Presenting Information in Sound

Sara Bly
Lawrence Livermore National Laboratory
Livermore, California 94550

In the past few years, the increase in interactive use of computers has led to an emphasis on human factors and the ways in which digital information can best be presented to users. Computer graphics has been at the forefront of this growth involving vision as an active aid in interpreting data. Bar charts, psuedo-color image processing, and 3-dimensional figures are but a few means of providing the viewer with data information.

As the use of computers increases, the need for a variety of alternatives of interacting with computers also increases. Computer-generated sound is one capability not being fully utilized in the computer/human interface. Just as an x-y plot reveals relationships in data, sounds might also reveal relationships in data. This report focuses on the potential for using computer generated sounds to present data information. The first section addresses multivariate data problems which might be aided by sound output. The second describes experiments performed to determine whether listeners can discriminate among data sets based on sound. The final section discusses ongoing work and future directions.

PRESENTING INFORMATION

Vast quantities of computer output are useful only if presented in ways which can be understood and utilized by human analysts. In problems involving several variables, it is particularly difficult to visualize the relationships among data samples and the differences among data sets. Computer graphics has become a widely used and sophisticated technique for the output of these digital calculations[5]. However, in areas in which the number of varying parameters exceeds our visual response to color, rotation and dimension or in areas in which the data does not correspond to our familiar three-dimensional perception, graphical displays are inadequate. When traditional methods of graphical plotting are used to display the data, information is often lost because only a few dimensions may be presented.

Two particularly interesting graphical methods use multivariate data without restricting the

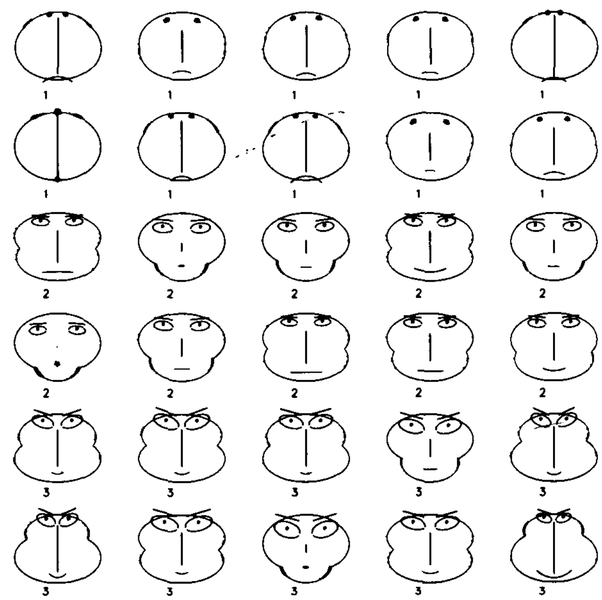

Figure 1: Chernoff's FACES

dimensionality. Andrews describes functions whose coefficients are the data variables of samples[1]. For data sets which are distinct, the resulting curves will cluster together accordingly. Chernoff depends on the ability of a human analyst to distinguish among faces[3]. The data sample variables are mapped to facial characteristics so that each data sample produces a corresponding face. Figure 1 plots ten data samples from each of three sets. The first two rows of faces are from set 1, the next two from set 2, and the last two from set 3. The differences in facial characteristics clearly separate the three sets of data.

Despite the success of graphical methods for presenting information, problems of adequately conveying multivariate data to a human analyst remain. In order to utilize graphical output, the human must first focus visual attention directly on the output, thereby restricting the information presented to that which is before the eyes. Secondly, graphics are generally limited to a finite number of dimensions, requiring that many multivariate data problems be reduced to fewer dimensions before analysis. Thirdly, when dealing with data about which very little is known, a variety of perceptions of the data is helpful. For many applications, current means of presenting information from the computer to the human are not sufficient.

Sound might improve these limitations by providing auditory cues of significant events when they occur (and thus easing the restriction of focused attention), by increasing the dimensionality of the graphics, and by presenting an alternative "view" of the data[13]. Computer generated sound is an available technology which offers a number of attributes for representing data dimensions. It is worth examining sound as a means of increasing the bandwidth of information from a computer to a human.

