

A DATA SONIFICATION DESIGN SPACE MAP

Alberto de Campo

Institute of Electronic Music and Acoustics (IEM)
University for Music and Dramatic Arts, Graz, Austria
decampo[at]iem.at

ABSTRACT

We propose a systematic approach for reasoning about experimental sonification designs for a given type of dataset. Starting from general data properties, the approach recommends initial strategies, and lists possible refinements to consider in the design process. An overview of the strategies included is presented as a visual (and mental) map, and the refinement steps to consider correspond to movements on the map.

The main purpose of this approach is to make implicit knowledge (often expressed in 'natural' ad hoc decisions by sonification experts) explicit and thus available for reflection, discussion, and learning.

This approach is the result of analysing design sessions which took place in the interdisciplinary sonification workshop 'Science By Ear' in Graz, March 2006, organised by the SonEnvir team [1].

1. BACKGROUND

When collaborations on sonification for a new field of application start, sonification researchers may know little about the new domain, its common types of data, and its interesting research questions; similarly, domain scientists may know little about sonification, its general possibilities, and its possible benefits for them. In such early phases of collaboration, the task to be achieved with a single particular sonification is often not easy to define clearly, so it makes sense to employ an exploratory strategy which allows for mutual learning and exchange. Eventually, the interesting tasks to achieve become clearer in the process.

Rheinberger describes in [2] that researchers deal with 'epistemic things', which are by definition vague at first (they can be e.g. physical objects, concepts or procedures whose usefulness is only slowly becoming clear); they choose 'experimental setups' (ensembles of epistemic things and established tools, devices, procedures), which allow for endless repetitions of experiments with minimal variations. The differential results gained from this exhaustion of a chosen area in the possibility space can allow for new insights. Then, an experimental setup can collapse into an established device or practice, and become part of a next experimental setup.

From this perspective, sonification designs start their lifecycle as epistemic things, which need to be refined under usage; they may in time become part of experimental setups, and if successful, eventually 'disappear' as established scientific tools.

1.1. Some Working Definitions

The objects or 'content' to be perceptualised can be well-known information, or new unknown data (or shades of gray in between).

The aims for these two applications are very different: for information, establishing easy-to-grasp analogies is central, for data, enabling the emergence of latent phenomena in the data. As working terminology for the context here, we propose to define the following three terms:

Auditory Display is the rendering of data and/or information into sound designed for human listening. This is the most general, all-encompassing term (even though the term 'display' has a visual undertone to it).

We propose to differentiate between two subspecies of Auditory Displays:

Auditory Information Display is the rendering of well-understood information into sound designed for communication to human beings. It includes speech messages such as in airports and train stations, auditory feedback sounds on computers, alarms and warning systems, process monitoring systems, etc.

Sonification or Data Sonification is the rendering of (typically scientific) data into (typically non-speech) sound designed for human auditory perception. The informational value of the rendering is often unknown beforehand, particularly in data exploration.

This paper focuses on Data Sonification in the narrower sense.

1.2. Common Sonification Strategies

The literature often classifies sonification approaches into Audification, Parameter Mapping [3] and Model-Based Sonification [4]. For the context here, we prefer slightly different categories, which will become clear along the way; so, our three most common approaches are: Sonification by Continuous Data Representation, Discrete Point Data Representation, and Model-Based Data Representation.

Continuous Data Representation treats data as quasi-analog continuous signals, relies on two preconditions: equal distances along at least one dimension, typically time and/or space; and sufficient sampling rate, so that interpolation between data points is meaningful. Both simple audification and parameter mapping onto continuous sounds belong in this category.

Its advantages include: subjective perceptual smoothness; interpolation can make the sampling interval (which is an observation artifact) disappear; perception of continuous shapes (curves) can be appropriate; audition is very good at structures in time.

Its drawbacks include: it is often tied to linear movement along one axis only; and events present in the data (e.g. global state changes in a system) can be difficult to represent well.

