Exploring Melodic Variance in Rhythmic Haptic Stimulus Design

Bradley A. Swerdfeger Jennifer Fernquist Thomas W. Hazelton Karon E. MacLean

Department of Computer Science
The University of British Columbia
{bswerd, adara, twhazelt, maclean}@cs.ubc.ca

ABSTRACT

Haptic icons are brief, meaningful tactile or force stimuli designed to support the communication of information through the often-underutilized haptic modality. Challenges to producing large, reusable sets of haptic icons include technological constraints and the need for broadly-applicable and validated design heuristics to guide the process. The largest set of haptic stimuli to date was produced through systematic use of heuristics for monotone *rhythms*. We hypothesized that further extending signal expressivity would continue to enhance icon learnability. Here, we introduce *melody* into the design of rhythmic stimuli as a means of increasing expressiveness while retaining the principle of systematic design, as guided by music theory. Haptic melodies are evaluated for their perceptual distinctiveness; experimental results from grouping tasks indicate that rhythm dominates user categorization of melodies, with frequency and amplitude potentially left available as new dimensions for the designer to control within-group variation.

Keywords: Haptic UIs, Multi-modal Interfaces, User Studies

Index Terms: H.5.2 [User Interfaces]: Haptic I/O;

1 INTRODUCTION

Information communication is a dominant element in computer interfaces. Today's computer users are inundated with visual and aural notifications, with the consequence of useful information becoming an irritating interruption. The haptic sense has the potential to support background communication that can be designed to reduce disruption in portable and embedded applications.

We anticipate that *haptic icons* – brief tactile or force stimuli associated with a meaning – will demonstrate the greatest utility in situations where other senses are occupied. For example, pilots must divide their attention among multiple visual representations of plane status to ensure a safe flight; one possible haptic representation indicates a needed pitch correction to the pilot through a graded vibration on the throttle stick [20]. At the consumer level, there are many opportunities in mobile devices for haptic notifications or directives (carrying information about e.g. incoming calls, navigation tasks or web browsing progress) when social or practical constraints make visual signals inconvenient or indiscreet. Unfortunately, haptic icons are a relatively new advent and we do not yet have heuristics or common design principles for communicating information reliably. For haptic icons to be useful, they must be: *distinguishable*, *learnable*, *appropriately salient* and *recognizable* in realistic conditions [16].

Previous work in the development of haptic icons has focused on the heuristic generation of rhythms at static frequencies and amplitudes appropriate for learnable, discrete message notification. Work by Brown and Brewster [3] and Ternes and MacLean [24, 25] used the constraint of monotonicity (single carrying frequency), to explore the vast rhythmic design space before embarking on melodic variation. However, use of the large rhythmic set described in [24, 25] revealed that as set size grows to 50 items or more, greater expressiveness is required for memorability. While a side-by-side comparison has not been made, it is likely that icons created with more expressive stimulus sets (i.e. containing more perceptually relevant variability) will also be more learnable when associated with meanings.

Haptic melodies[1] can be produced through note-by-note variation of frequency (tone) and amplitude (emphasis). (To imagine how frequency and amplitude could create a melodic vibration, consider touching the cage of a household electric fan. As the fan spins faster, its vibration increases in frequency. Now consider a larger fan, operating at the same frequency; the spinning of this fan will result in vibrations with a larger amplitude.) Allowing within-stimulus variation of other parameters explodes the design space. The present work examines how a sampling of haptic melodies are *perceived*, as a start towards heuristic guidance for future design.

The problem of making haptic icons more learnable has been approached by borrowing from psychological theory. 'Family-based' approaches to design [5, 6, 8] allow icons to be grouped by meaning along a display dimension. This approach facilitates learning by dividing the icon space into smaller, more manageable 'chunks' [18]. This previous work is limited by human perceptual capabilities aligned along the dimension provided by the haptic display, and the creativity of designers in devising metaphors. Our work aims to make family-based design more expressive and systematic by using *rhythm* as the primary grouping dimension for creating families of stimuli, while providing insight into what makes both rhythms and their melodic variants distinctive.

1.1 Approach and Overview

The primary goals of this work are to explore the use of melody in increasing the expressiveness of stimulus sets, while increasing usable set size (number and size of viable stimulus groups). In addition, we wish to determine and generalize the criteria that people use to group such stimuli, in support of future design.

To deal with the explosive increase in design space with the addition of melodic variation, we used heuristics derived from musical theory to compose our initial groups of stimuli, which were displayed on a mobile device with a piezo-actuated touchscreen (a modified Nokia 770T [13]).

Music composition rules are melodically or rhythmically based (Table 1). Ternes [24] states that for perceivability, stimuli should be limited to a small frequency range on the Nokia platform used here, and differences between frequencies must be relatively large. Thus, to enable creation of a large number of groups, rhythm was our main basis for group design. Most of our initial groups had related but not identical rhythms, while amplitude and frequency

[1]The term *melody* refers to a haptic stimulus that varies in rhythm, frequency and amplitude. The term *rhythm* refers to a haptic stimulus that only varies in rhythm. Note that a melody contains a rhythmic component.

varied within groups. To understand how users organize these stimuli (relevant to our secondary goal), we conducted user studies in which participants sorted stimuli into varying numbers of groups.

In the context of this melodic design space defined by rhythm, tone and intensity, we pose the following hypotheses: (H1) user-study participants' first-order stimulus groupings will follow *rhythm* in spite of frequency and amplitude variation; (H2) participants' first-order groupings will *not* follow either amplitude or frequency; and (H3) participants will demonstrate ability to discriminate stimuli which vary in amplitude and/or frequency but have consistent rhythm. Experimental validation of all three would confirm that rhythm should be used for primary groupings, and further imply that amplitude and frequency are suitable for within-group variation.

