A WEB GUIDE TO PERCEPTUALLY CONGRUENT SONIFICATION

Joe Fitzpatrick

Centre for Creative Informatics
Limerick Institute of Technology
Limerick, Ireland
joe.fitzpatrick@lit.ie

Flaithri Neff

Centre for Creative Informatics
Limerick Institute of Technology
Limerick, Ireland
flaithri.neff@lit.ie

ABSTRACT

Sonification is an increasingly popular mechanism for data exploration, promoting the need for a greater understanding of human auditory perception and of how sonified information is designed, presented, and interpreted. In this paper, perceptual modelling is used to explore and demonstrate how perceptual phenomena are accounted for in sonification design. The framework, extracted from a larger body of work, links perceptual phenomena such as stream segregation to sonification mappings to provide a systematic approach to identifying and addressing perceptually-driven problems in applied sonification. A web guide functions to situate and guide designers through the complex theoretical constituents of auditory perception incorporated in the Perceptually Congruent Sonification (PerCS) framework. This web guide (hosted on sonification.ie) highlights and summarises the perceptual phenomena most relevant to sonification design and uses simple audio-visual interactions to demonstrate their effect. Preliminary qualitative feedback from a brief survey elucidates a small number of end-user concerns and comments.

1. INTRODUCTION

In recent years, there has been significant discussion regarding various sonification mapping strategies, with particular emphasis on the impact of auditory perception in user-evaluations [1] [2] [3] [4] [5] [6]. Each of these authors share concerns regarding the robustness of sonification practices and propose various solutions ranging from greater clarification of sonification methodologies [3] to borrowing from existing fields of research, such as design research and linguistics [6].

When compared to data-visualization, sonification is significantly underutilized. One explanation for this may be because sonification applications often require users to learn and remember which data variable is being mapped to a particular abstract sound. In addition, users are able to glance at visualizations quickly to aid their contextualizing of data, while sonification is inherently serial in presentation, making it more difficult to link data relationships that are not side-by-side on the serial timeline. The visual display of information also acts as a reliable form of external memory, allowing the user to repeatedly reference data without needing to memorize it themselves, while auditory presentation relies heavily on user-memory beyond the most immediately-presented datapoints. Sonic Information Design [7] and the use of Embodied Cognition [6] are two approaches that aim to make auditory displays naturally easier to interpret and recall. Both frameworks use pre-existing non-auditory research that examines how humans naturally interact with stimuli, the findings of which are then applied to sonification.

In terms of evaluation of the field, at present, there are only a few contemporary reviews of the field that gauge the academic communitys disposition to sonification as a means of data exploration. One such review featured in the proceedings of the International Conference on Auditory Displays 2018 (ICAD) discusses the efficacy of sonification methods and includes qualitative data from interviews with prominent sonification practitioners [3]. Quinton et al. found the practitioners were concerned with the following: poor design frameworks being implemented, the indistinct line between artistic and scientific sonification, failures in identifying the needs of the end-user, and the challenges of including user experience evaluation through the design process.

Some potential solutions being proposed by interviewees in Quinto et al. include using simple mappings, mimicking nature where applicable, and identifying the efficacy of mappings based on their psychoacoustic elements. Although some opinions promote that sonification as a technique was in decline, the majority of practitioners agree that sonification is becoming more widely accepted and that with the proper implementation of effective mapping strategies, it could become a powerful tool for data exploration. Such strategies, like those mentioned above, must rely on a solid foundation of existing research. Auditory Scene Analysis (ASA) has been referenced as one such body of work that can be utilized in sonification design.

