Draft of Jamieson, J., & Boase, J. (2016). Listening to social rhythms: Exploring interactional, time stamped data using sonification. In L. Sloan and A. Quan-Haase (Eds.), *The SAGE handbook of social media research methods.* Thousand Oaks, CA: SAGE Publications.
Published version available at: https://us.sagepub.com/en-us/nam/the-sage-handbook-of-social-media-research-methods/book245370

Listening to social rhythms: Exploring logged interactional data through sonification

Jack Jamieson and Jeffrey Boase, University of Toronto

Abstract

The popularity of social media has given rise to a vast number of time-stamped logs of tweets, blog posts, text messages, status updates, comments, shares and other communications. These data sets can be explored to identify new types of interactional patterns and trends. Data sonification – converting data into sound – is particularly well-suited to exploring temporal patterns within time-stamped log data because sound itself is inherently temporal and the human auditory system has excellent temporal resolution. This chapter presents examples of sonifications of social media data, discusses considerations for performing sonification-based analyses, and describes a study in which sonification was used to explore temporal patterns in mobile text message log data. The intent is to allow readers who are unfamiliar with sonification to understand its capabilities and limitations, as well as how they may apply sonification in their own research.

Keywords

Sonification, exploratory analysis, social media, Twitter, asynchronous communication, response time, text messages.

Introduction

Social media interactions consist of a broad variety of activities, such as posting, retweeting, sharing, commenting and replying to status updates, personal messages, news articles and other forms of user engagement. Often, these interactions occur asynchronously, allowing participants to choose when they initiate, respond to, pause, ignore, and conclude interactions. As a result, the time at which events occur can be a non-verbal cue of eagerness, engagement, thoughtfulness, or other qualities (Döring & Pöschl, 2009; Kalman & Rafaeli, 2011; Quan-Haase & Collins, 2008; Walther & Tidwell, 1995). Moreover, asynchronous digital interactions such as tweets, Facebook updates, or mobile text messages are intrinsically temporal, and logs of this data are almost always time-stamped. By analyzing time-stamped interactional data, researchers can develop insights into a variety of topics, such as how interactions unfold over time, inequalities in the exchange of information, and personal network change over time.

The abundance of available logged data about social media interactions has created opportunities for new lines of inquiry and new styles of research. As a result, researchers can use this wealth of data to ask questions that are qualitatively different and develop novel analytic methods to address those questions. In this chapter we discuss the potential of data sonification – converting data into sound – for exploratory data analysis of time-stamped interactional data, such as that found in social media logs. Exploratory data analysis was popularized by Tukey (1977) as a set of methods for exploring statistical data. In contrast to other sorts of statistical analysis, exploratory analysis does not address specific hypotheses, but rather is used to identify patterns, relationships, and trends. Andrienko and Andrienko summarize that exploratory analysis 'is about hypothesis generation rather than hypothesis testing' (2006, p. 3). This makes exploratory analysis a suitable approach for discovering patterns and trends in the interactional data generated through social media and other forms of asynchronous digital communication. Exploratory analysis has been used by many researchers to study social media use, but most of this exploration has relied on visualization. Sound has particular potential for analyzing temporal patterns due to the inherently temporal nature of sound (Neuhoff, 2011), and sonification can offer a different and valuable perspective.

This chapter begins by explaining what sonification is and why it can be useful for analyzing interactional data. We then provide examples of existing social media sonifications to illustrate the current state of affairs. This is followed by a discussion of theoretical and methodological factors to consider when undertaking sonification-based research. Then we present a detailed description of a sonification-based study we conducted of mobile text message activity. In describing this study, we illustrate how the considerations identified in the previous section can affect the process of conducting research with sonification. This chapter concludes with a discussion of future directions for the use of sonification to explore social media and other interactional logs.

Why use sonification?

Data sonification is a method of converting data into sound. This allows researchers to listen to patterns, values, and relationships within data in much the same way that data visualization allows researchers to see them. The commonly accepted definition is that 'sonification is the use of non-speech audio to convey information' (Kramer et al., as cited in Hermann, 2008, p. 1). Since data visualization is a more common approach than sonification, it may be useful to consider sonification as a relative of visualization. Tufte's seminal work on data visualization, *The Visual Display of Quantitative Information*, begins by asserting, 'Data graphics visually display measured quantities by means of the combined use of points, lines, a coordinate system, numbers, symbols, words, shading, and colour' (2001, p. 10). Like visualizations, sonifications perform the function of conveying information, but do so using an auditory rather than visual set of representational tools. As such, the simplest way to conceive of sonifications is as sound-based analogues of charts, graphs, maps, or other visualizations. Where visualizations use points, lines, and other visual devices, sonification employs sounds with varying timbre, pitch, loudness, stereo position, timing and rhythm, consonance and dissonance, and other sonic properties.

The most significant advantage of sonification for exploring interactional social data is the inherent temporality of sound. While visual representations necessarily have a spatial dimension, sound always unfolds over time. Dayé and de Campo pointed out that this makes sonification excellent at conveying sequential information (2006). Using sonification, data representing events that unfold over time, such as asynchronous social interactions, can be conveyed along their natural dimension, time, instead of spatially as most visualizations would place them. Even though it is possible to create temporal visualizations that utilize animation, the human auditory system performs significantly better with rhythmic perception and temporal resolution than the visual system (Neuhoff, 2011). As a result, sonifications have the potential to effectively represent minute patterns along the dimension of time.

