

TOWARD SUPPORTING END-USER DESIGN OF SOUNDSCAPE SONIFICATIONS

KatieAnna E. Wolf

MeasuringU
Denver, CO, USA
katieanna.wolf@gmail.com

Rebecca Fiebrink

Department of Computing
Goldsmiths University of London
London, UK
r.fiebrink@gold.ac.uk

ABSTRACT

In this paper, we explore the potential for everyday Twitter users to design and use soundscape sonifications as an alternative, “calm” modality for staying informed of Twitter activity. We first present the results of a survey assessing how 100 Twitter users currently use and change audio notifications. We then present a study in which 9 frequent Twitter users employed two user interfaces—with varying degrees of automation—to design, customize, and use soundscape sonifications of Twitter data. This work suggests that soundscapes have great potential for creating a *calm technology* for maintaining awareness of Twitter data, and that soundscapes can be useful in helping people without prior experience in sound design think about sound in sophisticated ways and engage meaningfully in sonification design.

1. INTRODUCTION

1.1. Involving End Users in Sonification Design

The design of effective sonifications for a particular type of data and task can be challenging. Many approaches to sonification (e.g., parameter mapping sonification) present seemingly endless possibilities for ways sound may be manipulated in response to characteristics of the data [1]. A large body of work revolves around developing patterns and theories around representing data with sound [2, 3, 4, 5], and sound designers familiar with this work can benefit from such guidance in designing sonifications for end users of sonification systems. However, third-party designers often lack end users’ domain knowledge and understanding of the intended use of a sonification. Further, individuals may differ in their interpretation of audio representations of data [6].

For such reasons, user-centred design strategies have become more common in auditory display research [7]. As suggested by Walker and Nees, an effective sonification requires an understanding of the listener’s function and goals [5]. In practice, work by Verona and Peres shows that using a “task-based” approach—in which sonifications are designed based on the listener’s task rather than based on characteristics of the data alone—was found to increase listeners’ accuracy of working with sonifications [8].

However, including end users in the sonification design process is not an easy task. Collaboration between end users and expert sound designers can be labor intensive, and can involve con-

flicting priorities (as was found between dancers and designers in the design of a dancer movement sonification system [9]). Supporting independent end-user design is likewise challenging: for instance, non-experts may struggle to interpret specialist terminology used in sonification guidelines or design tools (e.g., terms like *frequency* and *timbre*).

We hypothesize that using real-world sounds that everyday people are already familiar with, like sounds in environmental soundscapes (sounds of the weather, animal vocalizations, and other natural sounds) might allow end users of sonification systems to design or refine sonifications while relying on terms they already know (like *bird tweet* or *running water*). Natural soundscapes have additional benefits: they are easily distinguished from background sounds, while still being able to fade out of attention without being tiring or obtrusive. Mauney and Walker found that users listening to such sonifications found the natural sounds to be “relaxing” [10]. Vickers et al. suggest that soundscapes can be “effective communication channels at the same time as being environmentally compatible and less fatiguing” [11].

1.2. Twitter Sonification

Twitter is a micro-blogging social media platform that allows users to broadcast short messages (“tweets”) to the world for anyone to view. Twitter can be thought of as a data monitoring platform—each user chooses specific other users who will appear on their “timeline” of recent tweets. While the choice of data itself can be highly customized, the presentation modality of that data can not. In addition to viewing the timeline on a Desktop or mobile device, Twitter offers sound and visual alerts to notify users of certain events of interest (e.g., a new tweet addressed to the user). Sound alerts can be muted, and the choice of alert sound can be changed, but this is the extent of sound customizability.

When enabled, auditory social media notifications can be obtrusive and—unless they drive the user to an app to view the triggering event—minimally informative. We hypothesize that using ambient soundscapes for Twitter data representation could be used as a form of *calm technology* that engages both the center and periphery of our attention, and is able to move back and forth between them [12]. That is, users could draw their attention to the soundscape and the data it represents when they wish to do so, and otherwise let the soundscape float at the edge of their periphery and maintain passive awareness. Yet it is impractical to pair everyday Twitter users with professional sound designers to create such soundscapes, and unworkable—given the highly individual information characteristics of each user’s feed—to create a one-size-fits-all sonification suitable for all (or most) Twitter users. There



This work is licensed under Creative Commons Attribution Non Commercial 4.0 International License. The full terms of the License are available at <http://creativecommons.org/licenses/by-nc/4.0>




Figure 1: The percent of participants who selected various reasons for changing their Twitter audio notifications on their computer.

fore, the design of appropriate ambient sonifications for Twitter users must leverage automation and/or interaction by the individual users themselves.

Some past work has used sound to represent Twitter data, mostly focusing on aesthetic presentations and performances. For instance in both *Tweetscapes* [13] and *I Hear NY4D* [14], real-time sonifications are created that utilize the content and geo-location of Twitter messages. Similarly the *Listening Machine*¹ presented a live sonification of 500 Twitter users around the UK. However, none of these projects provide users with control over the choice of Twitter that is sonified, or over the selection or design of sounds.

