

SMITH-KETTLEWELL DISPLAY TOOLS: A SONIFICATION TOOLKIT FOR MATLAB

Joshua A. Miele

The Smith-Kettlewell Eye Research Institute
2318 Fillmore St.
San Francisco, CA 94115
jam@ski.org

ABSTRACT

The Smith-Kettlewell Display Tools (SKDtools) is an accessibility toolkit for the popular modeling and data manipulation environment MATLAB. It provides blind and visually-impaired users with a variety of sonification and tactile display options, and is assembled from tools already existing within the MATLAB environment. Two distinct sonification methods are implemented, and one of the intentions of the project is to examine the relative effectiveness of these two methods using rigorous psychophysical techniques. SKDtools was developed at The Smith-Kettlewell Eye Research Institute's Rehabilitation Engineering Research Center under a grant from the National Institute for Disability and Rehabilitation Research.

1. INTRODUCTION

1.1 What is MATLAB?

MATLAB — a product of The MathWorks, Inc. — is a powerful number-crunching application that provides a multitude of data manipulation and display features. It includes a C-like interpreted language that can be used to create scripts and functions. A set of object-oriented graphics tools allows the user to draw plots, render images, and create graphical user interfaces. It is a cross-platform environment. Thus, tools developed on a UNIX machine are immediately portable to a Macintosh, Windows, or other MATLAB implementations.

In many ways MATLAB has become the computational lingua franca of science and engineering. Over the last ten years or so, the use of MATLAB in numerically intensive areas of research and development has become almost ubiquitous. Astronomers, computer scientists, economists, geophysicists, mechanical engineers, oceanographers, psychoacousticians, and many others have found MATLAB to be well-suited to their needs, with each community of users creating libraries of specialized functions and visualization tools tailored to their specific requirements. Collaborating researchers working in different labs will often share algorithms or data visualization techniques by exchanging m-files (MATLAB programs).

1.2 Project Motivation

Although access to information technology is an on-going challenge facing most blind and visually-impaired people, it is of particular importance to those involved in the highly technical fields of science and engineering. Inevitably, each discipline and sub-discipline has its own idiosyncratic hardware and software tools and modes of displaying data. These specialized tools and representations are rarely impossible to adapt into an accessible format. Nevertheless, the practicalities of doing so are almost always the first (and altogether too often the last) hurdle for blind people pursuing technical careers or advanced degrees in science or engineering.

In the days when each specialized tool or technique was implemented in dedicated hardware or software, the task of providing an accessibility solution required each research tool to be customized and adapted individually. This presented a daunting barrier to accessibility for each field's blind pioneers, and any hardware solutions developed were expensive to build and difficult to duplicate. However, with so many specialized tools now implemented in the MATLAB environment, it is possible to create a generalized set of accessibility functions that can be interfaced with existing MATLAB code to allow blind and visually-impaired researchers to more easily represent data in an accessible format. Such a generalized set of accessibility tools would allow each discipline's specialized MATLAB programs to be adapted readily and reliably. Furthermore, by developing the accessibility functions in an open-source environment, all MATLAB-using members of the blind and visually-impaired community may benefit and contribute.

Smith-Kettlewell Display Tools (SKDtools) is an example of such an accessible data representation toolbox. It was developed at the Rehabilitation Engineering Research Center (RERC) of The Smith-Kettlewell Eye Research Institute (SKERI) in San Francisco. The RERC is well-known for its early work in developing sonifying lab test equipment such as oscilloscopes, continuity testers, multi-meters, barometers, compasses, and much more, all with sonified output. SKDtools is one of the latest additions to SKERI's contributions to accessibility for tools needed by blind professionals involved in technical careers. It offers a starting point, both technically and conceptually, for others to begin designing sonification and other accessibility features for the MATLAB environment.

SKDtools was created and released with three main purposes. SKERI's primary intention was to provide a free, readily available set of tools for blind MATLAB users to produce sonified and tactile data representations from within their main work environment. Secondly, it was hoped that SKDtools could serve as a basis for other developers to create their own MATLAB-based sonification and tactile representation accessibility functions. A third intention was to implement two separate methods of data sonification — Discrete and Continuous Modes — in order to provide a practical platform for comparison of the effectiveness of the two methods for data sonification.

2. Description of SKDtools

SKDtools can be broken down into two main tool types: those for auditory data representation (sonification), and those for tactile data representation. The remainder of this paper is devoted to SKDtools' sonification features, although the tactile representation functions include a number of exciting and innovative aspects worthy of discussion in their own right.

SKDtools uses MATLAB's built-in ability to send digital signals to whatever standard sound hardware is connected to the work station. The sonified output is synthesized using a combination of SKDtools-specific dedicated algorithms and MATLAB's existing DSP features. The sonification scripts make no assumptions about the specific sound production hardware being used, thus preserving MATLAB's cross-platform convenience.

SKDtools is intended for presentation of two-dimensional (x-y) data. It uses a two-speaker presentation method and assumes that the listener is placed symmetrically between the speakers (headphones also satisfy this assumption). The first data dimension (the abscissa) is represented by lateralizing the sonified data trace in left-right space according to a simple amplitude panning algorithm [1]. The other dimension (the ordinate) is represented by varying the frequency of the sinusoidal output trace.

