Sonifications Sometimes Behave So Strangely

Paul Vickers

The comprehension of phenomena by analyzing and exploring data collected for the purpose is an old and established practice. Statistical methods have become quite sophisticated and are the bedrock of much modern scientific enquiry. Ever since William Playfair introduced the line, area, and bar charts (1786) and the pie chart and circle graph (1801) to the world, the field of information visualization research has refined and extended his ideas and has developed rules and heuristics for the visual representation of data. In all of this, it is not evident that the ontological nature of vision has been taken into account. And why would it be? Phenomenologists and anthropologists have presented varied and competing theories as to how we perceive the world visually, but it seems that much of that can be bracketed when it comes to choosing how to lay out a plot or a chart.

Sonification is a family of representational techniques that use non-speech audio to communicate data and data relations (think Geiger counter for data). With its recent use in the discovery of gravitational waves, sonification has begun to gain some cultural traction, but for the most part it lacks the ubiquity and acceptance of its graphical cousin, information visualization. The term "sonification" was adopted to describe the use of non-speech sound for communicating data and data relations, and when Greg Kramer established the International Community for Auditory Display and its associated conference series, the International Conference on Auditory Display in 1992, the emergent field of sonification research put down roots.

The idea of sonification at first seems so simple: take some data values and use them to control the properties of an acoustic signal such that listening to the signal reveals something about the data or the data relations that are driving it. Tools like the Sonification Sandbox (Walker and Cothran 2003) make this process very easy, generating auditory graphs that step through tabular data with each value altering the pitch of a chosen tone.

Following the emergence of affordable digital audio processing hardware in the 1980s and 1990s researchers began to investigate the possibilities afforded by the auditory modality for data and information analysis and exploration. As they began to explore more deeply the use of sound as a complement to (and in some limited cases, a replacement for) visual display techniques, it became evident quite early on that unlike visualizations, and to borrow from Diana Deutsch (Deutsch, Lapidis, and Henthorn 2011), sonifications “sometimes behave so strangely”. There was something about the auditory representation of data that meant issues of ontology and phenomenology kept raising their (often unwelcome) heads. Unlike graphs, which do not immediately come across as paintings or pieces of visual art, sonifications kept raising questions of their relationship to music and the sonic arts. From an engineer’s, computer scientist’s, or even psychologist’s point of view, all of whom in the early days of the field were trying to find good ways to map data to sound without any composerly intent, sonification is not music. And yet, as Deutsch (Deutsch, Lapidis, and Henthorn 2011) rediscovered (Pierre Schaeffer arguably being the first to document the phenomenon with his account of the *sillon fermé* (Schaeffer 1967)) the mind, regardless of our volition, sometimes adopts a musical orientation to listening (Vickers, Hogg, and Worrall 2017).

Sonification, it goes, ‘is not visualization for the ears, it follows completely different rules’ (Kosara 2009). At one level this is perfectly obvious and self-evidently true for vision is (primarily) spatial and hearing is (primarily) temporal. A graph persists over time, the whole can be seen at a single glance, and it may be compared side-by-side with another graph. But the physical phenomenon of sound exists only in its production. To experience an entire sound requires it be listened to as it unfolds over time, and comparing one sound with the memory of another is fraught with difficulty.

Further, Cartesian dualism holds that perception involves an outside that we see, hear, feel, smell, and taste which we then internally interpret by cognition to form an understanding of the world. This fits very well with a bottom-up account of sensory processing. But in recent years there has been a shift in understanding of perception, from the Cartesian dualism of body and mind to an embodied phenomenological account which involves the ‘whole organism in its environmental setting’ (Ingold 2000, 258), an understanding which has been embraced by the third wave of human-computer interaction (HCI) research.

Information visualization has gradually accreted conventions for the visual layout of data. Guided by writers such as Jacques Bertin (Bertin 1981) and Edward Tufte (Tufte 2001) standardized techniques and aesthetic heuristics have been adopted. In contrast, since the inaugural International Conference on Auditory Display in 1992 the question of how best to specify the data-to-sound mappings remains, to a larger extent, an open one in sonification research.

