INTERACTIVE SONIFICATION FOR VISUAL DENSE DATA DISPLAYS

Conference Paper · December 2016

2 authors:

Niklas Rönnberg
Linköping University
28 PUBLICATIONS   **119** CITATIONS

SEE PROFILE

Jimmy Johansson Westberg
Linköping University
73 PUBLICATIONS   **1,254** CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:

Master Thesis Project View project

Sonification View project

INTERACTIVE SONIFICATION FOR VISUAL DENSE DATA DISPLAYS

Niklas Rönnberg

Division for Media and Information Technology
Linköping University, Linköping, Sweden
Norrköping Visualization Centre-C, Sweden
`niklas.ronnberg@liu.se`

Jimmy Johansson

Division for Media and Information Technology
Linköping University, Linköping, Sweden
Norrköping Visualization Centre-C, Sweden
`jimmy.johansson@liu.se`

ABSTRACT

This paper presents an experiment designed to evaluate the possible benefits of sonification in information visualization to give rise to further research challenges. It is hypothesized, that by using musical sounds for sonification when visualizing complex data, interpretation and comprehension of the visual representation could be increased by interactive sonification.

This hypothesis is evaluated by testing sonification in parallel coordinates and scatter plots. The participants had to identify and mark different density areas in the representations, where amplitude of the sonification was mapped to the density in the data sets. Both quantitative and qualitative results suggest a benefit of sonification. These results indicate that sonification might be useful for data exploration, and give rise to new research questions and challenges.

1. INTRODUCTION

In order to reduce visual clutter and facilitate analysis of large data sets, it is common to employ renderings based on the data density [4]. This is typically achieved by rendering semi-transparent objects and additively blending them together. This can reveal structures in data that otherwise would have been missed. However, using density information has a drawback in that it is difficult to perceive the actual number of blended objects for different areas in the density representation, making it hard to find areas of similar density or find areas of highest density.

The focus of the present study was to explore sonification using musical sounds in visualizations, to give rise to new research questions and challenges for future work. As our research interest was in investigating the interplay between visualization and sonification, as well as the interaction between sonification and user performance, the choice was made to use the density of data sets as the main evaluation task to provide an easily controlled and measurable experiment setup. The study evaluates the use of sonification as an additional modality to visualization, for facilitating the analysis of large data sets that result in visual clutter. Two common visualization techniques are used in the evaluation: scatter plots and parallel coordinates (see Figure 1).
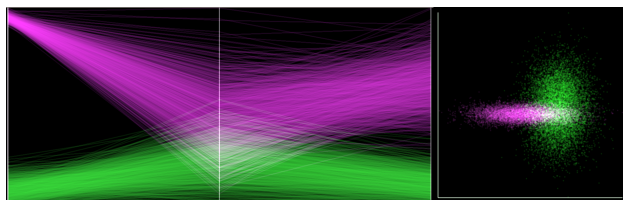

Figure 1. *An example of a parallel coordinates representation (left) and a scatter plot (right), as used in the experimental setup.*

Hence, this evaluation compares the visual representations with only density renderings, using additive blending, and density renderings with added sound (sonification). However, two data clusters with different colours might be perceived as having different intensities if the clusters' densities are mapped to the brightness of the representation. To counteract this a colour model giving equal perceptual lightness has been used [11].

2. RELATED WORK

Even though, the concept of sonification and data exploration is not new (see for example [6]), the combination of visualization and sonification has been sparsely evaluated in various fields of application, for example in connection to depth of market stock data [13], to augment 3D visualization [10], and to enhance visualization of molecular simulations [14]. All these studies suggest that there is a benefit of sonification in visualization. By combining the visual and the aural modalities it should be possible to design a more effective and efficient visualization [16].

When it comes to scatter plots there are not many research articles addressing density representations, however a few studies present interesting work [12, 2]. For parallel coordinates, on the other hand, there is significant research. For the interested reader, the authors refer to, for example, [17, 4, 1, 8, 9, 7]. However, none of this research exploits the use of sonification.

