Big Data Research 25 (2021) 100224

Contents lists available at ScienceDirect

Big Data Research

www.elsevier.com/locate/bdr

The Use of Smart Home Speakers by the Elderly: Exploratory Analyses and Potential for Big Data

Sandeep Purao *, Haijing Hao, Chenhang Meng

Bentley University, Waltham, MA 02452, USA

ARTICLE INFO

Article history:
Received 4 May 2020
Received in revised form 30 January 2021
Accepted 16 March 2021
Available online 2 April 2021

Keywords:
Smart speakers
Conversational devices
Interaction
Elderly
Visualization
Sonification

ABSTRACT

Contemporary smart home speakers (Amazon Echo™, Google Nest™) represent an opportunity to capture streams of big data that represent user interactions. Once captured, these data streams may be analyzed to understand use patterns for different demographics. In this research, we focus on one demographic, the aging-at-home elderly. We describe how such data streams may be captured with appropriate safeguards from smart home speakers deployed in the homes of the elderly individuals. We analyze the data to develop visual representations and sonifications to discover several patterns about aggregate use, use intensity, use progression over time, and types of use. Based on these, we draw tentative conclusions about the aging-at-home elderly make use of smart speakers at home, identify some problems, and describe how the empirical findings may be used to discover new voice skills that can contribute to the quality of life for the elderly. The paper concludes with a discussion of possibilities for applying the approach we have demonstrated for larger data sets.

1. Introduction

Conversational technologies have evolved [43] from early text-based software platforms (such as Eliza[1]) to specialized voice devices, described as 'smart home speakers' (such as Amazon Echo™ and Google Nest™). A smart home speaker is a type of speaker and voice command device that (a) utilizes Wi-Fi or other network, and (b) offers hands-free activation with one or more 'wake words,' to (c) perform user-specified tasks using built-in software, (c) access and manipulate data and services on the internet or (d) control home devices. The dominant examples of smart speakers in the marketplace (Amazon Alexa™ and Google Nest™) use these capabilities to: deliver news, announce calendar events, report weather, read out recipes, and control light switches, among other things. It is estimated that more than 25 million people today use these smart home speakers [31,26]. For the commercial technology companies, they represent the possibility of entering everyday space such as home and work [2] with the possibility of tapping into so-called behavioral surplus[2] [51]. Others point to the potential of such smart home speakers to become a normal, persistent and important part of households[3] [35]. Scholars and journalists suggest that the introduction of these smart home speakers raises issues of privacy [9,11,12,21], but also presents the possibility of anthropomorphizing, with the expectation that "we [would] communicate with them, not [only] through them" [2].

In spite of the almost ubiquitous presence of these smart home speakers, few scholars have reported empirical or conceptual investigations of how users *actually* interact with these devices, how they *perceive* these devices, and how they incorporate these devices in their lives [43]. Such studies are important because industry reports (e.g. [31]) cannot provide this in-depth understanding, and because these devices have the potential to inhabit out daily lives. These devices naturally capture ongoing data streams, which provide the potential for exploration with a big data perspective (beyond research methods such as interviews or surveys). Different user populations such as children, adults, the elderly and others present additional opportunities for such scholarly investigations. In this research, we focus on one segment, the elderly, aging-at-home [28].

The phrase 'elderly' or 'aging' describes individuals who are 65 and older, an increasingly larger fraction of the population across the planet (e.g. 15% of the US population and 27% of the population of Japan is 65 or older as of 2017 [50]) with growth rates that re-

* Corresponding author.
E-mail address: spurao@bentley.edu (S. Purao).

[1] See Eliza, early natural language processing system at https://en.wikipedia.org/wiki/ELIZA (Accessed 15 November 2020).

[2] The behavioral surplus describes how an organization can collect personal information and leverage it to develop new revenue streams.

[3] Amazon Echo™ claims to work with more than 20,000 smart-home devices representing more than 3,500 brands [45].

https://doi.org/10.1016/j.bdr.2021.100224
2214-5796/

main twice that of the overall population [3]. This is an important population segment because the smart home speakers represent a timely example of how the elderly interact with emerging technologies, and because of the potential for these devices to generate data streams that provide the potential for big data analytics. The question that drives this research is, therefore, the following:

- How do elderly individuals *actually* use the voice devices in their homes?

To carry out the investigation, this study uses one specific kind of smart home speakers, Amazon Echo™, as a representative example of location-locked[4] smart home speakers that are always-on and respond to commands that follow a wake-word (in case of Amazon Echo™, 'Alexa'). A typical interaction with such a device starts with a wake-word followed by a question (e.g., "Alexa, What's the weather today?") or a command (e.g. "Alexa, play music."). The few studies that have examined the use of these devices have provided descriptions such as the placement of devices at home, types of commands used and other similar statistics [35], and emerging analyses that are starting to theorize about use types [36,5]. In spite of the widespread market penetration of home speakers, such studies have remained rare, and none have explored the potential for analyzing the big data streams that these devices generate.

The primary goals of this work are, therefore, twofold. First, we explore how the aging-at-home elderly interact with and use such smart home speakers. To do this, we rely on visualization and sonification of the data gathered. Second, we illustrate how streams of data from these smart home speakers may be captured and analyzed without compromising privacy concerns to demonstrate the potential for big data analyses. This is important because these devices are often tied to personal information and shopping, making it difficult to obtain and analyze *actual* usage data (see, e.g. [43]). We respond to these goals by capturing data from smart home speakers deployed in the homes of the aging-at-home elderly, employing mechanisms to ensure the users' confidentiality; and illustrating how these data streams can be analyzed with visualization and sonification techniques. The key contributions of our work include: discovering patterns of use of smart home speakers use by the aging-at-home elderly; and describing techniques for obtaining and analyzing data streams from such devices while maintaining confidentiality, that scholars can use for future studies.

The remainder of the paper is organized as follows. Section 2 describes the background and related work to highlight gaps in prior work that we address. In section 3, we describe the research setting, data collection procedures, and data analyses. Section 4 details the findings with several visualizations and sonification of selected streams with a view to highlighting important patterns of home smart speaker use by the elderly. In section 5, we provide a discussion of the findings drawing connections to work elsewhere. Section 6 provides some concluding remarks and pointers to future research.

2. Background and related work

Over the last few years, smart home speakers have made significant inroads and are now commonplace in our homes [31,26,37]. However, in spite this market presence, there have been few empirical studies of their actual use [43]. As described above, a smart home speaker is a hardware device that works in an 'always-on'[5] mode, contains software triggered by a wake-word, listens to voice commands, converts these to text, develops a response that may require accessing resources on the internet, and speaks the response and/or carries out the command. They are designed with a set of built-in capabilities such as setting a timer or playing a radio station. More specialized capabilities can be added (called voice Skills for Amazon Echo™), similar to apps on mobile phones, which allow the smart speaker to respond to specialized commands and perform complex tasks such as controlling other devices such as light switches.

Few empirical investigations report how these smart home speakers are actually used by different users (e.g. [35]). Exceptions include some industry reports [31], likely driven by commercial interests, which cannot provide in-depth understanding of how different population segments use these devices. Such understanding is important for several reasons. First, as these devices are deployed in more homes, they are becoming part of our daily lives. It is, therefore, important to develop a deeper understanding of how they contribute to or hinder our lives [38], including any concerns related to privacy [45]. Second, an understanding of the modes of use can provide useful precursors to developing effective voice skills that respond to unique needs of different types of users. Third, a concern specific to the work reported here, is that these smart home speakers naturally capture data streams about how different individuals interact with these devices. The potential for big data analytics with these data streams remains much broader, compared to the design of specialized voice skills. They also provide a path to understanding the use patterns, separate from the use of other research methods such as interviews. Different user populations such as children, adults, the elderly and others – provide significant opportunities for scholarly investigations. In this work, we focus on one segment, users who are elderly, aging-at-home.

The phrase 'aging-at-home elderly' describes individuals who are 65 and older who have chosen to retire from their work and profession, moving on to volunteering with the community, spending time with family and friends, or pursuing hobbies or other endeavors. With advances in healthcare, better diet and exercise, this demographic makes an increasingly larger fraction of the population (15% of the US population, and 27% of the population of Japan as of 2017 [50]). Growth rates for this demographic remain twice that of the overall population [3], emphasized in the indicators such as rising costs of healthcare and services. This is an important population segment because individuals (although they age at different rates) eventually face physical and cognitive degeneration, limited mobility, visual and hearing impairments, and high illness susceptibility[6] [20]. These concerns shape and heighten the values (e.g. need for control, need for self-respect, need for a purpose in life) that these older citizens hold dear, often different from their younger counterparts as seen elsewhere in the design of service platforms [11,46].

Our study uses one specific kind of smart home speaker, Amazon Echo™, as a representative example. These devices respond a wake-word (for Amazon Echo™, they include words such as 'Alexa,' 'Echo,' or 'Computer'). When the smart home speaker is deployed at home and connected to the Wi-Fi, an aging-at-home elderly individual can interact with such a device using a wake-word followed by a question or a command. An example may be: "Alexa, What's the weather today?" This would result in a response from the home speaker: "In Boston, it is 80 degrees with a 20%

[4] Other voice-driven "agents" appear in different forms, such as those embedded in our mobile devices (e.g. Cortana™, and Siri™), which are different from the home-based voice devices we focus on.

