Consumption as a Rhythm: A Multimodal Experiment on the Representation of Time-Series

Catarina Maçãs*, Pedro Martins*, Penousal Machado
CISUC, Department of Informatics Engineering
University of Coimbra
Coimbra, Portugal
{*cmacas, pjmm, machado*}*@dei.uc.pt*

***Abstract*—Through Data Visualisation and Sonification models, we present a study of multimodal representations to characterise the Portuguese consumption patterns, which were gathered from Portuguese hypermarkets and supermarkets over the course of two years. We focus on the rhythmic nature of the data to create and discuss audio and visual representations that highlight disruptions and sudden changes in the normal consumption patterns. For this study, we present two distinct visual and audio representations and discuss their strengths and limitations.**

***Keywords*-Sonification; Visualisation, Multimodal representation, Time-series, Consumption rhythm.**

I. INTRODUCTION

Complex time-series are often represented through visualisation and sonification. While visualisation tends to be regarded as a more effective representation for most kinds of data, sonification is frequently presented as a suitable representation of time-varying data, as it provides two dimensions for the representation: the sound itself and the idea of time [1]. Additionally, the presence of regular repeated patterns favours the construction of rhythmical musical representations, which tend to be more pleasant to the listener.

We present a study on the representation of time-series using different sensory modalities. Our starting point is a collection of consumption data from a Portuguese retail chain and a previous work of ours on the same data [2]. The aim of this project is to create a representation that arouses the users' curiosity. These users are not the company's analysts, as such, it must be simple and compelling. We focus on the rhythmic nature of consumption to create visualisation and sonification models that can be further combined together to produce appropriate multimodal representations. This study is extended to include a discussion on the relationship between visualisation and sonification, namely when it comes to revealing different aspects of the data and to the understanding of some of the limitations associated with the proposed representations.

* Both authors contributed equally for this work.

II. MULTIMODAL REPRESENTATIONS OF TIME-VARYING DATA

Throughout the years, people enhanced their cognitive abilities, such as memory, thought, and reasoning, with the invention of external aids. One of the oldest and most decisive external aid is graphical representation [3]. Over time, the use of graphics for representing knowledge has revealed itself as an important and effective means for communicating quantitative information. Allied to our ability to perceive geometrical patterns, it enables and facilitates the detection of patterns and relationships in data [4].

Time-series can be visualised in several ways, as long as time is represented. Among the earliest graphical representations of time-varying data is an anonymous multiple time-series graph, from the 10^{th} century, of the changing position of the seven most prominent heavenly bodies over time and space [5].

Sonification is the practice of turning data into sounds. There are various sonification techniques, which can be broken down into the following categories: auditory icons, earcons, audification, parameter mapping sonification, and model-based sonification [6]. For this project, we focus on parameter mapping sonfication, in which data is responsible for varying different parameters of an audio signal [7].

An example of a sonification of time-varying data is *Quotidian Record*, a project created by Brian House, presented in 2012. In *Quotidian Record*, House represents through sound each place he had visited over one year. The basis for his project was the idea that our routines have inherent musical qualities and through music we are able to create an emergent portrait of each individual [8]. Other examples of sonification of time-varying data include *Climate Symphony* [9], a sonification of climate change based on the analysis of the chemical composition of an ice core drilled up in Greenland, or Living Symphonies [10], a musical installation that portrays the activity of the forest's wildlife, plants and atmospheric conditions.

In this section, we present a collection of projects that use both sound and visuals to represent time-series.

Two Trains
In 2015, Brian Foo created a musical sonification of income inequality on the New York City Subway's 2 Seventh Avenue Express, whose route includes three different boroughs: Brooklyn, Manhattan and the Bronx [11]. Foo's sonification is modelled after Steve Reich's *New York Counterpoint*, a minimalistic composition written for 11 clarinets and a bass clarinet that tries to capture the vibrant atmosphere of Manhattan. The sonification emulates a ride on the aforementioned train. At each moment, the quantity and dynamics of the instruments reflect the median household income of that area. For example, in wealthier areas, the instruments will increase in quantity, volume, and force. In addition to the sonification, Foo presents a simple visualisation of the train route, enabling the viewer to connect the sound to the respective geographic area. Foo's main intent in this project was to create a representation of the vibrant energy and orderly chaos of the New York City Subway system.

