

A design-driven sonification process for supporting expert users in real-time anomaly detection: Towards applied guidelines

S. Lenzi^{1,*}, G. Terenghi² and A. Moreno-Fernandez-de-Leceta³

¹Density Design Lab, Design Department, Politecnico di Milano, via Giuseppe Candiani, 72, 20158, Milano - IT

²Independent Researcher

³i3B, Instituto Iberoamericana de Innovación, Parque Tecnológico de Bizcaya, Ibaizabal Bidea, 48160, Derio - ES

Abstract

INTRODUCTION: We present early results of an ongoing investigation into the use of sound for real-time monitoring of anomalous behaviour in digital and digital/physical systems.

OBJECTIVES: We aim to define design guidelines to both support authors in the process of creating sonifications that are both efficient and engaging and the transition of sonification into a mass medium for the representation of data in everyday life.

METHODS: Through two Design Actions, we apply Design Research to the definition of the use case, the interaction paradigm and the experimental protocol for real-world evaluation of sonification tools.

RESULTS: Two Design Actions are described. Methodologies and results of the first experimental phase are presented in detail along with their influence on the second phase, currently ongoing.

CONCLUSION: We sketch a tentative design-driven process for sonifications for the real-time monitoring of anomalous behaviour in digital and digital/physical systems.

Keywords: data sonification, process monitoring, design research, sound design, soundscape

Received on 16 December 2019, accepted on 17 April 2020, published on 20 April 2020

Copyright © 2020 S. Lenzi *et al.*, licensed to EAI. This is an open access article distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi: 10.4108/eai.13-7-2018.163994

*Corresponding author. Email: sara.lenzi@polimi.it

1. Introduction

This article presents early results of an ongoing investigation into a designerly approach to the usage of sound for real-time monitoring of anomalous or malicious behavior in digital and digital/physical networks. We will illustrate two Design Actions built around the usage of sonification to represent data streaming from Artificial Intelligence-supported monitoring systems for expert users. These Actions aim to inform an applied approach to sonification able to answer key questions on the potential of sound as a complimentary sensory modality for supporting decision-taking in critical situations. We present a design-

driven approach to the representation of data through sonification that we believe will be able to address the main issues that are currently preventing this practice from being accepted as a standard tool available to designers. In particular, long time research goals include answering the following questions: in a data-intense society, why is the sense of hearing still underused in our daily relationship with data? How relevant is the design of the sound content in a sonification project? Can long-established practices such as sound design for film and gaming be of help in the design of usable and friendly interfaces for auditory monitoring and how does this relate to the so-called mapping problem in sonification? How can we validate the

success of such interfaces in a process that can lead from conceptualization to prototyping to real-world applications?

Through the description of two Design Actions (DA) focused on the sonification of data from a medium-sized water plant and the real-time sonification of a medium-sized Internet network, we aim to define a scalable design process in order to distil guidelines for the design of successful sonifications. We define successful, those sonifications that prove being of help to fulfil real needs, for real users, in the real world.

2. Data sonification today: an evolving discipline

The practice of using sound as a structured method to represent, analyse and communicate insights into data sets can be traced back to the 90s. ICAD, the International Community on Auditory Display was founded 25 years ago as a forum to discuss and define sonification as a discipline. The sonification of data and datasets was, at that time and for a large part of the past two decades, focussing on the representation of scientific data for purposes of expert analysis [1]. In fields such as seismology [2], geography [3], astronomy [4], social sciences [5], sound has been used as a complement or a substitute of visualization for the purpose of identifying patterns and outliers in (mainly) numerical datasets. Recently, with the exponential increase of both volume and impact of data in our society, the need to complement an overloaded visual channel with other sensory modalities has emerged, pushing the use of sound to represent data to gaining momentum. Expanding its reach beyond the borders of the scientific community, data sonification can now count on a diversity of use – cases that range from artistic expression to activism and data journalism through expert applications. At the same time, there seems to be an increasing debate within the community around emerging themes, such as the use of sound to engage non-experts audiences in socially relevant issues [6]; the potentialities of sonification for facilitating decision-taking in data-intensive tasks [7]; the role of aesthetics in sonification, with the accent on the relationship with the listener [8]; and finally, the interpretation of sonification as a design-driven practice in which the author intentionally places sound at the centre of a communication process [9].

Still, the real impact of sonification's artefacts on everyday life to improve or transform our everyday relationship with data is still extremely limited, if any. The reasons for this underrepresentation are certainly manifold. The lack of a standardised procedure to design and experimentally validate potential real-world applications of sonification has been recently pointed out at [7][10], but it is not a new concern. In 2011, the Sonification Handbook urged researchers in the field of Auditory Display (which includes data sonification) to familiarize with robust experimental protocols, described in detail in Chapter 6 of the Handbook [11] in order to validate the performance of

their sonifications. Literature presents a few cases where experimental research has been designed to investigate and possibly standardise how specific data dimensions can be successfully mapped into auditory dimensions [12][13] [3] a topic often referred to in the field of Auditory Display as the “mapping problem” [14][15]. Less effort has been dedicated to exploring how real users experience different mapping strategies in real-world situations, despite it being recognised as a crucial step by the very same authors. In their 2005 commentary [16] to their seminal article “Mapping and metaphors in Auditory Display” [12] Walker and Kramer conclude admitting “how crucial it is to empirically test an auditory display with listeners representative of the final users.” In this article, we argue that Design, as a discipline used to tackle complex, real-world “wicked” problems [17] can provide the methodological infrastructure sonification needs to engage in an iterative process with the goal of designing successful products, i.e. sonifications that can “transform the world from its current state to a preferred state” [18]. Such a position is not new in the sonification community. In 2012, Barrass already stated [19] that sonification could borrow from design, a discipline which is structured to “approach messy problems with multiple solutions whilst also allowing for critical reflection and empirical evaluation.” and built “through an iterative process of refinement of both aesthetics and functionality.”

With our research, we aim at substantiating such claims through the application of design research methodologies in order to feed an iterative process which would lead from ideation through prototyping to production of sonification as a real-world application. We interpret the mapping problem as a design (wicked) problem and as such, we understand it as being user-dependent as well as context-dependent and goal-oriented. In particular, in this article we focus on three areas in which sonification can be approached through the lenses of design: the definition of a use case i.e., of a real-world situation on which sonification can have a positive impact; the design choices made by the designer when it comes to the sound material of a sonification and their role in measuring the success of a sonification; the definition of a research protocol to validate the impact of a sonification in the real world and inform new solutions for similar use cases.

2.1. Defining a use case for sonification

An explicit use-case for a sonification project that aims to use sound to have an impact on how we relate to data in real life, must be defined beforehand. We hypothesize that the lack of a prior analysis of the characteristics of users, contexts and goals of a sonification project negatively affect the impact that sonification can have as a real-world application. Adapting a designerly way of knowing [20] to the design of data sonification artefacts, we propose that authors of sonifications engage, prior to focussing on issues of data-to-sound mapping, in an iterative process that would lead to the definition of what needs (including emotional

needs) the sonification would fulfil, for what user, in what context. Possible design constraints such as technical, economic, social and cultural aspects that might affect the outcome of the project should also be taken into account. This phase would cover, what a typical design thinking would define as the *research for design* phase, an “accepted ingredient of user/human-centred design methodologies” [21]. Its outcome should be the definition of a specific use case where the designer hypothesizes how sonification can facilitate an improved (more efficient and more engaging) relationship with data for a specific group of users, in a specific context. The ideation of a solution which outlines the sonification concept, the sound material to be used and the paradigms of data to sound interaction, would be sketched at this stage. To design professionals, these steps will sound familiar, even obvious. Disciplines such as architecture, product design, communication design, service design all apply such design thinking to their work process, from conceptualization to final production through iterative cycles of prototype adjustment. Introducing such a structured approach to sonification would help “reconfigure it from an instrument of scientific enquiry to a popular mass medium for a broad audience” [19].

