

HomeSound: An Iterative Field Deployment of an In-Home Sound Awareness System for Deaf or Hard of Hearing Users

Dhruv Jain¹, Kelly Mack¹, Akli Amrous⁴, Matt Wright², Steven Goodman³, Leah Findlater³, Jon E. Froehlich¹

¹Computer Science, ²Design, ³Human-Centered Design and Engineering
University of Washington, Seattle

⁴Bishop Blanchet High School
Seattle, WA

{djain, kmack3, matt3w, smgoodmn, Leahkf, jfroehli}@uw.edu

akliamrous2001@gmail.com



Figure 1: HomeSound is an in-home IoT prototype device aimed at improving sound awareness for people who are deaf or hard of hearing. In the image above, an occupant turns off the water faucet after being alerted to the “water pour” sound in HomeSound Prototype 2. Screenshots of the IoT display and accompanying smartwatch are also shown. See video.

ABSTRACT

We introduce *HomeSound*, an in-home sound awareness system for Deaf and hard of hearing (DHH) users. Similar to the *Echo Show* or *Nest Hub*, HomeSound consists of a microphone and display, and uses multiple devices installed in each home. We iteratively developed two prototypes, both of which sense and visualize sound information in real-time. Prototype 1 provided a floorplan view of sound occurrences with waveform histories depicting loudness and pitch. A three-week deployment in four DHH homes showed an increase in participants’ home- and self-awareness but also uncovered challenges due to lack of line of sight and sound classification. For Prototype 2, we added automatic sound classification and smartwatch support for wearable alerts. A second field deployment in four homes showed further increases in awareness but misclassifications and constant watch vibrations were not well received. We discuss findings related to awareness, privacy, and display placement and implications for future home sound awareness technology.

Author Keywords

Deaf and hard of hearing; smart home; sound awareness.

CSS Concepts

Human-centered computing—Accessibility technologies

INTRODUCTION

Sound awareness in the home has the potential to support deaf and hard of hearing (DHH) people with daily tasks (*e.g.*,

knowing the microwave beeped), safety-related information (*e.g.*, an alarm is sounding), and by keeping the user informed about the state of their home (*e.g.*, the shower is running) [18]. To date, several formative studies have explored DHH participants’ interest in and preferences for home-based sound awareness systems [6,18,24] showing that emergency alarms and appliance sounds are of high interest [6,24] and that people who identify as ‘Deaf’ or ‘hard of hearing’ are generally interested in similar sounds (*e.g.*, appliance alerts) with some exceptions (*e.g.*, alarm clocks) [24]. Most recently, a lab-based evaluation of a Wizard-of-Oz (WoZ) home sound awareness display [18] elicited a positive response: DHH participants preferred floorplan-based, localized views of sound, desired automatic sound classification and alert support, and expressed concerns related to trust, privacy, and information overload.

Building on these formative studies, we report on the design and field evaluation of a home-based sound awareness system called *HomeSound*—to our knowledge, the first such evaluation. Similar to other display-based IoT devices like the *Echo Show* or *Nest Hub*, HomeSound consists of a microphone and display, and multiple devices are installed per home; however, device interactions are controlled by touch rather than the user’s voice. We iteratively built and evaluated two versions starting with simple but accurate visualizations of sound feedback (Prototype 1) before adding more complex features (Prototype 2). We deployed each version for three weeks in homes with DHH occupants and conducted pre/post-interviews and weekly online surveys.

Prototype 1 was composed of 3-5 interconnected tablet-based displays encased in laser-cut wood frames that sensed and visualized sound characteristics such as loudness and pitch. Upon deployment in four homes, we found an increase

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 for Computing Machinery.
ACM ISBN 978-1-4503-6708-0/20/04...\$15.00
<https://doi.org/10.1145/3313831.3376758>

in the self- and home-awareness of the participants, who used context (*e.g.*, location, visual cues) to identify sounds from the display visualizations. However, similar to our previous WoZ-based lab study [18], participants expressed the need to incorporate automatic sound identification and alerts. In terms of privacy, the house occupants accepted the always-on monitoring, but some guests voiced concerns.

Informed from these findings, we extended Prototype 1 by adding a sound classifier for 19 common house sounds (*e.g.*, alarms, kitchen appliances) and a smartwatch app for providing alerts about sounds. We deployed this new prototype in four homes (two new, two repeat). Results show a further increase in participants' home awareness. However, misclassification of sounds and frequent smartwatch vibration alerts were not well received.

In summary, our research contributes: (1) two iterative prototypes of the first home-based sound awareness system for DHH occupants, and (2) insights from two three-week field deployments in six unique homes, including recommendations for future sound awareness technology. Our overarching aim is to help design future IoT devices like the Echo Show or Nest Hub while accounting for the needs and desires of DHH users.

RELATED WORK

We provide background on DHH culture, domestic computing research, and sound awareness needs and tools.

DHH Culture and Technology Adoption

Within the DHH population, individuals may identify as Deaf, deaf, or hard of hearing. Deaf (capital 'D') refers to individuals who follow certain values, practices and language of Deaf culture [21,28], irrespective of their audiological degree of deafness. Any individual can choose to associate with the Deaf community, including a hearing person, as is common with household members of Deaf individuals [8,28]. For hard of hearing and deaf (small 'd') people, the degree of hearing loss tends to closely determine communication strategies and the choice of assistive technologies (*e.g.*, hearing aids), and they may choose to interact more with either Deaf or hearing people [28].

Deaf people do not consider deafness to be a disability, and rely heavily on vision or haptic information (*e.g.*, flashing doorbells, vibratory alarm clock) for some sound-based applications compared to hard of hearing people [6]. The homes of some Deaf people are also designed to increase visual range (*e.g.*, by arranging furniture [19]) and may have distinct privacy norms (*e.g.*, more open to sharing personal information than hard of hearing people [17]). These cultural differences may influence the use and adoption of a home technology, which we examine in our work.

HCI Research in the Home

HCI research has evolved from optimizing technical efficiency of home systems to examining technology within social life and domestic practices [4,9]. For example, Forlizzi *et al.* [14] explored how a cleaning robot transforms the

practices and values of cleanliness in the home. Heshmat *et al.* [16] examined how a video recording system introduces tensions in family relationships and gender politics. Desjardins *et al.* [9] proposed seven genres of HCI home research. Our work intersects with four: how HomeSound is adopted in the home, shapes social routines, affects domestic practices, and promotes unexpected behaviors (*e.g.*, play).

For people with disabilities specifically, most HCI-related home technology efforts have examined the impact on disabled older adults' quality of life [10,23,33]. This research revealed differences between environmental augmentations (*e.g.*, walled displays) and individualized technologies (*e.g.*, wearables)—the former being more social but less private and personalizable [10]. However, the locations in which the devices are installed may mitigate privacy concerns [27]. More recently, a survey by Pradhan *et al.* [30] found that smart speakers increased the autonomy of visually and motor impaired users. Compared to these users, DHH people may have trouble interacting with smart speakers due to problems with Deaf accent [12]. Here, we investigate more accessible alternatives: screen and smartwatch-based prototypes.

