THE SONIC MAPPER:
AN INTERACTIVE PROGRAM FOR OBTAINING SIMILAIRTY RATINGS WITH AUDITORY STIMULI

Gary P. Scavone

Center for Computer Research in
Music and Acoustics (CCRMA)
Stanford University
Stanford, CA 94305, USA

gary@ccrma.stanford.edu

Stephen Lakatos and Colin R. Harbke

Department of Psychology
Washington State University Vancouver
14204 NE Salmon Creek Avenue
Vancouver, WA 98686, USA

lakatos@vancouver.wsu.edu

ABSTRACT

The Sonic Mapper is an interactive Linux-based graphical program that affords increased methodological flexibility and sophistication to researchers who collect proximity data for auditory research. The Sonic Mapper consists of a mapping environment in which participants can position and group icons in the two-dimensional plane of the screen. Options for collecting data concerning hierarchical groupings, category prototypicality, and verbal labeling provide additional opportunities to test hypotheses in a convergent manner. The Sonic Mapper also offers an environment for traditional pairwise comparisons, as well as one for performing free sorting tasks. A pilot study that attempts to use many of the Sonic Mapper's key features is described briefly below.

INTRODUCTION

Conventional multidimensional scaling typically requires participants to provide similarity ratings for all possible stimulus pairings. For a complete data matrix, this yields a total number of comparisons as computed using Equation 1,

$$C = \frac{N(N-1)}{2} \qquad (1)$$

where C is the total number of comparisons to be made, and N is the number of stimuli.

With even 50 stimuli, for example, participants are required to perform 1225 judgments. Aside from obvious considerations of fatigue, it is questionable whether participants can maintain a consistent set of criteria across so many comparisons. Primarily designed to address the issues of fatigue and criteria inconsistency associated with the collection of similarity ratings for large stimulus sets, the Sonic Mapper is an interactive Linux-based program through which listeners arrange sound stimuli in a manner that reflects their relative similarities. In addition to addressing the issues for which it was originally designed, the Sonic Mapper also provides some distinct advantages over conventional pairwise comparisons.

Constraining participants to make pairwise comparisons, for instance, may discourage them from generating and applying richer sets of perceptual and cognitive criteria in their ratings. For instance, if sounds A and B have attributes v, w, and x in common, and sounds A and C have attributes x, y, and z in common, it is less likely that a participant performing pairwise comparisons will discover that B and C have attribute x in common than if he or she is able compare A, B, and C simultaneously. The Sonic Mapper provides a visual display of how all sounds relate to one another simultaneously. In addition, allowing participants to group and colorize sounds with common attributes facilitates the identification of these attributes, especially over extensive testing. Static visual display of the stimuli (via icons) alone, although it may enhance participants' global appreciation of the stimulus set, does not allow participants to select which, and in what order, items are to be compared. Such control of comparisons is afforded by the interactive nature of the Sonic Mapper.

In a related context, conventional randomized presentation of pairwise comparisons makes it impossible for participants to exert control over the order of comparisons - to adjust their similarity rating, for example, for a previously presented pair in light of a new criterion that he or she may have generated on the basis of a later comparison – and thus may discourage the development of alternative decisional strategies. The interactivity of the Sonic Mapper leaves decisions regarding how stimuli are compared primarily to the participants. They are free to employ any strategy and change strategies at any time throughout the experimental process. We believe this interactivity also allows participants to develop a more global appreciation of the stimuli than they could achieve through pairwise comparisons alone. However, such interactivity reduces the amount of structure to the experimentation. For this reason, the Sonic Mapper can also enforce specific comparisons. The program can either read these comparisons from an external file or automatically generate a set of comparisons based on configurable parameters such as minimum comparisons per item and percentage of all possible comparisons. This feature ensures that listeners make certain comparisons in a pairwise manner, and yet, it still allows participants the freedom to choose how, and in what order, to approach these comparisons.