Another advantage is that representing data as sound can draw attention to regularly occurring patterns that might be difficult to discern using other methods. Sonification has particular merit for trend analysis in which listeners identify overall patterns of increases and decreases in quantitative data (Walker & Nees, 2011, p. 21). According to Ferguson, Martens, and Cabrera:

> Auditory representations can potentially extract patterns not previously discernible, and might make such patterns so obvious to the ear, that no-one will ever look for them with their eyes again. By capitalizing upon the inherently different capabilities of the human auditory system, invisible regularities can become audible, and complex temporal patterns can be "heard out" in what might appear to be noise. (2011, p. 178)

One of the main functions of exploratory analysis is to obtain a new perspective of data. By perceiving data in new ways, one can identify patterns and features that are not evident using traditional methods. Many exploratory analyses of social media data have tended to favour visual exploration. In contrast, sonification promises to illuminate different aspects of the data, particularly temporal dimensions for which it may be better suited than spatially oriented visual methods. By allowing researchers to perceive data in a novel way, sonification is conducive to generating new types of hypotheses.

Examples of social media sonification

Several researchers have used sonification to listen to social media data. This section presents a brief overview of significant works, illustrating the current state of social media sonification. These examples illustrate how sonification has been used for summarizing and analyzing logs of social media activity, and also point to areas for further development.

Detecting anomalous events with sonification: Ballora et al., 2012

Ballora et al. (2012) created an application that sonified stock market data alongside logs of tweets containing keywords related to those stocks. This sonification was used to make anomalous events detectable even to untrained listeners. Ballora et al. tested this system using stock market and Twitter data related to technology companies leading up to and following the Apple Worldwide Developer's conference in 2011. They found that test subjects who listened to this sonification could easily identify changes in the data around the time of the conference.

This study demonstrated that sonification can illuminate anomalous changes to data streams in a way that is apparent even to untrained listeners. Additionally, Ballora et al. demonstrated a novel approach to combining data sources with very different levels of precision. The Twitter data, which utilized keywords such as 'apple,' was less precise than the stock market data. Specifically, it could be ambiguous whether particular tweets containing the keyword 'apple' were related to Apple Corporation, while stock market data was clear in this regard. As a result, they determined that small-scale changes in the Twitter data were unlikely to be reliably significant, and designed the sonification of tweets to focus on large-scale changes.

This was accomplished by condensing the Twitter data into fifteen-minute histograms and allowing the sounds for each histogram to overlap somewhat, emphasizing overall trends rather than precise changes. Ballora et al. described how the two soundtracks differed:

> One soundtrack renders selected stock prices as rhythmically unique pulses, the pitches of which reflect stock price fluctuations. The second soundtrack maps selected keywords appearing in tweets to unique drone-like pitches, so that the appearance of a keyword is rendered as a simple, sustained tone at its associated pitch, at a particular amplitude. The result is a "sound cloud" of bell-like pulses and harmonically related drones. Periods of increased or decreased activity are easily perceptible as changes in the sound cloud's density, timbre, and rhythmicity. (2012, p. 1)

The use of two soundtracks illustrates how different sonification techniques can be suited to particular types of data. Ultimately, testing indicated that the sonification made significant anomalies apparent. Future work of this sort may benefit by exploring techniques for sonifying subtler patterns and changes.

Twitter and music: Ash, 2012; Bethancourt, 2012

In 2012 the International Community for Auditory Display (ICAD) – a central research community for the study of sonification – held a competition called ‘Listening to the World Listening.’ Entrants were invited to submit sonifications of the Twitter Music Trends data feed – a list of the top 50 trending artists on Twitter, updated every two seconds. The winning entrant was Kingsley Ash’s ‘Affective States,’ which represented each artist with a distinct tone, then modified several audio filters for that tone based on the presence of emotion keywords in blog postings about the artist (Ash, 2012). Another project was Matt Bethancourt’s ‘The sounds of the discussion of sounds,’ which played, in real-time, tones representing the amount of Twitter discussion about particular artists (Bethancourt, 2012). As an artist’s popularity waned, their tone would become quieter, potentially becoming silent if Twitter users stopped discussing them. Both examples produced a sonification that constantly shifted in relation to real-time Twitter discussions, and this allowed them to make use of the inherent temporality of sound.

Tweetscapes: Hermann, Thomas, Nehls, Eithel, Barri, and Gammel, 2012

The Tweetscapes project (Hermann, Thomas, et al., 2012) was a real-time sonification of Twitter activity in Germany, hosted at www.tweetscapes.de. The sounds produced by this sonification are described as ‘an interactive composition performed by Germany’s Twitter users’ (HEAVYLISTENING, 2012). The sonification ran for three years, from 2012 to 2015. Tweets originating from Germany were sonified in real time, with the sound for each tweet being modified according to parameters such as the number of followers for that tweet and its distance from the geographic centre of Germany (which determines reverberation and stereo panning). As well as a general stream, a hashtag stream is available in which hashtag keywords are distinguished by different sound samples and synthesis settings (Hermann, Thomas et al., 2012). One of the goals of Tweetscapes was to make sonification more publicly known, which was achieved in part through integrating its sonifications in the nationwide radio program *Deutschlandradio Kultur*. Although the project received significant media attention, the researchers acknowledge that the question of Tweetscapes’ practical use was often raised, and commented that ‘the practical use is very limited’ (2012, p. 119). One of the limitations of Tweetscapes was the lack of interactivity. Users were not able to filter the selection of tweets they listened to, but instead would listen to the entirety of German Twitter activity. On the website, the sonification was combined with a visualization of each tweet overlaid onto a map of Germany. This provided additional context and made the sonification easier to understand.