[5] Most have the option to be 'muted,' so they are no longer listening for the wake-word.

[6] These concerns have been further heightened with the covid-19 pandemic.

chance of rain." Another example may be: "Alexa, set a timer for 8 minutes." This would result in a response: "Eight minutes, starting now," and the timer would chime after eight minutes. Yet another example may be: "Alexa, play BBC News." This would result in connecting to a service deployed on the web and playing the news stream on the home speaker. Although empirical studies of smart home speakers have remained rare, a few have reported some findings and highlighted some concerns.

The few studies reported have focused on understanding general use in the home environment [32,36] or on the privacy and security aspects of this "always on" and "always listening" technology [24]. Pyae [38] lament this by pointing out that in spite of their widespread use, "little is known about ... user experiences ..." They report results from a web-based survey of a little more than one hundred participants (about use of the Google Home smart speaker) as challenges encountered by the users, including: not understanding non-English words, the need to repeat commands, lack of portability, difficulties integrating with other devices, problems related to background noise or multiple voices, and inability to issue multiple commands in a single transaction. Sciuto et al. [43] study a small number of users to provide descriptions of use such as the placement of devices at home, types of commands used and other similar statistics. Li [23] employ a different research strategy. They examine evaluation and reviews related to these smart home speakers and the associated software to analyze and extract topics that the users find important to describe. With the 2,000+ reviews, they identify fourteen topics, which reveal three dimensions: technical concerns (e.g. connections and challenges); daily use (e.g. music, radio, remote control functionality, weather); and specific use for children (e.g. children's song, children's story, fun). The difference between different user demographics is also reported by Druga et al. [8] who found that children, aged 3 to 10, interacting with text-based and voice-based agents (including Alexa) in a lab setting interpreted the intelligence of agents differently when the agent was able to speak.

One empirical study reported by Scuito et al. [43] points out that they were "unable to find any research that analyzed the logs produced by smart home speakers or performed in-home interviews to understand how households actually use the devices" (p. 859). They recruited users from an online forum (median age 32, mean age 33.5, 72% male) with longer term smart home speaker ownership (average ownership 368 days). Their usage reports showed average number of commands per day, where the households quickly settled into a stable usage level after some experimentation, so that the level of usage remained constant week-over-week. They also reported that bedrooms were the most common location for the smart home speakers, households with more devices used more commands for controlling home devices, and owning multiple home speakers resulted in a higher number of interactions per day. The daily patterns of usage included categories such as music, weather, timer and others. This study, although similar in spirit to the study we design and report here, did not address the demographic (aging-at-home elderly) that we focus on, and required the participants to have the ability to work with self-reporting usage logs for the smart speakers using a web extension the researchers had designed.

Some other empirical studies include Bentley et al. [4] who found that the median household performed 4.1 commands per day, and that the most frequent requests were for music (40%) and information (17%). Lopatovska et al. [26] identified information searches, content requests for entertainment (e.g. jokes, music or games) and the control of other home devices as the most common requests. Pradhan et al. [34] explored use of smart speakers by people with disabilities and reported that the smart speakers were found to be easier compared to other devices, improved independence, and described unique uses such as learning and memory support. The investigation by Porcheron et al. [33] showed that commands to smart home speakers were embedded in everyday activities. Other than these studies, which did not address the population segment we focus on or explore the potential for big data analytics, we were unable to find other empirical investigations.

We did, however, find emerging conceptual explorations about smart home speaker use. For example, Purington et al. [36] analyzed Alexa reviews of smart home speakers and found that owners who personified Alexa or had multiple household members indicated greater user satisfaction. Pradhan [35] also explore this by examining how users categorize these smart home speakers: as just another object (like a toaster) or as a social companion. To understand this possible move towards anthropomorphization, they deployed smart home speakers (Echo Dot™, a smaller of Amazon Echo™) in the homes of older adults for a three-week period. They found that participants do both – personify as well as objectify the smart home speakers, sometimes within the same sentence. The work from Brause and Blank [5] points to the most recent example of conceptual explorations and theorizing. They develop a typology of six use genres building on domestication theory and affordances, which include: companionship, self-control and productivity, sleep aid, health care, peace of mind, and increased accessibility. They argue: while a use genre is not technologically determined; instead, the users' agency guides their emergence. With these use genres, they suggest that it may be possible to go beyond simple characterizations such as convenience [27] and entertainment [26] have been previously identified in different contexts to identify long-term, stable, and more generalizable categories.

The review above points out that prior research has examined use of smart home speakers by multi-user households [35], children, people with disabilities [34], and in low-income regions [39] with methods such as surveys, analysis of online reviews, and deployments in homes. Our work is most similar to the latter in that it involves a deployment. However, the particular demographic we focus on, the aging-at-home elderly, have not been part of such empirical investigation of their use of smart home speakers. The lone study by Scuito [Ref] that has examined such usage, the participants have been younger (mean age 32). Our focus on this demographic is important. The potential for smart home speakers to support aging in place has been recognized by Choi [6] whose findings (based on interview data) indicate that older reports a positive experience with the smart speaker (asking practical questions, setting medication reminders) suggesting their potential for maintaining their independence. The work we report in this paper, thus, adds to this stream of research. As described above, we focus one population segment: the aging-at-home elderly, and explore a key characteristic of smart home speakers: generation of data streams amenable to big data analytics.

3. Methods

In this section, we describe the research setting along with the rationale and strategy for data collection (Section 3.1), and describe how the data was analyzed (Section 3.2). Our decision was to follow a direct data collection strategy, relies on logs generated by the smart home speakers, instead of interviews with users. Several reasons led to this decision. First, an interview-based approach can only provide a partial view of how the users think they are using the smart home speakers, and is known to be faulty because of problems with recall and precision [12,26]. Second, the specific demographic we are working with (the aging-at-home elderly) are even more susceptible to these cognitive limitations that prevent reliable recall [12]. Third, working with this demographic is fraught with problems such as simply having access on a regular basis for interviews (e.g., if we were to conduct these interviews periodically), further worsening the recall problem. Fourth, even if we were to conduct the interviews, many specific details such as the number of commands during a given day, the exact manner of use, times or days when these were given would be impossible to obtain. Fifth, much prior work [48,49] shows that what participants describe in response to questions may not be the same as what they *actually* do; e.g., they may state hypothetical instead of real situations, they may not be willing to say what they did (or would do). Finally, conducting interviews can be expensive, particularly with this demographic and with multiple visits to the participants. The specific techniques we describe next are driven by these considerations.

3.1. Research setting and data collection

Driven by our research goal, we collaborated with the Council on Aging in one of the cities surrounding greater Boston. The city has a nominal population of ∼60,000; one sixth of the population, ∼10,000+ are elderly. The city is home to multiple large corporate headquarters as well as universities. The Council on Aging is responsible for coordinating activities and support services for the elderly with programs such as meals on wheels, classes for the elderly, trips to different local facilities, visits to groceries and hospitals, and others. The Director of the Council participates in city administration along with the Mayor, Councilors, and appointed members of the government.

Working with this Council on Aging, we recruited elderly individuals for this study with the following qualifications: (a) have Wi-Fi available at home to use the Amazon Echo™ smart home speaker, and (b) not presently have this or any other smart home speaker at home. Following appropriate research protections, a sign-up sheet was made available at the Council. Individuals who expressed willingness were contacted by the research team. A total of 10 individuals expressed interest. Of these, two could not be a part of the deployment for various reasons. The smart speakers were, then, installed in the homes of 8 individuals. The research team visited the homes, activated the devices, and helped the elderly with a simple walk-through of how to use the device, and how to recover if it stops working. As a part of this, the researchers helped connect the smart speakers to WiFi at home. The smart speakers contained pre-installed apps (called voice skills) for doing tasks such as 'setting a timer' or 'setting an alarm' or 'telling a joke.' No changes were made to this base. However, the researchers disabled functionality related to individual information (e.g., physical address, payments, shopping for products or services). This was done to ensure that any data would remain anonymized.

The setup and deployment, therefore, did not use any personal information of the individuals such as location, identities or financial information. Instead, the research team registered the devices to dummy accounts with nominal identifiers. These precautions allowed complete use of the smart speakers but not for purposes such as shopping. We were then able to focus on the role of the device as a voice device or a conversational device, instead of other roles it may assume (e.g., 'a shopping assistant'). Minimal training was provided to the elderly. A one-page document described how to use the smart speaker, and showed some possibilities such as "Alexa, what's the weather today," "Alexa, should I carry an umbrella," "Alexa, is the pharmacy open today," and so on.

The research team retained access to the credentials, which allowed scraping of data from the server about how each individual used the device. This was different from the mechanism used in [Scuito Ref], which required direct involvement from the users (a difficult proposition for the aging-at-home elderly). Following the deployment, one of the co-authors periodically scraped data for each participant by using the normal login procedure (on average, every two to three weeks), and accumulated the scraped data. During the study, one participant opted out. For the remaining participants, the time span for the usage varied from 178 days to 413 days. Table 1 shows the seven participants and the time periods during which the data was scraped.

Table 1
Participants and duration of data scraping.