Understanding Sound Awareness Needs

As mentioned in the Introduction, Matthews *et al.* [24], Bragg *et al.* [6] and Jain *et al.* [18] explored DHH users' sound awareness needs and identified preferences for sound types (*e.g.*, a doorbell), sound characteristics (*e.g.*, loudness), display form factor (*e.g.*, tablets), visual elements (*e.g.*, floorplan), and themes that apply to the home (*e.g.*, privacy). Other studies have identified sound awareness preferences in other contexts that could apply to the home. For example, Findlater *et al.* [11] surveyed 201 DHH participants on sounds of interests for mobile and wearable devices, finding that urgent alerts (*e.g.*, alarms, safety-critical sounds) and "voices directed at you" received the highest priority. Mielke *et al.* [14] interviewed six DHH people, identifying hazard alarms, phone ring, and siren as preferred sounds for a workplace setting. While useful, however, these studies used surveys, semi-structured interviews, or single lab sessions. Our field deployment complements the above studies by extending key themes related to the home such as privacy, contextualized feedback, and cultural considerations.

Sound Awareness Tools for DHH People

Though past work has not built a sound awareness system for the home, relevant findings from other domains inform our work. Matthews *et al.* [24] conducted a lab evaluation of a desktop-based prototype in an office setting with 4 DHH participants, identifying the desired sound information to display (*e.g.*, sound source, location) and the visualization type (*e.g.*, spectrograph, rings). Bragg *et al.* [6] and Sicong *et al.* [32] used smartphones to recognize and display sound information, focusing only on conveying the sound identity (*e.g.*, phone ringing, sirens). The latter work also included a two-day field study with 86 DHH participants in two Deaf schools, which highlighted the importance of using both visual and vibration feedback to notify users about sounds.

In terms of wearable sound awareness solutions, Summers *et al.* [34] performed a controlled study of a wrist-worn vibrotactile aid to measure 19 DHH participants’ accuracy in identifying domestic sounds (*e.g.*, water running, door knock). Jain *et al.* [9] used a design probe method to explore sound visualizations on a head-mounted display with 24 DHH participants. Mielke *et al.* [26] conducted a WoZ study of smartwatch app designs with six DHH participants. The lattermost work showed preferences for smartwatch as a form factor for notifications due to its small, private display and its popularity, which alleviates concerns of stigma associated with using assistive technologies.

We build on the above work by designing IoT displays and smartwatch-based sound awareness prototypes specifically for the home and conducting a longitudinal field study (three-weeks) in the homes of DHH people.

HOMESOUND PROTOTYPE 1

HomeSound is inspired by commercially available display-based domestic IoT devices like the *Echo Show* or *Nest Hub* but designed specifically to provide sound information to DHH users. To create HomeSound, we followed a human-centered iterative design process starting with the construction and evaluation of a simple but accurate sound feedback prototype (Prototype 1) before building and deploying a more complex system (Prototype 2). With Prototype 1, our goal was to examine how DHH users and other home occupants would react to and experience a sound awareness system, which conveyed four sound properties: room-level location, loudness, duration, and pitch.

Prototype 1 consisted of 3-5 interconnected “picture frame” displays (Microsoft Surface tablets encased in a laser-cut wood frame). Each display continuously sensed, processed, and uploaded sound information in real-time, which was further processed by a backend server to produce a single across-home sound feedback visualization. Though the tablets were general purpose computers, the HomeSound displays were intended to function as IoT devices—no other tablet-based applications or interactions were possible. Below, we describe HomeSound’s privacy-preserving sound sensing pipeline, visualizations, and our implementation.

Sound sensing pipeline. For domestic IoT systems, privacy is a key concern [16,18]. While HomeSound relies on a distributed set of live microphones, from the onset, we designed our sensing pipeline to protect user privacy. Each device processes sound locally and only uploads non-reconstructable features. For signal processing, we take a sliding window approach: HomeSound samples the microphone at 44kHz and segments the data into 50ms windows (2200 samples). To extract *loudness* and *frequency*, we compute the average amplitude and maximum frequency in the window (FFT bin size: 20Hz; range: 0-22kHz) and upload the results. For each display, the backend server stores this information in a database and applies a simple *sound event* detection algorithm: when loudness crosses a minimum threshold (46dB), a ‘start’ event is marked, which

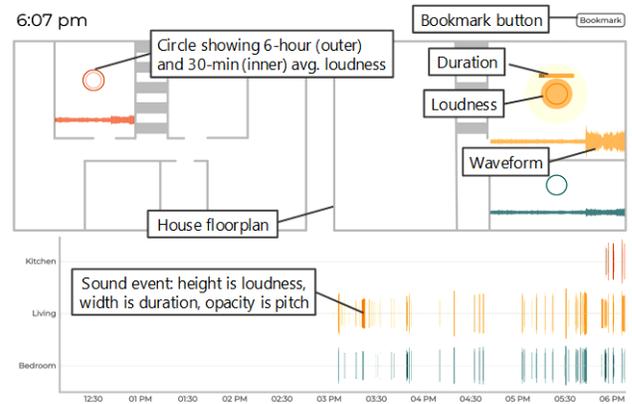


Figure 2: HomeSound Prototype 1 interface showing the floorplan view (top half) and history view (bottom half).

then ‘ends’ when loudness falls below 46db for one second. These thresholds were determined during a one-month pilot in home of the first author, who is hard of hearing.

Visual display. Informed by previous work [18,24], we designed the HomeSound display to be simple, glanceable, and require no direct interaction. The visuals are composed of two primary views: a *floorplan* (top half) and a *history* view (bottom half). In addition, a header bar displays the current time and a bookmark button, which allowed users to mark an event of interest for consideration by the research team (when pressed, the system took a screenshot and opened a form for typed feedback). See supplementary video.

The *floorplan* view showed a top-down blueprint of the home, which was overlaid by real-time sound information. For rooms with an installed HomeSound device, a ‘pulsing’ circle displayed in the room’s center depicted real-time ambient sound loudness—the circle’s radius was drawn proportionally to sound amplitude. At the top of each circle, we displayed a ‘sound event duration bar’, which visualized the length (in time) of the currently detected sound event. To enable comparisons across time, two circle outlines were drawn on top of the pulse showing average room loudness for the past 30 mins and 6 hours. In addition, we displayed a 30-second scrolling waveform at the bottom of each room, intended to help users detect visual patterns in sound activity.

For the *history* view, we created a custom time-series visualization, which showed per-room sound activity over the last six hours. Inspired by [36,37], sound events are displayed as rectangular blocks—block width represents duration, height is average loudness, and color opacity is pitch. The six-hour window was selected to enable recent comparisons across time (10 secs = ~1px) while balancing privacy concerns related to longitudinal patterns.

Implementation. We implemented HomeSound in Node.js [35] using a client-server architecture composed of three parts: a data client, a web client, and a backend server. For the displays themselves, we used Microsoft Surface Pro 3 tablets (i7 1.4GHz, 4GB RAM) encased in a custom wood frame, which ran both the data and web clients. The data

client sensed, processed, and uploaded sound data to the backend server while the web client visualized sound information downloaded from the server in a full-screen Chrome browser. We used pyAudio [38] for sound processing, socket.io for client-server communication, and D3 [5] and CSS for the visualizations. For the backend, we built a Node.js HTTP server on a Windows desktop computer (Intel i7 running Windows 10) using a pm2 process manager [39]. For each home, the server received pitch and loudness data from the data client, computed sound events (loudness, pitch, duration), and broadcasted this information in real-time to all web clients. Because the sound data (and state information) was stored on the backend, new web clients could be easily launched and supported.