2.2. Introducing sound design practices

Undeniably, the road for sonification to become a recognised mass medium for relating to data is still long. Even from within the auditory display community, sonification has been sometimes labelled as intrusive, distracting, causing listener fatigue, annoying, subject to issues of display resolution, precision, comprehensibility [14, p.154] and even “utterly ugly” [22]. Interestingly, the field of data visualization faced a similar challenge when transitioning from a mainly computer science discipline to a mass medium that needed to communicate to a lay, non – expert audience. Reflecting on the potential role of design to guide this transformation, authors such as Lau [23] and Manovich [24], proposed the term info-aesthetics when they called for a wider contribution of design to information visualization. This contribution would inject often overlooked considerations of the aesthetic dimension of the visual experience to prompt a more intuitive understanding by a non-expert audience. The aesthetic dimension of design and its role in the ideation process and, more importantly, in the experience of the final user is a long-standing debate which definition overcomes the scope of this article. According to Vande Moere [25], the aesthetic dimension of a visual representation of data is to be found both in data visualization as an artistic expression and in commercial products by professional designers as a parameter of “attractiveness” (appeal or beauty). The attractiveness of a visualization “may compel the user to engage with the data, enabling more effective communication of the information itself” [ibid., page 363]. Designing with attractiveness in mind (besides functionality, often defined in terms of effectiveness and efficiency) has to do with a process where the designer, following an exploratory phase (see Section

2.1) makes choices based on a designerly way of knowing and doing which is, still today, mostly tacitly acquired and developed through practice. The explicit acknowledgment of the rationale behind design choices in data visualization is, according to Vande Moere, a much-needed step for the development of guidelines for the creation of successful (efficient and compelling) visualization that intend to improve the engagement and meaning-making in the context of human-data relationships.

Opening the data sonification arena to contributions of other sound and music-related disciplines is not a new claim [8][14][19]. Here we propose to draw a parallel between the role of visual design practices for data visualization and the potential role of sound design practices for the transformation of data sonification into a medium of communication. We believe that the contributions of the field of sound design and in particular, sound design for film and gaming and sound branding, could form a corpus of practical knowledge that, if made explicit, would foster the design of more compelling data sonification experiences. Sound design for film, in particular, benefits from a long-established artisanal tradition now condensed in numerous educational curricula both as under- and postgraduate level where a mainly tacit knowledge has been structured to form new generations of sound designers. At the same time, authors like Michel Chion developed, in his seminal work “The Audiovision” [26] a rather sophisticated and influential theoretical model that attempts to systematize decades of tacit knowledge and successful practice. Sound branding, a mainly commercial field of application of sound design, reflects on its way of doing in more didactic publications such as *The Sonic Boom* [27] and through industrial arenas such as the annual Audio Branding Congress[†] organised by the Audio Branding Academy. ABA also confers industry awards where quality parameters are discussed and established, and publishes an annual yearbook where professionals reflect on their work and share the design rationale behind their projects. Sound design for multimedia, interactivity, gaming, film production, branding is also supported by numerous online and offline communities, forums and publications. Last but not least, soundscape studies, with their attention to the everyday experience of sound and the long-standing call [28] for an interdisciplinary approach to the design of everyday sounds could also provide design guidelines to approach the production of engaging sound contents. We argue, along the lines of Vande Moere, that making explicit the rationale behind these practices could inform design guidelines where *good design* [12] is defined through an iterative process that moves beyond what sounds “good” or “correct” based on the designer’s intuitions while acknowledging the aesthetic dimension of a sonification experience as a success factor to engage the audience in a more intuitive meaning-making process.

[†] <https://www.international-sound-awards.com/>

2.3. Experimental validation: designing a design research protocol for sonification

Walker and Kramer concluded their first investigation into the systematization of mapping metaphors for representing data dimensions acknowledging that “it is crucial to empirically test an auditory display with listeners representative of the final users” to inform iterative prototyping and verification of interface designs [12, p.411]. Twenty-five years later, the same need is still recurrent from within the data sonification community [10][14][15][19]. Admittedly, a design process that engages final users from the beginning of the design process (when decisions on mapping and sound material are taken) and ideally continues iteratively during prototyping until final delivery, is far from being the norm in the data sonification community. Moreover, there is no shared experimental protocol to assess the appropriateness of a sonification to real-world situations, to the detriment of its chances of becoming a more widely used data representation method. Again, we can turn to design to develop a path through which data sonification can reflect on its unique “way of doing”: as Ken Friedman puts it: “The challenge of any evolving field is to bring tacit knowledge into articulate focus. This creates the ground of shared understanding that builds the field.” [29]. In the past decades, the field of design has struggled to develop flexible and adaptable experimental protocols able to make justice of the designerly way of doing and knowing [20]. In the following Sections we will illustrate how we are borrowing from design research methods and methodologies [21] to define a design process for sonification projects from conceptualization to final delivery. In Section 4.3 we describe an experimental protocol that juxtaposes quantitative research with (qualitative) design research methodologies such as technology probes [30] and design probes [31].

3. Defining the use-case: using sound to disclose anomalous behaviour in real-time network monitoring

There seems to be a growing need for tools that can facilitate the communication between AI-based anomaly detection algorithms and human operators of SOC (Security Operations Centres). AI is being heavily introduced to monitor, in real-time, the behaviour of digital and digital/physical networks such as electric grids, water plants, Internet networks and so on, with the explicit goal of detecting at an early stage anomalous behaviour (mainly, but not only, due to cyber-attacks) and making predictions on future anomalies. Operators of data centres, already exposed to a heavy load of visual information, find themselves in need of yet another data representation interface with growing concerns on the impact on the operators’ performance due to information overload [32]. Additionally, AI algorithms are still exposed to a certain error rate which increases the chances of alarms resulting in false positives,

which in turns results in an increased load on the operators [33]. In the recent past, sonification has been applied as a complementary tool for the monitoring and disclosing of anomalies in fields such as medicine [34]; security [35]; industrial production [36], transportation [37]. In general, these applications leverage some well-recognised characteristics of sound that represent an added value for the detection of anomalies in real-time. These include, for instance: sound typically occupies the periphery of our attention, thus allowing for the user to carry out other, mainly visual, tasks [38]; we have a high capability to easily recognize recurring patterns or sudden changes in continuous sound events [39]; through sonification we can mitigate information overload of the visual channel [3].

The use case we defined for the application of our DA describes the following situation: in a medium-sized SOC, operators (normally three people) would be monitoring a system’s activity through about a dozen of screens. Real-time information is processed and displayed on the network’s SCADA (Supervisory Control and Data Acquisition) system. This includes both analytical data (for instance, details of a network’s single node status at a given time) and historical data (flowcharts of the network’s behaviour over a given time interval). We propose to integrate a real-time peripheral monitoring tool based on sonification. As a working hypothesis, we understand sound as a guide for a more efficient retrieval of analytical information in the existing visualization tools. Sound would be therefore used in order to complement visual information with the goal of facilitating an early response to emergencies. Additionally, we hypothesize that on the longer term, operators will be able to detect errors in the AI, integrating, thanks to continuous monitoring allowed by the sonification, their knowledge of the system with the information provided by the algorithm. At the end of our exploratory phase (see Section 2.1), we sketched a sonification concept where an interactive digital tool would generate a continuous background soundscape, which changes in the presence of anomaly would be determined by mapping rules. To these changes, operators are expected to react with an appropriate response, for example, searching for specific analytical information in the SCADA visualization system. This concept is not new in sonification literature [35][36][39][40]. With the early engagement of representatives of the final users, with the application of research protocols designed to build, through audience engagement, an iterative prototype development, with the focus on making explicit the implicit sound design choices, we intend to complement existing literature advancing in the design of a successful (efficient and compelling, see Section 2) sonification. In other words, we aim at designing a sonification artefact that ideally could become a product “successful in the marketplace” [18].