Listen to Wikipedia
Listen to Wikipedia is a project created by Stephen LaPorte and Mahmoud Hashemi, in 2013, to represent the most recent feed changes in Wikipedia articles [12]. These changes are represented by a bell if some entry is added, and by string plucks if someone removes an entry. The pitch is also manipulated depending on the size of the edit. In this project, the visualisation was generated to not overshadow the sonification. To represent visually the editions, and to add more insight about the data, the authors draw green circles to show edits made from registered contributors, white circles to represent unregistered users, and purple circles to represent edits made by automated bots.

Deep Into Animal Extinction
In 2014, a group from the IUAV University Communication Laboratory generated a sonification to represent how many animal species entered the IUCN Red List of endangered species. Each class (mammals, birds, amphibians, fishes, and reptiles) is represented by a note of a B minor chord and the number of species for the corresponding class is represented by the number of repetitions of every note [13].

Specimen Box
In 2014, the Office for Creative Research group created the Specimen Box, an interactive work that visualises and sonifies the Botnet activity. Botnets are distributed entities infected with malware that reach thousands of PCs. In this project, those networks are represented by their aggregated geographic position, creating visual and audible temporal patterns. The Office for Creative Research characterises this project as an exploratory tool that enables digital crime units' investigators to examine the different profiles of the Botnets through geographic position and time, and to understand their unique characteristics: their behaviour, how they propagate through PCs, and how they are adapting to the environment [14].

III. DATA

The data set for this project is concerned with the consumption values in 729 Portuguese supermarkets and hypermarkets of SONAE's chains, which cover the entire country. When shopping in these chains, the customers tend to use their loyalty cards to accumulate discounts and other benefits. Currently, the number of active cards is above 6 million, which can be considered an impressive figure, especially if we take into consideration that the Portuguese population is below 11 million (2011 census), and that the cards are issued by "household" and shared by the entire family. We had access to this data through a collaboration with SONAE. SONAE's aim with this project is to go beyond analytical visualisations, and explore more engaging and entertaining visualisations that comply SONAE's consumers to further analyse and explore the company's mark. This data set is relevant due to its richness, size, quality, and nature. We believe that the data set is a valuable asset of the work, allowing us to transform the Portuguese consumption patterns into several aesthetic artefacts.

We analysed all the transactions made on these supermarkets and hypermarkets from May 2012 to April 2014. Each transaction corresponds to one product bought and has attributes such as price, date, and time. Each product is placed in the product hierarchy of the company, which has six levels. For this work, we aggregate all the purchases in nine distinct categories: Grocery; Alcohol & Sweets; Health care; Beauty; Clothes; Furniture; House Care; Culture & Leisure; Pets & Nature Care. Additionally, we grouped those nine categories into groups of three by their type of consumption. The three types of consumption are defined as follows: essential (Grocery, Health Care, and Clothes); non-essential (Alcohol & Sweets, Beauty, and Culture & Leisure); and unknown (Furniture, House Care, and Pets & Nature Care).

IV. REPRESENTATION OF DATA

This project is based on the necessity to add a second level of information to a visualisation that emphasized the rhythm of shopping habits [2]. The intrinsic sense of rhythm in the consumption data and the necessity to explore it beyond visuality is the basis for the sonification and, therefore, for this project. The rhythmic pace of the dataset is caused by the periodic growth and decrease of consumptions during the week. Our main goal for this project is to demonstrate how the consumption in the different categories and types of consumption change over time. To accomplish it, we present two different approaches (which can be accessed through the following links: https://vimeo.com/270077726

and https://vimeo.com/270078256). In this section, we define the main concept behind each visual representation and how each attribute from our dataset is sonified.

The two percussive sonifications presented herein are not intended to highlight the rhythmical nature of the data. Instead, they assume the presence of such feature and exploit it to primarily emphasise the sudden changes and disruptions in consumption patterns. In terms of instruments, both sonifications are built around Taiko drums, a wide range of traditional Japanese drums, whose existence is speculated to date back from Ancient Japan. Taiko drums have served many purposes: for example, they were a common element in Buddhist and Shinto rituals or in the battlefield either to alarm the enemies or to issue commands [15].