STUDY 1: PROTOTYPE 1 DEPLOYMENT

To examine how DHH users react to and engage with a simple in-home sound awareness system, we performed a field study of Prototype 1 in four homes.

Method

Participants

We recruited DHH participants using email, social media, and snowball sampling. As our study focused on the home, we also recruited hearing household members of the DHH individuals. Six DHH and one hearing individual agreed to participate, and they were on average 62.4 years old ($SD=12.8$, $range=43-79$)—see Table 1. Of the six DHH participants, four reported congenital hearing loss, H1P1 reported onset at 3 years old and H3P1 at 4. Two participants used digital hearing aids and one used cochlear implants.

Procedure

The study, conducted by a hard of hearing author, had three parts: an initial interview and system installation, three-week system use, and a post-trial interview. Both pre- and post-interviews were held in the participants’ homes and audio recorded. A real-time transcriptionist attended all interviews, and five participants opted to also have a sign language interpreter. Consent and background forms were emailed in advance; written consent was taken in person.

Part 1: Initial session (1 hour). The initial session began with a 20-minute semi-structured interview (5 questions) with the DHH participants on experiences with sounds in the home, challenges faced, and coping strategies. The hearing participants then joined for a PowerPoint presentation on HomeSound, including the visualization overview and how user privacy is preserved. Afterwards, the participants gave a brief tour of their home and discussed the display placements. The initial number of displays was based on the home size (three for homes of <1000 sq. ft, four for 1000-1500 sq. ft, and five for >1500 sq. ft), but participants could ask to add or remove a display during the study. Though participants could choose any locations for these displays, we provided three suggestions: kitchen, living room and entryway. After the tour, the researcher took 20 minutes to draw the floorplan using online software and uploaded it to the server. The displays were then placed in the desired

H	ID	Age	Gender	Identity	Hearing Loss	#Rooms	#Occ.	#DHH
H1	P1	67	M	HoH	Severe	4-6	2	1
	P2	77	M	Hearing	N/A			
H2	P1	79	M	Deaf	Profound	7-9	2	2
	P2	60	F	Deaf	Profound			
H3	P1	56	M	HoH	Profound	10+	4	2
	P2	55	F	HoH	Severe			
H4	P1	43	M	Deaf	Profound	4-6	3	1
	P2	49	F	Deaf	Profound			
H5	P1	50	M	Deaf	Profound	7-9	4	2
	P2	49	F	Deaf	Profound			
H6	P1	22	F	HoH	Severe	4-6	3	1
	P2	21	F	Hearing	N/A			

Table 1: Homes for Study 1 (H1, H2, H3, H4) and Study 2 (H1, H2, H5, H6) with participant characteristics. Counts for occupants in the home (#Occ.) and for DHH occupants specifically (#DHH) include the participant.

rooms on a flat surface (*e.g.*, kitchen counter, bedside table) based on visibility and proximity to a power source. The researcher initialized and demoed the system by making some sounds (*e.g.*, clap, speech) in front of each display. Finally, participants were provided a UI reference sheet, and encouraged to give feedback using the bookmark form.

Part 2: Deployment period (3 weeks). During the three-week deployment, participants were instructed to perform their usual daily activities, interacting with HomeSound if desired. We emailed three weekly surveys (5 open-ended questions each) about overall experience, sound awareness, and any positive or negative incidents. If a participant did not complete a survey within 24 hours, a reminder was sent.

Part 3: Post-trial interview (1 hour). At the end of the deployment, we conducted another one-on-one interview with each participant (20 questions for DHH and 10 for hearing participants) on their system usage and experience, sound awareness, concerns, privacy issues and design suggestions. We also asked follow-up questions based on system logs, bookmarks, or survey responses. After the interview, we retrieved the displays.

Data Analysis

We conducted a thematic analysis [7] on the interview transcripts and weekly survey data. One researcher skimmed the transcripts to familiarize with the data, and conferred with the research team to generate an initial codebook. The researcher then iteratively applied codes to all transcripts while refining the codebook. The final codebook contained a 3-level hierarchy (11 level-1 codes, 54 level-2 codes, and 108 level-3 codes), of which the level-1 codes formed the high-level themes. Another researcher used this final codebook to independently code all transcripts. Interrater agreement between the two coders, measured using Krippendorff’s alpha [20], was on average 0.66 across all questions ($SD=0.30$, $range=0.47-1.0$); raw agreement was 86.3% ($SD=11.7$, $range=70.4-100$). Though the alpha value borders the acceptable minimum (0.667), the two coders resolved all disagreements through consensus.

Findings

We cover overall usage, sound awareness, display placement, privacy, and design suggestions. Throughout, we

refer to the six DHH participants and report quotes from the post-trial interview unless otherwise noted.

Overall usage patterns. On average, each home had 2411.8 total sound events/day ($SD=689.5$). Participants completed all weekly surveys, created 46 bookmarks (Table 2), and sent feedback using 9 email threads, and 21 text messages. Complementing this quantitative data, all participants reported viewing the HomeSound displays at least a few times a day, both explicitly (*e.g.*, to review past sounds) and incidentally (*i.e.*, noticing it during other activities). For explicit use, all participants reported checking a nearby display every few hours, and almost all ($N=5$ out of 6) reviewed sound activity when they came home from work. All participants also mentioned noticing sound information while walking around the house or engaged in other activities—particularly activities that generated sound (*e.g.*, cooking, conversation, laundry). Perhaps unsurprisingly, all participants reported decreased usage over time, as is evident from the logged bookmark data (Table 2):

“I looked at it a lot in the first week, but then not so much in the end. I got used to its presence and forgot it was there...” (H4P1)
“I felt like it was fun to look at it initially but then I am so used to living without sounds, that I used it less in the end.” (H2P1)

However, emphasizing the utility of the system for some people, H1P1, H3P2 and H4P1 mentioned feeling nostalgic about it after the deployment period ended, such as,

“I was waiting for the microwave to beep but was in the bedroom. So, I asked [my husband] if he could hear [it] [...] And he said: “where’s the system!?” (H1P1, text sent post-study)

Sound awareness. In terms of sound awareness, four participants reported how HomeSound made them realize that they were previously unaware of many sounds in their home. H3P1, for example, wore a hearing aid but said, *“I knew I was missing certain sounds. [But] I didn’t know how much I was missing”* (interview). All participants reported combining information from the HomeSound displays with contextual cues to determine sound activity:

“Every time I walked around the house, I saw disks [pulses] on tablets [emanating from] multiple rooms. I realized that my whole wooden home makes a lot of noise” (H3P1, week 1 survey).
“The peaks in waveform from kitchen meant that the microwave must have beeped, and my food was ready. [...] No one else [was] in the home.” (H4P1, week 2 bookmark, see Figure 3 below)



Figure 3: Partial snapshot of a participant’s bookmark showing the distinctive microwave beep pattern in the waveform.