User-focused sonification: Wolf, Gliner, and Fiebrink, 2015

In 2015, Wolf, Gliner, and Fiebrink proposed a model for data-driven sonification using soundscapes. Their goal was to facilitate end-user involvement in the process of designing a sonification so as to make sonifications more useful for those users. They prototyped this model by sonifying Twitter data, and expressed that their design would build upon the techniques of projects like Tweetscapes by allowing users to 'select the Twitter information they wish to monitor in real-time' (Wolf et al., 2015, p. 3). This effort to build a system in which users can select specific groups of Twitter data to be sonified in real time has significant potential for exploratory data analysis. Additionally, this system uses a simple sonification engine in which data is mapped to familiar sound samples, such as 'bird tweet' or 'running water' rather than potentially unfamiliar sound parameters such as frequency and timbre (Wolf et al., 2015). These examples demonstrate multiple purposes for sonification. The purpose of some of these projects is largely to evaluate and explore sonification's potential, and/or to make sonification familiar to a broad audience. These projects also illustrate sonification's usefulness for exploring, although there is room for advancement.

Considerations for sonification research

One of the most significant challenges for sonification is that interpretation can be difficult, particularly if the listener is unfamiliar with sonification. The following section discusses considerations for conducting exploratory analysis with sonification, particularly as pertaining to interactional data.

Tools for sonification

For researchers who want to utilize sonification techniques, the scarcity of available tools can be intimidating. Many sonifications are custom designed for specific research projects using complex software such as Max/MSP, SuperCollider or other sophisticated software or hardware synthesizers. However, it is also possible to build sonifications using tools such as *Sonification Sandbox* (2009) or the *E-Rhythms Data Sonifier* (2014).

Necessary criteria

If a sonification is to be useful for analyzing data, it must have consistently and clearly defined criteria. Hermann (2008) has argued that four criteria are particularly important to meet this standard. According to Hermann, a sonification must reflect *objective* properties or relations in the input data, the transformation must be *systematic*, the sonification should be *reproducible*, and the system should be flexible to work with multiple sets of *different data* (2008, p. 2). These four criteria are conducive for designing sonifications with data analysis in mind, since they emphasize the importance of representing data with rigour and reliability. To provide an example, a song whose composition is loosely inspired by a dataset would be unlikely to meet Hermann's criteria (since a songwriter's composition process likely involves subjective, creative decisions); however, an algorithmic transformation of that data into sound would qualify.

Appropriate tasks

Some analytic tasks are better suited to sonification than others. For example, point estimation for a particular datum (e.g. identifying that the value is 1.0 exactly, not 0.9 or

1.1) can be very difficult using sonification (Smith & Walker, 2005). In addition to it being difficult to identify the value of individual points, comparing multiple points poses another challenge. Point comparison requires estimation of two distinct points, plus a memory task of comparing the values of each. Walker and Nees (2011) theorize that point comparison should be more difficult than point estimation, but note that no empirical tests have examined point comparison with sonification. The difficulty of point estimation (and theorized difficulty of point comparison) with sonification is worthy of consideration because this task can be straightforward using a chart or other visualization. On the other hand, trend analysis, in which listeners assess overall patterns in a data source, is well suited for sonification. Walker and Nees (2011) suggested sonification is especially useful for trend analysis because 'sound may be a medium wherein otherwise unnoticed patterns in data emerge for the listener' (p. 21).

Furthermore, sonification is especially well-suited for exploring temporal and rhythmic patterns. As noted above, sound is inherently temporal, and the human listening system has far better temporal and rhythmic perception than the visual system. For example, humans are typically able to hear gaps in broadband noise stimuli as short as 2–3 milliseconds (Carlile, 2011) and at low frequencies it is possible to distinguish individual events with durations as brief as 20-50 milliseconds (Dombois & Eckel, 2011). This excellent temporal resolution bolsters sonification's utility for revealing patterns and anomalies that are difficult to perceive in other representations of the data.

Simultaneous streams and levels of analysis

The simultaneity of sonification is advantageous for representations of temporal data, but poses a challenge in how much data can be perceived at once. It is possible to listen to more than one stream simultaneously, but the difficulty of this task increases according to the complexity of the sound and the level of precision required to evaluate the sonification. In Ballora et al.'s (2012) sonification of Twitter and stock market data, monitoring multiple streams simultaneously was required. However, the purpose of the sonification was to detect significant anomalies, rather than to conduct precise analysis. Tweetscapes (Hermann, Thomas, et al., 2012) – a real-time sonification of German Twitter activity – also presented multiple streams of data by sonifying various hashtags separately, but is similarly not intended for analyzing minute patterns.

Listening to a very large number of streams simultaneously may lead the listener to perceive them as a group, rather than as multiple individual streams. This is particularly likely if individual streams sound similar to one another. Whether one listens to the patterns of individuals or of the group as a whole has a significant effect on the types of observations that are possible. Group level analysis is best suited for identifying patterns of activity based around fixed points in time. For example, one could observe overall trends such as large numbers of people tweeting about a particular event, exchanging text messages on New Year's Eve, or tending to be more active on social media during the day than late at night.

On the other hand, when conducting group level analysis with sonification, individual patterns may be obfuscated amidst the noise created by multiple overlapping streams. If for example, one individual sends 5 text messages each on Monday, Wednesday, and Friday, and another sends 5 text messages on Tuesday, Thursday, Saturday and, Sunday, a sonification in which these individuals' streams were merged would indicate that 5 text messages were sent every day, without capturing further details about the individual patterns. Any sonification that amalgamates multiple individuals into a group level analysis is likely to muddy individual patterns.

An alternative to listening to multiple audio streams simultaneously is to isolate individual streams. Listening to individuals engaged in dyadic interactions can reveal information that is obscured at the group level. When listening to individual streams, comparing streams to each other requires switching between them. When performing this sort of switching, listeners should be careful to account for the human listening system's need for time to adapt and become familiar with each stream (Hermann, Hunt, & Neuhoff, 2011). Ultimately, each level of analysis has its own advantages and disadvantages, and researchers should consider which is most appropriate based on the complexity of the sound, and the type of patterns one is searching for.