Discrete Point Data Representation creates individual events for every data point, one can easily arrange the data in different orders, choose subsets based on special criteria (e.g. based on navigation input), and when special conditions arise, they can be expressed well.

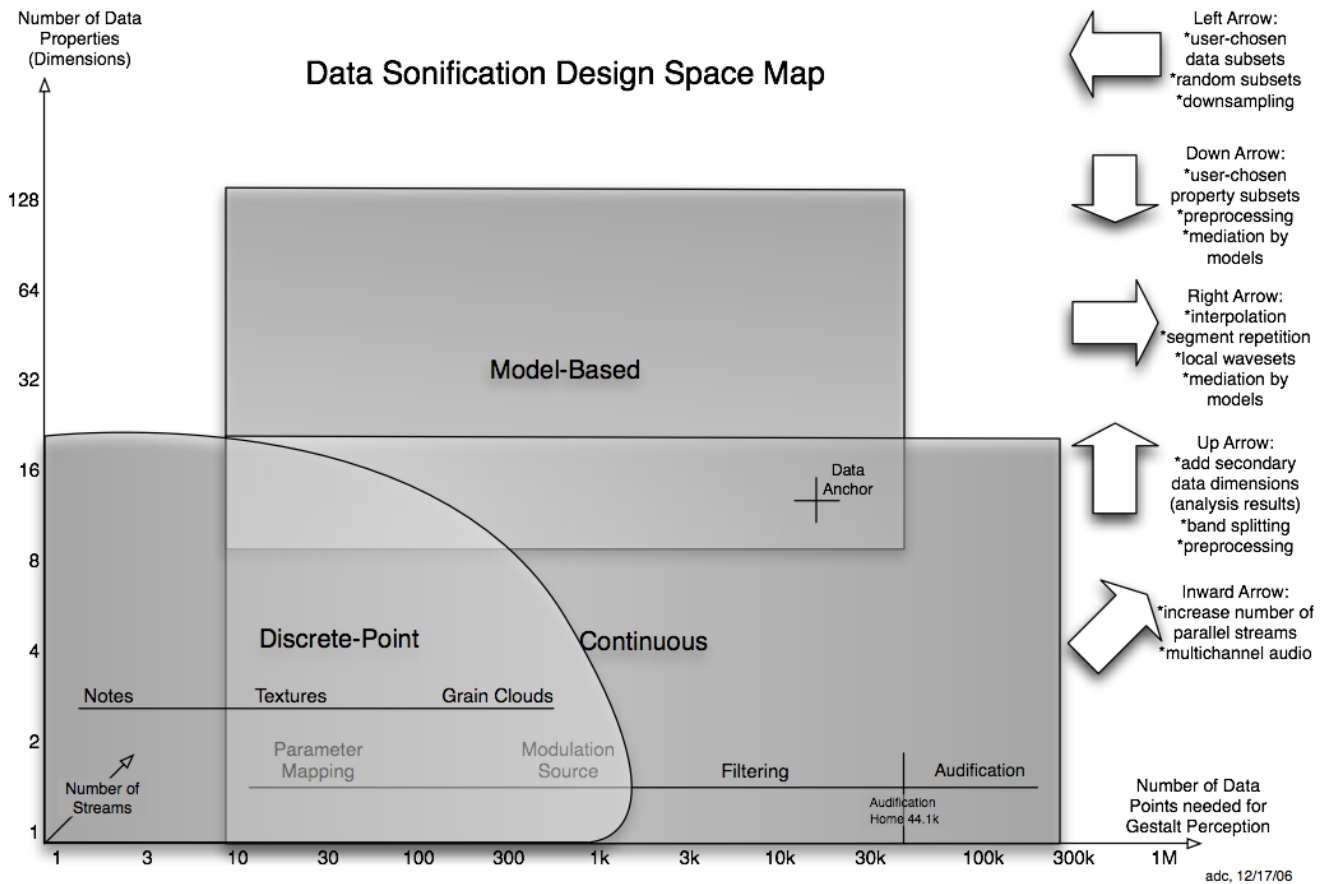


Figure 1: The Data Sonification Design Space Map.

