

EXPLORING SONIFICATION FOR AUGMENTING BRAIN SCAN DATA

Agnieszka Rogińska

Music and Audio Research Lab
Music Technology
New York University
35 West 4th St, New York, NY, USA
roginska@nyu.edu

Kent Friedman

Division of Nuclear Medicine
Department of Radiology
NYU School of Medicine
560 1st Ave, New York, NY, USA
kent.friedman@nyumc.org

Hariharan Mohanraj

Music and Audio Research Lab
Music Technology
New York University
35 West 4th St, New York, NY, USA
hm992@nyu.edu

ABSTRACT

Medical image data has traditionally been analyzed using visual displays and statistical analysis methods. Visual representations are limited due to the nature of the display and the number of dimensions that can be represented visually. This paper describes the use of sonification to represent medical image data of brain scans of patients with Alzheimer's dementia. The use of sonification is described as an approach to augment traditional diagnosis methods used to diagnose Alzheimer's dementia.

1. INTRODUCTION

A variety of methods and human sensory systems are put to use by the physician when examining a patient (verbal communication, visual examination, listening and touching a patient). In the recent past, the field of medical imaging has revolutionized diagnosis and management of disease. A variety of techniques (x-rays, mechanical vibration, fluorescence, rotation of atoms and radioactive decay) have been developed and produce multi-dimensional and time varying arrays of spatial information. Despite these advances, there are many limitations in the ability to perceive and analyze the vast amounts of data being generated by CT, PET scanners, MRI machines and newer combined devices (including PET/CT and PET/MRI). Diagnostic accuracy is higher than ever, but clinicians still are unable to detect certain conditions when the information provided by visual analysis or basic quantization does not uncover perceptible differences between disease and health. Displaying 3D data on visual displays is particularly challenging because of the many dimensions needed to be presented concurrently. Our omni-directional sense of hearing sense, however, allows us to hear sounds emanating from around us.

2. BACKGROUND

Across a broad range of scientific fields, sonifications and the auditory processing of image data identify patterns that are not detected by traditional means. Over the past 20 years, there have been many explorations and examples of sonifications. Sonifications of heart rate variability have been shown to have

diagnostic potential [3]. There is a strong history of sonifying EEG readings of brain activity (e.g. [1][2] [7]). Other relevant work includes EMG [15] and cell culture data [5]. These sonifications have subsequently proved useful as general introductions to physiological health, making distinctions between healthy and diseased states easy for uneducated listeners [8]. Sound has also been effective in combination with haptics as a surgical aid or training tool [9][12], and sound has functioned as an effective component in assistive rehabilitation therapy for stroke victims [16]. There has also been work done in sound rendering of tissue biopsy slides for diagnostic purposes [4].

The continuum line of the role of medical data sonification (Figure 1) ranges from aesthetics to diagnostics. By adding sound to augment visual medical data analysis, sonification's role can be to enhance the data analysis experience by providing generative music as a form of background sound. Sonification can be used to provide sound that mimics the visual data without necessarily adding any benefit in the data evaluation. The next step is to use sonification to augment data perception and evaluation, including reinforcing the data perception by adding data representation using a second sense (hearing), provide a different perspective using 3D sound, or place the observer in a fully immersive sonic environment that they can walk through or interact with, using immersive sonification. At the other end of the continuum, is sonification's use for diagnostics, where patterns and biological processes undetected using traditional visual and statistical analysis could be detected through sonification in order to diagnose disease or even predict or monitor response to therapy.

We refer to immersive sonification as a method of sonifying data in such a way as to place the viewer/listener of this data inside of an immersive 3D environment (using 3D sound) so that they may be able to navigate through the environment.

3. CHALLENGES OF IMAGING

A broad category of diagnostic testing, which has revolutionized diagnosis and management of disease, is medical imaging. A variety of techniques based on x-rays, mechanical vibration, fluorescence, rotation of atoms and radioactive decay

have been developed and produce multi-dimensional and time-varying arrays of spatial information.