Some research has considered sonification in connection to data exploration and scatter plots. For example, one study [5] evaluated visual as well as auditory scatter plots. However, this was not an evaluation of a simultaneous sonification of a visual representation but a comparison in performance between these two modalities, and there was no user interaction involved. Another study [15] presented a combined auditory and haptic interactive representation of scatter plots, that was successfully evaluated with visually impaired participants. Neither this study evaluated simultaneous visual and sonic representation of data.

The aim of the present study is to evaluate sonification in relation to information visualization of abstract data, to generate research questions and challenges for future work. As far as the authors are aware, this is the first evaluation of this kind.

3. THE SONIFICATION

In this study musical sounds were used. This decision was made because musical sounds when combined together create an emergent musical timbre that is a representation of the combined density of the data sets rendered on the computer screen. These musical sounds create a changing soundscape, which was assumed to bring meaning to the visualization without becoming too constant or repetitive. The two composed sounds, each sound sonifying one data cluster, used in the evaluation differed in pitch as well as in meter (i.e. rhythm), but were tuned and in the same

tempo. The sounds used were made up from synth strings, played softly with the high-frequency content slightly attenuated, with a soft synth bell-like sound accenting the meter of the sounds. The pitches used were C4 and G4, with fundamental frequencies of 261.3 Hz and 392.0 Hz respectively. This interval created a rather pleasant but still separable interval of a fifth [3]. The tempo used was rather slow, 70 bpm, and the different meters of the sounds created a combined rhythm that further enhanced the perception of, as well as the distinction between the two sounds.

A series of pilot tests were performed (n = 20) to verify discrimination between the two sounds, and to discern the just noticeable difference in amplitude for each sound, as well as to normalize the audibility of the two sounds. Of fifteen trials, and five practice trials, the participants successfully distinguished between the sounds 98% (*SD* = 6) of the time. The sounds were deemed to be sufficiently different for participants with different backgrounds and musical expertise, and the few incorrect answers could be explained by mistakenly given answers. The average just noticeable difference (JND) in amplitude for the two sounds was 2.02 dB (*SD* = 0.93). As the amplitude steps used in the sonification were larger than this level, the changes in amplitude were considered audible. When differences in audibility between the two sounds were tested, the participants had to discern the just noticeable difference in amplitude level of one sound while the other sound was held at a constant level. These tests showed that the C4 sound had an upward spread of masking of at average 5.24 dB (*SD* = 5.02), and consequently the C4 sound was attenuated with 2.62 dB to counteract the masking effect. After these tests and adjustments, it was assumed that these sounds could create responses by means of harmony, rhythm and amplitude; illustrating the density of, as well as the blend between, data clusters.

Each sound sonified one of the two data clusters, rendered in purple and green respectively (see Figure 1), and the amplitude of each sound was mapped to the density of the respective data cluster. The density level was mapped between no attenuation and quiet in eleven amplitude steps relative to the percentage of the data cluster's maximum density. This number of amplitude steps were chosen so that the change in amplitude level always were greater than the JND of 2.02 dB. As the user hovered with the computer mouse over the visualization, the amplitude level of the two sounds varied accordingly to the density of the data clusters currently covered by the mouse pointer. The sounds were then mixed together by the visualization software in real-time during the evaluation.

The reader is encouraged to listen to a short excerpt of the sonification to achieve a better understanding of the sonification. The excerpt covers when a user first explores one data cluster, then shifts focus to the other data cluster, and finally investigates the blend between the two data clusters. http://www.itn.liu.se/~nikro27/ison2016/ison2016_sonvis.mp3

4. COLOUR CORRECTION

In both representations, parallel coordinates and scatter plots, data cluster one was coloured using a shade of purple and data cluster two had a shade of green (see Figure 1). This ensured that the two data clusters had uniform perceptual contrast, as presented in [11]. These isoluminant colours are of equal perceptual lightness, which is necessary if the task of the evaluation is to determine different densities of two data clusters.

