

HYPOTHESIS-DRIVEN SONIFICATION OF PROTEOMIC DATA DISTRIBUTIONS INDICATING NEURODEGRADATION IN AMYOTROPHIC LATERAL SCLEROSIS

William L. Martens

Faculty of Architecture, Design and Planning
University of Sydney, Sydney NSW 2006 Australia
william.martens@sydney.edu.au

Philip Poronnik

School of Medical Sciences
University of Sydney, Sydney NSW 2006 Australia
philip.poronnik@sydney.edu.au

Darren Saunders

School of Medical Sciences
University of New South Wales, Sydney NSW 2052
Australia
d.saunders@unsw.edu.au

ABSTRACT

Three alternative sonifications of proteomic data distributions were compared as a means to indicate the neuropathology associated with Amyotrophic Lateral Sclerosis (ALS) via auditory display (through exploration of the differentiation of induced pluripotent stem cell derived neurons). Pure visual displays of proteomic data often result in "visual overload" such that detailed or subtle data important to describe ALS neurodegradation may be glossed over, and so three competing approaches to the sonification of proteomic data were designed to capitalize upon human auditory capacities that complement the visual capacities engaged by more conventional graphic representations. The auditory displays resulting from hypothesis-driven design of three alternative sonifications were evaluated by naïve listeners, who were instructed to listen for differences between the sonifications produce from proteomic data associated with three different types of cells. One of the sonifications was based upon the hypothesis that auditory sensitivity to regularities and irregularities in spatio-temporal patterns in the data could be heard through spatial distribution of sonification components. The design of a second sonification was based upon the hypothesis that variation in timbral components might create a distinguishable sound for each of three types of cells. A third sonification was based upon the hypothesis that redundant variation in both spatial and timbral components would be even more powerful as a means for identifying spatio-temporal patterns in the dynamic, multidimensional data generated in current proteomic studies of ALS.

1. INTRODUCTION

This study investigated three alternative approaches to the sonification of proteomic data distributions as a means to indicate the neuropathology associated with ALS. A local group of researchers routinely generate large complex proteomic datasets obtained from patient-derived cell lines

and animal models in efforts to understand the changes in the ubiquitin-proteasome system during the progression of ALS. It is common to attempt to interpret these data with the aid of visual displays, using graphics such as that shown in Figure 1. However, these visual displays often provide an unwieldy summary of the structure of complex proteomic datasets, and so it was of great interest to determine if sonifications could provide an additional useful approach to the exploratory analysis required for this data, both as an accompaniment to visual display (as suggested in [1]), and as an independent means by which a stand-alone auditory display might become regarded as potentially useful in its own right (this is not a new idea, such proposals appearing in the early 1980's [2]).




Figure 1: An example of a visual display of proteomic data of the sort that has been utilized to aid in understanding the changes in the ubiquitin-proteasome exhibited in studies of proteomic ALS neurodegradation.

The over-arching assumption here is that auditory enhancement of visually represented data can significantly increase the ability of researchers to detect subtle changes or anomalies in such numerical data sets. Although the motivation for this research was to develop potentially useful auditory display tools for practical applications in the medical sciences in general, this paper reports only preliminary results of a single case study of proteomic data associated with neuropathology in ALS patients. A primary goal of this particular study was to determine which of three sonifications would be judged by domain experts to be most successful in identifying differences in spatio-temporal patterns within the multidimensional proteomic data generated in a single exploratory study. In addition to polling the opinions of domain experts, however, this preliminary study assessed, using in two psychometric tasks, the success of three sonifications in aiding naïve listeners to identify differences between proteomic data distribution of three types of cells. The first task required naïve listeners to make dissimilarity judgments for all pairwise comparison of the set of nine cases defined by the factorial combination of three cell types and three sonification methods. The second task required those same naïve listeners to make ratings of each of the nine stimuli on a number of subjective attributes that might be related to the perceptual dimensions underlying their dissimilarity judgments.

Although these results reveal only the perceptual distinctiveness of the sonifications, they provide a basis for further exploration of the potential value of these sonifications, with the benefit that their perceptual distinctiveness has been established. Clearly, further work will be required to address important issues in sonification system usability for this application, and to determine whether the system will provide a real benefit to domain experts. Indeed, this paper reports preliminary evaluation results as an indication of progress on just one component of the larger medical research project in that it focuses only upon the effectiveness of the display technology used to aid medical scientists in interpreting and understanding their medical data (particularly proteomic data). Nonetheless, it is reasonable to assume that such preliminary examination of the perceptual distinctiveness of sonification system outputs will be a valuable first step in developing and exploring novel sonifications.