2. AUDITORY SCENE ANALYSIS

ASA examines how humans form an auditory scene from the sounds of the environment around them. Bregman first introduced the concept of such processes, which draw parallels with other perceptual findings observed in linguistics as well as Gestalt psychology [8]. His specific use of the word stream is a good manifestation of how Bregman conceptualizes the multitude of individual sonic dimensions often culminating for listeners as single perceptual units representing individual sound-source entities. It is also telling that, in retrospect, any other alternative like noise or just sound is misleading and inefficient to adequately describe what ASA pertains to. Another important term in Bregmans glossary is the Principle of Exclusive Allocation; an analogue to what Gestalt psychologists refer to as the Principle of Belongingness. Bregman

https://doi.org/10.21785/icad2021.014

makes considerable reference throughout his book to the similarities between visual perception and auditory perception, and about the advantages of applying some of the frameworks from the former to the latter. He explains that Belongingness refers to how humans perceptually group sensory objects as a means of interpreting chaotic stimuli in a perceptual scene. Exclusive Allocation is a necessary and direct consequence of Belongingness, in that if we perceptually place one object into a group of similar objects, it cant then be allocated anywhere else. In order to explain this further, Bregman references the writings of Koffka, a Gestalt psychologist who drew on the possible links between visual and auditory perception [9].

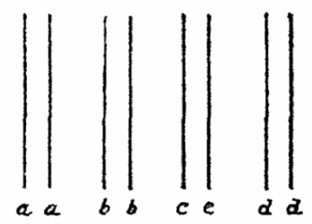

Figure 1: A visual example of perceptual grouping; the space between the letters acting as the principle cause of separation and integration [9]

Figure 1 is a visual example of perceptual grouping provided by Koffka. The most apparent grouping here is the stripes (the space between the 1st and 2nd a; 1st and 2nd b; and so on) that Koffka calls the figure. The rest of the image he refers to as the ground. This is contrasted to an audible pattern where each line represents a beat. In this auditory context, he refers to the ground as the stillness. The stillness referring also to the ever-present room noise level or what we know more commonly as the noise floor. These examples are both simple perceptual examples of how humans apply belongingness to organize both visual and auditory information from the senses.

According to Bregman, Perceived Continuity is another key concept that ties into both perceptual domains and is an example of perceptual closure. In the visual domain, Gestalt psychologists look at strong perceptual forms like squares and circles, referring to how the contour can be continued perceptually even if it disappears and reappears at various intervals. This accounts for the terminology behind Perceived Continuity in visual perception, however the term closure is what Bregman focuses on. In a visual representation where information/evidence is missing, and the overall shape is not strong, then the overall representation is hard to perceptually resolve. However, when the missing evidence is occluded with an arbitrary shape, the overall representation is suddenly made obvious (or resolved). The foreign and randomly-shaped occlusion perceptually completes the underlying image into one display or stream of information. This can be seen very clearly in one of his earlier papers [10]. Bregman then presents the auditory equivalent arranged by Dannenbring that compares a tonal glide interrupted by silence - or the Stillness as Koffka referred to it with a tonal glide interrupted by short broadband bursts [11]. The results are similar to those discovered by Gestalt psychologists whereby the unoccluded tonal glides fail to group as one stream. However, when they are connected by the broadband bursts, they stream as one and appear to continue on behind the tonal bursts. This sort of continuity provides an underlying structure to an auditory scene. Auditory demonstrations of this are currently available on Bregmans website. This website, and its auditory examples, has heavily influenced the design of the web guide proposed in this paper [12].

3. STREAM SEGREGATION IN SONIFICATION

Stream Segregation (SS) is one the core focuses of Bregmans work and is particularly relevant to sonification design. Even a low-dimensional dataset (e.g. weekly temperature) can have more than one column or data-stream (e.g. temperature in different locations). A sonification application used to represent such a dataset must decipher how to relay these independent data-streams in a way that a user can effectively distinguish each one of them. Say, for example, two independent temperature values are mapped to a frequency range. The simplest way to distinguish both is to only present one sonified data-stream at a time. However, in terms of complexity, this kind of parameter-mapping sonification - an auditory graph in this case - is as simple as it gets. When datasets are multidimensional (e.g. temperature and humidity), the resulting sonification application needs to be more complex if it seeks to be efficient or worthwhile. When tackling the challenges of using sonification for complex datasets, it is useful to have a basic understanding of not only how humans physically perceive sound but also how sound sources can perceptually merge. Stream Segregation as a concept refers to the ASA phenomenon where multiple streams either segregate or no longer segregate (merge) depending on how close they are in frequency, timbre, time, or space. For example, when the streams auditory attributes are varied or salient (in terms of tone complexity or amplitude) or distanced (in terms of frequency or time), the streams are more likely to remain segregated.

