Web Sonification Sandbox - an Easy-to-Use Web Application for Sonifying Data and Equations

Zachary Kondak*
Georgia Institute of Technology
Center for Music Technology
840 McMillan St. Atlanta GA 30332
zakondak@gatech.edu

Tianchu (Alex) Liang, Brianna Tomlinson
Georgia Institute of Technology
College of Computing
801 Atlantic Dr NW, Atlanta, GA 30332
liangtianchu@gmail.com, btomlin@gatech.edu

Bruce N. Walker
Georgia Institute of Technology
School of Psychology
654 Cherry Street, Atlanta, GA 30332-0170
bruce.walker@gatech.edu

ABSTRACT

Auditory and multimodal presentation of data ("auditory graphs") can allow for discoveries in a data set that are sometimes impossible with visual-only inspection. At the same time, multimodal graphs can make data, and the STEM fields that rely on them, more accessible to a much broader range of people, including many with disabilities. There have been a variety of software tools developed to turn data into sound, including the widely-used Sonification Sandbox, but there remains a need for simple, powerful, and more accessible tool for the construction and manipulation of multimodal graphs. Web-based audio functionality is now at the point where it can be leveraged to provide just such a tool. Thus, we developed a web application, the Web Sonification Sandbox (or simply the Web Sandbox), that allows users to create and manipulate multimodal graphs that convey information through both sonification and visualization. The Web Sandbox is designed to be usable by individuals with no technical or musical expertise, which separates it from existing software. The easy-to-use nature of the Web Sandbox, combined with its multimodal nature, allow it to be a maximally accessible application by a diverse audience of users. Nevertheless, the application is also powerful and flexible enough to support advanced users.

1. INTRODUCTION

Presentation of data using a combination of visual and auditory modalities (i.e., multimodal or auditory graphs) can allow for discoveries in a data set that are sometimes impossible with visual-only inspection. At the same time, multimodal graphs can make data, and the STEM fields that rely on them, more accessible to a much broader range of people, including users who have a range of disabilities, stemming from, but not limited to, vision impairment and other reading difficulties. There have been a variety of software tools developed to turn data into sound, including the widely-used Sonification Sandbox[1] [11, 4], but there remains a need for a simple, powerful, and more accessible tool for the construction and manipulation of multimodal graphs. Here we describe the development of exactly that, namely the Web Sonification Sandbox (aka, Web Sandbox[2]).

1.1 Auditory and Multimodal Graphs

Typical visual graphs have long been used by teachers at all levels to convey information to students, whether they come in the form of bar charts, plots of equations, scatter plots, or histograms. Just as appealing to students' visual intelligence through the use of visual graphs can help teach students a new concept, auditory graphs have been shown to be effective at doing so as well [12]. An auditory graph can be defined as "a class of auditory displays that are produced when sounds are mapped to quantitative data" [9]. This use of sound can stand alone, or can be used in conjunction with a visual graph, to form a multimodal display of data. Auditory graphs have been shown to have the capacity to be helpful to both sighted and visually impaired individuals [13], and can help reach students with different learning styles.

Another way in which the multimodal display is useful is in scenarios when a user may face information overload. Since most data comes in the form of visual information, professionals such as stock market analysts or data scientists may easily become overloaded by large amounts of visual informations. It has been shown that using non-speech audio instead of a new layer of visual information can reduce information overload [2].

Also, due to the very nature of multimodal graphs, one could learn or discover new skills and insights by relating the two modalities. If given appropriate tools, even an individual with little to no experience with music or sound synthesis could use multimodal data presentations to learn about patterns in the sounds by relating them to familiar

*Please direct all correspondence regarding this document to Zachary Kondak.

[1]http://sonify.psych.gatech.edu/research/sonification_sandbox/index.html
[2]http://sonify.psych.gatech.edu/research/web_sandbox/index.html

Web Audio Conference WAC-2017, August 21–23, 2017, London, UK.

concepts of visual graphs.

