

THE LOCAL HEAT EXPLORATION MODEL FOR INTERACTIVE SONIFICATION

Till Bovermann, Thomas Hermann, Helge Ritter

Neuroinformatics Group, Faculty of Technology,
Bielefeld University, Bielefeld, Germany.

[tboverma,thermann,helge]@techfak.uni-bielefeld.de

ABSTRACT

This paper presents a new sonification model for the exploration of topographically ordered high-dimensional data (multi-parameter maps, volume data) where each data item consists of a position and feature vector. The sonification model implements a common metaphor from thermodynamics that heat can be interpreted as stochastic motion of 'molecules'. The latter are determined by the data under examination, and 'live' only in the feature space. Heat-induced interactions cause acoustic events that fuse to a granular sound texture which conveys meaningful information about the underlying distribution in feature space. As a second ingredient of the model, data selection is achieved by a separated navigation process in position space using a dynamic aura model, such that heat can be induced locally. Both, a visual and an auditory display are driven by the underlying model. We exemplify the sonification by means of interaction examples for different high-dimensional distributions.

Keywords: Sonification, Exploratory Data Analysis, Interaction, Multi-modal Computer Interfaces

1. INTRODUCTION

The society we live in depends to a significant extent on the automatic collection and processing of data. Commercial market analysis as well as many other fields of research use data ascertainment and analysis to gain insight into market structures respectively other domain specific coherences.

Whereas conventional exploration methods operate well on low-dimensional data, nowadays the challenge shifts more and more to the analysis of large and high-dimensional data sets as they are subject to the research in Data Mining [1]. New exploratory data analysis techniques are required to discover unexpected or even unknown structure in such data. While mainly visualization was considered over the last decades, in particular auditory display techniques offer a refreshing alternative approach, since our auditory system is sensitive to a quite complementary set of features, and sonification may augment visual displays so that in the end perceptual limitations in a single modalities can be overcome.

This brings up the question how displays in different modalities can be bound together so that they are experienced as a whole rather than different parts. We see synchronisation of display activity, and continuous interaction as key aspects in multi-modal displays to attain this coherence. We start from *Model-based sonification* which provides a framework for the definition of interactive exploratory sonification systems [2] and generalize this framework towards multi-modal data displays (cf. Sec. 2). We illustrate this approach by means of a new sonification model for high-dimensional data with topographical organization. In the presented *Local Heat Exploration Model* (LHEM, cf. Sec. 3), we extend Model-based Sonification along two directions, (i) by introducing dynamic selection models, and (ii) by allowing parallel displays for different output modalities. In Sec. 4 we give some example sonifications for synthetically rendered structures, followed by a discussion about the benefits of LHEM.

2. MODEL-BASED EXPLORATION

The human perceptual system is highly optimized to cope with a specific mix of modalities (e.g. acoustic, tactile, visual, olfactory sensations, etc.) as they usually provide the sensory image of processes in the world. We observe synchronisation between visual and auditory changes, and stimuli are connected in a particularly coherent manner. It is by *interacting* with the world how we ourselves cause multi-modal responses to appear and it is closeby to expect that the human brain truly makes use of the combined perceptions to understand the world.

In contrast we frequently see in current systems for exploratory data analysis combinations of visualizations and sonifications that lack this linking, but that are rather arbitrary combinations of individual displays.

Model-based Sonification (MBS) proposes a framework how to bind interaction to informative acoustic responses about data under analysis, and puts into focus an underlying *dynamic model* that mediates between the data in their hostile high-dimensional space, and the sound space our ears are familiar to analyse [2, 3, 4]. We present a generalization that puts dynamic models in the center of a multi-modal

display so that – same as in the real world – the different displays are fed by model-specific information. Doing this, above motivated principles like synchronisation, coherent binding of modalities, etc. are automatically implemented in the displays. We name this extension *Model-based Exploration* (MBE) to emphasize the intrinsically multi-modal approach.

