Complementing Visual Data Mining with the Sound Dimension: Sonification of Time Dependent Data

Monique Noirhomme-Fraiture[1], Olivier Schöller[1], Christophe Demoulin[1], and Simeon J. Simoff[2,3]

[1] University of Namur, Institut d'Informatique,
5000 Namur, Belgium
mno@info.fundp.ac.be
[2] School of Computing and Mathematics, University of Western Sydney
NSW 2007 Australia
s.simoff@uws.edu.au
[3] Faculty of Information Technology, University of Technology, Sydney
PO Box 123 Broadway NSW 2007 Australia

Abstract. This chapter explores the extension of visual data mining by adding a sound dimension to the data representation. It presents the results of an early 2001 experiments with sonification of 2D and 3D time series data. A number of sonification means for these experiments have been implemented. The goal of these experiments was to determine how sonification of two and three-dimensional graphs can support and complement or even be an alternative to visually displayed graphs. The research methodology used the triangulation method, combining the automated generation of the sound patterns with two evaluation techniques. The first one included the assessment and evaluation of the sound sequences of the sonified data by the participants in the experiment via a dedicated server. The second one was based on the analysis of an evaluation questionnaire, filled by each participant that performed the tests. The chapter presents the results and the issues raised by the experiments.

1 Introduction

As defined earlier in this book, visual data mining is a part of the KDD process [1], which places an emphasis on visualisation techniques and human cognition to identify patterns in a data set. Ankerst [1] identified three different scenarios for visual data mining, two of which are connected actually with the visualisation of final or intermediate results and one operates directly with visual representation of the data. The design of data visualisation techniques, in broad sense, is the formal definition of the rules for translation of data into graphics. Generally, the term 'information visualisation' has been related to the visualisation of large volumes of abstract data. The basic assumption is that large and normally incomprehensible amounts of data can be reduced to a form that can be understood and interpreted by a human through the use of visualisation techniques. The process of finding the appropriate visualisation is not a trivial one. A number of works offer some results that can be applied as guiding heuristics. For example, Wickens and Carswell [2] defined the Proximity Compatibility

S.J. Simoff et al. (Eds.): Visual Data Mining, LNCS 4404, pp. 236–247, 2008.

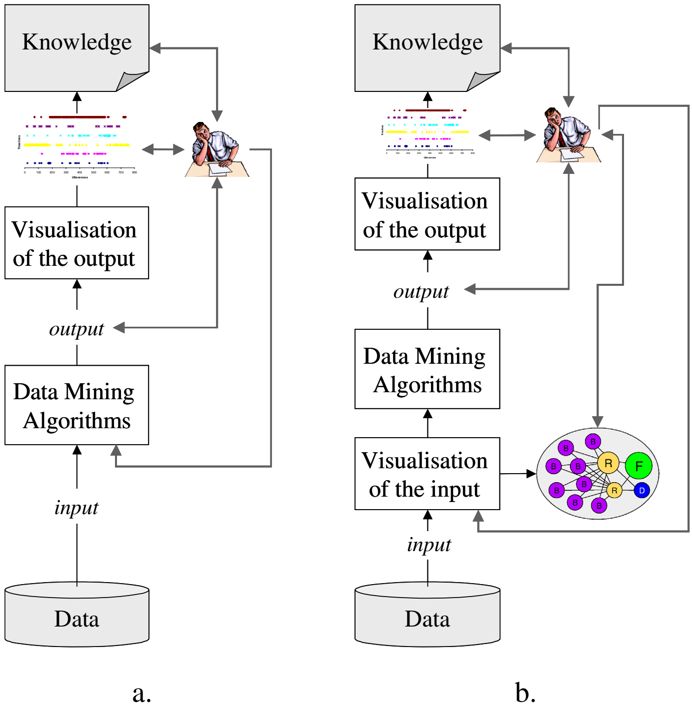

Fig. 1. Visualisation and visual data mining

Principles (PCP) for various visualization methods in terms of tasks, data and displays - if a task requires the integration of multiple data variables, they should be bundled in proximity in an integrated display. Based on this principle authors have concluded that 3D graphs do not have an advantage over 2D graphs for scientific visualisation (which may not necessarily hold for visual data mining).

