

# WAVE SPACE SONIFICATION

Thomas Hermann

Ambient Intelligence Group  
Center of Excellence in Cognitive Interaction Technology (CITEC)  
Bielefeld University, Bielefeld, Germany  
thermann@techfak.uni-bielefeld.de

## ABSTRACT

This paper introduces *Wave Space Sonification (WSS)*, a novel class of sonification techniques for time- (or space-) indexed data. WSS doesn't fall into the classes of Audification, Parameter-Mapping Sonification or Model-based Sonification and thus constitutes a novel class of sonification techniques. It realizes a different link between data and their auditory representation, by scanning a scalar field – defined as wave space – along a data-driven trajectory. This allows both the highly controlled definition of the auditory representation for any area of interest, as well as subtle yet acoustically complex sound variations as the overall pattern changes. To illustrate Wave Space Sonification (WSS), we introduce three different WSS instances, (i) the *Static Canonical WSS*, (ii) *Data-driven Localized WSS* and (iii), *Granular Wave Space Sonification (GWSS)*, and we demonstrate the different methods with sonification examples from various data domains. We discuss the technique and its relation to other sonification approaches and finally outline productive application areas.

## 1. INTRODUCTION

Sonification allows to experience data by listening, enabling listeners to understand variations in data as they manifest in systematic variations in auditory space [1]. It affords a systematic connection of the *data space* to the *sound signal space* which in turn, on listening, navigates the *auditory perceptual space* and ultimately leads to shifts in a *cognitive space of interpretations*, i.e., result in insights. The art of sonification design is to bring together knowledge of data, sound and perception to craft a suitable linkage between data- and sound signal space that allows meaningful patterns in data space to translate to distinguishable patterns in perceptual space and hence, give listeners insights about patterns in the data or awareness on their change, whatever the purpose of the auditory display might be.

Many sonification techniques have been introduced over the past decades, which illustrates how complex, rich and different a linkage can be achieved (s. icad.org and [2]). From the perspective of *structure*, these sonification techniques can roughly be organized into five different classes: (1) earcons, (2) auditory icons and parameterized auditory icons, (3) parameter-mapping sonification (PMSon), (4), audification, (5) model-based sonification (cf. corre-

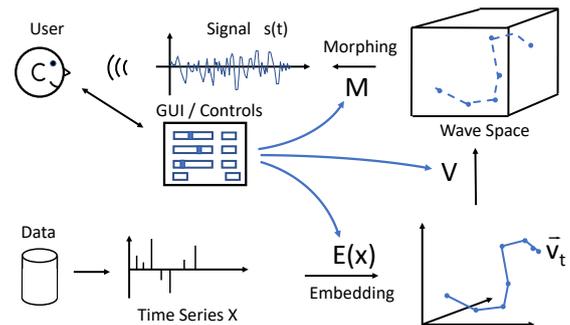


Figure 1: Wave Space Sonification: data-driven trajectories probe a static or data-refined wave space scalar field.

sponding chapters in [2]). Of course, the classes are never sharply separated and many sonifications combine different sonification techniques into hybrid sonifications.

Focusing on time-indexed data, as they occur in all walks of life and data science, from biomedical data to motion analysis, to name a few, all the techniques have offers to make, yet the dominating methods are audification and parameter-mapping sonification. Let's review how in these approaches, data space and sound space are connected.

In audification, the relation is most intimate: the data speak for themselves, as they are used as instantaneous sound pressure level that constitutes a sound signal as function of time. Although many useful processing steps such as resampling, filtering, non-linear warping and time-stretching allow to refine the transition between data and sound space, data *become* a signal.

In parameter-mapping sonification, the relation is mediated: data features are used as input for a mapping functions which computes synthesis parameters for a sound synthesis. While being highly flexible and adaptable, the nature of the relation between data and signal is that data *control* (rather than *become*) the sound signal.

Are there any more different ways of linkage? Model-based sonification offered yet another way to connect the spaces: here the data are neither becoming a sound nor used as control, but the data are turned into (or parameterize) a sound-capable system (in other words: the instrument), while interaction with them via the model is left to the user. As, however, relations in the data systematically translate to properties of the 'data-informed instrument', it offers



This work is licensed under Creative Commons Attribution Non Commercial 4.0 International License. The full terms of the License are available at <http://creativecommons.org/licenses/by-nc/4.0>

a versatile and highly interaction-builtin way to explore structures in data.

We see that all these approaches exercise a qualitatively different paradigm how data and sound signals are connected. This paper introduces yet another, novel mechanism for this connection: the data *define the navigation of high-dimensional sound signal spaces* – resulting in a new cluster of methods which are in between parameter-mapping sonification and audification, yet different in structure from both, as will be shown below. This approach is primarily intended for time- (or space) indexed data, and thus we start with a focus on multivariate time series data in the following section. In turn, WSS will be formally introduced and then illustrated by few example instances demonstrated with sonification examples. This should provide ample orientation to discuss how WSS relates to audification, parameter-mapping sonification and model-based sonification in Sec. 5. The paper concludes with an outlook on future work.