Each type of Taiko drum has its unique timbre and pitches. Despite the differences, the sound of these instruments can be characterised as powerful and dramatic, especially when bigger drums are used, such as the odaiko or the chu-daiko. We believe that these characteristics make Taiko drums an interesting basis for the sonification of consumption data gathered from hypermarket and supermarket chains. First, and as previously mentioned, the consumption patterns tend to present a cyclic behaviour over time, i.e., there is a regular repeated pattern, which favours the idea of a rhythmic representation. Secondly, sudden changes and disruptions can be signalled with intensity changes, which, due to the stirring sound of this type of percussion, are expected to effectively inform the listener about the significant data modifications. Finally, the vast variety of Taiko drums allows us to assign different types of consumption, or even categories, to the various groups of drums. Note that our rhythmic compositions do not strictly follow rules of Taiko drumming, as the main idea is to explore the characteristics of the timbre rather than using typical Taiko-based compositions.

The proposed sonifications follow a parameter mapping model, with the volume (gain) of the percussion instruments reflecting the numbers of transactions within the different types of consumption. For each type of consumption, we normalise the number of transactions to let the values vary between 0 and 127. Those values are then rounded to integers to set the volume of the drum beat. In both sonifications, a beat represents a day.

As our goal is to explore a multimodal representation, we tried to introduce a certain complementarity between the visual representations and the auditory displays. Both sonifications models disregard the consumption data at the category level. Furthermore, according to our model, not all types of consumption can be considered per beat. We followed this strategy not only to introduce complementarity between the different representations but also to study how informative or evocative this type of rhythmic sonifications can be.

The sonification process was developed in Max/MSP. A Max patcher is responsible for reading the data and generating MIDI notes according to the sonification model. The MIDI notes are then sent to Ableton Live. The instruments are played through Ableton Live using the AIR Xpand!2 VST plug-in, which includes a wide range of percussion instruments. It should be noted that the plug-in does not use the typical Japanese terminology to classify the different type of Taiko drums it provides. We use three types of Taiko drums: Taiko Big drums 1 and Taiko Big drums 2 (similar to chu-daiko), and Taiko drums (similar to okedo-daiko).

A. First Approach

For the first visual representation (Fig.1), our main goal was to represent the different categories and to visually perceive their behaviour over the two years. To explore the data, we created two different approaches. The first one was a classical small multiples representation, where each day is placed in a grid, horizontally organised by day of the week and vertically by month. The second approach was an animation of each cell of the grid previously described. For this project, we only used the second approach, as our intent is to assemble the visual representation with the sonification. Hereupon, we will only describe the latter.

For the representation of each category, and since their values were so sparse (i.e., some categories had a range of values too small, when comparing with other categories), we decided to normalise each consumption value by category independently. This enabled us to represent the peaks of consumption in each category, instead of enabling the comparison between them.

Each category is represented by a different circular shape, which is coloured depending on the type of consumption it represents: red, for non-essential, green for essential, and grey for unknown (Fig. 2). To represent the differences in consumption over time, we used the size of each shape. As the days go by, each shape grows or decreases in size depending on the consumption value. Every shape is centred in the middle of the canvas and morphs its size between days, so we can have a more continuous perception of each category consumption behaviour.

To enable the viewer to analyse the consumption behaviour over time, we added the date in the upper left corner of the visualisation. The user will not be able to analyse the consumption volumes of each day precisely, but will be able to compare the consumption behaviour in different times of the year. Furthermore, as the consumption rhythm is mainly caused by the differences in consumption between weekdays and weekends, in the upper side of the canvas, we represent visually the day of the week. This representation is composed of a row with the length of the canvas divided in seven equal parts. A rectangle is then positioned in one of the seven possible positions, accordingly to the current week day (the first position of this row represents Mondays, the second position represents Thursdays, and so forth).

