

# SOMson — SONIFICATION OF MULTIDIMENSIONAL DATA IN KOHONEN MAPS

*Simon Linke*

Ligeti Center  
Hamburg University of Applied Sciences  
Hamburg, Germany  
simon.linke@haw-hamburg.de

*Tim Ziemer*

Institute of Systematic Musicology  
University of Hamburg  
Hamburg, Germany  
tim.ziemer@uni-hamburg.de

## ABSTRACT

Kohonen Maps, aka. Self-organizing maps (SOMs) are neural networks that visualize a high-dimensional feature space on a low-dimensional map. While SOMs are an excellent tool for data examination and exploration, they inherently cause a loss of detail. Visualizations of the underlying data do not integrate well and, therefore, fail to provide an overall picture. Consequently, we suggest SOMson, an interactive sonification of the underlying data, as a data augmentation technique. The sonification increases the amount of information provided simultaneously by the SOM. Instead of a user study, we present an interactive online example, so readers can explore SOMson themselves. Its strengths, weaknesses, and prospects are discussed.

## 1. INTRODUCTION

Self-organizing maps [1], also known as Kohonen maps, are artificial neural networks that represent a high-dimensional feature space on a low-dimensional map. This unsupervised learning technique can serve for data browsing, exploration and knowledge acquisition, pattern recognition, clustering, and data classification. SOMs are utilized in the fields of medicine [2, 3], biology [4, 5], geology [6], musicology [7, 8], sustainability [9], ethnology [10, 11], material science [12, 13] and many more.

In contrast to other artificial neural networks, SOMs are not a black box. Instead, they are explicitly designed to inspect data sets and let users explore all network coefficients. Through analyzing the SOM, users gain an understanding of the training data. To date, this exploration and analysis is mainly based on visualization of single component planes or the somewhat condensed U-matrix (for a detailed explanation, see Section 2). While very useful and often intuitive, these visualizations cannot present the whole picture. As SOMs map a high-dimensional feature space to a two-dimensional grid, they do not simultaneously communicate all feature magnitudes of the underlying items. As a solution, we present SOMson, the sonification of a SOM based on a four-dimensional feature space. SOMson enhances SOMs by sonifying each node of the unit layer, allowing users to explore more aspects of the underlying data in an integrated fashion. Instead of evaluating the benefit of SOMson through a user study, we decided to provide a clear explanation of SOMs and SOMson and let the readers experience SOMson themselves in an interactive demonstrator.

In the remainder of the paper, we explain self-organizing maps, how they work, and how they are visualized. Then, we explain the psychoacoustic sonification, what data it sonifies, and how. After that, we present the SOMson user interface. Then, we give a guided tour through SOMson in an interactive online demonstrator, complemented by video examples. Finally, we discuss SOMson and provide a short conclusion.

## 2. SELF ORGANIZING MAP

Self Organizing Maps (SOMs) are explained in detail by their inventor in [1]. We briefly summarize it here to give a better understanding of SOMson. SOMs are artificial neural networks with just one input and one output layer, the *unit-layer*.

In our demonstrator, we analyzed 15 songs, generally referred to as items. From each song, we extracted the 4 features *PhaseSpace*, *ChannelCorrelation*, *PhaseSpaceHigh* and *bpm* [14, 8]. This way, each item is represented by a 4-dimensional feature vector. Each song would have its unique location in a four-dimensional space, and their constellation could be observed in terms of proximity, like Euclidean distance. However, visualizing a four-dimensional space is not straightforward. Instead, a SOM is trained to represent the spatial constellation on a two-dimensional grid.

The 15 items with their 4-dimensional feature vectors are the input layer of the SOM. Each item from the input layer is connected to each node in the output layer, the so-called *unit layer*. In our case, the unit layer is a quadratic grid with  $16 \times 16$  nodes. Every node holds a 4-dimensional vector, a pointer into the four-dimensional feature space. We call the combination of a node and its pointer a *unit*. Initially, the unit layer is randomized, i.e., each pointer points at a random location in the 4-dimensional space. This unit layer will be trained iteratively, whereby the (now two-dimensional) representation of all 15 items should preserve the original (high-dimensional) topography as closely as possible. This makes the SOM comparable to a projection of a high-dimensional space on a low-dimensional space, or to a multidimensional scaling approach.

To train the unit layer, we identify the node whose pointer is most proximate to the location of item 1. We can call this node the *winning node* and refer to the combination of this node and its pointer as the *Best Matching Unit* (BMU). Now, the item “drags” this node’s pointer toward the location of the item. This means we modify the pointer to become a weighted mean value of the item’s location and the original pointer. The weighting is called the *learning coefficient*, and the pointer modification is called *learning*. The pointers of all neighboring nodes are modified, too. The larger the



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/>

distance between a node and the winning node, the lower the learning coefficient. The decrease in learning over distance is called a *neighborhood function*, often modeled as a Gaussian function centered around the BMU. This implies that the first item already alters the landscape of the complete unit layer from random pointers to a somewhat noisy Gaussian curve around item 1.

The procedure is repeated with all items. This means that all 15 items will affect the pointers of all nodes. You can imagine that items proximate to each other strengthen the pointing towards them, at least from the nodes nearby. In contrast, items from very different locations “steal” pointers from their nearby nodes. Pointers of nodes in between are torn back and forth, therefore pointing towards their middle. In machine learning terminology, this is referred to as *competitive learning*.