## 2. MULTIVARIATE TIME SERIES IN STATE SPACE

A necessary requirement for WSS is that the data are indexed according to a scalar variable. The most frequent case for that is that of time-indexed data, i.e., time series. As most phenomena of interest deliver several measurements at each point in time, the more general class is that of multivariate time series. Examples are EEG, ECG, EMG, inertial sensor data, motion capturing data, climate models, to name a few. However, WSS will also work if the index is of other nature, e.g. a spatial variable, such as ice core or other geological probe properties varying over a spatial index. As time series are the more typical situation we anchor the mathematical treatment on an index  $t$ .

Let  $\vec{x}(t)$  be a  $d$ -dimensional measurement of a system over time  $t$ . A discrete measurement process of  $n$  steps yields a dataset  $X = \vec{x}(t_i)_{i=1\dots n}$ . Often, but not necessarily, timestamps are equidistant, i.e.,  $t_i = i \cdot T = i/\nu$ , where  $\nu$  is the sampling rate in Hz and we can more easily write it as  $\vec{x}_i$ . A frequently used visualization of such data is by plotting component functions of time, as depicted in Fig. 2 (left). If we interpret measurement tuples as points in a vector space, we can consider the data to describe a trajectory in a  $d$ -dimensional space. Fig. 2 (right) depicts for illustration the linear projection onto a 3D-subspace spanned by 3 features.

This perspective stimulates interesting paths to deepen the understanding of the phenomenon. For instance, if we see orbits, this corresponds to oscillations in the time series; if we see convergence to a point, we understand a fix point attractor. If we see that trajectories cross each other, i.e., there are divergent trajectories from one and the same point, we understand that the representation is not a *state*, but merely a low-dimensional projection of the truly higher-dimensional situation. In other words, knowing a vector doesn't allow us to infer in what direction it will continue.

For that sake, methods are used to extend the vectors so that they are *states*. One of them is delay-embedding, which augments a time series  $\vec{x}_i$  by its past version  $\vec{x}_{i-k}$  delayed by  $k$  time units, into a now  $(2 \cdot d)$ -dimensional vector. Figure 3 illustrates delay embedding using a 1-dimensional time series for two delays. While only low-dimensional projections (up to 3) are attractive for visualization, higher-dimensional embeddings can easily be created (and may be necessary to yield a *state*), yet visual inspection is clearly limited. Here lies an important selling point for WSS, as the audibility of the trajectories in the truly high-dimensional space is

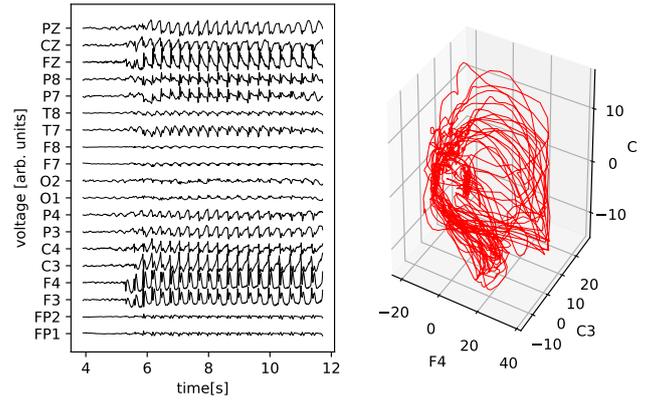


Figure 2: Typical plot of a multivariate time series (top) and a projection of the corresponding trajectory on three selected axis (right). The example shows human EEG during an epileptic seizure.

possible without problems, as we will see, resp. hear.

We see that details of the dynamics manifests in properties of the trajectories, such as velocity and its change, curvature and its change. In modern medicine and systems biology there is the notion of 'dynamical disease', which argues that we fail to understand disease by merely looking at averages (e.g. hormone levels, body temperature, etc.) but instead disease manifests in their rhythmical orchestration and thus how values change in concert with each other. From that perspective we can take additional motivation to strive at visualizations and sonifications which make such coherences explicitly perceivable as shapes that depend on the dynamics [3].

While time series analysis is a very elaborated research area, the above-introduced basics should suffice to unfold Wave Space Sonification.

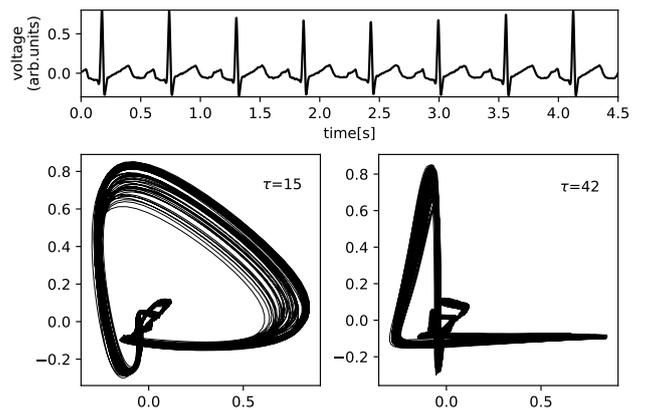


Figure 3: Delay-embedding of a one-dimensional time series (few cycles shown top) using delay  $\tau = 16$  (left) and  $\tau = 42$  (right), here showing human ECG.