Current methodologies used to "perceive" and interpret medical image data are largely based on the human visual system, in some cases enhanced by use of simple graphs or tables depicting numeric data. Such visual and traditional quantitative analysis methods have led to great advancements in nuclear medicine, radiology and other fields.

Despite these innovations, many limitations remain with respect to the medical community's ability to perceive and analyze the vast amounts of data now being generated by CT scanners, PET scanners, MRI machines and newer combined devices (including PET/CT and PET/MRI). Diagnostic accuracy is higher than ever, but clinicians still are unable to detect certain conditions when the information provided by visual analysis or basic quantization does not uncover perceptible differences between disease and health.

There are two possible causes for the existing limitations in the diagnosis of disease using medical image data. One possibility is simply that the techniques in use do not provide enough information to allow absolute differentiation between disease and health, no matter how sophisticated our methods of analysis might be or might become. Another possibility is that the information needed to diagnose disease is "hiding" within these images and is not perceptible because the means by which we have chosen to examine the data are insufficient for detection of complex patterns associated with disease. This second (and far more likely) possibility suggests that the medical community needs to find new ways to process and understand the data that is being acquired by advanced scanners and other testing equipment.

The concept of searching for new methods to understand complex medical image data is not without precedent. In recent years, researchers performing brain scans on patients found that basic visual analysis of the distribution of information was at times insufficient to detect subtle diseases (epilepsy being one example). In order to overcome the limitations of visual interpretation, new methodologies were developed whereby brain image data was segmented and compared to a database containing information about normal individuals. Such techniques have proven successful and provide useful complementary information when combined with visual inspection. [6][11][10].

Quantitative analysis techniques for examining medical image data represent a significant step forward beyond the traditional visual processing system of the human brain. However, they raise the question of whether or not even newer and perhaps less conventional methods for image data analysis might be able to improve our ability to diagnosis and monitor disease even further.

Advanced computer processing techniques generating increasingly complex visual reports may allow doctors to find new ways to detect disease. Multi-parametric predictive modeling [13] and machine learning [14] are examples of two evolving areas in medical image analysis.

When considering possible diagnostic techniques of the future, it is important to look back to the history of medicine and remember that visual analysis has not been the only way for a human to detect disease. It is worth exploring other ways to perceive patterns buried deep within complex data.

One potential alternate method for interrogating medical image data is by means of translation of information in to a format that can be processed by an alternate pathway: the human auditory system. Already regarded by experts as superior to vision in the domains of frequency and time, the human ear and associated auditory cortices present a compelling alternative system for perception of medical scan data. Independent of considerations of resolution, the complex neurological pathways of hearing offer a new perspective for understanding spatial, frequency and temporal data. Fortunately, a nascent field exists to allow for such a line of inquiry: sonification.

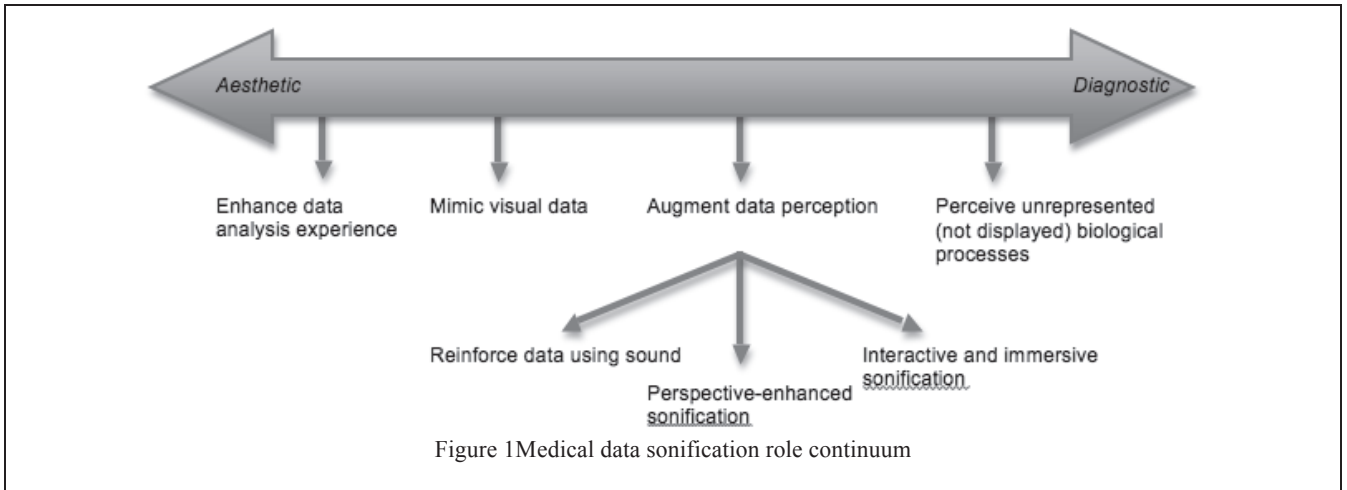
The central hypothesis of this paper is that auditory processing of medical image data will identify patterns associated with disease that cannot be detected by traditional means.

