The Role of Choice in Longitudinal Recall of Meaningful Tactile Signals

Mario Enriquez* and Karon MacLean[†]

Department of Computer Science, University of British Columbia

ABSTRACT

Haptic icons (brief tactile stimuli with associated meanings) have the potential to convey abstract information through touch; however, there has been little systematic investigation of how sets of perceptually distinct tactile signals can be best utilized to convey meanings, nor of how enduring these associations can be.

We hypothesized that when users can choose the signals which will represent specific concepts, their learning and recall will be eased and enhanced. Taking future embedded interfaces as context, we used two sets of 10 distinct tactile signals to compare recall of concept-meaning associations in two conditions: (1) arbitrarily assigned and (2) participant-chosen associations. Participants learned associations in under 20 minutes at 80% accuracy; at 2 weeks, recall of the associations previously learned was 86% with no significant effect of assignment condition. Subjective confidence levels sharply lagged actual performance, with zero expectation of ability to recall at 2 weeks.

To the best of our knowledge, this is the first study of either longitudinal recall or the role of user choice on synthesized stimulus-meaning learnability. Its results underscore the eminent practicality of using haptic icons in everyday interface design, suggesting high learnability and a surprising user ability to find their own mnemonics for carefully composed stimuli, regardless of how associations are assigned.

CR Categories and Subject Descriptors: H.5.2 User Interfaces: Haptic I/O.

Additional Keywords: Design, Experimentation, Human Factors, Languages.

1 INTRODUCTION

Sensory overload is a common problem with contemporary user interfaces, particularly for those that connect users to computation embedded in portable devices and non-desktop environments. These are increasingly pervasive, often have complex functionality, and are frequently used in contexts which pose multiple demands on a single sensory modality: e.g., accessing a cell phone while driving or in a theatre, or using a visually-dependent remote control while watching TV in a darkened room.

The complexity of the intelligent aids we currently use in our daily tasks contributes to sensory overload. For example, the display of most cellular telephones conveys information far beyond basic "caller ID", which means that its operation demands the visual sense for longer continuous periods of time exactly when visual attention is most fragmented [26]. Interface designers increasingly must struggle to simultaneously address the demand for more functionality and keep device size small.

On the assumption that there is some degree of modularity in attentional processing and that using a different sensory channel for communication can reduce interference with critical visual tasks [1, 12, 36], one possible approach to these constraints is to divert some information through the touch sense. Multiple-resource theory asserts that individuals can acquire information via multiple modes of sensory input; and further predicts that two simultaneous inputs are less likely to interfere with one another if the information is delivered via different rather than a single sensory modality [36]. MRT also models other kinds of interference in a multitasking environment, e.g. competition for a user's cognitive or linguistic resources. Interface designers must continue to develop sophistication in understanding the user's *full* context. However, addressing the currently overwhelming state of purely sensory overload is a good place to start.

To this end, we aim to convey information through the sense of touch in the simplest and most transparent fashion possible. We build upon the concept of *haptic icons*: brief tactile stimuli that have been associated with a meaning. Haptic icons present a new means of displaying information to people that can be discrete, convenient and informative, meanwhile decreasing the dependence on the visual and auditory channels and minimizing overall perceptual interference.

Haptic icons should be generally useful in contexts where the visual and auditory channels are already in use, but where information is ideally accessed in "layers" defined by type, criticality or momentary context. To begin with, they will soon be integrated into handheld devices, where they are expected to (at minimum) reduce frequency of scanning between device and environment, a considerable advantage in busy environments, or provide a private conduit in social situations.

At the lowest level, haptic icons can be used by devices and objects to notify users of an event, their identity or their current state or contents. Simplistic signals such as pager vibrators have existed for years. However, we argue that these binary or amplitude-graded signals contain far less information than is possible with systematic, perceptually guided design, and that humans are more able to quickly learn and remember them than they expect. In a future where icons are better designed and users have developed tactile acuity and familiarity with the communication concept, we anticipate that they will support expressive and nuanced communication.

The obvious concerns about this unfamiliar medium relate to (a) maximum recognizable size of stimulus set; (b) learnability of stimulus-meaning associations; and (c) longevity of those associations. All of these factors are exacerbated by the currently limited expressivity of tactile displays, particularly those suitable for embedded, mobile and wearable applications.

Information density of stimulus sets: The experiment reported here uses a modest set size of 10, and as such does not directly address the challenge of large sets. However, in other work we have demonstrated a distinguishable set which meets the limits of foreseeable human associative limits, with novice users, a moderately expressive display and close attention to stimulus design [33]. This promising situation will improve with better devices and an experienced user base.

We also point out that icons can carry varying amounts of information. Information transfer can be increased by increasing set size (while ensuring the set remains identifiable), or alternatively by increasing the perceivable information content of

*e-mail: enriquez@cs.ubc.ca
[†]e-mail: maclean@cs.ubc.ca

Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems 2008
13-14 March, Reno, Nevada, USA
978-1-4244-2005-6/08/$25.00 ©2008 IEEE

individual signals; or possibly both. The best approach depends on the need of the application. For both, the threshold of perceivability is modulated by characteristics of the anticipated operating environment, and complex icons might be harder to recognize under workload. Design heuristics and techniques for this aspect of icon construction is emerging in other work (e.g. [8, 9]), but is not the goal of the present project.