This increased awareness was useful at times for performing household tasks, such as when H2P2 used the system to monitor sound from her washer and dryer and was thus able *“to get my clothes done sooner”* (week 3 survey).

Home	Size	Floors	Displays	Busiest	Events	B1	B2	B3
H1	1060	1	4	Den	3383.4	8	4	1
H2	1740	2	5	Dining	2041.7	5	2	0
H3	2700	2	5	Family	1545.3	11	6	3
H4	950	1	3	Kitchen	2676.6	4	2	0

Table 2: Study 1 homes, with sizes (in sq. ft.), number of displays, the most active room, daily average of sound events, and total bookmarks for each of the three weeks (B1, B2, B3).

Additionally, H3P1 looked for door open and close sounds in the history view to see when his roommate left, so he could know when to take a shower. Two participants used the display to monitor the well-being of their family members:

“[I knew my autistic brother] was pacing around the kitchen [when] I looked at this [display]” (H3P1)
“It would help [to] recognize if somebody had a fall. If something happened to [my husband who has Parkinson’s], I would have no idea. I would just find him on the floor.” (H2P2 interview)

Five of six participants also noted how HomeSound provided insight into their own behaviors. For example:

“I practiced closing and opening the cabinet and monitor the sound waves with the history to improve my ability to be very quiet for [my spouse]. He is a light sleeper...” (H1P1)
“I can turn the kitchen fan off. Earlier my mom/dad used to tell me. [Now,] when they are away, I turn it off on my own.” (H4P1)

However, this increased awareness did not always produce positive reactions or insights. For example,

“I was shocked [to learn] how much [...] noise I create during meals.” (H2P1, week 1 survey)
“I felt embarrassed that this picked up my loud laugh and decided to be careful...” (H4P1)

Display placement. The physical environment of the home influenced where HomeSound units were placed. In general, participants chose the most active rooms and placed the displays in salient, highly-viewable locations (a shelf or a table). Nevertheless, placement sometimes posed a problem due to visibility or space concerns. For example, to preserve kitchen counter space, H4P1 decided to place the display on top of the fridge but then could not see it. Similarly, participants who opted for a bedroom installation (H1, H3 and H4) found the screen light disruptive at night, and either covered the displays or put them face down, which decreased their utility. Having the sensing (microphone) and display coupled on a single device also caused issues, some of which reflected the importance of DHH individuals being able to maximize their sightline [19]. For example,

“I usually sit over here in my dining room, which [is] a good vantage point for me to see the house. [So,] I placed [the display] here. But then its usefulness was reduced [as] it was far from the kitchen and I wanted the kitchen sounds. So, I moved it to the kitchen. But then, I wasn’t able to see it from the dining.” (H1P1)

Decoupling the sensing from the display would also address a desire to have a display where one may not want sensing (*e.g.*, bathroom), as suggested by past work [29] and H3P2:

“I wanted a monitor in the bathroom to see what was going on in my home but then I don’t want it to [hear] the private bathroom stuff. Can you make the mic and the display separate?”

Privacy. All occupants accepted the system after learning its privacy-protecting measures but interview responses indicate there could be future privacy concerns. In the initial interview before we explained how HomeSound preserves privacy, three DHH participants had expressed concern about recording conversations, such as: *“Is it recording my voice? Do I have to be concerned about what I am going to say when I am near it”* (H1P1). Surprisingly, no household members beyond the DHH participants expressed privacy concerns to us directly or indirectly. This openness may have been related to its assistive nature. For example, H1P1 said, *“[My hearing spouse] accepted it because it was an assistive technology and he knew this was necessary to help me,”* a sentiment that was also echoed by H4P1:

“My mom was concerned when she was cooking, and the system was showing all her cooking activity. But she knew it was important to be there for me to help recognize the sounds.”

However, this notion was not necessarily shared by guests, which included visits from friends and family in two homes:

“My friend asked his wife to not hold a conversation near a tablet [...] Then I explained that this [system] cannot display words and he seemed to be ok with it then. Although I must say he was a little put off initially.” (H3P1)

To mitigate this issue, three participants suggested adding the ability to turn off the recording in a room as needed.

Design suggestions. The most common design suggestions included updates to the visual designs, a notification feature, and automatic sound identification. When asked about the interface design, all participants appreciated the floorplan and history views, finding them easy to learn and use. In contrast, the waveform was seen as too abstract, although it offered an indirect benefit to some participants: H3P1 and H4P1 noted that they had used the waveform to begin recognizing visual patterns of recurring sounds; for example, H4P1 identified the microwave beep.

Four participants also made suggestions for adding features to the history view, such as accessing a stored waveform, daily or weekly summary of information, and the ability to see different timespans. For example, H3P1 wrote about the history view in a week 2 bookmark that, *“I would like to zoom in and see the sound signature [for the microwave beep].”* Finally, H3P1 suggested that different visual designs may be useful for different locations in the home, such as a floorplan view in the office where they sit close to the display, compared to a more generic alert design in the kitchen: *“just some kind of alert that there is a sound.”*

Because the visual information was only useful when it was within sight, all participants reported missing useful sounds when they were not close to the display (e.g., door knock, appliance alerts) and requested a notification feature.

“I decided if I have to watch the tablet for tea kettle whistle, I may as well watch the tea kettle. Pair with a watch that would vibrate to let you know what is happening.” (H1P1, week 2 text message)

“If the tablets could flash when some sound occurred, then I could check.” (H2P1)

All participants also expressed the desire for automatic sound classification, so that they did not have to rely so much on context or hearing roommates to guess the source of sounds.

“I would sometimes let the dog out and I always have to make sure to check... Sometimes I would forget that I put him outside and if he wants to come in, he would bark and bark [...] It would really help if it tells me the dog was barking.” (H2P2)

“I happened to notice [pulses] in the den and bedroom while I was in the dining area. I went to the den to find out what it was, nothing I could ascertain. So, I went to bedroom. Nothing there either. I asked [my hearing spouse] who said it was Siren from outside.” (H1P1, week 1 bookmark)

Building on the ability to identify sounds, five participants wanted to know safety-related sounds (e.g., fire alarms) from outside the home.

Summary. Participants appreciated HomeSound for its ability to increase self- and home-awareness. They used context such as location and visual cues to identify sounds from basic visualizations, which influenced some daily chores and increased awareness about other occupants' activities. In terms of privacy, house occupants accepted the always-on sound monitoring more than the guests. Finally, the need to constantly monitor the displays, and the lack of automatic sound identification were primary limitations.

HOMESOUND PROTOTYPE 2

Informed by our experiences with Prototype 1, we extended HomeSound in two ways: first, we added a real-time, deep-learning based sound classification engine to automatically identify and visualize sound events; second we designed and implemented a complementary smartwatch system that provided customizable sound alerts via visual+haptic notifications. We describe both extensions below as well as updates to the IoT display interface.