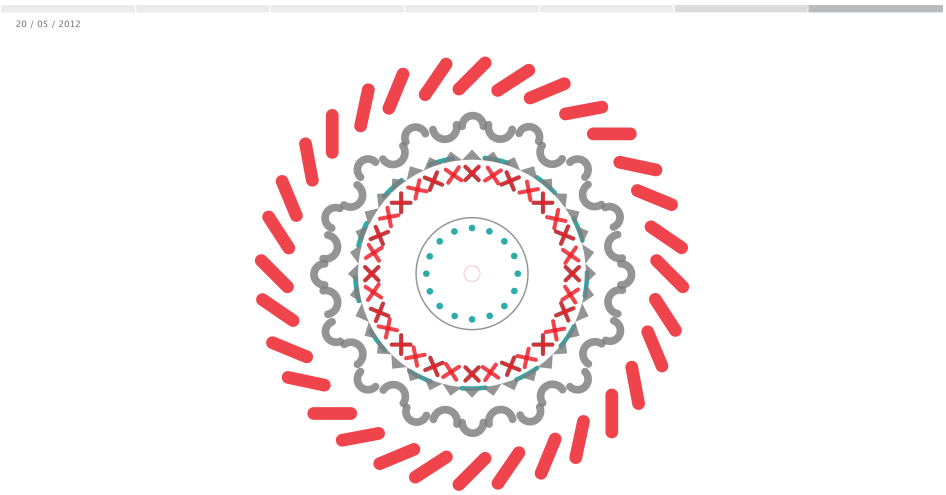

Figure 1: Screenshot of the first visualisation layout

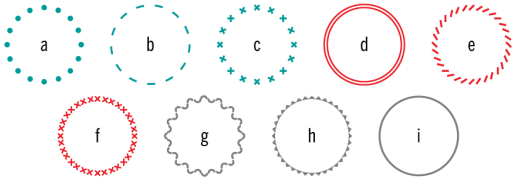

Figure 2: Representation of each category (shape) and their types (colour): (a) Clothes, essential; (b) Health Care, essential; (c) Grocery, essential; (d) Culture & Leisure, non-essential; (e) Beauty, non-essential; (f) Alcohol & Sweets, non-essential; (g) Pets & Nature Care, unknown; (h) House Care, unknown; (i) Furniture, unknown.

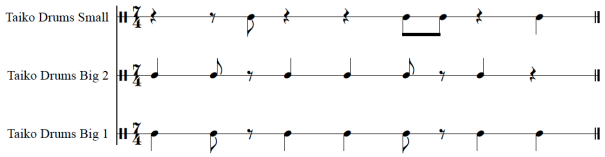

Figure 3: Basic Taiko drums beat for Representation #1

The first sonification has a time signature of $\frac{7}{4}$ and, as previously mentioned, each pulse represents a day of the week. Figure 3 presents the basic beat defined for this sonification: Taiko Drums Big 1 corresponds to essential goods, Taiko Drums Big 2 corresponds to non-essential goods, whereas the transactions from the remaining categories are associated with an ensemble of four Taiko drums.

As Fig. 3 suggests, this sonification focuses on the transactions of essential and non-essential goods. To accentuate even more the changes in consumption patterns, we added other instruments to the sonification: the transactions of essential goods are also associated with thunder drums; a orchestral snare drum is used for transactions of non-essential goods. The remaining goods are not associated with any other instrument.

In this approach, our intent was to use sonification to represent a different level of the data (the sonification of the types of consumption, instead of the sonification of the categories, as in the visualisation), enhancing the understanding of the consumption behaviours. As the sonification only sonifies consumption by type and the visualisation only represents consumption by category, it is difficult to perceive the relation between the two representations. This creates some discomfort when attempting to create connections between sound and image. Nonetheless, when all the values of the three categories from one consumption type rise, the relation between sound and image is better understood.

To improve the readability of the multimodal representation, we implemented a second approach, in which we represent the same level of aggregation.

B. Second Approach

For the second visual representation (Fig. 4), our main goal was to highlight the differences between types of consumption. As previously stated, in the previous approach it was not possible to compare the consumptions among themselves, as the values are normalised by category independently. As such, for this representation we summed all the categories in each type of consumption and normalised those values by the maximum value of all three types.

To distinguish the different types of consumption, and as in the first representation, we used three colours: green, red and grey to represent, respectively, the essential, non-essential and unknown consumptions. Each type is represented by a coloured circle. The size of the circle varies depending on the consumption value of each type of consump-