Models for creating and learning associations: As discussed below, common approaches to creating stimulus-meaning associations are *abstract* or *semantic*, with the first offering the ability to control, optimize and thus maximize the size of a usable stimulus set, while the latter seems as if it would be easier to learn and remember. These two approaches have not, however, been compared for efficacy side-by-side. To do so, the difference between them must be considered more broadly: we suggest that at issue is not whether the *designer* perceives a semantic association, but whether the *learner* does. Through our past experience with designing and deploying haptic icons, we have noticed that users often have personal opinions about appropriate associations which the designer cannot possibly predict. We therefore hypothesized that the safest way of supporting semantic associations is to let the user "roll their own."

Longevity of associations: The final and likely the most critical prerequisite for success is the potential for an enduring association. Once learned in an initial session, will they persist without reinforcement after an interval of time? For how long? Do different associative and learning mechanisms influence success?

The experiment reported here is a first effort to shed light on the last two questions.

1.1 Overview

The experiment presented here was designed to test the hypothesis that allowing participants to choose their own stimulus-meaning associations would, by permitting leverage of their own implicit mental models, improve various subjective and objective metrics relating to learning and retention of those associations. Furthermore, this experiment examines the degree to which users retain learned meanings after a two-week interval, without the benefit of any interim reinforcement.

Our typical practice for supporting users' learning of stimulus-association meanings, regardless of the mechanism used by the designer to build the set, has been an iterated reinforced learn-test cycle (e.g. [9, 16]). In the present experiment, we used this approach in a first condition termed *arbitrary associations*, using pre-assigned, randomly chosen stimulus-meaning matches for a set of 10 meanings. We also tested a second condition termed *user-chosen*, wherein users selected associations for the same 10 meanings from a set of 20 perceptually differentiable haptic stimuli. In both conditions, we otherwise used the same methodology for reinforcement learning, subsequent recall testing and eliciting subjective responses. To broaden the test's external validity, we evaluated these two conditions in two simulated interface scenarios: a *hand-held navigation unit* and an *automobile radio control*.

The 20-element stimulus set we used was constructed by varying rhythm patterns of two-second duration, presented at different amplitudes through a tactile display.

Our results suggest that subjects are able to learn and later remember stimulus-meaning associations after a brief learning period at 80% recall. Furthermore, association persistence at 2 weeks after the learning period was 86% of the originally learned associations. We saw no significant difference in average performance between arbitrary and user-chosen associations. Interestingly, many participants reported that they believed the arbitrarily-chosen associations had been designed with metaphorical intent; i.e. they discovered their own mnemonic associations, and this could explain the undifferentiated result. Post-experiment interviews also revealed that participant's expectations and confidence levels for their actual performance sharply lagged their actual performance. None believed that they could recall more than a few associations learned before.

These findings have important implications for the design of interfaces intended to communicate information through touch: they underscore the eminent feasibility of using haptic icons in everyday interface design, suggesting high learnability independent of designer assignment mechanism. They also improve our understanding of the cognitive steps employed by users in their learning process. A full grasp of this process is essential if we are to maximize this channel's potential, by designing icons optimally and supporting users in learning them.

2 Related Work

Past work which relates to the design of meaningful tactile signals, particularly for use in high-workload environments, include foundations for tactile perception and attentional processing, design and discrimination of "raw" (unassociated) haptic stimulus sets, embedded haptic feedback in high workload environments, and a catalog of specific projects employing haptics for abstract information display using either abstract or semantic approaches. We are not aware of any past longitudinal studies, or comparisons of different icon construction methods.

2.1 Tactile Perception and Attentional Processing

Over the years, we have gotten a better understanding of our tactile psychophysical capabilities through the studies which have documented our exquisite sensitivity to, e.g., texture felt through a probe [23]. Tan *et al.* [30] have measured information transfer rates of 2-3 bits/second for haptic stimuli presented to three fingers of a participant's hand, showing that appreciable content can be conveyed through this channel. Currently, the psychophysical research of greatest immediate relevance to tactile signaling relates to thresholds for resolving different excitation parameters (e.g. [23, 25]) and both temporal and spatial masking effects (e.g. [13, 19, 29]). The values thus determined are heuristically useful for avoiding conflicts in first-pass stimulus prototyping. However, it is difficult to predict how parameters will be perceived when used together, and further, how users will organize multidimensional stimuli within a group. The testing mechanisms described below therefore remain essential until our psychophysical and cognitive sophistication greatly improves.

As discussed earlier, some currently recognized attentional theories support the approach of "offloading" information display onto the haptic channel [36]. Other attentional research demonstrates *linkages* between vision and haptics (e.g. showing that haptic stimuli can be used to orient a user's attention in another sensory modality by using taps on the back to direct gaze [31]. The latter suggests ways that haptics and vision can be used synergistically in high-load environments, but could also undermine the idea that different modalities will not interfere with one another's processing. Further work in this area is needed to better understand the perceptual processing "pipeline".

2.2 Designing and Validating Stimulus Sets

Prerequisite to usable haptic signals are perceived distinctiveness and structural richness in the stimulus set. If stimuli feel too similar or vary along too few dimensions, users cannot create long-lasting associations to them (imagine a graphic icons set of 20 blocks, each a different shade of gray). Given that today's tactile displays can rarely be controlled in more than three dimensions if that, these dimensions must be exploited with care.