Certain physical properties of sound are well understood thanks to the extensive body of psychoacoustic literature. Equal loudness contours, the relationship between perceived pitch and loudness, and so on are well documented and can be factored into sonification designs. Rules for some types of sonification have been proposed, such as John Flowers' heuristics for successful auditory graph design (Flowers 2005), with pitch being used as the main carrier of data values. But, as Bruce Walker's program of work demonstrated (Walker, Kramer, and Lane 2000), there is no universal property obtaining to the polarity of data-to-pitch mappings; some data are better understood where a rise in value corresponds to a rise in pitch, while others work seem to work better the other way around. A partial explanation for this might be that we associate sounds with real world events. While we see objects, we do not hear them, rather we hear the sounds they make, that is we hear events (Rosenblum 2004). Further, the sounds objects and events produce give us knowledge about the objects' size, density, and type. Low frequency sounds typically belong to heavy, dense objects so an increase in weight might be sensibly sonified with an inverse pitch mapping. On the other hand, physical height conceptually works the other way around, so the greater the height, the higher the pitch of the sonification will be.

Psychoacoustics is based largely on a laboratory-based bottom-up information processing model (Clarke 2005) in which raw sounds are given meaning by attending to their context, what has been heard most recently, prior listener training, experience, and so forth. In this model the physical properties of sounds are decoded, then cognition is employed to classify the sounds according to their form, organization, rhythm, and so on. Then at the top level the listener applies social and cultural filters to attribute aesthetic value, meaning, and any referential properties (see Clarke 2005, 11–14). As the sensory interrelatedness of perception and our interactions with the environment lead us to needing to embrace an embodied account of perception, we discover that sonification becomes much more complex than we first thought. As John Neuhoff realized, we need to discuss real-world psychoacoustics in terms of ecology and embodied experience (Neuhoff 2004). Al Bregman's magisterial work on *Auditory Scene Analysis* (Bregman 1990) serves as a stepping stone between this bottom-up information processing Cartesian dualistic approach to perception and the rich embodied experience it is being seen as by many today.

Sonification listening may be said to be an embodied, interactional, and practically situated activity. Interaction can be with the sonification tool itself, as in the case of Interactive Sonification (see Weinberg and Thatcher 2006), but also with the environment and space in which the listening takes place. Sonification is a lot more interesting than lab-based stimulus-response tests. Within information visualization there are some established aesthetic principles which, if followed, are deemed to lead to more successful representations. That is, representations that the intended user is able to read and understand without confusion or ambiguity. For example, consider graph layout aesthetics, such as the goal of minimizing the number of edges that cross each other in order to reduce the visual complexity. At this point it is not yet clear what an aesthetics of sonification entails or even if such a thing exists. Music philosophy has several competing aesthetic accounts but, as has been pointed out repeatedly elsewhere, sonification is not music, that is, it is typically not designed with composerly intent or with the goal of producing a musical aesthetic experience. Indeed, if one looks at sonification through the various lenses of music philosophy it appears to inhabit the (musically) contradictory position of Referential Formalism. It is referential because its very purpose is to point the listener to something beyond itself (the data) yet also formal because the meaning of the sonification lies within its syntactic and organizational structures.

If the view is taken that aesthetics deals with sensory perception (Vickers, Hogg, and Worrall 2017; Barrass and Vickers 2011) (and this appears to be the reason why graph aesthetics have been developed) then a way to approach the question of sonification aesthetics is to come at it pragmatically in terms of how we might design sonifications that are, as Stephen Roddy puts it, ‘communicatively effective’ (Roddy 2015).

How do we choose the mapping?

How, then, do we choose the mapping? How does the translation of data into sound affect the data and our understanding of it and how do we come to decide to translate those data through particular sounds and not through others (which might influence how we attribute meaning to the data)? An ungenerous answer to the question (from looking at many of the sonifications put forward over the last quarter of a century) is that a great deal of thought was not always given to this aspect. This is, of course, unfair, and belies much serious consideration, but there is a sense in which much early sonification work was motivated by the novelty of simply being able to map data to sound. Questions of aesthetics were usually limited to whether or not the sonification sounded pleasant and there also appears to have been an underlying assumption that sonifications should be easy to use, that is, easy for the listener to understand that information being communicated (more on this later).