The predominant use of continuous rating scales for pairwise comparisons may encourage participants to think unidimensionally about the similarities between stimulus pairs. For example, when dissimilarities between two sounds are represented along a unidimensional slider, listeners are less likely to consider both amplitude and pitch when rating sounds. Instead, listeners can, and often do, make their ratings based on the more salient attributes of the sounds (i.e., the attribute with the greatest similarity), or attempt some "cognitive averaging" that is unlikely to remain consistent over multiple pairs and obscures the attributes from which the averages were derived. The Sonic Mapper allows for similarity ratings along two continuous dimensions represented by the Euclidean distance between individual icons. Although two-dimensional ratings may appear advantageous over the unidimensional scales employed in conventional pairwise comparisons, the employment of 2-D representation of distance limits the listeners' ability to rate beyond these two dimensions – to think outside the plane of the screen, so to speak. This limitation, however, is, one that the Sonic Mapper shares with pairwise comparisons; there is no evidence that participants will rate stimuli beyond the single dimension provided to them when using a unidimensional rating scale.

The majority of published MDS solutions seem to be two-dimensional, and it is indeed rare to see solutions in the literature with four or more interpretable dimensions. It is unclear whether this may be a by-product of MDS algorithms themselves (i.e., that they tend to generate optimal solutions that are low-dimensional) or whether participants cannot access more than two or three perceptual/cognitive dimensions simultaneously, even though more might potentially be available. If low-dimensional solutions predominate, we wonder whether it might be better to make this assumption explicit in an algorithm, and to treat the presence of higher dimensions as an exception, rather than the rule. Finally, because of the two-dimensional constraint imposed by the computer screen, continuous similarity ratings made using the Sonic Mapper must be restricted to two dimensions.

OVERVIEW OF SONIC MAPPER FEATURES

Sonic Mapper allows participants to provide similarity estimates for large stimulus sets in an efficient manner. Participants move icons corresponding to individual sounds from a docket into the two-dimensional space of the computer screen in order to reflect relative similarities within the sound set. Afterwards, the Euclidean distances for the full matrix of sounds are computed. The experimenter can set several limits on the number and nature of comparisons per stimulus that participants must make during a session. Mnemonic devices are included to allow participants to keep track graphically of the nature and extent of past comparisons. There are several features included in the Sonic Mapper that augment the experimenter's ability to detect situations in which a k-D ($k \geq 3$) solution is more appropriate.

Sonic Mapper provides a high degree of interactivity, including opportunities for multiple comparison strategies, availability of mnemonic cues to remind participants of past comparisons, and the presence of feedback concerning classification schemes. SonicMapper allows a participant to form multiple hierarchical levels of groupings and subgroupings of stimuli using colorization and boundary definitions, with self-generated labels for each group (see Figure 1).

The experimenter can define a number of parameters separately for each participant before running a sound mapping session. In addition to allowing the specification of stimulus sets and basic display features, the program provides the option of requiring pairwise comparisons – derived randomly or according to a predefined matrix – on a desired subset of the stimulus set (Figure 2). The required comparisons feature permits one to test the validity of the mapper's two-dimensional layout by deriving a second multidimensional solution from the same data set based on the pairwise comparisons. An optional backup feature saves the current state of a participant's mapper configuration at

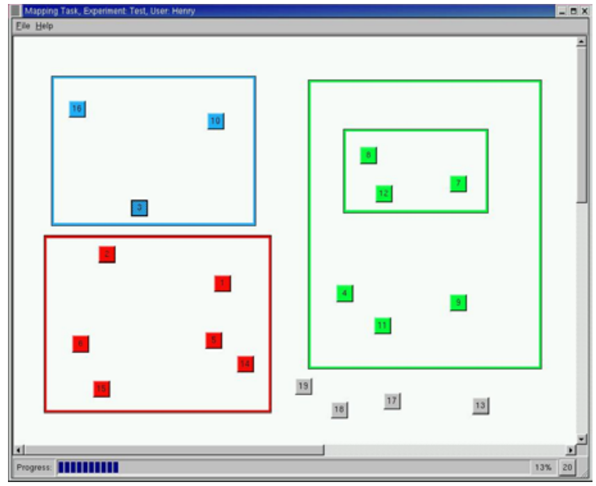

Figure 1. Main mapping window permits the participant to generate a two-dimensional similarity space with the option of superimposed hierarchical grouping boundaries.