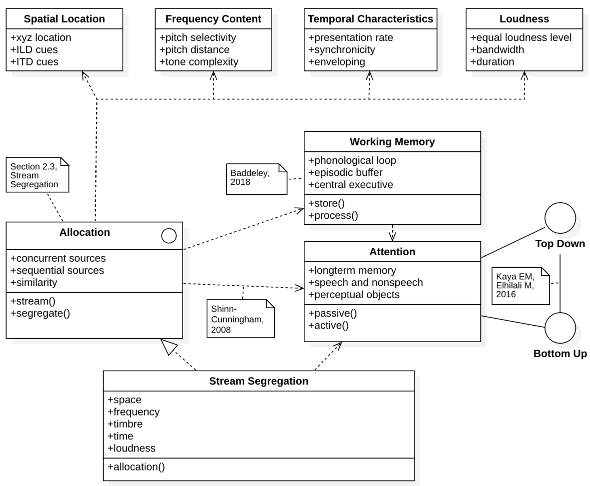

Figure 2: A visual model of SS. Implementation is denoted by a dotted line and hollow arrow; dependency is denoted by a dotted line and arrow

A proposed visual model of this phenomenon is presented above in figure 2 using the Unified Modeling Language (for its practicality in moving towards eventual model codification). In addition to the four above characteristics, loudness is also included as a core determinant of stream segregation. These factors can be described as attributes of the SS class. In other words, the SS class and whether it streams or segregates depends on these factors. Allocation is the fundamental role of SS and can be assigned as a method of the SS class. The attributes and methods of allocation can be illustrated further as an interface (denoted by a circle) that is implemented (or realized) by the SS class. The attributes include the number of concurrent sources, the number of sequential sources, and the similarity between the auditory characteristics of sources. The methods (or operations) of the allocation interface include streaming (fusion) and segregation. In addition to the SS class and its implemented interface Allocation, are the classes Working Memory and Attention. As denoted by the dotted line and arrow, the Allocation interface is also dependent on these classes, however, these are not the focus of this paper. The aim of the web guide is to identify and unpack what the Allocation interfaces attributes and methods are dependent on (spatial location, frequency content, temporal characteristics, and loudness), and examine how they might be relevant to sonification design strategies. This is reflected below in a second visual model.

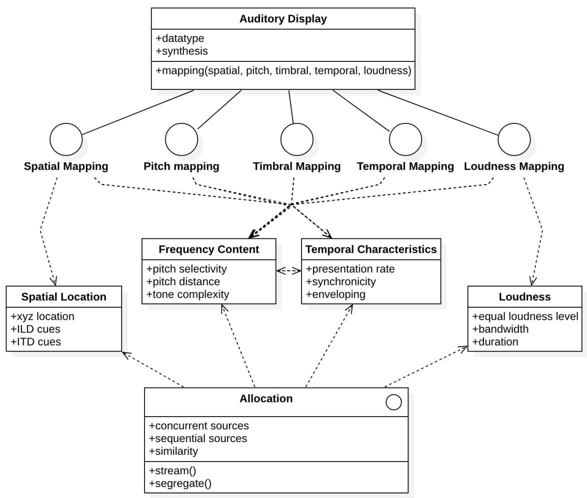

Figure 3: A visual representation of mapping-factor dependencies. Dependency is denoted by a dotted line and arrow. Note all mappings are affected by Frequency Content and Temporal Characteristics

In this visual model, the core determinants of stream segregation are linked to common sonification mapping strategies [1]. An extensive review of each determinant (displayed as classes in figure 3) and how they relate to sonification methods was recently completed as part of an ongoing PhD thesis. However, this research in its text format simply adds to the ever-accumulating body of sonification literature. In order to make the research more palatable to aspiring sonification practitioners, a web guide that highlights the most prevalent points was designed.

4. THE WEB GUIDE

HCI research suggests that interaction has a powerful impact on how humans learn and absorb information [13] [14] [15]. With this in mind, the author introduces the development of a web-based interactive PerCS guide, that utilizes the Web Audio API. The guide is designed to be accessible to aspiring sonification practitioners without a comprehensive background in auditory perception. As such, concepts that may appear obvious to sonification practitioners through experience or expertise are still addressed.