### 3. THE FRAMEWORK OF WAVE SPACE SONIFICATION

Wave Space Sonification, as motivated in the introduction, offers a novel approach to connect data space with sound (signal) space, by means of a data-driven sampling of audio signals organized in a high-dimensional scalar sound field, called here *wave space*. So what is a wave space? We are familiar with the idea of audio signals stored in a vector (i.e., buffer or table) as sequence of samples, if we look at it as a one-dimensional sound space. Using a phasor at constant speed and moving along the dimension we reproduce the stored sound at a specific playback rate, associated with the phasor increment. WSS generalizes this process by defining the sonification to be the scanning of a sound space according to the trajectory that a time series forms in its vector space. Let's define it formally before we continue to specify useful fillings of wave space.

Let  $\vec{x}(t)$  be a sampling of time-indexed data  $X$ , either as continuous function or time-discrete sampling, e.g.  $t_i = i/\nu$ . Let  $\vec{v}(t) \in \mathbb{R}^d$  be a trajectory, for instance as directly used by taking the multivariate time series  $\vec{x}(t)$  and/or augmented by selected appropriate delay embeddings. We obtain  $\vec{v}(t) = E(X)(t)$  as an embedding of  $X$  into a  $d$ -dimensional space. Furthermore, let  $V : \mathbb{R}^d \rightarrow \mathbb{R}$  be a scalar function which we call *wave space function*. Finally we need a warping or morphing function  $M(t)$  which allows to control how exactly we move along the trajectory, which could be different from how the original time series moves.

With the above assertions, we define a *wave space sonification* (WSS) as the transformation of data into a sound signal via

$$s(t) = V(E(X)(M(t))) = V(\vec{v}(M(t))) \quad (1)$$

which is basically created by sampling the wave space function while moving along the embedded trajectory by means of a chosen time-advancement morphing.

WSS defines sonification as a scanning process by means of (a) a trajectory definition in wave space, and (b) a suitable definition of a wave space function, and (c) a suitable way of moving along the trajectory. The method renders interesting *auditory views* already for simple and static choices for  $E$ ,  $V$ , and  $M$ . However, a particular benefit will arise from regimes where the wave space is filled or modified according to the trajectory itself, as will be demonstrated later on.

As to (a), the trajectory definition, Section 2 has already featured the most important ideas. However, it may be useful to additionally apply dimensionality reduction methods to limit the complexity of the space, e.g., by applying a principal component analysis (PCA) [4] to select only those dimensions in which the trajectory exhibits large variations. Often this (or other data mining methods for pre-processing) can remove irrelevant detail and facilitate focusing on key properties at the same time. The trajectory definition might furthermore involve some filtering to remove random processes, such as additive measurement noise, to get smoother paths through wave space, which in turn translates to clearer signals in auditory space.

As to (b), the definition of the wave space function can be achieved at hand of different approaches. Firstly, it can be explicitly specified by using an algebraic expression such as

$$V(x_1, \dots, x_d) = \sum_{i=1}^d h_i(x_i) \quad (2)$$

for given functions  $h_i : \mathbb{R} \rightarrow \mathbb{R}$  (see examples in Sec. 4.2). Secondly, it can be defined by help of given sound signals  $s_i[n]$ , us-

ing an interpolation/resampling function to define associated real-valued functions  $h_i$ , for instance to be used in the linear combination in eq. (2) above. Thirdly, it is possible to anchor specific sound signals at user-selected locations with a defined orientation in wave space by using a mixture of signals approach, as explained next.

Assume we wish to anchor  $K$  localized wave fields  $\phi_i()$  at given centers  $\vec{\mu}_j$ ,  $j = 1, \dots, K$  with given orientation defined by a transformation matrix  $\Sigma_j = U_j D_j U_j^T$ . We can then compute  $V$  by

$$V(\vec{v}) = \sum_{j=1}^K \phi_j((\vec{v} - \vec{\mu}_j)\Sigma_j(\vec{v} - \vec{\mu}_j)^T). \quad (3)$$

If the localized wave fields contribute only around the origin, for instance by using

$$\phi_j(\vec{v}) = s_j[v_1] \cdot \exp(\vec{v}_{2\dots d}^2/\sigma_j), \quad (4)$$

the resulting WSS will exhibit the selected timbre of  $s_j[\cdot]$  whenever the trajectory passes close to location  $\vec{\mu}_j$ . Depending on the chosen morphing, the sample is activated at higher or lower sampling rate.

As to (c), the suitable way for moving along the trajectory, the first and most basic choice would be to let the trajectory itself determine the speed. The instantaneous velocity is then proportional to the distance between subsequent embedding points, and so velocity  $\propto \vec{v}[n+1] - \vec{v}[n]$ . It makes sense to interpolate the segmentwise velocity linearly to avoid unsteady changes while moving along. This way, faster trajectories naturally translate to up-sampled signals. Note that this would mean an equal number of samples in wave space for any segment  $[\vec{v}_j, \vec{v}_{j+1}]$  in embedding.

In contrast, one could also define a morphing that moves along the trajectory but at *constant* velocity, which is called the natural parameterization (also: unit-speed or arc-length parameterization) in Differential Geometry of curves. We skip mathematical details, yet given a set of embedding points  $\vec{v}_i, i = 1, \dots, N$  we could just sample each segment  $[j, j+1]$  with a number of samples in wave space that is proportional to the length  $\|\vec{v}_{j+1} - \vec{v}_j\|$ .