In the first round, the learning coefficient is large, and the neighborhood function has a small decrease over distance. When all items have been used to train the SOM, the process is repeated. Every round, the learning coefficient is reduced, and the neighborhood functions becomes narrower, making the learning related to an iterative process. This way, the unit layer (re-)organizes itself. Hence the term *self organizing* map. Every round, the order of the items is chosen randomly to ensure the same influence of every item. Of course, the calculation of the BMU is renewed every round, too. Due to the reduction of learning coefficients and neighborhood functions, the unit layer converges to a nearly stable state. Thus, the training ends after multiple iterations (2000 in the given example) since every round would only slightly affect the BMUs, while the overall topology stays the same.

When the training is over, each node’s pointer points at another location. Each item with which the unit layer has been trained has one BMU. Items that are similar in terms of many features will have proximate BMUs (or even the same BMU). They will cluster. Moreover, all pointers of the nearby nodes will point towards them. Items very different from this cluster but similar to each other (in terms of some feature magnitudes) may cluster somewhere else on the map. Again, all pointers of nodes around them will point towards this cluster. The nodes in between will point to the middle of these clusters and not to one of them. These are separation lines between two clusters. Depending on the data, items may not cluster but distribute over a subregion of the map. In this case, pointers may gradually point from item to item, like an interpolation. Such gradients do not exhibit clear separation lines.

Note that the unit layer is high-dimensional. Each node holds a pointer that is as high-dimensional as the input feature vector. However, the whole point of a SOM is the reduction of a high-dimensional feature space to a low-dimensional map that maintains the original topology. To date, this is achieved through some low-dimensional visualizations. Most importantly, the trained unit layer is visualized through the so-called *U-matrix*. It is presented in Fig. 1. Instead of trying to visualize the 4-dimensional pointer of each node, the U-matrix only shows the mean distance between a node’s pointer and all neighboring nodes’ pointers. When a unit and its neighbors point at the same 4-dimensional location, it is plotted in black. The larger the mean distance between the node’s pointer and the neighboring nodes’ pointers, the lighter it is. This way, clusters appear as black islands, separated by white seas. The darker the island, the more similar the units. The lighter the sea, the larger the difference between the islands. On this map, each item from the training set is visualized by a colored dot on its respective BMU. Once trained, new items can be added to the SOM to analyze their relationship to the other items.




Figure 1: The *U-matrix* is the main output of a *Self Organizing Map*. Instead of visualizing the 4-dimensional pointer at each of the  $16 \times 16$  nodes of the unit layer, the *U-matrix* indicates the mean distance between each node’s pointer and the pointers of all its neighboring nodes. Nodes in the corners have 3 neighbors, the other nodes along the fringe have 5, and nodes in the middle have 8 neighbors. The single items used to train the SOM (techno music) are shown as colored dots, whereby different colors represent different techno music styles.

In our example, we analyzed four different styles of techno music and plotted them in red, cyan, green, and blue. However, the training of the SOM is unsupervised, i.e., the algorithm is neither informed about, nor affected by our manual categorization. As you can see, the green dots cluster well. So do the blue dots. A gray separation line separates both clusters. The cyan dots cluster well, too. They are separated from the blue dots by a slightly lighter separation line, indicating that these clusters are a bit more different than the green and the blue one. The separation between the green and the cyan clusters is even lighter. The only white separation line can be found between the red item in the lower-left corner and the green cluster, indicating that this red item is more distinct from the green cluster than, for example, from the other red dot on the left-hand side of the map.

The *U-matrix* provides us with a lot of information about the relationships between the items. But there are four things we cannot see:

1. Where in the feature space are our items allocated?
2. In how far are items on the same island different from one another?
3. In terms of which features are the islands different from one another?
4. How similar or dissimilar are those islands that are no neighbors?

The so-called *component planes* are an approach to answering these questions. Component planes map the magnitude of a single feature at each node to color. Here, dark blue represents the lowest value, and yellow represents the highest value. The four component planes of our SOM are illustrated in Fig. 2. The component planes reveal, for example, that the cyan items exhibit the highest

*Phase Space* value, the highest *Channel Correlation*, the highest *BPM* value, and a medium value of the *Phase Space High* feature. So, overall, this cluster is somewhat extreme. In other words, the component planes allow us to answer question 1. Moreover, we can see that cyan items mostly differ from each other in terms of *Phase Space High*, which answers question 2. When we compare the island of the blue items with the island of the cyan items, we see that both share a medium to high *Phase Space High* and mostly a high *Phase Space*. What distinguishes them the most is the *Channel Correlation*, followed by the *BPM*. This means we can answer question 3. Last but not least, we can compare the islands that are no neighbors. For example, the uppermost red item and the island of blue items share a similar *Channel Correlation*, while the *Phase Space* value of the red item is much lower, and the *BPM* value is just slightly lower. In contrast, the *Phase Space High* value of the red item is a bit higher than that of the blue items. This answers question 4.




Figure 2: The component planes plot the magnitude of each one feature at each unit from dark blue (minimum) to yellow (maximum).

Together, the U-matrix and the component planes present a lot of the information inherent in the SOM. The benefit is that once learned, they can be interpreted quite easily. Therefore, SOMs allow users to analyze data sets, identify clusters and distributions, e.g., to explore or presort new datasets, or browse through data based on similarity. Moreover, they allow researchers to study the explanatory value of single features and feature combinations, the intuitiveness of different distance measures, and much more.