Sound Classification Engine

To create a robust, real-time sound classification engine, we followed an approach similar to *Ubicooustics* [22], which uses a deep convolutional neural network (CNN) called *VGG16* [15] pre-trained on 8 million YouTube videos [2]. Because *VGG16* is developed for video classification, we used transfer learning to adapt the model for sound classification. For this, similar to [22], we use a large corpus of sound effect libraries—each of which provide a collection of high-quality, pre-labeled sounds. We downloaded 19 common home-related sounds (e.g., dog bark, door knock, speech) from six libraries—BBC [40], Freesound [13], Network Sound [41], UPC [42], TUT [25] and TAU [3]. All sound clips were converted to a single format (44Hz, 16-bit, mono) and silences greater than one second were removed, which resulted in 31.3 hours of recordings. We used the method in Hershey *et al.* [22] to compute input features. Finally, we fine-tuned the model by replacing the last fully connected layer with a fresh layer, retraining on only a subset of sound classes (Table 3) to generate per-room classification models.

Experimental evaluation. To evaluate our model, we collected our own ‘naturalistic’ sound dataset. We recorded

16 sound classes from five homes using the same hardware as HomeSound—a Surface Pro 3 with a built-in microphone. For each home, we collected sounds in three rooms (*bedroom, kitchen, living room*). For each sound class, we recorded five 10-second samples at three distances (5, 10, and 15 feet). We attempted to produce sounds naturally (*e.g.*, using a kettle or running water to wash hands). For certain difficult-to-produce sounds—like a fire alarm—we played snippets of predefined videos on a laptop or phone (33 total videos were used). Because we train per-room classification models, not all sound classes were recorded in each room; Table 3 shows the sounds per room (the three outdoor sounds were not recorded for this experimental evaluation). In total, we collected 1,200 recordings (3.3 hrs).

Before testing our model, we also added 20% sound data from other rooms in our test set that our model should ignore (called the "unknown" class). For our evaluation experiment, we classified data collected from each room using the appropriate per-room classification model. Our overall accuracy was 85.9% with small, per-room differences: the average in the living room was 88.5% ($SD=3.5\%$) followed by the kitchen (86.4%; $SD=3.0\%$) and bedroom (82.5%; $SD=7.3\%$). The best performing sounds included cutlery (100%, $SD=0$), door in use (98.7%, $SD=2.7\%$), and water pour (96.0%; $SD=5.3\%$) and the worst: phone ring (60.0%; $SD=32.0\%$), alarm clock (50.7%; $SD=24.1\%$), and dishwasher (45.3%; $SD=7.8\%$). For poor performing classes, understandable mix-ups occurred: *e.g.*, 24.0% of phone ring sounds were classified as doorbells and 38.4% of alarm clock sounds as a phone ring. Interestingly, accuracy was unaffected by recording distance: at 5ft $avg=85.5\%$ ($SD=16.2\%$), 10ft (85.2%; $SD=15.9\%$), and 15ft (87.1%; $SD=15.1\%$). We return to classification accuracy and its impact on users in Study 2 Findings and our Discussion.

Implementation. We built the classification engine in Python using *Google TensorFlow* [1], which ran locally on each HomeSound device (to protect occupant privacy): 1 second of microphone data was buffered (44,000 samples) and relevant features extracted and classified. Only the classified sound, classification confidence, loudness, and room location were uploaded to the server (no raw features were transmitted). On the server, all sounds below 50dB or 50% confidence were ignored; the others were broadcast to the web and smartwatch clients.

Smartwatch

To transform HomeSound from a passive awareness system to a proactive one and to eliminate line-of-sight requirements, we designed and implemented a complementary Android-based smartwatch application. The smartwatch displayed a notification along with a vibration alert whenever a classified sound event occurred. The display included sound identify, classification confidence, and room (Figure 4). Importantly, each user could customize which sound alerts to receive by clicking on a notification, opening a scroll list, and selecting snooze options (1 min, 5

Kitchen	Bedroom	Living room	Outdoors
Cutlery	Alarm clock	Cat meow	Hammer
Dishwasher	Cough	Dog bark	Drill
Microwave	Snore	Doorbell	Vehicle
Water pour	Door in use	Door in use	
Phone ring	Phone ring	Phone ring	
Speech	Speech	Speech	
Hazard alarm	Hazard alarm	Hazard alarm	
Kettle Whistle		Door knock	

Table 3: List of sounds recognized by our sound classifier.

min, 10 min, 1 hour, 1 day, or forever). In our deployments, we used the *Android Ticwatch E2* watch [43] running *WearOS 2.0*, which communicated with the backend server using WiFi. To enable notifications even when the watch was in a low-power sleep state, we used the firebase messaging service (FCM) for watch-server communication.

HomeSound Display Updates

For the HomeSound IoT display, we made three primary changes: first, to incorporate the real-time sound classification engine, we visualized sound identities and their confidence below each circle pulse in the floorplan view and as annotations on the ‘sound event’ blocks in the history view (Figure 4). Second, similar to the smartwatch application, we added a customization menu, which allowed users to select which sounds to show on each display. Finally, we added a pan-and-zoom feature to the history view to increase granularity of sound event visualizations.

STUDY 2: PROTOTYPE 2 DEPLOYMENT

To evaluate Prototype 2, we performed a second field deployment in four homes using an adapted Study 1 protocol.

Method

Participants

As before, we recruited DHH participants and other house members through email, social media, and snowball sampling. As an iterative deployment, we did not exclude repeat participants; hence, two of the four homes (H1, H2) were the same as in Study 1 (Table 1). Six DHH and two hearing people agreed to participate, whose age averaged 53.1 years old ($SD=20.9$, $range=21-79$). Four DHH participants reported onset of hearing loss as congenital, H5P2 reported at 2 years old and H1P1 reported 3 years old. Two participants used an assistive device: hearing aids.

Procedure

We followed the same process for weekly surveys, and data logging as Study 1 but made slight changes to the initial session and post-trial interview. For the initial session, we made three modifications. First, to generate the room-specific models, participants selected up to eight sounds for each room from the 19-sound list (Table 3). Second, to demonstrate the smartwatch app, the researcher produced and snoozed a sound (*e.g.*, speech); participants were also asked to charge the watch at night and wear it throughout the day, including outside the home if desired by connecting to a WiFi source. Third, for the initial interview, new participants responded to the same questions from Study 1 while repeat participants were reminded of their Study 1

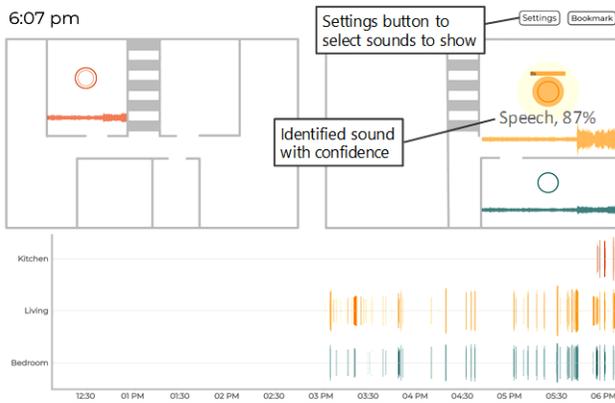


Figure 4: HomeSound Prototype 2 interface for displays.

responses and asked if anything had changed. Finally, for the post-trial interview, we added two questions on usage and experience with the sound classification and smartwatch app.