We also predict that participants' first-order stimulus groupings may be based upon ideas transferred from musical theory; although this conjecture is less certain and considered a path of investigation.

In the future, we also would like to test whether melodic variation enhances icon usability beyond that achievable with monotone rhythm; however, it will be challenging to devise an unbiased comparison.

In the remainder, we first outline several design approaches currently used to design to haptic stimulus sets, and justify the need for more expressive media and melody as a viable solution. Through an iterative design and evaluation sequence, we then show that even in the presence of melodic variation, users employ rhythm as a first-order grouping property of haptic stimuli; whereas melody is suitable (and important) for within-group stimulus variation. The following sections describe related work, our stimulus design process, the studies and their analysis, and conclusions we drew from them. We close with suggestions for future work.

2 Related Work

Iconography has been used in computer systems since the advent of WIMP interfaces. By taking advantage of the incredible human capacity for symbols [7], icons serve as a natural way to *represent information* in complex environments. With the recent advances in haptic display technology, researchers have been building on previous work in visual and primarily auditory icons to develop sets of *haptic icons*. A comprehensive overview can be found in [16]; here, we summarize the most relevant details.

2.1 Auditory Icon Design Approaches

Auditory and haptic icon design share a key attribute since both modalities are temporally sequential [11, 17]. Auditory icon design can be divided into two philosophies of design: *representational/metaphorical* and *abstract*.

Gaver [9] introduced auditory icons by representing information with a specific sensory experience that is directly related to the item being symbolized so that the link between stimulus and meaning is as intuitive and natural as possible. For instance, dragging an object might be accompanied with a 'scraping' sound. Unfortunately, this approach suffers from poor salience control. An unimportant event may be perceived as more salient or similar to a critical notification. This problem lies in conflict with Weiser's ideal of 'calm technology' [27] which serves as a widely accepted philosophy for non-visual interaction design.

Blattner *et al.* [1] take an abstract approach to designing structured 'Earcons' where 'motives' (a series of notes that differ in pitch and amplitude) can be combined to create compound icons. For example, one can combine the motives sequentially for 'destroy' and 'file' to represent 'delete file' abstractly. Brewster *et al.* [2] extended this work by examining how people can perceptually differentiate 'Earcons'. They found that the structured approach aided in differentiation, as did varying timbre rather than restricting stimuli to simple tones.

Our approach aims to provide guidance for designing *perceptually controlled* sets of icons that are expressive enough to approach Gaver's idea of creating representational icons while lending themselves to systematic design and salience management.

2.2 Haptic Icon Set Design Approaches

A variety of approaches to haptic icon design have been attempted since this is new ground that must be validated by research before it is applied commercially. These strategies include perceptual, musical and structural design.

2.2.1 Perceptual Design

MacLean and Enriquez [17] emphasize that the stimuli in haptic icon sets should be designed by first understanding how synthetic haptic signals are perceived and then later assigning meanings to these perceptually validated stimulus sets. Their design process involves fully exploring the output space of their device and then performing quick, iterative user studies to plot the perceptual space of their icons using MDS (**M**ulti**d**imensional **S**caling, elaborated upon in Section 2.3) in order to make their icon sets as distinguishable as possible. Using this method, they created 36 icons that vary in waveform, amplitude and frequency.

Using a similar method, Ternes and MacLean developed the largest stimulus set to date with 84 perceptually distinguishable tactile stimuli [25]. They created these stimuli by first using heuristics to choose 21 *rhythms*, then expanding this set with two variants each of amplitude and frequency, applied as a constant to the entire rhythm (4×21 monotone stimuli). This set is elaborated upon in Section 3.2. Although this is a very large stimulus set whose distinguishability was demonstrated, in deployment we have discovered that users find the rhythms easier to mentally assign to distinct meanings than their amplitude or frequency variations. These global adjustments, which register as changed intensity or tone, appear to be better suited to representing varying attributes for the icons (e.g. priority).

2.2.2 Musical Design

Van Erp and Spape [26] created a set of 59 haptic stimuli by translating music sequences from the auditory to the vibrotactile domain on the basis of note tone. They found that users distinguish these melodies on dimensions of *intrusiveness* and *tempo*. However, this investigation did not extend to meaning assignment, leaving open the question of whether designers can reassign arbitrary semantic associations to stimuli that might already have meaning for the user.

2.2.3 Structural Design

Attempts at creating structured sets of haptic icons have focused on family-based approaches: icons in each family share haptic features, increasing set learnability by allowing users to 'chunk' groups of items [18].

Chan *et al.* [5, 6] create a representational set of haptic icons in the context of remote collaboration, where the metaphor used for design reveals the family of the icon. They created seven icons by varying the order of tones, number of pulses and magnitude for families representing *changes* in control, *being* in control and *waiting* for control, respectively. Their set achieved 95% recognition rates under workload after three minutes of learning.

In Enriquez *et al.*, stimulus frequency corresponded to the icon's family and the waveform to its function [8]. They demonstrated a 76% recognition rate in completely arbitrary meaning-matches (for a conservative test) after 20 minutes of practice for a set of nine icons encoded as a 3^2 matrix.

Brown [4] found 73% user accuracy in identifying nine two-dimensional 'Tactons', where dimensions of priority and message type (3 priorities, 3 types) were encoded as roughness and rhythm.

Two of these studies [5, 8] also employed perceptual MDS to validate and refine the icons within their sets.

Although the family-based approaches appear to be effective, they are limited to relatively few families. Enriquez *et al.*'s approach limits its expressiveness for families to one dimension (frequency) and is therefore limited by perceptual acuity along that dimension. For instance, if the haptic device is capable of displaying a frequency range of 500 Hz and humans can only perceive differences of 100 Hz reliably, then the technique is limited to at most six families. Representational approaches illustrated by Chan *et al.* [5] and (in the auditory domain) Gaver [9] may be more learnable, but are weak both in repeatability (through reliance on designer creativity in generating good metaphors, which is particularly difficult for more abstract concepts) and salience management; and consequently in scalability.