The downside of these visualizations is that they do not provide the overall picture. In the U-matrix, most islands look the same, leaving the 4 open questions listed above. Even though the component planes answer these questions, we can only analyze them one by one. Users either plot them on one screen and can only focus on one at a time, or they skip between graphics. All

these visualizations do not integrate well. As a solution to this problem, we suggest SOMson, an interactive, multidimensional sonification of each unit’s feature magnitudes. SOMson can be considered an augmentation of the SOM visualizations to increase the informativeness of Kohonen maps.

### 3. SOMson SONIFICATION

The simultaneous auditory display of multiple data dimensions or variables is often referred to as *multidimensional* or *multivariate sonification* [15, 16]. Multidimensional sonifications have already been proposed in 1980 in [17], and later, e.g., in [18, 19]. In the course of the Sonic Tilt Competition 2023 [20], many new two-dimensional sonifications with 2 polarities each have been developed, like [21, 22, 23, 24, 25]. Important requirements for multidimensional sonification include

1. interpretability of dimensions
2. continuity of dimensions
3. linearity of dimensions
4. a high resolution of dimensions and
5. orthogonality between dimensions. [26]

The sound in our demonstrator is based on the psychoacoustic sonification as introduced for two dimensions in [27] and extended to three dimensions in [28]. Experiments with passive [29] and interactive users [30] have revealed that the psychoacoustic signal processing fulfills the above-mentioned requirements through a mapping of single dimensions to *chroma*, *roughness*, *sharpness* and *loudness fluctuation* of a Shepard tone. The orthogonality between these dimensions is also demonstrated in a YouTube playlist: <https://youtu.be/7EeB7AGJnpQ&list=PLVv3BMS8IIXGo-SkwwD9rSUQKCPly89kK>. Note that we treat roughness and loudness fluctuation as independent dimensions, while the cited literature uses them to represent two polarities of the same dimension.

Implementation of the psychoacoustic sonification is straightforward: The source code can be found in [31] and is based on [32], which has been implemented in the CURAT sonification game [33] and the Tiltification spirit level app [34] and other applications [35, 36]. Experiments with the psychoacoustic sonification have shown that training is not necessary [37] but helpful for the interpretation [27] and interaction [38].

What is sonified by the multidimensional sonification is the pointer at each node of the SOM. Here, each dimension of the pointer, i.e., each feature magnitude, is mapped to the magnitude of one dimension of the psychoacoustic sonification:

1. *PhaseSpace*  $\rightsquigarrow$  Carrier Frequencies  $\longrightarrow$  *Chroma*
2. *ChannelCorrelation*  $\rightsquigarrow$  Frequency Modulation Index  $\longrightarrow$  *Roughness*
3. *PhaseSpaceHigh*  $\rightsquigarrow$  Peak Position of Amplitude Envelope  $\longrightarrow$  *Sharpness*
4. *bpm*  $\longrightarrow$  Amplitude Modulation Frequency  $\longrightarrow$  Speed of *Loudness Fluctuation*

where “ $\longrightarrow$ ” indicates a linear mapping and “ $\rightsquigarrow$ ” indicates a non-linear mapping.

Each feature magnitude is normalized, resulting in a variable  $x \in [0, 1]$ , which is used to modulate the audible output signal  $y(\omega, t)$ , which is produced by nine sine-wave oscillators:

$$y(\omega, t) = \sum_{i=0}^8 \hat{A}_i \sin(2\pi \omega_i t) \quad (1)$$

Each frequency  $\omega_i$  is modulated by the *PhaseSpace*  $x_{ps}$  to produce a *Chroma* as follows:

$$\omega_i = 25 \cdot 2^{\left(i + \frac{4x_{ps}}{12}\right)} \quad (2)$$

The Carrier Frequencies of the *Chroma* are further modulated to add *Roughness* representing the *ChannelCorrelation*  $x_{cc}$ :

$$y(\omega, t) = \sum_{i=0}^8 \hat{A}_i \sin(2\pi \omega_i t + I(x_{cc}) \times \sin(2\pi \times 30t)) \quad (3)$$

Here, the modulation index  $I$  is given by a mixture of a logarithmic and a linear mapping:

$$I(x_{cc}) = 0.4 \times 5^{2.8 x_{cc}} + 0.6 x_{cc} 5^{2.8} \quad (4)$$

*PhaseSpaceHigh*  $x_{ph}$  is represented by modulating the Amplitude  $\hat{A}_i$  of each oscillator according to its Frequency  $\omega_i$  and thus, changing the *Sharpness*:

$$\hat{A}_i = \exp\left(-0.5 \left(6.66 \left(\log_2(\omega_i)/9 - (0.5 + 0.24 x_{ph})\right)\right)^2\right) \quad (5)$$

This Amplitude is further modulated by adding a signal  $A_2$  with the same Amplitude and a frequency responding to the *bpm*-feature  $x_b$ :

$$A_2 = \hat{A}_i \times \sin(2\pi \times 8x_b) \quad (6)$$

#### 4. SOMson INTERFACE

SOMson is designed as an interactive online demo that can be explored using a computer mouse [31]. It is programmed using *JavaScript*. For sound synthesis, the *p5.sound* library is used, which provides straightforward access to the *Web Audio API* [39].

The SOMson user interface is fairly simple. A screenshot is presented in Fig. 3. It consists of the SOM visualization on the left and sonification parameters on the right. Six buttons allow switching between the U-matrix (*Map*) and the component planes of the four components *PhaseSpace*, *ChannelCorrelation*, *PhaseSpaceHigh* and *bpm* as well as showing/hiding the training data.

