainder were only *somewhat* ($N=5$) or *slightly* (2) interested. All but one

participant (P14) saw potential in using a smartwatch for this purpose, although 10 participants also predicted limitations, most notably the small screen size. Below, we cover findings from the in-lab comparison of our three feedback designs, discuss themes that emerged through the *in situ* experiences, and synthesize topics overlapping both parts of the study.

Part 1: In-Lab Comparison

Participants used three feedback designs on the smartwatch: *visual alone*, *visual+simple vibration*, and *visual+tactons*; the latter design separated into three tacton sets to convey sound identity, direction, and loudness. We report reactions to the three designs, followed by confirmatory findings on general sound feedback preferences.

Reactions to and Preferences for the Three Designs

Overall, participants felt the designs with vibration were more useful than the visual-only design. Perceived utility ratings for all three designs are shown in Figure 7a, and a Friedman test showed that the impact of these design options on utility ratings was significant ($\chi^2_{(2,N=16)} = 6.107, p < .05$). No post-hoc pairwise comparisons were significant after a Bonferroni correction, but the qualitative findings below provide insight on the significant main effect.

Visual Alone. Eight participants mentioned unprompted that visual feedback offers higher information throughput compared to vibration feedback, and that the visual-only design also offers “*the option to not want to be bothered [by vibration]*” (P13). Thirteen participants, however, were concerned that without vibration, they would miss sounds: “*I’m not going to be looking at my watch every two minutes when I’m out and about*” (P15). Still, across all conditions, participants reaffirmed the importance of visual sound information, such as: “*It’s nice to have visual and the sensory input as well [but] I mean without the visual, I feel like there’s not really a point.*” (P10).

Visual with Simple Vibration. For *visual+simple vibration*, the key advantage identified by all but one participant (P12) was how it could push sound notifications to the wearer. For example, P9 liked that it would support “*alerting to a situation*”, while P5 said it would “*trigger you to look at your watch*”. However, frequent notifications were a concern ($N=4$), as captured by P16, “*I don’t want it to be constantly vibrating because it’s a noisy world.*” In terms of the specific vibration design, three participants wished it were more prominent, for example, “*because I think if you’re outside moving around or doing something, you might not feel it as much*” (P15). This finding suggests that the vibration intensity should be adjustable.

Visual with tactons. Participants felt that the primary benefit of tactons was in minimizing the need to look at the watch face, which eight participants viewed as speeding up their response time. Providing more non-visual detail than a simple vibration would also allow the wearer to, “*...determine whether it was worth looking at*” (P8) and could provide sound support in a socially acceptable manner:

“Let’s say I’m at a meeting, and I’m trying to listen to somebody, but I don’t want to be rude and look at my watch when they’re talking. [...] ‘Ah-ha!’ Somebody is calling my name, and then I don’t have to look [at the watch].” (P3)

At the same time, important design concerns arose. Some participants ($N=8$) worried about the effort and time required to use tactons. For example, P13 commented on the difficulty of interpreting tactons while physically active: “*unless you’re just sitting here, it’s going to be hard to know what’s happening.*” Further, P6 commented on the time required to absorb tacton-based information: “*I had to wait and wait and then the vibration had finished and by the time I finished decoding what it was [...] I’d missed whatever happened.*”

Participants were also concerned about learning tactons ($N=9$). P11 described this challenge while also appreciating the possible long-term benefits, saying, “*It might take you a while to learn it, but after you learn it then it would be automatic...*” As another example, P14 contrasted the difficulty of learning tactons to the ease of the visual feedback: “*It would work after I got used to it, whereas just the visual you wouldn’t need to get used to it.*”

To mitigate learning and recall problems, several participants ($N=8$) suggested using only a small, simple set of tactons: “*I think three is enough to memorize, but [more] would be hard to distinguish*” (P16). Another approach was to more intuitively match the tacton design to the semantics of the sound ($N=4$): “*There’s no reason to assume the door knock’s three quick bursts, while the phone is three longer bursts, and a name being called out is this other pattern*” (P2).

Finally, after participants used the three tacton sets (*i.e.*, for sound identity, for loudness, and for direction), we asked which of those three sound characteristics participants would most like conveyed via tactons. Responses were split between identity ($N=8$) and direction ($N=6$), with only two participants preferring loudness. For example, P10 chose tactons for sound identity, saying, “*If it’s a phone call and I’m busy right now, I can ignore it. Whereas, if it’s my wife calling [my name], I better check that out*” (P10). P16, in contrast, felt that direction was most immediately actionable, stating: “*When someone calls me [...] I don’t want to have to look at the watch and then look toward the sound.*”

Sound Characteristics in General

As shown in Figure 7b, most participants felt that all three sound characteristics were useful but to differing extents: a Friedman test showed that there was a significant effect of sound characteristic on utility ratings ($\chi^2_{(2,N=16)} = 8.24, p < .05$). After a Bonferroni correction, no post-hoc pairwise comparisons were significant, but the significant main effect adds evidence to past work showing that sound identity and direction are of greater interest than loudness [9,28].