Physical systems may offer many inspirations for such multi-modal exploration systems, think for instance of a wind chime, or a tree. Yet MBS/MBE is able to describe even unphysical systems which may, however, be more suited to intermediate abstract structures in high-dimensional spaces to a balanced mix in different displays. Conceptually we inherit most aspects from Model-based Sonification, extending it by replacing the uni-modal sonification rendering with a multi-modal display rendering. We now introduce a new sonification model *Local Heat Exploration Model* to give an example for MBS/MBE system design.

3. LOCAL HEAT EXPLORATION MODEL

In many situations the high-dimensional data under analysis can be divided into two parts: (i) topographically ordered attributes \mathbf{x}_p , referred to as position, and (ii) measured features at that position \mathbf{x}_f . Examples are medical volumetrical data, computer tomographics, geologic measurements, or satellite images. Even if there is no topographical information given, techniques like Kohonen’s Self Organizing Maps (SOM, [5]) can be used to assign a positional dimension in a principled way. Such a data set is sufficiently described by

$$\mathbf{X} = \{(\mathbf{x}_p^\alpha, \mathbf{x}_f^\alpha)\}_{\alpha \in \{1 \dots m\}}, \text{ with } \mathbf{x}_p^\alpha \in \mathbb{R}^{d_p}, \mathbf{x}_f^\alpha \in \mathbb{R}^{d_f}, \quad (1)$$

Given a data set \mathbf{X} the local distribution in feature space of an environment around a given point in position space can lead to possibly interesting structures such as clustering or locally linear dependencies in feature space. A standard method to examine the conditional distribution is to estimate it in an offline analysis process, and present the results by using data display tools like glyphs or scatter plots [6]. In contrast to this discretized examination, our approach enables the user to experience local distributions interactively and in real-time. For this the user navigates a selection aura (cf. Sec. 3.2) in position space with the help of a scatter plot which dynamically highlights selected items. At the same time, a sonification model processes user interactions and renders the acoustic response due to LHEM’s evolution. Specifically the excitory process is metaphorical given by induction of heat to the selected data (cf. Sec. 3.3). Motivated by the physical model of heat, this *local heat* causes interactions between the data items in feature space, depending on their local vicinity. Sound is the response to heat induced activity. In result the model allows

to explore neighborhood structures by listening to its output without any need for any dimensional reduction.

3.1. Exploration System Architecture

The LHEM description above suggests to discern three modules in the system’s design:

Data Selection The interactive process of navigating position space. Selected data are subject of exploration.

Exploration Model The dynamical model whose configuration is determined by the selected data.

Exploration Display The perceptual front-end to the user, links temporal evolution of the exploration model to perceivable units.

By design, interaction is an integral part of LHEM. Lack of interaction will cause the system to come to rest at a stable state, producing no audible output. It therefore does not disturb the user while possibly working on other things.

Selection and Exploration are separated interactions, for which generally any kind of continuous input device could be chosen, but preferably an intuitive coupling should be achieved. In LHEM we use a dynamic selection aura controlled by a graphics tablet with software velocity preprocessing allowing the user to *push* the aura literally. For the exploratory interaction we decided to use a finger pressure sensitive device such as the Audio-Haptic Ball described in [7].

One advantage of the current architecture is that interfaces can be easily exchanged due to the clear modularity.

3.2. Selection Aura

One goal of the module described here is to provide an easy to use but also highly interactive selection tool to mark the currently examined data items, since the later described exploration process should be the main focus of the user. For this a close link between the user’s intention and the data selection process has to be accomplished.

Inspired by the work of Fernström [8] who uses an aura to compute the general amplitude of a sonified data record, in LHEM an aura is used to interactively focus on parts of the data. This is done by providing a movable location and a radius in the position space of the data set to the user. For this let $C = (\mathbf{c}, r) \in (\mathbb{R}^n \times \mathbb{R})$ describe a sphere (the *aura*) in \mathbb{R}^n centered at \mathbf{c} with radius r . An item \mathbf{x}^α then is called *selected*, if $\|\mathbf{c} - \mathbf{x}^\alpha\| \leq r$.