In the following Sections, we describe the two DA for the detection, through sound, of anomalies in a medium-sized water plant and in a medium-sized Internet network. The first Action uses sonification to represent data streaming at hourly rate from a water distribution network to the control room, in order to detect and disclose anomalies due to

cyber-attacks. The project, a collaboration between Politecnico di Milano and Singapore University of Technology and Design, is described at length in [33]. After an iterative prototyping phase an experimental phase involving six expert users in a real setting was undertaken. Results were later used to substantiate the design approach to the second Design Action, which applies sonification to anomaly detection in a medium-sized fully digital (Internet) network. Data on the behaviour of the network and its components are streamed and sonified in real-time to monitor anomalous behaviour and detect issues of security before they become critical. This second case, still ongoing, will be described mainly focusing on the sound design phase as described in Section 2.3.

4. First Design Action: disclosing cyberthreats to water distribution systems through sonification

Due to the extensive usage of digital sensors and digital technologies, physical systems such as water plants and water distribution networks are more and more confronted with cyber-threats which expose these key infrastructures to unknown dangers. Artificial Intelligence is being deployed to support human operators in disclosing these threats and promptly react, to avoid potentially critical situations. Though more and more vital to the management and understanding of the vast amount of data streamed at each moment from such networks, machine learning and deep learning algorithms are still exposed to a certain degree of error which may mislead operators to take the wrong decision in a moment of danger. At the same time, as described in Section 3, SOC centres are already heavily reliant on visual screening for the storage, management and analysis of incoming data. Adding yet another visual interface dedicated to the real-time monitoring of cyber-threats would expose analysts to a higher risk of information overload and alarm fatigue. Additionally, during their normal working day operators of a medium-sized water plant are expected to carry out other (again, mainly visual) tasks such as reading and writing reports, analysing historical data sets, answering the occasional phone call, interacting with colleagues. The peripheral monitoring that sound allows for makes sonification a promising choice for the real-time monitoring of incoming data on cyber-threats.

4.1. Introduction to the Design Action

We conceptualised, designed, prototyped and tested in a real environment a tool that uses sonification to represent data collected by the SOC of a water infrastructure of a medium-sized imaginary town (C-Town). We hypothesize that the sonification tool is complementary, and not a substitute, for data visualization. Figure 1 represents the map of C-town where the plant is located. The imaginary network is organised in five districts.




Figure 1. Map of C-Town

Data is sonified at an hourly rate to allow peripheral monitoring of potentially malicious activity. Once alerted, the operator would react first of all searching for additional, more detailed information on screen where the existing data visualization, with details of the network present and past status, is displayed. Following the working hypothesis described in Section 3, we identified two main goals of the DA:

1. The sonification would complement the existing visualization and as such, we would not try to convey all the existing information through sound.

Borrowing from Lean Innovation methodology [40], we assumed the sonification tool to have the attribute of *good enough* i.e., the tool would not aim at delivering to the user all the existing information as obtained by the detection algorithm, but information that is good enough to quickly locate additional relevant information in the existing – mainly visual – data representation tools.

2. The sonification would try to put the human operator in the position of identifying errors in the anomaly detection algorithm.

For this purpose, the sonification tool is designed not as a binary alarm system which would communicate to the operator the presence or absence of malicious activity (alarm yes/no) according to a “ex post” judgement made by the detection algorithm, but as an alert system. Such a system would represent anomalies in the network as a continuum: the discrepancy between the expected behaviour and the real behaviour of the system at a given time is represented as is, without a prior judgement on whether the system should be considered under attack or not. This way, we aimed at limiting the impact of false positives thus bringing back the decisional responsibility to the human operator and his knowledge of the network.

4.2. First and second prototypes

A first round of prototypes was designed to explore different mapping hypotheses. Each prototype used different sound material which included musical instruments, concrete sounds and digital sound processing such as distortion and pitch shift. At the core of this first group of prototypes were a series of choices on the data-to-sound mapping and design of the sound content:

- Data-to-sound mapping reflected an approach based on the use of conceptual metaphors to map data dimensions into sound dimensions.

Conceptual metaphors, a concept introduced in 1980 by George Lakoff and Mark Johnson [41] are not new to the world of data sonification [12][13][42][43]. In this case, we designed the mapping strategy based on the metaphor that through distorted sound, a listener would gauge distortion in the data represented by the specific sound i.e., that hearing *something wrong in the sound* would metaphorically communicate the expectation of something wrong in the data. The usage of distortion has been recently attributed to the representation of uncertainty in geographical data [3], an additional experimental result in support of our approach. Moreover, conveying a sense of “wrongness” through distortion is coherent with a typical approach in sound design for film and gaming, where distortion is often used to amplify the spectator’s feeling of uneasiness or alteration in the narrative [44].

- Walker and Kramer [45] considered that the addition of contextual sounds “that do not represent data, but rather contribute to the interpretation of the sounds” can be very beneficial. We decided to design the sound material to carry additional contextual information.

Based on these two main design choices, we sketched a series of prototypes composed by two sound layers: a series of realistic background sounds were meant to represent, with various degrees of indexicality [46] the different components of the network (pumps, tanks, valves); on top of these sounds, more abstract “alert sounds” were added. These sounds would react to anomalies in each of these components, showing various degrees of digital distortion. Different distortion methods were explored, such as pitch shift, delay, filtering, addition of digital distortion such as clicks and noise. This first round of prototypes was then shared with experts for a first, closed-doors feedback session which led to discarding most if not all of them due to issues of usability, efficiency and pleasantness of the sound experience. In particular, the constraint of data streaming at an hourly rate made it technically very difficult to use continuous, evolving indexical sounds in a meaningful manner. It is here important to notice that some of the concepts of this very first approach was later integrated into the design of the second Action, following the experimental results described in 4.4.

The second round of prototypes was designed completely pivoting our approach. Four prototypes, based on four different data-to-mapping Scenarios, were sketched taking into account the experts’ feedback to the first round. In

particular, the experts felt that sonification should convey little but clear information on the anomalous status of the whole system and the gravity of the anomaly, while the first round of prototypes tried to represent the status of each component of the network. Additionally, the need seemed to emerge to be able to locate information on an ideal map of the city. We therefore decided to focus on the representation of the five districts, each one with a different sound. As for the sound material, in the absence of clear indications from existing literature on the performance of tuned sounds versus non tuned sounds in this context, we decided to decline each scenario in two versions: version A, using musical (tuned) sounds; version B, using concrete (non-tuned) sound samples. Figure 2 illustrates the iteration cycles of the prototypes. Each Scenario is labelled with the corresponding mapping strategy.

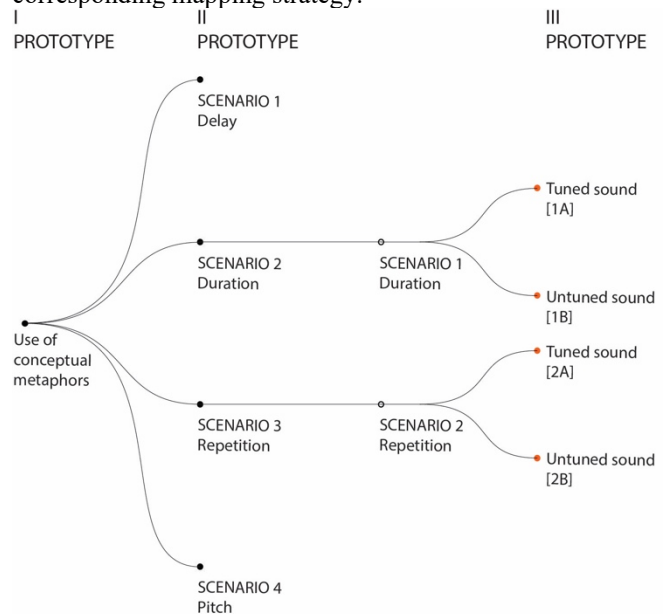


Figure 2. First Design Action. Prototypes iteration cycles

In all Scenarios, in a regular situation, the five different sounds for the five districts play synchronically at every hour. In case of anomaly, the behaviour of each sound would be altered. In particular:

- In Scenario 1, the sound of a district presenting anomalies would play later than the others, proportionally to the amount of anomaly.
- In Scenario 2, the sound of the anomalous district would last longer, proportionally to the anomaly.
- In Scenario 3, the sound of the anomalous district would loop. The speed of the loop cycle would increase based on the value of anomaly.
- In Scenario 4, the pitch of an anomalous district’s sound would be altered proportionally.