Any interpolation in between these two extremes, and also any extrapolation can be achieved, and they can all be characterized by means of a morphing function  $t' = M(t)$ , which could simply be said to *warp time*.

Let's look at some practical realizations of WSS next, before we discuss the method further.

### 4. WAVE SPACE SONIFICATION INSTANCES

To better understand how WSS translates multivariate time series into sound, let's explore a few examples. Same as Model-based Sonification (MBS), as overarching framework, pre-structures many different sonification models which can be called instances of MBS, WSS can have different instances, according to the choice of embedding, wave space and morphing. We start with a most basic WSS instance, that might be called *canonical wave space sonification*, proceed with a *data-driven localized wave space definition* and close with GWSS, the *granular wave space sonification*. Sound examples are provided via the following DOI: [10.4119/unibi/2919709].

#### 4.1. Canonical Wave Space Sonification

The most basic choice of constituents for WSS is to (a) use the data-driven trajectory itself for embedding, (b) define the wave space as a sum (or product) of mathematical functions, and (c) use a trajectory-controlled speed, i.e., a temporal morphing  $M(t) = t$ . A canonical choice for wave space is

$$V(x_1, \dots, x_d) = \sum_{i=1}^d \sin(2\pi\omega_i x_i). \quad (5)$$

If angular frequencies  $\omega_i = \omega_0 \cdot i$  are chosen as multiples of a fundamental  $\omega_0$ , a straight line trajectory along the  $i$ -th axis will translate into a sine signal of frequency  $\omega_i$ , i.e., a harmonic of frequency  $\omega_0$ .

Let's assume for example the multivariate time series  $\vec{x}(t) = (\cos(2\pi\omega_0 t), \sin(2\pi\omega_0 t))$ . The trajectory moves on a circle in wave space at constant speed. At  $t = 0$ , only  $x_2$  changes as  $x_1$  varies little around the maximum. Hence, the WSS produces a tone of frequency  $\omega_2$ . This frequency decreases as the sine's slope decreases while at the same time the frequency  $\omega_1$  of the  $x$ -component increases until  $t = \pi/2$ . This transformation between activation of the  $\omega_2$  and  $\omega_1$  sine continues and becomes the audible pattern as result of the sonification. Fig. 4 depicts the trajectory, the wave space and the sound signal. The sonification is provided as Sound Example S1<sup>1</sup>.

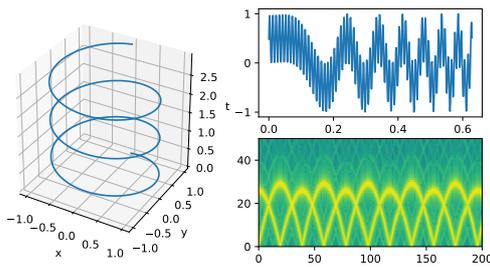


Figure 4: Canonical Wave Space Sonification - the trajectory is a circle, the 3D plot depicts it as spiral using  $t$  as  $z$ -axis. Right side plots show the beginning of the sonification (first, below the spectrogram).

Things become auditorily more interesting when more complex timbres are used to define or fill the wave space. For instance using the same frequency along  $x_1$  and  $x_2$ , but only even/odd harmonics along  $x_1/x_2$ . Then the trajectory manifests as timbre transition between two timbres. (sonification example S2).

The following sound examples S3.1-4 illustrate canonical WSS with the monthly number of sunspots dataset (measurements over 150 years which exhibits the 11-year solar cycle) [5]. The embedding of this one-dimensional time series into 2D-space is obtained by using a delay-embedding of  $k = 60$ , i.e. roughly 5 years. According to the choice of the rate parameter, which scales the velocity at which the trajectory is scanned, very different auditory views emerge, from sonifications that allow to understand

<sup>1</sup>All sound examples are provided with description on <https://doi.org/10.4119/unibi/2919709>

the overall evolution over time (e.g. how does the sunspot oscillation vary over centuries) (see sound examples S3.2 and S3.4) to auditory views that allow to understand individual orbits (e.g. fast/slow, small/large) and compare their inner shape and structure (see sound examples S3.1 and S3.3).

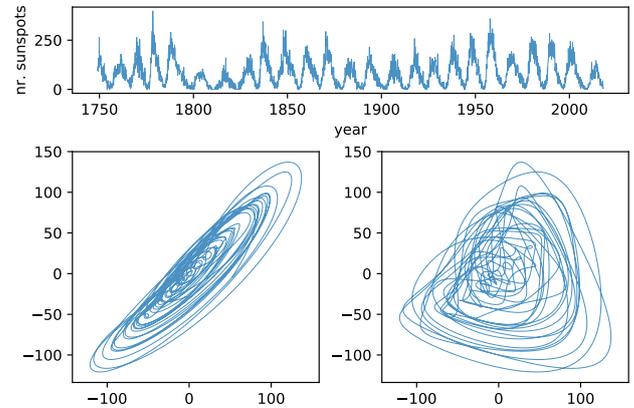


Figure 5: Wave Space Sonification of monthly sunspot data: top: time series of monthly sunspots and two delay-embedding of the filtered signal below, with  $k=6$  (left) and  $k=24$  (right). Listen to sound examples S3.1-4.