1.2 Auditory Graphing Tools

Given that there are many compelling reasons to support the development of tools to turn data into sound, it is not surprising that several sonification and auditory graph applications do currently exist. An exhaustive listing of related software is beyond our scope here. However, we can say that some software tools have specifically been designed to be used by artists or musicians, such as the DataToMusic API[3] [10]. These are not always generalizable to a range of STEM and educational purposes. Other software is intended to be used in a more curated or corporate context, such as the Taranos:Cloud Sonification Framework[4], which is aimed somewhat more at developers with coding experience and knowledge. Some software applications in this space have limitations in terms of accessibility to a broad range of users.

Indeed, the original Sonification Sandbox, previously mentioned, was implemented as a Java Application, which makes it less accessible to users with vision impairment, and more difficult to maintain and update the code. While this Java-based Sonification Sandbox (aka, Java Sandbox) supports an extensive range of features, moving toward a web platform would allow for a lightweight application with greater portability and flexibility to run on a variety of devices and browsers. Users would be able to use it without downloading or installing any software, which we feel will enable individuals who tend to shy away from downloading software to experiment with using it. Whereas the full Java Sandbox is feature-rich and powerful, it remains more complex than necessary for many users. We wanted to provide an interface that can be powerful, but also abstracts many of the difficulties and nuances associated with designing effective auditory graphs away from the user (and allow them to further explore those features if they wish).

2. WEB SONIFICATION SANDBOX

The Web Sandbox was designed with four main functional goals in mind and these goals have guided the design process at all levels. The goals are: maximize accessibility, maintain a high ceiling, maximize portability, and maximize utility.

Figures 1, 2, and 3 show screen shots of the Web Sandbox to demonstrate the simplicity of the web-based interface.

2.1 Maximize Accessibility

Accessibility for all individuals is a central focus of this project. Perhaps most importantly, the Web Sandbox should be completely accessible to individuals with vision impairment. Another design goal was for the Web Sandbox to be usable by students and adults of differing educational levels. As such, we decided to allow for graphs to be sonified out of the box using various default sound generators. This means that a user does not need any prior experience with sound design in order to use the Web Sandbox. The user interface (UI) of the Web Sandbox supports various intuitive (and learned) keyboard controls that allow the user to interact with a given auditory graph (e.g. one can explore individual points on the graph by using the arrow keys). Finally, the Web Sandbox allows for auditory graphs to be

[3]https://dtmdemo.herokuapp.com/
[4]https://github.com/taranos/taranoscsf

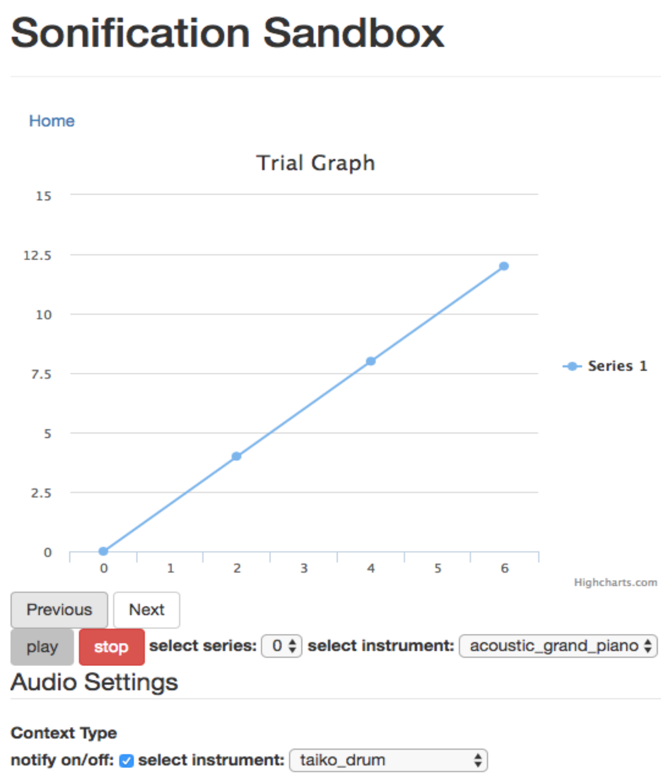

Figure 1: Sonify interface with a simple line graph.