More recently, there has been a deeper interrogation of how we listen to sonifications, what role the aesthetic plays in the experience and the nature of the relationship between sonification and the sonic arts (including music). This has been informed largely by the aesthetic turns in the field of HCI which moved from the functional approaches of traditional HCI through considerations of user experience informed by a pragmatist aesthetics (Barrass and Vickers 2011) to today's third-wave which deals with the phenomenological nature of embodied perception and interaction, for which Richard Shusterman coined the term "somaesthetics" (Shusterman 1999). Stephen Barrass and I put forth the case for sonification to consider these pragmatist experiential ideas in thinking about sonification aesthetics (Barrass and Vickers 2011) and Bennett Hogg, David Worrall and I (Vickers, Hogg, and Worrall 2017) took this further by directly addressing the question of embodied perception in sonification design. This was motivated by questions around the nature of sonification listening, the directness of a sonification, and the prior listening experiences of the sonification user. The question now becomes "how might we in future decide on the mapping?" Such an enquiry affords the opportunity to consider the factors involved in sonification as an embodied and interactional listening experience. Just as no 'widely accepted model of an aesthetic interaction' exists (Lenz, Hassenzahl, and Diefenbach 2017, 81) so is there no current definition of an aesthetics of sonification. However, as we move from the very functional view of early sonification research to considerations of the somaesthetic issues, then three factors become very important in the design of sonifications: directness, space, and listening, and I will address these below.

Directness

The choice of sound depends, in large part, on the type of sonification approach adopted. Sonification approaches span a continuum from the very direct, indexical processes involved in audification to the conventional representations (in semiotic terms) used in parameter mapping sonifications which can be very indirect and highly metaphorical. In audification the dataset defines the sonification as it involves transposing the frequencies of a time series dataset into the human audible range, together with any necessary filtering to remove unwanted linear distortions and occasional dynamic range compression to flatten out large variations in sound level. Because the data itself is transposed such that each data value effectively becomes an individual sample in a digital audio signal, the resultant auditory stream is very direct and tightly coupled to the dataset. The choice of what sound to use then becomes one of what filtering and scaling to apply to the signal in order to best make the audification "readable" and fit for purpose (see Dombois and Eckel 2011 for a fuller treatment of audification).

When it comes to sonifications in which there is no inherent link between the data and the chosen sounds, the directness of the representation is determined by the mapping strategy chosen by the sonification designer. Perhaps the most direct sonifications that use the data to drive the parameters of an audio signal are auditory graphs. They are so called because just like a visual graph maps one dimension (typically time) of the data to the abscissa and the values of the data to the ordinate, an auditory graph represents the abscissa by elapsed time and the data values by some change in the audio signal. The simplest way to effect this is to control the frequency of a sinusoidal oscillator with the data values. A high value gives a high pitch, a lower value a lower pitch. As each data value is plotted the pitch of the signal rises and falls accordingly. Historically, pitch has been most often chosen in auditory graphing and parameter mapping strategies alike. For auditory graphs it is a simple but effective mapping. For parameter mapping sonifications pitch seems to have been chosen as often for its ease of implementation as for any other reason.