Data Analysis

We followed the same Study 1 analysis process with the same two coders. In summary, the final codebook contained 9 level-1 codes, 21 level-2 codes, and 65 level-3 codes. K-alpha was 0.78 ($SD=0.14$; $range=0.62-0.94$) and raw agreement was 91.7% ($SD=4.3$; $range=85.8-97.2$). All disagreements were resolved via consensus.

Findings

We discuss new insights related to Study 1 themes (usage patterns, household tasks, placement, privacy) as well as new emergent themes (cultural differences, playful interactions).

Usage patterns. On average, each home had 3,297.4 sound events/day ($SD=819.1$); 65.9% of them (2174.2, $SD=525.4$) were automatically classified. Participants filled all surveys, created 41 bookmarks, sent 13 email threads and 32 text messages (Table 3). In the first week, email and text messages asked about system operation—particularly on how to snooze the smartwatch app or select sounds on the displays, indicating a higher learning curve than Study 1. H6P1 corroborated this: “at first you have to get used to it. Like a new computer [...] it took time to learn”.

Smartwatch: In general, participants appreciated the watch alerts and wore the watch consistently, except when sleeping, bathing, or going out. Three participants chose to wear the watch outside the home but only wanted to be alerted about urgent home sounds (e.g., a fire alarm) and felt irritated about snoozing other sounds, which indicates a need for location-aware customization. In the home, notifications diminished the need to actively monitor the IoT displays but the ‘alert’ vibrations could be distracting and overly persistent. Three participants reported using the snooze feature; two others became inured—for example, H2P1 “would just ignore them when working on my computer” (interview). One participant removed the watch itself when the noise levels were high:

“I had company last Sunday. All of a sudden, it began [vibrating] constantly. I couldn't take away my attention because I didn't want to be rude to my company [and] click [on the app] to snooze different sounds. It was easier to take it off.” (H1P1)

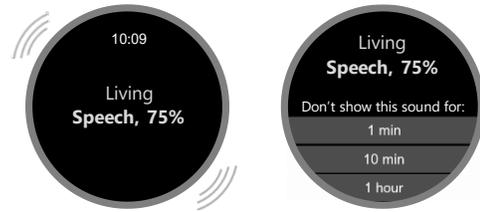


Figure 5: HomeSound Prototype 2 smartwatch interface. When a sound occurs, a vibration and a visual notification is received on the watch (left), which when clicked, opens the main app (right) that allows users to snooze the sound.

Displays: All displays ran continuously for three weeks, except in H5 where the bedroom display was closed every night for privacy. As smartwatches provided sound alerts, all participants looked at the displays only to recap the events of the day; consequently, for repeat participants, the perceived value of the IoT displays decreased. Only H2P2 and H5P1 described incidents of using the displays immediately after a sound occurred if one was nearby and in view. The watch also decreased the need to have multiple displays, as corroborated by a repeat participant during the interview:

“I check the tablet when I get home from work and see what had happened [...] For the historic information, having one tablet was sufficient. Having multiple tablets was overwhelming” (H2P2)

Household tasks. In some cases, HomeSound helped participants perform household tasks by alerting them to desired sounds in the house (e.g., someone knocking, dog barking, children’s shouts). For example,

“I was [...] working on my laptop, the watch showed my dog was barking [in another room]. I went and corrected my dog right away. This helps me train the dog over time [...] Also, the watch lets me know when the washer is done.” (H2P2, week 2 survey)

“The first day [when] the contractor would come over for the kitchen remodel. I was sitting close to the door. But the watch vibrated and [displayed] “door knock” and I thought, oh [from] now [on,] I don't have to sit and wait.” (H6P1)

However, system failures resulting from sounds being not supported (e.g., dryer beeps, garbage disposal) or misclassified also negatively affected daily routines. For example, H2P1 said: “a fan running in the kitchen kept identifying as microwave [...] and I will go and check again and again.” More concerning was when the same sound was only sometimes misclassified. For example, the system correctly identified door knocking initially, so the hearing children of H5P2 assumed that the system would always notify their parents of door knocks, creating issues:

“So, somebody was knocking at the door and [my kid] thought that [...] my watch will tell me and didn't bother to come up to me [...] And the guy was knocking, knocking, and finally [my kid] comes up to me and [signs], why are you not opening the door!?”

Interestingly, four participants found creative ways to compensate for some misclassifications using the help of house members, animals, or context:

“Someone knocked at the door [...] and [HomeSound] was not recognizing any of it. But [my dog] barked, and the watch alerted me to “dog bark”, so I went and looked.” (H2P2)

"I know microwave cannot run in my bedroom, so I ignored it." (H6P1, to a follow-up question on her bookmark)

"It said hazard alarm, but no lights were flashing. [My hearing spouse] confirmed it was from outside." (H2P1)

Playful interactions. Beyond household tasks, all occupants reported instances of deliberately initiating actions to record and explore their sound "footprint". For example, H1P1 would sometimes *"clap, hoot, or talk at the system for seeing them later in time."* Though this behavior mostly occurred in the first week, it resurfaced when guests arrived (in H2 and H6). In addition, H5P1 reported that his two children liked seeing the annotations in history view, and they would *"scream at [the displays] to see bubbles going up, down [pulsating in floorplan]"* or would *"play a variety of sounds"* to see how well the system would perform.

System improvements. Participants offered concrete ways to improve the smartwatch, display, and sound classification.

Smartwatch: To reduce the constant smartwatch vibration, five participants suggested alerting about a sound only once:

"If somebody is talking, it should tell me once, and that's it. This alerts me every second, and I [snooze for some time] and it comes back again. I can't [snooze] indefinitely, because what if somebody is actually calling me." (H1P1)

He added: *"also, if it showed speech, I would want to know [who] is talking."* This need to provide more details on the speech sound (e.g., speaker id, tone) was highlighted by three other participants as well.

Display: Because of the smartwatch, participants used the displays mainly to view past sounds and suggested enlarging the history view ($N=5$); four wanted to remove the floorplan and add the floorplan's characteristics (loudness, duration, pitch, color) to the smartwatch app ($N=5$) or the history view ($N=2$). For example, while sketching a new design during the interview, H1P1 said:

"[No] need to have the layout (floorplan) [...] Just show history [...] Make it larger to fit the entire screen [sketching]. Right now, the color codes are used to identify a room. But perhaps color codes could be used to [distinguish] sounds in the history".

Classification: To compensate for sound classification issues, participants suggested three technical improvements. First, three participants wanted more precise localization in areas with many sound producing appliances (e.g., kitchen) so they can identify sounds from their location. Second, three participants suggested dynamically adjusting the microphone sensitivity to increase feedback utility:

"Before remodeling the kitchen (in second week of deployment) we had a porcelain sink. We have stainless steel sink now. The water is quieter [in porcelain sink] but when it hits the stainless-steel sink [the sound] is amplified. So, too loud now and [...] I had to move the [display] a little farther." (H1P1)

Finally, participants suggested fine tuning the system ($N=3$) to the sounds of their home. For example, H6P1 said:

"better to record some sounds myself [...] I would prefer knowing [my spouse's] speech rather than knowing everybody else's"

Home	Size	Floors	Displays	Busiest	Events	Identified	Bookmarks
H1	1060	1	4	Den	3647.4	2417.5	12
H2	1740	2	5	Dining	1986.5	932.2	8
H5	1900	2	5	Living	4754.8	3213.6	17
H6	4453	1	4	Kitchen	2801.1	2133.3	4

Table 4: Study 2 Homes, their sizes (in sq. feet), number of displays deployed, most active room, average events logged and identified each day, and total bookmarks by participants.