SOMson is controlled using a computer mouse. The sonification is interactive. Left-click a node on the map and hold the mouse button to hear the sonification of the unit. It represents the magnitudes of the unit, i.e., the four-dimensional pointer. Note that this information is not visible in either of the maps. Only the sonification provides this information. You can move from node to node to hear their difference. Only the node at which the cursor points will be sonified. This way, you can explore the map interactively. Release the mouse button to stop the sonification. This way, you can compare dedicated nodes or items.

The sonification sounds the same, no matter whether you load the U-matrix or any of the component planes. However, when one of the *component planes* is loaded, activating the *ID* button will freeze all sliders except one. This way, the sonification is coherent with what you see. This may also help in learning to distinguish the different sonification parameters. Using the mouse to move the sliders without clicking on the map will also change

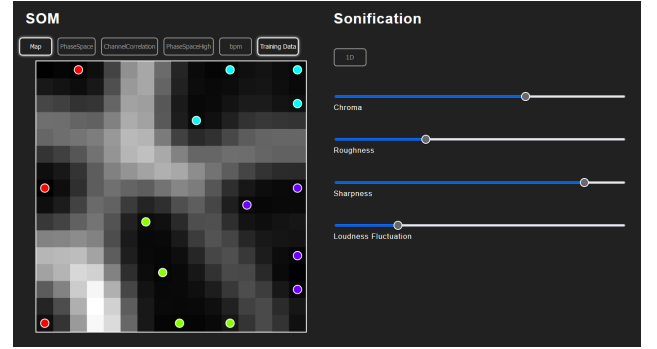


Figure 3: The SOMson interface with the visualizations on the left and the sonification parameters on the right. Buttons allow switching between U-matrix and the component planes, and showing/hiding the training data.

the selected parameter. Finally, while clicking on the map, the sliders on the right also move according to the magnitude of the selected units. This provided visual feedback may not significantly help explore the SOM's data, as it distracts attention from the map. Still, it provides helpful feedback when learning to distinguish the different audible parameters.

#### 5. SOMson: GUIDED TOUR

In this section, we guide you through SOMson. We recommend exploring our interactive SOMson project on <https://simon-linke.github.io/SOMson/simple/>. Alternatively, you can watch all single steps in our YouTube-playlist [https://www.youtube.com/watch?v=VqfHfaI\\_aVA&list=PLVv3BMS8IIXFoSn6p2svrIhNmPHupIW1c&index=1](https://www.youtube.com/watch?v=VqfHfaI_aVA&list=PLVv3BMS8IIXFoSn6p2svrIhNmPHupIW1c&index=1).

**Step 1:** Compare items between two islands on the U-matrix: the red item in the upper-left corner and the green items. You may start comparing the red item with the uppermost green item. Click on the red one, then the green one, then the red one again. Repeat if you need more time. Concentrate on what has changed, by how much, and in what direction: pitch, roughness, sharpness, and loudness fluctuation. You may write down your observations. Then, explore the sound of all green items. What do they have in common, what is different, and by how much?

Differences between the islands that you can hear are:

1. The red item has a much lower pitch than the green items
2. The red item sounds sharper/brighter than the green items
3. The loudness of the red item fluctuates much faster than the loudness of the green items
4. The red item sounds subtly less rough than the green items

Repeat your auditory inspection within all green items. Differences between items on a single island (the green ones) are:

1. They have a very similar pitch
2. They are audibly rough
3. They all sound fairly sharp/bright (but the uppermost one is less sharp than the others)
4. Their loudness fluctuates slowly, especially the item on the lower left



**Step 2:** Stay at the U-matrix: The gray level indicates how similar neighboring fields are regarding the features with which the SOM has been trained. Now, explore the different fields on the map.

What you can hear is how similar the fields sound to their neighbors. Instead of summarizing all features to a single attribute (like the gray level in the visualization), the sonification indicates the magnitudes of all 4 features with which the SOM has been trained. With some practice, you can hear out which and by how much the 4 features have changed from one field to the next.

**Step 3:** In the U-matrix visualization, all clusters look alike: black islands mildly separated by gray lines or clearly separated by white lines. But sonification tells you more. Please explore how the reference island of green items (green island) sounds different from the island of purple and cyan items (purple and cyan island). Take your time and write down how chroma (pitch), sharpness (brightness), roughness, and loudness fluctuations differ. You can see that the green island is only mildly separated from the purple island, while the cyan island is clearly separated from the green island. Thanks to SOMson, you can also hear that

1. the green island sounds more similar to the purple island than to the cyan island
2. on the purple island, the intensity of all sound attributes has raised a bit
3. on the cyan island, the chroma and sharpness have raised a bit, while roughness and loudness fluctuation have raised dramatically.

Recall the mapping between feature and sound attributes. The sound indicates that songs on the cyan island are faster (bpm) and more mono (ChannelCorrelation) than the others.

Note that the respective video has some high-frequency artefacts that are not present in the interactive demo.

**Step 4:** Pull all sliders to the left, and then gradually change the magnitudes of all 4 parameters. You can hear that changing the magnitude of one parameter exclusively changes the intensity of one single sound attribute. It does not affect the others. This is one important requirement of multidimensional sonification: orthogonal dimensions must not interfere perceptually. Another requirement is that you can distinguish many levels of each parameter. This means the dimensions have a high resolution. The third requirement is that the dimensions are continuous: No interruption or *jump* is heard when gradually moving the sliders. Last but not least, the relationship between the slider position and the perceived intensity of the respective sound attribute is linear: small motions always sound like small changes, no matter where the slider is located.