Directness is a multivalent term in sonification as different writers have used the word to express different ideas about the relationship between sound and data. For example, Till Bovermann et al (Bovermann, Tünnermann, and Hermann 2010) use directness as a measure of the responsiveness of an auditory display, such that user interactions lead to quick changes in output. By contrast, and taking a steer from semiotics, Bennett Hogg and I (Vickers and Alty 2006) viewed directness as the conceptual distance between the data and its mapping, that is, a measure of the arbitrariness of the data-to-sound mapping. For example, a symbolic mapping involving sonic metaphors that stand for features of the data (e.g., the use of real-world sounds such as bird song and frog croaks to represent features of network traffic (Debashi and Vickers 2018; Vickers, Laing, and Fairfax 2017)) is an arbitrary mapping in the sense that the sounds chosen bear no direct relationship to those data or phenomena represented. Contrast this with an audification in which the sound generated is directly caused by the scaling of the data. There arises, then, a question as to what sort of mapping is best (indexical or symbolic), a question which, at this point in time, remains unanswered. A representational view of sonification holds that the data being referenced should somehow be a part of how the sonification is properly experienced so that the sonification is experienced in terms of the data it represents (Vickers, Hogg, and Worrall 2017, 96). If Deniz Peters is correct in his assertion that 'an essential part of our listening experience draws on what our own body suggests might have gone into the making of [a] sound' (Peters 2012) then the directness of a sonification's mappings ought to play a very important role in how successful the sonification is at communicating its underlying data. On the face of it, the mappings from data to sound should be as direct as possible with the implication that the more symbolic a mapping is, the less successful it might be. To this end, Robert Höldrich and I have begun work to explore how to implement good direct mapping strategies which we call "Direct Sonification" (see https://paulvickers.github.io/DSSon/). But this view does not account for the occasions when a symbolic mapping might be considered by the listener to be direct. For example, if one wished to sonify the comings and goings of worker bees in a hive over the course of a week, sensors could be added to register each time a bee arrives and leaves and this data could be mapped to a buzzing sound that mimics that of a bee in flight. This is not a direct mapping in the sense that the data themselves are not the cause of the sound (in the way that they are in audifications, or in the Direct Sonification mentioned above); because the data are generated by the activity of bees, and the sounds are of bees, one could argue that the data have become part of the sonification experience and are thus an authentic representation.

The idea that the more (causally) direct a mapping is, the less conceptual distance there is between the data and the sonic parameters, the more likely a sonification is to be successful is an attractive one. The more complex and richer the mapping, the greater the possibility that artefacts of the sonic rendering will be mistaken for properties of the data. For example, the use of tonal musical frameworks and rhythms could lead to expectations and understandings on the part of the listener that are based in the listener's prior experience rather than pointing to characteristics of the data. Perhaps a particular chord sequence is generated by a particular combination of data, a sequence that calls to the listener's mind a meaning that is not intended and which leads to incorrect inferences being drawn. This is one of the reasons why Hogg, Worrall, and I began a program of work to explore how accounting for the subject position in sonification design might lead to clearer, less ambiguous renderings (Vickers, Hogg, and Worrall 2017).

So far, we have considered the translation between data and sound only as a one-way activity, but we do need to consider the effect the rendering might have on the data. Of course, the objection is immediately raised that such an effect is impossible; how can any sonification affect the data it represents? It cannot, in any real sense alter the data values, or the underlying phenomenon from which the data were measured. The user can, of course, on listening to a sonification, choose to change the phenomenon or system which was being sonified. For example, if I am sonifying my heart and respiratory rates during exercise, the feedback might cause me to increase or decrease my activity which will, in turn, lead to changes in my heart beat and breathing. But here the sonification is a messenger, not an actor. Alternatively, and this is perhaps the more interesting consideration, the sonification might influence the way we interpret the data, leading us to change the way we perceive it, a sort of auditory version of seeing something in a new light: it causes us to appreciate the data, or the phenomenon from which it was measured, anew. The phenomenon hasn’t actually changed, but it certainly appears different than before.

Space and listening

The act of listening to a sonification is always situated within a space. Sonifications can be designed for monophonic, stereophonic, or multi-channel sound, or three-dimensional playback. If headphones are used then virtual listening spaces and ambiences can be created using combinations of convolution reverberation, binaural recording and reproductions techniques, ambisonics, head-related transfer functions (HRTFs), surround sound, and so forth. To create multi-channel or three-dimensional sound fields without headphones requires multi-loudspeaker arrays, or sophisticated equipment such as Sonible’s IKO, an icosahedral loudspeaker that employs beamforming and ambisonics to create a three-dimensional sound image (Sonible Gmbh, n.d.).