Participant	Device installed	Data scraped till	Number of days
A	Apr 2019	Oct 2019	178
B	Feb 2019	Feb 2020	360
C	Feb 2019	Feb 2020	377
D	Feb 2019	Jan 2020	331
E	Jan 2019	Feb 2020	405
F	Jan 2019	Feb 2020	413
G	Jan 2019	Feb 2020	410

Table 2
An example of scraped data.

Date	Day	Time		Command	Voice recording[a]
7/18/19	Thursday	5:34	PM	Alexa	*Not captured*
7/18/19	Thursday	5:34	PM	Set a timer for 6 minutes	*Not captured*

[a] Although we did not capture these recordings in response to the privacy concerns, we indicate this as a possibility that other researchers may pursue with additional institutional review board protections.

The participants were all older than 65, aging-at-home, and retired from their careers and work. We did not collect additional data such as specific location, exact age, or other demographic details such as marital status to further protect the users' privacy. After the smart speakers were deployed, data scraping was done with a Python script that accessed the front-end, overcame problems such as differential display (text, time), and partial display (screen size, week). Privacy safeguards remained. The scraping effort did not capture any voice recordings. As part of our continued effort to ensure privacy, we also did not attempt to distinguish different users of a device (such as the individual or spouse or visitors to the house). Each record in the resulting dataset, thus, represented one interaction with Amazon Echo™ at a given time. Table 2 shows an example of data scraped from one of the participants.

3.2. Data analysis

To facilitate analysis, we constructed a simple conceptual model that showed the concepts and relationships across the concepts (see Fig. 1).

Relying on this conceptual model, several additional data elements could be computed such as days the device had been active, an indicator to capture whether the command was understood, time elapsed since the previous interaction, number of days of consecutive use, and others. The research team relied on the Tableau™ platform for loading the data and generating visualization charts. Selected data streams were extracted and loaded into the TwoTone™ platform for generating sonifications. The procedures followed for conducting these analyses involved the following steps.

First, raw data, loaded into Excel™ was moved to SAS™, organized into structured datasets, and then imported into Tableau™ to generate visualization charts. These included summarizations, descriptive statistics and analytical displays. Some outlier data was removed to generate additional visualizations, which we explored in the aggregate as well as for different participants. This, first round of analyses, pointed to several interesting patterns, that we explored for each participant. One example of this analysis was the frequency of usage, and how it changed over time. This pointed to additional analyses that could then be conducted, such as use and interaction at different times of the day, or during weekdays vs. weekends. Additional analysis possibilities were detected such as

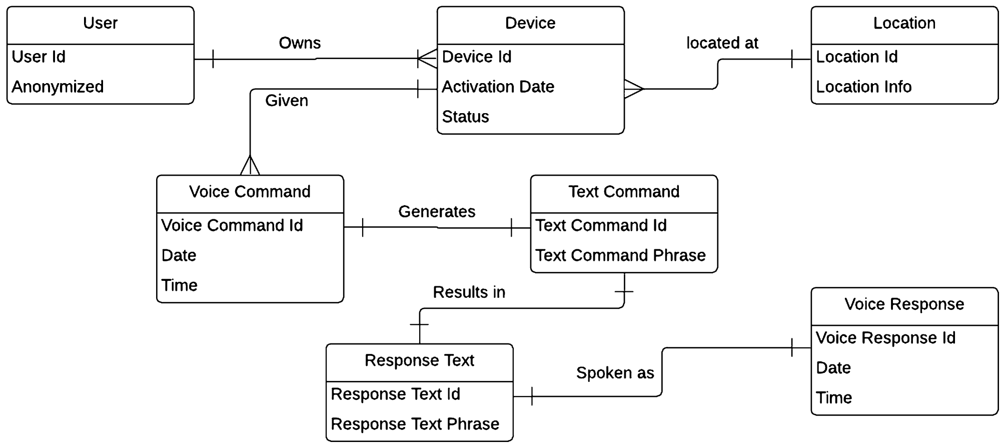

Fig. 1. A conceptual model of the scraped data.

Table 3
Participants and number of use instances.

Participant	# Use instances
A	939
B	254
C	1,410
D	882
E	1,887
F	747
G	714

the number of consecutive days of use, and the number of consecutive days of non-use. We exported the data as needed, and generated sonifications to overcome traditional constraints such as page size for analyzing datasets with time-dependent variables [22,47], treating it as an exploratory technique. Additional cycles of analyses continued with more nuanced representations of interactions and use such as the types of commands used, the length of interaction episodes and others. The findings are described in the next section.

4. Findings

The data set consisted of 6,833 use instances from 7 users, ranging from 254 instances (Participant B) to 1,887 instances (Participant E). Table 3 summarizes. Each represented an interaction (such as issuing a command or request from the user, aimed at the smart speaker).

4.1. Characterizing aggregate use

The first set of analyses examined how regularly the elderly individuals used the smart home speakers. To do this, we explored active vs. non-active days. Fig. 2 summarizes.

Next, we explored the occurrence of use on different days and at different times. Fig. 3 shows the occurrence of use at different times during the day.

Participant E had the most usage (1,887 instances), whereas participant B had the least usage (254 instances). In aggregate, 2,998 (43.9%) of use instances occurred during the hours before Noon, whereas 3,835 (56.1%) occurred during the hours after Noon. This split, however, varied significantly across the participants. Consider, for example, participant D for whom 22% (197) of the use instances occurred in the morning hours. The corresponding number for participant A was 65% (614) of the use instances. Participant A was also the only participant who interacted with the device more in the hours before Noon. We continued the exploration by examining usage across weekdays versus weekend days. In aggregate, 4,245 (63.6%) of use instances occurred during weekdays, whereas 2,488 (36.4%) occurred during the weekend. There were differences across participants. Figs. 4a and 4b show the distribution of use instances on weekdays versus weekend for each participant.

Participants B and C had more use instances during the weekend. This was not the case for others. Participant E had most use instances (75%, 1,412) during weekdays, whereas participant C had the least (44%, 621). Our analyses continued with an examination of use intensity. We calculated the average number of use instances per day (2.9, across the entire duration of use), and the average number of use instances for each *active* day (5.9 across the number of active days). Fig. 5 summarizes.

The analyses continued with an examination of each participant's use with descriptive statistics. Fig. 6 shows the box-plot displays for each participant with the quartiles, median, and the minimum and maximum. The scale has been adjusted to remove the display of outliers. A visual inspection shows that the profiles for the participants did not show dramatic differences, with median use between 3 and 6, and the average use between 4 and 7.5 use instances. Fig. 6 summarizes.

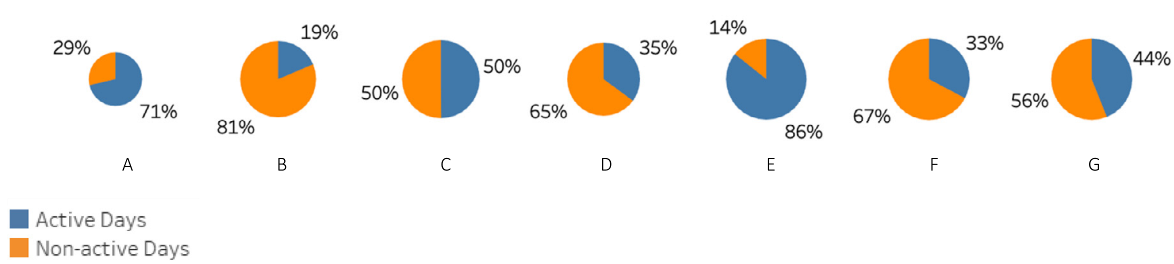

Fig. 2. Active vs. non-active days.

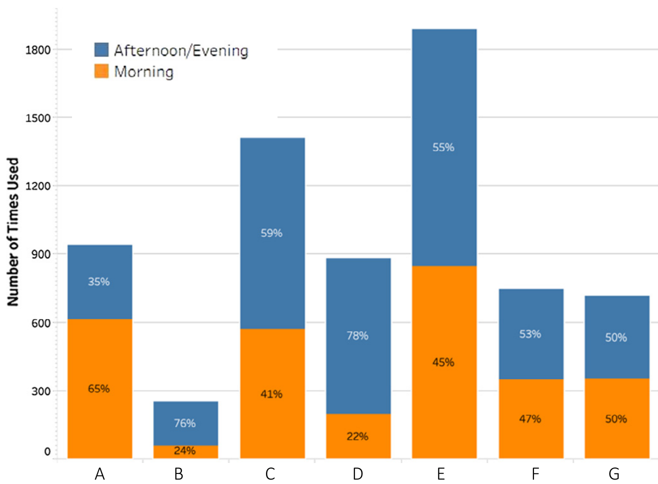

Fig. 3. Use at different times of day.

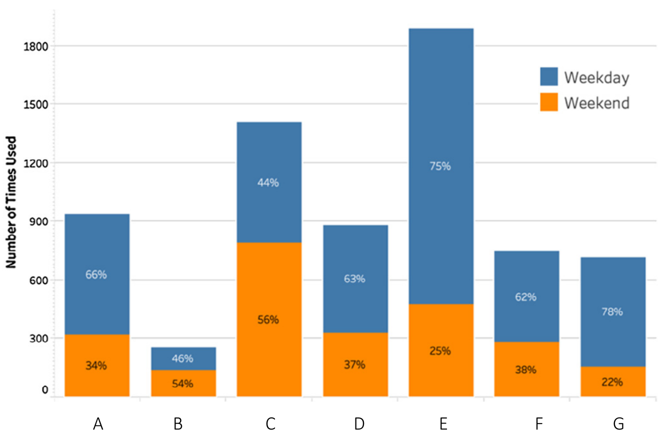

Fig. 4a. Use during Weekdays vs. Weekends.