2. METHODS

Three alternative proteomic data sonifications were developed and compared in terms of their potential to communicate to the user changes due to neuropathology associated with ALS. Without any prior experience with auditory display of such data, an approach was adopted that is here termed ‘hypothesis-driven’ design of these sonifications, with the expectation that the relative value of the sonifications could be evaluated by the domain experts, who would try to detect differences between the results for ALS cases versus control cases. Of course, ground truth was already available for these particular data, since the ALS cases were selected on the basis of established medical diagnostic procedures. Therefore, relative performance under blind testing conditions has been used to determine which sonifications were best able to communicate to the user regarding neuropathology associated with ALS via changes in sonifications that were directly driven by differences in distribution of the proteomic datasets.

One hypothesis to be tested was that variation in predominantly timbral attributes would be most effective in revealing differences in proteomic data distributions. An alternative hypothesis was that variation in spatial timbral attributes would be more effective in creating audible differences between the sonifications produced for each of the three cell types. Finally, a third hypothesis was that including redundant variation in both timbral and spatial attributes would be more effective than just one or the other of these two individual approaches in isolation. Preliminary results suggested that comparison of these three sonifications, based as they were upon this ‘hypothesis-driven’ design, should allow for the rejection of a hypothesis that had resulted in less effective sonifications, leaving for future consideration only those hypotheses that were not rejected through blind testing. This scientific approach to the initial evaluation of competing sonifications will be examined in more detail in subsequent sections of this paper; however, before describing further this study’s experimental design, the auditory display technology underlying the alternative sonifications will be presented.

2.1. Sound synthesis for the sonifications

In order to generate a sonification for the available proteomic data of interest, a strategy for synthesis that took into account the complexity of the large multivariate dataset was formulated based upon parameter mapping [3]. For nine distinct cases, an assembly of short-duration, temporally overlapping ‘grains’ of sound were created, the parameters of which were selected to approach approximately the minimum perceivable event time for distinct percepts of duration, frequency and amplitude (i.e., approaching auditory resolution of human observers in discriminating between identifiable attributes of loudness, pitch, and those component auditory attributes that are generally regarded as belonging to one of two collections termed timbral or spatial attributes). The ‘hypothesis-driven’ design approach taken here required sound synthesis technology that could offer independent variation of many synthesis parameters to provide identifiable variation in distinct auditory attributes. In the initial stage of this work, synthesis based upon a simple physical model [4] was tested for its versatility in producing a wide range of short sounds exhibiting audibly identifiable timbral variations that all had potential for evoking physical referents in the minds of the listeners. In the next section, the spatial positioning of grains is explained.

The synthesis technology that ultimately was adopted for this project resembles granular synthesis (see [5]), in that a multitude of short sound sources formed an ensemble output (likened to a ‘swarm’) rather than forming clearly separable events that might be heard as distinct in time and space. In all sonifications designed for the current work in this way, there was always a hypothesis to be tested regarding which parameters of the data were ‘mapped’ to particular synthesis parameters. It is beyond the scope of this paper to present the details of the synthesis technology that was developed and refined through experimentation with the available multidimensional proteomic data. Suffice it to say that swarms of percussive ‘grains’ (again, see [5]), synthesized with ‘parameter-mapped’ control over multiple timbral attributes, were distributed in time and space according to the distribution of proteomic data that featured 1815 variables observed over the nine cases to be examined.