Participants also provided open-ended reasoning for their ratings. While participants were generally positive about receiving sound identity information, some ($N=3$) expressed concerns about accuracy—a concern that also arose later in

discussions about sound filtering. For direction, three participants felt that this information would be useful because they can identify sounds by hearing but “*the direction is much harder to pick out*” (P2). Finally, participants mentioned specific ways in which loudness could be useful, most commonly related to safety ($N=5$): “*when something is loud, then it's something that you want to be alerted for*” (P13) and “[*...a smoke alarm*] is loud for some people but to me it's not” (P3).

Summary. Participants preferred to have vibrational and visual feedback rather than visual alone. Both vibrational designs show promise, though tactons may need to be limited to a small, simple set.

Part 2: Physical Contexts of Use

Following the lab evaluation, participants visited three locations to experience how the prototype might work in context: a student lounge, café, and bus stop. These visits helped participants consider additional aspects of sound feedback, with 11 participants mentioning new use cases and/or increased interest in smartwatch sound awareness. Emergent discussions focused on soundscape complexity and safety, primarily in public or semi-public contexts—but private use in the home also arose.

Soundscape Complexity. After visiting the three locations, participants remarked on the diversity and number of available sounds, as captured by P15:

“There's a huge variety of things that you could need to be aware of. [...] You're out there, you don't think about them. But in designing something that could be useful, you do have to think about it. And the complexity of it, of what the environment is, this is eye-opening.”

Despite the challenge of designing for this complexity, busy public locations may be particularly important for sound awareness support. For example, P14 returned from the visits feeling more positive about the watch than she had in the lab:

“The [café] is just phenomenal because it's the thing that really gives people anxiety. ‘Are they going to hear me? Am I going to hear them?’ There's so much ambient noise. In a place like [the student lounge] or in your house with the microwave and whatever, okay, it's quiet. But when you go to a place outside, bus stop, [café], outside your home, this is just... and again in your car, this is just incredible.”

As a result of their *in situ* experiences, four participants changed how they thought about haptic feedback. For P6, the value of the haptic tactons increased: “*There's just a lot going on, and so if it was a short vibration, I could know to ignore it.*” Similarly, for P14, who went from having no interest in vibration before the visits to saying, “*I realized that even though I do hear the sound, I want that vibration.*” Conversely, the complexity of the soundscape gave P8 and P16 a greater appreciation for the option to “*mute*” (P8) the vibration and only see visual information.

Situational Awareness and Safety. A second set of reactions emerged around situational awareness and safety. All participants except P4 liked the watch for mobile use,

with many ($N=9$) emphasizing personal safety. For example, P16 discussed the utility of sound notifications when walking alone at night: “*I want to know if there's some sort of noise, if someone might be following me.*” Similarly, P3 mentioned using the watch for traffic awareness: “*It could alert you when there's a hazard, like if I'm riding my bike: ‘There's a car coming.’*” Five participants mentioned using the watch during outdoor recreation, such as “*thunder*” (P3) and “*a mountain lion*” (P6) while hiking. P6 also imagined the watch could warn of danger in the warehouse where he works (“*say a shelf of product just fell down*”).

Not everyone agreed on the watch's value for situational awareness, however, with P4 expressing concern that the watch could be distracting and thus reduce safety: “*I would never use something like this to tell me about traffic. Ever. [...] Taking time or being distracted by vibration. To look at the watch, it takes me away from my environment.*”

Many participants ($N=8$) also raised the idea of using the watch in a professional or classroom setting, primarily to aid in social participation. For example, P16 thought the watch could help her participate in class: “*Sometimes my teacher will call my name, but I don't notice that it's happened, and then I miss the question that's being directed at me.*” and P7 wanted to discriminate “*softer versus louder*” sounds in her work with musicians.