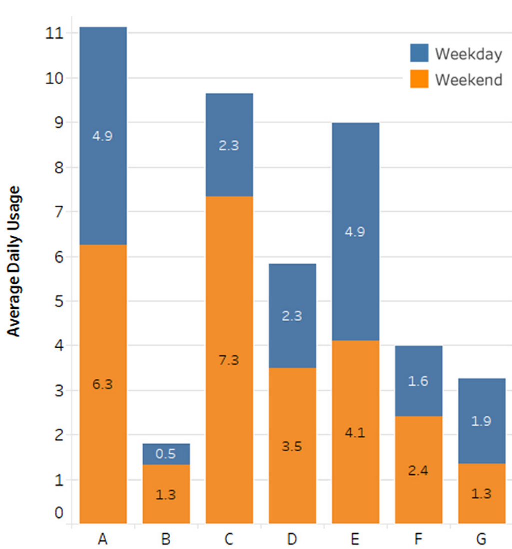

Fig. 4b. Use during Weekdays vs. Weekends (Normalized for Number of Days).

Although there are some differences (participant C had 196 use instances on one of the days; participant F had 84 use instances on another day), the median and the mean, as well as the lower quartile were comparable across all participants. Participants A and D had higher third quartile numbers (10 and 10 respectively), indicating higher number of use instances more frequently.

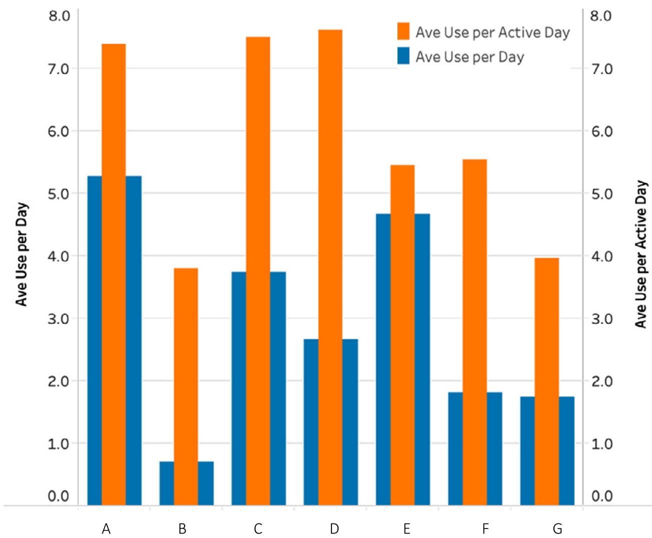

Fig. 5. Use Intensity (Active Days vs. All Days).

4.2. Characterizing use progression

The next set of analyses were focused on analyzing how the use progressed over time, and other, time-based characterizations of use such as continuous days of use or gaps in use. We start with Fig. 7, which shows the trajectories of use, mapped as the total number of use instances for every month for each participant over the complete timespan. Other than participant E, who shows sustained and even slowly increasing number of use instances over time, most others appear to indicate a high number of use instances followed by a dampening. The peaks in usage appear in the first few months of use. Whether this is due to a learning curve or an initial curiosity or enthusiasm cannot be detected based on just these frequencies of use instances. Participants B and F show particularly strong peaks at the beginning, and others show a similar use pattern. Here, participant E exhibits a different use pattern. The dotted line in the figure shows the slow decline in the frequency of use instances by all participants (except participant E). The trend line represents a polynomial function, which fits the seven users' monthly use trajectories, significant at $p <= 0.05$.

To explore these progressions, we sonified the data streams with a more granular view, the usage is aggregated for each week, instead of each month. Sonification is a novel approach to present datasets with time-dependent variables that is gaining traction as a data analysis technique for 'social issues' but with few guidelines [22], although others have suggested its application for representing the temporal context [47] and for exploratory scholarly communication [42]. These were the primary reasons for our use of sonification. The design space for sonification choices is still being explored in different domains, where current scholarship still relies on making intuitive choices [44]. Our work, therefore, builds on the contemporary mode of application for sonification techniques, to demonstrate how it can explore the data streams from smart home speakers, with the possibility for varying window sizes.

The sonification results we present rely on data at the granular level of each week. Although this analysis may be presented as a chart, these displays can become difficult to comprehend with more granular time windows and larger number of time periods. More specifically, the charts can become increasingly cluttered because of constraints such as page width. In contrast, sonification allows use of any time window (e.g. weekly or even daily) without such constraints because the audio can be played across time. With different time window sizes, one can then drill down to more detailed granularity level if needed. In this way, sonification can

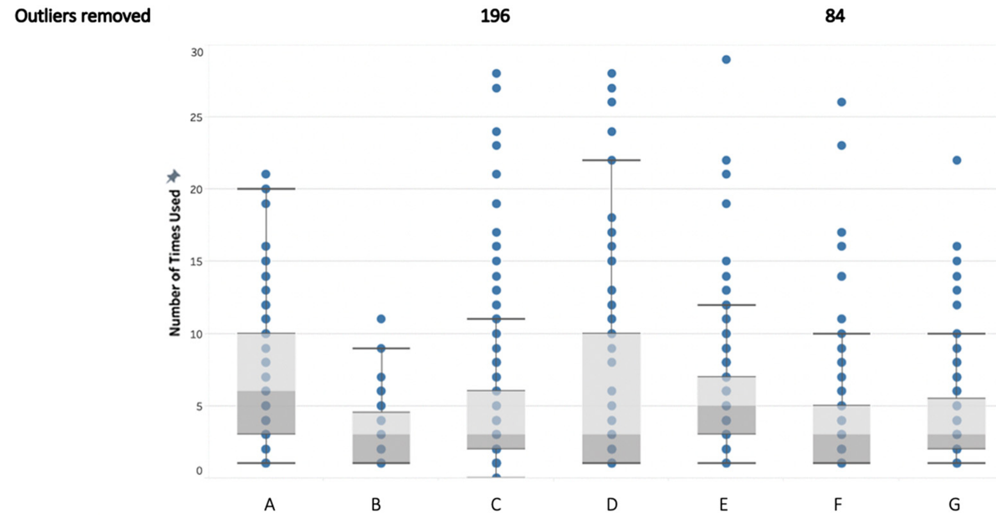

Fig. 6. Distribution of Use Instances (All Participants).

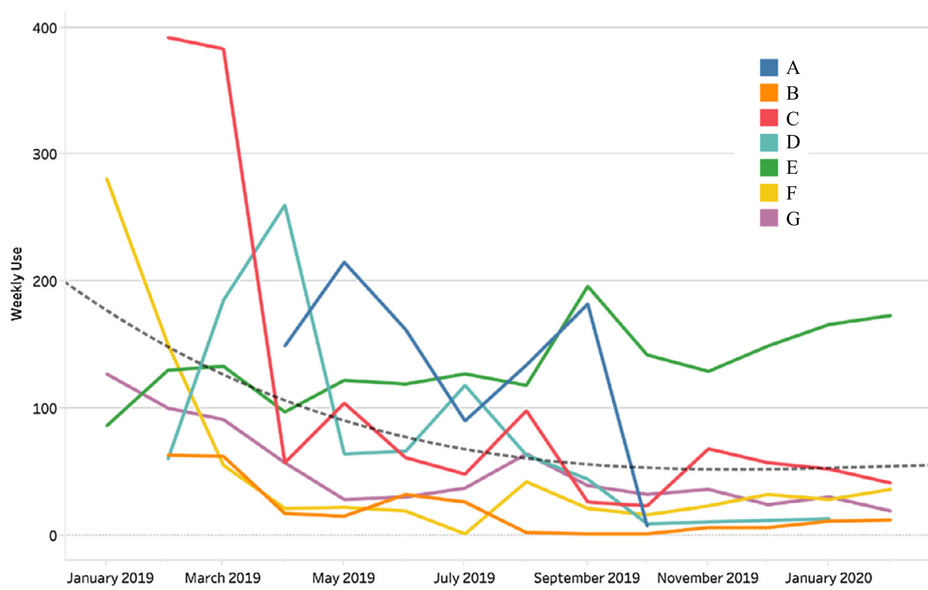

Dotted line in the graph represents a polynomial trend model of degree 7	
Use Instances =	-2.60776e-25*Month ^7 + 2.66924e-20*Month ^6 + -7.02555e-16*Month ^5 + 0*Month ^4 + 0*Month ^3 + 0*Month ^2 + 0*Month + 5.43437e+06
Number of modeled observations: 85 (filtered: 0); Model degrees of freedom: 4 (residual degrees of freedom: 81) SSE (sum squared error): 416319; MSE (mean squared error): 5139.74 R-Squared: 0.199685; Standard error: 71.692; p-value (significance): 0.0004093	

Fig. 7. Use over Time and Trend (*Monthly*, All Participants).

help overcome (at least partially) the challenge for science communicators attempting to balance data complexity, fidelity, and comprehensibility [42]. The sonification results, therefore, allowed us to ‘listen’ to the data streams, and compare these across different users. The links to the sonification appear in Table 4. The final column in the table indicates what to listen for in each audio track (14-seconds, representing several months of data for each participant). These sonified streams provide another angel to understand the differences of the usage patterns over temporal dimension across each participants.