#### 4.2. Static Sample-based Wave Space Sonification

Instead of using merely mathematical functions to define wave space, it is well possible and acoustically much more interesting to define wave space by means of available samples, i.e., recorded sound signals. The most straightforward way is to choose a sound file  $s_i(t)$  for each dimension  $i = 1, \dots, d$  of wave space and define the wave space function as

$$V(\vec{x}) = \frac{1}{d} \sum_{i=1}^d s_i(c_i \cdot x_i), \quad (6)$$

where  $c_i$  are scaling factors that allow to tune each dimension differently.

As a first sonification example, listen to the sunspot orbits where the  $x_1$  dimension is filled with the spoken vowels 'a-e-i-o-u' (vowels as in bath-bear-bee-beau-boo) with a short gap in-between, and along the  $x_2$  it is left silent. The sonification examples S4.1-2 (using the sunspots data set) and S4.3-4 (for ECG, realtime and /4 slowdown) illustrate that now much more understanding of the dynamics comes from the more structured layout. This WSS is called static as the wave space filling does not depend on time or on the data. Both extensions are possible and likely paths towards more useful sonifications. Of the many possible paths, the following section examines the manual adaptation of wave space to given data.

#### 4.3. Data-Driven Localized Wave Space Definition

In many situations time series form a 1D-attractor, i.e., a limit orbit, around which deviations occur depending on initial conditions or external perturbations. For instance, the human ECG signal of a

healthy subject will have some variations between individual heart beats but largely exhibit a structurally reproducible pattern, as depicted in Fig. 3. For understanding properties of the cycle in detail, it would be beneficial to craft *data-driven wave space functions* so that samples are centered at relevant locations and aligned to the dominant direction in the delay-embedding. For instance, the QRS-complex (the peak of the heart beat) manifests for the large slightly rotated D-shaped orbit in the lower left plot of Fig. 3, shown also in Fig. 6. A data-driven wave space definition could position a sound sample  $w_1(t)$  on the horizontal part, and likewise  $w_2(t)$  and  $w_3(t)$  on the vertical part and along the main diagonal at the origin. Naturally the sonification would exhibit a sequence of  $w_1$ ,  $w_2$  and  $w_3$  sounds, depending on the idiosyncrasies of individual spikes.

Specifically, we suggest a mixture of  $k$  localized *wave space hyper-cylinders*<sup>2</sup>, centered at locations  $\vec{\mu}_k$ , oriented along unit vectors  $\hat{u}_k$  by

$$V(\vec{x}) = \sum_{k=1}^K g_k \cdot w_k(r_k \hat{u}_k^T (\vec{x} - \vec{\mu}_k)) \cdot K_\sigma(\|(\mathbf{1} - \hat{u}_k \hat{u}_k^T)(\vec{x} - \vec{\mu}_k)\|),$$

where  $g_k$  is the gain of  $w_k(t)$ , which is a sound sample, and  $K_\sigma(y)$  is a window function that decays with increasing  $y$  such as  $K_\sigma(y) = \exp(-y^2/2\sigma)$ . The scalar factors  $r_k$  allow to stretch or squeeze the sound sample along  $\hat{u}_k$ .

The data-driven localized wave space sonification is of practical utility if the analysis focus is on deviations of orbits from a reference orbit, such as in monitoring situations, diagnosis or exploratory data analysis contexts.

Sound example S5.1 is a sonification of the ECG cycle depicted in Fig 6, using a sound sample of spoken letter 'r' and 's' for the outgoing (horizontal towards the right) and returning (vertical, downward) trajectory of the QRS-complex. A whistling sound is added to the origin with a 45° degree orientation, which is turned into a discernable pattern by the other parts of the ECG (u-wave, p-wave, etc.). The used samples were quickly recorded with the laptop microphone and are provided as sound example S5.0. Sound example S5.2 represents a sonification at 1/4 the rate, allowing to perceive the signal between the QRS-complexes in greater detail.

#### 4.4. Granular Wave Space Sonification

The instances of WSS introduced above use the state trajectory directly to sample the Wave Space, and thus have a highly direct coupling between trajectory dynamics and sound signal. However, this comes at some disadvantages. On the one side, with the data-driven velocity mode, the changing trajectory speed causes strong (often too extreme) resampling so that the timbre cannot anymore be properly recognized, yet at least the timing is correct, i.e., sonification time is proportional to data time. On the other side, using the 'unit speed' natural parameterization or arc-length velocity mode, the extreme sampling problem is solved but then the temporal pattern is strongly distorted, causing short large spikes (such as the QRS-complex in the ECG data example) to be represented by much more sonification time than it proportionally takes within the cardiac cycle.