Placement. The addition of the smartwatch and sound classification feature changed how occupants positioned the IoT displays in their home. While there was a decreased need for line-of-sight, participants felt that HomeSound needed to be close enough to sound activity to accurately detect and classify sound events. Consequently, three homes moved the displays closer or farther from sound sources, during which issues with space emerged:

"But then my kitchen is small... there's [only] so far I could move it. So, I placed it [a little outside the kitchen]" (H1P1)

And H5P2 said in the interview:

"There was nothing [no shelf space] closer to the front door, so I had to bring a table but then the door wouldn't [fully] open"

Another theme related to how sounds propagate through a home, which could be confusing or raise privacy concerns:

"I saw hammering in multiple rooms which surprised me [...] But then it occurred to me that these three rooms are all closer to the street, and so that must be the loud construction noise [from outside]" (H6P1, week 1 bookmark)

"Even after turning [the bedroom tablet] off [at night], we were concerned whether the tablets in other rooms [close to our room] would pick up the signal and the kids can see from downstairs [living room tablet]. The children might think it's weird to be having many sounds at this time of the night..." (H5P2)

Cultural differences. Though all participants reported that HomeSound helped them in some way, subtle differences emerged between Deaf (H2, H5) and hard of hearing (H1, H6) households, potentially related to cultural context (e.g., Deaf people tend to rely less on sounds than hard of hearing people [6]). Indeed, the three Deaf participants expressed that, apart from important cues (e.g., doorbell, alarm), other information was "nice to know" rather than "need to have":

"I've been Deaf since I was an infant, so I am used to life without sound. It is not really a big concern unless it is an emergency... I think the system could be better for people who became deaf later in life or hard of hearing but not necessarily for someone like myself." (H2P1).

However, Deaf individuals may find the sound information more valuable with longer-term exposure. For example:

"It was surprising that I'm making many noises such as closing the door, cabinets, talking to my dog, putting food and flatware on the counter (cutlery noises). [...] I felt awkward subjecting my [hearing] kids to all this noise and tried to change." (H5P2)

Privacy. Surprisingly, though with Prototype 2, participants received constant notifications about sound events (with more information), privacy concerns did not change from Study 1. Participants and their household members reported no issues except in the case of guests (H2, H6) who were

more curious than suspicious, and the issue of display placement with children (H5) (both are detailed above).

Summary. The smartwatch and sound classifications diminished the importance on the IoT displays while increasing general sound awareness. Key concerns included system failures, unpredictable sound classification errors, and overly persistent watch vibrations.

DISCUSSION

Though past work has identified preferences for in-home sound feedback through formative studies with DHH people [6,18,24], we designed and conducted the first field evaluation of two iterative real-time systems. Some of our results (*e.g.*, feelings about privacy, information overload) contextualize the prior findings from lab-based evaluations [6,18,24] but with higher ecological validity. We also report on new findings that are only possible via a real-world field deployment, such as usage patterns over time, reactions to varying locations of displays at home, and influence on social dynamics. Below, we discuss further implications of our findings and opportunities for future work.

UI design. Regarding specific UI elements, our findings both reinforce and contradict prior work—perhaps due to differences in how users respond in field evaluations. Similar to Jain *et al.* [18], sound history was appreciated, though unlike in Matthews *et al.* [24], the waveform was not. Reactions to the feedback design also differed between our two studies. While the floorplan was appreciated in Study 1 (and prior work [18,24]) as an indicator of current sounds and their location, the addition of the smartwatch alerts reduced this utility. Thus, participants suggested removing the floorplan and supplementing the sound information (*e.g.*, loudness, duration) on the smartwatch. Future work should further investigate how to balance this information between the small smartwatch face and tablets.

Sound misclassifications. Past work in home sound awareness technology [18] enumerates three possible classification errors: false positive (showing a sound that did not occur), misattribution (showing an incorrect sound) and false negative (not showing an occurred sound). In our work, no participants reported false positives, possibly because of the loudness thresholding. However, misattributions (*e.g.*, fan identified as microwave) and false negatives (*e.g.*, unsupported sounds like garbage disposal) were both problematic. Our findings also suggest a more concerning case: when these errors are *unpredictable*—which can cause users’ expectations and system behavior to be mismatched.

These misclassification issues suggest a need to further improve system accuracy or, at the very least, mitigate the potential downsides of inaccurate or unpredictable behavior. One possibility is to employ a customization approach such as that proposed by Bragg *et al.* [6] to allow participants to train the system for their home, though this training may be tedious and difficult if the sound itself is inaccessible to the user. While we conveyed classification confidence to users,

there may be opportunities to adapt how sound information is displayed based on confidence, such as displaying more ambiguous information (*e.g.*, a sound occurred) when confidence is low (as opposed to simply choosing *not* to show low confidence sounds at all, as in our design).

Social dynamics. Our participants used sound history to help manage social dynamics: coordinating schedules or monitoring well-being. House members seemed to accept the system because it was perceived as ‘assistive’ [31]. However, as indicated by past work [18], the system has the potential for privacy intrusion (*e.g.*, surveillance without consent). Thus, future work should consider: who should be able to view the sound history and what inferences can be drawn about human activity from the visual representations?

Self-awareness. An unintended effect of the system was increased self-awareness, leading to adaptations in behavior. Thus, future work should explore how the system can better support these feedback-based adaptations, for example, by showing volume graphs and notifications of personal sounds above a certain volume. Past work in tactile-based sound awareness examined wrist-worn devices for regulating personal voice levels [34], but this work could be extended to visual displays and smartwatches in the home context. Importantly, there can be negative implications to any system that explicitly or even inadvertently encourages users to change their behavior. For example, while most feedback was received positively, the visualizations of loud noises created by participants were sometimes associated with feelings of embarrassment.

Limitations. Our findings are based on two three-week deployments but with only four homes. Future work should include a larger and more varied set of households, which could yield additional insights, especially related to social dynamics. Second, though our controlled evaluation showed acceptable overall accuracy, we do not have quantitative data on how our sound classification system worked in practice; instead, we only have participants self-reported perceptions of the system. Finally, while our designs were informed by prior work [18], other visual and haptic representations are possible and should be explored.

CONCLUSION

In this paper, we presented the iterative design and evaluation of *HomeSound*—the first IoT-based sound awareness system for DHH households. Our findings demonstrate value, especially with regards to feelings of increased awareness amongst our DHH participants but also uncover important issues related to privacy, social dynamics, and classification accuracy. Our work has implications for future ‘smarhome’ displays such as the Echo Show.

ACKNOWLEDGMENT

We thank Gierad Laput, Karan Ahuja, Elliot Saba, Dominik Moritz, Liang He, Venkatesh Potluri, John Doherty, and Angela Lin for their help. This work was supported by NSF grant IIS-1763199 and by a Google Faculty Research Award.