Group analysis. Group level sonifications are most useful for identifying overall trends and patterns that are based around fixed points of time. Using group level analysis, it is possible to listen to large populations at once, and to gain a holistic perspective of certain trends within those populations. Because researchers can listen to a large number of individuals at once, results observed from group level sonifications are likely to be generalizable.

Individual analysis. Sonifications of individual patterns may allow researchers to observe patterns that would be obscured in group level analyses. Individual analysis offers the most precise resolution of the data, and may be useful for identifying subtle patterns or patterns oriented around points in time that are relative to each individual. The challenge posed by individual level analysis is that listening to large numbers of individuals may be time consuming. As a result, for individual level analysis it is important to have an effective sampling strategy.

Dyadic analysis. Listening to interactions between pairs of individuals may be particularly useful for interactional data, as it can allow researchers to investigate patterns of communication between two individuals. This has the same advantages and disadvantages as individual level analysis, but can also draw attention to features such as the speed at which individuals respond to each other or who usually initiates communication.

Combinations. In some cases, it may be appropriate to combine multiple levels of analysis. This has the potential to reveal ways in which individual patterns relate to patterns among the larger population. For example, researchers could compare group level activity to an individual stream as a method of identifying ways that the individual differs from the group. Alternately, one could listen to a sonification of an individual's social media activities alongside a group level of sonification of replies, shares and other responses to the individual's activities.

Mapping sounds

When listening to sonification it is necessary to consider how some sounds can suggest meanings independent of the source data (Grond & Hermann, 2011; Walker & Kramer, 2005). A sequence of notes in a major key may suggest a happier meaning than a minor key, regardless of whether a sense of happiness or sadness accurately reflects the information being conveyed. And a sonification that is thunderously loud suggests different emotional meanings than one that is meek, even if both represent the same data.

Supper (2014) discussed how sound design can introduce emotional meanings, using sonifications of natural phenomena as examples. In some cases, she argued, the meaning conveyed by sound design can be conflated with the information from the source data. This has the potential to skew interpretations, but sonification designers can also take advantage of this to create sonifications that illustrate rich meanings by seeming to be *true* to the phenomenon being represented:

> The sonification, for instance, of a volcano is different from the sounds that are emitted by the volcano itself. However, certain rhetorical, musical and technological strategies are used to suggest that the sonification represents something about the volcano that might not be immediately visible or audible from the volcano, but from deeper within it. It is not about sounding like a volcano per se, but about being true to the volcano – or rather, about allowing listeners to believe that the sonification is true to the volcano. (Supper 2014, p. 51)

Although Supper refers to a sonification of a volcano, the same principle is valid for sonifications of digital phenomena such as social media activities. One might choose to indicate that a particular post was shared many times by adding reverberation and echo, or by modifying its pitch or loudness. The choice of mapping can often have a strong effect on how listeners interpret meaning, and some mappings will seem truer to the data than others.

As a consequence, even when a sonification is systematic its designers have the ability to steer interpretations through aesthetic decisions. Sonification can convey a variety of subjective and emotional meanings, and this is exemplified in a special sonification issue of *AI & SOCIETY* that combined articles from both sonification researchers and artists (Sinclair, 2011). The artistic potential of sonification is in some respects at cross-purposes to its scientific analysis applications. However, this is no different from other perceptualization methods such as visualization. Just as a sonification can introduce bias through sound design, visualizations can suggest deceptive meanings through choice of colour, symbols, and scales. In both cases, researchers should endeavour to understand enough about the form to detect misrepresentations where possible.

Training

An important consideration is that sonification is unfamiliar to many researchers. Whereas most researchers are familiar with at least some visualization techniques, sonification is much less common. Generally, listeners with musical ability or training tend to be more accurate than musically untrained listeners when interpreting sonifications (Neuhoff, Knight, & Wayand, 2002). However, even among sonification specialists, there are many cases where a common vocabulary of representational techniques is lacking. As a result, usually at least some training is required for listeners to interpret a sonification (Walker & Nees, 2011). In some cases, sonification designers may include instruction manuals or specific training procedures. In all cases, it is advisable that listeners familiarize themselves with a sonification system before attempting to make new discoveries. One method is to listen to aspects of the data the researcher is already familiar with. Developing an understanding of how the sonification represents known patterns is a valuable step toward being able to discover new patterns.

Visual cues

Finally, accompanying a sonification with a visual component can be useful to provide context and make the sonification more easily understood. Presenting sonification alongside a corresponding visualization can allow researchers to utilize the temporal strengths of sonification alongside the spatial strengths of visualization. For example, in the E-Rhythms Data Sonifier, researchers can click on a visually presented timeline to navigate through the data. And in Tweetscapes, the geographic origin of tweets is visualized by overlaying graphical representations on a map.

Example: Sonifying asynchronous text message logs

In the following section we provide an example of sonification being used to explore time-stamped interactional data (Further discussions of this study was presented in Jamieson, Boase, & Kobayashi, 2015a, 2015b). We used data sonification to conduct exploratory analysis of non-identifying smartphone logs of text messaging. This data is similar to social media activity logs, which often catalogue asynchronous communication. The E-Rhythms Data Sonifier software (2014) was used to explore the data in several ways. The most fruitful exploration consisted of listening to the speed at which pairs responded to each other's text messages, and considered how this related to relational dimensions including relationship role (family, co-worker, or other), discussing important matters, and trust. We discuss our exploratory process as an illustration of the strengths and challenges of using sonification to explore communication logs.