The aim of the web guide is to enable designers to input what auditory parameters they intend to use and get in return constructive feedback. Specific to their application, this feedback would highlight what perceptual concerns the designer needs to be aware of according to the PerCS framework. An interactive and dynamic web guide with helpful visualizations is feasible using three core web-based programming languages: HTML, CSS, and JavaScript, alongside a few external JavaScript-based APIs. The Web Audio API, for example, is a tool that can be used to present auditory examples of ASA phenomena such as belongingness and streaming. Such examples aim to provide designers with an immediate understanding of why an awareness of these phenomena is important when designing a sonification application with multiple auditory dimensions. In addition to this auditory component, interactions also include visual aids to further help designers understand the phenomena outlined by PerCS. These visual aids are designed using Three.js, a JavaScript 3D library.

In its current form, the interactive audio-visual web guide is a concise incorporation of the PerCS framework as outlined in the previous section. The current iteration of the application is built using the typical web-programming tools mentioned above (HTML, CSS, JavaScript, JQuery) in conjunction with the Web Audio API and Three.js. The web guide is hosted using Google Firebase, a web and mobile app development platform backed by Google, Inc. As of the writing of this document, the PerCS web guide is accessible via the domain sonification.ie or sonification-asa.firebaseapp.com/. A Github repository is also available at github.com/JoeFitzpatrick/PerCS, where the code can be reviewed in its entirety, including the code used to present auditory phenomena as audio-visual interactions.

On the home page, the user is presented with the overall aims of the PerCS Guidelines as well as the visual models of sonification design and stream segregation (presented in a slideshow). These visual models provide a brief insight into how sonification mapping strategies are directly linked to the determinants of stream segregation as outlined in figure 3.2. This introductory page also explains that the PerCS guidelines are complimentary to pre-existing sonification techniques and practices. Lastly, this introduction notes that while the guidelines are helpful to those implementing artistic sonification; they are most relevant to designers of more practical or scientific applications/displays. The guideline modules themselves are found on another page that is easily navigated to from each page via the menu-bar at the top of the page (Guidelines). These guidelines are broken down into five modules: Frequency and Spectral Content, Loudness, Spatial Content, and Multimodal Applications as seen below in figure 4.

The source material for these modules includes the authors PhD research as well as audio-visual interactions that aim to enhance the users understanding of the content. In addition, there are a number of links to other online tools that may be of use to the user. The language in these sections remains largely academic, however, deviates a little from this form to be more user-friendly. In terms of referencing sources, the modules follow the Harvard Referencing System and references can be found at the bottom of this page or on the References page. The most important module of the guidelines is the Frequency and Spectral Content module, as it is relevant to virtually all sonification applications and auditory displays regardless of the technique or strategy implemented. It

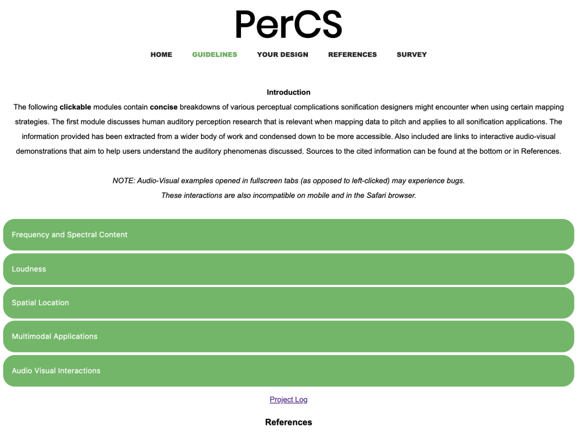

Figure 4: Screenshot of the guideline modules on the PerCS Web Guide

is, as such, the most content-heavy and where all the audio-visual examples are found. These interactions are here to help the designers understand how the perceptual phenomena discussed occur and how they might relate to sonification design.

The Your Design page simply asks what auditory parameters the designer intends to use and whether there is more than one modality being used. Using this information, the interactive web guide only brings up the relevant PerCS guideline modules. This page also gives brief examples of what is meant by mapping-pairs in the form of three tables. For example, it explains that a display where temperature is mapped to a pure tone, and humidity is mapped to a square tone, is utilizing pitch and timbre. An application where the location of a clock hand is mapped to a horizontal L/R location (panning value) is using space. The last example is if the height of a wave is mapped to amplitude, then the display is using loudness.