Its advantages include: more flexibility, e.g. random iterations over data subsets; and the lack of illusion of continuity may be more accurate to the data.

Its drawbacks include: attention may be drawn to data independent display parameters, such as a fixed grain repetition rate; and at higher data rates, some of the advantages may not apply.

Model-Based Data Representation employs more complex mediation between data and sound rendering by introducing a model, whose properties are informed by the data.

Its advantages include: domain knowledge can be captured and employed in the model; and models may be applicable to multiple types of datasets.

Its drawbacks include: assumptions built into models may introduce bias leading away from domain understanding; there may be a sense of disconnection between data and sound results.

2. THE SONIFICATION DESIGN SPACE MAP

Task/Data Analysis [5] focuses on solving well-defined auditory information design problems: How to design an Auditory Display for a specific task, based on systematic descriptions of the task and the data. Here, the phenomena to be perceptualised are known beforehand, and one tries to render them in an easy to grasp way.

The Sonification Design Space Map given here addresses a

similar but different problem: The aim to be achieved here is to find transformations that let structures/patterns in the data (which are not known beforehand) emerge as perceptual entities in the sound which jump to the foreground, i.e. as identifiable 'interesting audible objects'; in the electronic music field, these are called 'sound objects' (from '*objets sonores*' [8]), in psychoacoustics literature, 'auditory gestalts' (e.g. [10]).

In other words, the most general task to achieve in sonifications is to detect auditory gestalts in the acoustic representation, which can be assumed to correspond to the patterns and structures in the data one aims to find.

2.1. The Map Axes

To facilitate this search for the unknown, the Design Space Map enables a designer or researcher to engage in systematic reasoning about applying different sonification strategies to his/her task or problem, based on data dimensionality and perceptual concepts.

Especially while the task is not yet clearly understood and defined (which is often the case in exploratory contexts), reasoning about data aspects, and making well-informed initial choices based on perceptual givens can help to develop a clearer formulation of useful tasks.

So, the proposed map of the Sonification Design Space (see

figure 1) has these axes:

X-axis : the number of data points estimated to be involved in one gestalt, or 'expected gestalt number';

Y-axis : the number of properties of interest of each data point, i.e. the number of data dimensions to be employed;

Z-axis : the number of streams estimated to be suitable for meaningful data representation.

To ensure that the auditory gestalts of interest will be easily perceptible, the most fundamental design decision is the time scale: In auditory gestalts (or sound objects) of 100 msec and less it becomes more and more difficult to discern meaningful detail, while following a single gestalt for longer than say 30 seconds takes great concentration; thus, a reasonable first order of magnitude for a good time frame for single gestalts is the duration of echoic memory, i.e. roughly 1-3 seconds [6]. The 'expected gestalt number' is the number of data points (of the dataset under study) that should be represented within this chosen time frame to allow for perception of individual gestalts at this data subset size. Note that this does not limit the maximum size: the micro-time scale is a fascinating area for creating expressive sound [7].

2.2. The Map Zones

The zones shown in the figure 1 do not have hard borders; their extensions are only meant to give an indication how close-by (and thus meaningfully applicable) which strategies are for what data 'gestalt number'. Similarly, the number ranges given below are only approximate orders of magnitude, and mainly based on personal experience in electronic music and sonification.

The Discrete-Point Zone ranges roughly from gestalt numbers 1 - 1000 and from properties numbers 1 - 20; the transition shown in the map from Note-like percepts via Textures to granular events which merge into Clouds is mainly perceptual.

The Continuous Zone ranges roughly from gestalt numbers 10 - 100.000 and from properties numbers 1 - 20; the main transition here is between Parameter Mapping and Audification, with various technical choices indicated along the way, such as using the continuous data signal as Modulation Source, Band Splitting, and applying Filtering.