Data description

The data for this study was collected using the Network Navigator application, which respondents installed on their Android smartphones. The application collected non-identifying voice, text, and email log data and correlated these logs with responses to on-screen survey questionnaires. These surveys included questions about recently contacted ties or communication partners, such as whether they were family members, whether respondents trusted them, and whether respondents discussed important matters with them. The full data set contained logs collected from 132 adults living in the United States in 2011 who explicitly consented to participating in the study. Our study focused on text messages and relied on responses to survey questions to provide information about ties. We focused on text messages because our sonification method represented the number of events that occur, but did not indicate the duration of events. Moreover, the data did not contain the content of calls or other information that could have made it possible to infer communication patterns within each telephone conversation. Consequently, text messages, which are discrete asynchronous communications, were better suited to our study. We limited our study to communications with ties where 1) the respondent answered at least one pop-up survey about that tie and, 2) at least one text message was exchanged with that tie. This narrowed our selection to 77 respondents, who exchanged a total 11,215 text messages with 149 ties.

The E-Rhythms Data Sonifier

Exploratory sonification analysis was conducted using the E-Rhythms Data Sonifier software, which was designed by the authors (as of this chapter's publication, the Data sonifier software is available for free download at http://individual.utoronto.ca/jboase/software.html). A time-stamped data file is loaded into the software, and a sonification is created to represent the amount of activity over time. Data can be filtered according to its contents and sent to distinct sounding tracks. For example, incoming text messages can be represented with a different sound than outgoing text messages, or communication among family could be distinguished from communication with coworkers. After filtering which data will be represented by each sound, the researcher chooses a length of time to be represented by each beat. Time is condensed, so a researcher might set each beat to represent one hour of activity, then play back the sonification at one beat per second. At each beat, a sound is triggered, representing the number of events that took place during that period. This makes the sonification akin to a histogram where each beat indicates the number of events that occurred over a given period of time. The more events that take place, the more intense the sound. Researchers can choose to indicate this intensity with either loudness or pitch, depending which they think is most suitable for their data. In our study, the number of events was mapped to loudness; loud sounds indicated many text messages were exchanged, soft sounds indicated fewer text messages, and silence indicated that no

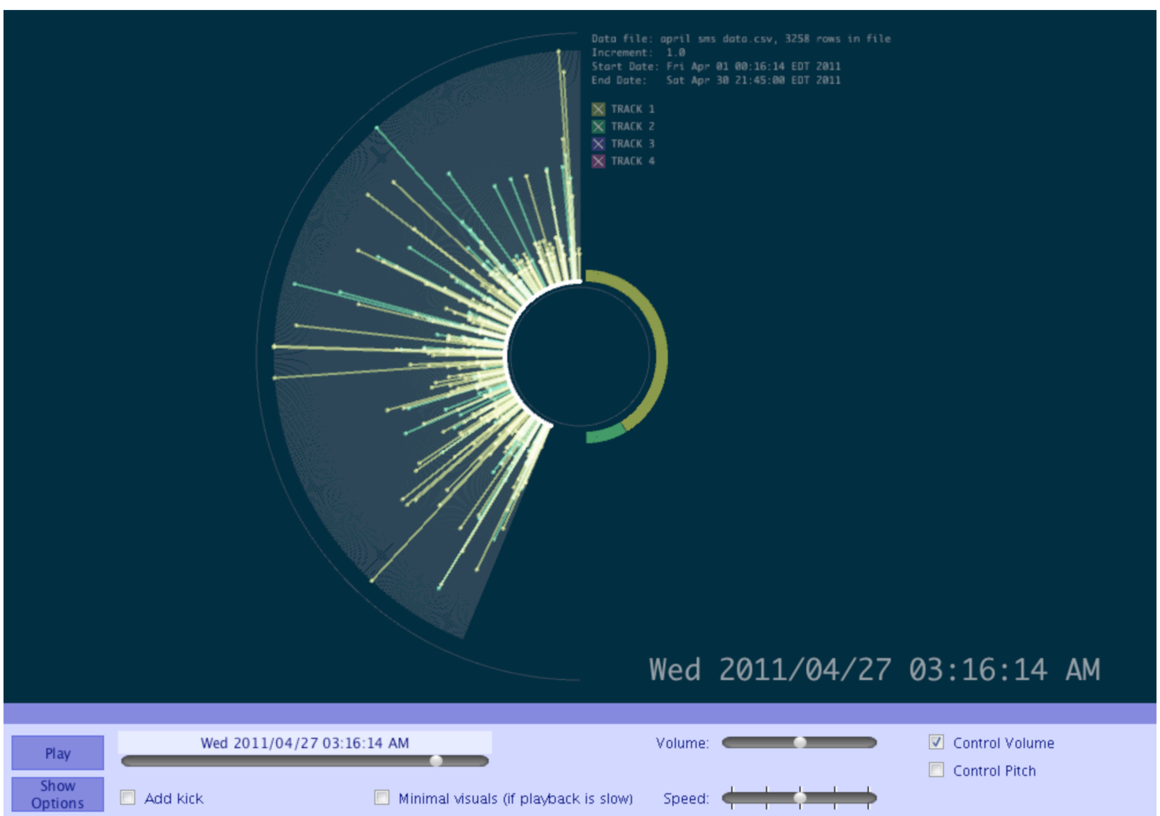

Figure 1: Screenshot of E-Rhythms Data Sonifier

messages were exchanged.

In addition to sonification, the software includes a visualization component. The visualization makes it easier to contextualize each sound in relation to overall trends and aids temporal navigation through the data. For example, our text message log data sometimes included periods of inactivity, and we used this visualization to quickly locate periods of activity without having to listen to long passages of silence.