Our work aims to increase the expressiveness and effectiveness of family-based approaches by coding family temporally, through the use of rhythm and melodic variance. This will increase the number of possible families (here, interchangeably expressed as 'groups') and allow for repeatable design since the process requires little creativity.

2.3 Perceptual Multidimensional Scaling

Multidimensional Scaling is an established technique for visualizing how users perceptually organize a set of stimuli. The algorithm takes a dissimilarity matrix of the stimuli and reduces this large dimensional space to N dimensions where the variance along these dimensions is maximized. These dimensions may give insight into complex, perceptual dimensions. For instance, Hollins *et al.* [12] found dimensions such as hard/soft and slippery/sticky for real tactile surface textures. In our work we use the relatively efficient and accurate cluster-sorting method adapted by MacLean and Enriquez [17] and further analyzed and validated by Pasquero *et al.* [21] and Luk *et al.* [15].

3 Stimulus Design

3.1 Design Parameters

Creative experimentation cannot be the sole basis for designing reusable sets of haptic melodies, limited as it is by designer creativity and situation-specific concerns. Instead, with an approach of 'perceptual design' we seek to understand how humans classify, compare and respond to melodies [16] by designing stimuli with informed heuristics and validating stimulus sets through user participation.

3.2 Stimulus Design Space

Haptic melodies are defined by rendering characteristics such as duration, tempo, frequency, note density and amplitude ranges, which in turn are constrained by hardware capabilities and human perception. Our haptic melodies are displayed on the modified Nokia 770T Internet Tablet, which has a piezo-based tactile 90×54 mm touchscreen displaying 800×480 resolution for stylus interaction (Figure 1; [13]). This device restricts vibration rendering to ≈ 50 discretely timed vibrations each at uniform frequencies and amplitudes per programmed script; therefore continuously varying amplitudes (e.g. complex sinusoids) are not feasible.

Constraints due to the need for a consistent and usable signal structure (e.g. overall duration, repetition to achieve rhythmic sense, empty space for note definition, etc) are taken from [25]. Ternes' 21 haptic rhythms for the Nokia 770T had the following characteristics:

- Each stimulus is 2 seconds long, with 4 identical, consecutive 500 ms repetitions.
- Each iteration is divided into 16 equally spaced segments.

Figure 1: The Nokia 770T Internet Tablet

- Each iteration is comprised of notes which occupy 16 (*whole*), 12 (*three-quarter*), 8 (*half*), 4 (*quarter*), 2 (*eighth*) and any number (rests) of segments.
- Spacing between notes is 31.25 ms (after eighth and quarter notes) or 62.5 ms (all others).
- Distinct 500 ms stimuli were devised based on heuristics.

Ternes expanded the 21 rhythms to 84 stimuli by playing each at one of 2 amplitudes and 2 frequencies (i.e. every note in a given 2-second stimulus has the same tone and intensity); and demonstrated set discriminability with an MDS visualization[25]. The mutual uniqueness of these stimuli is based solely on rhythmic variance.

To move from haptic *rhythms* to *melodies*, we developed new heuristics to systematically apply melodic variation (note-by-note changes in tone and intensity) to this base set of heuristically determined, perceptually validated rhythms. Ternes [24] suggests that two amplitudes and two frequencies best suit the capabilities of the human perceptual system.

3.3 Design Heuristics

The inherently rhythmic nature of haptic melodies suggests a potential for design guidance from musical theory. Evidence that audio and haptic signals are correlated and complementary [11, 14] suggests that musical techniques for eliciting human affect might transfer to the haptic sense.

Characteristic	Classification	Response(s)
Major tonal scale	Melodic	Happiness
Minor tonal scale	Melodic	Sadness
Slow tempo	Rhythmic	Serenity, Fear
Rapid tempo	Rhythmic	Jollity, Anger
Regular Spacing	Rhythmic	Calmness
Syncopation	Rhythmic	'Jarred' feelings

Table 1: Common associations between musical characteristics and human affective responses

Associations between melodic and rhythmic composition techniques (e.g., specific variations of pitch, tempo, use of syncopation, etc.) and human affective response are well documented. Table 1 lists associations derived e.g. from [10, 23, 22], which we used as a starting point for our own heuristic set. Since many rely on differences in rhythm, we also used them to produce new groups of stimuli centered around rhythms that depart from the original 21. Altogether, our initial heuristics were:

- **Ensure syncopation differences between melodies** (i.e., ensure that some rhythmic bases are perceived as 'off the beat' as opposed to being regularly spaced)
- **Ensure note density differences between melodies** (i.e., rapidity vs. languidness – resulting from the number of notes per melody – should differ between melodies)

- **Ensure differences in frequencies and amplitudes between melodies** (i.e., strength and variance of vibrations should differ between distinct melodies)
- **Ensure rhythmic differences between melodies** (i.e., rhythmic base of melodies should vary between groups)

3.4 Melodic Stimulus Design Tool

A haptic melody visualization, authoring and playback tool (called the *Melodic Stimulus Design Tool* or MSDT) was developed using Java 1.6 and the related Swing Project. It supports drag-and-drop note placement according to the constraints of the design space and melody structure (Figure 2). We used MSDT to produce haptic stimuli according to the heuristics above, by modifying frequency and amplitude composition of the 21 starting rhythms as well as producing new rhythms in line with the heuristics not captured in the original set.