Following another experts’ review, only Scenario 1 and 3 were brought forward for validation with users.

4.3. Experimental design

Once selected the two Scenarios in two Versions, we turned to investigating design research methods for the experimental phase. As discussed in Section 2.2, we believed the concept of technology probes [30] and design probes [31] to be particularly promising. Probes are a method used to inspire developments in a design process [47]. Usually, probes are introduced at a very early stage of the conceptual development in order to gain fresh, unstructured insights that could even radically change the course of the prototyping. In the words of Hutchinson [30] probes satisfy three main goals “the social science goal of understanding the needs and desires of users in a real-world setting, the engineering goal of field-testing the technology, and the design goal of inspiring users and researchers to think about new technologies.” We believed that these goals well matched our objectives. At the same time, we needed to assess the performance of the user against the specific task of disclosing cyber-threat to the network using the sonification tool. We therefore designed an experimental protocol which involved both qualitative and quantitative methods. The experiment included three phases:

- a preliminary questionnaire, to assess the competence of each user in the fields of water management and cyber-threats and their familiarity with music and/or sound production.
- A task-oriented test to be completed during the experiment, to assess the performance of the users understanding the status of the network (regular/anomalous); rate the level of anomaly on a 1 to 5 scale; identifying how many districts were anomalous.
- A final, semi-structured interview where we aimed at gauging feedback on the overall experience, design and technological issues and at gathering suggestions for future developments of the prototypes.

Six experts from the field of water infrastructure management and cybersecurity (three men and three women, from different countries and cultural backgrounds) were selected. Scenarios 1 (A and B) and 2 (A and B) were made

available via real-time streaming to the testers in order for them to integrate them into their typical job routine. Our users would switch on the sonification web application at the beginning of their working shift and this would automatically play every hour (as mentioned, according to the current rate of data streaming from the network to the SCADA system) during eight consecutive hours. In line with the probes approach, we decided to provide users with a general introduction to data sonification and to the specific problem our sonification was meant to solve, but not to require them to go through a training phase nor did we provide the testers with a demo of the sonification Scenarios. Only the keys to the mapping along with the individual sounds associated with each district, for each Scenario, were provided. This way, we hoped to inspire users to find their own way in the relationship with the sonification in order to collect as much unexpected feedback as possible. Having four different sonifications to test over a few days, we considered the risk that the users would learn from the first sonification and use this knowledge to better perform with the following ones. We therefore decided to distance the testing of each scenario over a period of two weeks, in order to have a few days' interval between each test, and to randomize the order of each Scenario's versions. All the prototypes, the preliminary instructions and an extended description of the project can be found online at <https://ginevraterenghi.github.io/sonifying-cyber-attacks/index.html>.

4.4. Results

Figure 3 shows the results of quantitative testing for each Scenario at every hour, for each tester. If the identification of the anomalous status of the system does not seem to pose any issue in any of the Scenarios, testers performed more irregularly in the identification of the number of anomalous districts. In general, Scenario 2 seems to give better results while on the other hand, the same Scenario performs less solidly as for the attribution of a level of anomaly.

- (i) the hanging indent. This is a list, note the hanging indent.




Figure 3. Analysis of quantitative testing

As mentioned in Section 4.1, the level of anomaly was not embedded by design, but was obtained retroactively while analysing the testers’ answers, to be able to check their performance. The reason for this is that we did not want to predetermine a level of anomaly attributing artificial thresholds that are not in the original dataset. Ideally, we wanted such thresholds to emerge progressively as the operator becomes familiar with the sonification tool. It was nonetheless important to be able to judge the coherency of each tester’s subjective evaluation of gravity to gauge their understanding of the sonification. Figure 4 shows examples of the analysis of qualitative results, which are more of interest for the scope of this article as they informed the design approach for the second DA. Results were obtained transcribing the final interviews and clustering the answers based on a series of sub-themes. As mentioned in Section

4.1, the level of anomaly was not embedded by design, but was obtained retroactively while analysing the testers’ answers, to be able to check their performance. The reason for this is that we did not want to predetermine a level of anomaly attributing artificial thresholds that are not in the original dataset. Ideally, we wanted such thresholds to emerge progressively as the operator becomes familiar with the sonification tool. It was nonetheless important to be able to judge the coherency of each tester’s subjective evaluation of gravity to gauge their understanding of the sonification. Figure 4 shows examples of the analysis of qualitative results, which are more of interest for the scope of this article as they informed the design approach for the second DA. Results were obtained transcribing the final interviews and clustering the answers based on a series of sub-themes.