Listening method and findings

When listening, we filtered data according to the results of survey questions. This made it possible to compare communication patterns among family to those among coworkers, or communications with trusted ties to those with untrusted ties. At first, we listened to entire groups at once. When listening at the group level, it was possible to identify patterns based around fixed points in time. For example, we created a sonification that divided each day into four beats, and observed a rather musical pattern of three beats followed by a pause – *one, two, three, (pause), one, two, three, (pause)*. The pause indicated that participants rarely exchanged text messages in the middle of the night. Observing this expected pattern helped us to familiarize ourselves with some of the types of patterns that sonification could draw forth. Group level sonifications were successful at revealing consistent patterns such as a day/night cycle or bursts of activity around holidays, but could not illuminate individuals' patterns. As discussed above, group level sonification tends to allow the patterns of individuals to become lost in a sea of group activity.

To be able to listen to communication between individuals, we randomly selected 16 pairs, each consisting of one respondent and one of their ties, and listened to their texting activity at a dyadic level. To ensure we had enough activity to be able to listen to distinct patterns, we only selected pairs who had exchanged at least 100 texts between each other over the course of their saved logs. Within each pair, we assigned each individual a distinct sound, so it was possible to distinguish between messages sent by each. The first pattern that became evident upon sonifying activity at this dyadic level was that response time varied considerably between each pair. Some pairs consistently replied to each other within a few minutes, while others took up to several hours to reply to text messages. A limitation of our listening was that we lacked information about the content of messages, which made it impossible to assert which messages were replies and which started new conversations. However, it was possible to infer that, for example, a message that occurred after several days of no communication was likely to indicate a new conversation, and communications with only a few minutes between them were almost certainly part of the same conversation.

Previous research has suggested that temporal cues such as time of day or the speed at which people respond to each other can be indicative of the intimacy of their relationship (Döring & Pöschl, 2009; Kalman & Rafaeli, 2011; Quan-Haase & Collins, 2008; Walther & Tidwell, 1995). Notably, Walther and Tidwell (1995) found that shorter response times to task messages such as work communications are likely to indicate more intimacy and eagerness than long response times, but that longer response times to social communications may indicate more intimacy than quicker responses. This is likely because intimate social partners may feel less pressure to respond quickly. Building from this research, we recorded an estimated average response time for each pair then noted

Pearson correlations between that estimated average response time and various measures of tie strength represented through the survey results. This preliminary analysis suggested three findings, which we formed into hypotheses for further testing.

> H1: Family members will have shorter response times than non-family.
>
> H2: Ties who are trusted by respondents will have longer response times than untrusted ties.
>
> H3: Pairs who discuss important matters will have shorter response times than those who do not.

Statistical testing

As stated earlier, exploratory analysis is better suited to hypothesis generation than to hypothesis testing. We conducted a statistical analysis to assess the validity of our sonification-generated hypotheses. First we identified messages that could potentially be replies (i.e. where the direction of communication changed) and generated a response time variable noting the time between these messages (N = 4,687). Because this variable included long gaps where the pair went up to days or weeks without communicating, we used a complete linkage cluster analysis on the response time variable to focus only on texts that could reasonably be considered direct replies. The cluster analysis allowed us to create groups of similar response times, and distinguish shorter response times that were likely to indicate replies from longer response times that indicated silences between conversations. Without viewing the content of messages, clustering was an approximate method of distinguishing replies from messages that initiated new conversations. We used this method because it provided a reasonable approximation to the more intuitive distinction that was made when listening to the sonification. Moreover, Kalman and Rafaeli (2011) stated that a pause of ten times the average latency constituted silence (and therefore a lack of reply) in online asynchronous communications. Our clustering method identified a cluster of 74% of events (n=3,487) with a mean response time of 216 seconds. The next largest cluster (n=216) had a mean response time of 2,305 seconds,

Table 1: Mean and median response times among sonified sample.

	Response time in seconds			
	Mean		Median	
Category	Incoming	Outgoing	Incoming	Outgoing
Family members	214	224	208	243
Not family members	476	205	214	159
Trusted ties	469	190	247	198
Not trusted ties	175	252	179	266
Discuss important matters	347	185	184	205
Do not discuss important matters	408	260	244	223

over ten times that of the larger group, suggesting that this was a reasonable albeit rough method of distinguishing between replies and new conversations.

We then calculated the mean incoming and outgoing response time (in seconds) for each pair, and then calculated mean and median response times for the independent variables noted in our hypotheses. Since survey questions were answered only by one member of each pair (the respondent), incoming refers to messages sent to that respondent, and outgoing refers to messages sent by that respondent. These mean and median response times are listed in Table 1, below.

The mean response times observed through statistical analysis were consistent with our hypotheses, although H1 and H2 only appear to be valid for incoming responses. Median response times were less varied, and do not exhibit significance for H2 and H3, but are fairly consistent with H1. This suggests that the sonification testing emphasized variances among response times that were relatively long, and that differences among shorter response times were less correlated with how participants responded to the survey questions.

Statistical analysis was then used to examine the extent to which these findings applied to the large sample of 149 tie pairs. The results supported H2 and H3, but not H1. In fact, statistical analysis indicated that mean response time for family members was higher than for non-family – the opposite of what occurred among the 16-pair sample.

Discussion

Our analysis consisted of two stages. Exploratory sonification was used to generate hypotheses, which were then tested using statistical analysis. The fact that the hypotheses generated through sonification were supported by the initial statistical analysis of the 16-pair sample indicates that sonification was generally successful at identifying patterns within the data. It may have been possible to perform similar testing without sonification, but sonification was effective at highlighting response time as a variable for further investigation. In the dataset, each event was time-stamped, making response time an implicit variable. However, sonifying text message exchanges between pairs made response time explicit. This supports the notion that exploratory sonification can be an effective method for hypotheses generation. This also emphasizes the importance of considering different levels of analysis. Listening to the data at group, individual, and dyadic levels highlighted different types of patterns in the data, but each level also had its

Table 2: Hypotheses generated through sonification compared to results of statistical analysis

	Mean response time for each independent variable		
	Hypotheses from sonification	Results of statistical analysis	
Category	Sonified sample (16 tie pairs)	Sonified sample (16 tie pairs)	Full sample (149 tie pairs)
Family members	Shorter	Shorter	Longer
Not family members	Longer	Longer	Shorter
Trusted ties	Longer	Longer	Longer
Not trusted ties	Shorter	Shorter	Shorter
Discuss important matters	Shorter	Shorter	Shorter
Do not discuss important matters	Longer	Longer	Longer

own set of associated challenges.