Using a model of molecular brain imaging whereby detection of small molecules is accomplished through the use of radioactive decay of injected "tracers" (PET imaging), we propose that patterns, when sonified, will emerge from the data and show information that was not previously detected by visual and pre-existing computer analysis techniques. These patterns, when heard and processed by the human brain, might one day allow the medical community to detect diseases that are presently invisible by currently existing methodologies. To our knowledge, this is the first report of the use of sonification in the analysis of three-dimensional medical image data.

4. IMMERSIVE SONIFICATION

Displaying 3D data on visual displays is particularly challenging because of the many dimensions needed to be presented concurrently. Our eyes are limited with their "field of view", and cannot see behind objects that obstruct the view. Our hearing, however, is omni-directional, and we can hear sounds emanating from around us, and in distance.

We refer to immersive sonification as a method of sonifying data in such a way as to place the viewer/listener of this data inside of an immersive 3D environment so that she may be able to navigate through the environment. The sonified data is rendered constantly. However, the listener will be closer/farther to a region of the data depending on their virtual location in the immersive data environment. For example, we can think of an immersive sonification of 3D brain data as putting the listener inside of the scanned brain. Through the use of spatial sound, we are able to create an immersive environment in order to present a spatially distributed sonification.



5. SONISCAN

SoniScan is a sonification tool developed in the Matlab technical computing software. Matlab is a flexible and versatile computing package that facilitates the manipulation of data, and synthesis and processing of sound. SoniScan provides a graphical interface through which a user can load, manipulate, and sonify Digital Imaging and Communications in Medicine (DICOM) data - a standard format for viewing and distributing medical imaging data. The program was constructed with a modular approach in mind at both a macro and micro level, in order to allow maximum flexibility for exploring sonification methods that are conducive to brain data display. The data and signal flow are illustrated in Figure 2.

The following four modules have been designed to map the data to sound: control, path, Sound synthesis, mappings, and audio controls. DICOM data is read into SoniScan, and data manipulation is performed in the Controls module. After passing through the Sonification Path module, which specifies the data path the sonification will follow, the Sound Synthesis module defines the details of the sound synthesis mechanism and passes to the Mappings module. Sound is generated using these parameters and the resulting audio is played using the Audio Controls module.

5.1. Control module

The Control module reads in and prepares the DICOM data to be sonified. This module has five main functionalities: data read, data selection and zoom, sonification duration, and data adjustment.

SoniScan loads DICOM-formatted files. A 3D matrix of the scan data is constructed and used for the sonification and the basic visual display provided. A single frame of data may be selected for sonification, a subsection of one frame, or a full three-dimensional scan. Using HRTF processing, the data selected is spatialized, a process by which directly correlates

the location of each data pixel to an apparent sound location heard by the listener.

Transposing the DICOM data values that are suitable for auditory presentations requires that all data be easily rescalable. Additional controls of range, shift and volume allow the mapping ranges to be adjusted. The range adjusts the data range. It allows a lesser or greater differentiation between the largest and smallest values in the data. The shift operation transposes pitches to higher or lower part of the musical scale (i.e. like shifting to higher or lower notes on a piano).