In the next section of the paper, we describe results of a survey that reveals how 100 Twitter users currently use audio notifications and what their primary objective is for using Twitter. Then, we describe a study of nine people engaging with a new tool for end-user design and customization of soundscape sonifications for Twitter data of interest. In this study, we explore how participants felt about and used soundscape sounds for representing their Twitter data. We discuss how properties of sound, user intention, and personal associations impact users' experience of soundscapes. We also discuss users' rationale for sonification design decisions. This work contributes to a better understanding of the utility of soundscapes for creating ambient, personal data displays, as well as a better understanding of how to support end users in designing bespoke soundscape sonifications.

2. SURVEY OF TWITTER USERS

We conducted a survey of active Twitter users to better understand the type of information people seek when they check Twitter, and how and why people use and customize Twitter audio notifications. We posted the approximately five-minute survey on Amazon Mechanical Turk and asked for active Twitter users to "Answer questions about your use of Twitter". We paid each participant \$0.50.

2.1. Survey Results

We collected responses from 100 self-described active Twitter users. 53 were female, 47 were male, and their ages ranged from

¹<http://www.thelisteningmachine.org>

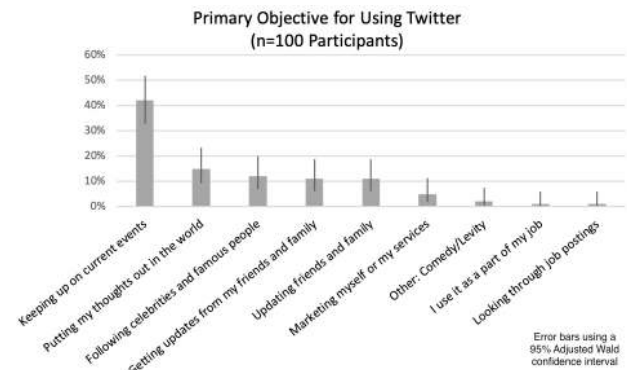


Figure 2: This figure shows the percentage of participants who selected the different primary objective for using Twitter.

19 to 78 years old (mean = 35.51, $\sigma = 12.05$). We collected participants' self-reported information on: the number of years they had been active on Twitter (mean = 4.67, $\sigma = 2.22$), the average amount of time they spent on Twitter per day (mean = 68.87 minutes, $\sigma = 61.22$), the average number of times they accessed Twitter per day (mean = 8.29, $\sigma = 9.01$), the number of accounts they followed (mean = 426.81, $\sigma = 714.50$), and the number of Twitter followers they had (mean = 456.22, $\sigma = 829.40$).

Of the 100 participants, 64 used their desktop/laptop at least some of the time to access Twitter. Of these 64, 41 never enabled audio notifications and 23 used them at least some of the time. When asked why they didn't enable audio notifications (a multiple-choice questions in which participants were asked to select all options that applied), over half (24 participants) indicated that they didn't want to be interrupted, while just under half (18 participants) indicated that they used visual notifications. Other responses included "I don't want to disturb others" (8 responses) and "The information is not important" (6 responses).

We asked participants who enabled Twitter audio notifications to tell us reasons why they had ever changed the built-in sound notifications on their device (another multiple-choice question asking participants to check all options that applied). Results for the 23 laptop/desktop participants appear in Figure 1. The three main reasons for changing the audio notifications were: (1) to be quieter, (2) to be more relaxing, and (3) to be less annoying. Each were selected by eight respondents (34.8%).

We also asked participants to specify their primary objective for using Twitter, by selecting an option from a pre-compiled dropdown list or by writing in their own. The results appear in Figure 2 and show "Keeping up on current events" was the most common objective selected by 42 out of 100 participants (the next most selected was "Putting my thoughts out in the world" with 15). Finally, we asked participants to provide an estimate of the amount of time they spent on various activities (writing tweets, responding to other tweets, reading tweets on their timeline, reading tweets on trending topics, and reading tweets on the timelines of other accounts). While answers varied, each activity drew a response from at least 74 of the 100 participants. "Reading tweets on my home timeline" was the most popular activity overall; 87 participants reported that they did this at least some percentage of the time, and 33 participants spent at least 50% of their time on this activity.

The survey findings further support our intuition that using

soundscapes to represent Twitter data may be useful, as soundscapes have been found to be relaxing [10] with the potential to be less fatiguing than other sound interfaces [11]. A data presentation modality with the capacity to be less disruptive and slightly more informative than existing audio notifications also seems attractive. These findings have also informed the design and participant selection of our subsequent study on end-user Twitter sonification design, described in the next section.

3. SOUNDSCAPE SONIFICATION DESIGN STUDY

We next conducted a study to determine how Twitter users feel about using soundscapes to represent Twitter data of interest to them, what may be the benefits and challenges of using soundscapes to represent this data, and how people reason about and design with soundscapes when given the ability to design their own sonifications.

3.1. Summary of ESCaper Application

To support the study, we designed a GUI-based web application called ESCaper (Environmental Soundscape Creator). ESCaper allows Twitter users to identify specific Twitter data (e.g., accounts or hashtags of interest) that they are interested in monitoring with sound, and to create a sonification of that data using soundscapes. ESCaper uses the WebAudio API² to play and manipulate environmental sounds selectively chosen from freesound.org. ESCaper provides sound samples for two soundscapes: a Forest (crickets, stream, thunder, frogs, a wolf howling, a nightingale, a fly, and an owl) and a Beach (ambient people, sea wind, waves, flock of seagulls, single seagull, foghorn, single wave crash, ambient birds).