In the early days the majority of sonifications were designed for stereo playback either with headphones or the small loudspeakers commonly used with desktop computers. The focus here was on producing the data-to-sound mappings with little regard given to the listening experience. Headphones provide convenient isolation to reduce the effect of environmental noise during listening tests and also allow experiments to be conducted with multiple participants in a single laboratory. Experimental hypotheses revolved around whether the use of sound (either on its own or in conjunction with a visual display) improved participants’ ability to construct knowledge about the data. Even when spatial audio reproduction systems were used, the focus was largely on whether spatial audio could be used to communicate information rather than on the listening experience as an interactional embodied activity.

When we consider the subject position and think about designing for embodied experience, we begin to realize that the sonification designer’s past experiences, listening skills, and frames of reference could be very different from those of the intended listener. As Karin Bijsterveld observes, sonification designers tend to have ‘trained ears’ (Bijsterveld 2019, 104) and it is not always going to be the case that the intended listener will have developed their listening skills to the same extent. In the case where the listener and the sonification designer are not the same person, such as when designers and domain experts come together to collaborate on producing sonifications for the domain experts it is entirely possible that what the designers are able to infer from the sonification is not the same as the listeners whose data is being sonified.

Not only does the mapping itself affect how we perceive and experience a sonification, but the spatial aspects of the presentation also play a role. Gerriet Sharma's concept of the "Shared Perceptual Space" (Sharma 2016) provides a framework for exploring the sculptural aspects of spatial audio and how to approach the perceptual issues that arise during spatial audio production (Wendt et al. 2017). The Shared Perceptual Space is the space 'within which the perceptions of composers, scientists and audience intersect in respect of three-dimensional sound objects' (Sharma 2016, 3). With it, Sharma discovered that he could construct generalized descriptions of sound objects and that the 'collisions of perceptions gradually informed the ensuing compositional process and led to an expanded understanding and a different practice of artistic work with these phenomena' (*ibid*). The idea of 'situated perspective' (Harrison, Tatar, and Sengers 2007) has gained traction in the wider field of HCI but sonification research has not yet caught up. Even if a sonification is to be designed for stereo headphone presentation, it would still be instructive to consider the situated perspective of the listening and to use concepts such as Sharma's to explore how better to design and construct sonifications. When moving to more ambitious spatialized presentations we can ask questions such as what is the impact of spatial attributes (foreground/background, inside/outside, high/low, 2D/3D, direction) on perception of spatial sound-textures produced by mapped data? How can an understanding of shared perception inform and improve sonification design?

Listener experience

All of the above inexorably draws us to consideration of the listener, both in terms of the embodied experience that occurs during listening, as well as the listener's past experience, skill, and knowledge. The subject position is the stance a listener adopts towards the objects of perception (Clarke, 2005). Designing for the subject position is about careful direction of the listener towards what the sonification designer desires to reveal about the underlying data. That is, the 'aesthetic enters at the point of constructing the subject-position such that … something in the aesthetic of the sound has to match the phenomenon being revealed' (Vickers, Hogg, and Worrall 2017, 105). This, coupled with knowledge gained from understanding the shared perceptual space, lets us focus on the embodied interactional experience of sonification listening.

However, in our endeavors to address the complexities of embodied listening experience it is easy to fail to deal with listener skill. It has often been assumed that sonification should be designed so as to be as easy as possible to listen to, to require as little training as possible to use. Sometimes this is because the experiments to evaluate the usefulness of a sonification are designed to be run over short periods with large groups of listeners who are typically not domain experts (undergraduates are often recruited as participants for this purpose). Other times it may be motivated by the fact that sonification still often fails to be treated as a serious field of scientific research and enquiry, and so designers have felt that sonifications that are not simple to use will be quickly dismissed. However, it has long been accepted in other fields that sound-based exploratory tools require skill to use well. In the hands of an adept physician, a stethoscope can be used to diagnose heart conditions; sonar operators need to be trained to use their equipment to be able to distinguish between different underwater objects and structures; and a skilled mechanic can often troubleshoot a car engine by listening to the sounds it makes (Bijsterveld 2019, 2). So why should we insist that sonifications require little skill to use? If we are using sonification to explore complex data then there is every reason to expect that the subtle differences in the sounds produced will require a degree of training to detect. Complex tools require training and skill to use well and if we are to go beyond the very simple sonifications (that are also often not very interesting) the issue of listener training needs to be tackled. Of course, someone joining the navy as a sonar operator would have the expectation of receiving training on how to listen to sonar signals. A climate scientist interested in modelling the effects of pollution on global temperatures, on the other hand, might not reasonably have the expectation that they will need to develop analytical listening skills in order to do their job. But, if sonification users can be trained to listen more analytically than they might be used to, can we choose richer, more subtle, data-to-sound mappings that allow deeper and more valuable sonic exploration of data than has been hitherto accomplished? It will be necessary, then, to determine how “ordinary” users can be trained to listen in a skillful manner and, hence, to use sonifications more effectively. It will be interesting to discover what the practical limitations and the implications of such training for sonification design are.