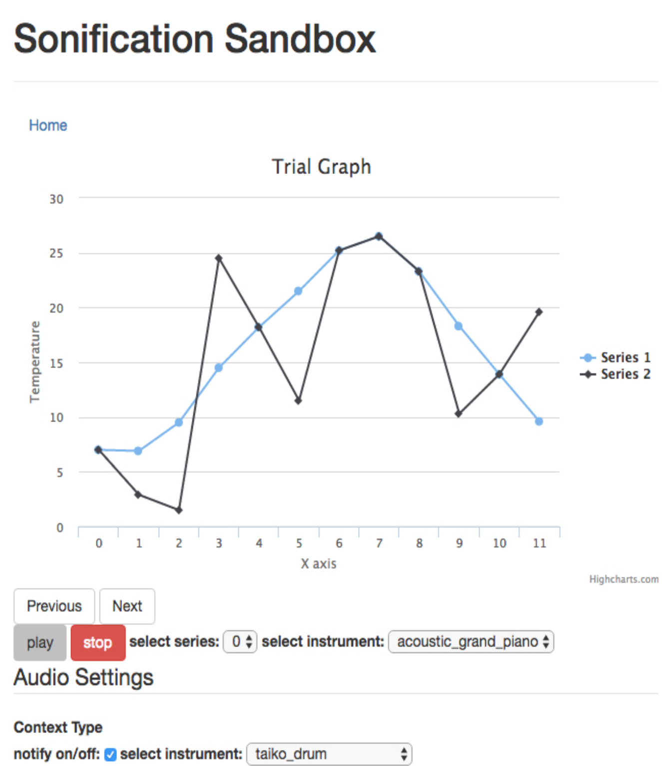

Figure 2: Sonify interface with visual representation of a graph with two series.

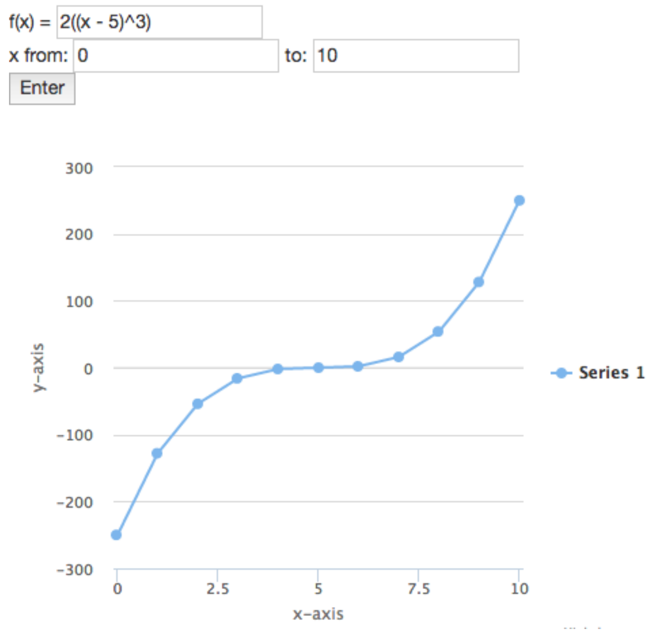

Figure 3: Interface for plotting equations.

generated by equations that are input in standard calculator format. This means that the user does not need to learn any special code or scripting to input equations if she already knows how to do so on any standard calculator.

2.2 Maintain a High Ceiling

Building on the idea of accessibility and usability for individuals from all educational backgrounds, we wanted to minimize the amount of time it takes to learn how to use the Web Sandbox. This does not mean that the Web Sandbox is not powerful; rather, we want it to represent a LTHC (low threshold, high ceiling) platform. Both beginners and trained professionals should find the sandbox immediately useful [8]. For instance, along with the built-in instruments that allow individuals without sound design experience to use the Web Sandbox, we also support a synthesizer, which allows the more experienced users to tweak the available synthesis parameters to their liking. Furthermore, the power user can also take advantage of the fact that she can compare any number of data series within a single auditory graph.

2.3 Maximize Portability

The Web Sandbox is designed for maximum portability in two ways. First, we designed the web app to support users, meaning that a user's data can be saved on the server side for retrieval from any other location. As far as we know, there is no other sonification tool that is structured in this fashion. Second, we have also implemented an import/export feature that allows users to import data from any major spreadsheet program into the Web Sandbox as well as export data from the Web Sandbox to be read by another program or another Web Sandbox user. This supports easy and functional connections to data analysis processes which users would already be employing.