As we listen to these audio tracks, the differences across participants are easier to appreciate. Consider, for example, two audio tracks. The continued high notes (indicating sustained use throughout) are evident for participant E, compared to the initial high notes followed by lower notes for participant C (indicating little use after the enthusiastic start). The table describes (in the last column) what a listener may discern from the sonification efforts. As another example, consider the distinction between participants B, C and F (initial high notes, indicating frequent and heavy use) compared to participant A (initial beeps, indicating infrequent early use) and participants D and G (mid-notes, indicating average use, whether regularly or periodically). The field of sonification is in its infancy [42], and therefore, our ability to listen to and interpret sonified data analysis is also being developed (compared to our ability to read and interpret various charts). Nevertheless, the sonification results we present, along with the interpretations suggested in the table suggest important possibilities.

Table 4
Sonification of *Weekly* Use over Time.

Participant	Audio track (0:14 each)	Listen for
A	https://bit.ly/Participant-A	Initial beeps, ongoing mid-notes, final beeps Late start, sustained use, early stop
B	https://bit.ly/Participant-B	Initial high notes, petering out after that Enthusiastic start, minimal use after that
C	https://bit.ly/Participant-C	Initial high notes, moderate tempo after that Enthusiastic start, some use after
D	https://bit.ly/Participant-D	Mid-notes, high notes, few low notes after Slow start, higher use, minimal use after that
E	https://bit.ly/Participant-E	Continued high notes throughout Significant sustained use throughout
F	https://bit.ly/Participant-F	Initial high notes, ongoing mid-notes after Enthusiastic start, some ongoing use after that
G	https://bit.ly/Participant-G	Periodic mid-notes, with some beeps Somewhat sustained use, with some gaps

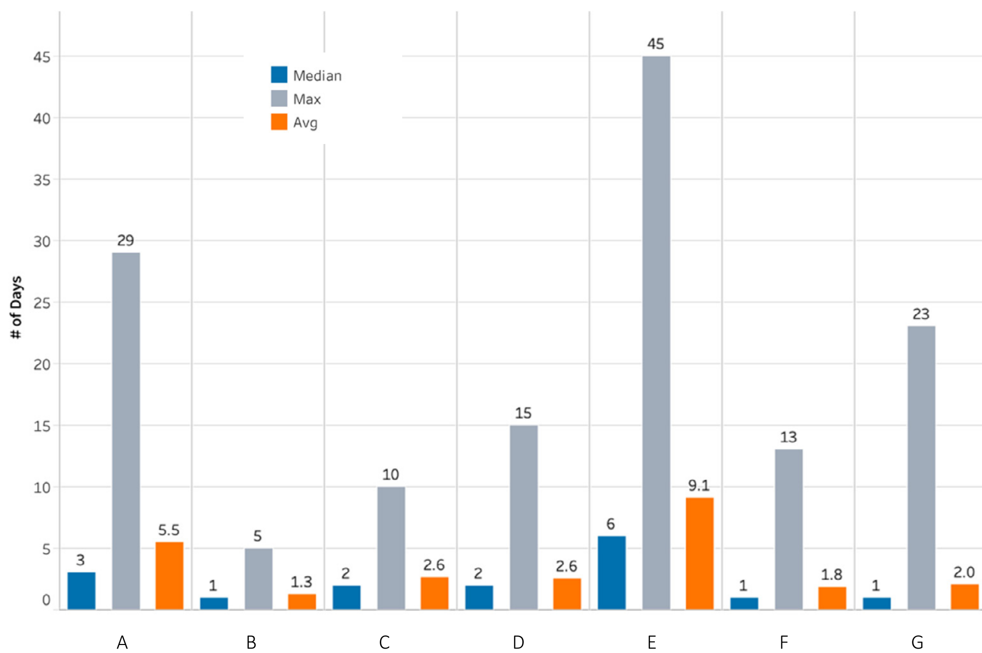

Fig. 8. Episode Lengths (All Participants).

To further investigate the use patterns, we examined the number of days of continuous use, defining it as an 'episode.' The word 'episode' is used in its dictionary sense to indicate[7] "an event or a group of events occurring as part of a larger sequence; an incident or period considered in isolation." Our intent is not use of the word in a medical sense, where it can describe an occurrence of a pathological condition. Instead, episodes represent parts of a larger narrative [, 2] akin to episodes of a television series. If a break occurred, the number of days of continuous non-use was defined as a 'gap.' Once the participant resumed the use, we treated that as the start of a new episode. If a participant were to use the device every day, the data would show one long episode; if a participant were to persist with non-use for several days, the data would show long gaps. Fig. 8 shows the median, max and average episode lengths for each participant.

With this analysis as well, we note that participant E had the longest episode (length: 45 days), although participants A and G also had episodes that were long (29 and 23 days respectively). The average episode length was, in fact, fairly high for participant A (5.5 days), which almost matched the median episode length for participant E (6 days). The concepts of Episode and Gap are naturally amenable to a longitudinal analysis. To do so, we chose two participants, Participant D (longest episode: 15 days, average: 2.6 days, 45 episodes), and Participant E (longest: 45 days, average: 9 days, 38 episodes) over approximately the same duration. Fig. 8 summarizes.

The differences are clearly evident. The use progression of participant D is marked by a very long (92 days) towards the end of the time duration, although there is a return to use episodes again. In contrast, the use progression of participant E is marked by increasingly longer use episodes. The two longest episodes (45 days and 40 days) appear towards the end of the time duration. Because of the time dimension, we sonified these to examine the progression. Table 5 shows these outcomes.

The visualization in Fig. 9 and the sonification in Table 5 draw on the same data with decisions such as colors and display (visualization) and instrument and tempo (sonification) by the research

[7] (Oxford English Dictionary, accessed Jan 2021.)

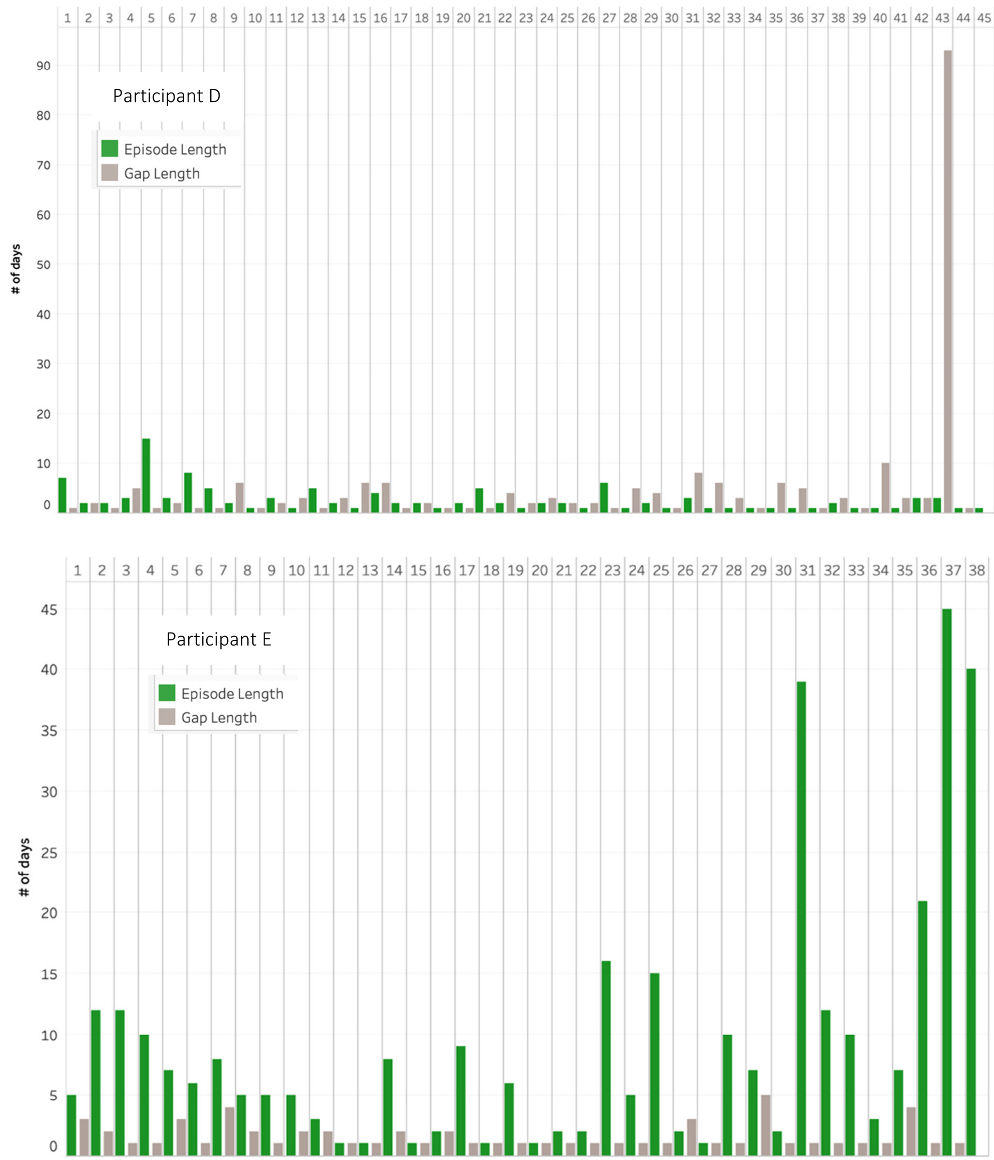

Fig. 9. Episodes and Gaps (Two Participants).