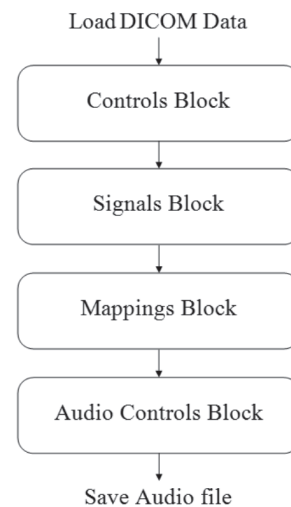


Figure 2 SoniScan schematic

5.2. Sonification path module

The sonification of the data stored in a matrix can be done by sonifying one data voxel at a time, by row, by column, or sonifying the full 2D or 3D matrix simultaneously. The Sonification Path module allows the user to define the exact path through which to scan the selected data. SoniScan contains three preset paths: from left to right, top to bottom, or all data simultaneously, as seen in Figure 3. If the left-to-right path is

selected, each data column is sonified and played concurrently, followed by the next column of data. If the right-to-left path is selected, each row is sonified, and rows are played concurrently. If simultaneous is chosen, effectively the path is removed: the sum of each row is taken, resulting in a single column, with each value representing the total values of the corresponding row. We refer to these three mapping trajectories as conventional paths.

In addition to the three conventional paths, we also allow sonification along a split-path. The selected data frame is split along a specified line (split line) and sonified as two halves. The split-line is specified by the cartesian x and y coordinates of any two points on the line. The motivation behind the split path is to facilitate detection of asymmetries between the two halves of the data, which correspond to the two brain hemispheres.

There are two types of split paths. The first is a hard left/right, where the selected data is split into two halves, the left side is sonified and played to the left ear, and the right side is sonified and played to the right ear. The second type is the difference scan, where the mirror of the right side of the data is subtracted from the left side. This can be thought of as folding the halves onto each other and taking a difference between the two halves. Resulting asymmetries between the two sides of the data are the only audio signals that are audible, since symmetric data is cancelled out. In both types of split paths, the sonifications generated are based on the mappings and scan paths specified, as described above.

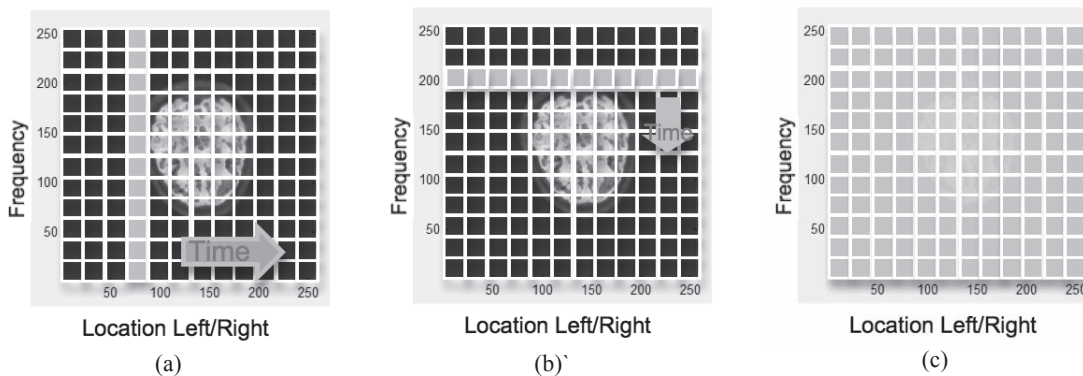


Figure 3 Sonification paths: (a) left-to-right; (b) top-to-bottom; and (c) all data.

5.3. Sound synthesis module

The Sound Synthesis module specifies the base signals to use for the sonification. It is modular in structure and can easily support additional signal types in the future. Any combination of signals contained in this module can be used simultaneously in the sonification. In our current implementation, the sonification can use band-passed noise, triangle wave tone, or a plucked string model. For example, when band-passed noise is selected, a noise signal is used for the sonification that is band-passed at the centre frequency that corresponds to a data value.