The only such evaluative mechanism we are aware of is based on perceptual Multidimensional Scaling (MDS) to extract perceptual axes for complex synthetic haptic icons ([25], and also used in [35]). The hardest part about this use of MDS is efficiently collecting high-quality difference data from users for relatively large stimulus sets. [25] employs an efficient cluster-sorting technique for this purpose, showing that a 36-item stimulus set constructed by varying frequency, magnitude and shape of 2-sec, time-invariant haptic wave shapes map to two perceptual axes. It suggests that expressive capability is maximized in one frequency subspace (7-25 Hz) for that particular force-feedback knob.

By comparison, MDS applied to a wide range of real stimuli suggest up to four perceived dimensions, some highly complex [3]. This might represent an upper limit for synthetic stimuli given improved display technology; alternately, synthetic approaches might enable designers to create new dimensions not present in natural sensations, and exceed this.

Brown *et al.* [6] created 27 abstract tactile signals by varying rhythm, roughness and spatial location, based on prior indications of 3 differentiable levels for each of the 3 parameters individually. The design of this set did not consider parameters interactions, potentially explaining low recall performance (below).

Most recently, Ternes and MacLean have devised a variant of the MDS methodology to handle larger sets, demonstrating its use on a set of 84 stimuli [33]. This set was created through a careful analysis of rhythm perception, used with frequency and amplitude [34]; new perceived sub-dimensions of rhythm were revealed.

2.3 Haptic Cues in Distracting Environments

Examples of simple haptic signaling can be found in commercial products. Some cell phones use distinctive vibration patterns [22]: a customized vibration can transmit e.g. caller identity with less intrusiveness than a ring tone.

In 2001, BMW introduced automotive haptic feedback with the iDrive™, a force-feedback knob for accessing secondary vehicle functions such as audio and climate-control [21] by mapping them to different programmed feels, i.e. contextual navigation cues and fixtures. The product initially struggled due to strangeness, poor usability and steep learning curve [10], but design iterations and developed user experience met with increased appreciation.

In the handheld domain, tactile feedback has similarly added context and cues for application navigation [24, 27], building our knowledge how mobile activities can benefit from this modality.

2.4 Semantic vs. Abstract Icons, and Stress Testing

Semantic icons represent objects or notions through a literal, direct symbol: for example, using the sound of a paper being crushed to indicate deleting a computer file. Gaver *et al.* [17, 18] defined “Auditory Icons” as auditory representations of real objects and actions. The proposed advantage of using a semantic presentation is intuitiveness. Conversely, abstract approaches are similar to the auditory model used by Brewster *et al.* [4]: “Earcons” are sounds and rhythms with no intrinsic or cultural meaning; their target or meaning has to be learned to be effective.

In a first instance of applying a semantic approach to tactile information display, Chan *et al.* [8, 9] developed 7 haptic icons to facilitate application sharing among distributed members of a group, by indicating request urgency in a custom turn-taking protocol. In both abstracted and situated environments, they found that the haptic icons (designed to be intuitive) could be learned to a high degree of accuracy in under three minutes and remained identifiable even under significant cognitive workload. The associations used were carefully explained to the users prior to the test.

Tang *et al.* [32] prototype a representational numerosity display and test it under visual overload. Their experiment shows that people can perceive and accurately process haptically rendered ordinal data while under cognitive workload, with accuracy ranging from 75-93% depending on representational model.

In [7], Brown *et al.* used a magnitude representation for 9 signals composed from 3 levels each of “roughness” and pager motor intensity, to indicate respectively 3 cellphone message types and 3 priority levels. The idea was that the different intensity levels would intuitively represent different urgency levels. They found a recognition rate of 52% for roughness and 70% for intensity level.

Several examples lie mid-way along the abstraction spectrum. Van Erp *et al.* propose that familiarity with tactile rhythms drawn from popular music will aid in recalling concepts arbitrarily assigned to these tactile patterns; i.e. the abstract association would benefit from more memorable stimuli [35].

Allen *et al*. conducted an exploratory study to measure usability of eyes-free music playlist navigation based on symbolic haptic renderings of key song parameters. Users learned force-feedback mappings for music parameters to usable accuracy with 4 minutes of training. Song tempo was modeled by detent frequency, and a measure of a song’s energy level by detent amplitude [2].

In Brown *et al.*’s 27-item rhythm-based set, participants were asked to identify concepts associated to the 3 different levels of the 3 different parameters used for each presentation of the stimuli [6]. results showed an overall identification rate of 47.8% (chance performance would be 30% correct).

One study exhibits thoroughly abstract links. Enriquez *et al.* demonstrated an encouraging ability of users to learn deliberately arbitrary associations for a family-organized set of 9 icons presented through a haptic knob, with 77% average recall performance (chance = 33%) [16].

The idea of creating sets of intuitive representational haptic icons which do not require their users to learn their meaning is very attractive. Intuitive *visual* icons can be found everywhere, from computer desktops to signs in nature parks. However, creating these intuitive representational tactile icons is not an easy task. We are limited by both current display technologies as well as a poor understanding of how tactile signals are perceived and processed. It is for this reason that we decided to investigate whether users could build enduring mnemonic representations for abstract synthetic tactile signals.