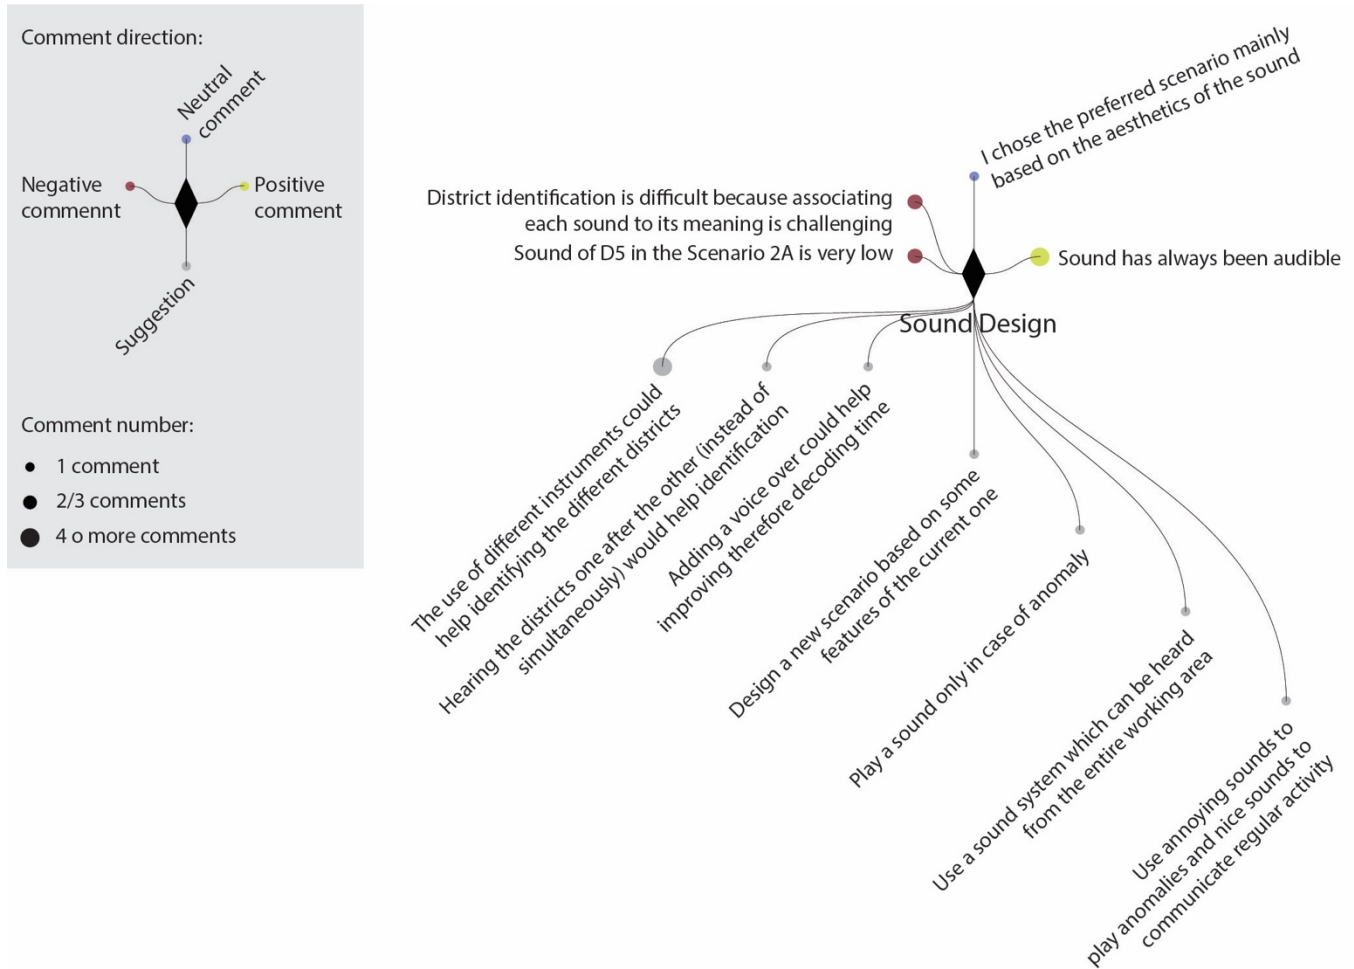


Figure 4. Analysis of qualitative testing

The theme *Sound Design* (shown in Figure 4) is the most interesting for this article. Users suggested a series of design variants to the prototype, including the use of musical instruments and voice, the use of silence and the use of emotional attributes to define anomalous/regular activity. The theme of an aesthetic judgement and its influence on the

choice of a preferred scenario also emerges. In Figure 5 we can see as users seem to attribute meaning to each sound while also relating feelings of pleasure and displeasure, happiness, fun, uneasiness, irritation and even fear, towards specific Scenarios.

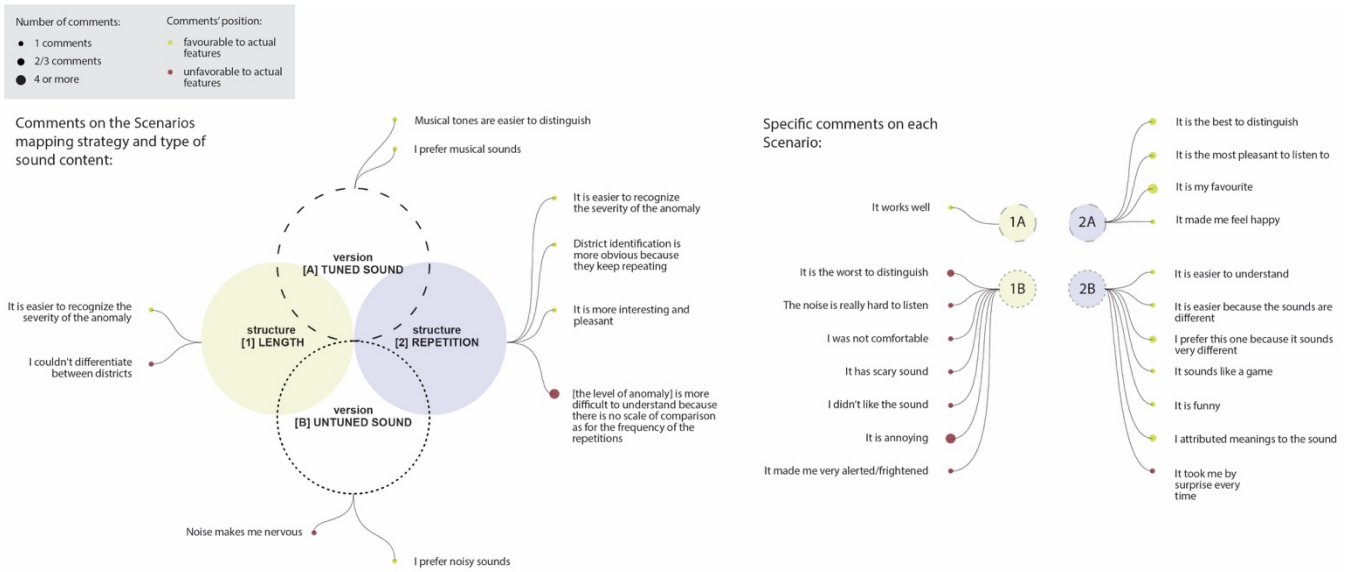


Figure 5. Analysis of clustered results for the four Scenarios

No Scenario emerged as the clear “winner” of the user’s preferences, with feedback being inconsistent throughout the experiment, as shown in the left section of Figure 5. On the overall experience, all six users agreed on the importance of the design of the sound material, along with the request for sound to be able to integrate not only with the work environment but also with tastes, needs and habits of the users. Additionally, users would welcome a background sound that is non-intrusive and potentially relaxing, as long as its design characteristics do not interfere with the analyst’s capacity to focus on their tasks. For example, volume should be reasonably adjusted in order for them to hear requests from colleagues or phone calls, sounds should not repetitive, sounds should be familiar enough to remain at the periphery while still attracting attention in case of a change in their behaviour without increasing the cognitive load - such as familiar, recognizable sounds would do. These results were used to inform the design of the second Design Action: using sound to monitor and disclose anomalous behaviour in fully digital networks.

5. Second Design Action: digital networks as dynamic natural soundscapes

The second DA, which development is still ongoing, applies sonification to a fully digital, medium-sized Internet network property of the Spanish company Ibermática. Ibermática also served as a partner of the sonification project through its research institute i3B[‡]. A fully functioning prototype was designed that would use data from the proprietary digital platform PLAGEMODA, which collects, stores and analyses data received from the

[‡] <https://ibermatica.com/innovacion/>

company’s Internet network’s activity. A customized AI algorithm completes the platform allowing for the disclosure of anomalous behaviour. The project is described in detail in [48]. For the goals of this article, after briefly describing the prototype, we will discuss how the experimental results of the first Design Action informed the design choices of the second Action and the plan for its experimental validation. While the data-to-sound mapping strategy was adjusted to adapt to the specific constraints of this Action (see paragraph 5.1), the design of the sound contents was both informed by the results of the experiment described in 4.3 and by existing literature. In particular:

- Feedback from users of the first DA stressed the importance, in real-time peripheral monitoring, for the sound content to be unobtrusive and to locate itself at the background of office activity without interfering, unless needed. Additionally, users seemed to have a tendency to attribute a familiar meaning to concrete sounds in order to facilitate identification and recollection. They also highlighted how known sounds could minimize their cognitive effort and learning curve and how aesthetic considerations and feeling of pleasure in hearing the sounds encouraged them to look for insights into the data those sounds represent.
- Overall results lacked a clear indication of both aesthetic preference or performance advantage between tuned or untuned sounds. We therefore felt free to introduce new sound material to experiment with.
- Literature presents a few cases where natural soundscapes have been introduced in peripheral monitoring of real-time data streaming [36][49][50]. In [50], Vickers states that “we are already used to dealing with everyday background sound and quickly deciding what sounds need attending to and what sounds can be pushed to the attentional background. A soundscape

offers the sonification designer the potential to leverage this innate information processing capacity in such a way that important changes in the cyber environment become salient in the soundscape.”

We therefore considered promising to focus on a strategy in which a dynamic natural soundscape evolves based on real-time data and seamlessly adapts to the context of a shared office environment.

5.1. The DynAmb prototype

A digital tool was designed that receives data from the PLAGEMODA platform to which the custom anomaly detection algorithm is applied. At the granularity we are looking at, data from the network are clustered around three main Fields (see Table 1) which relative weight (the Field Impact) determines the overall behaviour of the network.