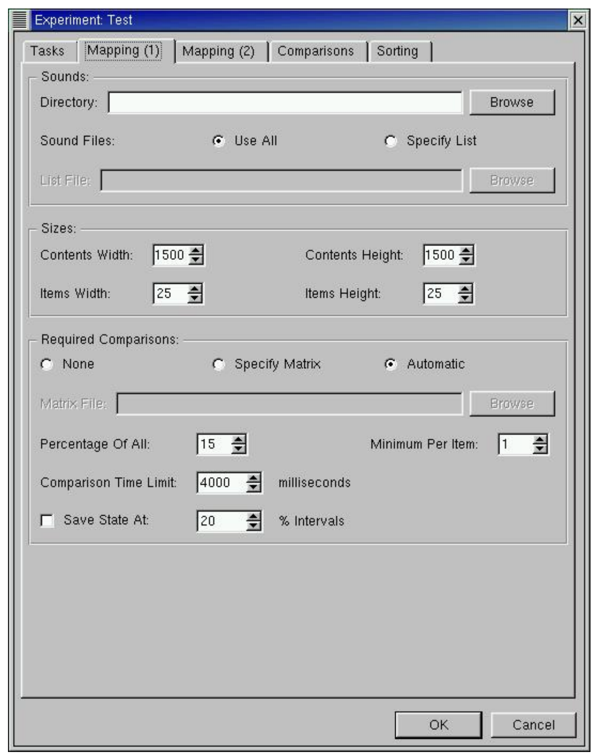

Figure 2. The first experimental specification window for the mapping task. Sound stimuli, visual display parameters, and required pairwise comparisons can be specified here.

specified intervals; if states are saved to separate files, a "time-lapse" view of a participant's arrangement of icons, consisting of an arbitrary number of sequential snapshots, can be obtained across the course of a complete experimental session.

Three supplementary data collection options allow the experimenter to test additional hypotheses about dimensionality and categorization (see Figure 3). After a participant completes his or her stimulus map, the experimenter has the option of requesting confidence ratings for the final location of each stimulus. Stimuli with low confidence ratings can then be tested further to determine whether they signal the need for additional dimensions or whether they represent “outliers” that have unique properties not easily represented by common dimensions. An option for permitting hierarchical groupings can also be enabled by the experimenter. Upper and lower boundaries can be placed on the number of groups or categories a participant is allowed to generate. Each group can be colored and labeled by the participant (labeling can either be forced or left optional by the experimenter). A separate menu item also allows the participant to view the entire grouping scheme for the mapper screen in a graphical tree format. For those interested in obtaining more specific information about group membership, Sonic Mapper provides the option of collecting prototypicality ratings for each stimulus relative to its group membership.

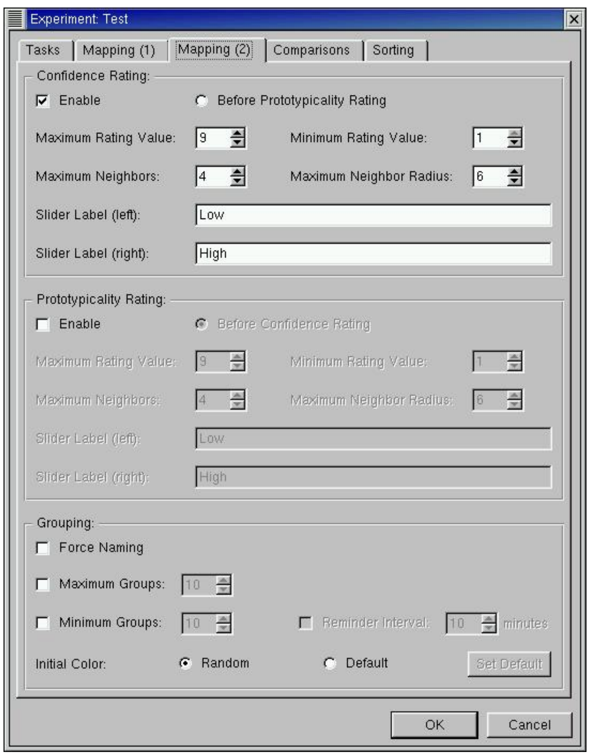

Figure 3. The second experimental specification window for the mapping task. Control is provided over several options for data collection, including hierarchical grouping, as well as confidence and prototypicality ratings.

To provide additional opportunities for collecting alternative comparison data within a single experimental session, Sonic Mapper has utilities for obtaining proximity data in the form of pairwise similarity ratings as well as ordinal sorting data. Figure 4 illustrates the interface used to collect pairwise similarity data. Participants adjust a continuously variable slider whose poles can be labeled by the experimenter in a separate specification window. The experimenter can also specify whether a full or half matrix is to be used to generate stimulus pairs (for half-matrix usage, an option to include or exclude the matrix diagonal is provided). Sound stimuli may be obtained by directing Sonic Mapper to a particular directory or to a text-file containing a list of the sounds’ names.