Instead of direct control of the position \mathbf{c} , we propose a *dynamic* aura, so that it performs a damped motion if pushed into a certain direction \mathbf{p}_u . Technically this is achieved by the computation of the new aura center by

$$\begin{aligned} \mathbf{c}[t + \Delta t] &= \mathbf{c}[t] + \mathbf{v}[t] \Delta t \\ \mathbf{v}[t + \Delta t] &= \lambda \mathbf{v}[t] + \mathbf{p}_u \Delta t \end{aligned} \quad (2)$$

In fact we use a slightly modified model, which prevents the aura from running wildly through the data set by pushing it multiple times into the same direction [9].

We add a bounding box to the selection model which surrounds the data items. When the aura touches an edge it bounces back like a ball, so it remains within the data set.

The resulting aura selection model then consists of the user's input, a data set to operate on, an initial internal state (being no speed and radius zero, resulting in no initial selection), the update algorithm of the internal states (Eq. (2)) and its output being the current selection.

3.3. Heat Model

The selected data items are processed by an exploration model. The here described Local Heat Exploration Model only considers the feature part of the selected data items. For each feature \mathbf{x}_f^α a point mass is attached in model space at coordinates x_f^α . In addition a heat parameter h^α is associated with it indicating the amount of external heat affecting it. It decays exponentially over time to $h^\alpha = 0$, and can be interpreted as proportional to stochastic movement of a virtual element from its center position $\mathbf{x}_{feature}^\alpha$. If h^α is big enough collisions between neighboring data elements occur at a rate that depends on their summed heat parameters. Modelling this behavior results in a so-called *heat vector* for each pair of selected elements ($\mathbf{x}_f^\alpha, \mathbf{x}_f^\beta$) given by

$$h_{\alpha,\beta} = \begin{cases} \frac{(2r-d)^2}{2r} & d < 2r \\ 0 & d \geq 2r, \end{cases} \quad (3)$$

where $d = d(\mathbf{x}_f^\alpha, \mathbf{x}_f^\beta)$ is an arbitrary but fixed distance function defined on $\mathcal{D} \ni e^\alpha$ controlling the sonification described in the next subsection. The computation is systematically shown in Fig. 1.

As shown in Sec. 4, using the computational heat results in distinguishable output depending on the actual data density of the current data selection. Feature vectors similar to each other produce high heat values, whereas dissimilar features result in lower ones.

3.4. Sonification Display

Instead of modeling stochastic interactions of the *data molecules* explicitly, we render the audio stream by superimposing lots of short *grains* to compose a *grain cloud* [10, 11]. Each of the cloud forming grains lasts for approx. 5 ms, approaching the minimum time for frequency and amplitude discrimination. In a superposition of hundreds of these short-duration grains even minor variations in their duration can cause strong side effects in the spectrum of the cloud sound texture. Grain clouds therefore are predestined to be used as a synthesis technique for auditory displays of complex data sets where the availability of discriminating sound




Figure 1: The heat of an object is expressed by a radius r (red) of a surrounding sphere (cf. Sec.3.4). The overlap of two heats results in output $h > 0$ (green).

controlling parameters are appreciated. The advantage of using grain clouds over other sound rendering techniques is the ability to change the sound quality according to a mass of nearly arbitrary definable low level parameters in short time intervals. This is because the source of the single grains can be rendered using approximately all possible synthesis techniques modulated by the multiplication of an amplitude envelope, e.g. a triangle function forming a short time lasting grain. The rendering technique, its control values and the envelope's quality can be controlled for every rendered grain. For LHEM an additive grain cloud is synthesized by superimposing grains

$$g(t) = \sum_{i=0}^{m-1} a_i (e(t) - o_i) \sin(2\pi f_i t) \quad (4)$$

where

- $m \in \mathbb{N}$ number of mixed sine oscillators,
- $e(t) : \mathbb{R} \rightarrow \mathbb{R}$ the amplitude envelope,
- $f_i \in \mathbb{R}$ oscillators frequencies,
- $a_i \in [0, 1]$ maximum amplitude of frequencies,
- $o_i \in \mathbb{R}$ times delaying the amplitude envelopes.

The linking of LHEM to synthesis then proceeds over the following controls:

Cloud Density d average number of grains per seconds.¹,

Grain Duration u the time, in which $e(t) \neq 0$,

Grain Oscillator Frequencies f_i

¹definition according to [11]

Grain Amplitudes a_i

Onset Delays o_i the onset of the amplitude envelopes of the single frequencies could be delayed from the real triggering event.