Table 1. Example of the dataset used in the second Design Action

orig_h	Network status	Field Impact	Anomaly Index	Anomaly
	Field 1					
192...63	Record_Count_transf	0,1536			2,930634	F
192...97..	orig_h	0,5368			1,525231	F
192...102	resp_h	0,28205			1,518315	F
192...63	Record_Count_transf	0,16843			3,472138	T

At any given time, the algorithm would assign relative weight to each of the Fields, which would in turn give us an indicator of what components of the system are showing a higher level of activity. Based on this, a global Anomaly

Index is determined by the algorithm to assign a level of anomalous behaviour to the network. Only if this level is understood by the algorithm as a danger, an indication of Anomaly “Yes” would also be given to the operator. Unlike the context of the first DA, data is streamed (at this granularity) every minute: we therefore needed a continuous soundscape instead of punctual sound events played every hour as in the first Action. This constraint was the main difference with the previous DA. Coupling this consideration with existing literature and feedback from users (see Section 4.3), we drafted a design concept where the digital network is represented by an evolving natural soundscape, in particular, the soundscape of a forest. We understand the soundscape as a simple metaphor (the Internet network as a forest) through which the user/listener makes sense of the behaviour of the network and its components interpreting the behaviour of the forest both as a whole and in its individual components (birds, insects, wind through the trees, rain and the occasional thunder). An approach based on conceptual metaphor was used in the first, discarded prototype of the previous DA, as described in 4.1. We therefore decided to recover that concept, in the belief that thanks to data streaming practically in real - time we would overcome the main technical obstacle that had doomed that first prototype (see Section 4.2).

Data-to-sound mapping also reflected the approach to the first prototype of the first DA. The conceptual metaphor of distorted sounds as representative of anomalies in the data was re-considered, with the belief that this time, it would be successful. Figure 6 shows how each parameter of the dataset is mapped to sound material representative of the soundscape of a typical European forest, with birds, insects and wind representing the networks’ Fields and rain and thunders representing anomalous behaviour.

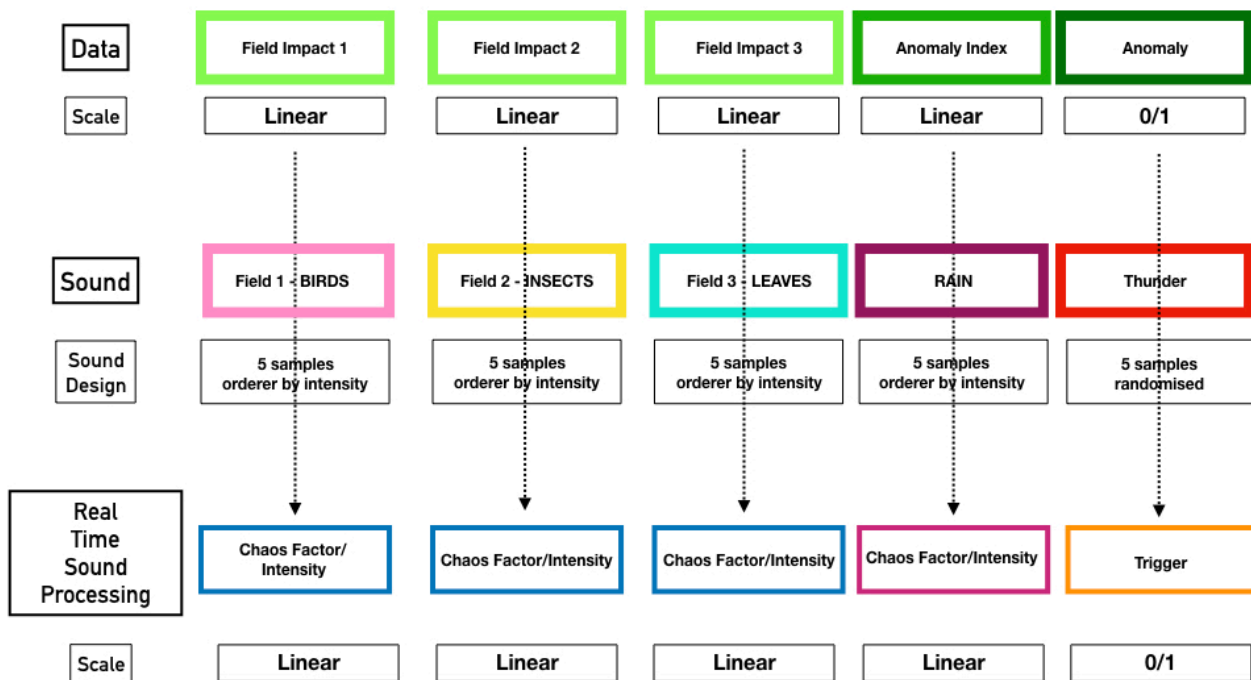


Figure 6. Second Design Action. Data-to-sound mapping matrix

We interpret the forest soundscape as a metaphor that stands for the Internet network. The digital network is seen as a forest where each element contributes to the overall status of the ecosystem. When one – or more than one – element of the ecosystem behaves anomalously (growing too much, growing too little, disappearing due to a disease or to the action of humans, and so on) the overall system is affected. Usually, a certain degree of anomalous behaviour is embedded in any living system – and so it is in a typical Internet network. When certain conditions are met, though, the anomaly becomes explicit and the system needs to act in order to avoid serious consequences. Once we defined the embracing metaphor, we moved on to sketching the interaction paradigm in detail. As shown in Figure 6, elements of the forest ecosystem such as birds, insects and leaves (sonified as wind blowing through trees) represent each Field, while rain and thunder represent anomalies. Metaphorically, the approaching of a storm signifies an increase in the level of anomalous behaviour. A thunder would be the sign that the anomaly is recognized as such by the algorithm i.e., it is a malicious anomaly. Leveraging the conceptual metaphor of distortion in sound as an anomaly in data, we defined a “chaos factor”: an auditory element

meant to alter the perception of the listener and alert her on anomalies in the network. The chaos factor is proportionally mapped to data i.e., the more data show an anomalous behaviour, the more the behaviour of the corresponding sounds will become chaotic. In the current version of the prototype, we designed the chaos factor as a proportional increase in the intensity of sounds. We define “intensity” as the combination of the acoustic density and volume of each sample in a specific category, on a scale 1 to 5. For example, at a low intensity, Field 1 (birds) will sound as Sound Example 1[§]: where birds’ calls show a behavioural pattern connected to a sense of distance, calm. At maximum intensity (i.e., maximum anomaly level), birds’ calls will sound denser, louder, overall more chaotic (listen to Sound Example 2^{**}). We implemented an off-line prototype - called DynAmb (Dynamic Soundscape) - using Python^{††}

§ <https://soundcloud.com/user-467084838/low-intensity?in=user-467084838/sets/field-1>

** <https://soundcloud.com/user-467084838/high-intensity?in=user-467084838/sets/field-1>

†† <https://www.python.org/>

and Max/MSP by Cycling '74^{‡‡} with the goal of testing it with real users in a real environment to validate our working hypothesis. A Python script simulates data streaming at a 2 seconds' rate from the PLAGEMODA platform to DynAmb, where a Max/MSP patch translates it into sounds. The two tools communicate via OSC (Open Sound Control) protocol^{§§}. A demo of the Max/MSP tool receiving and sonifying both regular and anomalous data can be found at <https://vimeo.com/409007148>.

5.2. Working hypothesis and experimental design prototype

Two rounds of adjustments have been done to the prototype following experts' advice. At this stage of the research, we formulated a working hypothesis to be tested with real users in a real work environment, and we are starting to define the experimental protocol. We hypothesize that the current prototype:

- meets the requirements emerged in previous testing of a sonification that blends with context (an office working environment), is pleasant to hear and intuitively understood, minimizes cognitive effort to decode information, leverages familiar sounds with a high sound quality.
- Allows for peripheral monitoring.
- In the long run, it helps the operators predict anomalous behaviour, thus supporting the prevention of mistakes due to errors in the AI.

We are designing an experiment to validate these assumptions and further explore the relationship of real users with the design of the sonification tool. For this second Design Action, we plan to:

- maintain a task-based quantitative questionnaire to evaluate users' performance.
- Replace an approach based on probes with phenomenological research. Barrass and Vickers already identified IPA (Interpretative Phenomenological Research) as a potential complement to "traditional metric- and task-performance based technique" with the aim to "gain insights into how a person experiences, or makes sense of, a phenomenon." [14, p.19].