Usage in the Home. Finally, while not directly asked, all participants mentioned potential benefits in the home. Almost half ($N=7$) introduced emergency alerts while sleeping. P4 said, “*To be able to go to bed, put this on, and know if somebody was trying to break down the door or the fire alarm is going off, or maybe the baby is crying.*” Other notable home uses for the watch included awareness of family voices (7) and responding to non-urgent sounds (6), like appliance alerts. For example, P15 recounted a story about forgetting to turn off his alarm clock one day, “*and both neighbors on both sides of me in the houses, they checked to see that everything was okay. I was mortified. But if I'd had a watch, it would've said, ‘Hey, there's a sound going on,’ and that would really have been nice to have.*”

Summary. The *in situ* experience with the watch highlighted sound support challenges in busy, public contexts. The watch showed promise for safety while mobile and at home, and for social support in school/professional settings. However, negative impacts of distraction need to be considered.

Part 3: Synthesizing Cross-Cutting Themes

Finally, we present cross-cutting themes that emerged across the entire study session, including sound filtering, social contexts of use, privacy concerns, and design suggestions.

How to Deal with Soundscape Complexity: Filtering

Upon returning to the lab, all participants were against conveying every detected sound, reflecting past work [9,28,31]. For example, P4 said, “*I don't think I would want [the watch] constantly telling me every sound that came in with directions and arrows,*” while P15 said, “*I think there*




Figure 7. Utility of (a) haptic complexity, (b) sound characteristics, and (c) filtering options, with conditions ordered by usefulness.

needs to be some way to filter what you do want to pay attention to. And that's going to differ for everybody." A few participants ($N=3$) mentioned that they would want to have visual feedback for all sounds but that vibrational feedback was different. P7, for example, linked the need for filtering to the ability of hearing people to ignore or attend to sounds:

"[Hearing people] have the ability to screen out the sounds because you guys are used to hearing. [...] The vibration, I think, could be a lot for me because it doesn't actually have the ability to filter out the sounds, which is why I prefer to see the visual and I can pay attention to it and then decide."

Importantly, one participant (P4) was hesitant about adding filtering at all, expressing concern about allowing the device to choose what to filter:

"You might be filtering out other awareness that you have built up over years in favor of, 'Well, this thing knows, and in fact this thing might know better than me, so I'm just gonna ignore my instinct, I'm not going to bother looking because this will tell me.' [...] I want to hear it all, and I want my own, I want to be able to choose what's more important."

Filtering as a Function of Sound Characteristic. While prior work has explored notifications for noise above a certain threshold [43] or specific sounds [2,28], our study is the first to compare several filtering options. At the three *in situ* locations, participants experienced what it would be like to filter vibration notifications based on loudness, direction, or sound identity. Utility ratings for these three filtering options (Figure 7c) did not differ significantly, $\chi^2_{(2,N=16)} = 3.05$, $p > 0.05$, although qualitative comments highlight tradeoffs and possible applications of each.

All but one participant (P4) wanted to filter by sound identity, reflecting past work [2,28]. Many discussed how they would prioritize identified sounds; for example, P2 was excited to see nature sounds: "I think different people want to know about different sounds. I would like to pick up, like, bird calls, bird chirping." Three participants, however, were concerned that filtering sounds by type would be technically infeasible, especially in contexts where sounds are unpredictable. P4, for example, said of the café: "You [the user] can't program in breaking of glass because you wouldn't know that was gonna happen."

Most participants responded positively to filtering by direction ($N=13$), especially for sounds occurring behind

them: "The deaf population, we see things ahead of us and know what's happening, pretty much" (P8) but "things happening behind me, that would be desirable [to know]" (P13). Safety was often given as the primary reason to filter by direction ($N=5$).

Finally, eleven participants wanted to filter by loudness, emphasizing both the relationship between volume and importance, as well as the need to consider different ambient noise levels. P1 was enthusiastic about loudness filtering, while recognizing its limitations:

"It might be annoying because there are a lot of loud sounds that aren't necessarily that important, [...] but at the same time, loud sounds are often loud for a reason, so I feel like it's still necessary."

Other participants felt loudness filtering would be useful only in certain contexts; for example, loudness could be useful for P7 in "a music room", while for P5, being notified of loud sounds while at work could be problematic because, "I work close to a fire station." As another example, P6 experienced loudness filtering at the café and was concerned that the loud ambient noise would create distracting notifications: "if I'm paying attention to my watch and it keeps vibrating, I might miss my drink come up."

Summary. Participants requested that sound feedback be filtered, and all three types of filtering (identity, loudness, direction) had value. However, questions arose over being able to predict what sound identities to filter and whether to trust the device's filtering decisions.