3 Approach

We have previously tested people’s ability to learn and remember both semantic and deliberately arbitrary stimulus-meaning associations and found encouraging results for both. Here, we sought to compare these two cases directly, and to use longitudinal recall as one important new metric.

To test this proposition, we needed a mechanism for learning the associations as well as a performance measure indicating how well the user is able to perceive and process the signals and relate them to their respective meaning

Thus, we ran this experiment in two sequential sessions separated by two weeks, using the same participants in both. In one of his/her sessions, the participant chose which tactile signals would represent the set of 10 meanings for that session. In the other, he or she were given arbitrary tactile signal-meaning associations for a different 10-meaning set. The order of these sessions was randomized among subjects. The two sets of meanings were drawn from two embedded contexts and counter-balanced with individual and condition order.

Each session had three stages (Figure 2). The first was self-guided learning, where users used a graphic interface that showcased a set of tactile stimuli and allowed them to learn associations between them and their meanings [9, 16]. This was followed by a computer-guided reinforced learning stage, where

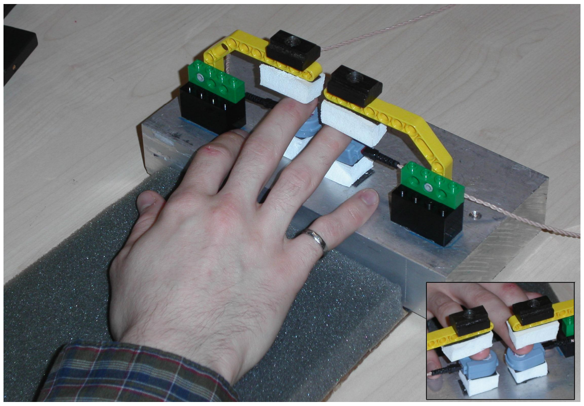

Figure 1: Participants placed the index finger of their non-dominant hand on a tactile display

users were asked to identify a series of randomly presented tactile signals and drag them into boxes labeled with their respective meanings, while receiving feedback about errors. In the final test stage, users again performed this learning task but without error reinforcement.

In this experiment, we use objective measures of recall performance for both arbitrary and user selected associations (both immediately after training and 2 weeks after) as well as subjective measures of participant opinion regarding task difficulty and confidence levels for the learned associations.

4 Methods

Experiment methods consist of display and setup, stimulus design, a 3-phase experiment protocol and recorded measures.

4.1 Tactile Display

Our experiments were carried out using a custom display integrating one Audiological Engineering (AE) tactile display (www.tactaid.com, visible in Figure 1). These voice-coil-based transducers, which are used commercially in hearing aids, are capable of producing precisely timed (on/off within 2 ms) waveforms at a useful range of frequencies and amplitudes, with maximum efficiency at 250 Hz; and can be driven directly by a computer's sound card. Tactile displays using similar technology can be found in commercially available mobile phones, PDA's and GPS navigation units.

4.2 Stimulus and Meaning Set Designs

The design for the tactile stimulus set used for this study took as a starting point the 84-element set of rhythm, frequency and amplitude mentioned above [34]. Although we only required 20 stimuli here, we felt it would be interesting to add additional textural diversity through the addition of some more naturalistic stimulus, i.e. more broad-spectrum than the single-frequency tones of the source set; prior work has shown that having richer timbre in auditory signals aids in recalling meanings associated with them [5]. We hypothesize that similarly, richer tactile signals would lead to higher recognition rates as well. Therefore, we chose 11 disparate stimuli from the 84-element rhythm set (drawn from various distinctive areas of that set's MDS map). These were complemented with 9 additional signals created manually by auditory recording of sounds through a microphone, such as taps on a microphone and scratching over a rough surface, to reach our target set size of 20.

Table 1. Simulated Interface Scenarios Used for Functions

GPS Navigation	Left, Right, Forward, Back, Up, Down, Faster, Slower, Stop, Go
Automotive Audio System	Volume, Balance, Bass, Treble, Mid Range, Fader, Mute, Tuner, CD, AM

Signal meanings were drawn from two scenarios representative of the type of multi-tasking, attentionally demanding contexts where we anticipate haptic icons will be most useful: a hand-held navigation unit, and an automobile radio control. 10 were used, a group size felt to represent a reasonably broad utility; Table 1 lists their specific values. We took this approach to avoid unattractive alternatives of meanings that were completely unrelated (unfeasible) or which could not be semantically organized by participants in some unexpected way – that is, we asserted the connections among the group, leaving it unambiguous.

4.3 Physical Setup and Instructions

The design of the apparatus (Figure 1) was driven by needs for consistent hand position and finger pressure, as well as vibration isolation to prevent crosstalk between the stimulus sites. It utilized one AE display mounted on a 3 cm thick aluminum plate and insulated with 1 cm thick latex foam rubber commonly used to mechanically isolate sensitive electronic equipment from vibration. The participants placed their non-dominant hand on another foam pad which was attached to the aluminum plate; weights mounted on articulated plastic arms held his/her index finger against the transducers with a constant pressure of 30 grams.