EXPERIMENT

To verify that sound does have potential for presenting data information to analysts, I ran several experiments. These experiments served two purposes; the first to determine that sound does convey accurate information about the data and the second to examine the potential of using sound to convey more information than the usual means of data analysis. The following sections describe one method of encoding data into sound, the methods for the experiments, and the results of each phase of those experiments.

Procedures

In order to present data information in sound, there must be a means of encoding data values in sound characteristics. Consider a mapping between a particular n-dimensional data sample and a discrete sound or note. Characteristics of a note include pitch, volume, duration, waveshape and envelope. Each of these may vary over a well-defined range. A single data sample can then be encoded into a single note by using the value of each variable to determine the level of a sound characteristic. Thus, each data sample produces a corresponding note which is determined by a mapping from data dimensions to sound characteristics.

For the experiments, I used six-dimensional data and mapped it to six characteristics of sound (pitch, volume, note duration, fundamental waveshape, attack envelope, and overtone waveshape). Pitch varied over the 48 notes in four octaves of a piano scale, volume varied from very soft to very loud in twelve increments, and duration varied from 50 msec to 1050 msec in increments of 5 msec. The waveshape of a fundamental frequency varied from a pure sine to a random buzz. Figure 2 shows four of the ten variations of waveshape used. Similarly, a fifth overtone varied in waveshape and was added to the fundamental. An attack envelope varied from a long, slow attack to a constant envelope.

Each trial of the experiment consisted of two sets of data which differed from one another in a well-defined manner. Fifty samples made up each set of data and each sample was six dimensional. The subjects were given no information about the ways in which the two sets differed. In the experiment, samples were presented to the subject as computer generated sounds. A subject first heard ten notes identified as being from set 1 and ten notes identified as being from set 2. These were considered the training samples. After repeating the notes as often as desired, a subject then entered the testing portion of the experiment. For a given trial, the goal of the experiment was for the subject to correctly identify those test samples which were in set 1 and those which were in set 2. That is, the subject heard a note and was asked to determine whether that note belonged to set 1 or to set 2. For each test sample, the subject's response was recorded. A subject's final score was the number of correct responses in determining whether test samples belonged in set 1 or set 2.

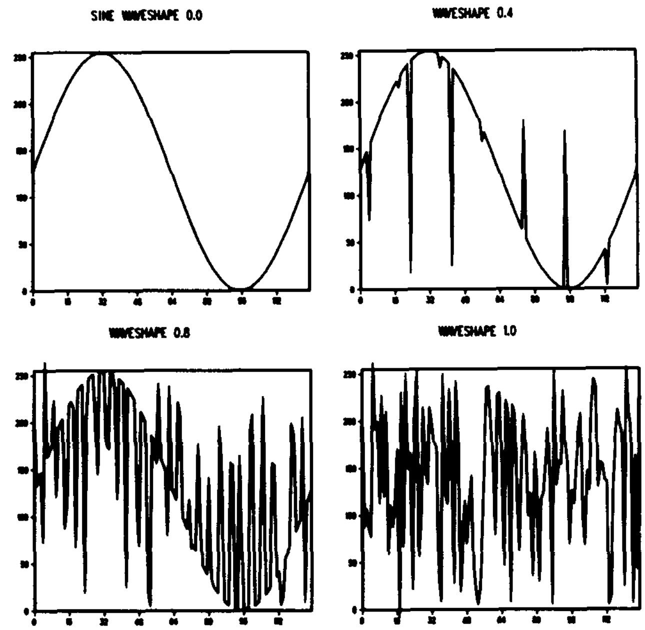

Figure 2: Waveshape Variations

Phase 1

To establish a basis for the fact that sound does convey data information, data bases were generated in which set 2 variables were transformed relative to set 1. Initially sets 1 and 2 were sets of six-dimensional random normal data. For one group of subjects, set 2 was translated relative to set 1. For another, set 2 was scaled relative to set 1 and, for the third, the variables in set 1 were strongly correlated relative to set 2. For simplicity, all variables were equivalent; that is, each dimension of a transformed set was translated, scaled, or correlated equally.