Social Contexts of Use

In terms of social acceptability, all but one participant (P6) felt they would be comfortable using the smartwatch around other people, although some key considerations arose. In general, participants commented about not caring what others thought ($N=9$) and/or that they were excited to show the watch off ($N=6$). For example, P1 said, "Any tool to help me access my surroundings is better than no tool, and, honestly, if people think it's weird that's their issue," while P2 said, "it wouldn't make you stand out, because most people are accustomed to some people wearing watches." The watch was also seen as useful for spoken communication ($N=11$), for example, "If someone's behind me in a store, and they say, 'Excuse me,'" (P14) and "being able to pick up where a voice first starts coming in from" (P2).

That said, eight participants mentioned how social context may impact usage. P3 said, for example, “I’d be more likely to use it when I’m alone, because when I’m with my friends or my family, then I would depend on them.” Use in a Deaf cultural context was also discussed, with P16 acknowledging that a sound awareness technology may not be appropriate around other Deaf people, a finding that reflects past work [9]. Similarly, P5 said: “No matter who I’m with, I’d like to have the environmental information. [...] But] with hearing people, they’re trying to speak with you. Whereas around deaf people, they’re signing.” Finally, six participants mentioned not wanting to negatively impact others by appearing distracted by the watch or vibration notifications.

Privacy Concerns

We asked participants if they had any privacy concerns using a watch with an always-on microphone—none did. P15 argued that the smartwatch supports user privacy because of its unobtrusive and commonplace interaction: “Because all you have to do is turn your wrist.” We did not probe for data collection and security issues, and the topic emerged for only two participants, such as: “Where does all that information go? I would be interested in knowing” (P12).

Design Suggestions

Participants provided design suggestions throughout the study, including user interface ideas, alternate haptic methods, and customization support. Though participants were not asked to critique our visual design, responses were generally positive—e.g., P14 found it, “very easy to read” and “very clean.” For new user interface ideas, participants suggested icons, having the screen flash or change brightness to indicate some sound characteristic, and sound history support. Adding color ($N=7$) was most commonly requested to improve glanceability and encode other sound information. For example, P10 compared color to words, saying, “it’s faster. [...] Like if it’s blue, it’s like, ‘this is happening.’ If it’s red, ‘this may be important’,” while P6 mentioned that color would be good “for privacy reasons.” Participants were also invited to sketch out design ideas; five chose to do so. Figure 8 shows three examples, including: using iconography, color, and visualizing co-occurring sounds. There were fewer haptic-related ideas; the most common suggestion was to adjust vibration intensity to convey sound information ($N=5$). Other ideas included using Morse code for tactons, providing directional information through multiple vibration motors, and continuously vibrating for emergencies.

For customizability, participants wanted to prioritize specific sounds, switch between filtering options, create personal tacton sets, and add preset settings for specific contexts. When participants were asked if the watch should automatically adapt to different contexts, responses were mixed. Several participants ($N=7$) were receptive to the idea, suggesting that the watch could adapt based on location, noise level, or the wearer’s activity. P8 said, for example, “if it knows I’m driving, that would be great.” Some participants who were against the idea were worried about being




Figure 8. Three example participant design sketches: icons for sound identities, such as a telephone (left, P5), use of colors and other changes (center with color words emphasized, P1), and co-occurring sounds (right, P6).

confused by the changes ($N=3$) and/or cited negative experiences with automatic adjustment in other devices ($N=4$). P5, for example, mentioned that her hearing aid had “...just kept changing on its own and filtering out sounds.” To mitigate these issues, several participants suggested being able to override any automatic changes ($N=4$).

Summary. The smartwatch form factor generally met expectations of privacy and social acceptability, but use may vary by social context (confirming [9]). Glanceable visuals are preferred, for which icons or color may be useful. Customization was deemed important, although opinions were split over whether settings should adjust automatically.

DISCUSSION

This study confirms DHH users’ preferences for having both visual and haptic feedback in a wearable sound awareness system [9,30,31], but also: (1) extends our understanding of how to design these feedback modalities in combination, (2) demonstrates the potential for small sets of haptic patterns, and both (3) highlights user reactions to soundscape complexity in busy environments as well as (4) identifies promising methods for filtering that complexity (i.e., based on sound characteristics and context). Here, we reflect on how to combine visual and haptic feedback for smartwatch sound awareness feedback, considerations for managing soundscape complexity, and limitations of our work.

Complementary Roles of Visual and Haptic Feedback

Visual and haptic feedback offer complementary roles for wearable sound awareness systems, and their combination provides users with flexibility. Advantages of visual feedback include high information throughput and ease of understanding. However, the small smartwatch screen is limiting, so simple and glanceable designs are preferred. Suggestions for increasing glanceability include using icons or color to encode sound identity—changes that could also be designed to preserve the wearer’s privacy in the presence of others. Further, the visual designs in our study showed only a single sound at a time, with no notion of history. How (and if) to provide this more complex information on the watch is an open question.