The tactile display was interfaced through the sound card in a 2.5 GHz Pentium 4 computer running Windows XP. To mask auditory noise from the haptic display, participants wore headphones and listened to white noise throughout the test session. They received graphical feedback from a Dell 22" LCD monitor positioned approximately 60 cm away, and made responses with a standard mouse and by typing on a keyboard in front of the monitor.

At the beginning of each experiment run, participants read instructions presented on the computer screen and were queried for questions. Following the three experiment stages, at the end of the experiment, participants were debriefed about their experience, and solicited for subjective reactions.

4.4 Protocol

Every participant carried out two sessions of the experiment (U:

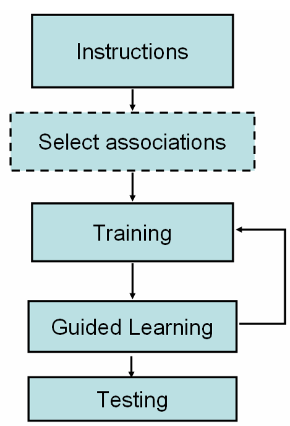

Figure 2: Experimental procedure

Table 2. Allocation of Experiment Run Types. Participants were assigned to one of the four experiment run types.

Session type	1st Session	2nd Session (+2 wks)
1	User-Chosen+Radio	Arbitrary + Nav
2	User-Chosen + Nav	Arbitrary + Radio
3	Arbitrary + Radio	User-Chosen + Nav
4	Arbitrary + Nav	Chosen + Audio

User-Chosen associations and A: *Arbitrary* associations), conducted 2 weeks apart. For each session, participants had to learn associations between 10 tactile signals and 10 meanings. In the *user-chosen* session, participants were given 10 meanings drawn from one of the two contextual scenarios used, and asked to choose their favored tactile signals from the full set of 20 tactile signals described above. For their *arbitrary* session, participants were presented with arbitrarily chosen associations between the set of 10 meanings and 10 of the 20 tactile signals (the same associations for all participants). The order of the two different sessions was counter-balanced among participants.

Our two conditions (*arbitrary* or *user-chosen* associations) and two meaning scenarios (*navigation* or *radio*) thus resulted in 4 participant types (Table 2).

For the session run second (regardless of type), participants began the session with a brief recall test to measure how well they could recall the tactile signals associated to meanings learned two weeks prior.

Each session had three phases: training, guided learning and testing (Figure 2). Participants were allowed to switch between the self-guided and reinforced learning stages as many times as they required until they decided to proceed to the testing phase. Once in the testing phase, they could not return to either the self guided or reinforced learning phases.

Training Phase: Using the GUI shown in Figure 3, participants could repeatedly click on each of 10 different buttons labeled with meanings and feel the corresponding tactile stimulus. Participants were allowed to return to the self-guided learning interface from the reinforced learning phase if desired.

Guided Learning Phase: In the Identification and Reinforcement views of the Guided Learning GUI (Figure 4), participants were presented with 10 labeled meaning boxes, purposefully ordered along a different dimension from Figure 3, along with 10 draggable stimulus tiles. A left mouse-click on a tile triggered playback of the corresponding stimulus. With a right mouse-click, the tile could be dragged into a meaning box.

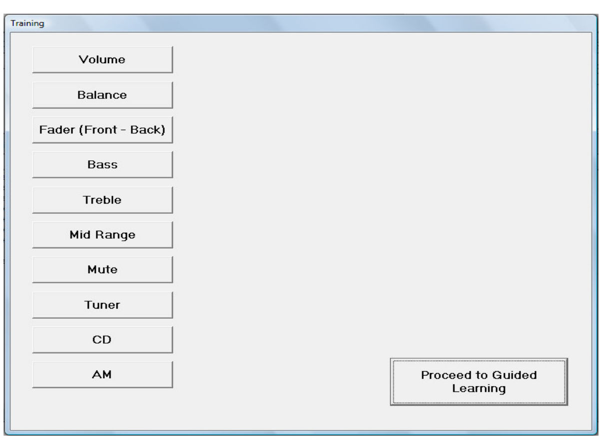

Figure 3: Participants learned pre-defined meanings for 10 haptic icons, by feeling and matching to targets

Figure 4: Guided Learning and Testing phase GUI. In training, participants tested their knowledge of the signals' meanings and received feedback as to their placement of the tactile signal tiles (green for correct placements, red for incorrect). Placement feedback was not given during the Test Phase.

Participants could feel a given tile's stimulus any number of times before placing it. To discourage participants from grouping tiles based on relative comparisons (as opposed to absolute recall of associations) a tile could not be moved or played once it had been placed. At the end of each trial, participants were given visual feedback regarding any errors (correct tiles turned green and incorrectly placed tiles turned red) and could return to the self-guided learning phase if they wished to do so. Once a participant had sorted all 10 tiles into their meaning boxes 3 times, they could proceed to the testing phase.

All participants proceeded to the test phase regardless of performance in the reinforced learning phase. Given that we intended to measure recall performance as a function of condition, we chose this approach rather than having participants train to a preset performance level.

Test Phase: A test trial was identical to the identification step of a Guided Learning trial, with the exception that participants did not receive feedback on their performance, which was recorded. The test phase consisted of 10 randomized trials. To minimize fatigue, a 5-minute break was enforced after Trial 5.