**Step 5:** Switch to the *PhaseSpace* component plane, check the *ID* box, and explore the first dimension. Before you do so, you may move all sliders towards the left for a more pleasant sound.

When browsing through the map, you hear the pitch change from low (dark blue) to high (yellow). You can simultaneously see and hear the magnitude of this dimension: some fields are the same, some are very similar, and some gradual changes occur.

You can make very many observations. For example,

1. some purple items are more closely related to cyan items than to the remaining purple items
2. even though far apart on the map, the red items are fairly similar to each other
3. the red islands have a much lower pitch than all other fields.

**Step 6:** Switch to the *ChannelCorrelation* component plane, check the *ID* box, and explore the second dimension.

When browsing through the map, you hear the clean sound in the dark blue region. Subtle differences between dark blue fields of different shades are audible. When moving from blue via turquoise to yellow, the roughness increases. Fields with similar colors also exhibit a similar degree of roughness.

**Step 7:** Browse through the *PhaseSpace* component plane in *ID* mode. Most fields on the map have a large value, yielding similar colors between lime green and yellow and a bright timbre between quite sharp and really shrill. Visually and auditory, the largest contrast can be found in the lower-left corner. Here, the color and sharpness do not gradually fade but exhibit obvious steps. From lime green to dark blue and from quite sharp to dull. This effect stays audible, no matter what chroma, roughness level, or loudness fluctuation frequency you choose.

**Step 8:** Explore the *bpm* component plane. The feature magnitude of bpm is mapped to the speed of loudness fluctuation. In the dark blue region, the loudness fluctuates very slowly. The fluctuation is getting faster over blue, turquoise, and lime green to yellow. Through this mapping, you can easily compare different fields on the map. Even though some turquoise fields look the same, you can hear which one fluctuates faster.

**Step 9:** Hear where seeing tricks you. Go, e.g., to the *PhaseSpace* component plane and click on two spatially separated fields that you think look the same. You will realize that they often do not sound the same. You can hear which one has a higher magnitude, i.e., a higher pitch.

Such comparisons are of particular difficulty in vision, as human color and lightness vision is affected not only by the focused color but also by the relation to its neighboring colors and lightness levels, as exemplified in Figs. 4 and 5 in color (comparable to the component planes) and in grayscale (comparable to the U-matrix). This is one of the reasons why interactive sonification has been proposed as a complement for lightness, color, and contrast enhancement of visualizations [40].




Figure 4: *Bezold effect: Colors and shades may appear different depending on their surrounding colors. Here, all three notes have the same (single) color, even though the one on the left may appear darker, and the one in the middle may seem to have a color gradient.*

**Step 10:** Now that you have experience using and interpreting SOMson, you should explore and understand the underlying, invisible data of the SOM. For example, click on the red item on top to hear its feature magnitudes. If you have a feeling for the sounds already, you may realize that its pitch is on the lower side, its sharpness is very high, and the other attributes are somewhere in the middle.




Figure 5: *The Bezold effect also holds for lightness: All notes have the exact same (single) color and lightness level, even though the one on the left may appear darker and the one in the middle appears as if it had a lightness gradient.*

Next, you should compare it to the red item in the middle. In what respects is it different?

- It has a (slightly) higher pitch
- It sounds (slightly) rougher
- It sounds less sharp
- Its loudness fluctuation is slower

Now, compare the red item in the middle with the red item at the bottom. How is the one at the bottom different?

- It has a much lower pitch
- It sounds rougher
- It sounds way less sharp
- Its loudness does not fluctuate (magnitude of 0).

When you listen with care, you will be able to interpret and compare the feature magnitudes of all items and nodes.

## 6. DISCUSSION

Even though one of the authors had no previous experience with psychoacoustic sonification, both of us could work with it right away. It may take a while to a) learn on what sound aspects to concentrate on, b) get a feeling for the absolute magnitude of single features, and c) manage to integrate all sound attributes to get an overall picture of the feature magnitudes. But we could instantly hear a) which items were similar and which were not, b) where islands and seas were, and c) whether items differed in pitch, the speed of loudness fluctuation, or were alike. Some information added to the SOM via sonification integrates seamlessly into the workflow, so gathering the necessary experience is only a matter of practice time. The sliders provided for each sonification parameter can further help reduce a user’s adaptation time by providing visual feedback while exploring the map and allowing users to manipulate each parameter individually.

Of course, there are attempts to visualize multiple features of a SOM in a single representation. E.g., [41] proposes to map the magnitude of the features to specific color channels of the RGB colorspace. An interactive example of this approach is shown here: <https://musicai.uni-hamburg.de/en/how-does-a-kohonen-map-work/>. Nevertheless, even this straightforward approach must be learned, as some experience is needed to recognize the single color channel from any given

color. Further, it limits the dimensions of the feature space to a maximum of three.

Note that when using the auditory instead of visual parameters, a four-dimensional sonification is not a limit at all. For example, amplitude-based panning can be utilized to localize the Shepard tone at different azimuth angles, referred to as *auditory event angle* in psychoacoustic terms [42]. This would yield a five-dimensional sonification. Another dimension that has already been suggested [28] and evaluated [30] is the auditory *fullness* that can be implemented by reshaping the spectral envelope of the Shepard tone. Fullness is largely independent of the other five dimensions and fulfills the above-mentioned requirements for multidimensional sonification.