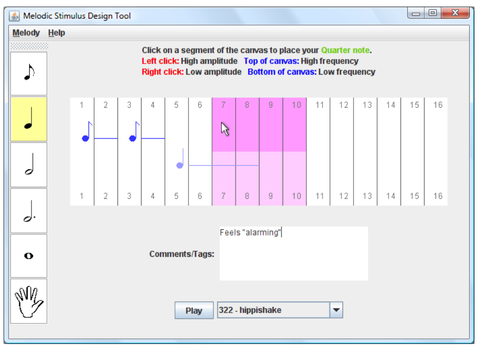

Figure 2: The *Melodic Stimulus Design Tool* showing a haptic melody consisting of two high frequency, high amplitude eighth notes followed by a low frequency, low amplitude quarter note.

MSDT's note canvas has two vertically stacked segments. The medium gray horizontal bar in Figure 2 under the cursor indicates that a mouse click will place a *high frequency* quarter note (4 segments) starting in the 7th segment of the bar. If the user moves the mouse downward to the light gray horizontal bar, the colours will swap and clicking will result in *low frequency* note placement. In both cases, a left click will place a *high amplitude* note, and a right click a *low amplitude* note. Opacity differences visually indicate note amplitude. Because a note starting in segment 7 would overlap with the quarter note already starting in segment 5, the pre-existing quarter note would be automatically removed to make room. The 'Comments/Tags' box has been used to note that the melody 'feels alarming', a tag which will be saved with the stimulus. Visual and auditory playback allowed stimulus preview; encoding was required to play them on the Nokia display.

4 Initial Stimulus Set

We ultimately produced 6 stimulus groups, each containing 6 stimuli. Based on the design heuristics laid out in Section 3.3, we developed perceptual distance metrics to mutually compare the initial 21 rhythmic bases. Results from this process provided a space from which we could evenly sample rhythmic bases. We chose 4 bases from the initial set of 21 to provide adequate coverage, allowing the two remaining groups to explore heuristics not captured in the initial set. Specifically, the groups based on concepts (syncopation or descending amplitude) employ new rhythmic bases not represented in the original 21 stimuli. The rhythm-based groups have identical note composition, varying only in frequency and amplitude of individual notes. The concept-based groups vary in rhythm, fullness (the number of segments in which a note is present), frequency and/or amplitude. The initial set is described as follows and illustrated in Figure 3:

Rhythm-based groups:

- **3Q**: 3 quarter notes followed by a rest (1–6)
- **1H2Q**: 1 half note followed by 2 quarter notes. (7–12)
- **2Q**: 2 quarter notes followed by 2 rests. (13–18)
- **4-6E**: 4 or 6 eighth notes; either 6 eighth notes; or 2 eighth notes, 2 rests, and 2 eighth notes. In the case that 6 consecutive eighth notes are used, the centre two are low amplitude to reduce their impact. (19–24)

Concept-based groups:

- **S**: all melodies are uneven or syncopated, with varying note densities and types. (25–30)
- **HLA**: high-to-low amplitude; the whole bar is filled with various note types. The first half always has high amplitude notes, while the second half always has low amplitude. The melody is always low frequency despite changes in note composition. (31–36)

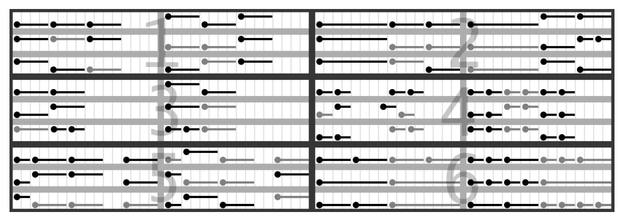

Figure 3: All 36 melodies from the initial six groups. Grey notes are low amplitude, raised notes are high frequency. Stimuli are listed left-to-right and top-to-bottom.

5 Study 1

5.1 Hypotheses

Group Number. Because we designed our stimuli in 6 groups, we hypothesized that users would also find 6 the most natural number of groupings for the full set. **Rhythmic Grouping.** We also hypothesized that users would sort melodies based on our rhythmic grouping ideas, outlined in Section 4. **Concept Grouping.** Finally, since this work is exploratory in nature, we hypothesized — tentatively — that users would likely find syncopation and consistent amplitude composition to be groupable criteria. Although this lies in conflict with the hypotheses laid out in Section 1.1, this seems somewhat reasonable based on musical theory transfer. The reconciliation of this hypothesis will provide additional design guidance, reduce potential sources of Type II error and avoid confirmation bias.

5.2 Participants

The initial study had 7 participants, a number in line with recommendations for MDS analyses [16]. All were Computer Science graduate students aged 18–25 (6 male).

5.3 Apparatus, Task and Design

Using the Nokia N770T tablet described in Section 3.2, participants were asked to group the melodies together into a specified number of bins based on their own similarity criteria. An application was created on the Nokia device to facilitate the grouping task (Figure 4). In the application, all 36 melodies appeared at the bottom of the screen as buttons. Their order was randomized each time a new grouping screen was loaded. The buttons could be clicked on to feel their melody, or dragged into a bin. Participants were asked to hold the device in their non-dominant hand and use the stylus for interaction.

Due to limitations of the device, only 16 of the buttons (indicated by an '!' symbol beside their number) had haptic feedback loaded at any given time. To load a button's feedback, a user would click on an unmarked button and, after a few seconds, feedback would load for that and the surrounding buttons. While this inconvenience increased the difficulty of the task, we have no reason to believe that it biased the final outcome.

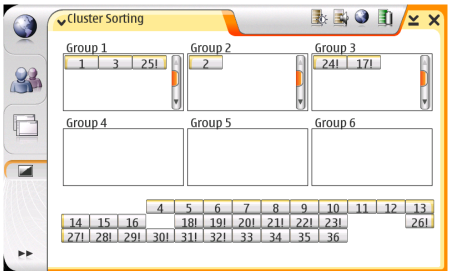

Figure 4: Screenshot of the sorting application on the Nokia 770T. A user is sorting the 36 stimuli into 6 groups.

Each participant performed 4 grouping tasks for the same 36 stimuli. The group sizes were 3, 6, 9, 12, presented in random order. The sorting data was recorded as well as subjective information obtained through various questionnaires.