Users first sign in to ESCaper using their Twitter account, which authenticates our application and gives it read-only access to their Twitter data. The interface then asks users to select the data they wish to passively monitor using soundscape sounds. Specifically, users select four of the Twitter accounts they are following and additionally specify two keywords or hashtags of interest.

ESCaper uses an automatic mapping technique from previous research [15] that restricts how data can be mapped to soundscape sounds. Specifically, ESCaper allows each of the four selected Twitter accounts to be mapped to one short-duration (“instant”) sound, such as a wolf howl or seagull call; the playback of a given sound will indicate that the corresponding account has just tweeted. Additionally, ESCaper allows each keyword/hashtag to be mapped to a longer duration (“continuous”) sound (such as continuous crickets or ocean waves); one specified aspect of the sound playback (speed, gain, or left-right panning) will then represent changes in the number of tweets each second that contain that keyword/hashtag.

ESCaper provides two user interfaces for specifying these mappings from Twitter accounts and keywords to the selection of corresponding sound samples. Interface 1 requires users to manually select a sound to correspond to each selected account or keyword. Interface 2 is nearly identical, with two differences: (1) after the user specifies the accounts and keywords of interest, the interface automatically populates itself with initial choices of sounds for each one; and (2) it includes a ‘Randomize Sonification’ button that, when pressed, randomly assigns a unique sound to each of the Twitter accounts or keywords not currently mapped to a sound. See Figure 3 in the Appendix for an image of Interface 2.

²https://developer.mozilla.org/en-US/docs/Web/API/Web_Audio_API

Once a mapping has been fully or partially specified, users can listen to it by running it on “new” Twitter data and listening to the sound generated from the data. For the purposes of this study, however, we generated the sonification using a pre-recorded Twitter data stream that was identical for each participant (with accounts and keywords in this historical data re-mapped to the each participant’s chosen accounts). This allowed us to keep properties of the data generating the soundscapes (e.g., frequencies of tweets) consistent across all participants, no matter which accounts and keywords participants chose.

3.2. Participants

As ESCaper was designed as a computer application, we used the survey in Section 2 to identify active Twitter users who used Twitter on their computer at least 30% of the time and also met at least two of the following three criteria: (1) They spent at least 50% of their time using Twitter on a computer; (2) Their primary objective on Twitter was to receive information rather than post information (i.e., they selected either: “Keeping up on current events”, “Getting updates from friends and family” or “Following celebrities and famous people”); (3) They spent at least 50% of their time on Twitter reading tweets rather than writing them.

We contacted 13 of our survey participants who met this criteria as well as an additional 14 people who took a separate survey we posted with the same screening questions. Of these 27 people, nine consented to participate in our study. Seven were female, two were male, and their ages ranged from 26 to 56 (mean 39). Two participants had some past experience with sound design: one had five years of experience working with oscillators and synthesizers for tone layering and sound mixing, while another had 1.5 years of experience editing audio clips with the Audacity sound tool. Participants’ musical abilities ranged from 0 to 35 years of experience.

3.3. Study Procedure

A pre-study survey asked participants which Twitter account they would be using for our study and to list seven Twitter accounts that they were most likely to check on an average day, as well as five hashtags or keywords that they were most likely to search for on an average day or that they had searched for in the past month. We used this information to enable ESCaper to pre-populate its drop-down lists for selecting accounts.

For the study itself, we video chatted with participants for one hour and had them share their screens using the *appear.in* communication tool. We recorded the screen and audio of each session (with consent from the participant), so we could reference the transcript afterwards. After a brief introduction of the facilitator and the research, the facilitator read the task scenario below:

“Imagine that you are on your computer doing your normal tasks, such as checking your email, reading online articles, online shopping, paying bills, etc. As you are focused on these tasks, you also want to be able to passively monitor specific Twitter information. Your goal with this user study is to design an audio representation of your Twitter data using environmental soundscapes (animal vocalizations, sounds of the weather, etc.) that will allow you to stay informed about the data on Twitter while your main focus is on another task. Please talk out loud and describe your thought process as you interact with the interfaces, with a specific focus on your design process.”

The study was designed with both soundscape (Forest, Beach) and interface (Interface 1, 2) as within-subjects variables. Each

participant used Interface 1 once and Interface 2 once, with a different soundscape for each; we randomly assigned each participant an order of the two interfaces and an order for the soundscapes. This enabled each participant to explicitly compare the soundscapes and interfaces.

After the scenario text was read, the participant was told to use the first ESCaper interface until they felt they had created a design that accomplished the task, after which they were encouraged to open a new tab on their web browser and to spend a minute doing a secondary task with the soundscape in the background. As our primary focus of the study was to observe how participants designed their sonifications with soundscapes (rather than on how accurately they were able to monitor the data) the single minute was for the user to assess and iterate on their initial design. After a minute had passed, the facilitator asked the participant if they thought their soundscape accomplished their goal or if they wanted to make changes. Once they did not want to make any changes, the facilitator then asked the participant the following about their experience with that interface in an unstructured interview format:

- Overall, how easy or difficult was it for you to complete the task? (1-7 Very difficult to Very Easy)
- Overall, how enjoyable was it for you to complete the task?(1-7 Not at all enjoyable, to Very Enjoyable)
- How satisfied are you with the sound of your final design? (1-7 Not at all satisfied to extremely satisfied)
- How confident are you that you would be able to use this sonification for passively monitoring your Twitter data? (1-7 Not at all confident to extremely confident)
- What were you focused on most while you were designing your sonification?
- What was the most challenging part of creating a design using this interface?