One challenge was illustrated by the fact that mean response times among family and non-family were different between the 16-pair sample that was sonified and the larger 149-pair sample. This indicates the randomly selected 16-pair sample was not representative in regards to H1. In order to make response time between pairs apparent using sonification it was necessary to listen at a dyadic level, which increased the potential for sampling errors. Verifying our sonification-generated hypotheses using statistical analysis provided an opportunity to identify this error by analyzing the larger population of pairs at once. This allowed us to test the hypotheses among the whole sample in a way that would not have been feasible with sonification alone.

Future Directions

Sonification demonstrates potential for exploratory analysis of time-stamped interactional data, but there is much room for future work. Ferguson, Martens, and Cabrera (2011) reflected upon the current state of exploratory statistical analysis using sonification:

> It must be said that the current state of the art must be considered to be quite immature as yet, with many challenges for sonification research to tackle in the future. In fact, it might be proposed that the best approach to take in designing and developing statistical sonifications in particular would be one that includes critical evaluation of the results at each attempt. (p. 192)

While sonification has demonstrated analytic potential – especially with temporal information – there is still much to be learned. For this reason, sonification has more immediate potential for hypothesis generation than hypothesis testing. Testing hypotheses generated using sonification will lead to a better understanding of the strengths and limitations of sonification.

Among existing sonifications of social media it is common to present a real-time overview of activity. Typically these sonifications allow one to monitor activity and to notice large shifts or sudden changes. Some of these projects, such as Tweetscapes, were created with the goal of increasing awareness of sonification. For this purpose emphasizing broad, easily observable trends is a suitable strategy. Other techniques may be used to allow researchers to conduct deeper analyses. For example, encouraging temporal navigation such as rewinding, fast-forwarding, and looping would make it easier for researchers to analyze specific temporal passages in detail. Additionally, allowing researchers to listen at different levels of analysis such as individual and dyadic would facilitate the discovery of different types of trends than can be revealed through group analysis alone.

Currently, a large portion of research papers about sonification have been written by researchers who are themselves engaged in creating or evaluating sonification methods (Supper, 2012). To support the field’s efforts at outreach into research domains, several designers have attempted to create tools that can be used by researchers who are not sonification specialists. For example, Grond designed a sonification tool for molecular structures and dynamics, and distributed the tool as a plugin for a software package that his potential audience of chemists was already familiar with (as cited in Supper, 2011, p. 256). This reduced the learning curve and allowed molecular researchers to use Grond's sonification alongside their existing toolset. Forthcoming sonification plugins for statistical packages such as R may be successful in allowing researchers to incorporate sonification (see for e.g. Stone & Garrison, 2013). Additionally, as discussed earlier Wolf et al.'s work (2015) on involving end-users in sonification design has potential for broadening the field by allowing users to design sonifications with particular applications in mind. Lastly, the E-Rhythms Data Sonifier (2014) is designed to use a limited number of relatively simple sound properties (primarily time and volume), and can be used with almost any time-stamped data. This makes it a viable tool for social researchers who do not have prior expertise with sonification.

Researchers studying online communication benefit from an abundance of interactional data, such as time-stamped logs of activities. The growth of generalized sonification tools provides an opportunity to explore this data in new ways. The high level of detail and large quantity of this data gives it the potential to illuminate patterns and trends that have not been apparent in other representations of data. Exploratory analysis is an important technique for uncovering these potential patterns, and the time-dimension present in much of this data can be well explored through sonification.

References

Andrienko, N., & Andrienko, G. 2006. Exploratory analysis of spatial and temporal data a systematic approach. Berlin: Springer.

Ash, K. 2012. Affective states: Analysis and sonification of Twitter music trends. In Proceedings of the 18th International Conference on Auditory Display (pp. 257–259). Atlanta, GA.

Ballora, M., Cole, R. J., Kruesi, H., Greene, H., Monahan, G., & Hall, D. L. (2012). Use of sonification in the detection of anomalous events. In J. J. Braun (Ed.), *Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications*. http://doi.org/10.1117/12.918688

Bethancourt, M. 2012. The sounds of the discussion of sounds. In Proceedings of the 18th International Conference on Auditory Display (pp. 254–256). Atlanta, GA.

Carlile, S. 2011. Psychoacoustics. In T. Hermann, A. Hunt, & J. G. Neuhoff (Eds.), The sonification handbook (pp. 41–62). Berlin: Logos Verlag.

Dayé, C., & de Campo, A. 2006. Sounds sequential: Sonification in the social sciences. Interdisciplinary Science Reviews, 31(4), 349–364. http://doi.org/10.1179/030801806X143286

Dombois, F., & Eckel, G. 2011. Audification. In T. Hermann, A. Hunt, & J. G. Neuhoff (Eds.), The sonification handbook (pp. 301–324). Berlin: Logos Verlag.

Döring, N., & Pöschl, S. 2009. Nonverbal cues in mobile phone text messages: The effects of chronemics and proxemics. In R. S. Ling & S. W. Campbell (Eds.), The reconstruction of space and time: Mobile communication practices (pp. 109–135). New Brunswick, N.J: Transaction Publishers.