5.4 Procedure

Participants completed a pre-session questionnaire to reporting demographic information such as country of origin, years of experience playing a musical instrument, and sense of rhythm.

Before completing the sorting tasks, participants were given a quick demonstration of the software and task. They were encouraged to develop their own similarity criteria with which to sort the stimuli, and asked to keep the bins filled at similar levels. After each of the four sorts, participants filled out a post-group questionnaire where they indicated the difficulty of choosing melody groups given the number of groups allowed in that sort.

After all four sorts, participants completed a post-session questionnaire asking for their preferred number of groups, the overall difficulty choosing a group for stimuli, sorting strategies, and meanings that came to mind for any of the stimuli.

Dissimilarity values for each user and stimulus pair were calculated identically to [17].

5.5 Results and Preliminary Findings

5.5.1 Qualitative MDS Analysis

The first study's results are shown in Figure 5. The most groupable characteristic appears to be consistent, non-syncopated rhythm, which is mostly in line with our Rhythmic Grouping hypothesis. Most melodies in the rhythm-based groups are clustered fairly close together in while those in the concept-based groups are not.

We carried out a detailed qualitative analysis by closely examining the MDS stimulus clusters.

Group 3Q. 3Q melodies 1, 2, and 4 were sorted together, but 5, 6 and 3 were not. While the 3Q melodies all have the same sequence of notes with varying frequency and amplitude, 3 and 6 have a varying amplitude that participants described as 'rolling' or 'pulse-y'. Conversely, 1, 2, and 4 had either no or one amplitude change so they did not feel the same as 6 and 3. These melodies were described as 'abrupt' or 'hard'. Melodies 1,2,3,4 and 6 ended in a high frequency or maintained a consistent frequency, in contrast to 5 which ended in a low frequency, possibly explaining why it was not sorted near other 3Q melodies.

Group 1H2Q. Melodies 7, 8, 10 and 12 were sorted together. As for 3Q melodies, similarity was based on a change in amplitude and frequency that ended in a high or consistent selection of each. One outlier, 11, had a low amplitude in the middle of the melody giving it a more 'rolling' feeling than the others. The other, 9, was the only one in this group to begin with a high amplitude and frequency and end with a low value of each.

Group 2Q. All 2Q melodies except 17 were sorted together. Their simple, consistent rhythm made them easy to distinguish from other melodies while feeling similar to each other. Melodies 13 – 16 had 2 quarter notes. Melody 17 had 1 quarter then 2 eighth notes, and 18 had 2 eighth and then 1 quarter note. 17 probably felt the most dissimilar due to its 2 concluding eighth notes which were emphasized by their high amplitude. This accentuated the different number of notes between this melody and the others in group 2Q.

Group 4-6E. Melodies 20, 22, and 24 were most closely grouped. Most of these stimuli felt 'rolling' as they had either no or low amplitude notes in their middle. Even though two melodies ended in a low amplitude and/or frequency, the 'rolling' characteristic was strong enough for most participants to group them together.

Group S. Participants did not group the syncopated melodies together. Although the sequence of notes was consistent, we suspect that the unevenness of the melodies coupled with widely varying amplitudes and frequencies made them feel dissimilar.

Group HLA. This group, based on the concept of descending amplitude, had widely varying rhythms while amplitude, frequency, and fullness were held constant. No melodies in this group were found near one another, suggesting that rhythm and note type may be a more groupable characteristic.

5.5.2 Questionnaire Results

The post-group questionnaires indicated that sorting the melodies into 12 groups was the most difficult (average=1.7, on a scale of 1-5 with 1 being most difficult) and 3 groups easiest (average=2.6).

The results for the post-session data revealed the average difficulty overall to be 1.9. 3/7 participants thought 6 was the best number of groupings, 3/7 preferred 3, and 1/7 preferred 9 groups. All participants indicated that they correlated stimuli based on the rhythm or tempo of notes. 3 said that amplitude similarity played a factor, and 3 participants also sorted based on frequency or intensity of notes.

Although demographic information such as musical experience and country of origin was collected, due to the small number of participants, it was not feasible to analyze the effect of these variables on the task.

5.5.3 Preliminary Findings

From this study, we summarize several observations of stimulus grouping:

- Changes in both frequency and amplitude were perceived qualitatively as changes in *intensity* – consistent with [5].
- A high amplitude note followed by a low amplitude note of the same frequency is perceived as 'rolling' or 'pulse-y' – in contrast to 'hard' or 'abrupt' sequence. Such notes feel continuous and seem to 'roll' into one another rather than feeling distinct.
- Low amplitude sustained notes surrounded by *staccato* notes are often grouped with those that have rests in the same position as the sustained notes.
- Syncopated melodies are not grouped together.
- Melodies consistent in amplitude composition but varying in rhythm and note density are not grouped.