Then, the participant was asked to repeat the same task with the second interface assigned to them. The same questions above were asked of the participant upon completing the task with the second interface.

Finally, at the end of the study the facilitator asked a last set of questions comparing the two interfaces and soundscapes:

- How was your experience different between the two interfaces? Which of the two interfaces did you prefer? Why did you prefer that interface?
- What did you like best about the soundscape sounds? What did you like least about the soundscape sounds? Which soundscape did you like best? Why?
- (In Interface 2) How helpful was the starting sonification to your design process? How helpful was the randomization to your design process?
- How would your use of the sounds change if you were going to be listening to them over a long period of time?
- In what contexts could you imagine using this interface? How often could you imagine using this interface? Are there other phenomena on Twitter that you would want to use sound to represent?
- Did you feel like you were able to interpret the Twitter information from the soundscapes? What made that easy or difficult to do?

4. SONIFICATION STUDY RESULTS

4.1. Interface and Soundscape Ratings and Preferences

Participants' ratings of difficulty, enjoyment, satisfaction, and confidence for each task are presented in Table 1. For each ESCaper interface, soundscape, and order of presentation we ran a paired-samples t-test for each rating question and found that there was no statistically significant difference between the ratings for the two soundscapes, between the ratings for the two interfaces, or between the ratings for the order in which the tasks were presented.

Table 1 shows each participant's preferred ESCaper interface and soundscape. Six participants preferred Interface 1 and three preferred Interface 2, while six people preferred the Beach soundscape and three preferred the Forest. Five of the six participants who preferred Interface 1 also preferred the Beach soundscape. Running the Fisher exact test for the interface and soundscape preferences we did not find a statistically significant difference.

4.2. Factors Influencing People's Experiences of Using Soundscape Sounds to Represent Twitter Data

Through the interview questions as well as through observations of participants thinking aloud while interacting with ESCaper, we learned about many of the factors that they identified as influencing their experience of hearing Twitter data represented as soundscape sounds. In this section, we present common themes that arose.

4.2.1. Relationships Between Interface Sounds and Real-World Sounds

Several participants mentioned how the sounds in the real world, might interfere with their ability to detect the sounds in the ESCaper interface, or how the sounds in the interface may lead them to react as if the sounds were occurring in the real world.

For instance, when discussing what would cause them to change the sounds in the sonification, one participant stated: *"If there was rain sounds [in the interface] and it was raining outside, then I might switch it because then I might be distracted 'Is that the rain on my computer? Is that the rain outside? Or maybe I am just sick of listening to rain', So then I might switch it to something else"*. While determining which sounds to use in their design, another participant stated: *"This thunder will have to compete with my outside thunder, my real-life thunder"*. Additionally, two participants specifically reflected that they might react to the fly sound in the interface as if it were actually there in the real world: *"Not too long ago we had a problem with flies in here. So every time I hear a fly I instinctively duck my head..."* and *"I live in Florida and yesterday there was fly in the house. To me it may not be as ambient, it might sound like an actual fly"*.

4.2.2. Personal Associations with Sounds

In addition to participants confusing real-world sounds with the interface sounds (as with the fly samples above), participants also discussed how their personal associations with the sounds played a role in their selection process. Again the fly sample was one that drew a lot of personal associations: *"I hate bugs"*, and *"The fly reminds me too much of my past work... I used to do research with flies. It would just probably make me feel like I am at work"*. Similarly, when selecting forest sounds, one participant ruled out

Participant No.	Task No.	Task Interface	Task Soundscape	Difficulty Rating	Enjoyment Rating	Satisfaction Rating	Confidence Rating
1	1	Interface 1	Beach*	7	7	6	6
1	2	Interface 2*	Forest	7	7	5	6
2	1	Interface 2*	Forest*	7	4	5	6
2	2	Interface 1	Beach	5	4	3	5
3	1	Interface 1*	Forest	4	5	6	7
3	2	Interface 2	Beach*	6	5	4	7
4	1	Interface 2	Beach*	7	7	7	6
4	2	Interface 1*	Forest	7	7	7	7
5	1	Interface 2	Forest	6	6	7	7
5	2	Interface 1*	Beach*	6	7	7	7
6	1	Interface 1*	Beach*	7	6	6	7
6	2	Interface 2	Forest	7	6	5	5
7	1	Interface 1*	Forest	7	7	7	7
7	2	Interface 2	Beach*	7	7	6	7
8	1	Interface 2*	Beach	6	7	6	7
8	2	Interface 1	Forest*	7	7	6	7
9	1	Interface 2	Beach	5	7	6	6
9	2	Interface 1*	Forest*	7	7	7	7

Table 1: Each participant completed two tasks (one for each interface and soundscape). In this table, we present the interface and soundscape they preferred (marked with an ‘*’) and their ratings for each task on: how difficult it was to complete the task (Difficulty), how enjoyable it was to complete the task (Enjoyment), how satisfied they were with the sound of their final design (Satisfaction), and how confident they were that they would be able to use the sonification for passively monitoring their Twitter data (Confidence).

the crickets sample, stating: “I definitely don’t like the crickets. I have timmitus and it reminds me of the cicadas in my ears”.