Input	$\mathbf{x}^\alpha = (\mathbf{x}_h^\alpha, \mathbf{x}_d^\alpha, \mathbf{x}^\alpha) \in \mathbb{R}^{d_h+d_d}$
Grain	$g^\alpha = (s_0^\alpha, \dots, s_{m-1}^\alpha)$
Cloud Density	$d^\alpha = \frac{\lambda_d}{\sum_{i=0}^{m-1} x_{d,i}^\alpha + 1}$
Sub-Grain	$s_i^\alpha = (f_i^\alpha, o_i^\alpha, u_i^\alpha, a_i^\alpha)$
Grain Partial Frequencies	$f_i^\alpha = \exp(\lambda_{f0} + (\lambda_{f1} x_{d,i}^\alpha))$
Grain Partial Attack	$o_i^\alpha = \lambda_{o0} + \lambda_{o1} (1 - e^{-\lambda_{o2} \hat{p}(x_{h,i}^\alpha)})$
Partial Amplitude	$a_i^\alpha = \frac{\lambda_a x_{h,i}^\alpha}{\max_j(x_{h,j}^\alpha)}$

Table 1: Summary of grain cloud parameters and its corresponding functions customized to LHEM. $\lambda_d, \lambda_{f0}, \lambda_{f1}, \lambda_{o0}, \lambda_{o1}, \lambda_{o2}, \lambda_a$ are scaling factors ensuring that the resulting sound is in an audible range.

Despite its only few controls, this simple grain cloud technique is able to produce a wide range of dynamical changing sounds. The actual mapping of heat vectors to sound parameter is shown in Tab. 1.

3.5. Implementation

LHEM was implemented in EmSie [9], a framework for distributed computation of multi-modal and interactive exploratory data analysis tasks with special focus on Model-based Sonification and Model-based Exploration.

The architecture of EmSie enforces a modulated structure of LHEM by dividing it into the parts listed next. The resulting application structure is illustrated in Fig. 2.

Selection Model Input Module to connect controllers with a selection model. As example we use a graphics tablet to push the selection aura.

Exploration Model Input Only two parameters, the induced heat and its damping factor are needed to be fed into the system. We currently use a haptic controller [7].

Selection Model For LHEM this module implements the dynamic selection aura (cf. Sec. 3.2). As input it processes the Selection Model Input.

Exploration Model Here the Heat Model described in Sec. 3.3 is located.

Selection Model Display We currently display the selection in a dynamic scatter plot visualization.

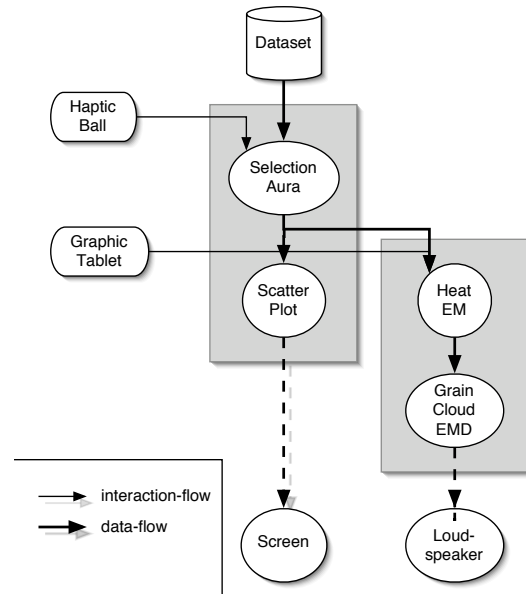


Figure 2: Schematic figure of LHEM.

Exploration Model Display This module implements sound rendering, so in LHEM the grain cloud sonification.

An advantage of using EmSie for LHEM is the opportunity to reuse the implemented modules in other exploration tasks.