In the early days, it was largely sufficient to show that sonification *could* be done, and some preliminary heuristics on how to map certain types of data to sound were produced. The underpinning theory was drawn from music philosophic accounts of listening (particularly those of Pierre Schaeffer, Michel Chion, and R. Murray Schafer) and from psychoacoustics. More recently, the role of aesthetics has become a branch of sonification research in its own right as researchers have started to tackle the rich issues associated with sonification listening as an embodied and interactional experience. It is hoped that this recent program of research, with a particular focus on sonification directness, listener skill, and the space(s) in which sonification listening takes place will yield valuable insights into how to successfully map rich and complex (and increasingly “big”) data to sound.

References

Barrass, Stephen, and Paul Vickers. 2011. "Sonification Design and Aesthetics." In *The Sonification Handbook*, edited by Thomas Hermann, Andrew D. Hunt, and John Neuhoff, 145–72. Berlin: Logos Verlag.

Bertin, Jacques. 1981. *Graphics and Graphic Information Processing*. Walter de Gruyter.

Bijsterveld, Karin. 2019. *Sonic Skills: Listening for Knowledge in Science, Medicine and Engineering (1920s–Present)*. London: Palgrave Macmillan UK. https://doi.org/10.1057/978-1-137-59829-5_3.

Bovermann, Till, René Tünnermann, and Thomas Hermann. 2010. "Auditory Augmentation." *International Journal of Ambient Computing and Intelligence* 2 (2): 27–41.

Bregman, Albert S. 1990. *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, Massachusetts: The MIT Press.

Clarke, Eric F. 2005. *Ways of Listening: An Ecological Approach to the Perception of Musical Meaning*. OUP USA.

Debashi, Mohamed, and Paul Vickers. 2018. "Sonification of Network Traffic Flow for Monitoring and Situational Awareness." *PLoS One* 13 (4). Public Library of Science: 1–31. https://doi.org/10.1371/journal.pone.0195948.

Deutsch, Diana, Rachael Lapidis, and Trevor Henthorn. 2011. "Illusory Transformation from Speech to Song." *Journal of the Acoustical Society of America* 129 (4): 2245–52.

Dombois, Florian, and Gerhard Eckel. 2011. "Audification." In *The Sonification Handbook*, edited by Thomas Hermann, Andrew D. Hunt, and John Neuhoff, 301–24. Berlin: Logos Verlag.

Flowers, John H. 2005. "Thirteen Years of Reflection on Auditory Graphing: Promises, Pitfalls, and Potential New Directions." In *Proceedings of 11th International Conference on Auditory Display (Icad2005)*, edited by Eoin Brazil, 406–9. Limerick, Ireland.

Goudarzi, Visda, Katharina Vogt, and Robert Höldrich. 2015. "Observations on an Interdisciplinary Design Process Using a Sonification Framework." In *ICAD 15: Proceedings of the 21st International Conference on Auditory Display*, edited by Katharina Vogt, Areti Andreopoulou, and Visda Goudarzi, 81–85. Graz, Austria: Institute of Electronic Music; Acoustics (IEM), University of Music; Performing Arts Graz (KUG).

Harrison, Steve, Deborah Tatar, and Phoebe Sengers. 2007. "The Three Paradigms of HCI." In *Alt. Chi. Session at the SIGCHI Conference on Human Factors in Computing Systems San Jose, California, USA*, 1–18.