To overcome these limitations we finally introduce *Granular Wave Space Sonification*, which decouples the trajectory from the sound-collecting process to some degree. While moving along

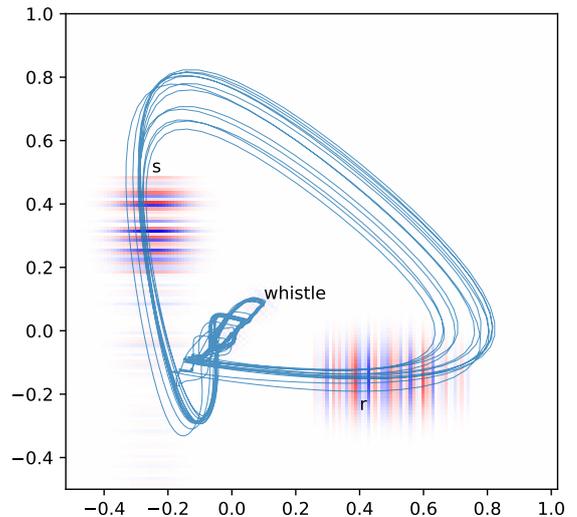


Figure 6: Localized Wave Space Sonification of ECG: the spoken sound 'r', 's' and a whistling sound at constant pitch are inscribed into wave space at the visible location and orientation. The whistling sound is high pitched and more quiet and thus less visible here unless zooming in.

the trajectory we regularly spawn 'particles' to probe wave space along the trajectory. These particles have a limited range and can propagate at their own velocity, independently of the moving trajectory. According to the morphing function (or velocity model) this spawning could happen equally often per second, or whenever a determined distance has been moved along the trajectory. The *Granular Wave Space Sonification* (GWSS) is then simply the superposition of all particles' individual wave space sonifications. However, as usual in granular synthesis, we also apply to each grain a temporal amplitude envelope window (hanning window) to reduce transient effects.

Compared to standard WSS we gain three control parameters: (i) the particle rate (in particles per second), the duration (in seconds) of the particles' life time, allowing the sound to become a mixture of slightly different particle views on the time series, and finally (iii) the propagation speed of particles in unit length per second, which can be adjusted between a narrow local spatially confined focus at slow speed and long-ranging particles, even up to half a cycle or more at large speed.

The following sound examples illustrate GWSS for the ECG data set used before. A sonification of the Lorenz attractor and EEG data will be made available as supplementary material and shown at the conference. Sound example S6.1 is a GWSS rendered with a particle rate of only 3 Hz, particles of 20 msec, using the static WSS wave space function. You can hear the well isolated grains and distinguish their timbre. Compare this to sonification S6.2., where the particle duration has been increased to 100 ms. Here the temporal contour in each grain becomes perceivable. In sonification S6.3 the short grains are kept but the particle rate is

<sup>2</sup>'hyper' because the definition is not only for 3D but any dimension

increased to 20 Hz. Try to attend to the rhythm of the heart beat (0.5 Hz, as a 2x slow-down is used). This is much easier in sonification S6.4 where the particle duration is increased to 3 ms and the particle rate has been increased to 40 Hz so that grains blend and you can follow the 5 heart beats. If more details are needed without affecting the spectrum, the rate can now easily be changed, as for instance to a 6x slow-down in sonification S6.5.

Of course, the GWSS can also be used with any of the other proposed wave space fillings, such as the localized wave space samples. Examples S6.6 and S6.7 illustrate sonifications for a used spoken 'a-e-i-o-u' sound file laid out along the  $x_0$ -axis.

## 5. DISCUSSION

The above presented WSS instances demonstrate that Wave Space Sonification is a structured yet still quite versatile and flexible framework to investigate time- (or space-) indexed data. Depending on the selection of the rate, different phenomena in given data can be explored – from differences in-between individual orbits to the overall structure of the dataset as a whole, including trends and changes in rhythmical organization. The possibility to define the wave space function at hand of user-selected timbres enables to rather easily create sonifications that meet expectations in timbral complexity or pleasantness. For instance, very precise but slightly annoying wave files which might be perfect for analysis and diagnosis could then easily be replaced by less precise but much more pleasant sounds, whatever is regarded as such, for using the system in a real-time monitoring scenario.

The number of parameters to control WSS is nicely limited, as too many parameters would only make it overly difficult to adjust the system to one's needs. They include the embedding, the rate, spatial scale, selection of morphing function, selection resp. definition of wave space. Depending on the latter, a number of additional parameters might be needed. For the GWSS, we furthermore have introduced particle duration, particle rate and particle speed, which all are self-explanatory when you understand the sonification process.

The chosen datasets for sonification examples have in common that their trajectories are repetitive with variations, i.e., they are quasi-periodic. This is a particularly useful situation as then the ear can easily understand the cycle, learn and memorize the basic pattern and subsequently focus on differences therein. However, non-stationary patterns would also make sense for WSS, for instance in the case of monitoring: assume that due to a trend or change of parameters the trajectory shifts to another area of wave space which is occupied by a completely different timbre – the listeners will very likely notice the salient change corresponding to such a non-periodic shift.

The second commonality in the datasets is that the trajectories are sufficiently smooth. This makes both the visualization easier to follow and the sonifications smoother. Smooth trajectories can usually be obtained for any data by applying a low-pass filter. However, it often makes also sense to apply a high-pass filter to remove DC-offsets or any slow waves that would manifest in delay embeddings as spatial translations of the orbits, so that practically a band-pass filter with user-selected low- and high-frequency cut-off is the best.

How WSS performs without this filtering, and in how far the filtering may remove relevant information from the underlying data depends on the data and remains to be learnt in long-term experience with – and further investigations of – WSS.