Some participant’s personal associations went beyond the sample level as participants mentioned their familiarity with the soundscape as a whole. When asked which soundscape they preferred, one participant responded: “I mean, I like the beach sounds, people love beach sounds, but I grew up in the country, so I am more used to sounds like [the forest]. It was comforting, it made me think of home, and I like that it was running in the background, that is was something that I could do while I am working, and it would keep me calm and keep me present.” When asking another participant what they liked best about the soundscape sounds, she mentioned the familiarity of the Beach soundscape as it sounded like her home on a Saturday morning. She even specifically ruled out the birds sound because it didn’t sound “beachy” enough for her and “didn’t feel as familiar”. She also thought that the Forest soundscape would work well on a day where she wanted to be alerted, since the sounds are not as familiar to her and would not fade out of attention as easily as the Beach sounds.

4.2.3. Desire to Alert or to Passively Monitor

When asked which soundscape participants preferred, three preferred the Forest soundscape, while six preferred the Beach. Those who preferred the Beach soundscape often stated that it would be better for passive monitoring as the Forest was more alerting:

- “[The beach] is more ambient sound, it’s more soothing. It is nice to have in the background because it is more mellow. I feel like with the forest sounds, the animals chirping and wolves howling, it was a bit loud and more distracting, so when I am doing a secondary task I prefer the more mellow in the background.”
- “For passive monitoring I would much prefer the beach. It is much more passive to me. The forest is, as the forest is, it is

alive and active and wants your attention.”

- “Beach soundscape was soothing. The forest definitely alerts you.”
- “I am a beach person. I wish there was more variety in the soundscape, but the sounds were pleasant. The forest was more annoying with the animals and the birds, [whereas] these at the beach, they were more peaceful.”
- “The [beach] was less audibly distracting when I was reading, because the sounds were softer... The sounds were overall more pleasant. The sounds were more distinct in the forest, less ambient.”

However, similar to the woman in the previous section who would use the forest soundscape on a day when she wanted to be alerted, other participants also mentioned their interest in being able to switch to the forest interface: “For passive monitoring I would definitely use the beach, but if was in a mode where I wanted to be alerted I would use [the forest] sounds”. In fact, all participants who preferred the Forest soundscape did so because the sounds were more easily distinguished: “There was more diversity in the [forest] sounds. I feel like I could actually use them more. They were distinct enough I could actually tell what they were trying to do”, “The sounds were more distinct, more discrete, the mix was good” and “I was able to hear all the sounds”.

4.2.4. Properties of Soundscape Sounds

Eight out of the nine participants stated (while they were thinking out loud or when they were asked about the soundscape sounds) that at least some aspect of the soundscapes were *soothing, peaceful, or relaxing*. In fact from the quotations above, we can see that one of the main reasons participants felt that the Beach soundscape was better for passively monitoring data was because of its relaxing and soothing properties. Two participants specifically men-

tioned using the Beach soundscape to for meditation/relaxation: “It’s almost like I can meditate to this, relax” and “These sounds are almost soothing. They are almost like a relaxation tape”.

Participants also discussed how the soundscape sounds compared to the sounds of other interfaces. For instance, one participant reflected on the notification sounds used in other applications at her work. She noted that the soundscapes didn’t sound like the other devices, which were more “mechanical”. She even stated that she had asked someone that day to turn down their notifications because of their sounds. Three participants specifically contrasted the soundscape sounds with other application sounds that were “jarring” and “jolting”:

- “The kind of standard g-chat or whatever “ding” is a very attention grabbing sound, but you can’t really diversify that too much ‘cause if you are just doing higher “ding” versus a lower “ding”, that is just not going to work. The natural sounds you have more diversity in it and they are not super alarming, it is not like a big alarm that is going off in your head. For me, the nature sounds, they don’t induce any anxiety whereas some sort of buzzer would jolt my mind.”
- “[The soundscape sounds] weren’t jarring, like alerts, like an alert sound like a bell or chime. It is just something that is going on in the background and if you want you can just tune it out. Pay attention to it when you want to.”
- The third participant stated that for the software she used at work (messaging and video-chatting software), the sounds were more distracting (“boings” and “ding dongs”), so she doesn’t use them because they are “jolting”. She thought that since the soundscapes were more natural that she would like them better and could see herself using them.

Only one participant expressed concern with the soundscape sounds: “Nature sounds seem generic. Not modern enough”.

4.3. Participants’ Motivations for Sound Selection

The main action in designing the sonification of the Twitter data in ESCaper is to assign particular sounds in the soundscape to either particular Twitter accounts or to Twitter hashtags/keywords. Below we present the common themes we observed as people thought out loud about why they were choosing particular sounds and as they responded after each task to the question “What were you focused on most while designing your sonifications?”