5. METHOD

For the present study, 20 participants with a median age of 30 (range 20 to 60) with normal, or corrected to normal, vision and self-reported normal hearing were recruited.

Two stimuli were used in the study. One stimulus display was comprised of parallel coordinates, and one of a scatter plot (see Figure 1). Each visual representation included 2 data clusters of various shapes and densities, and in different colours. The data clusters were created using normal distributions with 5,000 to 10,000 samples per cluster. The application was implemented in C++, OpenGL and OpenAL.

The participants had to identify and mark the density areas in both representations using a standard computer mouse. The tasks were performed in two setups: (1) with visual modality alone, and (2) multimodal with both visual and aural modalities. Both test setups had the same visual information, but in the multimodal modality the densities of the two data clusters were sonified by corresponding amplitude levels of the two sounds as the computer mouse hovered over any position in the representation. The order of modalities was balanced between participants to avoid order and learning effects.

Specifically, the tasks were to find the highest density areas in the data clusters, and to find a matching density level in one data cluster from a given density level in the other data cluster. The test was initiated with a practice trial for familiarization, and was then continued with ten trials for each visualization (parallel coordinates and scatter plots) in each modality (visual modality and multimodal). The order of visual representations was balanced to avoid order effects. In the end of the test, the participants were given a questionnaire about their experience of the sonification, and answers were given via a 5-point Likert scale with ratings that ranged from 1 (strongly disagree) to 5 (strongly agree).

Visual stimuli were presented on a 15" computer screen and auditory stimuli through a pair of Beyerdynamic DT-770 Pro headphones. The output of the headphones gave an auditory stimulation of approximately 65 dB SPL. The experiment took place in a single session in a quiet office. Even if there were ambient sounds, the test environment was deemed quiet enough not to affect the tests conducted.

The experiment yielded objective measures of sonification benefit, accuracy and response time automatically recorded in the visualization program, as well as subjective measures manually marked with a pen by the participant.

6. RESULTS

When accuracy for finding the highest density area in the data clusters was analysed using a repeated measures ANOVA with one within-subject factor, sonification (no sonification and sonification), a main effect of sonification was found ($F(1,39) = 12.34$, $p = 0.001$), where accuracy was significantly higher with sonification compared to without. The mean performance for accuracy was 82% (*SD* = 8) without sonification and 86% (*SD* = 5) with sonification (see Figure 2).
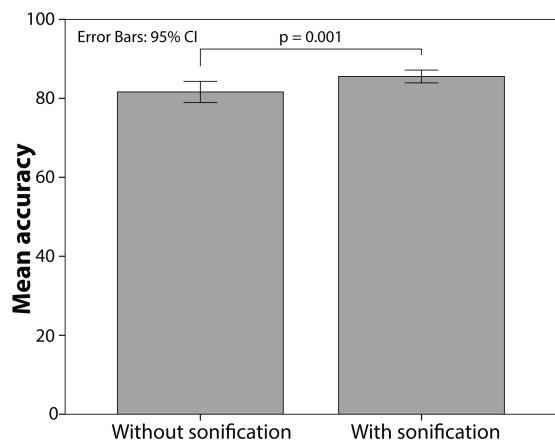
Figure 2. *The mean accuracy when determining the highest density area was significantly higher when sonification was used compared to without sonification.*

When response time for finding the highest density area was analysed using a repeated measures ANOVA with one within-subject factor, sonification (no sonification and sonification), a main effect of sonification was found (F(1,39) = 49.16, $p < 0.001$), where response times were significantly longer with sonification compared to without sonification (see Figure 3). The mean response time was 9.6 seconds ($SD = 6.4$) when no sonification was used, and 20.5 seconds ($SD = 11.0$) when sonification was used.
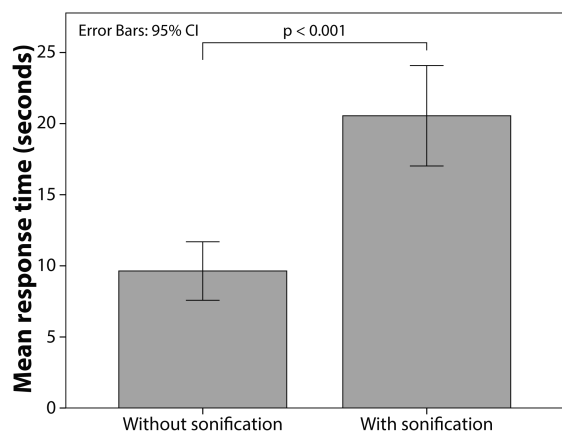