Visual data mining relies heavily on human visual processing channel and utilises human cognition overall. The visual data mining cycles are shown in Fig. 1. In most systems, visualisation is used to represent the output of conventional data mining algorithms (the path shown in Fig. 1a). Fig. 1b shows the visual data mining cycle when visualisation is applied to the original or pre-processed data. In this case, the discovery of the patterns and dependencies is left to the capacity of the human visual reasoning system. The success of the exercise depends on the metaphor selected to visualise the input data [3].

Although human visual processing system remains a powerful 'tool' that can be used in data mining, there are other perceptual channels that seem to be underused. Our capability to distinguish harmonies in audio sequences (not necessarily musical ones) is one possibility to complement the visual channel. Such approach can be summarised as 'What You Hear Is What You See'. The idea of combining the visual and audio channels is illustrated in Fig. 2.

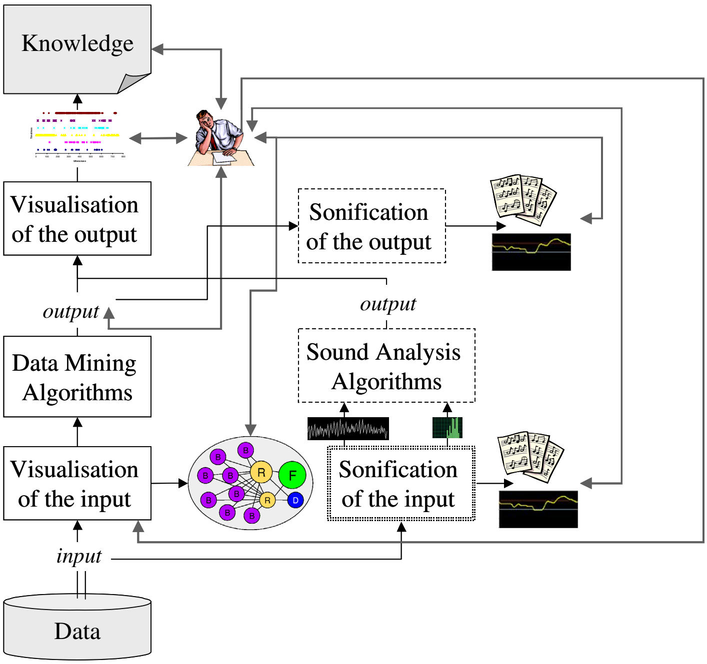

Fig. 2. Combining visual data mining and sonification

The conversion of data into a sound signal is known as sonification. In this sense the operation is similar to visualisation in that it converts abstract data into a form that we can perceive. Hence, similar to the application of visualisation techniques in Fig. 1b, sonification can be used both for representing the input and/or the output of the data mining algorithms. It can be the basis of mining techniques on its own (for example, see the Sonify system for Excel[1], developed in 2005 by Accentus), as well as in complementing visualisations.

This chapter looks in the second case, when sonification complements visualisation techniques. In this case sonification is expected to be synchronised with the visualisation technique. Further, in this chapter we discuss the issues connected with designing such data mining techniques and present an example of a practical implementation of combined technique. We briefly discuss the characteristics of the sound that are suitable for such approach, the actual sonification process, the design of the overall combined technique and the results of the experiments conducted with proposed technique.

2 Characteristics of Sound for Time Dependent Data Representation

Several researchers have investigated the use of sound as means for data representation [4-11] and interfacing computer systems [12] . In this context, the important

[1] http://www.sonifyme.com/

feature of the sound is that it has a *sequential nature*, having particular duration and evolving as a function of time. A sound sequence has to be heard in a given order, i.e. it is not possible to hear the end before the beginning[2]. Similarly, a time series data depends on time and have the same characteristics of a sequence. This similarity leads to the assumption that sound provides good means to represent time series data.

3 Sonification

A straight-forward way to transform time dependent data into sound is to map the data to frequencies by using linear as well as chromatic scale mappings. We call this process a pitch-based mapping. We compute the minimum and maximum data values from the chosen series and map this data interval into a frequency range, chosen in advance. Each value of the series is then mapped into a frequency. To avoid too large, non-realistic intervals, we first discard outliers.