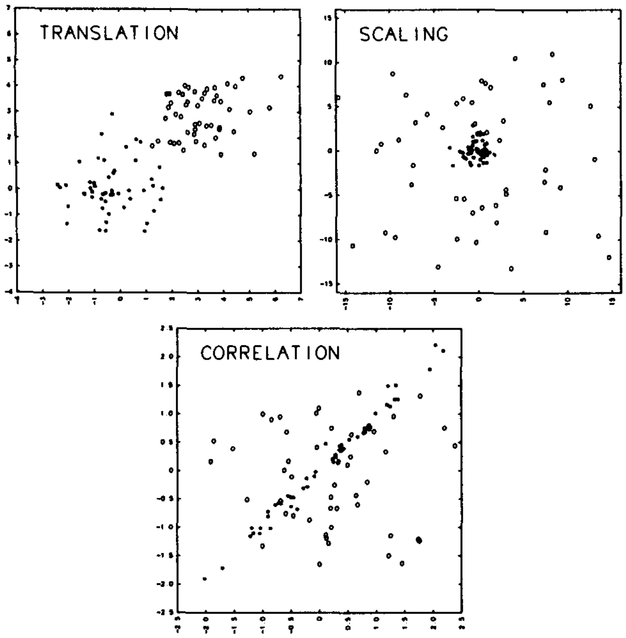

Figure 3: Phase 1 Data Sets

Figure 3 shows 2-dimensional graphical representations of three pairs of data sets. In the first plot, set 2 is translated in all six dimensions relative to set 1. In the second plot, set 2 is scaled in all six dimensions relative to set 1. In the third plot, set 2 variables are not correlated while set 1 variables are. Note that subjects did not see a display of the data but rather <u>heard</u> six-dimensional representations in sound.

Table 1 is a summary of the results for the experiments in determining whether the sound encoding of multivariate data would convey

	#subj	#test	%correct
Translation			
by 3	7	40	92%
by 1	7	40	70%
by 0.5	7	40	53%
Scaling			
by 8	7	40	69%
by 4	7	40	74%
by 2	2	40	55%
repeat by 8	5	40	76.5%
Correlation			
.99	4	60	60%

Table 1

consistent information.
The first group of seven subjects was first tested on data in which set 2 was translated 3 standard deviations from set 1. The subjects were then tested on data in which set 2 was translated 1 standard deviation from set 1 and finally on data in which set 2 was translated .5 standard deviations from set 1. Similarly, the second group of subjects was tested on data in which set 2 was scaled by 8 standard deviations and on data in which set 2 was scaled by 4 standard deviations. The third trial of the second group was data in which set 2 was scaled by 2 standard deviations for two of the subjects. The remaining 5 subjects were again tested on data in which set 2 was scaled by 8 standard deviations. The third group of subjects was tested on data in which the variables of set 1 were strongly correlated. As expected, generally the further separated the two sets, the better the subjects were able to distinguish between them.

<u>Phase 2</u>

During the second phase of the experiment, the intent was to discover if, for some data, sound could add more information than other methods alone. It was desirable to provide a data base whose characteristics were dependent on the multivariate nature of the data so that discrimination by usual methods was more difficult. One set of samples of 6-dimensional data was obtained from a multivariate normal random deviate generator. Samples were then separated into two sets such that a sample, s = (x1, x2, x3, x4, x5, x6), belonged to set 2 if and only if

$$x2^2 + x3^2 + x4^2 + x5^2 + x6^2 \le 1.5^2$$
or
$$x1^2 + x3^2 + x4^2 + x5^2 + x6^2 \le 1.5^2$$
or
$$x1^2 + x2^2 + x4^2 + x5^2 + x6^2 \le 1.5^2$$

Only samples in which all six variables had positive values were included. Thus, at least five of the six variables in each sample of set 2 had value less than 1.5 and at most one of the variables x1, x2, or x3 could have value greater than 1.5. In a complete positive and negative space, one can think about the data in set 2 being contained in three "cylinders", each of radius 1.5 about the first three axes. Fifty samples were used from each of the two sets. The experiment then consisted of three different means of presenting the data information — visually, aurally, and with a combination of visuals and sound.