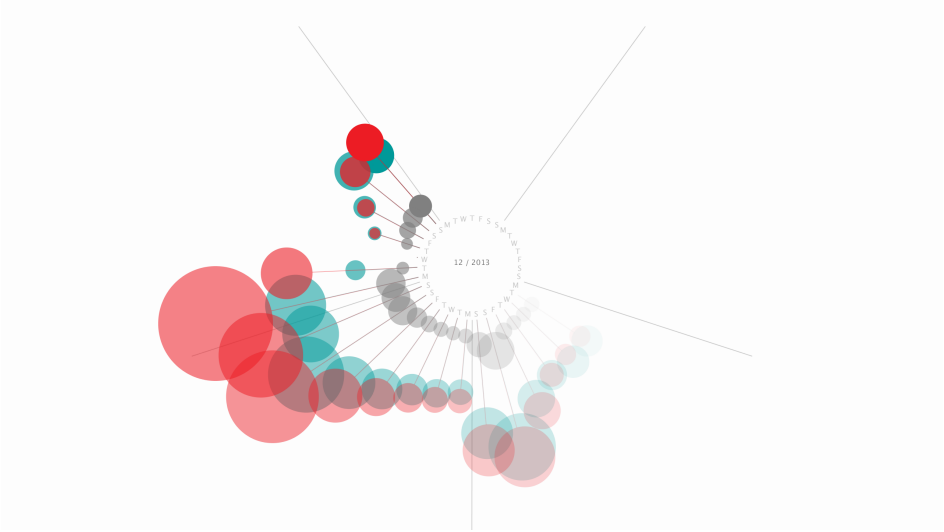

Figure 4: Screenshot of the second visualisation layout

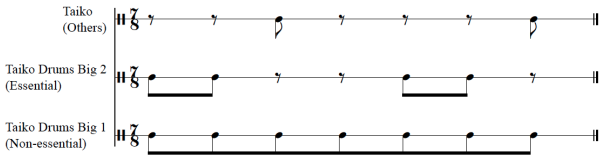

Figure 5: Basic Taiko drums beat for Representation #2

tion at a given day. To emphasise the different consumption values, we also distance each circle from the centre of the canvas. The lower the summed consumption value, the closer the circle is to the centre of the canvas.

For this representation, and apart from the use of size and colour to distinguish the types of consumption, we wanted to explore the movement and positioning of elements to augment the information represented. As it was difficult to distinguish days in the previous approach, we wanted to make it more legible for the second one. Hence, as time passes, the circles change position in a invisible circle. Each lap represents a month, and we visually divided the lap in 5 slices (as can be seen in 4), representing 5 distinct weeks. At the beginning of each month, the circles start in the upper side of the circle (in the first slice), and as time passes, they circle around the canvas. Each slice is further divided in the seven days of the week, so the position in which the month begins is dependent on the day of the week of the first day of the month. To facilitate the understanding of which month is, we placed a legend in the interior of the circle. To better associate each circle to the respective day of the week, a line is also drawn connecting the circle and the legend. Finally, as time passes, we maintain the representation of some of the previous days, with decreasing transparency, to create a sense of movement and to be able to view some of the previous days, giving some context over the sonification.

The second sonification, which was designed to focus on the transactions of non-essential and other goods, has a time signature of $\frac{7}{8}$. The drums beat for this sonification is depicted in Fig. 5: Taiko Drums Big 1 corresponds to non-essential goods, Taiko Drums Big 2 corresponds to essential goods, and the transactions from the remaining categories are associated with a group of two Taiko drums.

As in the previous sonification, other percussion instruments were added to the model: essential goods are associated with timpani, non-essential goods are associated with thunder-drums, other goods are associated with chimes, which are played at every beat. Chimes were introduced to represent the consumption patterns in terms of pitch, using the same strategy as the one that we defined for the gain of the instruments.

With this second approach, we could enhance the growth of consumptions between the three types of consumption. As the data used for the visualisation and sonification where the same, now the "reader" can more easily establish relationships between the representations and be more attentive.

V. Discussion

The multimodal representations described herein stress, first and foremost, the rhythmic pace of the data. The visual artefacts can arguably be regarded as more informative, as they translate consumption to a finer detail. Nevertheless, the proposed multimodality was created with the intent of becoming a richer experience in terms of understanding. The auditory counterpart contributes to a better comprehension of consumption peaks over time. While the sonification model was designed to not consider all types per pulse, the final representations become sufficiently informative as the result of complementing sound with the visual artefacts.

This study suggests that percussive sonifications are an interesting approach to represent the type of data that we worked with. The percussion can portray the hustle and recurrence associated with shopping habits. However, it also has some limitations, especially in terms of informativeness, if the data is sufficiently complex. A possible solution to circumvent this problem is to use pitched percussive instruments and consider pitch as a parameter to be mapped.