5. AUDIO-VISUAL INTERACTIONS

At this point in time, the PerCS web guide includes four audio-visual interactions that can all be found in the Frequency and Spectral Content module - a module that is suggested as relevant to all sonification applications. As mentioned, the interactions were designed using HTML, CSS, JavaScript and two additional APIs: Three.js and Web Audio. The URL links to these interactions are made noticeable using GIFs that preview the functionality of each one. All the interactions can be experienced here: sonification.ie/resources.html under the Frequency and Spectral Content module (that provides the necessary context) or under the Audio Visual Interactions module (https://sonification.ie/resources.html).

The first interaction is a two-channel audio-visual example that demonstrates 'streaming'. All the interactions are triggered by clicking and holding the left mouse click. In this interaction, the cursor is used to change the pitch of the left channel and can be moved along the x-axis towards, away or beyond the middle line where the right channel plays at around 1kHz. Two channels are used in this interaction to avoid demonstrating phasing and instead demonstrate how two streams can still perceptually merge as one despite being binaurally separated. This is a consequence of binaural fusion and perceived pitch, which are discussed in the spatial location module of the web guide. The main purpose of this interaction, however, is to show that using multiple pure tones for data variables that need to be individually distinguishable in applications such as auditory graphs can be problematic.

The second interaction shown in figure 5 is a simple interactive audio-visual example of temporal groupings where tones perceptually group because of how close they are to each other (temporally). Once again, the tone sequence is started by clicking and holding the left mouse click. Within this interaction there is a toggle in the top left that can be toggled on with the left mouse click or by pressing the tab key followed by the spacebar (standard browser tab-selecting). This toggle allows the user to contrast the temporally-paired tones with temporally-unpaired tones. The visual element of this interaction borrows from Koffkas visual example of belongingness depicted earlier in figure 1. In this interaction, the 'stillness' or ever-present noise-level is represented by noise.

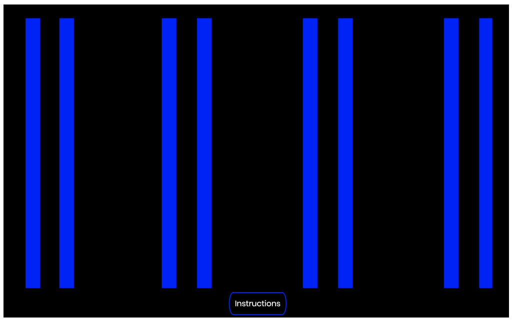

Figure 5: Audio-visual interaction demonstrating belongingness

The third interaction shown in figure 6 is an audio-visual tool that demonstrates the effects of various presentation rates. The cursor's location on the x-axis dictates the starting frequency. The cursor's location on the y-axis dictates the tone-sequences tempo. In this interaction, the toggle in the top left allows the user to loop the predefined tone sequence, which is recommended when listening for apparent frequency groupings. The users are instructed to start by listening to the sequence at a lower tempo first, where the order of the notes should be easily perceived, and then to try a faster tempo. It is at the higher speeds that they are instructed to listen for any frequency groupings. This interaction also shows the link between the visual phenomenon apparent motion and auditory grouping.

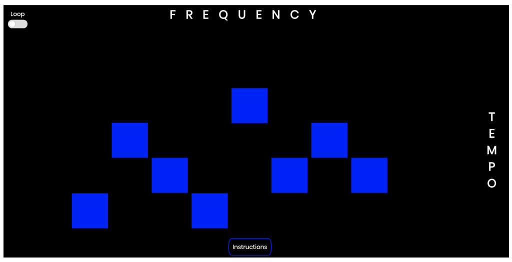

Figure 6: An audio-visual interaction that enables users to experience the effects of varying presentation rates on frequency groupings

The last interaction shown in figure 7 is an audio-visual tool that can be used to demonstrate the streaming effect apparent in concurrently sonified data-streams, as well as the use of a saw-tooth wave to focus or isolate a particular trend. Each channel can be toggled on or off using the toggles at the top of the display. Each colored channel may also be 'focused' - i.e. converted from a sine wave to a sawtooth wave. The users can also refresh the page (CMD/Ctrl-Shift-R) to randomize the data points on the graph. Users can also test their ability to isolate and identify the focused trend by clicking the top right box in the display. This separate interaction lets them guess which trend is isolated before providing the answer.