Figure 3. *When searching for the highest density area, the mean response time was significantly longer with sonification compared to without.*

The mean accuracy for matching a given density area in one data cluster with a chosen area by the participant was 16.4 ($SD = 6.4$) without sonification, and 15.5 ($SD = 6.4$), the lower the value, the better match between the density areas. When these accuracy data were analysed with using a repeated measures ANOVA with one within-subject factor, sonification (no sonification and sonification), no significant difference in accuracy was found ($p = 0.533$).

However, when response times for finding the matching density area were analysed using a repeated measures ANOVA with one within-subject factor, sonification (no sonification and sonification), a significant difference was found (F(1,39) = 39.21, $p < 0.001$), where response times were longer when sonification

was used. The mean response time was 11 seconds ($SD = 6.7$) without sonification, and 18.2 seconds ($SD = 10.2$) with sonification (see Figure 4).
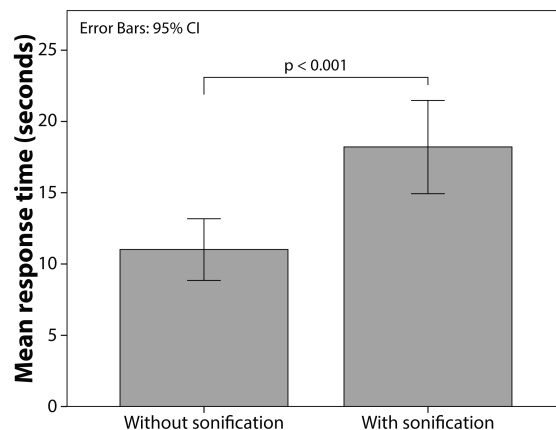

Figure 4. *When the participant had to match two density areas the mean response time was significantly higher with sonification compared to without sonification.*

Regarding the subjective measures from the Likert scale, the participants agreed that there was a benefit of sonification (mean value 4.1, range 3 to 5), see Figure 5. Also, the participants were neutral towards appreciating listening to the sounds (mean value 3.7, range 2 to 5).
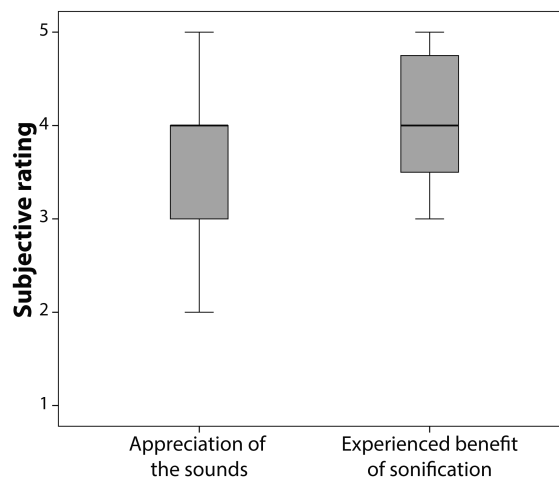

Figure 5. *The participants were neutral towards appreciating listening to the sounds (to the left), and the participants agreed that there was a benefit of sonification (to the right).*

7. DISCUSSION

The results presented in this study suggest that sonification can improve perception of density in visualization of complex data in parallel coordinates as well as in scatter plots. This was shown by improved accuracy when sonification was used when determining the highest density level in a specific data cluster. Consequently, the sonification seemed to supply additional information to the user, thus improving accuracy. However, response time did not improve by sonification. This indicates that

the participants used sonification to improve accuracy of the response, rather than responding faster but with the same accuracy as without sonification. Furthermore, sonification might be more beneficial for tasks where accuracy and precision is more important than swift responses.