Another pre-processing and transformation procedure is the smoothing of the time series. In fact, if we map all the points of a series into a sound, we will hear rather inconsistent sounds. Hence, as a first step we smooth the series by a standard mean, for example, using the moving average method. As a second step, we map the smoothed curve into pitch. Beat drums can be used to enhance the representation of the shape of the curve, as shown further in the chapter.

3.1 Detection of Outliers

To detect statistically the values of the outliers, a confidence interval is computed at each time t, based on the normal distribution. Once a data value is detected outside the confidence interval, the corresponding time value is stored and sonified.

3.2 Beat Drums Mapping

The mapping between a “beat drum” sound and the time series is as follows: the rhythm of a beat drum increases with respect to the rate of growth of the smoothed curve (i.e. the first derivative).

3.3 Stereo Panning

Variation of the stereo acoustics is introduced, for example, an increase of the volume of the right speaker and decrease of the volume of the left speaker.

3.4 3D Curve

When the time series is given at each time, not by a single value but by a function of values, one can choose to “hear” at each discrete time the function. One can also choose to cut the surface at a certain level and to hear “continuously” the obtained curve as a

[2] Images and drawings do not have such constraint. Strictly speaking digital sound recording can provide access to an arbitrary section of the sound fragment, even reproduce the sound in reverse order, which is beyond the scope of this paper.

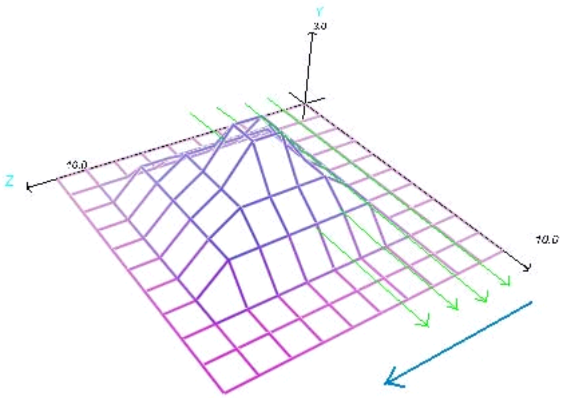

Fig. 3. Example of a surface that can be sonified

function of time. We call these transformations respectively horizontal and vertical travelling. An example of a 3D data surface for sonification is shown in Fig. 3.

4 Experimental Workbench for Data Sonification and Mining

The prototype experimental workbench has been implemented in Java programming language, using the MIDI package of the Java Sound API[3] [13]. The MIDI sequence is constructed before the actual playback. When the designer starts the sonification, the whole sequence is computed. Then computed sequence is sent to the MIDI sequencer for playback.

5 Design of the Experiment and Methodology of Data Collection

The purpose of the experiments is to determine how the sonification of two and three-dimensional graphs can complement or be an alternative to visually displayed graphs. As visual data mining complemented with the sound dimension is a human-centred process, a Web site has been created, where sound sequences are presented and can be evaluated by the participants in the experiment. The site contained a questionnaire that had to be filled in by each participant performing the test. The questionnaire included questions for collecting demographic and skills data about each participant, information about the techniques, data mining questions and questions for evaluation of the individual sonification techniques and the overall approach. The structure of the questionnaire and the accompanying information is as follows:

A. Participant identification: name, age, gender, title/position, e-mail address. These data are used to identify the subject and to validate the answer.

---

[3] API – **A**pplication **P**rogramming **I**nterface.

B. Participant abilities: field of activity, musical experience (instrument played, practicing period), self-evaluation of musical level (from ‘no experience’ to ‘expert level’).

C. 2D evaluation: 2D evaluation is divided into three subtasks:

C.a: **Explanation of sonification techniques:** Each of the four sonification techniques used - pitch based only, beat drums, stereo and extreme values detection, was briefly described and at least one example was given. Participants in the experiments needed to understand which specific features in the data were manifested by each of the four sonification technique.

C.b: **Application test:** Time series data sets were presented to each participant. Each time precise questions were asked. For the 2D and 3D series four and three questions or group of questions were asked, respectively. Question 1 aimed to evaluate if a subject could perceive a global trend in the series and to understand if the relation with the time scale could be detected. For each sequence, beat drums and stereo mapping were added to enhance the pitch-based sonification. Question 2 aimed to identify whether extreme values were detected. Question 3 aimed to identify whether seasonal trend could be detected. Question 4 was focused on general trend identification. For each question, the subject needed to specify the number of times he had listened to the sonified data before answering the question.