Our study shows that haptic feedback is critical too. DHH people make strong use of visual cues for environmental awareness [28]. Haptic notifications (whether simple or pattern-based) are thus important because they get the user’s

attention without interfering with existing visual awareness strategies. A related benefit is that haptic feedback could be particularly useful for safety-related notifications while the wearer is sleeping. Despite these positives, however, overly frequent or obtrusive vibrations were seen as problematic, reflecting early work on tactile sound awareness systems [36]. In noisy situations especially, it may be best to allow users to turn off or reduce the haptic feedback and to explore visual descriptions of the soundscape.

Many projects [14,37] have explored tactons for haptic communication, though our work is the first to apply them to sound awareness. While participants expressed concern about the learning curve and the time required to interpret a tacton, the overall response was positive. Due to the limited control we had over our smartwatch’s vibration output, our designs were preliminary, meant to assess the general idea of using tactons. Thus, future work will need to focus on more specific design attributes, such as intuitive tacton sets. Our findings suggest that tactons should be limited to a small set, and that, due to individual preferences for what information to convey via tactons (loudness, identity, direction), users should have the ability to configure how they are used.

How to Manage Soundscape Complexity

The sheer complexity of the *in situ* soundscapes impacted participant responses to sound awareness feedback. While a previous survey showed that only 63% of DHH respondents predicted they would want sounds filtered [9], *all* participants in our study desired filtering—a disparity we attribute to our participants’ exposure to realistically complex soundscapes. To manage this complexity, past work has limited notifications to specific sounds [27] or, for vibrotactile feedback alone, to sounds above a loudness threshold (e.g., [43]). Our study is the first to compare different filtering options, and to examine filtering based on direction. Positive responses to all three options (identity, loudness, direction) and varied ideas about how each would be useful suggest that future work should continue to evaluate these options and to further refine their designs. Specifically, future work should examine more realistic sound pacing than our study’s consistent stimuli pace (one sound every 3 seconds) and should explore the ability to switch between filtering presets depending on the context.

The feasibility of the filtering designs we evaluated is an important factor to consider. Filtering based on *loudness* can be done with a simple threshold loudness level, though our findings suggest that the threshold may need to automatically adapt to background noise levels and/or be controllable by the user. Sensing sound *direction*, in contrast, is more difficult and would likely need additional hardware, such as a wearable microphone array [20]. Reliably *identifying* open-ended sounds is also complicated by factors such as overlapping sounds, background noise, and differences across locations [2,11]. With respect to sound identification, our study focused on the idea of being able to identify a small set of sounds reliably, such as Bragg *et al.* [2].

Automatic filtering also introduces ethical and practical considerations of how much trust a DHH user should put in a sound awareness system and what constitutes an appropriate and accurate representation of the surrounding soundscape. Trust, for example, was highlighted by P4 in our study, who preferred to rely on her existing sound awareness strategies than to trust a system in unfamiliar locations. An important complication is that a DHH user may have limited means of judging the sound awareness system’s accuracy for themselves, such as noticing that there are errors in identification or filtering. While researchers will need to continue grappling with these issues, system transparency offers a potential path forward: systems should be transparent in how they make decisions, should provide real-time information about what is being filtered/identified, and allow the user to modify those decisions as necessary.

Limitations

We enumerate three primary limitations. First, our volunteer participants may have been biased toward sound awareness technologies compared to others in the highly diverse DHH population; for example, larger surveys of DHH participants show that some segments of the population are less interested than others [2,9]. Second, limited exposure to tactons may have reduced their perceived utility compared to if participants had had more time to learn and use them. Third, our *in situ* exploration was brief, did not show real sounds, and occurred within a small radius. While this setup allowed us to identify new considerations that purely lab-based evaluations had not previously seen, work in other contexts and study of longer-term use is needed to understand the tool’s broader utility and adoption/abandonment issues.

CONCLUSION

Our work explored DHH users’ preferences toward smartwatches employed for sound awareness and implications of their use *in situ*. In a qualitative study with 16 DHH participants, we conducted a Wizard-of-Oz prototype evaluation in the lab, exploring options for combined visual and haptic sound feedback, then presented brief sound scenarios in three *in situ* locations, demonstrating different options for filtering. Confirming an overall preference that portable sound awareness systems include both visual and haptic feedback, we also characterize vibration utility, both for push notification, and as actionable sound information displayed through vibrational patterns (tactons). All participants requested filtering—particularly to limit haptic feedback—as a method for managing soundscape complexity, with varied advantages seen for filtering by sound identity, direction, or loudness. These findings should influence the design of smartwatch-based systems, but also sound-awareness technologies generally.