When the task was to match density areas, the statistical analysis did not show a significant relationship between sonification and accuracy. The reason for this might be threefold. First, the density areas provided by the test were randomly placed within a data cluster, and had an average density of 49% (range 19-76%). As the sonification gets more attenuated due to less density in the data cluster, the participant's hearing ability takes a more important role for perceiving the sonification. If the participant had problem hearing the sonification, this ought to make the density matching task more difficult. Furthermore, as the sonification gets more attenuated the test situation gets more vulnerable to background noise and disturbances. Second, matching amplitude levels of the two sounds might require more trained ears and experiences that none or few of the participants possessed. Third, this kind of task might simply not be something that sonification can improve.

These possible explanations were to some degree supported by the longer response time when sonification was used. The prolonged response time suggest that the participants tried to use the sonification to improve the accuracy. However, the extended amount of time was in vain as there was no improvement. Consequently, it might rather be reasoned that sonification in this task, at least in the task's current form, decreased overall performance.

The subjective measures supported the objective accuracy measurement, which might suggest that sonification simplified finding the highest density level in each of the data clusters. However, the experience of the sounds was slightly more divergent, but regardless if the participant liked the sounds used for sonification or not, there was a stated experienced benefit of sonification. As always with subjective ratings, the exact reason for how a participant rated, for example, the experienced benefit of sonification might be uncertain. The test situation may prime a certain response, when the participant tries to be helpful, which gives rise to false positive results.

The tasks in this study might be considered as a bit artificial, since a mathematical test could mark the highest density area automatically, rather than demanding a user interaction with visual inspection of the representations. However, at this stage of exploring sonification and visualization, density in a data cluster was a simple parameter to both sonify and visualize. It should be kept in mind that the aim of the present study was to investigate sonification in visualization to generate future research questions, and not to determine the best way to provide information about density levels.

As also stated above the participants' experiences and trained ears, and even musicality might be of importance for distinguishing between the sounds and the amplitude balance between the sounds. In the present study this has not been taken into account, but for future studies the questionnaire should be further developed to include such parameters.

Even though the two sounds were rather clearly separable according to the pilot tests of the sounds, there was a similarity between them since they were tuned and in the same tempo. It is plausible that more diverse sounds might improve performance in tests like the ones that have been used in the present study.

Overall, these results suggest that sonification might be a useful tool for data exploration. The results found in this study are therefore encouraging and give rise to new research challenges.

8. FUTURE WORK

For future work, the first step will be to investigate sonification for a wider range of information visualization representations and techniques for data exploration. This should show where the benefit of sonification is at its greatest, as well as when the visual modality is less loaded or when it is highly loaded. The second step is to investigate whether the benefit of sonification translates from accuracy to response time as well. By creating an evaluation setup that demands fast response times, it should be possible to investigate the benefit of sonification on response time rather than on accuracy.

When these studies have given a basic understanding of how sonification relates to visualization, a third research challenge arises from the possibilities of interactive sonification for data exploration. The choice of evaluation tasks should be further evolved, and the amount of interaction could be increased to further enhance the sonic experience of the visualization, for example by means of changed timbre or harmony to sonify relations between data clusters.

The fourth research challenge is to further explore different kinds of sounds for sonification in connection to the user's musicality, such as different timbres, different tempo and rhythm, as well as different harmonies and intervals. This in turn leads to a fifth research challenge in personification of sounds used for sonification. Most probably, users have different abilities to comprehend and distinguish between musical sounds, as well as respond differently to them and have different taste for the sounds. This leads to a more user experience orientated evaluation setup.

Finally, more research will be needed to explore if and how the experiences and knowledge from these future studies translates to other areas and applications for sonification.

