

DRAWING BY EAR: INTERPRETING SONIFIED LINE GRAPHS

Lorna M. Brown and Stephen A. Brewster

Glasgow Interactive Systems Group
Department of Computing Science
University of Glasgow G12 8RZ, UK
{lorna, stephen}@dcs.gla.ac.uk
www.multivis.org

ABSTRACT

The research presented here describes a pilot study into the interpretation of sonified line graphs containing two data series. The experiment aimed to discover the level of accuracy with which sighted people were able to draw sketches of the graphs after listening to them. In addition, it aimed to identify any differences in performance when the graphs were presented using different combinations of instruments—either with piano representing both data series (*same-instruments* condition), or with piano representing one data series and trumpet representing the other (*different-instruments* condition). The drawings were evaluated by calculating the percentage of key features present. The results showed that accuracy was high (over 80% on average) in both conditions, but found no significant differences between the two. There were indications of some differences between the two conditions, but a larger study is necessary to discover whether these are significant. The results indicate that graph sonification systems should allow users to choose between these two presentation modes, depending on their preference and current task. The study showed that sonified graphs containing two data series can be interpreted, and drawn, by sighted people, and that evaluation with blind users (our target users) would be worthwhile.

1. INTRODUCTION

Blind and visually impaired people are often deprived of access to information due to the use of visualisations such as graphs. Graphs are present in everyday life, in newspapers and magazines, and are used regularly in the study of mathematics and the sciences. The MultiVis Project aims to make graphs accessible to blind and visually impaired people using the senses which are available to them, namely hearing and touch.

Research has shown that people are able to interpret line graphs that are sonified by representing each data point with a musical note [1-3]. Most research has focused on line graphs containing a single data series, but it has also been shown that it is possible to sonify graphs containing two data series [4-6].

A previous study [4] showed that, when two data series were presented simultaneously, blind people were able to locate intersection points and global maxima and minima. The research presented here describes a pilot experiment which aimed to discover whether, using the same system, users were able to draw

a sketch of a graph containing two data series after listening to the sonified version of the graph.

It has been shown that it is possible to match sonified graphs containing two data series to visual graphs through multiple choice [5]. However, whether people can interpret these sonified graphs without any visual cues being supplied and draw them has not been tested. The ability to draw these graphs is important because it shows the user has constructed a mental model of the shape of each line, and any points of interaction between the two lines. The experiment is also more valid for the target audience if this methodology is used, as blind people do not have access to visual cues.

Since drawing is difficult for blind people, the participants involved in this experiment were all sighted. If the results were to show that sighted people were unable to draw the graphs, it would indicate that blind people would also have difficulty constructing a mental model of the graphs. If, however, it were shown that sighted people could draw these graphs, it would indicate that a future study with blind people (perhaps drawing on swell paper with heat pens) would be worthwhile.

An additional aim of this experiment was to establish whether performance would change if, instead of using piano for both data series, a different instrument was used to sonify one of the series. Using different instruments is potentially advantageous as it may allow the user better perception of the individual shape of each line whilst attending to both simultaneously. A potential disadvantage is that users may find it difficult to compare pitches due to the different tonal qualities of the two instruments.

2. LINE GRAPH SONIFICATION

Sonification is the use of non-speech audio to represent data. A line graph is sonified by representing each data point with a musical note [1]. The y-axis is mapped to pitch, thus the higher the y-value of the data point, the higher the musical note. Moving along the x-axis causes the musical note representing the corresponding y-value to be played. Data is converted to sound by mapping data values to MIDI (Musical Instrument Digital Interface) notes. The range of MIDI notes used is from 35 (B1) to 100 (E7). On a Soundblaster Live soundcard, notes outside this range can be difficult to perceive and differentiate from one another [4], and they are inaudible on some other soundcards.

2.1. Sonification of Graphs Containing Two Data Series

In order to sonify line graphs containing two data series, each data series is sonified as described above. Through headphones, one series is panned to each ear, thus perceptually separating the two. Users listen to both data series simultaneously (one in each ear), such that at any point on the x-axis they will hear the y-value of one data series in their left ear and the y-value of the other data series in their right ear [4]. Stereo panning has been shown to be successful for separating simultaneously playing earcons [7], and this technique can be applied to simultaneously playing data series.