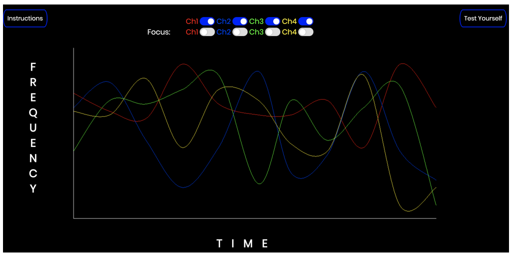

Figure 7: An audio-visual interaction that enables users to listen to concurrently sonified trends and isolate particular trends

It is worth noting here that this interaction is provided along-side research that shows that the first four harmonics of a sawtooth wave are more than enough to isolate a trend [16]. A note in the instruction box clarifies that the sawtooth wave in this interaction is not simply the first four harmonics of a tone, however, but is simply used to demonstrate one way of segregating sonified streams.

6. PRELIMINARY FEEDBACK

In order to collect feedback on the PerCS web guide, a short Google Forms survey was added under the Survey page of the website. In addition, the survey and web guide were circulated via a number of email lists, including both the International Conference on Auditory Displays list and the more wide-reaching Auditory List (initially started by Albert Bregman in 1992 and now maintained by Dan Ellis). The following results are from twelve participants who claimed to be familiar with sonification practices. Two questions included in the survey ask how accessible the overall guide is and how relevant they think the written content and figures are to sonification practices. Participants were provided a five-point Likert scale to respond.

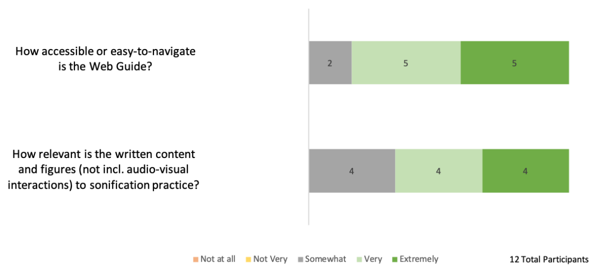

Figure 8: Likert responses to questions on the web guides accessibility and content relevance

As shown in figure 8, most of the participants found the web guide easy to navigate, suggesting that the overall layout and design of the website doesnt require additional work. In relation to how relevant the research content of the site is to sonification, the majority of participants again responded positively. There is, however, room for improvement; perhaps by relating the auditory perception research more directly to sonification mapping strategies or by providing specific examples. Figure 9 displays the participants Likert responses to two questions that aim to draw out qualitative feedback on the audio-visual interactions specifically. Like the last pair of questions it pertains to the accessibility of the interactions and their relevance to sonification.

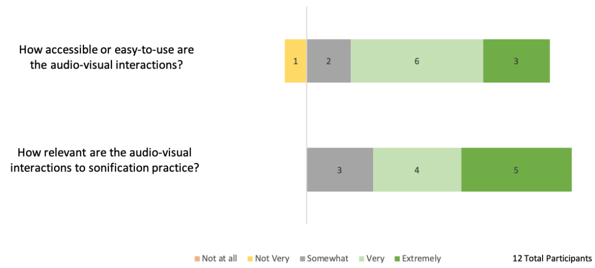

Figure 9: Likert responses to questions regarding the accessibility and relevance of the audio-visual interactions

Although these preliminary responses remain positive, participants found the interactions less accessible than the other written and visual content of the web guide overall. This response is somewhat expected as the interactions require more engagement and understanding of the context in which the interactions are found. This result might be improved with more detailed instructions on how to engage with the interactions or by adding instructional video demonstrations. Although no technical issues were specified with these responses, browser-version compatibility with the Web Audio API or Three.js may also be related. As seen by the Likert response to the second question, the majority of participants were again satisfied with how relevant the audio-visual interactions were to sonification. This preliminary result suggests that the context provided sufficiently explains how the interactions pertain to sonification practices.