2.2. Spatial sound processing for headphone display

Although discrimination of frontward from rearward incidence of sonification components could be well supported if binaural processing were to be coupled with head-tracking technology (see [6]), the experimental stimuli generated in the current investigation were not modified by active sensing of the listener's head turning. Without such tracking of head movements, the sonification designer should not expect the listener to be able to clearly identify whether a source presented at a given lateral angle is being presented with frontward or rearward incidence. Due to the difficulty in supporting reliable front/rear distinctions using uncoupled binaural processing for headphone-based spatial auditory display (again, see [6]), only a simplified model of head acoustics was employed here to move sonification components along the listener's interaural axis. The acoustical cues that were simulated in order to accomplish this manipulation of sound source incidence angle included the interaural time delay (ITD) and the head shadow that generally grows larger at the listener's contralateral ear as the incidence angle of the source is offset laterally from the listener's median plane. This approach offered an advantage over a single-user headphone display in that several listeners could use the system simultaneously without the unexpected variation that would occur if the spatial processing were coupled with head movement of just one of multiple listeners. Of course, using head-coupled updating of headphone-based binaural rendering technology could be added for single-user exploration of the spatial configuration of sonification components (including sensitivity to the listener's translation movements as well as changes in head orientation); however, for the initial studies reported here, only non-head-tracking headphone technology was employed.

2.3. Spatial versus timbral emphasis in sonification

There were nine sonifications created from the factorial combination of three synthesis solutions¹ applied to data from three cell types. So for each of three types of cells that should produce an identifiably different sound, each of three unique parameter-mapping synthesis solutions were applied for presentation. The first of these synthesis solutions was termed the 'Timbral-only' approach, which put emphasis upon timbral differences resulting from spectral variation between grains. The second approach was termed the 'Spatial-only' approach, which held grain spectra constant, and only distributed the grains spatially along the listener's interaural axis. The third approach was termed the 'Spatial-Timbral' approach, and combined redundant variation in the output sound based upon the simultaneous application of both of these parameter-mapping approaches. These sonifications were chosen as candidates for best allowing the differences between cells to be appreciated by any observer, not just those with domain knowledge.

¹ The synthesis solutions employed here were all programmed within the Matlab™ environment. Although the details of the synthesis approaches taken would no doubt be of interest to a subset of readers, those details are considered to be beyond the scope of this paper. The code itself provides the most enabling description of the synthesis approaches. In order to enable interested researchers in replicating the approach taken in the project described in this paper, the employed Matlab code will be provided online (please send an email request to the first author for the URL).

2.4. Experimental Tasks

While discrimination between sonified cases was examined in pilot tests that were run informally during development of competing sonifications, the formal study that allowed more comprehensive analysis of similarities and differences between sonifications comprised two tasks. The first task was a pairwise dissimilarity-rating task, in which the global differences between nine sonification outputs were examined, without respect to particular identifiable attributes. The second task was an attribute-rating task, in which the particular character of each of the nine sonification outputs were examined with respect to identifiable attributes that were exemplified by anchoring stimuli found to be positioned at the extremes of each continuum for those attributes that seemed most distinctly varying within the set of nine stimuli. In fact, for the initial exploration of the characteristics of the nine stimuli, only the sonification developer engaged in the selection of adjectives describing the stimuli through informal discrimination tasks, and so no profiling of the stimuli was done by the five listeners who were naïve regarding the purpose of the experiment. The two formal tasks were completed by these five naïve listeners, but the adjectives used to describe the attribute rating scales were only introduced after the completion of the pairwise dissimilarity-rating task, in order to avoid drawing attention to the experimenter-identified attributes. The instructions for the first task indicated to the listeners that global dissimilarity ratings were required, rather than differences between sonifications based upon particular auditory attributes).

All pairs of nine sonifications were presented to five listeners for their evaluation via Sennheiser HD600 headphones at a comfortable listening level (approximately 75 dBA). Each listener completed one block of 72 trials, which is the number of paired comparisons resulting from the exclusion of the diagonal entries of the 9 x 9 matrix of dissimilarities (i.e., excluding all comparisons between identical stimuli). The sonifications in each pair were presented twice, in two separate trials, with order of presentation reversed for the second presentation, and always separated by a 1-s delay. For each pair of sonifications, listeners recorded their inter-stimulus dissimilarity ratings using a horizontal slider incorporated into an onscreen Graphical User Interface (GUI). On-screen instructions prompted listeners to indicate how similar they thought the sonifications sounded, with the leftmost response indicating that the sonifications sounded most similar, and the rightmost response indicating that the sonifications sounded maximally dissimilar. Each listener had to develop his or her own criterion for the anchoring point of maximal dissimilarity during an initial practice run in which 12 representative pairwise comparison trials were completed. After the initial practice run of 12 trials, each listener completed the formal run of 72 trials. The dissimilarity data matrices produced by each listener in these 72-trial runs could have been averaged to produce a single dissimilarity data matrix for group analysis, however a more powerful analysis using Individual Differences SCALing (INDSCAL) was employed to examine how the five listeners differed from each other, in addition to the summary that is available via examination of the group result.