The parameters of these signals, namely amplitude, frequency, spatial location, and time of occurrence, are governed by the data to be sonified. The manner of mappings is specified in the Mappings block (next section). Any number of base signals can be simultaneously selected for sonification, and doing so will layer the signals on top of each other.

5.4. Mappings module

The Mappings module defines how the image data is mapped to different parameters of the audio. Three different mapping techniques have been used: amplitude mapping, frequency mapping, and spatialization. These are explained in detail in the following sections.

5.4.1. Amplitude mapping – Fourier synthesis

In the case of amplitude mapping, each pixel's intensity is mapped to the physical amplitude of the audio signal corresponding to that pixel. The greater the data value, the greater the amplitude of the audio signal resulting in a louder sound. Therefore, highly active regions of the data result in louder regions.

When the mapping is performed, different frequencies are assigned to different pixels based on their location in the image. Frequencies are distributed between 500Hz and 5kHz. How the frequencies are distributed depends on the sonification paths. For Path 1 (left to right), frequency varies from top to bottom,

with the highest frequency assigned to the topmost rows in the data (Figure 3a). For Path 2 (top to bottom), frequency varies from left to right, with highest frequency assigned to the rightmost column (Figure 3b). Path 3 (all data) follows the same frequency assignment as Path 1 (Figure 3c). While Paths 1 and 2 present each row or column sequentially, Path 3 renders an integrated spectrum of the entire image, since the same sets of frequencies for every column are presented simultaneously. A total sum of values is taken for each row/frequency, and this sum controls the amplitude of each frequency component.

5.4.2. Frequency mapping

In the case of frequency mapping, pixels are assigned frequency based on their intensity, with all amplitude values being kept at a constant level. The assignable frequencies are quantized to integer multiples of 500Hz to a maximum of 5kHz, so all generated audio contains only harmonic content.

The motivation behind performing frequency mapping is to judge the intensity distribution of a particular dataset through timbre. Datasets containing more high-intensity pixels will have more high-frequency harmonic content. This may in some cases result in a harsh timbre, if there is a large range of values being presented and many frequencies are sounding, or a ringing, if most values are high, and thus each pixel is rendering the same frequency. Hence, through frequency mapping, the user can gain a quick idea of the composite intensity at each step along the selected scan path.

5.4.3. Spatialization

The Spatialization module distributes sonified sound spatially around the listener. Assuming users would employ off-the-shelf headphones of reasonable fidelity, a great range of apparent stereo locations may be synthesized. Spreading the sound around the listener can result in a better differentiation of sonified regions, and thus a better distinction of features pertaining to sections of the data.

There are four spatial mapping methods employed in SoniScan: intensity panning, vertical spatialization, horizontal spatialization, and full 3D spatialization.

The intensity panning method uses interaural intensity differences between the two ears to pan sounds from left to right. When the sounds sent to the left and right ears are of equal level, the virtual auditory image appears to be located in the center. As the sound level of one of the ears increases, the location of the source appears to be originating from the side with the greater sound level. Intensity panning is effective for creating changes along the horizontal plane, but not the vertical plane. Using the intensity panning method, the audio corresponding to each pixel is panned to an apparent position based on the Cartesian x-coordinate of the pixel. Pixels that are in the center of the image will be panned and perceived to be coming from the center, those that are on the right side of the image will be panned to the right, and so on. A left-to-right image path will result in the sound containing all the vertical image data moving from left to right, while in a top-to-bottom image path the sound constantly surrounds the listener from left to right, and the vertical data is presented consecutively.

Vertical HRTF spatialization utilizes head-related transfer functions to map the selected dataset onto a two-dimensional (up-down, left-right) vertical aural image space directly in front of the user. The spatialized audio correlates with the spatial distribution of the visual image. For example, data on the top left corner of the image is sonified and perceived to be coming from a high elevation, the left data that is in the middle of the image is presented at ear level, and data in the lower part of the image is spatialized below ear level.

Horizontal HRTF Spatialization utilizes head-related transfer functions to map the selected dataset onto a two-dimensional horizontal aural image space (front-back, left-right) that places the user in the center. Effectively, the image is laid flat on the horizontal plane, and the listener is placed in the center of the image. The audio corresponding to the sonified dataset is spatially placed all around the user on the horizontal plane. For example, data that is in the upper left corner of the image would be sonified and presented in the left front of the audio image.