2.4 Maximize Utility

A clear goal of the Web Sandbox is to provide effective auditory graphs that actually convey information well. One way we ensure that auditory graphs are useful is to provide the user with a wide array of sound options for each series in a given auditory graph. This allows users to differentiate between the multiple series more easily due to variations in timbre [6]. Another important feature that we added to increase the Web Sandbox's utility is the addition of auditory context cues. Walker and Nees showed how the relative loudness of auditory context information correlated with individuals' effective understanding of an auditory graph [9]. Based on this research, we decided to include various auditory context options in the Web Sandbox. For example, one such auditory context cue consists of playing a percussive sound at an interval that represents 1 on the x-axis of the auditory graph. This context helps users keep track of where data points are located with respect to one another on the x-axis. Future implementations of the Web Sandbox will include context cues such as letting the user play a separate stream of audio representing the local or global minimum or maximum, as well as other easily-calculated descriptive statistics, like mean, median, or the mode of a given graph. Context together with visualizations of graphs provide and cement understanding of a dataset in a meaningful way [1].

3. TECHNICAL DETAILS

As stated earlier, the Web Sandbox is designed with platform independence in mind. As such, the technologies that the Web Sandbox relies upon were chosen based on their platform independence and flexibility.

3.1 Basic Architecture

The Web Sandbox is a MEAN application. MEAN stands for MongoDB, Express.js, AngularJS, and Node.js. MEAN is a stack of applications that serve as a framework on which one can build a web application. We chose to use the MEAN stack because of its flexibility and relatively high popularity. A large portion of the web development community is interested in utilizing these newer web technologies, and each of them have a broad base of support. We felt that this approach would allow us to more easily launch development while also providing a framework that could successfully support the app after it started to become more complex and complete.

3.2 Database

The database is implemented using MongoDB, a flexible NoSQL database architecture that is easily scalable. MongoDB uses a data model similar to JSON (JavaScript Object Notation), allowing for document structure to be more flexible and reducing the amount of database changes necessary compared to the more complex development processes needed for traditional relational databases. At the top level, the basic architecture of the Web Sandbox's database is a collection of users. Within each level, every user entry contains a collection of graphs. Finally each graph contains all data and sound information about a given auditory graph. Figure 4 illustrates the database structure of the Web Sandbox. The graph information structure present at this lowest level of the database functions as the model in the MVC (model view controller) architecture of the app.

3.3 Visual Graphing

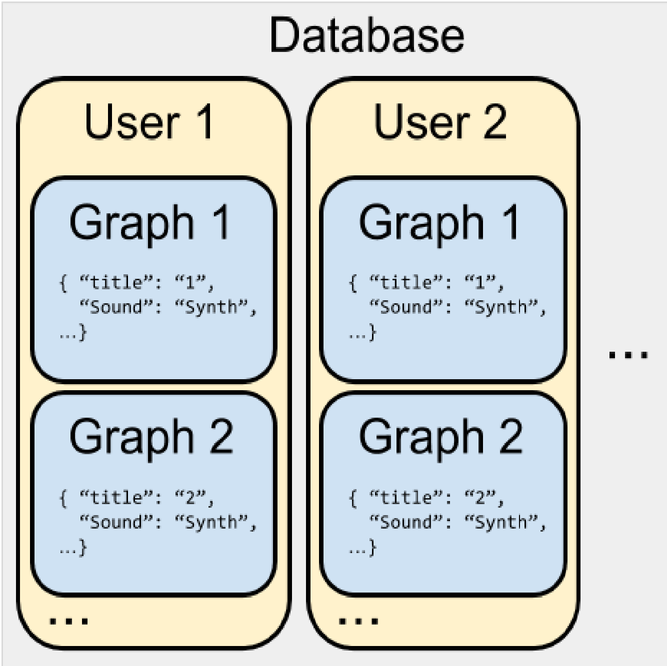

Figure 4: Graphical representation of the database structure.

It was necessary for us to display the data in a visual format along with the audio since individuals may choose to use the Web Sandbox to examine data through the modalities of sight and sound simultaneously. In order to accomplish this, we decided to use a JavaScript visualization library called Highcharts[5]. We chose to use Highcharts because of its simplicity and its large and active community. Highcharts can flexibly load data from CSV files (excel) and JSON data, making it integrate perfectly with our database structure. It also allows for the highlighting of specific points on a graph, which is a vastly important feature when playing the auditory graph; helping users visually track where in the graph they are at all times during the progression of the auditory graph's playback.