The combined collected dissimilarity ratings from the group of five listeners were submitted to INDSCAL to obtain two useful outputs: First INDSCAL produced a two-dimensional (2D) spatial configuration of cases (a group 'Stimulus Space'

derived for five listeners taken together) in which each sonification was given coordinates along two dimensions so that the Euclidean distances between the points corresponded to the dissimilarity ratings. The INDSCAL analysis also produced estimates of the differences in weighting that each of the five listeners placed on the resulting dimensions (which weightings are captured by INDSCAL in terms of a 'Subject Space'). Further details of the analysis are given in the next section of this paper (see the book on *Modern multidimensional scaling* by Borg and Groenen [7] for a more complete explanation of INDSCAL analysis).

The instructions for the second task indicated to the listeners that attribute ratings were required, based upon than differences between sonifications that could be identified with particular auditory attributes. These attribute scales were anchored by adjectives that had been selected by the experimenter to represent the most distinct differences within the set of nine stimuli that seemed likely to be understood by the naïve listeners without much explanation. The selected anchors included the following pairs of adjectives:

Sparse ↔ Dense
Tense ↔ Relaxed
Smooth ↔ Rough
Compact ↔ Scattered
Simple ↔ Complex

3. RESULTS

The results of the INDSCAL analysis of the obtained data are shown in Figures 2 and 3. The Stimulus Space shown in Figure 2 uses plotting symbols that indicate the type of cell ('Control', 'ALS', and 'Fibroblast') for which each sonification was generated, as indicated in the legend located in the upper right corner of the graph. Line segments connect the plotting symbols in order to group together the results for the three cell types that were associated with each type of sonification that was employed for the group. The interpretation of this graph may not be obvious at first glance, but it is actually quite straightforward: The three groups of connected symbols will be plotted close to one another if the perceived differences between them is relatively small. For example, the smallest cluster of symbols that are grouped near the origin of the graph (i.e., the [0,0] point) are associated with 'Spatial-only' sonifications that were heard to be more similar to each other than those associated with the other two groups of sonifications. The value of this plot is that the relative distance between plotting symbols can be interpreted as providing a uniform indication of both within-group differences and between-group differences (the term 'uniform' is used here to indicate that all distances here are based upon a common Euclidean scale). Yet it remains to be asked, what can be concluded from such results. The primary conclusion would be that the sonification type used for the group of cells associated with 'Spatial-Timbral' sonifications are showing the greatest inter-stimulus distance of all three groups, and therefore this sonification solution would be preferred according to the criterion that these cells should produce sonifications that are perceptually different as possible.




Figure 2: Stimulus Space resulting from the INDSCAL analysis of nine sonifications. Line segments connect the three sets of points associated with each sonification approach, and symbol shapes indicate the cell type being sonified (as identified in the inset legend).

As mentioned in the methods section of this paper, INDSCAL analysis also produces estimates of the differences in weighting that each listener places on the resulting dimensions. This INDSCAL-derived 'Subject Space' is shown in Figure 3 for the group of five listeners who participated in this exploratory study. Although the obtained dissimilarity-rating data requires different weights on the two dimensions of the group 'Stimulus Space' for each listener, the advantages of INDSCAL is that these differences make it possible to separate such individual differences from the group solution, which shows the common underlying configuration that fits best to all the data. Note that three of the listeners put roughly equal weights upon the two 'Stimulus Space' dimensions, indicated by vectors drawn at around 45° from the origin of the graph in Figure 3. One listener put slightly more weight upon Dimension 2, while the remaining listener put more weight upon Dimension 1. Nonetheless, the results are consistent with the hypothesis that for the nine sonifications presented, the five listeners share a common underlying perceptual space that admits of two salient dimensions (although conjecture about the existence of a third underlying dimension might be tempting to consider, comparisons between just nine stimuli do not provide an adequate basis for supported such a conclusion, as explained in the Borg and Groenen [7] book).