4.3.1. Associations Between Data and Sound

Five of the nine participants mentioned that certain aspects of the sound were reflective of certain personalities of the Twitter accounts they were selecting them to represent. For instance, one participant described that she selected the wolf sound for a particular Twitter account, as she thought of them as being the pack leader, and she wanted to make sure she could hear that sound as they were Tweeting updates in the upcoming week. Another participant kept the gull sound that was preselected because it “*Might be appropriate for him. He is very different in appearance with [tattoos] and stuff*”. The three other participants very heavily relied on personalities to assist them in making decisions about which sounds to use for the particular accounts:

- *What I was focused on most was really the personalities of the people and also the personalities of the hashtags, kind of like*

what I think of them. So like for Julian Assange, someone for WikiLeaks, the singing bird was kind of the perfect thing. For a sports reporter to be a seagull, and just be squawking all the time also fit perfectly for Jeff Howe. And Nina just alerting people, just bringing attention to stuff that people may ignore and may not realize is going on, I think is a great use of the foghorn. So I was really just matching personalities and kind of visualizing what these people sound like to me.

- *I was thinking of the nightingale for Joy Reid because although she is calm and discusses things calmly, when she gets excited about something she gets happy, and I think this is a very happy sound. So I am going to choose [the nightingale] for her. Wolf I am going to choose immediately for Stonekettle, because he is a veteran, ex-military guy, he is very masculine. So I think Wolf will be very good for him. Joe I think of “Fly in the Ointment”. Like the idea of getting under someone’s skin, or spoiling something, and since he’s a comedian and since he likes to heckle other users, or annoy the crap out of them, I feel like that would be a good representation for him. He doesn’t really annoy me, but I feel like that’s his person, that he likes to do that to other users. So I feel that will work for him.*
- *The wolf howls to get attention and that is what Donald Trump usually does when he is on Twitter or making his speeches. For Netflix...the owl. The reason why I am going to choose the owl for Netflix is that basically when you want to watch a movie or a show you can watch anytime, but with Netflix people are usually watching it over night hours. So with the owl usually it stays up all night, as the same as the Netflix we’ll be able to watch movies over night.*

However, two participants did mention that they had difficulty forming associations to soundscape sounds: “I don’t know if I would associate any animals with this account” and “The most challenging thing is with a lot of abstract sounds, making sure I remember what each one is.” These two participants both suggested that being able to add in their own sounds would help to create a more memorable association to the data. For instance one participant stated “I would be curious to try [uploading my own sounds]. Like for example the Korean [hashtag], I would probably upload a sound clip of either a song, or a snippet of a song, or someone speaking Korean, like that sort of thing. Something that is super duper customized.” Another participant stated: “Being able to add my own music or clips, not necessarily just nature sounds. Nature sounds seem generic. Not modern enough.”

A third participant was very interested in using song clips from television shows and movies to create notifications to Twitter accounts related to those shows and movies. Specifically, she mentioned having a song from the Star Wars films (the *Imperial March*) be associated with Twitter accounts related to the those films, and having the *Stranger Things* TV show theme song be associated with Netflix Twitter account, which is the streaming service that distributes that show. She believed that because these songs were so familiar to her, she would be able to tune them out and use them for passively monitoring the data.

4.3.2. Importance of Data

Three participants discussed selecting sounds that would be louder or more prominent in the soundscape to represent the Twitter accounts that they were more interested in detecting:

- *Basically, I was looking to make certain ones stand out. I wanted Raw Story and The Hill to stand out because they are actual news sources, so they would be most likely to be breaking news. The other two were comedy accounts that I just like to follow.*
- *I like Jaclynhill so I gave her a little louder [sound]. The ones I really care about I would give louder [sounds].*
- *Funko does a lot of giveaways if you retweet their stuff, so I want to make sure that I heard that one.*

4.3.3. Ability to Interpret Sounds

Five participants mentioned how the properties of the sounds themselves would affect their ability to detect/interpret the data through the sounds:

- *The first thing I was thinking of was just how distinct the sounds were. So I could tell which ones were which without any effort... If I can't detect differences in sound then it isn't going to do its job.*
- *For the first couple ones [I was focused] on the louder animals, because if I am going off and looking at something I would want something that would get my attention.*
- *Those are both birds, so I am going to change one of them. Some of these [sounds] are too similar to the others. With all of the different birds, all of the birds would get jumbled over each other. It would be hard to tell them apart.*
- *Waves is louder so I am going to pick that. I am mainly picking it because it is louder than the other options.*
- *I will choose thunder. It is more attention grabbing than crickets and stream.*

4.3.4. Sound Preference

Finally, most participants ruled out or selected particular sounds due to their preference of the sounds:

- *I would not want to listen to the fly. I hate bugs... I love wolves, so I wanted Funko and wolves to be associated together. Things that I like I wanted to hear more of, I know that it would call my attention to more.*
- *Secondary was how pleasant or unpleasant the sounds were, so that is why I eliminated the fly sound.*
- *I was avoiding water sounds. They were not as desirable.*
- *[I focused on] the sounds that I liked and how soothing they were.*
- *I like crickets on a summer night in a field, but this just doesn't feel comfortable.*
- *Oh no - not flies. To me [the fly] sound is annoying*

In some instances, participants were actually drawing on perceptual properties of the sounds themselves to describe why they preferred or did not prefer certain sounds:

- *I definitely like the thunder more [than the crickets], it is more of a lower rumble, instead of the crickets which are high pitched. I feel like something that is lower is better for me personally. If it is too high pitched, it is drawing all my attention. I can't even read a twitter post, because it is captivating too much attention.*

4.4. Creativity and Control in the Sonification Design Process

The main differences between the two ESCaper interfaces were that Interface 2: (1) started participants with each Twitter account/keyword/hashtag being pre-assigned a soundscape sound, and (2) contained a 'Randomize Sonification' button to randomly select sounds for data groups that were not yet selected by the participant. By asking each participant which interface they preferred and why, we were able to gain information about how useful these functionalities were to the participants and how these functionalities affected the participants' design process, with a specific focus on the amount of control participants had on the designs.