We exemplarily implemented a 7-dimensional sonification available on <https://simon-linke.github.io/SOMson/extended/>. Here, a second auditory stream based on a noise generator is added. The first dimension of the noise is its color. It varies from brown over pink, white, and blue to purple. This mostly affects its brightness. Panning is utilized to implement the second dimension of the noise, which affects the auditory event angle. In the terminology of auditory scene analysis, the Shepard tone and the noise are segregated auditory streams [42]. This means that not only the sound generator but also the noise generator are distinct sound sources. More importantly, it means that the Shepard tone and the noise tend to be perceived as individual sound sources. The strategy of mapping multivariate data to attributes of various auditory streams has already been proposed in [18]. The benefit of this segregation is that we can add more dimensions to the sonification without producing perceptual interference. A disadvantage is that you cannot accurately interpret several auditory streams’ attributes simultaneously. To hear details, you have to concentrate on one stream and, if required, switch attention to the other. Sometimes, you can easily control your focus of attention. However, especially when drastic changes occur, the sound itself may capture your attention. Adding even more dimension is undoubtedly possible. But at some point, these pieces of information do not integrate well, meaning that the benefit of SOMson gets lost, which is presenting an integrated overview of all features of the SOM and its items. Furthermore, interpreting more dimensions requires better listening skills, more cognitive resources, and, certainly, more training.

In the 7-dimensional demo, we also added the features to a *Modulation Matrix*: 1.) Some mappings between data features and audio parameters are particularly intuitive, such as mapping *bpm* to the speed of loudness fluctuations, where faster fluctuations mean faster music. We, therefore, allow mapping between features and sound parameters to be reassigned. 2.) Sometimes, it makes sense to invert the polarity, e.g., mapping the magnitude of a (hypothetical) *darkness* feature to auditory sharpness/brightness from bright to dull instead of dull to bright. We, therefore, allow each mapping polarity to be inverted. 3.) Sometimes, listening to 7 parameters at once is overwhelming, or the presence of one sound attribute distracts from the others. We, therefore, allow muting selected features.

So far, we have implemented SOMson [31] using the p5.sound library [39] and SOM data in JavaScript Object Notation (JSON), as Web Audio is widely accessible [43]: <https://github.com/Simon-Linke/SOMson> As using Python for data sonification is becoming more and more popular these days [44, 45, 46, 47], we are considering implementing SOMson as a Python package, too.

## 7. CONCLUSION

In this paper, we introduce SOMson, a sonification of self-organizing maps. As neither the U-matrix nor the single component planes provide all information about the underlying feature magnitudes, we provide them by means of interactive sonification. Based on a four-dimensional sonification, we guide readers through SOMson. With this interactive demonstration, readers can experience its benefits rather than imagining it based on demo videos or experiment results. We think SOMson's benefit is so obvious that a formal evaluation is obsolete, especially since most of the dimensions have been implemented and evaluated before [37, 38, 30]. We have developed SOMson for up to 7 dimensions, but with increasing dimensionality, the interpretability requires more listening expertise, becomes more cognitively demanding, and not all aspects of the sound integrate well. Being an interactive sonification, SOMson is an engaging data augmentation, motivating researchers to explore the data in all its facets (see, e.g., [48]), which is exactly what SOMs are made for.

## 8. ACKNOWLEDGMENT

This research was funded under the Program "Innovative Hochschule" (innovative university) by the Federal Ministry of Education and Research (BMBF) of Germany (Grant No.13IHS232C) and the City of Hamburg. We thank Michael Blaß for his amazing Apollon package that includes a Self Organizing Map, which we used for generating our map.