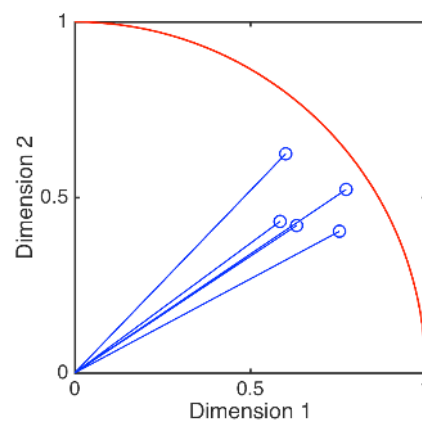


Figure 3: Subject Space resulting from the INDSCAL analysis of nine sonifications for five listeners.

Having concluded that the nine sonifications can be configured within a common perceptual space for the five listeners, it remains to be determined how those two salient dimensions might best be interpreted. One answer to this question would be to find out whether INDSCAL-derived Stimulus Space coordinates could be related to ratings of those stimuli along identifiable auditory attribute scales anchored by the adjective pairs selected by the experimenter. Therefore, the attribute ratings made by those same five listeners, when presented with those nine stimuli individually on separate trials, were submitted as competing predictor variables in a stepwise regression analysis. Although the correlations between sets of ratings could be fairly high, such as that between the smooth-rough and the tense-relaxed rating data, the stepwise regression analysis showed a single set of ratings as the best predictor for INDSCAL Dimension 2, and that was the set associated with the compact-scattered anchoring adjectives (with $R^2=0.66$). After excluding the compact-scattered set from consideration for interpreting INDSCAL Dimension 1, no one of the four remaining predictors showed a particularly high correlation with the coordinates of the nine stimuli on Dimension 1. However, when the smooth-rough and the tense-relaxed rating data were combined to form a new composite predictor, that new predictor accounted for a more of the variance in Dimension 1 coordinates (with $R^2=0.45$).

Taken together, the results of the two tasks serve to show how big the differences were between nine sonifications, and also suggest how one might describe the nature of those differences. The differences between the outputs of the three sonification techniques were best described as varying along a compact-scattered dimension (in the vertical direction of the graph). Clear differences also existed in the configurations derived for the sonifications of the three cell types, which differences were associated primarily with variation in both the smooth-rough and the tense-relaxed ratings.

4. DISCUSSION AND CONCLUSION

The results obtained in this exploratory study only scratch the surface of the problems that must be addressed in developing and evaluating sonifications in this domain of complex, multidimensional proteomic data generated by research studies in medical science. Of the three types of sonifications presented, it seems that sonifications mapping from data to both timbral and spatial parameters provide more distinguishable results than mapping to either timbral or spatial alone, although these results must be regarded as quite preliminary. Nonetheless, the results seem quite promising when compared to the results of typical attempts to visualize such data. One such attempt utilizes multivariate analysis to reduce the complexity of the data to a more easily digestible form. Of course, similar data reduction procedures can be used as a pre-processing step for sonification as well (see [8]). What most such analyses typically attempt to do is to capitalize upon redundancy in the data to find a lower-dimensional perspective on the patterns of underlying variation.

The fact that the 1815 variables are somewhat correlated with each other means that a good deal of the variance in the data is shared, and that shared variance might be represented by a projection of the cases onto a single axis or two through the 1815-dimensional space defined by the proteomic

variables. The most common multivariate analytic technique that seeks out such a projection is Principal Component Analysis (aka PCA). PCA effectively rotates the axes in a multivariate space to find the principal axis along which the variance in the dataset is maximized, taking advantage of the covariance between all the variables. The analysis also finds a second axis, orthogonal to the first, that accounts for the greatest proportion of the remaining variance (see [9]). Figure 4 shows the scores on the principal components resulting subspace projection for the nine cases that were examined in the current study. While the simplicity of the graph in Figure 4 suggests that a simple difference might exist between the three groups of three items here, there is no way of learning from the graph what the meaning of the underlying components might be. Nonetheless, the PCA does provide a potentially more satisfying look at what is going on in the data, even though this involves a somewhat unwieldy graphical analysis of the weights involved in constructing the linear combinations on which the scores shown in Figure 4 are based.




Figure 4: Principal Component (PC) Scores resulting from the multivariate analysis of the proteomic data that featured 1815 variables observed over nine cases.