The Model-Based Zone ranges roughly from gestalt numbers 10 - 50.000 and from properties numbers 8 - 128; because the approach is so general, there are no further orientation points in it yet. Existing varieties of model-based approaches are still being analysed in the terms of this Sonification Design Space, and will then be integrated in appropriate locations on the map.

3. REFINEMENT BY MOVING ON THE MAP

In the evolution of a sonification design, the intermediate versions can be conceptualised easily as locations on the map, based on how many data points are rendered into the basic time frame, how many data dimension are being used in the representation, and how many perceptual streams are in use. A step from one version to the next can then be considered analogous to a movement on the map. This mind model aims to capture the design processes we could observe in the Science by Ear workshop [1].

3.1. Data Anchor

For exploring a dataset, one can start by putting a reference point on the map, which we call Data Anchor: This is a point on the map corresponding to the full number of data points and data dimensions. A first synopsis, or better *Synakusis*, of the entire dataset (within the time frame of ca. 3 seconds) can then be created with one of the nearest sonification strategies on the map. Subsequent sonification designs and sketches will typically correspond to a movement down from this point (i.e. toward less dimensions) and to the left (toward less than the total number of data points).

3.2. Shift Arrows

Shift Arrows allow for moving one's current 'working position' on the Design Space Map, in order to employ different sonification strategies in the exploration process. Note that some shifting operations are used for 'zooming', and leave the original data untouched, while others employ data reduction, extension, and transformation; in any sonification design one develops, it is essential to differentiate between these and document the steps taken clearly. Finally, one can decide to defer such decisions and open them for interaction, so that e.g. subset are selected interactively.

A *left-shifting arrow* can be used to reduce the assumed 'gestalt number', in effect using less data points within the presentation time frame. Some options are: investigating smaller, user-chosen data point subsets (this can be by means of interaction, e.g. 'tapping' on a data region and hearing that subset); downsampling, e.g. linear, averaging, bandlimited, or by random subsets; and other forms of data preprocessing.

A *down-shifting arrow* can be used to reduce the 'properties number', i.e. to employ less data properties (or dimensions) in the presentation. Some options are: dimensionality reduction by preprocessing (e.g. statistical approaches like Principal Component Analysis (PCA) or Multidimensional Scaling (MDS), or using locality-preserving space-filling curves, e.g. Hilbert curves); and user-chosen data property subsets, keeping the option to explore others later. (Model-based sonification concepts may also involve dimensionality reduction techniques, yet they are in principle quite different from mapping-based approaches.)

An *up-shifting arrow* can be used to increase the number of properties used in the sonification design; e.g. for better understanding of mixed signals, or to increase 'contrast' by emphasizing aspects with relevance-based weighting. Some options are: band-splitting time series data into frequency bands can increase detail resolution; using (extracted and smoothed) amplitude of signal to accentuate its dynamic range; other domain-specific forms of preprocessing may be appropriate for adding secondary data dimensions to be used in the sonification design.

A *right-shifting arrow* can be used to increase the number of datapoints used, which can help to reduce representation artifacts. Some options are: interpolation of signal shape between data points; repetition of data segments (e.g. granular synthesis with slower-moving windows); local waveset audification; and model-based sonification strategies can be used to create e.g. physical vibrational models informed by few original data points.

Interpolation in time-series data is often employed habitually without further notice; the model given here allows for notating this transformation as a right-shifting arrow. If one is certain that the sampling rate used was sufficient, using cubic (or better) interpolation instead of the actually measured steps creates a smoother

signal which is nearer to the phenomenon measured than the sampled values. When such a smoothed signal is used for modulating an audible synthesis parameter, the potentially distracting presence of the time step unit should be less apparent.