Table 5
Sonification of Episodes and Gaps.

Participant	Audio Track (0:27; 0.:28)	Listen for
D	https://bit.ly/Participant-D-Episodes-and-Gaps	A steady beat (pitch indicates gap length), overlay with few high notes (shorter episodes)
E	https://bit.ly/Participant-E-Episodes-and-Gaps	A steady beat (pitch indicates length of gaps), overlay with several high notes (longer episodes)

team. They present an interesting comparison of different ways of presenting data summaries.

4.3. Nature of use

The final set of analyses was based on the content of use instances. We first defined different 'types,' e.g., simply the wake-word 'Alexa,' a specific command, or an indication that the command was not understood. Fig. 10 summarizes.

The chart on the left shows the aggregate for use of the wake-word (2,011 instances), commands (1,729 instances), wake-word plus commands (1,774 instances), and commands not understood (206 instances) with the size of the bubble indicating the relative frequency. The smallest bubble indicates the frequency of commands not understood. The visualization on the right shows this information for each participant. One disturbing number in this chart is the 12% (87) commands not understood for participant D. Participant E again stands out as an effective user (highest occurrence of wake-word plus command). We note that the bubble 'wake word plus command' may indicate that the pace at which an utterance is understood by the device. For example, if a participant request (e.g., Alexa, what's the weather today?) is delivered in a halting way, it is possible that the device would record the first part (Alexa) as the wake-word, followed by the request (what's the weather today), i.e., two use instances. On the other hand, if a participant were to deliver the request quickly, the device may record the complete utterance as a single use instance. Across all instances, just the wake word 'Alexa' appeared as an utterance (i.e. it showed up in the data scraped as a single word on a line) separated from the command that followed a few seconds later (e.g.

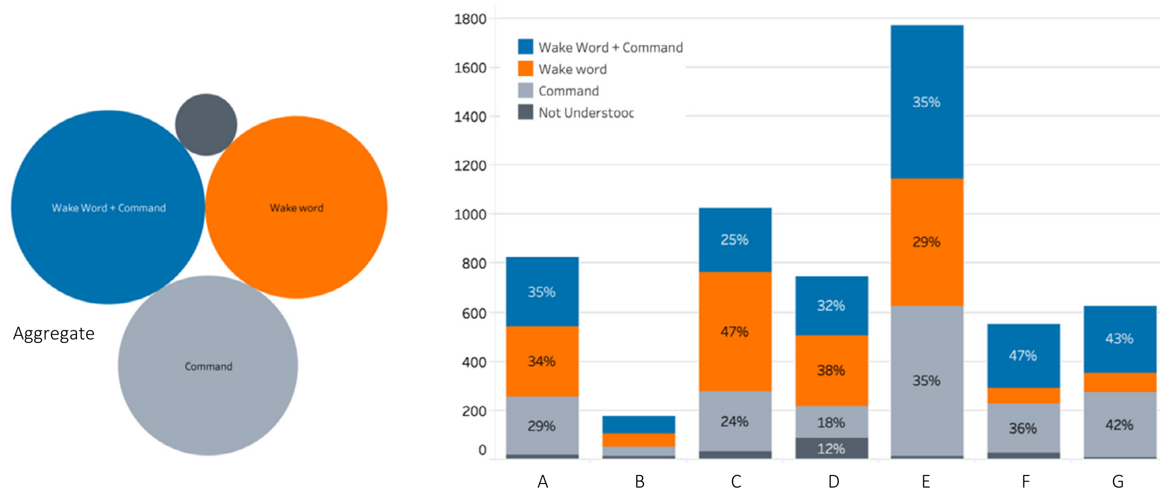

Fig. 10. Types of Use Instances (Aggregate and All Participants).

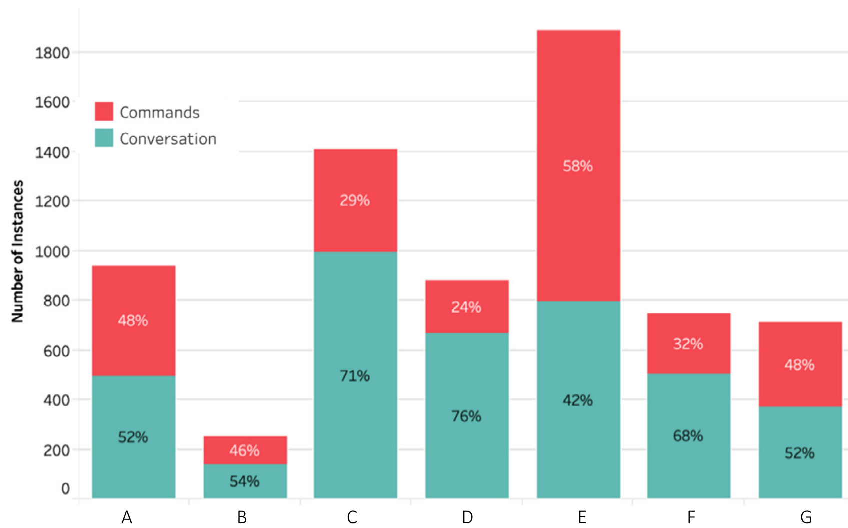

Fig. 11. Commands vs. Conversations (All Participants).

‘set a timer’) for 53.1% of the total occurrences of the wake word. In contrast, the wake word ‘Alexa’ appeared as a complete command (i.e. it showed up in the data scraped as the wake word followed by a command, e.g. ‘Alexa, set a timer’) for 46.9% of the total occurrences of the wake word.

To further understand whether such halting use would allow the participants to engage in a ‘conversation,’ we identified conversations, that is, a set of lines close to each other. We used a generous time period to determine this: if use instance Y followed use instance X within three minutes, we construed that the two lines were part of a single conversation. In the absence of such line Y, we construed that line X was a single command. Fig. 11 shows this analysis.

In aggregate, 2,827 use instances (42%) were commands. In contrast, 3,961 use instances (58%) were part of a conversation. Like other analyses, the percentages varied across the participants. Participant C appeared to deliver only 29% of the use instances (415) as commands. In contrast, participant E delivered as many as 58% of the use instances (1,093) as commands. We note that six participants showed more use instances that were part of a conversation (except participant E).

As a final analysis, we examine the different types of use instances. Here, we coded the data to identify use types such as playing music or news (indicated by the word ‘play’) or asking for

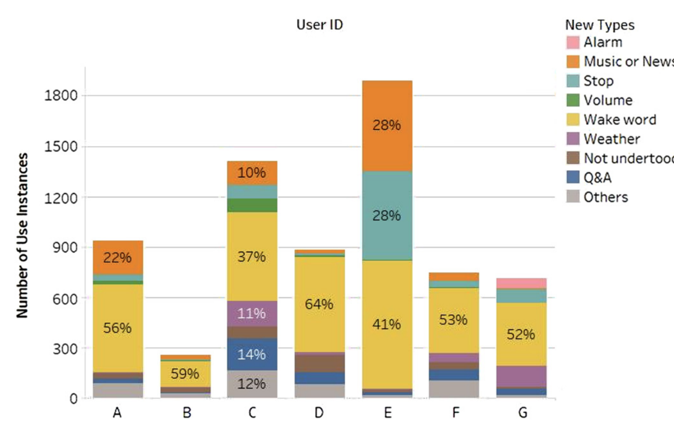

Fig. 12. Types of Use Instances (All Participants).

weather (indicated by the words ‘weather’ or ‘temperature’), and so on. Fig. 12 shows this analysis, which can be seen as the next step following the analysis in Fig. 10.

We note that use of the ‘wake word’ (i.e. appearance of just the wake word ‘Alexa’ as a single word on a line) separated from the command that followed a few seconds later (e.g. ‘set a timer’) remains a significant element in this analysis. Other than this, the categories 'alarm' and 'music or news' were frequent for several participants (e.g., Participants E, A, and C). Many also used the 'stop' command frequently (28% of total use instances for Participant E, and others including participants G, C, F and A). Interestingly, the 'weather' category was not as frequent (somewhat of an exception were Participants C, D and G). The small gray areas at the bottom of each bar indicate use instances in the 'other' category. These were instances such as 'start the timer' or 'tell me a joke' and several others. In essence, the differences across use patterns were also seen in the types of use instances.

Together, these analyses allowed us to visualize and sonify the patterns of voice device use at home by the elderly drawing on difficult to capture data from seven elderly, aging-at-home individuals who used these devices at home over a span more than a year on average. Although many of these may be seen as descriptive, we note that such descriptive, yet rich analyses have been largely absent [Refs], and therefore, provide an important window into how this population segment uses smart home speakers. In the next section we reflect on the conduct of the study, and suggest some interpretations of the findings.