9. CONCLUSIONS

By using interactive sonification when visualizing complex data, the accuracy in finding density areas could be increased. Furthermore, response time increased as participants spent more time in achieving, or trying to achieve, higher accuracy. The current study suggests that sonification is suitable for some aspects of visualization of complex data, like finding the highest density area, but maybe not others. It has also given rise to interesting research questions and challenges for future work.

10. ACKNOWLEDGEMENTS

11. REFERENCES

[1] A. O. Artero, M. C. F. de Oliveira, and H. Levkowitz, "Uncovering clusters in crowded parallel coordinates visualizations", in *Proc. of the IEEE Symposium on Information Visualization (InfoVis)*, Austin, USA, April 2004, pp. 81-88.

[2] S. Bachthaler and D. Weiskopf, Continuous scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1428-1435, 2008.

[3] I. Deliòge and J. Sloboda, *Perception and Cognition of Music*, volume 32, 1997.

[4] G. Ellis and A. Dix, A taxonomy of clutter reduction for information visualisation. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1216-1223, 2007.

[5] J. H. Flowers, D. C. Buhman, and K. D. Turnage, Cross-Modal Equivalence of Visual and Auditory Scatterplots for Exploring Bivariate Data Samples. *Human Factors*, 1997, 39(3), pp. 341-351.

[6] J. H. Flowers, D. C. Buhman, and K. D. Turnage, Data Sonification from the Desktop: Should Sound Be Part of Standard Data Analysis Software?, ACM Transactions on Applied Perception 2:4, 2005, pp. 467-472.

[7] J. Johansson and C. Forsell. Evaluation of Parallel Coordinates: Overview, Categorization, and Guidelines for Future Research. *IEEE Transactions on Visualization and Computer Graphics,* 22(1):579-588, 2016.

[8] J. Johansson, C. Forsell, M. Lind, and M. Cooper, Perceiving patterns in parallel coordinates: Determining thresholds for identification of relationships. *Information Visualization*, 7(2):152-162, 2008.

[9] J. Johansson, P. Ljung, M. Jern and M. Cooper. Revealing Structure within Clustered Parallel Coordinates Displays, in *Proc. of the 11th IEEE Symposium on Information Visualization (InfoVis)*, Minneapolis, Minnesota, USA, October 2005, pp. 125-132.

[10] M. Kasakevich, P. Boulanger, W. Bischof, and M. Garcia, Augmentation of visualisation using sonification: A case study in computational fluid dynamics, in *Proc. of the 13th Eurographics Symposium on Virtual Environments (IPT - EGVE)*, Weimar, Germany, July 2007.

[11] S. Maureen, In color perception, size matters. *Computer Graphics and Applications*, pp. 8–13, 2012.

[12] A. Mayorga and M. Gleicher, Splatterplots: Overcoming overdraw in scatter plots. *IEEE Transactions on Visualization and Computer Graphics*, 19(9):1526-1538, 2013.

[13] K. Nesbitt and S. Barrass, Evaluation of a multimodal sonification and visualisation of depth of market stock data, in *Proc. of the International Conference on Auditory Display (ICAD)*, July 2002.

[14] B. Rau, F. Frieß, M. Krone, C. Müller, and T. Ertl, Enhancing visualization of molecular simulations using sonification, in *Proc. 1st IEEE International Workshop on Virtual and Augmented Reality for Molecular Science (VARMS)*, March 2015, pp. 25-30.

[15] E. Riedenklau, T. Hermann, and H. Ritter, Tangible Active Objects and Interactive Sonification as a Scatter Plot Alternative for the Visually Impaired, in *Proc. 16th International Conference on Auditory Display (ICAD-2010)*, Washington DC, USA, June 9-15.

[16] M. W. Rosli and A. Cabrera, "Gestalt Principles in Multimodal Data Representation", *IEEE Computer Graphics & Applications*, 35(2):80-87, 2015.

[17] E. J. Wegman, Hyperdimensional data analysis using parallel coordinates. *Journal of the American Statistical Association*, 85(411):664- 675, 1990.