There is some debate as to whether separating simultaneously playing sounds by spatial location alone is sufficient. Deutsch [8] suggests that “pitch grouping” will occur, such that people will group all the high pitches together and all the low pitches together, overriding the grouping by spatial location. This could cause a problem when interpreting graphs, as listeners may get confused between the two lines after an intersection point. Deutsch indicated that pitch grouping may occur even when the sounds are differentiated by the use of different timbres. Bregman [9] states that two sounds separated by spatial location alone will be heard as one combined sound, rather than two individual sounds, but that this will only occur when they are very close in frequency.

Since separation by spatial location alone may not be sufficient to enable users to differentiate between the two data series, it may be beneficial to represent each data series with a different instrument (timbre). The disadvantage of this method is that it can be difficult to compare two pitches played by different instruments, making it harder to locate specific points such as intersections and global maxima and minima.

2.2. Selecting Instrument Combinations

Very little work has been carried out to identify which instruments, and instrument combinations, are successful for presenting graphs containing two data series. Researchers have either chosen to represent both series with the same instrument [4], or to represent each with a different instrument [5], but have not compared these methods. This research aimed to discover which of these methods was more successful.

The Grand Piano sound (MIDI Instrument 0) has been proven to be successful for graph sonification [4] so it was present in both conditions of the experiment described below. In the *same-instruments* condition the piano was used to represent both data series, while in the *different-instruments* condition one series was represented by the piano and the other by the trumpet (MIDI Instrument 56).

Brewster [10] recommends the use of instruments which are subjectively easy to tell apart. Rigas [11] divided instruments into groups, or families, based on people’s ability to recognise them. He placed the trumpet in a different family from the piano, indicating that the piano and trumpet can easily be distinguished from one another. In addition, it was necessary to choose an instrument that could successfully reproduce the same pitch range as the piano. The synthesised trumpet was able to produce this range with reasonable quality.

3. EXPERIMENT

3.1. Hypotheses

The hypotheses for this pilot study were as follows:

1. Drawings of the individual lines on the graph will be more accurate in the *different-instruments* condition than in the *same-instruments* condition as it will be easier to differentiate between the two data series.
2. Drawings of the interactions between the two lines (e.g., intersection points, global maxima/minima, relative y-positions) will be more accurate in the *same-instruments* condition than in the *different-instruments* condition as it is easier to compare the pitch of musical notes when they are played by the same instrument.

3.2. Software Interface

The computer keyboard is a familiar input device for many blind people, therefore the interface for SoundVis (the graph sonification software developed for this research) uses the numeric keypad. The key functions are defined in Table 1. Each key press will cause the notes representing the current y-values to be played.

Key	Function
6 key (right arrow)	move one step right along the x-axis. Holding down this key will play an overview of the graph
4 key (left arrow)	move one step left along the x-axis
5 key	stay on the same position
7 key (Home)	jump to start of graph
9 key (PgUp)	jump to end of graph

Table 1: Key Functions in SoundVis

3.3. Participants

This pilot study was carried out with six sighted subjects (third year, fourth year and M.Sc. Computing Science Students, and an Engineering Ph.D. student). The group consisted of five males and one female, aged between 20 and 26. Two participants described themselves as musicians, two said they had basic musical skills and two said they had no musical experience.

3.4. Method

A within-participants design was used, and the order in which participants took part in the two conditions was counterbalanced. Before starting the experiment, participants received training on using the interface and on interpreting sonified graphs. This included two tasks of the same type as those in the experiment itself, after which the experimenter provided feedback in order that participants could judge their performance.

Each condition in the experiment consisted of ten tasks. A different set of graphs was used for each condition, and all the graphs were generated from basic mathematical functions (sine, cosine, straight lines, quadratics, etc.), such as the graph shown in Figure 1.

In each task the participants explored a sonified graph containing two data series for two minutes in order to build up a mental picture of it. Participants could explore the graphs at their own pace using the interface described in Section 3.2. At the end of the two minutes participants were given one further minute in which to draw a sketch of the graph. Each line was drawn with a different coloured pen, so that the experimenter could differentiate between them. Participants were not allowed to draw, or take notes, during the exploration of the sonified graph, and no feedback on performance was provided during the experiment.