The full 3D spatialization method can be used for data that is three-dimensional. This would be a VR version of the data, with a series of horizontal or vertical scans all active simultaneously. HRTFs and distance mapping are used to position the data points around the listener – front, back, left, right, up, down. In this case, all the data can be sonified concurrently and spatialized to reflect the position of each data point relative to the listener. The listener can place herself in the middle of the data, or listen to the data from another perspective. For example, the listener can be center of a 3D brain scan and listen to all the data in the scan concurrently. By selecting different listening locations, the listener can effectively “walk through” the data in a fully immersive manner.

5.5. Audio control module

The Audio Control module manages the sonification process and handles the generated audio files. Once the relevant sonification and data parameters are set, the sonification is performed, and the rendered sonification can be played or stored to disc.

There is an additional A/B playback functionality, which facilitates the serial playback of two different sonifications as an A/B comparison. This is useful when comparing two sets of data that may have been taken at different points in time, or comparing two cross-sections of the same dataset in order to perceive the difference between the two. This capability for A/B comparison through sonification function is particularly important since the perceptual auditory system is more acute than the visual system at detecting temporal, spectral, and spatial changes. When small changes in data occur, they may not be immediately noticeable on the visual display, but may be more easily observed using sonification.

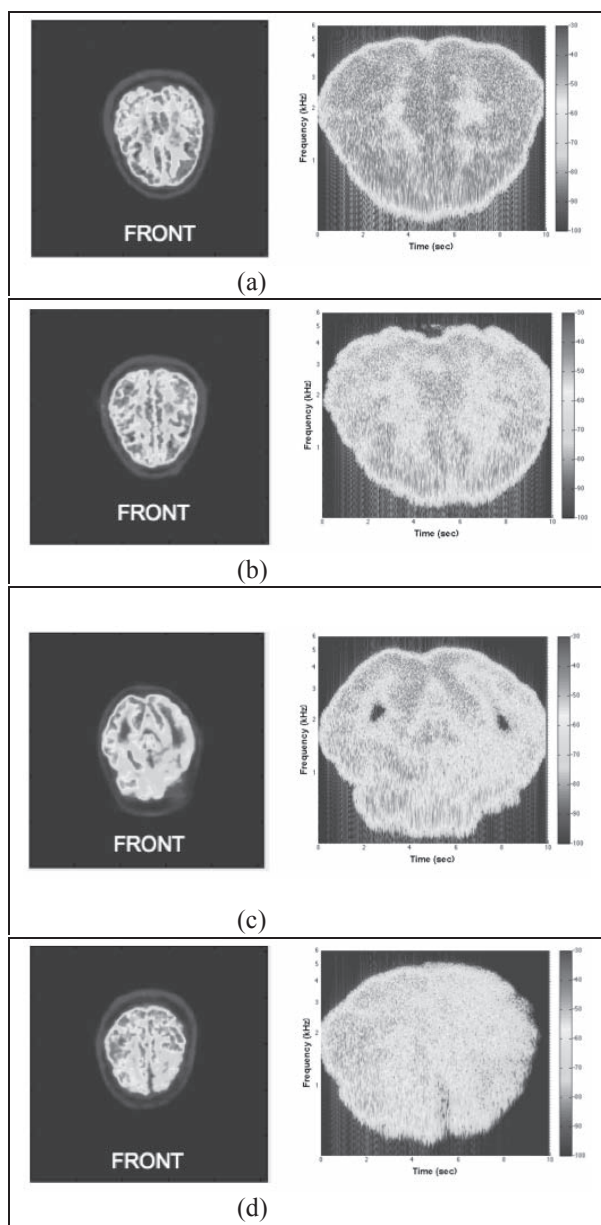


Figure 4 Images of brain scans (left column) and spectrograms of sonifications (right column) of healthy brains (a and b), and unhealthy brains with Alzheimer's dementia (c) and frontotemporal dementia (d).