ACKNOWLEDGMENTS

We thank Augustina Liu and Lucy Jiang for support during study sessions, and Kelly Mack for help with figures. This research was funded in part by the National Science Foundation under grant IIS-1763199 and by a Google Faculty Research Award.

REFERENCES

- [1] Tanja Blascheck, Lonni Besancon, Anastasia Bezerianos, Bongshin Lee, and Petra Isenberg. 2019. Glanceable Visualization: Studies of Data Comparison Performance on Smartwatches. *IEEE Trans. Vis. Comput. Graph.* 25, 1 (January 2019), 630–640. DOI:https://doi.org/10.1109/TVCG.2018.2865142
- [2] Danielle Bragg, Nicholas Huynh, and Richard E Ladner. 2016. A Personalizable Mobile Sound Detector App Design for Deaf and Hard-of-Hearing Users. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '16)*, 3–13. DOI:https://doi.org/10.1145/2982142.2982171
- [3] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* 3, 2 (2006), 77–101.
- [4] Stephen Brewster and Lorna M Brown. 2004. Tactons: Structured Tactile Messages for Non-visual Information Display. In *Proceedings of the Fifth Conference on Australasian User Interface - Volume 28 (AUIC '04)*, 15–23.
- [5] Yang Chen. 2017. Visualizing Large Time-series Data on Very Small Screens. In *EuroVis 2017 - Short Papers*. DOI:https://doi.org/10.2312/eurovisshort.20171130
- [6] Mohammad I. Daoud, Mahmoud Al-Ashi, Fares Abawi, and Ala Khalifeh. 2015. In-house alert sounds detection and direction of arrival estimation to assist people with hearing difficulties. In *2015 IEEE/ACIS 14th International Conference on Computer and Information Science (ICIS)*, 297–302. DOI:https://doi.org/10.1109/ICIS.2015.7166609
- [7] Harvey Dillon. 2008. *Hearing aids*. Hodder Arnold.
- [8] Johan Engström, Nina Åberg, Emma Johansson, and Jakob Hammarbäck. 2005. Comparison Between Visual and Tactile Signal Detection Tasks Applied to the Safety Assessment of In-Vehicle Information Systems. In *Driving assessment 2005 : proceedings of the 3rd International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design*, 232–239. DOI:https://doi.org/10.17077/drivingassessment.1166
- [9] Leah Findlater, Bonnie Chinh, Dhruv Jain, Jon Froehlich, Raja Kushalnagar, and Angela Carey Lin. 2019. Deaf and Hard-of-hearing Individuals' Preferences for Wearable and Mobile Sound Awareness Technologies. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–13. DOI:https://doi.org/10.1145/3290605.3300276
- [10] Karyn L. Galvin, Jan Ginis, Robert S. C. Cowan, Peter J. Blamey, and Graeme M. Clark. 2001. A Comparison of a New Prototype Tickle Talker™ with the Tactaid 7. *Aust. New Zeal. J. Audiol.* 23, 1 (May 2001), 18–36. DOI:https://doi.org/10.1375/audi.23.1.18.31095
- [11] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio Set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 776–780. DOI:https://doi.org/10.1109/ICASSP.2017.7952261
- [12] L. Giuliani, L. Brayda, S. Sansalone, S. Repetto, and M. Ricchetti. 2017. Evaluation of a complementary hearing aid for spatial sound segregation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 221–225. DOI:https://doi.org/10.1109/ICASSP.2017.7952150
- [13] Benjamin M. Gorman. 2014. VisAural: A Wearable Sound-Localisation Device for People with Impaired Hearing. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS '14)*, 337–338. DOI:https://doi.org/10.1145/2661334.2661410
- [14] J. Harkins, P. E. Tucker, N. Williams, and J. Sauro. 2010. Vibration Signaling in Mobile Devices for Emergency Alerting: A Study With Deaf Evaluators. *J. Deaf Stud. Deaf Educ.* 15, 4 (October 2010), 438–445. DOI:https://doi.org/10.1093/deafed/enq018
- [15] F. Wai-ling Ho-Ching, Jennifer Mankoff, and James A. Landay. 2003. Can You See What I Hear?: The Design and Evaluation of a Peripheral Sound Display for the Deaf. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*, 161–168. DOI:https://doi.org/10.1145/642611.642641
- [16] InnoCaption. Real Time Mobile Captioning for Hearing Loss. Retrieved August 12, 2019 from https://www.innocaption.com/
- [17] Dhruv Jain, Leah Findlater, Jamie Gilkeson, Benjamin Holland, Ramani Duraiswami, Dmitry Zotkin, Christian Vogler, and Jon E Froehlich. 2015. Head-Mounted Display Visualizations to Support Sound Awareness for the Deaf and Hard of Hearing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, 241–250. DOI:https://doi.org/10.1145/2702123.2702393
- [18] Dhruv Jain, Rachel Franz, Leah Findlater, Jackson Cannon, Raja Kushalnagar, and Jon Froehlich. 2018. Towards Accessible Conversations in a Mobile Context for People Who Are Deaf and Hard of Hearing. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '18)*, 81–92. DOI:https://doi.org/10.1145/3234695.3236362
- [19] Dhruv Jain, Angela Lin, Rose Guttman, Marcus Amalachandran, Aileen Zeng, Leah Findlater, and Jon Froehlich. 2019. Exploring Sound Awareness in the