Two of the three participants who preferred Interface 2 explained that the automated sound selection helped relieve the mental burden and stress of having to make a decision: "*Sometimes if I am unsure which one to select, the randomization would choose for me and avoid the confusion and be less time consuming by picking it for me*" and "*The suggested layout was easier to use, because I didn't go through and click and decide if I want to listen to all of [the sounds]. Overall, it was a lot faster and just less mental energy going into making a decision. It felt very effortless. I preferred how easy and not much energy to pick what I want*".

However, all of the participants who chose Interface 1 explained that being able to make up their own mind and have the creative control was the reason they preferred Interface 1:

- *I would rather make up my own mind based on what I want to hear. I liked having more choices.*
- *[Interface 2] seemed to be a simpler interface, and I didn't actually like it as much because of that. I like to have more options. The ability to be creative. You get some satisfaction when you are listening back to it to know that you put some work into it.*
- *I enjoyed [interface 1] more, even though it was more work for me, I felt more in control and I felt like it was more personal.*
- *For someone like me who is picky, Interface 1 would be better.*
- *Seeing them already randomized, it put in my mind that Oh, I can't play around with this.'*
- *That is exactly what I personally as a user want to see: 'Here is everything, make your choices'.*

5. DISCUSSION AND CONCLUSIONS

In this paper we explored how people use and feel about using soundscape sounds for representing social media data, specifically for communicating information about the occurrence and density of tweets of interest.

One of the challenges that participants identified with using soundscape sounds was that they could be difficult to separate from sounds in the real world. Real-world sounds could be mistaken for data, while sounds within the interface could be mistaken for events in the real-world and could trigger reactions from people as if those sounds were really occurring. Additionally, some found that certain sounds were annoying, too alerting, or just difficult to listen to. However, our study showed that when users are given control over the choice of sounds, they can avoid using sounds that they do not prefer in their soundscapes. Also, our participants described that they envisioned taking advantage of the ability to dynamically adjust their sonifications to mitigate these challenges.

For instance, some participants described using the Forest soundscape when they wanted their sonification to be more alerting and using the Beach soundscape when they wanted to more passively monitor their data.

Participants noted that a benefit of soundscapes was the relaxing and soothing quality of the sounds, especially those in the Beach soundscape. Even with these relaxing qualities, some participants still appreciated the distinctness of the sounds, and in comparison to the sounds of other applications, participants found soundscape sounds as less “jolting” and “jarring”. Additionally, some participants described the soundscapes as being familiar to them (“made me think of home”), which made it easier for the sounds to fade into the background. As described by Mark Weiser and John Seely Brown “*The result of calm technology is to put us at home, in a familiar place. When our periphery is functioning well we are tuned into what is happening around us, and so also to what is going to happen, and what has just happened*” [12]. All of our participants felt that they would use soundscapes as a way to passively monitor their Twitter data in some form or another. In particular, participants often described the soundscapes as being something they could listen to in the background while doing other tasks (working, writing, reading articles, etc.), yet still be able to draw their attention when they wanted. Soundscapes clearly have potential to turn Twitter into a calm technology that is less intrusive and more passive. One participant described this perfectly: “[The soundscape] is just something that is going on in the background and if you want you can just tune it out. Pay attention to it when you want to.” While a longitudinal study would be necessary to explore the use of soundscapes for passive monitoring of Twitter data over long periods of time, clearly it is worth exploring in the future.

Our study also demonstrated the feasibility of enabling end users to personalize soundscape sonifications representing their data, using a simple GUI interface. This interface gave each user the ability to individually decide what mattered most to them in the design process. We saw some participants prefer to make fewer decisions with the use of Interface 2, while others were interested in the creative process and specifically choosing sounds for particular aspects of the data. Even with only nine participants, we saw that there were several methods participants used to make design decisions including: (1) associating specific personalities or traits of their Twitter data with the sounds, (2) presenting the data they were most interested in more clearly than other data, (3) using sounds that were easiest for each of the participants to individually interpret, and (4) using the sounds that they preferred while avoiding sounds they disliked. Clearly, our participants were able to think about sound design in a sophisticated way, even with little or no experience in sound or sonification design.

6. ACKNOWLEDGMENT

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship under grant number 1148900. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Dr. Wolf is also grateful for the support of the Microsoft Research Womens Fellowship and the Gordon Y.S. Wu Fellowship from Princeton University’s School of Engineering and Applied Science in enabling this work.