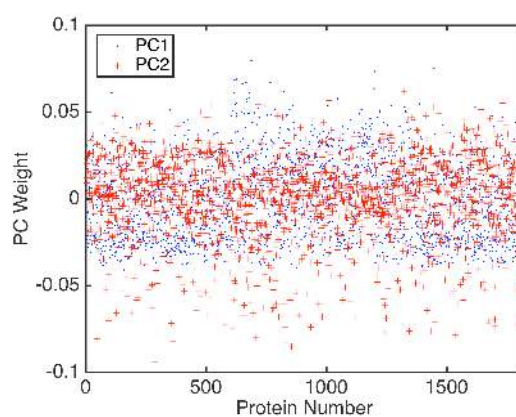


Figure 5: Weights placed upon the 1815 variables that resulted from the Principal Component Analysis (PCA) of the proteomic data.

The weights that were placed upon each of the 1815 variables are illustrated in Figure 5. It is difficult to imagine seeing a pattern here, but it is not so difficult to imagine hearing a change in the underlying pattern of sound ‘grains’ that might be generated through the spatiotemporal

distribution of those ‘grains’ in an appropriately constructed sonification of these data. It is precisely the expectation that such patterns might be appreciated by ear more readily than by eye that motivated the current work.

It is instructive to compare the current results with those of other developers of sonification systems that used similar granular synthesis techniques. An example of early work using granular synthesis, and that also was notable in that it shared a similar psychometric testing perspective, was that of Smith, et al [10]. Their approach was to use a three-alternative forced choice (3AFC) task to track a user’s ability to discriminate between ‘clouds’ of granules varying in frequency with controlled mean and standard deviation within two clouds being compared. Thresholds for hearing a comparison cloud to be ‘higher’ than a reference cloud were measured in the presence of distracting frequency-dependent amplitude modulation resulting from granules closely adjacent in frequency. In presenting their seminal work, Smith, et al [10] identified the types of perceptual discrimination that might be studied as either detection, recognition, or discrimination. The current work extends their psychometric testing approach to include global-dissimilarity-based perceptual scaling with attribute scaling to aid in the interpretation of underlying dimensions. The current results offer the advantage of determining which of a number of competing sonifications created for a small group of listeners the greatest overall perceptual differences between cases for which discrimination is desired (i.e., between cell types). Furthermore, the magnitude of the perceptual differences so observed could be compared to the magnitude of the perceptual differences existing between the multiple outputs of the competing sonification techniques, since these both types of differences were scaled in a common multidimensional space.

In a 1999 review paper, Barrass and Kramer [11] have provided a comprehensive survey of approaches for designing sonifications, and also have outlined ongoing concerns with the existing sonification practice. Of particular relevance is their discussion regarding how knowledge about auditory perception can allow sonification designers to predict how listeners will be able to perceive (if not understand and interpret) variations in novel sonifications. As in the current study, such knowledge can be derived for competing sets of sonifications, however, the point made by Barrass and Kramer [11] is well taken that the theoretical evaluation of new, untried designs requires more than psychoacoustic data. This is because psychoacoustic theories do not involve issues of representation that are central in sonification, since listeners need to hear the underlying data relations in the sounds, rather than just the auditory attributes that are modulated by them.

Thus, the current work in evaluating ALS-related proteomic data sonifications must be regarded as work that is still in the early exploratory stage. As the work enters into a second more confirmatory stage, it will become critically important to establish the means whereby progress and success can be ascertained. Therefore, in a manner that was thoroughly discussed in Bonebright and Flowers’ [12] chapter in the *Sonification Handbook* on ‘Evaluation of Auditory Display’ the initial (and ongoing) evaluation of the current sonifications has been focused upon whether the auditory distinctions displayed in each of the three case studies are in fact as audible and intelligible as the sonification system

developer has designed them to be. The evaluation methods that were used in this regard are those borrowed from perceptual science, and include psychophysical testing for detection, discrimination, and identification of displayed auditory attributes. In addition, the perceptual distinctions displayed in the sonifications presented in the current study, which were intended to distinguish differences that exist in the proteomic data, were assumed to grow larger in perceptual magnitude as differences in the data grew larger. This assumption was directly tested in the experiments reported here, in that known differences in the data (according to the medical science) were used to predict the heard differences reported by observers based strictly upon blind pairwise comparisons between sounds (i.e., the estimated perceptual differences/similarities formed by listeners blindly, on the basis of the auditory display alone). Other simple psychophysical tests involving pairwise discrimination in terms of identifiable attributes can certainly be considered, but have not been executed as yet.