The actual implementation of these modules was done by creating classes in SUPERCOLLIDER [12], an object-oriented programming language designed primarily for algorithmic interactive composition and sound rendering, but also very suitable for rapid prototyping of sound based exploration systems. Additional modules, e.g. the input layer model using the Audio-Haptic Ball are build up in NEO/NST, a graphical programming environment developed at Bielefeld University [13]. The communication between the single modules is realized by a network interface defined in EmSie and based on *Open Sound Control* (OSC, [14]).

3.6. Convergence of Sonifications in Performance Scaling

A data set can be regarded as a limited sample of an underlying unknown distribution. The field of statistics provides means to estimate unknown parameters (like densities, variances etc.) with the nice property of convergence: the larger the sample, the better the estimate. We propose to orient the design of sonifications along the same principle so that sonifications converge with increasing sample size *acoustically* towards a sound that represents the exact distribution. This principle is applied here to decouple the sonification of large

data sets from the available computation power, and is a key ingredient to ascertain real-time operability of LHEM.

Specifically, a problem occurs on the computation of the local heat for all records, which requires to evaluate all pair distances and thus $O(N^2)$ operations where N is the cardinality of the current selection. Since the Local Heat Exploration Model should enable the user to change the controls like the induced heat at control rate, this computationally complex task has to be done in real-time. When exploring a large data set, the 'just-in-time' computation is infeasible. Motivated by the fact that in data exploration accentuating on data inherent parameters such as the local feature density, only a qualitative view on the data is intended, a simplification of the algorithm can be made, by computing only a limited number of heat values for each data item like shown exemplary in Fig. 3. Altering the subset of corresponding data records each time step in a random manner assures only limited artifacts that converge acoustically with subset size to the original task. Setting the number of computed heats to the number of given data items results in the computation of all heats, i.e. the task equals to the one without performance scaling.




Figure 3: Choose n arbitrary points out of all available and compute the distances between them and the actual data item. For each item this process is done for different randomly chosen data items. The light painted cross marks the actual processed data item.

The advantage of this method consists in the reduction of computational steps to be done in one time step: It is now in linear complexity concerning the involved data records. On average all pairs of data records are minded in the exploration process, and therefore observed by the user. Therefore it could be interpreted as a time domain load balancing.

The performance scaling can easily be mapped onto other computationally expensive data processing algorithms to adapt them to a real-time task for exploration.

4. INTERACTION EXAMPLES

For evaluation, synthetically rendered test data sets are used, since they allow to control accurately the data distribution being the basis of the analysis process of the systems dynamic behavior. In addition the user can learn to interpret LHEM's output for later exploration of unknown data sets.

4.1. Test Data Sets

A test data set \mathbf{X} is created in two steps.

Position Rendering Due to the fact that the selection process is less important to evaluate, since its functionality is easy to understand by using it, the data positions are two-dimensional uniform drawn from a distribution on $[0, 1] \times [0, 1]$.

Feature Rendering The feature vectors \mathbf{x}_f have to be chosen independent from each other, but depending on their position. So a transfer function $f : \mathbb{R}^{d_p} \rightarrow M$ (with a suitable set M) must be defined such that it is possible to generate a probability density p_f on \mathbb{R}^{d_f} via f .

Let $f : \mathbb{R}^{d_p} \rightarrow [0, 1]$ be well defined. The probability density p_f may be defined as

$$p_f(\mathbf{x}_f) = \theta_{0, f(\mathbf{x}_p)}(\mathbf{x}_f), \quad (5)$$

with $\theta_{a,b}$ the uniform distribution density on $[a, b]$.

All following data sets consist of 2000 data items with a three-dimensional feature part and are rendered according to Eq. (5).

4.1.1. Gaussian Feature Distribution




Figure 4: The transfer functions $f(x)$ of (a) a Gaussian feature distribution and (b) a hollow sphere distributed test data set. Darker points represent lower variances. The red line shows the moving center of a selection aura (cf. Sec.4.2).

The Gaussian distributed data sets are rendered using

$$f(\mathbf{x}) = 1 - e^{-(\lambda \|\mathbf{x} - \mathbf{x}_0\|)^2}, \quad (6)$$

where \mathbf{x}_0 defines the center of the distribution and $\lambda \in \mathbb{R}$ is proportional to the width of the Gaussian's bell. It is shown in Fig. 4(a). The nearer a vectors position is to \mathbf{x}_c , the more similar a feature vector is to its neighbors. Therefore the feature distance between two likely centered vectors is even lesser than the distance between two vectors far away from \mathbf{x}_0 .