7. REFERENCES

- [1] F. Grond and J. Berger, *The Sonification Handbook*. Logos Publishing House, 2011, ch. Parameter Mapping Sonification, pp. 363–397.
- [2] A. de Campo, “A data sonification design space map,” in *Proceedings of the International Workshop on Interactive Sonification (ISon)*, 2007.
- [3] S. Barrass, “Sonification design patterns,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2003.
- [4] J. Anderson, “Creating an empirical framework for sonification design,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2005.
- [5] B. Walker and M. Nees, *The Sonification Handbook*. Logos Publishing House, 2011, ch. Theory of Sonification, pp. 9–39.
- [6] B. N. Walker and B. S. Mauney, “Individual differences, cognitive abilities, and the interpretation of auditory graphs,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2004.
- [7] S. D. H. Cornejo, “Towards ecological and embodied design of auditory display,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2018.
- [8] D. Verona and S. C. Peres, “A comparison between the efficacy of task-based vs. data-based semg sonification designs,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2017.
- [9] S. Landry and M. Jeon, “Participatory design research methodologies: A case study in dancer sonification,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2017.
- [10] B. S. Mauney and B. N. Walker, “Creating functional and livable soundscapes for peripheral monitoring of dynamic data,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2004.
- [11] P. Vickers, C. Laing, M. Debashi, and T. Fairfax, “Sonification aesthetics and listening for network situational awareness,” in *Proceedings of the Conference on Sonification of Health and Environmental Data*, 2014.
- [12] M. Weiser and J. S. Brown, “Designing calm technology,” *PowerGrid Journal*, vol. 1.1, pp. 75–85, 1996.
- [13] T. Hermann, A. V. Nehls, F. Eitel, T. Barri, and M. Gammel, “Tweetscapes: Real-time sonification of twitter data streams for radio broadcasting,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2012.
- [14] B. Boren, M. Musick, J. Grossman, and A. Roginska, “I hear NY4D: Hybrid acoustic and augmented auditory display for urban soundscapes,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2014.
- [15] K. E. Wolf, G. Gliner, and R. Fiebrink, “End-user development of sonification using soundscapes,” in *Proceedings of the International Conference on Auditory Display (ICAD)*, 2015.

8. APPENDIX

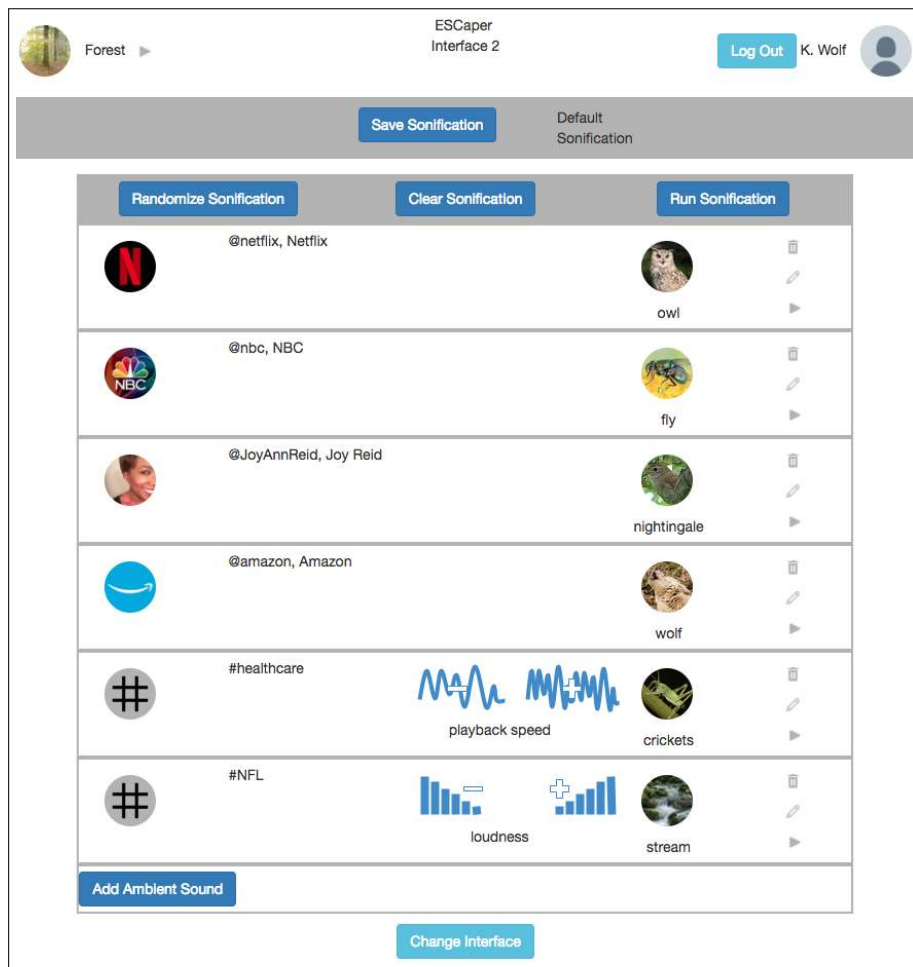


Figure 3: This image shows Interface 2 of the ESCaper application using the Forest soundscape sounds. In this sonification design the Twitter accounts for Netflix, NBS, Joy Reid, and Amazon are represented by the owl, fly, nightingale, and wolf, respectively. The change in the number of tweets including the hashtag ‘healthcare’ are represented by the change in playback speed of the crickets, while the change of the number of tweets including the hashtag ‘NFL’ are represented by the change in loudness of the stream.