Beyond these more elementary psychological measurement techniques, future development of the sonification systems under test will employ a broad range of evaluation methods, which have been chosen to address the most important issues in sonification system usability. In the final analysis, however, the completed sonification system must meet explicit acceptance criteria before its success is demonstrated. As outlined by Schneiderman and Plaisant [13], these criteria for evaluating system performance might include the following:

- Time for users to learn specific functions
- Speed of task performance
- Rate of errors by users
- User retention of commands over time
- Subjective user satisfaction

In addition to the overall satisfaction with the displayed sonification that may be expressed by system users with domain knowledge, which satisfaction may diminish with time, a more objective evaluation is to be recommended. It is not enough that users think that they can use a system effectively; rather, it is important to determine whether users can reliably make accurate judgments about the information being displayed as part of a typical use-case analysis. Thus, otherwise satisfying sonifications, which are nonetheless finding no support from the results of double blind testing, will eventually be rejected. Ultimately, it is hoped that such an approach will contribute to the formulation of a more general theory of sonification. Empirical results such as these might allow a sonification theory to evolve through a somewhat natural ‘winnowing out’ of unsuccessful approaches, supporting a general approach to sonification with the potential to fill ‘ecological niches’ with truly winning applications.

5. REFERENCES

- [1] C. Scaletti, and A. B. Craig, “Using sound to extract meaning from complex data,” *Electronic Imaging ’91*, San Jose, CA. International Society for Optics and Photonics, 1991.

- [2] J. J., Mezrich, S. Frysinger, and R. Slivjanovski. "Dynamic representation of multivariate time series data," *Journal of the American Statistical Association*, 79(385), pp. 34-40, 1984.
- [3] F. Grond, and J. Berger, "Parameter Mapping Sonification," In Hermann, T., Hunt, A., Neuhoff, J. G., editors, *The Sonification Handbook*, Chapter 15, pp. 363–397. Logos Publishing House, Berlin, Germany, 2011.
- [4] K. Karplus, and A. Strong. "Digital synthesis of plucked-string and drum timbres," *Computer Music Journal*, 7(2), pp. 43-55, 1983
- [5] P. Dutilleux, G. De Poli, A. von dem Knesebeck, U. Zölzer, "Time-segment processing," In *DAFX - Digital Audio Effects* (2nd ed.). John Wiley & Sons Ltd., 2011.
- [6] W. L. Martens, "Perceptual evaluation of filters controlling source direction: Customized and generalized HRTFs for binaural synthesis," *Acoustical Science and Technology*. 24(5), pp. 220-232, 2003.
- [7] Borg, I., and Groenen, P. *Modern multidimensional scaling: Theory and applications* (2nd ed.). New York, NY: Springer, 2005.
- [8] S. Ferguson, W. L. Martens, and D. A. Cabrera, "Statistical Sonification for Exploratory Data Analysis," In Hermann, T., Hunt, A., Neuhoff, J. G., editors, *The Sonification Handbook*, Chapter 8, pp. 175-196. Logos Publishing House, Berlin, Germany, 2011.
- [9] G. H. Dunteman. *Principal components analysis*. SAGE Publications, Inc, Thousand Oaks, CA, USA, 1989.
- [10] S., Smith, H., Levkowitz, R. M., Pickett, and M., Torpey, "System for psychometric testing of auditory representations of scientific data," in *Proceedings of the International Conference on Auditory Display ICAD '94*, Santa Fe, New Mexico. 7–9 Nov, 1994.
- [11] S. Barrass, and G. Kramer, "Using sonification," *Multimedia Systems*, 7(1), 23-31, 1999.
- [12] T. L. Bonebright, and J. H. Flowers, "Evaluation of auditory display," In Hermann, T., Hunt, A., Neuhoff, J. G., editors, *The Sonification Handbook*, Chapter 6, pp. 111–144. Logos Publishing House, Berlin, Germany, 2011.
- [13] B. Schneiderman, and C. Plaisant, *Designing the user interface: Strategies for effective human computer interaction* (5th ed.), Boston, MA: Addison-Wesley, 2010.