In the introduction, we advertised WSS to be a new class of sonification technique, on the same level as audification, model-based sonification and parameter mapping sonification, yet not well fitting into any of these. Thus, the relation of WSS to these techniques should be clarified more thoroughly.

### 5.1. WSS is not Audification

Audification is indeed a special case of Wave Space Sonification: If we take the wave space definition to be  $V(\vec{x}) = V_{\text{aud}}(\vec{x}) = x_1$ , then the WSS is exactly the audification of the feature represented as the first feature, i.e.,  $x_1$ . If the state space is defined by a delay-embedding, then  $x_1$  is just directly the series of measurements.

This makes it clear that in turn WSS can not be a special case of audification. Audification, in its original form, doesn't generally include the application of arbitrary transfer functions apart from probably dynamic compression which is sometimes considered as signal conditioning [6]. In WSS, an audification with dynamic compression is likewise obtained by using  $V(\vec{x}) = h(x_1)$  where  $h(x)$  is a sigmoidal transfer function. More advanced approaches for signal conditioning in audification is the application of a phase vocoder or buffer granulator (similar to TGrains UGen in Supercollider3) to decouple time progression to some degree. This would also be a special case of WSS using  $V_{\text{aud}}(\vec{x})$  with the GWSS introduced above.

Looking at audification as special case of WSS illustrates that WSS is a more general approach, and only in very special situations a WSS might appear to be an audification. The most important differences are that (i) typically not the data series alone but a data-orchestrated trajectory in data- (or state-) space is used for probing the wave space, that (ii) the wave space is a multi-dimensional signal and thus much more complex than a linear or smooth transfer function, and consequently (iii), that thus already slowly oscillating signals such as 1 Hz rhythms can deliver audible results – whereas audification would require significant temporal compression to shift oscillations to the audible frequency range which would in turn reduce the temporal resolution.

Audification and WSS share the problem that a speedup (i.e., resampling in audification resp. scaling of the embedding in WSS) can easily result in aliasing and cause misleading spectral reflections at the Nyquist frequency  $\nu$ . As the wave space might usually be already filled with many frequencies (up to  $\nu$ ), ways to realize band-limited sonification would be a very useful and relevant focus of ongoing research.

Extensions of Audifications, such as *Augmented Audification* by Vogt and Höldrich [7] perform frequency shifting using a single-side-band modulation. This is an extension that shares with WSS the ability for flexible adjustment of the sonifications frequency range and at the same time and independently, to interactively zoom into temporal structures at any scale, yet it is achieved with a different method (multiplying vs. concatenating functions), so that the two techniques do not relate easily to each other.

### 5.2. WSS is not Parameter Mapping Sonification

In Parameter Mapping Sonification (PMSon), data are mapped to synthesis parameters of a sound synthesis algorithm, thus allowing to route different data features (i.e., components of a data vector) to different auditory qualities. If we would use a synth that taps into a wavetable and map our given time series data to the phasor (or tapping location) we would look at WSS from the perspective

of mapping. So, isn't WSS just a PMSon? The described PMSon is possible but w.r.t. the key idea of PMSon, it is a bit pathological: usually in PMSon we would define our parameters to systematically vary an auditory quality, e.g., pitch, level, duration, attack time, etc. This is not the case with mapping on a phasor.

Also, with the same argument we could explain an audification to be but a PMSon – which is obviously not helpful, as in terms of the central ideas these methods are clearly distinct.

Looking at WSS as a mapping doesn't do justice to its rather signal-near (to the level of sample-based control) nature of how data relates to the sonification – in parameter mapping this is usually much more indirect and usually driven by the wish to map a variable's change to likewise perceivable change of a perceptual quality.

Furthermore, WSS abstracts from the underlying data towards the definition of an embedding which is then subjected to a morphing before accessing the wave space, so we can't really say that the data index directly a wave table. Finally the WSS is non-parametric with regards to data and sound parameters: while in PMSon problems arise when there are more features than synthesis parameters – such as which of them to ignore – WSS is immediately applicable for data of any dimension.

### 5.3. WSS is not Model-based Sonification

WSS uses the idea of movement in space, which may remind to the idea of model space in Model-based Sonification. So can we regard WSS as a sonification model? The key idea in Model-based sonification is that the data *become* a dynamic system, capable of showing acoustic behavior. This is typically achieved by using the data to determine the setup of elements in the model, e.g. mass-spring systems, etc. An MBS is usually in a state of equilibrium and thus silent until excited by interaction, and this interaction is usually initiated by the data-exploring user. So, this is structurally quite different from WSS, where the data can more be said *to operate on or to drive* wave space than *to become* wave space. Thus WSS is clearly not a MBS.

---

In summary, we conclude that Wave Space Sonification is best regarded as a new class of sonification methods, conceptually closest to audification, and equally distant to parameter mapping sonification and model-based sonification, sharing some characteristics from the latter.

### 5.4. Related Sound Synthesis Sonification Methods

Key ingredients of WSS are the definition of wave space and the use of a trajectory that slowly taps into it and a scanning of wave space. We don't know of any other related work that uses this specific approach for sonification of multivariate time series. So we start with the more distant area of computer music.