5.6. Examples

Due to the limitations of presenting audio examples in a written paper, we are using spectrograms (time-varying representations of the spectrum of an audio signal) to show the correlation between the PET scans and their sonifications. Figure 4 contains four examples of the spectrograms of sonifications that were created from normal and abnormal brain data. The left side of the figure contains single slices from three-dimensional PET scans depicting sugar utilization in the brain. On the right side are the spectrograms of the sonifications. In these

examples, the x-axis of the spectrogram represents time, while the y-axis shows the frequencies from low (bottom part of the graph) to hi (top part of the graph). The intensity of the color represents the amplitude of the spectral content at each point in time – with blue indicating low amplitude, and red high amplitude.

The top two examples – Figure 4(a) and Figure 4(b) - are of healthy brains exhibiting a homogenous and symmetric pattern of sugar metabolism. As can be seen in the spectrograms (and heard in the sonifications), the symmetry of the spectral content as well as the full bandwidth reflects normal glucose uptake in the brain. Conversely, in Figure 4(c), there is asymmetric and decreased signal intensity in the brain of a patient with Alzheimer's dementia. A lack of low frequency content is heard in the sonification, and is visible in the second part of the spectrogram, starting at approximately 13 seconds. This region on the spectrogram corresponds to the most severely diseased portion of the brain, resulting in an audible hole in the frequency spectrum. Likewise in Figure 4(d), an image of a brain scan of a patient with frontotemporal dementia is presented and demonstrates a lack of low frequency spectral content corresponding to the frontal lobes of the brain. This leads to an unusual sonic representation of that hemisphere.

6. DISCUSSION & CONCLUSIONS

Brains scans contain complex, highly variable data and present a challenge to the interpreting physician. Imaging experts spend years learning to properly read such studies, yet detection of subtle disease remains difficult. Compounding matters, many disease processes remain invisible even to the best observers, either due to lack of meaningful information or undiscovered means by which to identify the relevant data. Visual quantitative techniques have improved matters but there is room for further improvement. It is suspected that as-yet-undiscovered information exists within these images and has diagnostic and therapeutic relevance.

Development of new ways to understand complex, multi-dimensional data from brain and other body scans is important for advancement of the medical imaging sciences. Sonification seems to be an appropriate target for further studies in this area, harnessing the auditory system's spatial acuteness and omnidirectional hearing in order to identify patterns in data that may not be apparent by other means. The examples presented above confirm that even preliminary work in sonification can begin to differentiate disease from health. Increasingly complex auditory modelling might one day reveal sophisticated and relevant medical information.

The tool described in this paper, SoniScan, is a research system developed for exploring the benefits of sonification for medical data and researching sonic parameters that are promising in guiding physicians as they diagnose disease and monitor treatment. With the sonification mappings we have explored thus far, we see a significant and audible correlation between the image data and the sonification. These correlations serve as a foundation upon which refinements can be explored to pick out increasingly subtle variations in image data that might one day carry diagnostic relevance and ultimately impact patient care.

We are still in early stages of evaluating our renderings and exploring new mappings between PET data and audio parameters. The parameter mappings will significantly affect the audibility of patterns and possibility of extending sonification as providing value in monitoring and/or diagnosing disease. Initial reactions among the authors included that of an immediate, pleasurable sense of recognition when hearing the simultaneous ascending/descending pitch patterns, which clearly corresponded to the outline of the skull in the image and in/active areas of the brain. While not of any diagnostic value, this still made the “first handshake” with the environment a pleasant and engaging one. Areas of diseased brain are clearly audible as distinct sounds contrasting with the sonifications of healthy patient scans.

During a cursory comparison of normal and abnormal brain scans, it was immediately apparent that there was much more scattered activity with the dementia scans than with the normal scans, and there was a certain sonic quality that might be described as more “strident” in the central area of the scan. By zooming in on the central area of greatest activity, it became apparent that the dementia scans consisted of activity farther forward in the brain, due to the later starting time of audible activity in the L-R scan.