The last part of the survey gives the participants a chance to suggest further research they think might be relevant to PerCS, while the last question looks for suggestions for other potential audio-visual interactions that relate to sonification and auditory perception:

- *What research relevant to auditory perception and sonification is missing from this web guide?*
- *Can you suggest any other potential audio-visual examples or interactions that relate to auditory perception and are relevant to the design of auditory displays?*

In response, several participants suggested incorporating music-perception research as they have found such research benefits auditory display design just as much as traditional auditory perception research. In particular, it was proposed that research regarding timbre perception in music and the work of McAdams and his colleagues might significantly contribute to the PerCS web guide [17] [18] [19]. While the web guide does include some of this research (under the Frequency & Time module), a more comprehensive collection of works relating to music perception as a whole might identify potentially relevant information to aid sonification design.

Other suggestions included providing an independent module within the guide that highlights audio evaluation strategies and tools that are currently available online. Furthermore, it was proposed that a brief review of each and their core strengths in relation to sonification design would be a useful and relevant asset within the PerCS web guide. More specifically, it was suggested that a review that identified and outlined which evaluation approaches were best for certain sonification would be invaluable.

In relation to potentially helpful audio-visual interactions, one participant suggested that including interactive examples and explanations of more classic ambiguities and illusions might expose some useful information for sonification practitioners. Examples in their suggestion included the McGurk Effect, Shepard tones, and the Tritone Paradox. Such examples are similar to the auditory illusions presented by Diana Deutsch [20] [21]. Additional research that reviews and explains such illusions might indeed be a useful addition to the web guide, perhaps again in its own separate independent module.

Beyond the survey, a number of suggestions and comments were made by sonification practitioners and researchers via email. One suggested including more information on the role of the spectral centroid in timbre perception as it is a relatively easy-to-manipulate auditory parameter that can be mapped in auditory displays. Such research is included in the web guide to a certain extent, however this could be more comprehensively discussed in relation to sonification design.

In addition to feedback regarding the current iteration of the web guide, some researchers offered advice on the future direction of the project. In their emails they proposed the next logical step from PerCS (and other similar research) might be to incorporate functions and models of such research directly into synthesis software solutions under the banner of Auditory Scene Synthesis. They add that the early groundwork for justifying and implementing Auditory Scene Synthesis in future sonification projects has already been proposed and might benefit from further development [22]. The SoniPy framework, for example, is an early example of Auditory Scene Synthesis being applied [23].

7. FUTURE WORK

The overall objective of the PerCS framework is to encourage developers in sonification design to more heavily consider and discuss prevalent phenomena in auditory perception. The web guide aims to facilitate not just experienced practitioners but also data analysts without a strong background in sound who still want to pursue sonification as a means of data exploration. More specifically, it proposes using ASA principles to further investigate such phenomena (like SS) that are relevant when designing robust auditory displays. The aim is that if more publications relating to sonification share how practitioners tackle the phenomena highlighted by the PerCS framework, the robustness of all future auditory displays can improve. Additionally, if the sonification community appears to be supportive of such principles, the idea of using sound to relay data might be more broadly accepted.

In its current state, the interactive web guide aims to make the overall perceptually congruent sonification framework more accessible by supplying only pertinent research-supported information to designers using written, visual, and audio feedback. The survey outlined in the previous section is still available on the website (https://sonification.ie) and will continue to collate feedback on the web guide for future analysis. Early qualitative results are sparse, yet promising, and have included helpful suggestions for future iterations. In summary, the core components of future work on PerCS will include further examination and implementation of music perception research and the integration of auditory scene synthesis functions supported by the PerCS guidelines.

The next phase of this research is to expand on and review the overall Perceptually Congruent Sonification framework as well as ascertain more background information on how perceptual conflicts are considered in current sonification projects. This will be in the form of a systematic review of such considerations and will include feedback from the practitioners on how they self-evaluated their design. The ultimate aim of this phase will be to identify what aspects of auditory perception are largely unaccounted for in common sonification designs and how relevant practitioners believe such information to be.