- Home for People who are Deaf or Hard of Hearing. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–13. DOI:<https://doi.org/10.1145/3290605.3300324>
- [20] Yoshihiro Kaneko, Inho Chung, and Kenji Suzuki. 2013. Light-Emitting Device for Supporting Auditory Awareness of Hearing-Impaired People during Group Conversations. In *2013 IEEE International Conference on Systems, Man, and Cybernetics*, 3567–3572. DOI:<https://doi.org/10.1109/SMC.2013.608>
- [21] Hamed Ketabdardar and Tim Polzehl. 2009. Tactile and visual alerts for deaf people by mobile phones. In *Proceeding of the eleventh international ACM SIGACCESS conference on Computers and accessibility - ASSETS '09*, 253–254. DOI:<https://doi.org/10.1145/1639642.1639701>
- [22] Ki-Won Kim, Jung-Woo Choi, and Yang-Hann Kim. 2013. An Assistive Device for Direction Estimation of a Sound Source. *Assist. Technol.* 25, 4 (October 2013), 216–221. DOI:<https://doi.org/10.1080/10400435.2013.768718>
- [23] Raja S. Kushalnagar, Gary W. Behm, Joseph S. Stanislaw, and Vasu Gupta. 2014. Enhancing caption accessibility through simultaneous multimodal information. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS '14)*, 185–192. DOI:<https://doi.org/10.1145/2661334.2661381>
- [24] Seungyon “Claire” Lee and Thad Starner. 2010. BuzzWear: Alert Perception in Wearable Tactile Displays on the Wrist. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*, 433–442. DOI:<https://doi.org/10.1145/1753326.1753392>
- [25] Live Caption. Live Caption. Retrieved August 12, 2019 from <http://www.livecaptionapp.com/>
- [26] Patrizia Marti, Iolanda Iacono, and Michele Tittarelli. 2018. Experiencing sound through interactive jewellery and fashion accessories. In *Congress of the International Ergonomics Association*, 1382–1391.
- [27] Tara Matthews, Scott Carter, Carol Pai, Janette Fong, and Jennifer Mankoff. 2006. Scribe4Me: Evaluating a Mobile Sound Transcription Tool for the Deaf. In *Proceedings of the 8th International Conference on Ubiquitous Computing (UbiComp '06)*, 159–176. DOI:https://doi.org/10.1007/11853565_10
- [28] Tara Matthews, Janette Fong, F. Wai-Ling Ho-Ching, and Jennifer Mankoff. 2006. Evaluating non-speech sound visualizations for the deaf. *Behav. Inf. Technol.* 25, 4 (July 2006), 333–351. DOI:<https://doi.org/10.1080/01449290600636488>
- [29] Tara Matthews, Janette Fong, and Jennifer Mankoff. 2005. Visualizing non-speech sounds for the deaf. In *Proceedings of the 7th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '05)*, 52. DOI:<https://doi.org/10.1145/1090785.1090797>
- [30] Matthias Mielke and Rainer Bruck. 2016. AUDIS wear: A smartwatch based assistive device for ubiquitous awareness of environmental sounds. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 5343–5347. DOI:<https://doi.org/10.1109/EMBC.2016.7591934>
- [31] Matthias Mielke and Rainer Brück. 2015. A Pilot Study about the Smartwatch as Assistive Device for Deaf People. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility (ASSETS '15)*, 301–302. DOI:<https://doi.org/10.1145/2700648.2811347>
- [32] Matthias Mielke and Rainer Brück. 2015. Design and evaluation of a smartphone application for non-speech sound awareness for people with hearing loss. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 5008–5011. DOI:<https://doi.org/10.1109/EMBC.2015.7319516>
- [33] Chulhong Min, Seungwoo Kang, Chungkuk Yoo, Jeehoon Cha, Sangwon Choi, Younghan Oh, and Junehwa Song. 2015. Exploring current practices for battery use and management of smartwatches. In *Proceedings of the 2015 ACM International Symposium on Wearable Computers - ISWC '15*, 11–18. DOI:<https://doi.org/10.1145/2802083.2802085>
- [34] Yi-Hao Peng, Ming-Wei Hsi, Paul Taelle, Ting-Yu Lin, Po-En Lai, Leon Hsu, Tzu-chuan Chen, Te-Yen Wu, Yu-An Chen, Hsien-Hui Tang, and Mike Y Chen. 2018. SpeechBubbles: Enhancing Captioning Experiences for Deaf and Hard-of-Hearing People in Group Conversations. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, 293:1--293:10. DOI:<https://doi.org/10.1145/3173574.3173867>
- [35] Erik Pescara, Alexander Wolpert, Matthias Budde, Andrea Schankin, and Michael Beigl. 2017. Lifefact: Utilizing Smartwatches As Tactile Heartbeat Displays in Video Games. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia (MUM '17)*, 97–101. DOI:<https://doi.org/10.1145/3152832.3152863>
- [36] A. J. Phillips, A. R. D. Thornton, S. Worsfold, A. Downie, and J. Milligan. 1994. Experience of using vibrotactile aids with the profoundly deafened. *Int. J. Lang. Commun. Disord.* 29, 1 (January 1994), 17–26. DOI:<https://doi.org/10.3109/13682829409041478>
- [37] Martin Pielot, Benjamin Poppinga, and Susanne Boll. 2010. PocketNavigator. In *Proceedings of the 12th*