2.1 Attentional and Psychoacoustical Considerations

Lateralization of Sonified Trace: The horizontal component (x-axis) of the sonified trace is lateralized in left-right auditory space by use of a simple level-panning algorithm. Without full control over speaker placement, room acoustics, and head placement, to say nothing of head-related transfer functions, the exact perceived position of the trace in auditory space cannot be predicted. At best, it can be stated that there is some continuous, lateralized, auditory space with a one-to-one relationship between input value and trace location. Such a system makes it difficult for a naive subject to "read" exact x-values from a sonified trace. However, given a stable head/speaker configuration, minimal training may enable listeners to form

associations between lateral position and x-values with a reasonable degree of accuracy [2].

Single Auditory Trace: In a visually presented x-y space, it is possible to display many traces simultaneously to facilitate comparisons between different sets of data. The auditory domain is poorly suited to attending to multiple streams simultaneously. One reason for this is that it is possible for signals of different frequencies to mask one another in the auditory periphery [3]. This leads to interference between multiple simultaneous auditory traces. In addition, if peripheral (energetic) masking is ignored as a factor, there may still be issues related to informational masking ([4], [5]) and divided attention [6] that may result in interference among simultaneous auditory traces. For these reasons, SKDtools presents only a single auditory trace at a time. If comparisons among multiple data sets are desired, it is easy to make the appropriate manipulation to reveal the relationship of interest. For example, in order to observe the difference between two data sets, one may subtract one set from the other and sonify the result.

Frequency Limitations: As with any digital signal, the upper limit of the output signal frequency is set by the sampling rate ($F_{\max} = \text{SamRat}/2$). Because pitch becomes indistinct at frequencies greater than approximately 4 kHz, and in order to minimize unnecessary computational overhead, the default sampling rate for SKDtools is set at 11025 Hz — a standard sampling rate manageable by most audio hardware. This puts the maximum allowable output frequency at 5512.5 Hz. The lower limit of pitch perception is around 80 Hz, however, the threshold for low frequency tones increases rapidly with decreasing frequency. Thus, in order to ensure that the sonified signal lies well within the flat region of normal auditory perception, SKDtools sets a lower frequency limit of 200 Hz [3]. It is also worth noting that most small speakers such as may be attached to a desktop computer have poor response in the low frequency region [7]. This is another reason to limit the bottom end of the auditory display space.

2.2 Discrete and Continuous Modes

SKDtools implements two main methods of data sonification: a Continuous Mode which modulates a single sinusoidal tone in proportion to the input value, much like a voltage controlled oscillator (VCO); and a Discrete Mode which uses a set of fixed-frequency sinusoids, presenting only the tone or tones nearest to the corresponding output frequency. In both cases, a linear increase in the input values results in an exponential increase in the output frequency. In other words, a given increase or decrease of input value will result in the output frequency changing by a fixed musical interval.

The SKDtools Continuous Mode is similar to a number of existing sonification implementations on other platforms. However, one of the classic problems associated with this kind

of continuous, frequency-modulated (FM) sonification is that unless a listener has absolute pitch, it is impossible to interpret exact ordinate values from the sonified trace. Indeed, unless the listener has some ability to distinguish relative pitch, and the sonified data uses a tonal frame of reference, interpreting exact data values from the trace will be impossible. Additionally, continuous FM tone sonification schemes do not lend themselves to accurate analysis of slope or acceleration information. For these reasons, the continuous VCO-style sonification paradigm has most often been offered as a method of qualitative data inspection, rather than as a means of extracting quantitative conclusions[8].

The Discrete Mode was developed in an effort to improve the ability of listeners to extract quantitative information such as ordinate value, slope, and acceleration from sonified data representations. Whereas Continuous Mode presents a single tone that modulates in frequency within the allowable range to represent different values, Discrete Mode uses a set of fixed-frequency tones spaced at even exponential intervals throughout the full range of frequencies — in other words, each tone is spaced at equal musical intervals from its adjacent higher- and lower-frequency neighbors. At any given instant, the Discrete output signal consists of one, or a mixture of two, of these constant-frequency tones. Where the Continuous Mode would use a single frequency between two of the adjacent Discrete output tones, the Discrete Mode uses a combination of the two tones nearest in frequency to that output value. When an input value does not correspond exactly to one of the frequencies in the Discrete display, the two nearest tones are heard in combination. Their amplitudes are weighted in proportion to the distance from the input value to each adjacent tone's frequency. Thus, when an input value corresponds exactly to the frequency of one of the possible output tones, only that tone is heard in the output signal. By the same token, when an input value corresponds to a frequency at the half-way point between the two Discrete tone frequencies (in log space), the two tones have equal power in the output signal.

Although difficult to describe in words, the benefits of the Discrete Mode are readily apparent to the listener. With each adjacent tone representing an equal step in input space (for example one unit), it is easy to count steps from a known value — such as a zero point — to obtain the trace value at some other point. Similarly, as the trace changes between Discrete frequencies, temporal information in the form of relative tone duration provides a distinct cue to slope and acceleration. It is proposed that these features will result in an improvement in performance relative to the Continuous Mode.