The reasons why we are planning to replace a probe-based qualitative investigation is that the current prototype is more than a probe, it is a fully functioning sonification tool. Additionally, we do not have different versions of the tool to compare between (as in the first DA and in [31]) therefore we want to observe the users more closely, focusing only on the tool and the development of a relationship with it as a specific phenomenon. The experiments will take place in a real work environment where expert users will integrate the prototype into their daily routine for a longer period of time.

^{‡‡} <https://cycling74.com/products/max/>

^{§§} <http://opensoundcontrol.org/osc>

6. Towards design guidelines for applied sonification in network monitoring and anomaly detection

In this article, we illustrated preliminary results of an investigation to assess the potential role of design for supporting the transition of the field of data sonification from a niche method for scientific analysis into a mass medium for human-data interactions. We believe that authors of sonification projects could benefit from applying a designerly way of doing to three main areas: the definition of a specific use-case prior to taking decisions on data to sound mapping; the design of the sound material; and to the definition of an experimental protocol for testing sonification in the real world. We described two Design Actions where we apply sonification to the real-time detection of anomaly in digital systems. Figure 7 sums up the results of this first phase of our investigation.

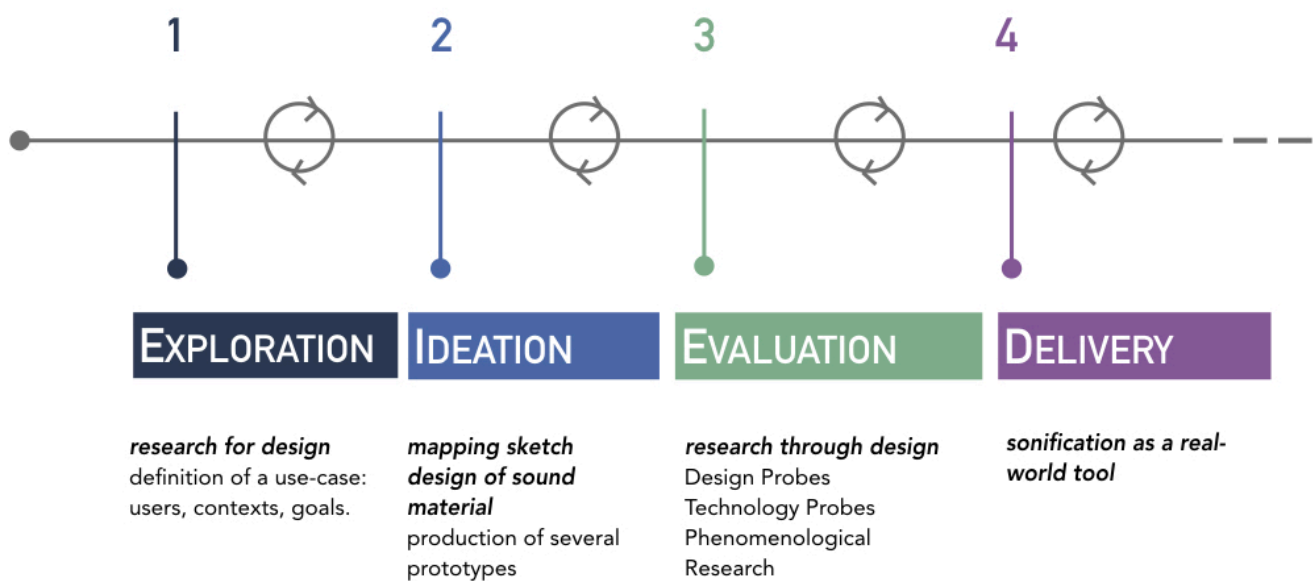


Figure 7. Sonification design process mapping matrix

We used a *research for design* process to define the specificities of the use-case. For the first Design Action, we designed several prototypes of which two were tested through a methodology inspired by technology and design probes. We used the results of this experiment to inform the design of a prototype for the second Design Action and in particular: in our use-case, users expect a sonification to be unobtrusive in the context of daily routine; using familiar sounds might help minimizing cognitive effort and the need for training; users recognize the importance of the aesthetics of a sonification to enhance user engagement and facilitate its adoption in real life. The resulting prototype builds upon the soundscape of a forest in which anomalies in the data are metaphorically represented by distortion (our “chaos factor”) in the sound content. The experimental phase for the second Design Action is currently ongoing: through phenomenological research we are investigating how users integrate the sonification in their daily work routine to evaluate the success of the prototype. The results of this

new experimental round will inform a new iteration of the prototype that we will apply to other contexts with a similar user-case, such as healthcare monitoring, industrial manufacturing and cybersecurity. We believe that the design-inspired work process represented in Figure 7 can form the basis of a framework for the design of successful (efficient and compelling) sonifications that real users will be willing to adopt in real life situations.

Acknowledgements.

Sara Lenzi and Ginevra Terenghi wish to thank Prof. Stefano Galelli from Singapore University of Technology and Design (SUTD) and Prof. Riccardo Taormina from TU Delft for first triggering and later tirelessly contributing to the development of first Design Action. They also wish to thank their supervisor Prof. Paolo Ciuccarelli for continuous advice and support. Sara Lenzi and Aitor Moreno wish to thank Prof. Damiano Meacci (Conservatorio di Bologna) for their help in the development of the second Design Action prototype in Max/MSP.