5. Discussion

We start by emphasizing that the work we have reported, including the actual deployment of smart home speakers in the homes of the aging-at-home elderly, developing mechanisms to ensure confidentiality, and capturing the data streams in an anonymized manner – represent important lessons that provide a platform for larger studies. By doing so, we have demonstrated how such studies may be conducted, including the feasibility of capturing the data streams on an ongoing basis. The description of mechanisms for data capture along with the possibilities for visualization and sonification can provide a platform to further automate some of the phases with a view to examining a larger number of individuals with continued capture of data streams and analyses. Our research can provide the template for such efforts.

The analyses and the findings summarized in the previous section themselves provide some interesting pointers as well. We note that although some broad patterns may be discerned in the aggregate, the elderly individuals clearly have different profiles. As an example, note the differences across different times of day or days of week (see Figs. 3, 4a and 4b). Such differences are also noted in other patterns such as the use progressions (see Fig. 7 and the sonification in Table 4), as well as key differences in the use episodes and gaps (see Fig. 9 and the sonification in Table 5). The differences persist in the nature of use such as content of commands and length of conversations (see visualizations in Figs. 11 and 12). These findings appear to correlate to the heterogeneity noted in research related to life-course theory [28] in studies of aging. These differences are important also because they point to design possibilities for specialized voice skills[8] for the elderly. Although the path from empirical findings to identifying such specialized skills can require additional steps such as creative visioning, exploring feasibility, and developing blueprints [25]. Our contributions should, therefore, be seen as the first step in this process, by providing empirically grounded findings that show the use patterns of this specific demographic: the elderly – that are not based simply on perceptions but rather, actual usage of these devices.

Nevertheless, we outline some possibilities for future designers, particularly for the use of smart speakers as contributors for resourceful aging [10]. This may, for example, consist of simple medication reminders, searching for information about different ailments, entertainment such as playing music, social support such as calendar reminders, or companionship such as simple chatting (see, e.g. [13]). The importance of searching was also noted by Lopatovska et al. [26]. Our findings (see Fig. 12, types of use instances) suggest that specific support for such activities may be helpful for this demographic. Further, it may also be possible to devise algorithms that detect unusual change in usage (e.g., not checking in over several hours or multiple days depending upon the threshold set) as a signal of health risk to caregivers or family members, and as a possible trigger for a family member to use the drop-in feature in some smart home speakers.[9] We note that the differences in usage patterns can make such simple algorithms problematic. On the other hand, specialized skills (e.g. proactive prompts that require responses) may help to ensure that the elderly individuals remain active and minimize long gaps in use (see Figs. 8 and 9).

Another possibility is to accommodate potentially halting speech patterns (see Fig. 10 and the following interpretation); and they may need to be provided specialized skills appropriate for their unique needs (see Fig. 12). A similar finding (problems with non-English words or the need to repeat some commands) was also noted by Pyae [38]. Our findings emphasize this for the aging-at-home elderly, and add to this. More specifically, we find that separation of the wake word from the actual command is a clear indication of the halting pattern with which the elderly are able to issue commands to the device. It also suggests a possible obstacle to carrying out back and forth interactions (i.e. conversations) with the device for individuals from this demographic (different from the issues pointed out by Li [23]). The relatively few occurrences of conversations (in spite of a generous time break of 3 minutes) compared to commands (see Fig. 11) is also an indication that some technical and research challenges will need to be overcome to support longer conversations. As an example, consider the need to represent the context (e.g., semantic content, common ground, and discourse state such as the sequence of utterances and responses so far), to facilitate conversations [12]. Another possibility would be to explore approaches such as automatic methods for activity recognition [30] that may help recognize ongoing interactions and activities. Conversations (as interactions between the user and the device that goes beyond issuing a command) can be beneficial and important because they can support more complex tasks, and even respond to other considerations such as concerns of isolation for the elderly [14]. The need to support conversations has led to some emerging solutions [1]. We note that the long timespan that the participants used the devices provides a key pointer to the need for such solutions in the context of deployment of use of voice devices in the homes of the elderly. These concerns are different from previous empirical studies such as Scuito et al. [43] who focused on a participant pool that was much younger, and often had multiple smart home speakers deployed at home. On the other hand, our findings correlate well with those by Pradhan et al. [34] who explored the use of smart speakers by people with disabilities who found them to be easier compared to other devices.

We also note that our findings showed continued even in the absence of some key functionalities, such as personal calls or shopping, which we de-activated in favor of an anonymized and confidential approach to data capture and scraping. This is a noteworthy finding. With these analyses, we find further support to question the stereotyping of the elderly [4,15,18,19], further supporting the

[8] Similar to specialized 'apps' that we are familiar with on the mobile devices and phones.

[9] This feature allows an incoming call to be connected automatically, which may be useful for checking in on elderly parents (see https://amzn.to/390zIZJ, Accessed 15 November 2020).

heterogeneity hypothesis in life-course research about aging [28]. We find that the different use patterns we observed align with contemporary descriptions of the elderly such as: healthy, active, and independent individuals who possess a very different image of themselves [7]. Many of them are aware of and are effective users of technologies, having spent a better part of their lives working with and learning new technologies, challenging the traditional lessons of "gerontechnology" [20]. For this population segment, it is not sufficient to think of technologies merely in terms of 'assisted living technologies' to improve independent living and aging at home [3]. Examples include robots to carry out specific tasks [4], technologies to aid with cognitive impairment [29], and social robots to improve emotional and relational wellbeing [10,40]. In contrast, a broader conceptualization can include other technologies such as mobile phones and laptops and tablets, as well as home automation devices and voice devices like the once we have analyzed. With this recognition – that there is great variety in the everyday lives, needs and motivations of the elderly – our findings support the view that technology design must shift from making things that are following one certain path in a "foolproof" manner [15,16] to making things that lead to 'more resourceful aging' [10] by providing capabilities that can support multiple us patterns. This, more enlightened view, describes technology design for the elderly in a manner that is ethical, in response to the concerns and values of the elderly [10,19].

6. Concluding remarks

In this paper, we have explored data streams from smart home speakers (with Amazon Echo™ as the representative example) deployed in the homes of the aging-at-home elderly. Our intent is to report exploratory analyses of usage patterns by this demographic, not the strengths or weaknesses of specific devices (see [17,41]). The key contributions of our work include findings related to the patterns of use of smart home speakers by the aging-at-home elderly, presented as visualizations and sonifications. Compared to prior work that has explored such usage of smart home speakers [43,35,36,35], our investigation is focused on a new demographic that has remained unexplored so far. This in an important demographic because of the potential for these devices to support resourceful aging and improve quality of life of the elderly [6]. Our findings include characterizations of aggregate use, use intensity, use progressions, and different types of use. The conceptual model we have created provides the basis for structuring the data scraped and inferring/computing additional data. We have carried out the investigation in a way that maintains the participants' privacy by employing two mechanisms: (a) using accounts that were expressly created for this research to activate and deploy the devices, and (b) ensuring that that data collection or analysis did not include any voice recordings, location, personal or financial data of the participants. With these constraints, the data we gathered and analyzed still provided insights into how the elderly used the smart home speakers. This is an important outcome because it points to the possibility of exploiting the big data potential of the data streams generated by these smart home speakers to directly contribute to improving the quality of life for the aging-at-home elderly (e.g. as pointers to designing specialized voice skills for the smart home speakers). Both the approach for data capture, and the approach for visualization and sonification can be applied to larger datasets, while ensuring confidentiality. Our work is planned to progress by complementing these empirical findings with interviews and other methods such as surveys to further understand the rationale behind the findings we report.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We acknowledge participation from the elderly, and the Waltham Council on Aging. We also appreciate comments from the review team at the MoBiD'19 Workshop shaped our work. The work reported has been funded by the National Science Foundation under award number 1641148. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