E-Rhythms Data Sonifier [Computer software]. 2014. Accessed 2 Nov 2015 from http://individual.utoronto.ca/jboase/software.html#e-rhythmsdataanalysis

Ferguson, S., Martens, W., & Cabrera, D. 2011. Statistical sonification for exploratory data analysis. In T. Hermann, A. Hunt, & J. G. Neuhoff (Eds.), The sonification handbook (pp. 175–196). Berlin: Logos Verlag. Accessed 2 Nov 2015 from http://sonification.de/handbook/download/TheSonificationHandbook-HermannHuntNeuhoff-2011.pdf

Grond, F., & Hermann, T. 2011. Aesthetic strategies in sonification. AI & SOCIETY, 27(2), 213–222. http://doi.org/10.1007/s00146-011-0341-7

HEAVYLISTENING. 2012. #tweetscapes - Listen to Twitter [Video file]. Accessed 2 Nov 2015 from https://www.youtube.com/watch?v=0lKSFlB_-Q0

Hermann, T. 2008. Taxonomy and definitions for sonification and auditory display. In Proceedings of the 14th International Conference on Auditory Display. Paris, France.

Hermann, Thomas, Nehls, A. V., Eithel, F., Barri, T., & Gammel, M. 2012. Tweetscapes – Real-time sonification of Twitter data streams for radio broadcasting. In Proceedings of the 18th International Conference on Auditory Display (pp. 113–120). Atlanta, GA.

Hermann, T., Hunt, A., & Neuhoff, J. G. (2011). *The sonification handbook*. Berlin: Logos Verlag. Accessed 2 Nov 2015 from http://sonification.de/handbook/download/TheSonificationHandbook-HermannHuntNeuhoff-2011.pdf

Hermann, T., Hunt, A., & Neuhoff, J. G. 2011. Introduction. In T. Hermann, A. Hunt, & J. G. Neuhoff (Eds.), *The sonification handbook* (pp. 1–6). Berlin: Logos Verlag. Accessed 2 Nov 2015 from http://sonification.de/handbook/download/TheSonificationHandbook-HermannHuntNeuhoff-2011.pdf

Jamieson, J., Boase, J., & Kobayashi, T. 2015a. *The faster the response, the stronger the bond? Dimensions of relational closeness and texting response time*. [Poster] presented at the 65th Annual conference of the International Communication Association., San Juan, Puerto Rico.

Jamieson, J., Boase, J., & Kobayashi, T. 2015b. *Using sonification to explore texting response time in time stamped interactional data.* Presented at the CAIS/ACSI & LRI 2015 Conference, Ottawa, Canada.

Kalman, Y. M., & Rafaeli, S. 2011. Online Pauses and Silence: Chronemic Expectancy Violations in Written Computer-Mediated Communication. Communication Research, 38(1), 54–69. http://doi.org/10.1177/0093650210378229

Neuhoff, J. G. 2011. Perception, cognition and action in auditory displays. In T. Hermann, A. Hunt, & J. G. Neuhoff (Eds.), The sonification handbook (pp. 63–87). Berlin: Logos Verlag. Accessed 2 Nov 2015 from http://sonification.de/handbook/download/TheSonificationHandbook-HermannHuntNeuhoff-2011.pdf

Neuhoff, J. G., Knight, R., & Wayand, J. 2002. Pitch change, sonification, and musical expertise: Which way is up. In Proceedings of the 2002 International Conference on Auditory Display. Accessed 2 Nov 2015 from http://www.icad.org/Proceedings/2002/NeuhoffKnight2002.pdf

Quan-Haase, A., & Collins, J. L. 2008. I'm there, but I might not want to talk to you. Information, Communication & Society, 11(4), 526–543. http://doi.org/10.1080/13691180801999043

Sinclair, P. (Ed.). 2011. Sonification: What where how why artistic practice relating sonification to environments [Special issue]. AI & SOCIETY, 27(2). http://doi.org/10.1007/s00146-011-0346-2

Sonification Sandbox [Computer software]. 2009. Accessed 2 Nov 2015 from http://sonify.psych.gatech.edu/research/sonification_sandbox/

Stone, E., & Garrison, J. 2013. audiolyzR: Give your data a listen (Version 0.4-9). Accessed 2 Nov 2015 from http://cran.r-project.org/web/packages/audiolyzR/index.html

Supper, A. 2012. "Trained ears" and "correlation coefficients": A social science perspective on sonification. In Proceedings of the 18th International Conference on Auditory Display (pp. 29–35). Atlanta, GA.

Supper, A. 2014. Sublime frequencies: The construction of sublime listening experiences in the sonification of scientific data. Social Studies of Science, 44(1), 34–58.

Tufte, E. R. 2001. The visual display of quantitative information (2nd ed.). Cheshire, CT: Graphics Press.

Tukey, J. W. 1977. Exploratory data analysis. Reading, MA: Addison-Wesley Pub. Co.

Walker, B. N., & Kramer, G. 2005. Mappings and metaphors in auditory displays: An experimental assessment. ACM Transactions on Applied Perception (TAP), 2(4), 407–412.

Walker, B. N., & Nees, M. A. 2011. Theory of sonification. In T. Hermann, A. Hunt, & J. G. Neuhoff (Eds.), The sonification handbook (pp. 9–40). Berlin: Logos Verlag. Accessed 2 Nov 2015 from http://sonification.de/handbook/download/TheSonificationHandbook-HermannHuntNeuhoff-2011.pdf

Walther, J. B., & Tidwell, L. C. 1995. Nonverbal cues in computer-mediated communication, and the effect of chronemics on relational communication. Journal of Organizational Computing, 5(4), 355–378. http://doi.org/10.1080/10919399509540258

Wolf, K. E., Gliner, G., & Fiebrink, R. 2015. A Model for Data-Driven Sonification Using Soundscapes (pp. 97–100). ACM Press. http://doi.org/10.1145/2732158.2732188