References

- [1] Kramer, G. & Walker, B. & Bonebright, T. & Cook, P. & Flowers, J. & Miner, N. & Neuhoff, J. & Bargar, R. & Barrass, S. & Berger, J. & Evreinov, G. & Fitch, W. & Grohn, M. & Handel, S. & Kaper, H. & Leivkowitz, H. & Lodha, S. & Shinn-Cunningham, B. & Simoni, M. & Tipei, S.. (1999). *The sonification report: Status of the field and research agenda*. Report prepared for the National Science Foundation by members of the International Community for Auditory Display.
- [2] Dombois, F. (2001). Using audification in planetary seismology. In *Proceedings of the 7th International Conference on Auditory Display (ICAD)*, (Helsinki, Finland, July 29th-August 1), 227-230.
- [3] Ballatore A., Gordon, D., Boone, AP. (2018) Sonifying data uncertainty with sound dimensions. *Cartography and Geographic Information Science*, **46**(5): 85-400.
- [4] Diaz Merced W. L. (2013) *Sound for the exploration of space physics data*, PhD thesis, University of Glasgow.
- [5] Dayé C. and A. De Campo (2006.) Sounds sequential: sonification in the social sciences, *Interdisciplinary Science Reviews*, **31**(4): 349-364.
- [6] Sobliye, G. and Mortada, L. (2017). Uneasy listening: El subutilizado arte de sonificar datos para contar historias de injusticias sociales. *Revista Diseña*. Vol.1, pp. 202-219.
- [7] Hildebrandt, T. and Rinderle-Ma, S. (2015), Server sounds and network noises. In *6th IEEE International Conference on Cognitive Infocommunications (CogInfoCom)*, Gyor, 2015, pp. 45-50.
- [8] Roddy, S. & Furlong, D. (2014) Embodied Aesthetics in Auditory Display. *Organised Sound*. **19**(Special Issue 1. Sonification):7-77.
- [9] Lenzi, S. and Ciuccarelli, P. (2020). Intentionality and Design in Data Sonification of Social Issues. *Big Data and Society*, In Press.
- [10] Axon L, Nurse JRC, Goldsmith M, Creese S. A Formalised Approach to Designing Sonification Systems for Network-Security Monitoring. In: *International Journal on Advances in Security*. Vol 10., 2017:26-47.
- [11] Bonebright, T. L. and Flowers, J. H. (2011). Evaluation of auditory display. In Hermann, T., Hunt, A., Neuhoff, J. G., (eds.), *The Sonification Handbook*, (Berlin, Germany. Logos Publishing House,) ch. 6.
- [12] Walker, B. N. and Kramer, G. (2005). Mappings and metaphors in auditory displays: An experimental assessment. In *ACM Trans. Appl. Percept.* **2**(4):407-412.
- [13] Roddy, S. (2015). *Embodied Sonification*. PhD Thesis, Trinity College Dublin. School of Engineering, Department of Electronic and Electrical Engineering.
- [14] Barrass, S., & Vickers, P. (2011). Sonification design and aesthetics In Hermann, T., Hunt, A., Neuhoff, J. G., (eds.), *The Sonification Handbook*, (Berlin, Germany. Logos Publishing House), ch.11.
- [15] Roddy, S., & Furlong, D. (2015) Sonification listening: An empirical embodied approach. In *Proceedings of the 21st International Conference on Auditory Display (ICAD 2015)*. (Graz, Austria, July 6-10, 2015).181-187.
- [16] Walker, B. N. and Kramer, G. (2005). Sonification Design and Metaphors: Comments on Walker and Kramer, ICAD 1996. In *ACM Trans. Appl. Percept.* **2**(4): 413-417.
- [17] Buchanan, R. (1992). Wicked Problems in Design Thinking. *Design Issues*, Vol. 8, No. 2, (Spring, 1992), pp. 5-21.
- [18] Zimmermann, J., Forlizzi, J. and Evenson, S. (2007). Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. (New York, NY, USA, Association for Computing Machinery), 493-502.
- [19] Barrass, S. (2012). The Aesthetic Turn in Sonification. Towards a Social and Cultural Medium. *AI & society*. **27**. 177-181.
- [20] Cross, N. (1982) Designerly way of knowing, *Design Studies*, Volume 3, Issue 4, October 1982, Pages 221-227.
- [21] Stappers, P. & Giaccardi, E. (2017) Research through Design., In Soegaard, M. & Friis-Dam, R. (eds.), *The Encyclopedia of Human-Computer Interaction*, 2nd edition.
- [22] Kramer, G. (1994), An introduction to auditory display. In Gregory Kramer, (ed.), *Auditory Display, volume XVIII of Santa Fe Institute, Studies in the Sciences of Complexity Proceedings*, (Reading, MA, Addison-Wesley), pages 1-78.
- [23] Lau, A. and Vande Moere, A. (2007) Towards a Model of Information Aesthetics in Information Visualization. In *Proceedings of the 11th International Conference Information Visualization (IV '07)*. (IEEE Computer Society, USA) 87-92.
- [24] Manovich, L. (2008) Introduction to Info-Aesthetics. <http://manovich.net/content/04-projects/060-introduction-to-info-aesthetics/57-article-2008.pdf> (accessed 31st March 2020)
- [25] Vande Moere, A. and Purchase, H.C. (2012) On the role of design in information visualization. *Information Visualization*, **10**(4), pp. 356-371.
- [26] Chion, M., (1994) *Audiovision: Sound on Screen*, New York, U.S.: Columbia University Press.
- [27] Beckerman, J. and Gray, T. (2015). *The Sonic Boom: How Sound Transforms the Way We Think, Feel, and Buy*. (New York: Mariner Books).
- [28] Truax, B., (2001). *Acoustic communication*. (Greenwood Publishing Group).
- [29] Friedman, K. (2008) Research into, by and for design, *Journal of Visual Art Practice*, **7**:2, 153-160.
- [30] Hutchinson, H., Mackay, W., Westerlund, B., et al. (2003), Technology probes: Inspiring design for and with families. In: *CHI03 Proceedings of the Conference on Computer-Human Interaction*, 17-24.
- [31] Hogan, T., Hornecker, E. (2016), Feel it! See it! Hear it! Probing Tangible Interaction and Data Representational Modality., *Des Res Soc.*, (Fort Lauderdale, Florida, 5-10 2003): 1-13
- [32] Roetzel, P.G. (2019), Information Overload in the Information Age: A Review of the Literature from Business Administration, Business Psychology, and Related Disciplines with a Bibliometric Approach and Framework Development. *Bus Res* **12**, 479-522.
- [33] Lenzi, S., Terenghi, G., Taormina, R., Galelli, S., Ciuccarelli, P. (2019) Disclosing cyber-attacks on water distribution systems. An experimental approach to the sonification of threats and anomalous data. In *Proceedings of the 25th International Conference on Auditory Display*, (June 23-27, Northumbria University). 125-132.
- [34] Winters, R.M., Kalra, A., & Walker, B.N. (2019). Hearing Artificial Intelligence: Sonification Guidelines & Results from a Case-study in Melanoma Diagnosis. In *Proceedings of the 25th International Conference on Auditory Display*, (June 23-27, Northumbria University), 262-267.

- [35] Axon, L., Happa, J., Goldsmith, M., and Creese, S., (2019) Hearing Attacks in Network Data: An Effectiveness Study In *Computers & Security*. **83**:367-388.
- [36] Hildebrandt, T., Hermann, T., Rinderle-Ma, S., (2014) A sonification system for process monitoring as secondary task. In: *5th IEEE Conference on Cognitive Infocommunications (CogInfoCom)*, (Salerno, Italy, 5-7 November 2014). 191–196.
- [37] Adhity, S. (2019). London Bus Tunes: Using Sound to Improve the Safe Navigation of London’s Bus System. In *Proceedings of the 25th International Conference on Auditory Display*, (June 23-27, Northumbria University). 287-290.
- [38] Bakker, S., van den Hoven E, Eggen B. (2012), Knowing by ear: Leveraging human attention abilities in interaction design. *J Multimodal User Interfaces*, **5**:197-209.
- [39] Vickers, P. (2011), Sonification for Process Monitoring., In Hermann, T., Hunt, A., Neuhoff, J. G., (eds.), *The Sonification Handbook*, (Berlin, Germany. Logos Publishing House), ch.18.
- [40] Ries, E. (2011). *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. (New York: Crown Business).
- [41] Lakoff, G., & Johnson, M. (2008). *Metaphors we live by*. (Chicago: University of Chicago press.)
- [42] Roddy, S. & Bridges, B. (2016). Sounding Human with Data: The Role of Embodied Conceptual Metaphors and Aesthetics in Representing and Exploring Data Sets. In *Proceedings of MusTWork16 – Music Technology Workshop 2016*. (10-16 June, Dublin, Ireland).
- [43] Vogt, K. and Höldrich, R., (2010). Metaphoric Sonification Method-Towards the Acoustic Standard Model of Particle Physics. In *Proceedings of the 16th International Conference on Auditory Display (ICAD2010)*. (June 9-15, Washington D.C., U.S.A.), 271-278.
- [44] Hilmann, N. & Pauletto, S. (2014). The Craftsman: The use of sound design to elicit emotions. *The Soundtrack*. **7** (1): 5-23.
- [45] Walker, B. and Kramer, G. (2005). Sonification design and metaphors: Comments on Walker and Kramer. *ICAD 1996., ACM Trans. Appl. Percept.* **2**(4): 413–417.
- [46] Vickers, P., & Hogg, B. (2006). Sonification Abstraite/Sonification Concr'ete: An 'Aesthetic Perspective Space' for Classifying Auditory Displays in the Ars Musica Domain. In *Proceedings of the 12th International Conference on Auditory Display (ICAD2006)*. (London, UK, June 20-23). 210-216.
- [47] Gaver, B., Dunne, T., and Pacenti, E. (1999). Design: Cultural probes. *Interactions* **6** (1): 21-29.
- [48] Moreno Fernandez De Leceta A., Lenzi S., Galán García P., Tiso D., (2019) A sonification model for real-time anomaly detection in machine learning supported cybersecurity, DOI:10.13140/RG.2.2.11419.21287.
- [49] Debashi, M., and Vickers, P. (2018) Sonification of network traffic flow for monitoring and situational awareness. *PLoS ONE*, **13**(4): 1-31.
- [50] Vickers, P., Laing, C., Debashi, M., & Fairfax, T. (2014). Sonification Aesthetics and Listening for Network Situational Awareness. In *Proceedings of SoniHED-Conference on Sonification of Health and Environmental Data*, (York, U.K., 12th September).