3.4 Plotting Equations

Along with the ability to create auditory graphs via data input, users also have the option to create auditory graphs via mathematical equations. In order to implement the equation plotting feature, we relied on a JavaScript library called Math.js which integrates with Node.js server[6]. We decided to use this library because it supports a wide array of mathematical functions (polynomial, trigonometric, exponential, etc.). When a user enters a mathematical equation, the Sandbox will automatically generate a set of data points and display the plot. After generating the plot, the user is able to sonify the new data immediately. The user can then save the new plot into the database with the equation name as the title, to support easy retrieval of these data. Once saved into the database, the graph can be edited and viewed in the exact same manner as any other tabular data can be within the Web Sandbox. Due to this, equation data and imported data can easily be merged into a single graph, facilitating the comparison between some real world data and a mathematical equation. Users could also compare within families of equations ($y = x^2$ and $y = 4x^2$), or between different families ($y = x^2$ and $y = x^3$).

3.5 Program Architecture

The program is implemented with the MVC architecture in mind. As stated above, the graph data stored in the database acts as the model for the Web Sandbox's main sonification functionalities. There are three main ways to view the graph model: one is to view from the plot view (where a user can create new graphs by plotting equations); another is to view the graph model in a tabular format in the edit menu; and the last is to view the graph model from the main sonification menu. Each of these views has a corresponding controller code, which handles all updates and changes to the graph data.

3.6 Sound Synthesis and Data Mapping

In the current state, the Web Sandbox has a set mapping from data to sound. High valued data points are mapped to higher pitched sounds. The minimum and maximum sonification frequencies are currently arbitrarily hard-coded to D3 (146.8Hz) and F#6 (1480.0Hz) respectively. These minimum and maximum frequencies are mapped to the minimum and maximum data-points contained within a given graph. All data points between the minimum and maximum valued are are logarithmically mapped to the frequency range. Each graph can have any number of series, which can each be assigned a different sound source.

The Web Sandbox is able to produce sound via two separate mechanisms. The first mechanism is through the use of soundfonts. We used a JavaScript library called soundfont-player[7], which uses the WebAudio API[8] to produce sounds specified via the soundfont standard. Soundfont-player supports both JSON and MIDI formats, without the extra bulk of the MIDI.js library. Soundfonts provide the user with a variety of ready-made instruments (like pianos and organs) that are widely known and cover a diverse range of timbres. However, there are several logistical complications to using soundfonts to sonify data. Using soundfonts means that you must play a discrete sound to sonify each individual point on the graph. The usage of soundfonts also restricts the flexibility the user has in modifying the sound to deciding which soundfont to choose. While this lack of options for sound modification options is good for supporting a low barrier of entry, it also prevents individuals with sound design experience from having fine-grained control over the sonification output. Due to this restriction, we decided to implement an additional, even more flexible mechanism for producing sound.

The second sonification method generates the sounds using classic subtractive synthesis techniques. This method allows graphs to to be sonified in a continuous fashion, so that rather than playing a distinct sound for each point on a graph, a continuous sound can be modulated smoothly by interpolating between each discrete point on the graph. In addition to providing more control over the synthesis, this synthesizer creates a more accurate representation of a line

[5]http://www.highcharts.com/products/highcharts
[6]http://mathjs.org/
[7]https://github.com/danigb/soundfont-player
[8]https://www.w3.org/TR/webaudio/

as an infinite set of points (an important math concept to support for students with vision impairment who are learning about lines and graphing in school). The user has control over various parameters of the synth. The synth consists for four oscillators, which can each be either a sine, square, saw, or triangle waveform, and can be independently detuned by the users. The synth also incorporates a basic low pass filter with cutoff control. One potential issue of using the synth is that it requires the user to have some basic knowledge of subtractive synthesis, but this could also be used to help introduce those concepts to users as well. The synthesizer was implemented entirely using only the WebAudio API's native features.