C.c: **Subject preferences:** Each participant had been asked another group of four questions, aiming to evaluate subject preferences:

1. *Choice of instrument for pitch mapping*. Each subject was asked to hear the sonification of the same data, but with different MIDI instruments for the pitch mapping. Then the subject was asked to grade on the scale 0 to 10 the different instruments (acoustic grand piano, steel string guitar, violin, synthstrings 2, pan flute). These instruments were very different and belonged to a specific MIDI group including piano, guitar and the string group.
2. *Choice of instrument for drums mapping*. Participants were presented with different instruments (Celesta, Slap Bass 1, Timpani, Tinkle Bell, Woodstock) for the beat drums mapping, which had to be evaluated.
3. *Choice of sonification technique*. The subject had to give his opinion on the different mapping techniques: pitch, beat drums, stereo mapping, extreme values detection and on the sonification in general. The four levels for evaluation included: useless, sometimes useful, always useful and essential.
4. *Open question*: “Please tell us what you think about our project, our applications, this web page or anything else that comes to mind”.

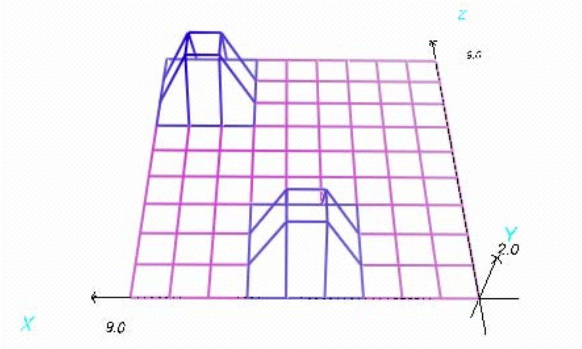

Fig. 4. Data for vertical travelling

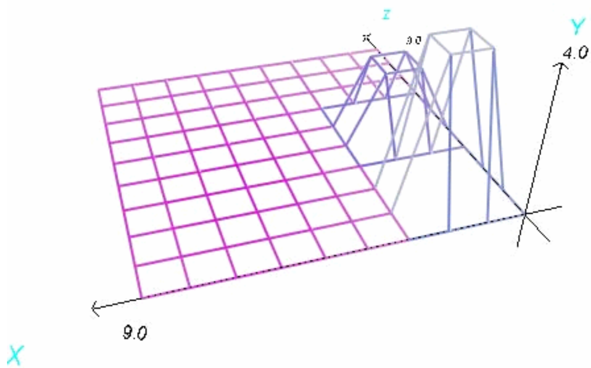

Fig. 5. Data for horizontal travelling

D. 3D Evaluation. Similar schema is used for 3D curves, which includes explanation about the sonification method, the application test (3 questions), and the subject preferences. The sonification method is based on slicing of the initial surface following some direction: vertical (see Fig. 4), horizontal (see Fig. 5) and diagonal.

Each 2D line obtained by the slicing process can be heard, similar to the 2D case. The different lines, proposed one after the other, are separated by a specific sound. Beat drums can still be added to pitch mapping.

6 Results of the Experiments

Below we present the results of the experiments. We describe the data sample that was obtained from the run of the experiment and the results for the 2D and 3D cases.

6.1 The Sample of Participants

23 participants answered the questions for the case of a 2D visualisation and 18 participants - for the case of 3D visualisation. A large part of the sample (9 participants)

includes people working in the computer science area, who have limited musical experience or no experience at all. To see whether musical experience is beneficial when dealing with sonified data, e.g. if people with musical experience get a better score, we have compared the average score in both groups. The influence of the musical and computer science background on the results is presented in Table 1 and Table 2, respectively. The score obtained for each question is based on the level of each answer, normalised on a score out of 100. The average score is the mean of the scores for the different questions.

Having a musical or a computer science background gives a minor advantage in using sonification of 2D curves. The differences are not significant. The experiments did not display any difference for the 3D sonification case. The way 3D case has been implemented is rather complex and uses a good spatial representation. Musical or computer experience has a minor influence on the result in this case.

Table 1. The influence of musical background on results

	2D sonification		3D sonification
	Av. score/100	**N**	**Av. score/100**
Musical background	76	9	40
No experience	66	14	42
		23	

The same comparison was done for computer scientists.