With *Scanned synthesis*, Verplank et. al. [8] introduced a similar idea in computer music, however, more related to allowing users to interactively control timbre for musical expression. Their aim was to improve the human interface in controlling complex timbre using dynamical systems whose 'state' becomes the wave form that is scanned.

The Wave Space, if filled with sound files, is related to *Wave Terrain Synthesis* by Roads [9], which exercises a similar generalization from 1-dimensional wavetables to 2 or 3-dimensional wave fields, but again only for the sake of computer music applications.

Note that the Wave Space Sonification here is generally introduced as a scalar field of  $d$  dimensions, where  $d$  is the dimensionality of the multivariate data. The definition was tailored to match properties of time series in state space, which can be high-dimensional and visualization falls short above 3 dimensions. Here, Wave Space Sonification might help to give extra information, particularly when combined with interactive and dynamic visualization.

The idea of following a trajectory in data space for defining the sonification is for instance used in Principle Curve Sonification by Hermann et al. [10]. But being a MBS, here some other acoustic processes are needed to explain the audio signal.

Williamson and Murray-Smith introduced Sonification of probabilistic feedback using granular synthesis [11]. Here the idea is to facilitate the operation of a control space by using granular synthesis to experience the probabilities of likely goals, using display quickening. While the combination of using state space, particles and granular synthesis makes this approach to appear similar, it's main idea is interactive control space feedback.

Among the many papers on parameter mapping sonifications, one finds granular synthesis but not with much overlap to this definition of WSS here.

## 6. CONCLUSION

In summary, this paper has introduced Wave Space Sonification (WSS) as a new class of sonification techniques, starting from a formal definition, then presented concrete instances of WSS as new sonification techniques that apply the WSS logic to represent a number of test data sets from various areas. The sonification examples showed that very different aspects, ranging from detailed information of individual cycles to overarching structures such as trends manifest in sound. While the core concepts are outlined here, much research remains to be done to better handle issues such as aliasing in canonic WSS or providing guidelines how to best define Wave Space in case of higher-dimensional systems. This ongoing work will hopefully also provide some body of knowledge and experience how to shape wave spaces so that the resulting sounds are both informative and aesthetically pleasant.

As for applications, we see the highest potential in two areas: (i) in biomedical data contexts (analysis, differential diagnostics, monitoring), for instance in fields such as ECG and EEG, and (ii), in interaction scenarios such as in sports and rehabilitation, since in these domains we often have repetitive movement cycles that would naturally translate into orbits in wave space. Since WSS is realtime-capable, the sound can be perceived as immediate feedback during the activity, allowing users to interactively explore the interrelations between their movement execution and its consecutive Wave Space Sonification.

## 7. ACKNOWLEDGMENT

This research/work was supported by the Cluster of Excellence Cognitive Interaction Technology 'CITEC' (EXC 277) at Bielefeld University, which is funded by the German Research Foundation (DFG). Thanks to A. L. Aldana Blanco for selecting and preparing the ECG data sample from PhysioNet. Thanks to the developers of python/numpy/sciopy/matplotlib/jupyter for their amazing tools.

## 8. REFERENCES

- [1] T. Hermann, "Taxonomy and Definitions for Sonification and Auditory Display," in *Proceedings of the 14th International Conference on Auditory Display (ICAD 2008)*, P. Susini and O. Warusfel, Eds. IRCAM, 2008.
- [2] T. Hermann, A. Hunt, and J. G. Neuhoff, Eds., *The Sonification Handbook*. Berlin, Germany: Logos Publishing House, 2011. [Online]. Available: <http://sonification.de/handbook>
- [3] G. Baier and T. Hermann, *Temporal Perspective from Auditory Perception*. World Scientific, 2008, pp. 348–363.
- [4] I. T. Jolliffe, *Principal Component Analysis*. Springer, New York, 1986.
- [5] B. WDC-SILSO, Royal Observatory of Belgium, "Sunspot data," <http://www.sidc.be/silso/datafiles>.
- [6] F. Dombois and G. Eckel, "Audification," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Berlin, Germany: Logos Publishing House, 2011, ch. 12, pp. 301–324. [Online]. Available: <http://sonification.de/handbook/chapters/chapter12/>
- [7] K. Groß-Vogt and R. Höldrich, "Augmented audification," in *Proceedings ICAD 2015 (ICAD 2015)*, Graz (Österreich), 07 2015.
- [8] B. Verplank, M. Mathews, and R. Shaw, "Scanned synthesis," *The Journal of the Acoustical Society of America*, vol. 109, no. 5, pp. 2400–2400, 2001. [Online]. Available: <https://doi.org/10.1121/1.4744477>
- [9] C. Roads, *Computer Music Tutorial*. Cambridge, MA, USA: MIT Press, 1996, ch. Wave Terrain Synthesis, pp. 163–167.
- [10] T. Hermann, P. Meinicke, and H. Ritter, "Principal Curve Sonification," in *Proceedings of the International Conference on Auditory Display*, P. R. Cook, Ed. International Community for Auditory Display (ICAD), 2000, pp. 81–86.
- [11] J. Williamson and R. Murray-Smith, "Sonification of probabilistic feedback through granular synthesis." *IEEE Multimedia*, vol. 12, pp. 45–52, 2005. [Online]. Available: <https://doi.org/10.1109/MMUL.2005.37>