8. REFERENCES

[1] G. Dubus and R. Bresin, "A systematic review of mapping strategies for the sonification of physical quantities," *PloS one*, vol. 8, no. 12, p. e82491, 2013.

[2] A. Supper, "Sublime frequencies: The construction of sublime listening experiences in the sonification of scientific data," *Social Studies of Science*, vol. 44, no. 1, pp. 34–58, 2014.

[3] M. Quinton, I. McGregor, and D. Benyon, "Investigating effective methods of designing sonifications," in *Proc. ICAD*. Georgia Institute of Technology, 2018.

[4] J. Ferguson and S. A. Brewster, "Evaluation of psychoacoustic sound parameters for sonification," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, 2017, pp. 120–127.

[5] S. Chabot and J. Braasch, "A framework for evaluating perceptual interactions of various dimensions of sound for data sonifications," *The Journal of the Acoustical Society of America*, vol. 143, no. 3, pp. 1750–1750, 2018.

[6] S. Roddy and B. Bridges, "Mapping for meaning: the embodied sonification listening model and its implications for the mapping problem in sonic information design," *Journal on Multimodal User Interfaces*, vol. 14, no. 2, pp. 143–151, 2020.

[7] S. Barrass, "Sonic information design," *Journal of Sonic Studies*, vol. 17, 2018.

[8] A. S. Bregman, *Auditory scene analysis: The perceptual organization of sound*. MIT press, 1994.

[9] K. Koffka, "Perception: an introduction to the gestalt-theorie." *Psychological Bulletin*, vol. 19, no. 10, p. 531, 1922.

[10] A. S. Bregman, "Asking the what for question in auditory perception," *Perceptual organization*, pp. 99–118, 1981.

[11] G. L. Dannenbring, "Perceived auditory continuity with alternately rising and falling frequency transitions." *Canadian Journal of Psychology/Revue canadienne de psychologie*, vol. 30, no. 2, p. 99, 1976.

[12] A. Bregman. (1996) Demonstrations of auditory scene analysis: The perceptual organization of sound. [Online]. Available: http://webpages.mcgill.ca/staff/Group2/abregm1/web/downloadstoc.htm

[13] E. Gaudioso and J. G. Boticario, "Towards web-based adaptive learning communities," in *Proceedings of the 11th International Conference on Artificial Intelligence in Education (AIED2003). Sidney, Australia*, 2003.

[14] S. Polovina and W. Pearson, "Communication+ dynamic interface= better user experience," in *End-User Computing: Concepts, Methodologies, Tools, and Applications*. IGI Global, 2008, pp. 419–426.

[15] G. Stahl, *Essays in computer-supported collaborative learning*. Lulu. com, 2017, vol. 9.

[16] J. Fitzpatrick and F. Neff, "Stream segregation: Utilizing harmonic variance in auditory graphs," *Proceedings of SMC*, vol. 18, 2018.

[17] A. Caclin, S. McAdams, B. K. Smith, and S. Winsberg, "Acoustic correlates of timbre space dimensions: A confirmatory study using synthetic tones," *The Journal of the Acoustical Society of America*, vol. 118, no. 1, pp. 471–482, 2005.

[18] G. Peeters, B. L. Giordano, P. Susini, N. Misdariis, and S. McAdams, "The timbre toolbox: Extracting audio descriptors from musical signals," *The Journal of the Acoustical Society of America*, vol. 130, no. 5, pp. 2902–2916, 2011.

[19] S. McAdams, "Musical timbre perception," *The psychology of music*, pp. 35–67, 2013.

[20] D. Deutsch, "The octave illusion revisited again." *Journal of Experimental Psychology: Human Perception and Performance*, vol. 30, no. 2, p. 355, 2004.

[21] ——, "Illusions and research." *Authors Website*, 2004. [Online]. Available: http://deutsch.ucsd.edu/psychology/pages.php?i=201

[22] D. Worrall, "The sonipy framework: Getting started," in *Sonification Design*. Springer, 2019, pp. 181–211.

[23] D. Worrall, M. Bylstra, S. Barrass, and R. Dean, "Sonipy: The design of an extendable software framework for sonification research and auditory display," in *Proc. ICAD*, 2007.