4.5 Measures

Our objective metrics were number and identity of tactile signals placed in each meaning box by experiment condition, and reinforced learning and test phase durations. We also measured the level of association recall from the first session after 2 weeks (before the second experiment run). That is, for n participants, we obtained $2n$ observations of immediate recall (n for each condition), and n observations of 2-week recall ($n/2$ for each condition).

To obtain subjective responses, we conducted an open-ended interview with each participant after each experiment run, in which we inquired about their thoughts about the tasks they had performed and the level of difficulty for the two different experimental conditions. Just before the short recall test at the start of the second session we asked participants how many of the associations they thought they could remember after 2 weeks. For the interview after the second and last run of the experiment, we also inquired whether choosing associations helped in learning them.

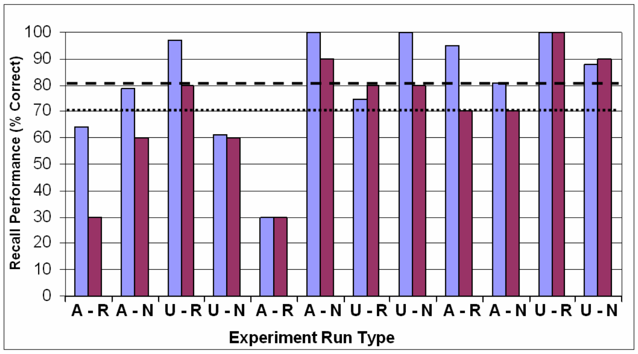

Figure 5: Longitudinal recall. Identification performance for first set of associations learned, by participant: immediately after learning (left blue-shaded bar) and as recalled 2 weeks later (right red-shaded bar). Bar labels indicate session type (U: *User Chosen* or A: *Arbitrary* associations; R: *Radio* or N: *Navigation* scenario). The dashed and dotted lines represent average performance immediately after learning and recall 2 weeks later respectively.

5 RESULTS

5.1 Participants

10 male and 2 female science graduate students were paid for their participation in the study (22-34 years, all right handed).

5.2 Identification Performance

Figure 5 shows the results for identification performance for the first session (regardless of session type) for each of the 12 participants, paired with recall of those same associations two weeks later (tested before the 2nd session with a new set). On average, participants correctly recalled 80.1% of the stimulus-concept associations immediately after this first learning period (left bar in each pair). Two weeks later, participants correctly recalled 70.1% of the same associations (right bar), or 86% of the associations recalled in the first test.

Figure 6 presents average recall performance immediately after learning for all sessions, grouped by simulated interface scenario type and user choice of associations. Figure 7 shows the same data broken apart more specifically A single-factor within-subjects ANOVA shows a statistically significant difference for scenario ($p<0.002$, $F=18.857$) and session order ($p<0.002$, $F=20.056$). The difference between User-Chosen and Arbitrary associations was not significant at $p<= 0.05$.

5.3 Subjective Opinions

5.3.1 General Comments after First Session

Following their first session, participants expressed that learning the associations was easier than they had expected. Some commented that the arbitrarily chosen associations were a good match to the functions in the interface (they did not learn that the associations were completely arbitrary until after the second session). One individual had trouble learning his first session's associations (arbitrary), stating that they "did not make sense". This participant had the lowest score for the first session's associations and their 2-week recall (both 30% correct). However, in the second session's associations, where this individual chose his own associations, his recall performance was 100%.

5.3.2 Subjective Estimate of Recall

At the start of the 2nd session, none of the participants believed that they could recall *all* the associations learned two weeks prior. Most participants predicted recall of zero associations; two expected they could recall 3 or more.

5.3.3 General Comments after Second Experiment Run

After the second session, participants were asked to comment on whether having a choice for the stimulus-function associations made a difference. Most indicated a preference for choice. One (the individual noted in Section 5.3.1) commented that choice was a necessity, given that the associations should make sense if you are to learn them. Only one participant believed that having a choice did not matter.

All participants were astonished to learn of their recall performance and commented that they did not believe their performance could be so good after such a short learning period. Furthermore, most participants reported that learning the associations was easier than they expected and that choosing the tactile signals themselves simplified learning even further.

6 DISCUSSION, CONCLUSIONS AND FUTURE WORK

Our primary research questions were:

1. Can a user learn and later recall a set of 10 concepts associated with haptic stimuli?
2. Do these associations endure after 2 weeks?
3. Are associations easier to recall when users choose the tactile stimuli-concept associations themselves?

6.1 Immediate and Longitudinal Performance

We found that both arbitrary and user-selected associations

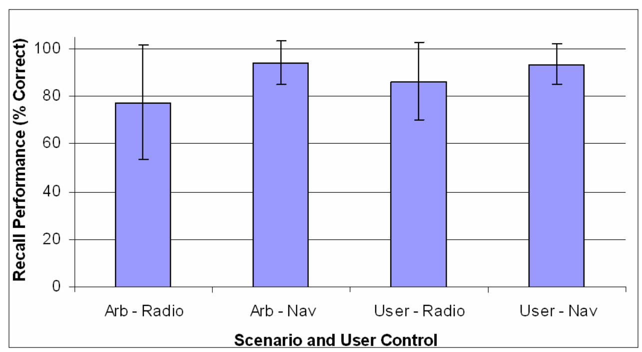

Figure 6: Identification performance immediately after learning, grouped by interface scenario and user control over associations. 24 sessions are represented (2 per participant; 8 values per bar).