Next steps for this work includes subjective studies with junior and senior-level nuclear medicine physicians to systematically compare the sonification performance with traditional techniques and determine if it reveals additional information that may improve upon existing analysis methods.

7. ACKNOWLEDGMENT

This work was supported by the NYU CTSI Award Number NIH/NCATS UL1 TR000038.

8. REFERENCES

- [1] Baier, B., Hermann, T., 2004. The Sonification of Rhythms in Human Electroencephalogram. Proceedings of the 10th International Conference on Auditory Display.
- [2] Baier, B., Hermann, T., & Stephani, U., 2007. Multi-Channel Sonification of Human EEG. Proceedings of the 13th International Conference on Auditory Display, p. 491-496.
- [3] Ballora, M., Pennycook, B., Ivanov, P. C., Glass, L. & Goldberger, A. L., 2004. Heart rate sonification: A new approach to medical diagnosis. LEONARDO, 37, p. 41-46.
- [4] Cassidy, R. J., Berger, J., Lee, K., Maggioni, M., & Coifman, R. R., 2004. Auditory display of hyperspectral colon tissue images using vocal synthesis models. Proceedings of the 10th International Conference on Auditory Display.
- [5] Edwards, A.D.N., Hines, G., Hunt, A. 2008. Segmentation of biological cell images for sonification. Proceedings of the 2008 Congress on Image and Signal Processing, Vol. 2, p. 128-32.
- [6] Ferrie, C.D., Marsden, P.K., Maisey, M.N., Robinson, R.O. 1997. Visual and semiquantitative analysis of cortical FDG-PET scans in childhood epileptic encephalopathies. J Nucl Med, 38(12), p. 1891-1894.
- [7] Hermann, T., Meinicke, P., Bekel, H., Ritter, H, Muller, H.M., Weiss, S. 2002. Sonifications for EEG data analysis. Proceedings of the 8th International Conference on Auditory Display.
- [8] Hong, S. L., 2007. Entropy Compensation in Human Motor Adaptation. Ph.D thesis in Kinesiology, Penn State University.
- [9] Jovanov, E., Starcevic, D., Wegner, K., Karron, D., & Radivojevic, V., 1998. Acoustic rendering as support for sustained attention during biomedical procedures. Proceedings of the 5th International Conference on Auditory Display, p. 1-4.
- [10] Kono, A.K., Ishii, K., Sofue, K., Miyamoto, N., Sakamoto, S., Mori, E. 2007. Fully automatic differential diagnosis system for dementia with Lewy bodies and Alzheimer's disease using FDG-PET and 3D-SSP. Eur J Nucl Med Mol Imaging 34, p. 1490-1497.
- [11] Minoshima, S., Frey, K.A., Koeppe, R.A., Foster, N.L., Kuhl, D.E. 1995. A diagnostic approach in alzheimer's disease using three-dimensional stereotactic surface projections of fluorine-18-FDG PET. J Nucl Med 36(7), p. 1238-48.
- [12] Müller-Tomefelde, C., 2004. Interaction sound feedback in a haptic virtual environment to improve motor skill acquisition. Proceedings of the 10th International Conference on Auditory Display, p. 1-4.
- [13] Najafi, M., Soltanian-Zadeh, H., Jafari-Khouzani, K., Scarpace, L. Mikkelsen, T. 2012. Prediction of glioblastoma multiform response to bevacizumab treatment using multi-parametric MRI. Plos One 7(1), p. 1-7.
- [14] Oh J, Wang Y, Apte A, Deasy J. 2012. SU-E-T-259: A Statistical and Machine Learning-Based Tool for Modeling and Visualization of Radiotherapy Treatment Outcomes. Med Phys. 39(6), p. 3763.
- [15] Pauletto, S., Hunt, A. 2006. The sonification of EMG data. Proceedings of the 12th International Conference on Auditory Display. p. 152-157.
- [16] Wallis, I., Ingalls, T., Rikakis, T., Olsen, L., Chen, Y., Xu, W & Sundaram, H., 2007. Real-Time Sonification of Movement for an Immersive Stroke Rehabilitation Environment. Proceedings of the 13th International Conference on Auditory Display, pp. 497-503.