Table 2. The influence of computer science background on results

	2D sonification		3D sonification
	Av. score/100	**N**	**Av. score/100**
Computer science	72	11	41
Others	69	12	40
		23	

6.2 Results in 2D

The following questions were included in the questionnaire, targeting 2D sonification:

Q.1 Given the annual data about sheep population in England and Wales between 1867 and 1939 (see Fig. 6), please, answer the following:

1.1 Were there more sheep in 1867 than in 1939?
1.2 About which year did the sheep population reach the minimum?

Q.2 Given the daily record of the morning temperature of an adult woman for two months, please, answer the following:

2.1 Did she have fever during the period?
2.2 If yes, for how long did she have the fever?

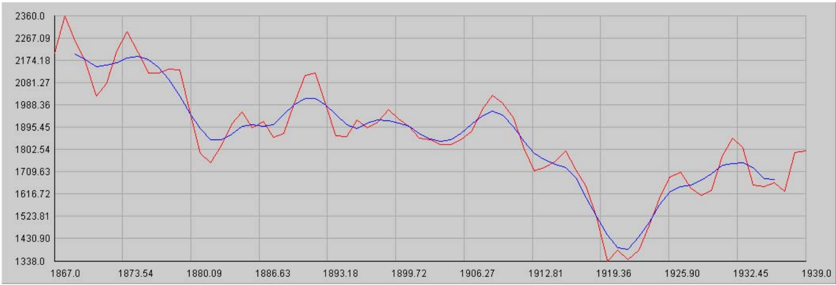

Fig. 6. Annual sheep population in England and Wales between 1867 and 1939

Q.3 Given the data about the monthly electricity production in Australia between January 1956 and August 1995, please, answer the following:

3.1 Is the electricity production in Australia lower in 1956 than in 1995?
3.2 How would you categorise the evolution of electricity production in Australia: as linear or as exponential?
3.3 Is the evolution of electricity production in Australia characterised by seasonal trend?

Q.4 Given the data about the monthly Minneapolis public drunkenness intakes between January 1966 and July 1978 (151 months), please, answer the following:

4.1 Were there more intakes in 1966 than in 1978?
4.2 Is the evolution of public drunkenness intakes linear?
The results are summarised in Table 3.

Table 3. Summary of the results for 2D

		Correct	Wrong	No idea
Q.1	1.1	17	5	1
	1.2	17	6	-
Q.2	2.1	23	0	0
	2.2	13	10	0
Q.3	3.1	22	1	0
	3.2	12	11	0
	3.3	16	6	1
Q.4	4.1	20	2	1
	4.2	19	3	1

6.3 Results in 3D

The following questions were included in the questionnaire, targeting the 3D sonification:

Q.1 A 3D graph containing 2 bumps has been sonified. The selected mapping is the vertical travelling and the sonification starts from the bottom right corner.

- If the grid below (3 x 3) represents the graph, where are these 2 bumps located?
- Do they have the same height?

Q.2 Same questions with respect to the horizontal travelling mapping.
Q.3 Same questions with respect to the diagonal travelling mapping.

The results are summarised in Table 4.

Table 4. Summary of the results for 3D

		Correct		Wrong	No idea
		2 correct	*1 correct*		
Qu1	1.1 Trend	4	6	8	-
	1.2 Value	11		4	3
Qu2	2.1 Outlyers	4	6	8	0
	2.2	12		33	0
Qu3	3.1 General trend	1	11	6	-
	3.2	10		6	2

7 Discussion

Though the data sets and the visualisations have been carefully selected to fit the purpose of the experiments, it is worth mentioning couple of “difficulties” introduced in the design of the experiments that could have had some influence on the outcome of the experiments:

1. In a graphical representation, if one wants to identify a particular point, s/he needs to find the information concerning that point on each axis. In the experiment, we provided little information about the scale (for example, see Fig. 4 and Fig. 5). The wording gives the limits for the time period. The lack of scaling information could have caused some difficulty in identifying particular points or sub-periods.
2. The outcomes in the case of sonification of a 3D graph are worse than in the case of 2D. When considering this result, it is worth taking in account that the sonification of a 3D graphical representation is more difficult then the sonification of a 2D graphical representation. A possible reason could be that the sonification technique is based on visual representation and does not use sound properties, but surface properties, as seen in a three axes referential.