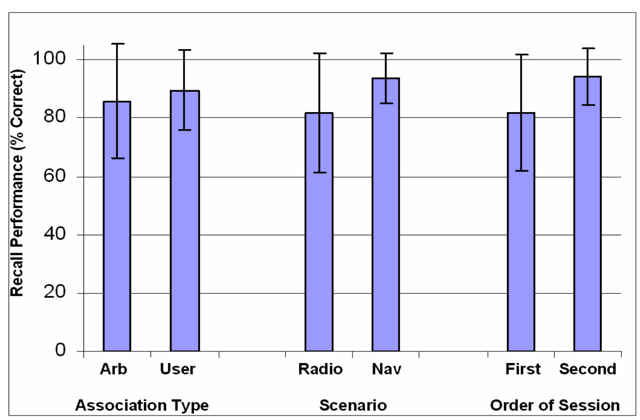

Figure 7: Performance grouped separately by: association type, scenario, and order of session (0 or 2 weeks). Performance improved significantly after a 2 week gap, and differed for the interface scenarios.

between tactile signals and meanings can be learned to a usable performance level (80%, where chance would be 9%)) after a 10-15 minute training period. This corresponds to 8/10, or a single switched assignment. Associations were remembered consistently over time (86% recall after 2 weeks of the initially learned associations) even without further reinforcement. To the best of our knowledge, this is the first report of long term recall of synthetic tactile signal associations.

These are promising results, which are consistent with positive answers to our first two research questions. While the evaluation described here was small, it was designed to be representative of realistic use contexts. If users can indeed learn the meanings of a set of tactile signals this quickly and persistently, haptic icons could be a practical method to present information such as device state or function identity in everyday contexts.

Further, we hypothesize that with regular and pervasive reinforcement, larger set sizes could be learned to better accuracy regardless of associative method. There are two basic stages of learning: acquisition (initial learning) and maintenance (repeated exposure to a single type of stimulus). The latter enables the organism to learn about stimulus properties and allows for longer term recall [20]. More work is needed to establish this here.

6.2 Learning and Performance by Association Method

Arbitrary associations were learned to the same level of recall as those with user-chosen associations, for the conditions present here. These results seem to show that the theoretical advantage of having the kind of metaphorical cue as presented by Chan *et al*. or more literal representation of the real world [18] was not necessary for the given haptic icons.

However, based on subjective responses it appears that participants often created their own metaphors for the arbitrarily composed sets; they were not dependant on a designer to build and explain the associations. The fact that they were *able* to do so might be one of the most interesting and unanticipated results of this experiment. It begs the immediate follow-up question of whether they would have been less able or willing to create their own mnemonics if they were explicitly told up front that the stimulus-meaning associations were ungrounded in any intentional meaning (at least some apparently believed that designer-created meanings existed).

Setting aside the tantalizing issue of *how* people are able to carry out what we thought would be a very difficult task: if borne out in more extensive evaluations on larger sets, the implication is greatly simplified icon design. The downside of a semantic approach to icon design is the difficulty of making it scale.

However, more intuitive associations, when available, can probably be learned more easily still; this seemed to be the view of our participants, regardless of their performance. Chan *et al*. found higher levels of success for haptic icons when users were given an explanation for their design and their related associations [9], although this was not a controlled condition. An interesting direction to explore is whether what we observed was a similar learning *performance* but at differing levels of *effort*; and whether, with more difficult learning tasks, this effort differential might translate to a measurable performance deficit.

6.3 Actual versus Subjective Recall Performance:

Participants did not believe they could recall stimulus-meaning associations at 2 weeks, and yet objective recall was only 10% (1/10 matches) less than immediately after training. A related effect has been observed in the vision literature and termed “mindsight”, where tested individuals guess that they have not seen a briefly displayed or masked stimulus, but when tested with an objective performance measure, their responses indicate a strong influence by the “invisible” stimulus [28].

Two areas of future investigation come to mind. On the cautionary side, it is conceivable that a lack of confidence could be detrimental in real usage; i.e. an untrusted warning might be dismissed; at least one study has suggested this with respect to warning signals [14]. Taking a more positive view, we wonder whether the disparity we observed is simply a matter of unfamiliarity; i.e. will trust in one’s “tactile intuition” come with regular, reinforced use?

6.4 Individual Differences

Even though results show that most participants could learn arbitrarily assigned associations, one individual struggled (30% recall, compared to 100% for 2nd-session user-chosen associations); stating that they “did not make sense to him”. Based on many similar anecdotal observations such as this, it would be unsurprising to find a wide and possibly bimodal distribution in human tactile acuity and/or higher level signal processing. And just as there are “visual” and “auditory” learners, perhaps some individuals will easily learn tactile associations in their own right, while some others will require a metaphorical reference to another modality to ground them.

6.5 Scale: Information Density and Larger Icon Sets

We also need to explore ways of increasing the amount of information that can be encoded in a single tactile information module. For example, the set of 20 tactile signals tested in this project could theoretically be used (through concatenation) to create a larger set of haptic phrases, each of which could convey more complex meanings and perhaps open the way for the development of a far-reaching haptic language. What is the limit to the complexity that can be perceptually and cognitively decoded from tactile messages?