VI. CONCLUSIONS

The creation of visual models that are able to repeatedly readjust to changes of dynamic data is a major and crucial challenges in Data visualisation and Sonification.

In this project, we focused on the representation of time-varying data through sound and image. Our data set consisted of the the consumption in 729 Portuguese supermarkets and hypermarkets of SONAE's chains. Knowing in advance that this data has an intrinsic rhythm associated with the differences in consumption during the weeks, the main goal for this project was to represent the variation on consumptions over time through visualisation and sonification. We aggregated the data in nine categories that were subsequently grouped into three types of consumption. To represent visually this data, we presented two visualisations, one more concerned with legibility, and another one, which was more exploratory, concerned with the representation of information through movement. As the first visualisation only visualises the consumptions in the nine categories, the second visualisation, in addition to the representation of the categories, also represents their values summed by type of consumption.

To better emphasise sudden changes or disruptions in consumption patterns, we complemented the visual representations with two percussive sonifications, which were built around Taiko drums. Both sonifications adopted a parameter mapping model, with the volume of the percussion instruments reflecting the numbers of transactions within the different types of consumption. Pitch was also considered as a parameter in the second sonification.

As for future work, we believe that there are several promising directions to be explored. At the multimodal level, we aim to further study the benefits and disadvantages of representations that are based on strongly complementary visualisations and sonification (in terms of information). If we restrict ourselves to the sonification part, there are several possibilities that deserve to be explored. A sonification model that allows filtering or emphasising certain types is something worth to be explored. To make the sonification more informative, we can also consider the use of pitched percussion instruments. Finally, the use of temporal zoom, i. e., the representation of data grouped by different units of time (e.g. hour, day, week or month) is something to be explored at the visual and auditory levels.

ACKNOWLEDGMENT

Catarina Maçãs is funded by Fundação para a Ciência e Tecnologia (FCT), Portugal, under the grant SFRH/BD/129481/2017.

REFERENCES

[1] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. H. Flowers, N. Miner, and J. Neuhoff, *Sonification report: Status of the field and research agenda*. International Community for Auditory Display, 1999.

[2] C. Maçãs and P. Machado, "The rhythm of consumptions," in *Expressive 2016 Proceedings of the Joint Symposium Expressive 2016, co-located with the Eurographics 2016*. ACM, 2016.

[3] D. A. Norman, *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1993.

[4] W. S. Cleveland and R. McGill, "Graphical perception: Theory, experimentation, and application to the development of graphical methods," *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 531–554, 1984.

[5] M. Friendly, *A Brief History of Data Visualizatio*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008.

[6] T. Hermann, A. Hunt, and J. G. Neuhoff, Eds., *The Sonification Handbook*. Berlin, Germany: Logos Publishing House, 2011.

[7] P. Vickers, B. Hogg, and D. Worrall, "The aesthetics of sonification," in *Body, Sound and Space in Music and Beyond:Multimodal Explorations*, C. Wöllner, Ed. New York, USA: Routledge, 2017, ch. 6, pp. 89–109.

[8] B. House. (2012) Brian house, quotidian record. [Online]. Available: http://brianhouse.net/works/quotidian_record/

[9] M. Quinn, "Research set to music: The climate symphony and other sonifications of ice core, radar, dna, seismic and solar wind data," in *Proceedings of the 7th International Conference on Auditory Display (ICAD 2001)*. Georgia Institute of Technology, 2001.

[10] J. Bulley and D. Jones. (2015) Living symphonies. [Online]. Available: http://www.livingsymphonies.com/

[11] B. Foo. (2015) Data-driven dj — two trains. [Online]. Available: https://datadrivendj.com/tracks/subway

[12] M. Wilson. (2013) Listen to the orchestra of users updating wikipedia. [Online]. Available: http://listen.hatnote.com/

[13] R. Gioria, A. Nogare, and A. Rao. (2014) Vanishing / digital / data sonification. [Online]. Available: https://www.behance.net/gallery/18558125/VANISHING-Digital-Data-sonification

[14] G. J. Smith. (2014) Specimen boxanalyzing botnet activity. [Online]. Available: http://www.creativeapplications.net/news/specimen-box-analyzing-botnet-activity/

[15] S. M. Bender, *Taiko Boom: Japanese Drumming in Place and Motion*. University of California Press, 2012.