4.1.2. Hollow Sphere Feature Distribution

The hollow sphere distributed data sets are computed according to

$$f(\mathbf{x}) = 1 - e^{-(\lambda(\|\mathbf{x}-\mathbf{x}_0\|^2 - r^2))^2} \quad (7)$$

with $r \in \mathbb{R}$ being the radius of the resulting d_p -dimensional sphere shown in Fig.4. Although the rendered data sets look similar to the Gaussian distribution, the feature parts of data records located near the origin are less similar to their neighbors. Near the sphere surface data records are rendered similar to the Gaussian distributed data set at this location. This leads to a distance hole in the middle, whereas viewed from outside it looks like a Gaussian data set (cf. Fig. 4).

4.2. Test Scenarios

Different interaction scenarios are used to explore the data sets. The sounds and their belonging spectrograms described next are provided on our web site [15].

4.2.1. Determined Selection Aura Path

To illustrate the system's ability to distinguish between feature vector distances, this scenario varies the selection's position in time, i.e. the Selection Aura's center. For all examples listed below the same trajectory was used, shown in Fig.4. All other input parameters remain fixed, i.e. the aura's radius is $r = 0.04$, the positional damping factor $d_p = 0$, the induced heat $h = 1$ and the heat damping factor $d_h = 0$. This controlling scheme was applied to the following data sets:

Gaussian Distribution with $\lambda = 3$ The selection of vectors near the center results in higher sounds separable from each other, while selections far away from the center are more deep and fuzzy.

Gaussian Distribution with $\lambda = 30/\lambda = 300$ The smaller the Gaussian, the smaller is the region in which the exploration system calms down, and its sound is lowered. For $\lambda = 300$ only a really small area of equally distances can be recognized.

Hollow Sphere Distribution with $\lambda = 5, \lambda = 50, \lambda = 100$ When moving through the sphere surface into the center ($t = 15$ secs) the exploration sounds like flying

through a curtain of deeper pitched sounds, caused by the high similarity of the feature vectors in the selection aura, followed by a much clearer sound. The smaller the sphere's rim, the more salient the effect.

The advantage of LHEM over classic visualization displays is emphasized in the last scenario. Since the visualization techniques only let the user see the surface of the sphere, and therefore do not show the inner distribution of feature vectors, a Gaussian-like distribution is expected (c.f. Sec. 4.1.2). With the selection aura of LHEM the user is able to navigate right into the center of the sphere, and hear the significantly different sounds comparing to a Gaussian distributed data set.

4.2.2. Interactive Exploration




Figure 5: (a) Interactive exploration of a Gaussian data set. (b) The induced external heat as a function of time. Each peak indicates a new heat induction.

This scenario shows how the system behaves, when used interactively. Here a Gaussian distributed data set with $\lambda = 30$ was explored, using a graphic tablet for pushing the aura and the Audio-Haptic Ball for inducing heat. As can be seen in Fig. 5 the trajectory produced by the interaction of the user with the system is not that straight as in the scenario described in Sec. 4.2.1. This is caused by the fact that the user has refined his activity by means of the system's auditory feedback to navigate the selection aura.

There are at least two different principles, how to explore a data set using LHEM, like shown in this exploration example. They differ in the way, how the user induces external heat to the exploration model. One possibility is to continuously vary the external heat and therefore neglecting the systems ability of changing the heat dynamically (4s-11s). Another possibility is to induce only at one time a heat value and then let the system fading it out. This produces a sound feedback varying according the heat damping parameter (11s to 50s). If the heat damping factor is set to $d_d = 0$ (50s to end), a sound similar to the one described in the variable position scenario is produced. In this case the user's main control is given by the navigation of the aura.

5. DISCUSSION

We have presented a new technique for interactive sonification of topographically structured high-dimensional data. The design shares with other sonification models that the sonification can be computed without any manual parameter tuning for all such data, independent upon dimensionality or size. This *generality* is a particular benefit which enables the user to profit from learning effects that automatically appear during ongoing use.