References

[1] Amazon, Building conversations with Alexa (Beta), https://developer.amazon.com/en-US/blogs/alexa/alexa-skills-kit/2020/07/introducing-alexa-conversations-beta-a-new-ai-driven-approach-to-providing-conversational-experiences-that-feel-more-natural, 2020. (Accessed 15 November 2020).
[2] A.D. Baddeley, Is working memory still working?, Am. Psychol. 56 (11) (2001) 851–864, https://doi.org/10.1037/0003-066X.56.11.851.
[3] L.L. Barrett, Healthy @ Home, American Association of Retired Persons (AARP) Foundation, Washington, 2008.
[4] F. Bentley, C. Luvogt, M. Silverman, R. Wirasinghe, B. White, D. Lottrjdge, Understanding the long-term use of smart speaker assistants, in: Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 2018, pp. 91:1–91:24.
[5] S.R. Brause, G. Blank, Externalized domestication: smart speaker assistants, networks and domestication theory, Inf. Commun. Soc. 23 (5) (2020) 751–763.
[6] Y. Choi, G. Demiris, H. Thompson, Feasibility of smart speaker use to support aging in place, Innov. Aging 2 (2018) 560, Supplement 1, Gerontological Society of America (GSA) 2018 Annual Scientific Meeting.
[7] Barbara Czarniawska, Narratives in Social Science Research, Introducing Qualitative Methods, SAGE Publications, Ltd, London, 2004.
[8] S. Druga, R. Williams, C. Breazeal, M. Resnick, Hey Google is it OK if I eat you?: initial explorations in child-agent interaction, in: Proceedings of the 2017 Conference on Interaction Design and Children - IDC'17, ACM Press, 2017, pp. 595–600.
[9] Eoghan Furey, Juanita Blue, Alexa, emotions, privacy and GDPR, in: Proceedings of the 32nd International BCS Human Computer Interaction Conference (HCI), July 2018.
[10] E. Giaccardi, L. Kuijer, L. Neven, Design for resourceful ageing: intervening in the ethics of gerontechnology, in: P. Lloyd, E. Bohemia (Eds.), Proceedings of DRS 2016, Design + Research + Society Future-Future-Focused Thinking: 50th Anniversary International Conference, Brighton, UK, 27–30 June 2016, 2016, (Vol. 1).
[11] H. Gil, F. Amaro, Active ageing and the role of ICT and assistive technologies: reflections and discussion for their use in Portugal, in: E-case & E-tech, vol. 2010, 2010, pp. 2750–2760.
[12] J. Gogan, M.D. McLaughlin, D. Thomas, Critical Incident Technique in the Basket, 2014.
[13] M. Green, Context and conversation, in: The Wiley Blackwell Companion to Semantics, 2020, pp. 1–23.
[14] R. Harris, Why the former surgeon general worries about loneliness and older adults. Forbes, https://www.forbes.com/sites/nextavenue/2020/04/29/why-the-former-surgeon-general-worries-about-loneliness-and-older-adults/?sh=174030d74d91, 2020. (Accessed 15 November 2020).
[15] P.W. Harvey, I. Thurnwald, Ageing well, ageing productively: the essential contribution of Australia's ageing population to the social and economic prosperity of the nation, Health Sociol. Rev. 18 (4) (2009) 379–386.
[16] S. Hyysalo, Representations of use and practice-bound imaginaries in automating the safety of the elderly, Soc. Stud. Sci. 36 (4) (2006) 599–626.
[17] F.L. Ile, G. Ţigu, Medical tourism market trends: an exploratory research, in: Proceedings of the International Conference on Business Excellence, vol. 11 no 1, 2017, July, pp. 1111–1121.
[18] H. Kendig, C. Browning, Directions for ageing well in a healthy Australia, Acad. Social Sci. 31 (2) (2011) 23–30.
[19] Y. Khoury, S. Purao, M. Duffy, The Influence of Values on the Use of Citizen Services: The Elderly Perspective. International Conference on Information Systems (ICIS TREO), 2018.
[20] S. Kwon (Ed.), Gerontechnology: Research, Practice, and Principles in the Field of Technology and Aging, Springer Publishing Company, 2016.

[21] J. Lau, B. Zimmerman, F. Schaub, Alexa, are you listening?: privacy perceptions, concerns and privacy-seeking behaviors with smart speakers, Proc. ACM Hum.-Comput. Interact. (2018), https://doi.org/10.1145/3274371, November. Article 102. (Accessed 20 Nov 2020).
[22] S. Lenzi, P. Ciuccarelli, Intentionality and design in the data sonification of social issues, Big Data Soc. 7 (2) (2020) 1–8.
[23] X. Li, H. Lim, H. Hwang, Daily life with Clova smart speaker: topic modeling of Clova review data, in: Proceedings of the XX International Conference on Human Computer Interaction, 2019, June, pp. 1–2.
[24] Y. Liao, J. Vitak, P. Kumar, M. Zimmer, K. Kritikos, Understanding the role of privacy and trust in intelligent personal assistant adoption, in: Information in Contemporary Society, in: Lecture Notes in Computer Science, 2019, pp. 102–113.
[25] C.H. Lim, K.J. Kim, Information service blueprint: a service blueprinting framework for information-intensive services, Serv. Sci. 6 (4) (2014) 296–312.
[26] I. Lopatovska, K. Rink, I. Knight, K. Raines, K. Cosenza, H. Williams, P. Sorsche, D. Hirsch, Q. Li, A. Martinez, Talk to me: exploring user interactions with the Amazon Alexa, J. Librariansh. Inf. Sci. (2018) 0961000618759414.
[27] E. Luger, A. Sellen, “Like having a really bad PA”: the gulf between user expectation and experience of conversational agents, in: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, 2016, pp. 5286–5297.
[28] K.U. Mayer, New directions in life course research, Annu. Rev. Sociol. 35 (2009) 413–433.
[29] D. McColl, L. Wyg, G. Nejat, Brian 2.1: a socially assistive robot for the elderly and cognitively impaired, IEEE Robot. Autom. 20 (1) (2013) 74–83.
[30] S. Ntalampiras, I. Potamitis, Transfer learning for improved audio-based human activity recognition, Biosensors 8 (3) (2018) 60.
[31] S. Ong, A. Suplizio, Unpacking the breakout success of the Amazon Echo. Experian.com, Available at: https://www.experian.com/innovation/thought-leadership/amazon-echo-consumer-survey.jsp, 2016. (Accessed 27 February 2017).
[32] M. Porcheron, Joel E. Fischer, Moira McGregor, Barry Brown, Ewa Luger, Heloisa Candello, Kenton O'Hara, Talking with conversational agents in collaborative action, in: Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW'17 Companion, 2017, pp. 431–436.
[33] M. Porcheron, J.E. Fischer, S. Reeves, S. Sharples, Voice interfaces in everyday life, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, 2018, pp. 640:1–640:12.
[34] A. Pradhan, K. Mehta, L. Findlater, 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, 2018, pp. 459:1–459:13.
[35] A. Pradhan, L. Findlater, A. Lazar, Phantom friend or just a box with information: personification and ontological categorization of smart speaker-based voice assistants by older adults, in: Proceedings of the ACM on Human-Computer Interaction, 3(CSCW), 2019, pp. 1–21.
[36] A. Purington, J.G. Taft, S. Sannon, N.N. Bazarova, S.H. Taylor, Alexa is my new BFF": social roles, user satisfaction, and personification of the Amazon Echo, in: Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA'17, 2017, pp. 2853–2859.
[37] S. Purao, A. Wu, Towards values-inspired design: the case of citizen-centric services, in: International Conference on Information Systems (ICIS), Milan, Italy, December 2013, 2013.
[38] A. Pyae, T.N. Joelsson, Investigating the usability and user experiences of voice user interface: a case of Google home smart speaker, in: Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, 2018, September, pp. 127–131.
[39] S. Robinson, J. Pearson, S. Ahire, R. Ahirwar, B. Bhikne, N. Maravi, M. Jones, Revisiting “hole in the wall” computing: private smart speakers and public slum settings, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI'18), 2018, pp. 498:1–498:11.
[40] A. Reis, D. Paulino, H. Paredes, J. Barroso, Using intelligent personal assistants to strengthen the elderlies' social bonds, in: International Conference on Universal Access in Human-Computer Interaction, Springer, Cham, 2017, July, pp. 593–602.
[41] M. Royo-Vela, U. Hünermund, Effects of inbound marketing communications on HEIs' brand equity: the mediating role of the student's decision-making process. an exploratory research, J. Mark. High. Education 26 (2) (2016) 143–167.
[42] N. Sawe, C. Chafe, J. Treviño, Using data sonification to overcome science literacy, numeracy, and visualization barriers in science communication, Frontiers Commun. 5 (2020) 46.
[43] A. Sciuto, A. Saini, J. Forlizzi, J.I. Hong, Hey Alexa, what's up?: a mixed-methods studies of in-home conversational agent usage, in: Proceedings of the 2018 Designing Interactive Systems Conference, 2018, June, pp. 857–868.
[44] Y. Shiomi, S. Sakai, H. Terasawa, Intuitive representation of traffic flow dynamics: application of data sonification, in: 2020 Forum on Integrated and Sustainable Transportation Systems, IEEE, 2020, November, pp. 184–191.
[45] J. Shulevitz, Alexa, should we trust you? The Atlantic, Available at https://www.theatlantic.com/magazine/archive/2018/11/alexa-how-will-you-change-us/570844/, Nov. 2018. (Accessed 20 January 2021).
[46] R. Stephen, Using conceptual metaphors to represent temporal context in time series data sonification, Interact. Comput. 31 (6) (November 2020) 555–576.
[47] K.E. Skouby, A. Kivimäki, L. Haukiputo, P. Lynggaard, I.M. Windekilde, Smart cities and the ageing population, in: The 32nd Meeting of WWRF, May. 2014.
[48] R. Stephen, Using conceptual metaphors to represent temporal context in time series data sonification, Interact. Comput. 31 (6) (November 2020) 555–576.
[49] Kenneth Train, Wesley W. Wilson, Estimation on stated-preference experiments constructed from revealed-preference choices, Transp. Res., Part B 42 (2008) (2007) 191–203.
[50] Kenneth Train, Discrete Choice Methods with Simulation, 2nd edition, 2009, pp. 152–156.
[51] World Bank Group, Population ages 65 and above 2018. Retrieved from: https://data.worldbank.org/indicator/SP.POP.65UP.TO.ZS, 2018. (Accessed 15 January 2021).