Another important issue that needs to be taken in account, when designing visual data mining systems with sound dimensions is the correspondence between visual and audio representations of the data. Consistent representations should provide audio representations that allow transitions from 3D to 2D projections in terms of corresponding sound representations.

8 Conclusions

Overall, the results of the experiments on sonification of time dependent data leave optimism for further investigation of sound as medium for presenting information. The original experiments have been conducted in the spring of 2001. The development of sonification systems like Sonify mentioned in the introduction of the chapter, oriented towards identifying patterns in high volume real-time series data, as in stock markets supports this optimism.

The sound can be an effective complementary interface to the visual interface for data representation. Similar results were presented by Alty [5] for people with disabilities. Extending the human perception of data through multi-sensory human interfaces has been the goal of making line graphs accessible to blind students through auditory and haptic media [14]. On the other hand, the experiments with sonification of surfaces in 3D space showed that sonification did not efficiently support the visual representation and certainly could not replace it at least in the form it had been implemented in the experiments.

In general, this early experimental work contributes to the research efforts on bringing other (non-visual) channels for information and data processing. This research area extends visual data mining, shifting the focus towards interactive systems that support rich perceptual – visual, audio, tactile – interaction between the human and the data representation. Such systems are expected to play significant role in assisting data understanding and supporting pattern discovery process, utilising human information processing capabilities [15].

Acknowledgements

We thank our colleagues Anne de Baenst and Florence Collot for their assistance in this research work.

References

1. Ankerst, M.: Visual Data Mining, in Ph.D. thesis. Dissertation.de: Faculty of Mathematics and Computer Science, University of Munich (2000)
2. Wickens, C.D., Carswell, C.M.: The proximity compatibility principle: Its psychological foundation and relevance to display design. Human Factors 37(3), 473–495 (1995)
3. Simoff, S.J.: Towards the development of environments for designing visualisation support for visual data mining. In: Proceedings Int. Workshop on Visual Data Mining, 12th European Conference on Machine Learning and 5th European Conference on Principles and Practice of Knowledge Discovery in Databases ECML/PKDD2001, Freiburg, Germany (2001)
4. Anrijs, K.: The use of sound in 3d representations of symbolic objects. Facultés Universitaires Notre-Dame de la Paix, Namur (1999)
5. Alty, J.L., Rigas, D., Vickers, P.: Using music as a communication medium. In: Proceedings of CHI 1997 (1997)
6. Brewster, S.A., Wright, P.C., Edwards, A.D.N.: An evaluation of ear-cons for use in auditory human-computer interfaces. In: Proceedings of Inter CHI 1993 (1993)

7. Conversy, S., Beaudouin-Lafon, M.: Le son dans les applications interactives. Laboratoire de Recherche en Informatique, Université de Paris-Sud, Paris (1995)
8. Hermann, T.: Data exploration by sonification (1999)
9. Kramer, G.: An introduction to auditory display, Auditory Display: Sonification, Audification, and Auditory Interfaces. Santa Fe Institute Studies in the Sciences of Complexity (1994)
10. Noirhomme-Fraiture, M.: Le son dans les interfaces IHM: Application à la représen-tation de données multivariées complexes, in Actes des 2èmes Journées Multimédia. Namur: PUN (2000)
11. Sahyun, S.C.: A comparison of auditory and visual graphs for use in physics and mathematics. Oregon State University (1999)
12. Kramer, G., Walker, B.N.: Sound science: Marking ten international conferences on auditory display. ACM Transactions on Applied Perception (TAP) 2(4), 383–388 (2005)
13. Java, Javasound API programmer's guide. Sun Microsystems (2000)
14. Ramloll, R., Yu, W., Brewster, S., Riedel, B., Burton, M., Dimigen, G.: Constructing sonified haptic line graphs for the blind student: First steps. In: Proceedings of the Fourth International ACM Conference on Assistive Technologies, Arlington, VA, 13-15 November (2000)
15. Zhao, H.: Interactive Sonification of Abstract Data - Framework, Design Space, Evaluation, and User Tool. PhD Thesis. College Park: University of Maryland, 288 (2006)