2.3 Other selected features

SKDtools includes a number of extremely useful features. The entire complement is not possible to describe due to space considerations. However, a selected set are of particular importance and are worthy of brief mention.




Figure 1. The diagonal line represents a Discrete Mode sonification of a linearly increasing data set. Horizontal lines represent pure tones of different frequencies. The amount of a given shading type in the diagonal trace represents the amplitude of that frequency component in the Discrete display at that specific point in time.

Zero Marker: Each time the sonified trace crosses the “x-axis,” the output signal includes a brief (100 ms) burst of broadband noise. This indicator serves to anchor the display. The listener is thus able to count units in the Discrete display to obtain value estimates of the output trace. The zero marker is also helpful in determining periodicity of a trace.

Off-scale Indicators: As already mentioned, the output frequencies of the sonified trace are limited at the top and bottom of the spectrum by psychoacoustics, electromechanics, and sampling rate. Thus, it is possible to include input values that are either too high or too low for SKDtools to sonify. When values are too low to be represented, the auditory trace for those points consists of a low-pass noise with a cutoff of approximately 300 Hz. When the input value is too large to be represented, the output trace at the off-scale points consists of a high-pass noise with a cutoff at approximately 4 kHz. This makes it easy to know when to rescale one’s sonification to better fit the auditory display space.

Horizontal Markers: Listeners are generally more accurate in localizing broadband signals than pure tones. In order to improve "readability" of x-axis locations, SKDtools includes the option of adding "ticks" at fixed intervals, or specific locations,

along the horizontal axis. By default the ticks are short (50 ms), square-windowed, 1 kHz tone bursts. The x-axis ticks are easy to lateralize, and are quite distinct from the broadband noise burst used to indicate x-axis crossings.

3. Proposed Experiments

The following experiments are currently in progress with results to be described in detail in subsequent publications. All are being conducted using headphones in a sound-proof booth. Subjects with a range of scientific, mathematical and visual experience are being used.

3.1 Trace reproducibility

Listeners will hear randomly generated curves of fixed duration and will be asked to reproduce the sonified trace using a stylus on a 2-dimensional touch tablet. Responses will be scaled and correlated to the input data. Thus, an estimate of trace reproducibility will be obtained. Effects of subject experience, trace complexity, and display mode will be analyzed.

3.2 Trace Height Estimation

Subjects will listen to sonified traces of a single period of a triangle wave. Wave amplitudes (max frequency) will be varied over a pre-determined range. Frequency starting points will be randomized. Subjects will be asked to estimate the amplitude of the sonified wave. The response will be subtracted from the actual amplitude value to produce an error distribution around the correct value. Data will be collected for Discrete and Continuous Modes and the two response distributions will be compared.

3.3 Acceleration Detection

Sonified traces of data that differ in curvature (concave up vs. concave down) will be presented in a two-interval, two-alternative forced choice experiment. A threshold will be estimated for acceleration detection for both Discrete and Continuous presentation modes, and a comparison between these thresholds will be made.

4. Future Directions

Although it is already finding use among blind MATLAB users, SKDtools is still in its infancy. The possibilities for additional sonification schemes are limited only by the time and creativity of the community committed to the growth and success of THIS open-source sonification toolkit for MATLAB. As the

SKDtools' features expand, the opportunities for blind and visually-impaired MATLAB users will also grow. In addition, the Auditory Displays community will gain a common platform for the demonstration and usability testing of new and innovative sonification methods. The Smith-Kettlewell Eye Research Institute invites collaborators to contribute their modifications and additions to the existing SKDtools codebase by keeping in touch via the SKDtools home page at <http://www.ski.org/skdtools>.

[1] Pulkki, V. (2001). Localization of Amplitude-Panned Virtual Sources, Part 1: Stereophonic Panning. *J. Audio Eng. Soc.*, 49(9), 739-752.

[2] Wightman, F. L., & Jenison, R. (1995). Auditory spatial layout. In E. William Epstein & E. Sheena J. Rogers (Eds.), *Perception of space and motion*. (pp. xix, 499): Academic Press, Inc.

[3] Moore, B. C. J. E. (Ed.). (1995). *Hearing*: Academic Press, Inc.

[4] Leek, M. R., Brown, M. E., & Dorman, M. F. (1991). Informational masking and auditory attention. *Perception & Psychophysics*, 50(3), 205-214.

[5] Yost, W. A. E., & Watson, C. S. E. (Eds.). (1987). *Auditory processing of complex sounds*: Lawrence Erlbaum Associates, Inc.

[6] Bonnel, A.-M., & Hafter, E. R. (1998). Divided attention between simultaneous auditory and visual signals. *Perception & Psychophysics*, 60(2), 179-190.

[7] Hartmann, W. M. (1997). *Signals, sound, and sensation*. Woodbury, N.Y.: American Institute of Physics.

[8] Kramer, G., & Santa Fe Institute (Santa Fe N.M.). (1994). *Auditory display : sonification, audification, and auditory interfaces*. Reading, Mass.: Addison-Wesley.