## 9. REFERENCES

- [1] T. Kohonen, *Self-Organizing Maps*. Springer, 1995.
- [2] A. Ultsch, "Self-organizing neural networks for visualisation and classification," in *Information and Classification*, O. Opitz, B. Lausen, and R. Klar, Eds. Berlin, Heidelberg: Springer, 1993, pp. 307–313. [https://doi.org/10.1007/978-3-642-50974-2\\_31](https://doi.org/10.1007/978-3-642-50974-2_31)
- [3] E. A. Fernandez and M. Balzarini, "Improving cluster visualization in self-organizing maps: Application in gene expression data analysis," *Comput. Biol. Med.*, vol. 37, no. 12, pp. 1677–1689, 2007. <https://doi.org/10.1016/j.combiomed.2007.04.003>
- [4] H. S. L. P. Qiu, "Jsom: Jointly-evolving self-organizing maps for alignment of biological datasets and identification of related clusters," *PLoS Comput. Biol.*, vol. 17, no. 3, p. e1008804, Mar. 2021. <https://doi.org/10.1371/journal.pcbi.1008804>
- [5] S.-B. Lee, Y. Choe, T.-S. Chon, and H. Y. Kang, "Analysis of zebrafish (*Danio rerio*) behavior in response to bacterial infection using a self-organizing map," *BMC Vet. Res.*, vol. 11, no. 1, 2015. <https://doi.org/10.1186/s12917-015-0579-2>
- [6] N. Xu, W. Zhu, R. Wang, Q. Li, Z. Wang, and R. B. Finkelman, "Application of self-organizing maps to coal elemental data," *Int. J. Coal Geol.*, vol. 277, p. 104358, 2023. <https://doi.org/10.1016/j.coal.2023.104358>
- [7] R. Bader, A. Zielke, and J. Franke, "Timbre-based machine learning of clustering chinese and western hip hop music," in *Audio Eng. Soc. Convention 150*, May 2021, p. 10473.
- [8] T. Ziemer, "Goniometers are a powerful acoustic feature for music information retrieval tasks," in *49. Jahrestagung für Akustik*, Hamburg, Germany, Mar. 2023, pp. 934–937. [https://pub.dega-akustik.de/DAGA\\_2023/data/articles/000600.pdf](https://pub.dega-akustik.de/DAGA_2023/data/articles/000600.pdf)
- [9] Željko D. Vlaović, B. L. Stepanov, A. S. Anđelković, V. M. Rajs, Z. M. Čepić, and M. A. Tomić, "Mapping energy sustainability using the kohonen self-organizing maps - case study," *J. Clean. Prod.*, vol. 412, p. 137351, 2023. <https://doi.org/10.1016/j.jclepro.2023.137351>
- [10] M. Blaß and R. Bader, "Content-based music retrieval and visualization system for ethnomusicological music archives," in *Computational Phonogram Archiving*, R. Bader, Ed. Cham: Springer, 2019, pp. 145–173. [https://doi.org/10.1007/978-3-030-02695-0\\_7](https://doi.org/10.1007/978-3-030-02695-0_7)
- [11] R. Bader, M. Blaß, and J. Franke, "Computational timbre and tonal system similarity analysis of the music of northern myanmar-based kachin compared to xinjiang-based uyghurethnic groups," *arXiv*, 2021. <https://doi.org/10.48550/arXiv.2103.08203>
- [12] F. Aquistapace, D. Castillo-Castro, R. I. González, N. Amigo, G. García Vidable, D. R. Tramontina, F. J. Valencia, and E. M. Bringa, "Plasticity in diamond nanoparticles: dislocations and amorphization during loading and dislocation multiplication during unloading," *J. Mater. Sci.*, 2023. <https://doi.org/10.1007/s10853-023-09223-7>
- [13] J. Qian, N. P. Nguyen, Y. Oya, G. Kikugawa, T. Okabe, Y. Huang, and F. S. Ohuchi, "Introducing self-organized maps (som) as a visualization tool for materials research and education," *Results Mater.*, vol. 4, p. 100020, 2019.
- [14] T. Ziemer, P. Kiattipadungkul, and T. Karuchit, "Acoustic features from the recording studio for music information retrieval tasks," *Proc. Mtgs. Acoust.*, vol. 42, no. 1, p. 035004, 2020. <https://doi.org/10.1121/2.0001363>
- [15] T. Ziemer and H. Schultheis, "PAMPAS: A PsychoAcoustical Method for the Perceptual Analysis of multidimensional Sonification," *Front. Neurosci.*, vol. 16, 2022. <https://doi.org/10.3389/fnins.2022.930944>
- [16] S. Ferguson, W. Martens, and D. Cabrera, "Statistical sonification for exploratory data analysis," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Berlin: COST & Logos, 2011, pp. 175–196. <https://sonification.de/handbook/chapters/chapter8/>
- [17] E. S. Yeung, "Pattern recognition by audio representation of multivariate analytical data," *Anal. Chem.*, vol. 52, pp. 1120–1123, 1980. <https://doi.org/10.1021/ac50057a028>
- [18] S. Barrass and V. Best, "Stream-based sonification diagrams," in *Int. Conf. Auditory Display (ICAD 2008)*, Paris, Jun 2008. <http://hdl.handle.net/1853/49945>
- [19] D. Black, J. A. Issawi, C. Hansen, C. Rieder, and H. Hahn, "Auditory support for navigated radiofrequency ablation," in *CURAC — 12. Jahrestagung der Deutschen Gesellschaft für Computer- und Roboter Assistierte Chirurgie*, W. Freysinger, Ed., Innsbruck, Nov 2013, pp. 30–33.
- [20] T. Ziemer, "The sonic tilt competition," in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, June 2023, pp. 203–206.