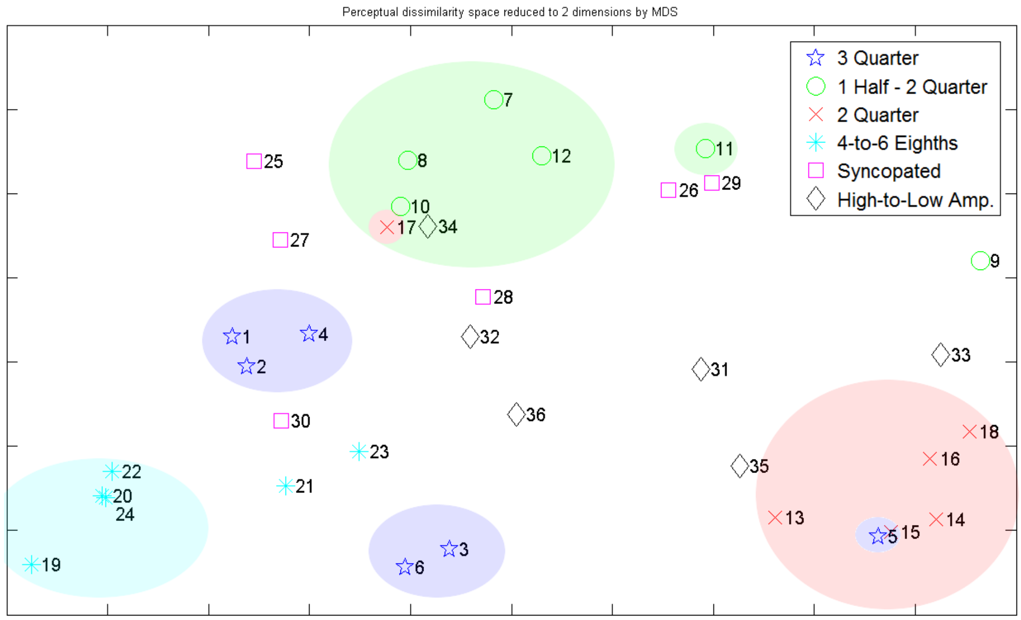

Figure 5: Visual representation of the icon sortings for all group sizes. The exact nature of the dimensions are unclear and complex. However, the horizontal axis seems to reflect changes mostly in fullness and the vertical axis seems to reflect changes in note length.

- Stimuli are generally not grouped by amplitude or frequency.
- The most groupable stimulus characteristics are note density and rhythm, as long as the rhythms are not syncopated.

With these observations in mind, we redesigned the stimulus set (Section 6) and conducted a new study with the goal of formulating more definite design heuristics.

6 Icon Redesign

To create a better set of stimuli and learn more about how melodies are perceptually grouped, we develop a new set of melodies by iterating on the previous set with our new information.

First, we removed the two concept-based groups. We then modified the anomalous melodies in the other four groups to make them more perceptually uniform (e.g. making Group 1 more 'abrupt' rather than 'rolling' by avoiding low-then-high amplitude notes on the same frequency).

Finally, we created two new groups in an aim to learn more about potential design heuristics. Both were derived from qualitative feedback from the first study. For the first new group [**All eighths (AllE)**], the entire bar is filled with eighth notes since dense stimuli were perceived as urgent and we wanted to maximize this characteristic for discriminability. The second group was created to feel like a horse 'gallop' [**Gallop (G)**] since some previous participants associated stimuli with patterns of locomotion ('walking', 'limping', 'running', etc.)

Figure 6 illustrates representative stimuli for these six groups.

7 Study 2

Hypotheses, apparatus, task, design and procedure are identical to Study 1 (Section 5). Six Computer Science graduate students (5 male) participated in our second study. None of the participants were involved in Study 1.

7.1 Results

Inspection of the MDS visualization from the second study (Figure 7) show that the melodies of group AllE, with one exception, were highly distinguished from all other melodies.

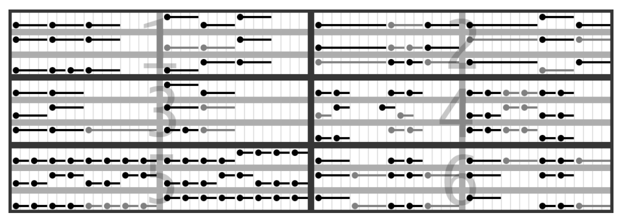

Figure 6: All 36 melodies from the re-designed six groups. Grey notes are low amplitude, raised notes are high frequency. Stimuli are listed left-to-right and top-to-bottom.

A closer inspection – without the dominating AllE group – was desired to analyze the results more thoroughly. Pasquero *et al.* [21] state that hidden patterns within the data can be uncovered by performing MDS on sub-matrices of the original dissimilarity matrix. A visual representation of the results excluding group AllE was produced (Figure 8).

7.1.1 Qualitative MDS Analysis

A detailed qualitative analysis was again carried out on the stimulus clustering in the MDS perceptual visualization.

Group 3Q. Every melody in this group was sorted together except for 5, which was the only one with 2 emphasized eighth notes in place of the center quarter note.

Group 1H2Q. All melodies were grouped together.

Group 2Q. All melodies except for 17 were sorted together. 17 was the only one with a half note taking up the whole last half of the bar. While designing this group we predicted that the half note would feel like a rest, but here it clearly did not. Instead, this rhythm (2 quarter notes followed by the single half note) was more often sorted with notes of group 1H2Q (one half note followed by 2 quarter notes); stimulus 17 was distinct from the stimuli from group 1H2Q only by *phase*.

Group 4-6E. The melodies in this and group G were sorted together with little distinction between the two. Melody 23 is close to group 2Q in these results because, we suspect, the two sets of