At the end of each condition participants completed NASA TLX scales [12] indicating the subjective workload experienced. In addition, they were asked to state whether, overall, they preferred the *same-instruments* condition, the *different-instruments* condition, or had no preference.




Figure 1: Graph 1—a graph used in the experiment

4. EVALUATION

The drawings were evaluated by the experimenter, who marked each drawing according to the number of key features of the graph included in the drawing. The key features of each graph had been identified prior to the experiment. Performance was measured as a percentage of key features present so as not to bias towards stimuli containing more features. This method of evaluation is subjective, but objective methods would be difficult to employ due to the large amount of data. The features of each graph were classified by whether they were features of the individual lines, or features of the interaction between the two lines. Table 2 lists the key features of the graph shown in Figure 1.

Line 1 (black)	Line 2 (grey)	Interaction
1. Curve	5. Straight line	7. One intersection point
2. Generally decreasing	6. Continuously increasing	8. Intersection at halfway point
3. Level section at start		9. Maxima roughly equal
4. Slight increase at end		10. Minima roughly equal

Table 2: Key Features of Graph 1

Figure 2 shows a drawing of this graph by one of the participants. It received a score of 9/10, with the only key feature missing

being the level section at the start of line 1 (feature 3). Although the maxima are not exactly equal, they are reasonably close, and the intersection was considered to be sufficiently close to the halfway point.




Figure 2: A participant's drawing of Graph 1

Figure 3 shows another drawing of the same graph that received a score of 7/10. The participant mistook Line 1 for a straight line (feature 1), and omitted the level section at the start and the increase at the end (features 3 and 4).




Figure 3: Another participant's drawing of Graph 1

5. RESULTS

The average percentage of features drawn correctly was very high (over 80%) in both conditions (Figure 4). In addition, the average percentage of features drawn correctly in any individual graph was never below 68%. The most common errors in the drawings were: level sections too short, or omitted completely; features misplaced on x-axis; and features drawn at the wrong relative y-positions.




Figure 4: Average Percentage of Features drawn correctly in each condition

No significant difference was found in the overall number of features drawn correctly between conditions, or in the number of features of the individual lines, or features of the interaction between the lines, drawn correctly between conditions (Figure 4). However, a closer look at the errors in the drawings indicates that

there may be differences between the two conditions that would be revealed through running the experiment with a larger group of participants.

In the *different-instruments* condition, 11 intersection points that did not exist in the graphs were present in the drawings compared to just four in the *same-instruments* condition. This may be because participants were unsure when both instruments were playing the same pitch, and were therefore inclined to guess that intersections were present when the pitches sounded similar. Participants also drew the relative y-positions of global maxima and minima incorrectly more often in the *different-instruments* condition (28 times in *same-instruments* condition, 44 in *different-instruments* condition), again indicating that participants had difficulty comparing pitches when they were played by different instruments. While no significant difference has been shown, these results indicate that representing both data series with the same instrument might make it easier to identify features of the interaction between the two series (as suggested in Hypothesis 2).

Common errors in the *same-instruments* condition, which occurred less frequently in the *different-instruments* condition, were the misplacement of maximum and minimum points of the individual lines on the x-axis (27 times in *same-instruments* condition, 10 times in *different-instruments* condition), and the drawing of the start and end points of a line at the wrong equivalent y-positions (25 times in *same-instruments* condition, 8 times in *different-instruments* condition). These results indicate that it may be easier to distinguish the features of each line when different instruments are used (as suggested in Hypothesis 1).

The problem of pitch-grouping, as discussed in Section 2.2, appears to have only occurred twice in this experiment. On both occasions (one in each condition) the same participant became confused between the two lines after an intersection point. Figure 5 shows one of these graphs, while Figure 6 shows the participant's drawing of it. As this has only occurred in two out of 120 drawings, it does not appear to be a significant problem.