- [21] K. Groß-Vogt and C. J. Rieder, “A-e-i-o-u — tiltification demo for ICAD2023,” in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, June 2023.
- [22] P. Kuchenbecker, “Voice balance — a spirit level based on vocal sounds,” in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, June 2023.
- [23] J. Niestroj, “Level assistant — a sonification-based spirit level app,” in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, June 2023.
- [24] S. Barrass, “Soniclevel-pobblebonk app for the ICAD 2023 sonic tilt competition,” in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, June 2023.
- [25] T. Ziemer, “Mixing levels — a rock music spirit level app,” in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, June 2023, pp. 233–235.
- [26] T. Ziemer, N. Nuchprayoon, and H. Schultheis, “Psychoacoustic sonification as user interface for human-machine interaction,” *International Journal of Informatics Society*, vol. 12, no. 1, pp. 3–16, 2020. <http://www.infoc.org/journal/vol12/12-1>
- [27] T. Ziemer and H. Schultheis, “A psychoacoustic auditory display for navigation,” in *Int. Conf. Auditory Displays (ICAD 2018)*, Houghton, MI, USA, June 2018, pp. 136–144. <http://doi.org/10.21785/icad2018.007>
- [28] T. Ziemer and H. Schultheis, “Psychoacoustical signal processing for three-dimensional sonification,” in *25th Int. Conf. Auditory Displays (ICAD2019)*, Newcastle Upon Tyne, UK, June 2019, pp. 277–284. <https://doi.org/10.21785/icad2019.018>
- [29] T. Ziemer and H. Schultheis, “Three orthogonal dimensions for psychoacoustic sonification,” *ArXiv Preprint*, 2019. <https://doi.org/10.48550/arXiv.1912.00766>
- [30] T. Ziemer, “Three-dimensional sonification as a surgical guidance tool,” *J. Multimodal User Interfaces*, vol. 17, no. 4, pp. 253–262, 2023. <https://doi.org/10.1007/s12193-023-00422-9>
- [31] S. Linke and T. Ziemer, “SOMson - Interactive sonification of self-organizing maps,” Apr. 2024. <https://github.com/Simon-Linke/SOMson/>
- [32] M. Asendorf, M. Kienzle, R. Ringe, F. Ahmadi, D. Bhowmik, J. Chen, K. Hyunh, S. Kleinert, J. Kruesilp, X. Wang, Y. Y. Lin, W. Luo, N. Mirzayousef Jadid, A. Awadin, V. Raval, E. E. S. Schade, H. Jaman, K. Sharma, C. Weber, H. Winkler, and T. Ziemer, “Tiltification/sonic-tilt: First release of sonic tilt,” in <https://github.com/Tiltification/sonic-tilt>, 2021.
- [33] T. Ziemer and H. Schultheis, “The CURAT sonification game: Gamification for remote sonification evaluation,” in *Int. Conf. Auditory Display (ICAD 2021)*, Virtual conference, June 2021, pp. 233–240. <https://doi.org/10.21785/icad2021.026>
- [34] M. Asendorf, M. Kienzle, R. Ringe, F. Ahmadi, D. Bhowmik, J. Chen, K. Huynh, S. Kleinert, J. Kruesilp, Y. Lee, X. Wang, W. Luo, N. Jadid, A. Awadin, V. Raval, E. Schade, H. Jaman, K. Sharma, C. Weber, H. Winkler, and T. Ziemer, “Tiltification — an accessible app to popularize sonification,” in *Int. Conf. Auditory Display (ICAD 2021)*, Virtual Conference, June 2021, pp. 184–191. <https://doi.org/10.21785/icad2021.025>
- [35] S. Schwarz and T. Ziemer, “A psychoacoustic sound design for pulse oximetry,” in *Int. Conf. Auditory Display (ICAD 2019)*, Newcastle, June 2019. <https://doi.org/10.21785/icad2019.024>
- [36] A. Greindl, P. Heidegger, K. Groß-Vogt, and M. Weger, “Expergefaktor: Sonic interaction design for an alarm clock app,” in *15th Int Audio Mostly Conf (AM’20)*, Graz, Sept. 2020, p. 4 pages. <https://doi.org/10.1145/3411109.3411149>
- [37] T. Ziemer, “Visualization vs. sonification to level a table,” in *Proceedings of ISON 2022, 7th Interactive Sonification Workshop*, Delmenhorst, Germany, 2022.
- [38] T. Ziemer and H. Schultheis, “Psychoacoustic auditory display for navigation: an auditory assistance system for spatial orientation tasks,” *J. Multimodal User Interfaces*, vol. 13, pp. 205–218, 2019. <https://doi.org/10.1007/s12193-018-0282-2>
- [39] The Processing Foundation, “p5.js-sound,” in <https://github.com/processing/p5.js-sound>, 2024.
- [40] N. Rönnerberg, “Sonification supports perception of brightness contrast,” *J. Multimodal User Interfaces*, vol. 13, no. 4, pp. 373–381, 2019. <https://doi.org/10.1007/s12193-019-00311-0>
- [41] R. Ponnmalai and C. Kamath, “Self-organizing maps and their applications to data analysis,” Lawrence Livermore National Lab., Livermore, CA (United States), Tech. Rep., 2019.
- [42] T. Ziemer, “Sound terminology in sonification,” *J. Audio Eng. Soc.*, vol. 72, no. 5, pp. 274–289, 2024. <https://doi.org/10.17743/jaes.2022.0133>
- [43] H. Lindetorp and K. Falkenberg, “Sonification for everyone everywhere: Evaluating the webaudioxml sonification toolkit for browsers,” in *Int. Conf. Auditory Display (ICAD 2021)*, Virtual Conference, June 2021, pp. 15–21. <https://doi.org/10.21785/icad2021.009>
- [44] D. Worrall, *Sonification Design. From Data to Intelligible Soundfields*. Cham: Springer, 2019. <https://doi.org/10.1007/978-3-030-01497-1>
- [45] D. Reinsch and T. Hermann, “Interacting with sonifications: The mesonic framework for interactive auditory data science,” in *Proceedings of the 7th Interactive Sonification Workshop*, N. Rönnerberg, S. Lenzi, T. Ziemer, T. Hermann, and R. Bresin, Eds., Delmenhorst, Germany, Sept. 2022, pp. 65–74. <https://doi.org/10.5281/zenodo.7552242>
- [46] D. Reinsch and T. Hermann, “sonecules: a python sonification architecture,” in *Int. Conf. Auditory Display (ICAD 2023)*, Norrköping, Sweden, 2023, pp. 62–69. <https://doi.org/10.21785/icad2023.5580>
- [47] J. W. Trayford and C. M. Harrison, “Introducing strauss: A flexible sonification python package,” in *Int. Conf. Auditory Display (ICAD 2023)*, Linköping, Sweden, June 2023, pp. 249–256. <https://doi.org/10.21785/icad2023.1978>
- [48] A. Hunt and T. Hermann, “Interactive sonification,” in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Berlin: Logos Publishing House, 2011, pp. 273–298. <https://sonification.de/handbook/chapters/chapter11/>