Ingold, Tim. 2000. *The Perception of the Environment: Essays on Livelihood, Dwelling and Skill*. Routledge.

Kosara, Robert. 2009. "EagerEyes Blog." https://eagereyes.org/blog/2009/new-sister-site-eagerears.

Lenz, Eva, Marc Hassenzahl, and Sarah Diefenbach. 2017. "Aesthetic Interaction as Fit Between Interaction Attributes and Experiential Qualities." *New Ideas in Psychology* 47: 80–90. https://doi.org/https://doi.org/10.1016/j.newideapsych.2017.03.010.

Neuhoff, John G., ed. 2004. *Ecological Psychoacoustics*. Elsevier Academic Press.

Peters, Deniz. 2012. "Touch: Real, Apparent, and Absent — on Bodily Expression in Electronic Music." In *Bodily Expression in Electronic Music: Perspectives on Reclaiming Performativity*, edited by Deniz Peters, Gerhard Eckel, and Andreas Dorschel, 17–34. London: Routledge.

Roddy, Stephen. 2015. "Embodied Sonification." Ph.D. Thesis, Trinity College Dublin.

Rosenblum, Lawrence D. 2004. "Perceiving Articulatory Events: Lessons for an Ecological Psychoacoustics." In *Ecological Psychoacoustics*, edited by John G. Neuhoff, 219–48. London: Elsevier Academic Press.

Schaeffer, Pierre. 1967. *Traité Des Objets Musicaux*. Rev. Paris: Seuil.

Sharma, Gerriet K. 2016. “Composing with Sculptural Sound Phenomena in Computer Music.” PhD thesis, Graz, Austria: Kunst Uni Graz.

Shusterman, Richard. 1999. “Somaesthetics: A Disciplinary Proposal.” *The Journal of Aesthetics and Art Criticism* 57 (3). [Wiley, American Society for Aesthetics]: 299–313. http://www.jstor.org/stable/432196.

Sonible Gmbh. n.d. “IKO.” http://iko.sonible.com/.

Tufte, Edward R. 2001. *The Visual Display of Quantitative Information*. 2nd ed. Cheshire, CT: Graphics Press.

Vickers, Paul, and James L. Alty. 2006. “The Well-Tempered Compiler: The Aesthetics of Program Auralization.” In *Aesthetic Computing*, edited by Paul A. Fishwick, 335–54. Boston, MA: MIT Press.

Vickers, Paul, Bennett Hogg, and David Worrall. 2017. “Aesthetics of Sonification: Taking the Subject-Position.” In *Body, Sound and Space in Music and Beyond: Multimodal Explorations*, edited by Clemens Wöllner, 89–109. SEMPRE Studies in the Psychology of Music. Routledge.

Vickers, Paul, Christopher Laing, and Tom Fairfax. 2017. “Sonification of a Network’s Self-Organized Criticality for Real-Time Situational Awareness.” *Displays* 47 (April): 12–24. https://doi.org/10.1016/j.displa.2016.05.002.

Walker, Bruce N., and Joshua T. Cothran. 2003. “Sonification Sandbox: A Graphical Toolkit for Auditory Graphs.” In *ICAD ’03 9th International Conference on Auditory Display*, 161–63. Boston, MA, USA.

Walker, Bruce N., Gregory Kramer, and David M. Lane. 2000. “Psychophysical Scaling of Sonification Mappings.” In *ICAD 2000 Sixth International Conference on Auditory Display*, edited by Perry R. Cook, 99–104. Atlanta, GA: International Community for Auditory Display.

Weinberg, Gil, and Travis Thatcher. 2006. “Interactive Sonification: Aesthetics, Functionality and Performance.” *Leonardo Music Journal* 16. The MIT Press: 9–12.

Wendt, Florian, Gerriet K. Sharma, Matthias Frank, Franz Zotter, and Robert Höldrich. 2017. “Perception of Spatial Sound Phenomena Created by the Icosahedral Loudspeaker.” *Computer Music Journal* 41 (1): 76–88. https://doi.org/10.1162/COMJ_a_00396.