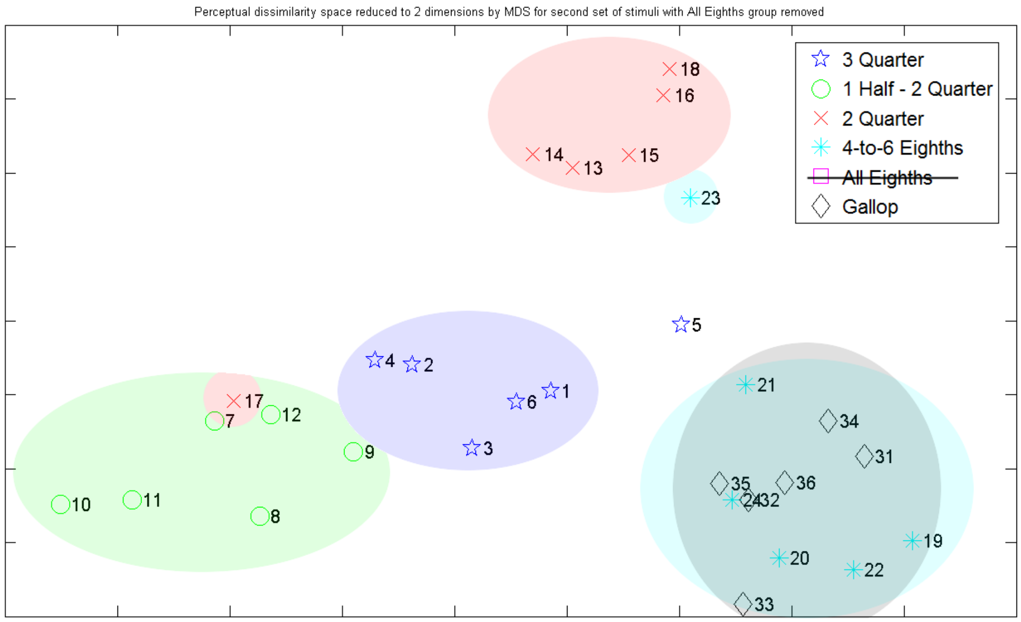

Figure 8: Subset MDS map stimulus sortings for groups 3Q, 1H2Q, 2Q, 4-6E and G, to better illustrate mutual differences between groups. The horizontal axis seems to reflect changes mostly in note length and the vertical axis seems to reflect changes in fullness.

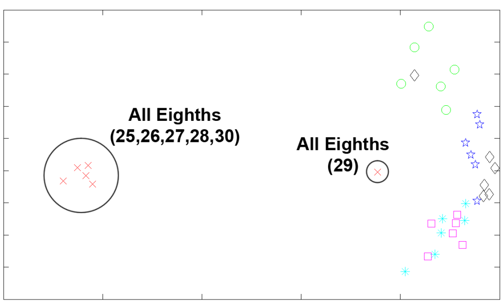

Figure 7: Thumbnail of MDS map of the stimulus sortings for all group sizes showing group position. 5/6 stimuli in group AllE are displayed in the left of the diagram, while all other stimuli are displayed on the right.

2 eighth notes with a rest in between feels like just 2 notes (the defining rhythm of group 2Q *was* just 2 notes). This suggests that participants were sensitive to the number of notes in a melody.

Group AllE. This group was nearly always sorted together by every participant. The rhythm for this group was very distinctive likely due to the saturation of notes in the bar and its continuous buzzing feeling. Melody 29 stands alone since it felt ‘rolling’ as opposed to ‘abrupt’.

Group G. As mentioned above, these melodies were indistinguishable from group 4-6E. We suspect that the inclusion of eighth notes in the second half of the rhythm may have resulted in one of two things: groups consisting of rhythms with eighth notes are perceptually salient, lending them to be sorted with other groups consisting of such notes; or the inclusion of eighth notes in a rhythm makes it perceptually too complex to distinguish among its different melodies, therefore if more than one group is considered complex their melodies are sorted together.

7.1.2 Questionnaire Results

The post-group questionnaires for this second study indicated that, like the first, sorting the melodies into 12 groups was the most difficult (average=2.0), 6 and 9 groups were the same (average=2.7), and 3 groups was the easiest (average=4.0).

Overall the average difficulty was reduced to 2.7 from 1.9 in the initial study. 4/6 participants thought that 6 was the best number of groups while 2/6 thought 3 was best. The strategies to sort melodies in this study centered around note rhythm. Only one user said amplitude played a factor and there were no indications that frequency or intensity of notes was used as a grouping strategy. Indeed, the number of notes, rhythm, pacing, and unevenness [25] were all considerations used by participants to sort the melodies.

As in study 1, demographic information was not considered.

8 Discussion

Generalizable results can be inferred from the detailed analyses of the two studies. Although the analysis is primarily visual, and therefore subjective, this work is exploratory in nature and geared toward design guidance and mapping out demonstrably fruitful areas for future work. In particular, the following trends appeared to dominate participants’ perception of melodic haptic stimuli:

- *Rhythmic differences between melodies dominate other distinctions*. Our results show that participants consistently grouped stimuli based on non-syncopated rhythm, even if such stimuli have different amplitudes and/or frequencies (H1, H2).
- *Perceived quantity of notes is a major grouping factor*. Despite variance in rhythm, amplitude and frequency, melodies which contain a similar number of notes were often grouped together.
- *Groups of rapid (eighth) notes are perceptually salient in rhythms.* Distinct groups in our second study that contained eighth notes were often confused by participants. The authors conjecture that the complexity of these rhythms makes it difficult to identify mutual differences, leading to confusion. This supports Newman’s claim that simple haptic stimuli are preferable [19]. The rhythmic domination is also supported as complex melodies contain quantities of notes that could be difficult to count.

- *'Abrupt' melodies are perceptually segregated from 'rolling' melodies even if they are devised from the same rhythm.* The holistic feeling of a note depends on its internal composition. Alternating between high and low amplitude notes on the same frequency makes a stimulus feel 'rolling' (Section 5.5.1).
- *Items that only differ in phase are grouped together.* The rhythm of a stimulus is perceived rather than its onset.
- *In some circumstances, replacing a quarter note with two eighth notes can increase expressiveness while maintaining groupability.* When we replaced a non-emphasized quarter note with two eighth notes, it was often grouped with its counterpart. This is not true for the case where the emphasis is on the two eighth notes.
- *Stimuli that only vary by frequency and amplitude are perceptually discriminable.* These stimuli were often, but not always grouped together (no occluding points in MDS plot). This indicates that people can perceive differences between these stimuli. This lends evidence to H3, however the exact degree is not conclusive and further investigation is required.