- international conference on Human computer interaction with mobile devices and services - MobileHCI '10* (MobileHCI '10), 423.
DOI:<https://doi.org/10.1145/1851600.1851696>
- [38] Stefania Pizza, Barry Brown, Donald McMillan, and Airi Lampinen. 2016. Smartwatch in vivo. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, 5456–5469.
DOI:<https://doi.org/10.1145/2858036.2858522>
- [39] Steven Schirra and Frank R. Bentley. 2015. “It’s kind of like an extra screen for my phone.” In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '15*, 2151–2156.
DOI:<https://doi.org/10.1145/2702613.2732931>
- [40] Caitlyn Seim, Rodrigo Pontes, Sanjana Kadiveti, Zaeem Adamjee, Annette Cochran, Timothy Aveni, Peter Presti, and Thad Starner. 2018. Towards Haptic Learning on a Smartwatch. In *Proceedings of the 2018 ACM International Symposium on Wearable Computers* (ISWC '18), 228–229.
DOI:<https://doi.org/10.1145/3267242.3267269>
- [41] Kristen Shinohara and Jacob O Wobbrock. 2011. In the Shadow of Misperception: Assistive Technology Use and Social Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '11), 705–714.
DOI:<https://doi.org/10.1145/1978942.1979044>
- [42] Liu Sicong, Zhou Zimu, Du Junzhao, Shangguan Longfei, Jun Han, and Xin Wang. 2017. UbiEar: Bringing Location-independent Sound Awareness to the Hard-of-hearing People with Smartphones. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 2 (June 2017), 17:1--17:21.
DOI:<https://doi.org/10.1145/3090082>
- [43] I. R. Summers, M. A. Peake, and M. C. Martin. 1981. Field Trials of a Tactile Acoustic Monitor for the Profoundly Deaf. *Br. J. Audiol.* 15, 3 (January 1981), 195–199.
DOI:<https://doi.org/10.3109/03005368109081437>
- [44] Martin Tomitsch and Thomas Grechenig. 2007. Design Implications for a Ubiquitous Ambient Sound Display for the Deaf. In *Conference & Workshop on Assistive Technologies for People with Vision & Hearing Impairments Assistive Technology for All Ages (CVHI 2007)*.
- [45] Eddy Yeung, Arthur Boothroyd, and Cecil Redmond. 1988. A wearable multichannel tactile display of voice fundamental frequency. *Ear Hear.* 9, 6 (1988), 342–350.
- [46] Hanfeng Yuan, Charlotte M. Reed, and Nathaniel I. Durlach. 2005. Tactual display of consonant voicing as a supplement to lipreading. *J. Acoust. Soc. Am.* 118, 2 (August 2005), 1003–1015.
DOI:<https://doi.org/10.1121/1.1945787>