3.3. Third Dimension Shifts

So far, all arrow movements have concerned movement in the front plane of the map, where only a single auditory stream is used for data representation. After the time scale, the number of streams is the second most fundamental perceptual design decision. By putting some data dimensions into parallel auditory streams (especially data dimensions of the same type, such as time-series of EEG measurements for multiple electrodes), overall display dimensionality can be increased in a straightforward way, while dimensionality in each individual stream can be lowered substantially, and thus becomes easier to perceive. (The equivalent movement is difficult to represent well visually on a 2D map, but easy to imagine in 3D space.) For multiple streams, all previous arrow movements apply as above, and two more arrows become available:

An *inward arrow* can be used to increase the number of parallel streams in the representation. Some options are: multichannel audio presentation; and setting one perceptual dimension of the parallel streams to fixed values with large enough differences to cause stream separation, thus in effect labelling the streams.

An *outward arrow* can be used to decrease the number of parallel streams in the representation. Some options are: selecting fewer streams to listen to; intentionally allowing for perceptual merging of streams.

Fusion between streams can be an appropriate expression of data features, e.g. in EEG recordings, massive synchronisation of signals across electrodes may cause the streams to fuse, which can represent the nature of some epileptic seizures well.

4. CONCLUSIONS

Conceptualising the sonification design process in terms of movements on a design space map, one can experiment freely by making informed decisions between different strategies to use for the data exploration process; this can help to arrive at a representation which produces perceptible auditory gestalts more efficiently and more clearly. Understanding the sonification process itself, its development, and how all the choices made influence the sound representation arrived at, is essential in order to attribute perceptual features of the sound to their possible causes: They may express properties of the dataset, they may be typical features of the particular sonification approach chosen, or they could be artifacts of data transformation processes used.

Note that this map is open to extensions: As new sonification strategies and techniques evolve, they can easily be classified as either new zones, areas within existing zones, or as transforms belonging to one of the arrows category; then their appropriate locations on the map can easily be estimated and assigned.

5. FUTURE WORK

There are several ways to extend the map and make it more useful:

Gaining a more detailed understanding of model-based sonification, and expressing that understanding in the terms of the conceptual framework of the map.

More and richer detail can be added, e.g. by analysing the steps taken in observed design sessions, classifying them as strategies, and adding them if new/different.

Expertise can be integrated by interviewing sonification experts, tapping into their experience, inquiring about their favorite strategies, or decisions they remember that made a big difference for a specific design process.

One can imagine building an application that lets designers navigate a design space map, on which simple example data sets with coded sonification designs are located. When one moves in an area that corresponds to the dimensionality of the data under study, the nearest example pops up, and can be adapted for experimentation with one's own data. The examples should obviously capture established sonification practice and guidelines, e.g. concerning mapping [9].

Finally, many of the strategies need not be fixed decisions made once; being able to make many of the strategic choices interactively when exploring a dataset would be both exciting and extremely valuable.

6. ACKNOWLEDGEMENTS

Big thanks go to the SonEnvir team, who put a lot of effort into making the 'Science By Ear' workshop as successful as it was, and who have generally been a joy to work with.

Big thanks also go to the workshop participants, who have been very open to this particular experimental setup in interdisciplinary collaboration; especially for the ongoing post-workshop discussions, including draft versions of this paper.

7. REFERENCES

- [1] <http://sonenvir.at/>, and <http://sonenvir.at/workshop/>
- [2] H.-J. Rheinberger, "Experimentalsysteme und Epistemische Dinge", Suhrkamp, Germany, 2006.
- [3] G. Kramer, ed., "Auditory Display", Addison-Wesley, 1994.
- [4] T. Hermann, "Sonification for Exploratory Data Analysis", PhD Thesis, Univ. Bielefeld, 2002.
- [5] S. Barrass, "Auditory Information Design", PhD Thesis, Australian National University, Sydney, 1997
- [6] B. Snyder, "Music and Memory", MIT Press, 2000.
- [7] C. Roads, "Microsound", MIT Press, 2002.
- [8] P. Schaeffer, "Traité des objets musicaux", Le Seuil, Paris, 1977.
- [9] B. Walker, "Magnitude estimation of conceptual data dimensions for use in sonification", PhD Thesis, Rice University, Houston, 2000.
- [10] S. Williams, "Perceptual Principles in Sound Grouping", in G. Kramer, ed., "Auditory Display", Addison-Wesley, 1994.