Figure 5: Graph 4—a graph used in the experiment




Figure 6: A participant's drawing of Graph 4

No statistically significant differences were found in the workload between conditions. However, four participants expressed a preference for the *different-instruments* condition and two expressed no preference. One participant who expressed no preference explained that, while she preferred the overall sound of the same instruments and found it easier to find intersection points in that condition, it was easier to distinguish the individual graphs when different instruments were used. These comments support both the hypotheses. A participant who preferred the *different-instruments* explained that it was sometimes hard to tell the two data series apart when they were both represented by the same instrument, as stated in Hypothesis 1. A larger scale study might reveal some improvement in the recognition of individual line features when different instruments are used.

6. CONCLUSIONS AND FUTURE WORK

The pilot study reported here has shown that it is possible for sighted people to draw sketches of graphs containing two data series with high accuracy after listening to sonified versions of the graphs. This shows that people are able to build up an accurate mental model of graphs while listening to them, and that testing with blind people would be worthwhile.

Although no significant differences were found, the study has indicated that using the same instrument for each data series might make it easier to identify interactions between the two lines (e.g., intersection points, relative y-positions of global maxima and minima). However, the majority of users expressed a preference for the representation of each data series with a different instrument, and there are indications that this may improve performance for identifying features of the individual lines. Further testing is required in order to investigate these differences in performance. Since both presentation modes have different advantages and disadvantages, graph sonification systems should allow users to choose between the two, depending on their preference and current task.

In future different instrument combinations should be tested, in order to find the optimum combination of instruments. In addition, the system should be tested with some real-world data sets, rather than with simple mathematical functions.

This research has shown that sonified graphs are successful for communicating information about basic graph shapes, and that these can be interpreted when two data series are presented simultaneously.

7. REFERENCES

- [1] D. L. Mansur, *Graphs In Sound: A Numerical Data Analysis Method for the Blind*, M.Sc. Thesis, Department of Computing Science, University of California, 1975
- [2] S. Sahyun, *A Comparison Of Auditory And Visual Graphs For Use In Physics And Mathematics*, Ph.D. Dissertation, Department of Physics, Oregon State University, Oregon, 2000
- [3] J. H. Flowers and T. A. Hauer, "Musical versus Visual Graphs: Cross-Modal Equivalence in Perception of Time Series Data", *Human Factors*, vol. 37, pp. 553-569, 1995.

- [4] L. Brown, S. Brewster, R. Ramloll, B. Riedel, and W. Yu, "Browsing Modes for Exploring Sonified Line Graphs", In *Vol. II Proceedings of HCI 2002*, pp. 6-9, 2002.
- [5] T. L. Bonebright, M. A. Nees, Connerley, T.T., and G. R. McCain, "Testing the Effectiveness of Sonified Graphs for Education: A Programmatic Research Project", In *Proceedings of ICAD 2001*, Espoo, Finland, pp. 62-66.
- [6] J. H. Flowers, D. C. Burham, and K. D. Turnage, "Cross-Modal Equivalence of Visual and Auditory Scatterplots for Exploring Bivariate Data Samples", *Human Factors*, vol. 39, pp. 341-351, 1997.
- [7] S. A. Brewster, P. C. Wright, and A. D. N. Edwards, "Parallel earcons: Reducing the length of audio messages", *International Journal of Human-Computer Studies*, vol. 43, pp. 153-175, 1995.
- [8] D. Deutsch, "Grouping Mechanisms in Music", in *The Psychology of Music, Cognition and Perception*, D. Deutsch, Ed. London: Academic Press Inc., pp. 99-134, 1982.
- [9] A. S. Bregman, *Auditory Scene Analysis: The Perceptual Organization of Sound*, 1st. ed: Massachusetts Institute of Technology Press, 1994.
- [10] S. A. Brewster, P. C. Wright, and A. D. N. Edwards, "Experimentally derived guidelines for the creation of earcons", In *Adjunct Proceedings of HCI'95*, Huddersfield, UK., pp. 155-159, 1995.
- [11] D. I. Rigas, *Guidelines for Auditory Interface Design: An Empirical Investigation*, Ph.D. Dissertation, Department of Computer Studies, Loughborough University, Loughborough, 1996
- [12] S. Hart and L. Staveland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research", in *Human Mental Workload*, P. Hancock and N. Meshkati, Eds. North Holland B.V., Amsterdam, pp. 139-183, 1988.