4. DISCUSSION

There are many scenarios in which we believe the Web Sandbox can be put to effective use. For instance, individuals with vision impairments could use the tool to help elucidate features of data and equations that are not easily or efficiently transmissible by text description or tactile information. Alternatively, the Web Sandbox can also be useful to those without vision impairments. For instance, a stock market analyst could use the tool to sonify some data that they cannot fit into their already busy field of view [7]. Since, as stated in the introduction, the Web Sandbox is not explicitly designed for musicians, it would be entirely feasible for those lacking a background in music or sound design to us the Web Sandbox as a tool for aesthetic exploration. These individuals could use the Web Sandbox for an artistic project like a musical composition or sound installation since using it requires no musical knowledge. These are just a few examples of the many cases where the Web Sonification Sandbox is useful.

5. CONCLUSION

Many advancements in web technology are centered around visual and textual information; supporting sonification and auditory displays supports inclusiveness and access to these new technologies to individuals with vision impairments. With the rapid advancements in web audio technology, researchers and developers can now invent systems and paradigms, like the Web Sandbox, to enrich the web experience for individuals with and without vision impairments through the use of audio.

6. FUTURE WORK

We are currently beginning to evaluate the Web Sandbox's usability and accessibility with a diverse group of users, to iteratively refine and further enhance the intuitive user interface. During these evaluations, we will test how quickly people can learn to use the Web Sandbox to perform simple baseline tasks such as sonifying a dataset and making observations based on the created graphs. These studies will include user experience surveys like the System Usability Scale (SUS) [3] and the Usability Metric for User Experience (UMUX) [5], and recording the number of errors when interacting with the Web Sandbox. Note that since the Web Sandbox is built with best practices in the field of sonification in mind, we feel that it is unnecessary to experimentally validate the sonification design on the whole.

We plan to expand the Web Sandbox's feature set to include additional functionality: ability to export and embed sonifications into other media (i.e. web pages), ability to intuitively sonify higher-dimensional data (i.e. images), expanding the synthesizer's options, more flexibility and control over the data to sound mapping (i.e. map data values to individual synthesis parameters in place of pitch alone), and support for additional accessibility features to support a diverse group of users (e.g., FLOE[9]).

7. REFERENCES

[1] F. J. Anscombe. Graphs in statistical analysis. *The American Statistician*, 27(1):17–21, 1973.

[2] S. A. Brewster. Using non-speech sound to overcome information overload. *Displays*, 17(3-4):179–189, 1997.

[3] J. Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.

[4] B. K. Davison and B. N. Walker. Sonification sandbox overhaul: Software standard for auditory graphs. pages 509–512. Proceedings of the 13th International Conference on Auditory Displays (ICAD 2007).

[5] K. Finstad. The usability metric for user experience. *Interacting with Computers*, 22(5):323–327, 2010.

[6] P. Iverson. Auditory stream segregation by musical timbre: effects of static and dynamic acoustic attributes. *Journal of Experimental Psychology: Human Perception and Performance*, 21(4):751, 1995.

[7] B. S. Mauney and B. N. Walker. Creating functional and livable soundscapes for peripheral monitoring of dynamic data. Proceedings of the 10th International Conference on the Auditory Display (ICAD 2004), July 2004.

[8] B. Myers, S. E. Hudson, and R. Pausch. Past, present, and future of user interface software tools. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(1):3–28, 2000.

[9] M. A. Nees and B. N. Walker. Relative intensity of auditory context for auditory graph design. pages 95–98. Proceedings of the 12th International Conference on Auditory Display (ICAD 2006), June 2006.

[10] T. Tsuchiya, J. Freeman, and L. W. Lerner. Data-driven live coding with datatomusic api. Web Audio Conference (WAC-2016), April 2016.

[11] B. N. Walker and J. T. Cothran. Sonification sandbox: A graphical toolkit for auditory graphs. pages 161–163. Proceedings of the Ninth International Conference on Auditory Display (ICAD 2003), July 2003.

[12] B. N. Walker, G. Kramer, and D. M. Lane. Psychophysical scaling of sonification mappings. Georgia Institute of Technology, 2000.

[13] B. N. Walker and M. A. Nees. An agenda for research and development of multimodal graphs. pages 428–432. Proceedings of the First International Symposium on Auditory Graphs, in the Proceedings of International Conference on Auditory Display (ICAD 2005), July 2005.

[9] http://floeproject.org/