These trends suggest the following design heuristics for family-based design of melodic haptic icons:

1. **Rhythm primary, amplitude and frequency secondary.** Group stimuli based on simple and distinctive non-syncopated rhythms, and modify the amplitude and frequency of individual notes for within-group variation.
2. **Abruptness**. Design groups to be either *abrupt* or *rolling*.
3. Be mindful of **periodicity** when dealing with rhythmic stimuli.
4. As long as they are not emphasized, **quarter notes can be replaced with two eighth notes** for within-group variation.

This work provides the community with guidance for developing melodic haptic stimuli, as well as a perceptually validated, groupable, and expressive set of haptic stimuli for which learnability can now be determined.

9 Conclusion and Future Work

In this paper we presented a process for the design of melodic haptic stimuli with the goal of making haptic stimulus sets larger, more expressive, and more learnable. To assess how the designed stimuli were perceptually grouped, we described two iterative user studies.

From these studies, we suggest several design heuristics for family-based icon design. Stimuli from the same family should have a common, non-syncopated rhythm and should be perceptually similar in terms of abruptness (versus being 'rolling'). Designers should be mindful of periodicity and can use un-emphasized groups of eighth notes to expand the expressiveness of their melody.

Despite these findings, the learnability and the exact degree of within-group discriminability of melodic haptic stimuli has not been evaluated. Future work will examine the learnability of these stimuli by taking advantage of their groupable nature. Performance during a learning task will provide insight into within-group discriminability since learnability relies on a unique and distinct percept. This work should also attempt to take advantage of the added expressive power of melody by determining how to assign meanings as metaphorically as possible in order to facilitate learning.

References

[1] M. Blattner, D. Sumikawa, and R. Greenberg. Earcons and icons: Their structure and common design principles. *Human-Computer Interaction*, 4(1):11–44, 1989.

[2] S. Brewster, P. Wright, and A. Edwards. Experimentally derived guidelines for the creation of earcons. In *HCI'95*, 1995.

[3] L. Brown, S. Brewster, and H. Purchase. A first investigation into the effectiveness of tactons. In *1st Worldhaptics Conference (WHC 2005)*, pages 167–176, Pisa, Italy, 2005.

[4] L. M. Brown. *Tactons: Structured Vibrotactile Messages for Non-Visual Information Display*. PhD thesis, University of Glasgow, April 2007.

[5] A. Chan, K. MacLean, and J. McGrenere. Learning and identifying haptic icons under workload. In *WHC*, pages 432–439, 2005.

[6] A. Chan, K. MacLean, and J. McGrenere. Designing haptic icons to support collaborative turn taking. *Int'l J Human Computer Studies*, 66:432–439, 2008.

[7] T. W. Deacon. *The Symbolic Species*. W.W. Norton & Company Ltd., 1997.

[8] M. Enriquez, K. MacLean, and C. Chita. Haptic phonemes: Basic building blocks of haptic communication. In *8th Int'l Conference on Multimodal Interfaces*, pages 302–309, 2006.

[9] W. Gaver. Auditory icons: Using sound in computer interfaces. *Human-Computer Interaction*, 2:167–177, 1986.

[10] K. Hevner. The affective character of major and minor modes in music. *American Journal of Psychology*, 47:103–118, 1935.

[11] E. E. Hoggan and S. A. Brewster. Crossmodal icons for information display. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 857–862, 2006.

[12] M. Hollins, R. Faldowski, S. Rao, and F. Young. Perceptual dimensions of tactile surface textures: A multidimensional scaling analysis. *Perception and Psychophysics*, 54:697–705, 1993.

[13] T. Kaaresoja, L. Brown, and J. Linjama. Snap-Crackle-Pop: Tactile Feedback for Mobile Touch Screens. *Proceedings of Eurohaptics 2006*, pages 565–566, 2006.

[14] S. Lederman. Auditory texture perception. *Perception*, 8(1):93–103, 1979.

[15] J. Luk, J. Pasquero, S. Little, K. MacLean, V. Levesque, and V. Hayward. A role for haptics in mobile interaction: Initial design using a handheld tactile display prototype. In *ACM Conference on Human Factors in Computing Systems*, Montreal, Canada, 2006.

[16] K. MacLean. Foundations of transparency in tactile information design. *IEEE Transactions on Haptics*, page 12, 2008.

[17] K. MacLean and M. Enriquez. Perceptual design of haptic icons. In *Eurohaptics*, Dublin, UK, July 2003.

[18] G. Miller. The magical number seven, plus or minus two: some limits on our capability for processing information. *Psychological Review*, 63:81–97, 1956.

[19] S. E. Newman, A. D. Hall, D. J. Foster, and V. Gupta. Learning as a function of haptic discriminability among items. *The American Journal of Psychology*, 97(3):359–372, 1984.

[20] T. Nojima and K. Funabiki. Cockpit display using tactile sensation. In *Eurohaptics*, pages 501–502, 2005.

[21] J. Pasquero, J. Luk, S. Little, and K. MacLean. Perceptual analysis of haptic icons: an investigation into the validity of cluster sorted mds. In *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (IEEE-VR)*, 2006.

[22] I. Peretz, L. Gagnon, and B. Bouchard. Music and emotion: perceptual determinants, immediacy and isolation after brain damage. *Cognition*, 68:111–141, 1998.

[23] M. Rigg. The mood effects of music: A comparison of data from four investigators. *Journal of Psychology*, 58:427–438, 1964.

[24] D. Ternes. Building large sets of haptic icons: Rhythm as a design parameter, and between-subjects mds for evaluation. Master's thesis, University of British Columbia, 2007.

[25] D. Ternes and K. MacLean. Designing large sets of haptic icons with rhythm. In *Eurohaptics*, pages 199–208, 2008.

[26] J. van Erp and M. Spape. Distilling the underlying dimensions of tactile melodies. In *Eurohaptics*, pages 111–120, 2003.

[27] M. Weiser and J. Brown. Designing calm technology. *Powergrid Journal*, 1.01:94–110, 1996.