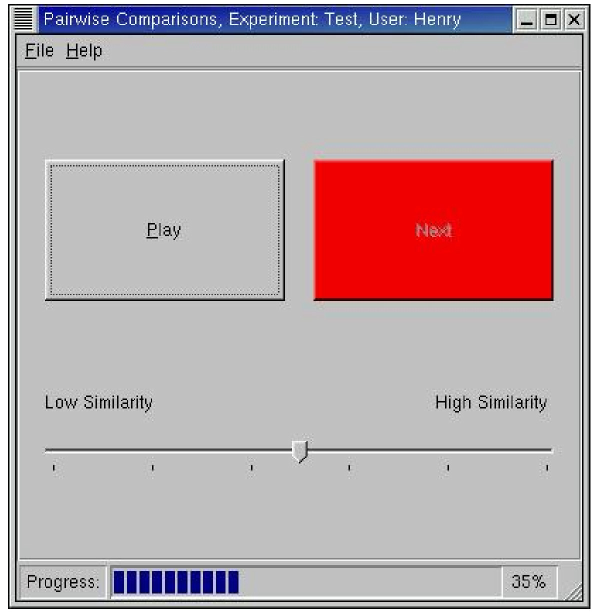

Figure 4. Sonic Mapper window for pairwise comparison task.

For the sorting task, sound stimuli are presented sequentially when the participant clicks on the speaker icon at the bottom of the main sorting window; stimuli are then place in one of *n* folders or bins specified by the experimenter. Through a separate specification window, the experimenter can place upper and lower bounds on the number of sorting groups or “bins” that the participant is permitted to create. Sorting groups comprise folders whose contents can be accessed at any time during an experimental session, so that the participant may reallocate stimuli to other groups at any time. Verbal labels may be provided for the groups by either the participants or, *a priori*, by the experimenter.

PILOT TESTING OF THE SONIC MAPPER

In our first use of Sonic Mapper, we obtained similarity judgments for 150 complex sounds from 22 participants. Stimuli were sound effects recorded from a variety of effects libraries (e.g., BBC). The sound effects all involve human-object interactions – a person walking on gravel, a person typing on a typewriter, or someone clinking a cup and saucer together. Since most of these stimuli evoke strong mental images of the sources or objects generating them, we compared participants who were instructed to focus on the timbre of each stimulus with participants who focused on the mental image generated by each stimulus. Participants were asked to generate between 5 and 15 hierarchical stimulus groupings using the interface, and participants in the mental imagery condition were also asked to provide verbal descriptors for each stimulus. In order to test the validity of a two-dimensional assumption for this space - since Sonic Mapper allows arrangement of stimuli in two dimensions only - we subsequently obtained pairwise similarity ratings from each participant for all within-category comparisons, as well as a randomly selected subset of across-category comparisons, and submitted the resulting dissimilarity matrices to MDS analysis.

Figure 5 shows the MDS solution for pooled similarity matrices for the "mental imagery" condition. In most cases, items clearly tend to group according to shared physical properties (e.g., "door" sounds at top of figure, "walking" sounds at middle right), in contrast to the timbre condition (figure not shown), where source-based clustering is not apparent. One general observation regarding how participants create their spatial arrangement: We were surprised by the enthusiasm that participants showed for the mapper task, working for several hours at a stretch. Most seemed to have a vested interest in thinking through carefully the complex problem of organizing large numbers of stimuli into a coherent perceptual representation, and many seemed genuinely proud of what they accomplished by the end. We conducted tape-recorded interviews with participants at the conclusion of the experiment, and were impressed with their ability to describe in clear terms the global aspects of their perceptual space. In short, the interactive character of the Sonic Mapper seemed to turn the task into an engaging problem-solving one.