For sound output, all six dimensions were used by mapping the variables to pitch, duration, volume, waveshape, envelope, and overtone. For graphical output, only two dimensions were used in an x-y plot. The points were not identified as belonging to set 1 or set 2. When a sample was used for training or testing, the point was highlighted on a display screen. Note that the intent was not to compare six-dimensional sound with two-dimensional graphics but to determine whether the sound could in fact increase the dimensionality of the graphics.

75 subjects participated, 25 in each of three groups. Each subject's trial consisted of 10 training samples randomly selected from set 1 and 10 from set 2. The subject then had 40 test samples to identify. The first group of 25 heard

each sample, the second group both heard and saw each sample, and the third group only saw each sample.

The average percentage of samples correctly identified was 62% for the graphics only presentation, 64.5% for the sound only presentation, and 69% for the presentation combining sound and graphics. Figure 4 is a plot of the raw data in which the subjects have been ordered from low score to high score within each

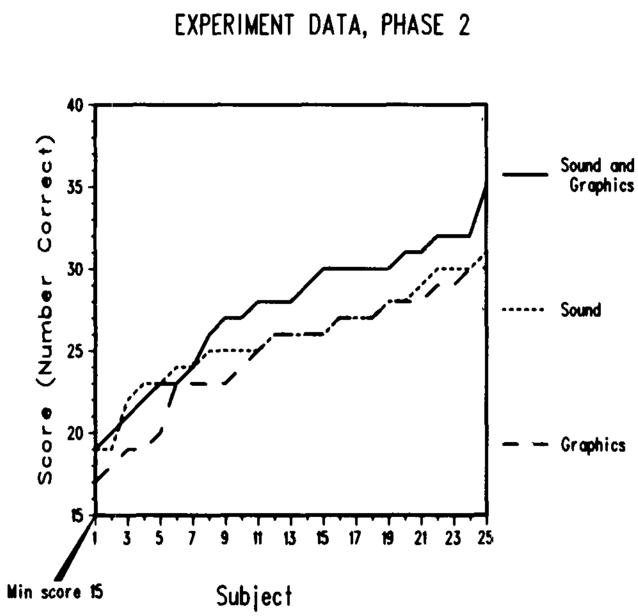

Figure 4: Phase 2 Results

group. Note that for those subjects correctly identifying at least 22 of the 40 test items, the group using both sound and graphics had consistently higher scores. The group with sound only performed as well as those with graphics output. These scores suggest that sound does add useful information to graphical output and that even sound alone is useful in discriminating between different sets of data.

Several months after the original experiment, 10 of the 25 subjects who participated in the sound-only trials were randomly selected to participate in a second experiment. The goal of this second experiment was to find out whether further training would improve a subject's performance. The subjects were given more experience using the facility with a variety of sound output applications. In addition, the training samples of the data were available to the subject throughout the experiment providing a reference at any time. The subjects stilll had no knowledge of the data base itself. These subjects who had further training correctly identified 74% of the test samples.

Overall the results are a positive indication that sound can indeed increase the information about multivariate data when it is presented to a human analyst. The responses from all groups were better than could be expected from guessing. The fact that further training improved the subject's ability to discriminate data samples based on sound alone suggests that the maximum performance level may not have been reached. It is also possible that improved methods for providing a reference for comparing sounds (much like the role of axes in an x-y plot) could increase the performance. Several areas of using sound synthesis for data presentation open themselves for futher exploration.

ON-GOING WORK

One implication of the results of the experiment is that sound not only conveys meaningful information about multivariate data but also has potential for increasing the information presented to a human analyst. It is possible in some cases that sound may yield more information to a human about multivariate data than graphical methods. However, the experiments made no attempt to compare the various methods of presenting multivariate data with graphics alone. Instead it seems most exciting to consider the potential of using both sound and graphics together. At this time, however, discovering more about the use of sound alone is critical to future studies involving both sound and graphics.