REFERENCES

1. Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, and others. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 265–283.
2. Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*.
3. Sharath Adavanne, Archontis Politis, and Tuomas Virtanen. 2019. TAU Moving Sound Events 2019 - Ambisonic, Anechoic, Synthetic IR and Moving Source Dataset.
4. Ryan Aipperspach, Ben Hooker, and Allison Woodruff. 2008. The heterogeneous home. In *Proceedings of the 10th international conference on Ubiquitous computing*, 222–231.
5. Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D3 data-driven documents. *IEEE transactions on visualization and computer graphics* 17, 12: 2301–2309.
6. 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*, 3–13.
7. Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2: 77–101.
8. Anna Cavender and Richard E Ladner. 2008. Hearing impairments. In *Web accessibility*. Springer, 25–35.
9. Audrey Desjardins, Ron Wakkary, and William Odom. 2015. Investigating genres and perspectives in HCI research on the home. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 3073–3082.
10. Pier Luigi Emiliani and Constantine Stephanidis. 2005. Universal access to ambient intelligence environments: opportunities and challenges for people with disabilities. *IBM Systems Journal* 44, 3: 605–619.
11. 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 *SIGCHI Conference on Human Factors in Computing Systems (CHI)*. In Submission.
12. Raymond Fok, Harmanpreet Kaur, Skanda Palani, Martez E Mott, and Walter S Lasecki. 2018. Towards More Robust Speech Interactions for Deaf and Hard of Hearing Users. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*, 57–67.
13. Eduardo Fonseca, Jordi Pons Puig, Xavier Favory, Frederic Font Corbera, Dmitry Bogdanov, Andres Ferraro, Sergio Oramas, Alastair Porter, and Xavier Serra. 2017. Freesound datasets: a platform for the creation of open audio datasets. In *Hu X, Cunningham SJ, Turnbull D, Duan Z, editors. Proceedings of the 18th ISMIR Conference; 2017 oct 23-27; Suzhou, China.[Canada]: International Society for Music Information Retrieval; 2017. p. 486-93*.
14. Jodi Forlizzi and Carl DiSalvo. 2006. Service Robots in the Domestic Environment: A Study of the Roomba Vacuum in the Home. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction (HRI '06)*, 258–265.
15. Shawn Hershey, Sourish Chaudhuri, Daniel P W Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, and others. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 131–135.
16. Yasamin Heshmat, Carman Neustaedter, and Brendan DeBrincat. 2017. The Autobiographical Design and Long Term Usage of an Always-On Video Recording System for the Home. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*, 675–687.
17. Janice Humphrey and B Alcorn. 2001. *So You Want To Be An Interpreter? An Introduction to Sign Language Interpreting*. Texas: H & H Publishing Company.
18. Dhruv Jain, Angela Carey Lin, Marcus Amalachandran, Aileen Zeng, Rose Guttman, 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*, 94:1-94:13.
19. Charlene A Johnson. 2010. Articulation of Deaf and Hearing Spaces Using Deaf Space Design Guidelines: A Community Based Participatory Research with the Albuquerque Sign Language Academy. Retrieved from https://digitalrepository.unm.edu/arch_etds/18
20. Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
21. Paddy Ladd and Harlan Lane. 2013. Deaf ethnicity, deafhood, and their relationship. *Sign Language Studies* 13, 4: 565–579.
22. Gierad Laput, Karan Ahuja, Mayank Goel, and Chris

- Harrison. 2018. Ubicoustics: Plug-and-play acoustic activity recognition. In *The 31st Annual ACM Symposium on User Interface Software and Technology*, 213–224.
23. Lili Liu, Eleni Stroulia, Ioanis Nikolaidis, Antonio Miguel-Cruz, and Adriana Rios Rincon. 2016. Smart homes and home health monitoring technologies for older adults: A systematic review. *International Journal of Medical Informatics* 91: 44–59.
24. Tara Matthews, Janette Fong, F. Wai-Ling Ho-Ching, and Jennifer Mankoff. 2006. Evaluating non-speech sound visualizations for the deaf. *Behaviour & Information Technology* 25, 4: 333–351.
25. Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen. 2016. TUT Sound events 2016. <https://doi.org/10.5281/zenodo.45759>
26. Matthias Mielke and Rainer Brueck. 2015. Design and evaluation of a smartphone application for non-speech sound awareness for people with hearing loss. In *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*, 5008–5011.
27. Alex Mihailidis, Amy Cockburn, Catherine Longley, and Jennifer Boger. 2008. The acceptability of home monitoring technology among community-dwelling older adults and baby boomers. *Assistive technology* 20, 1: 1–12.
28. Matthew S Moore. 1992. *For Hearing people only: Answers to some of the most commonly asked questions about the Deaf community, its culture, and the "Deaf Reality"*. Deaf Life Press.
29. Benjamin Petry, Thavishi Illandara, Don Samitha Elvitigala, and Suranga Nanayakkara. 2018. Supporting Rhythm Activities of Deaf Children Using Music-Sensory-Substitution Systems. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, 486:1--486:10.
30. Alisha Pradhan, Kanika Mehta, and Leah Findlater. 2018. Accessibility Came by Accident: Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 459.
31. Halley Profita, Reem Albaghli, Leah Findlater, Paul Jaeger, and Shaun Kane. 2016. The AT Effect: How Disability Affects the Perceived Social Acceptability of Head-Mounted Display Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2016)*, 4884–4895.
32. 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. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2: 17.
33. Dimitar H Stefanov, Zeungnam Bien, and Won-Chul Bang. 2004. The smart house for older persons and persons with physical disabilities: structure, technology arrangements, and perspectives. *IEEE transactions on neural systems and rehabilitation engineering* 12, 2: 228–250.
34. I R Summers, M A Peake, and M C Martin. 1981. Field trials of a tactile acoustic monitor for the profoundly deaf. *British journal of audiology* 15, 3: 195–199.
35. Stefan Tilkov and Steve Vinoski. 2010. Node.js: Using JavaScript to build high-performance network programs. *IEEE Internet Computing* 14, 6: 80–83.
36. EventDrops: A time based / event series interactive visualization using d3.js. Retrieved September 15, 2019 from <https://github.com/marmelab/EventDrops>
37. Luke Knox - Visualoop. Retrieved September 15, 2019 from <http://visualoop.com/blog/33248/portfolio-of-the-week-luke-knox>
38. PyAudio. Retrieved September 15, 2019 from <https://people.csail.mit.edu/hubert/pyaudio/>
39. PM2 - Advanced Node.js process manager. Retrieved September 15, 2019 from <http://pm2.keymetrics.io/>
40. BBC Sound Effects. Retrieved September 18, 2019 from <http://bbcsfx.acropolis.org.uk/>
41. Network Sound Effects Library. Retrieved September 15, 2019 from <https://www.sound-ideas.com/Product/199/Network-Sound-Effects-Library>
42. UPC-TALP dataset. Retrieved September 18, 2019 from <http://www.talp.upc.edu/content/upc-talp-database-isolated-meeting-room-acoustic-events>
43. Mobvoi TicWatch E2. Retrieved September 15, 2019 from <https://www.mobvoi.com/us/pages/ticwatche2>