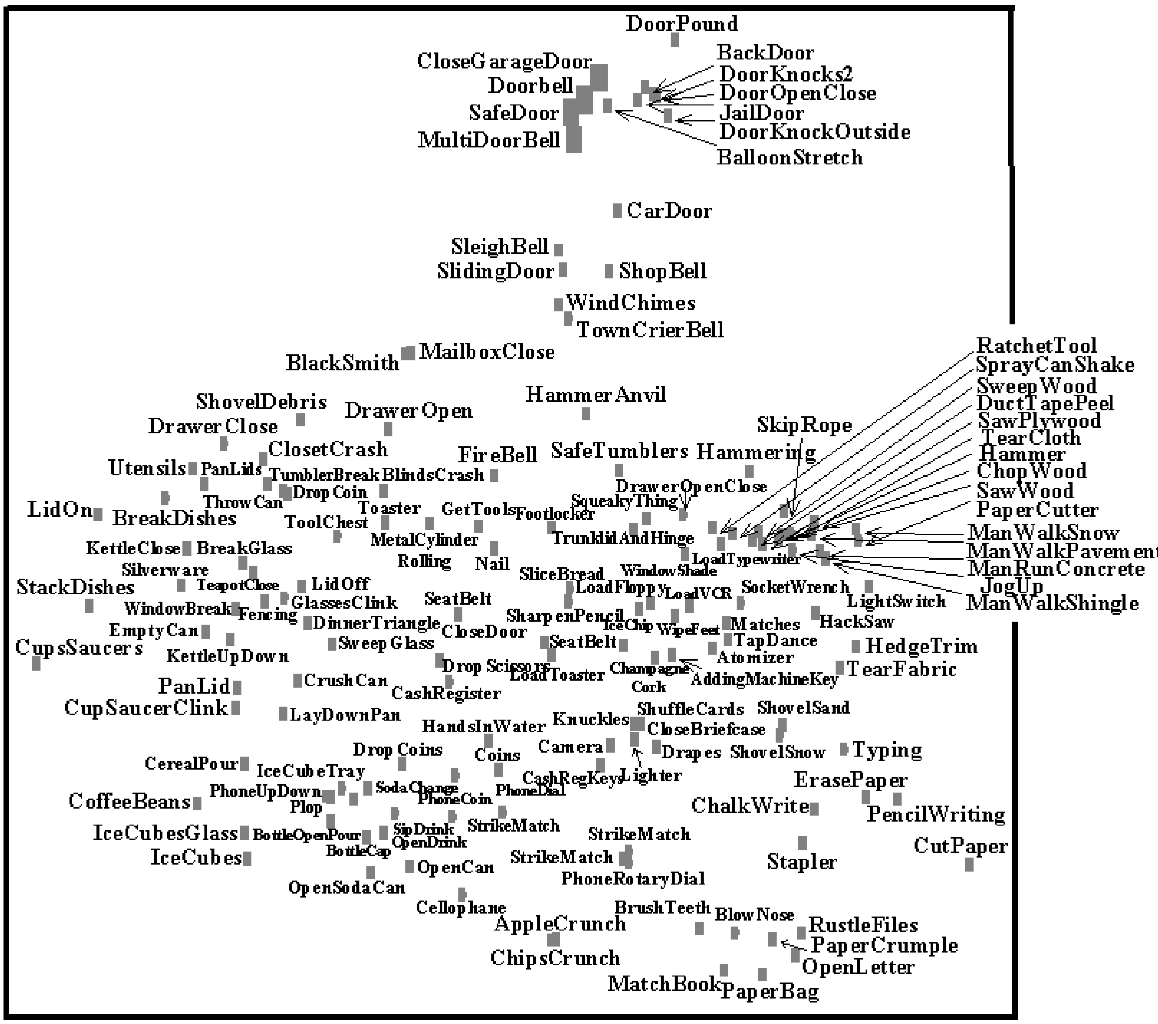

Figure 5. Multidimensional scaling solution (CLASCAL) derived from pooled similarity matrices obtained using Sonic Mapper.

CONCLUSIONS

The Sonic Mapper provides a viable alternative to conventional pairwise comparisons for collecting similarity data. In addition to being appropriate for large stimulus sets, the Sonic Mapper provides simultaneous presentation and manipulation (interactivity), as well as a two-dimensional rating scale. However, the Sonic Mapper is not without its limitations. Use of the Sonic Mapper in situations that require more than two dimensions is questionable, although still possible using an ordinal hierarchical sorting task, albeit with a loss in precision. Classification and grouping of stimuli along with confidence ratings are also incorporated in the Sonic Mapper's methodological repertoire to help the experimenter detect situations where a *k*-D solution would be most appropriate. Future versions of the Sonic Mapper will allow for visual stimuli sets and cross-platform use. A demonstration of the Sonic Mapper and its features will be included in the presentation.