<u>Applications</u>

Since sound is not static but varies with time, it seems particularly appropriate to consider applications in which one of the variables is time. A real-time simulation is one such application currently being examined with sound output[2]. Simulations which run for several hours generate pages of statistics describing each time step of the model. In most cases, there are several variables involved and very different intermediate results may cause the same end result. By encoding the variables at each time step into the parameters of sounds, each run of the simulation then has a corresponding "song". Listening to the "songs" for several runs of a simulation can aid a human analyst in determining those stages of the simulation which are particularly meaningful or in which a significant change occured.

In addition to the multivariate applications described, at least two other areas offer potential for sound exploration. Since sound has the characteristic of not requiring a particular focus, auditory cues would be useful in drawing attention to specific events. Sound might also be appropriate for different types of data, such as logarithmic data, which is often difficult to judge visually.

<u>Sound Characteristics</u>

Although the applications for sound encoding are important, the characteristics of sound deserve further attention. First, sound parameters are not independent and thus affect the relationships among data variables. For example at constant volume, a high frequency note will sound softer than a low frequency note. Second, the timbre of a note is especially significant in determining the perception of that note. More work should be done in considering the possible uses and variations of timbre. Likewise, location or stereo effects of sound could add dimensionality. Third, it is not clear how many sound parameters can be distinguished and which of those parameters have the most meaning.

SUMMARY

All of the work thus far has indicated that sound can be a useful means of presenting information to a human user. As visual feedback from computers improves human use of digital calculations, it seems appropriate to examine other human means of obtaining information. Although experiments have been done with touch[6], most computer research involving senses other than vision has been with sound. Since sound technology already exists for mass reproduction and playback, it seems a natural choice for further study of its uses to convey digital information. The computer/human interface can only be enhanced by another means of presenting information.

ACKNOWLEDGEMENT

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-Eng-48.

REFERENCES

[1] Andrews, D.F. "Plots of High-dimensional Data", *Biometrics*, Vol. 28, p. 125, 1972.

[2] Bly, S.A. "Sound and Computer Information Presentation", PhD Dissertation (work in progress), University of California, Davis.

[3] Chernoff, H. "The Use of Faces to Represent Points in k-Dimensional Space Graphically", *Journal of the American Statistical Association*, Vol 361, p. 361, 1973.

[4] Chowning, J.M. "The Synthesis of Complex Audio Spectra by Means of Frequency Modulation", *Journal of the Audio Engineering Society*, Vol. 21, No. 7, p. 526, 1973.

[5] Everitt, B.S. *Graphical Techniques for Multivariate Data*, Heinemann Educational Books, London, 1978.

[6] Geyer, K.E. and Wilson, K.R. "Computing With Feeling", *Proc. of the Conference On Computer Graphics, Pattern Recognition, and Data Sructure*, IEEE Catalog No. 75CH0981-1C, p. 343, 1975.

[7] Mathews, M.V. *The Technology of Computer Music*, The M.I.T. Press, Cambridge, Mass. 1969.

[8] Moorer, J.A. "Signal Precessing Aspects of Computer Music--A Survey", *Computer Music Journal*, Vol. 1, No. 1, p. 4, 1977.

[9] Risset, J.C. "Synthesis of Sounds by Computer and Problems Concerning Timbre", *Music and Technology*, Proceeding of the Stockholm Meeting, 1970.

[10] Rosenstein, M. "Computer-Generated Pure Tones and Noise and Its Application to Psychoacoustics", *Journal of the Audio Engineering Society*, Vol. 21, No. 2, p. 121, 1973.

[11] Taylor, C.A. *The Physics of Musical Sounds* American Elsevier Publishing Company, Inc. New York, 1965.

[12] Tenney, J.C. "The Physical Correlates of Timbre", *Gravesaner Blatter*, Vol. 7, No. 26, p. 106, 1965.

[13] Yeung, E.S. "Pattern Recognition by Audio Representation of Multivariate Analytical Data", *Analytical Chemistry* Vol. 52, No. 7, p. 1120, 1980.