We'd like to stress the high interactivity in LHEM, all interactions result in immediate auditory feedback within few milliseconds. The user can at any time profit from insights and refine the exploratory activity by means of the continuously increasing knowledge. Performance Scaling has been motivated as one principle to achieve this. The model-based sonification approach offers the chance of automatically create action-responses that we are familiar with from everyday interactions (like here the decay after interaction stops). Particularly if physical dynamics are adapted for use in the model, sound structures usually result that we even denote as natural – and we expect that this has an positive influence on learning time, interpretation effort, error rate, etc. The extension from sonification models to multi-modal exploration models has been motivated in Sec. 2. LHEM makes first steps towards this goal by combining a visual selection model display and an auditory exploration model display by means of an underlying model. An obvious benefit is that the modalities mutually support each other (here sound and spatial navigation are coupled). The user can attend spatial structures and acoustic structures simultaneously whereas in uni-modal displays the display selection disrupts the exploratory flow of interactions.

Modeling LHEM with EmSie is advantageous since it demands clear logical separation and enforces a modular implementation of software components for specific sub-tasks. While this may be harder during implementation we profit from (i) a continuously growing set of reusable software classes, and (ii) a clearly facilitated distribution on several computers in the network, which is particularly desirable in MBS/MBE due to the typical high computational complexity. The application of LHEM to real-world data, and the extension of EmSie by diverse interfaces and interaction types is subject to our ongoing work.

6. REFERENCES

- [1] U. M. Fayyad et al., Ed., *Advances in Knowledge Discovery and Data Mining*, MIT Press, 1996.
- [2] T. Hermann, *Sonification for Exploratory Data Analysis*, Ph.D. thesis, Bielefeld University, Bielefeld, February 2002.
- [3] T. Hermann and H. Ritter, "Listen to your data: Model-based sonification for data analysis," in *Advances in intelligent computing and multimedia systems, Baden-Baden, Germany*, G. E. Lasker, Ed., pp. 189–194, Int. Inst. for Advanced Studies in System research and cybernetics, 1999
- [4] T. Hermann and H. Ritter, "Sound and meaning in auditory data display," *Proceedings of the IEEE, Special Issue Engineering and Music*, 2004.
- [5] T. Kohonen, *Self-Organizing maps*, Springer Series in Information Sciences Vol. 30. Springer, Berlin, Heidelberg, New York, third edition, 2001.
- [6] M. Berthold and D. J. Hand, Eds., *Intelligent Data Analysis: An Introduction, second edition*, chapter 11, pp. 403–427, Springer-Verlag New York, Inc., 2003.
- [7] T. Hermann, J. Krause, and H. Ritter, "Real-time control of sonification models with an audio-haptic interface," in *Proc. of the Int. Conf. on Auditory Display*, R. Nakatsu and H. Kawahara, Eds., Int. Community for Auditory Display, pp. 82–86, 2002.
- [8] M. Fernström and C. McNamara, "After direct manipulation - direct sonification," in *Proc. ICAD '98*, 1998.
- [9] T. Bovermann, "A multi-modal exploration framework with application to interactive sonification," M.S. thesis, Bielefeld University, July 2004.
- [10] I. Xenakis, *Formalized Music*, Indiana University Press, Bloomington, 1971.
- [11] C. Roads, *Microsound*, MIT Press, Cambridge, 2002.
- [12] "Supercollider Hub," <http://supercollider.sourceforge.net/>, July 2004.
- [13] H. Ritter, "The graphical simulation toolkit neo/nst," http://www.techfak.uni-bielefeld.de/ags/ni/projects/simulation_and_visual/neo/neo_e.html, 2000.
- [14] M. Wright, "Opensound control specification," <http://www.cnmat.berkeley.edu/OpenSoundControl/OSC-spec.html>, March 2002.
- [15] T. Bovermann, T. Hermann and H. Ritter, "Model-Based Sonification - demonstrations," <http://www.techfak.uni-bielefeld.de/~thermann/projects/index.html>, 1999.