To be broadly useful, haptic icon set size must be somewhat scalable, and certainly larger than 20 items; we have suggested that discernable sets of 75-100 with today’s tactile hardware will allow us to concentrate on cognitive bottlenecks while awaiting hardware improvements [34]. This raises two questions: Can participants remember associations for a “large” set of haptic icons, and how do we systematically design both the stimulus sets and the associations to them to optimize learnability?

Most participants in this study seemed to have little difficulty in creating metaphors to remember the haptic icons’ meanings; but how will this scale to larger sets? To increase scalability, more work is required to determine what underlies intuitive associations. Furthermore, if these haptic icons are to be used in different applications or interfaces, we must set standards to ensure that their meanings remain consistent throughout.

6.6 High Workload Environments

Another important consideration, given the likelihood of multitasking / time-and-safety-critical working environments, is the robustness of haptic icons to workload. Could users *utilize* the tactile signals used here in a real-world situation, while driving an automobile or using a handheld GPS while walking down a busy urban street? Methodologies for exploring these questions are being developed [9, 11, 15, 32] but the general concept of designing interfaces for high workload is one with an open future.

7 Conclusions

Results from this experiment have provided some initial data on the degree to which humans can learn and retain tactile stimulus-meaning associations, in a somewhat situated context, as well as some subjective observations on *how* they might be performing this learning feat. In general, our results are very encouraging; suggesting that everyday use of haptic icons with current tactile display technology is feasible on the basis of learnability.

We have suggested many directions for future work. Of these, the most immediately essential have to do with scalability of set size, a more detailed look at longitudinal learning, and use in realistic, attentionally demanding contexts.

REFERENCES

1. Alais, D., Morrone, C. & Burr, D. Separate attentional resources for vision and audition. *J. Neuroscience*, *1-12*.
2. Allen, M., Gluck, J., MacLean, K., Tang, E., An Initial Usability Assessment for Symbolic Haptic Rendering of Music Parameters. in *ICMI 2005, 7th International Conference on Multimodal Interfaces*.
3. Bergmann Tiest, W.M., & Kappers, A. M. L. Analysis of haptic perception of materials by MDS and physical measurements of roughness and compressibility. *Acta Psychologica*, *121* (1). 1-20.
4. Brewster, S., Wright, P. & Edwards, A. A detailed investigation into the effectiveness of earcons. in Kramer, G. ed. *Auditory Display: Sonification, Audification and Auditory Interfaces*, Addison-Wesley, 1992, 471-498.
5. Brewster, S.A. Providing a Structured Method for Integrating Non-Speech Audio into Human-Computer Interfaces, University of York, UK, 1994.
6. Brown, L.M., Brewster, S.A. and Purchase, H.C., Multidimensional Tactons for Non-Visual Information Presentation in Mobile Devices. in *Mobile HCI*, (Espoo, Finland, 2006).
7. Brown, L.M.a.K., T., Feel who's talking: using tactons for mobile phone alerts. in *CHI '06 extended abstracts on Human factors in computing systems*, (Montréal, Québec, Canada, 2006), 604-609.
8. Chan, A., MacLean, K. E., McGrenere, J. Designing haptic icons to support collaborative turn-taking. *Accepted with minor revisions to International Journal of Human Computer Studies, 2007.*
9. Chan, A., MacLean, K., McGrenere, J, Learning and Identifying Haptic Icons Under Workload. in *World Haptics*, (Pisa, Italy, 2005).
10. Day, J. Can BMW's iDrive Pass Its Road Test Now? *Electronic Design*.
11. Driver, J.a.S., C. Crossmodal Spatial Attention: Evidence from Human Performance. *Crossmodal Space and Crossmodal Attention*. 141-177.
12. Duncan, J., Martens, S. and Ward, R. Restricted attentional capacity within but not between sensory modalities. *Nature*, *387*. 808-810.
13. Enriquez, M.J., MacLean, K. E. Backward and Common-onset Masking of Vibrotactile Stimuli. *Brain Research Bullettin* (Special Issue on Robotics and Neuroscience.).
14. Enriquez, M.J., MacLean, K. E., Impact of Haptic Warning Signal Reliability in a Time-and-Safety-Critical Task. in *12th Annual Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, IEEE-VR2004*, (Chicago, USA, 2004).
15. Enriquez, M.J., MacLean, K. E., Neilson, H. Interplay of Tactile and Visual Guidance Cues under Multimodal Workload, University of British Columbia, Computer Science, Vancouver, 2007.
16. Enriquez, M.J., MacLean, K. E., Chita, C., Haptic Phonemes: Basic Building Blocks of Haptic Communication. in *Eighth International Conference on Multimodal Interfaces (ICMI'06)*
17. Gaver, W.W. Everyday Listening and Auditory Icons *Dept. of Computer Science*, University of California, San Diego, 1988.
18. Gaver, W.W. The SonicFinder: An Interface that Uses Auditory Icons. *Human Computer Interaction*, *4* (1). 67-94.
19. Gescheider, G.A.M., N. Some temporal parameters in vibrotactile forward masking. *Journal of the Acoustical Society of America*, *98* (6). 3195-3199.
20. Hall, G. *Perceptual and Associative Learning*. Oxford University Press, 1991.
21. Haller, R. The Display and Control Concept iDrive - Quick Access to All Driving and Comfort Functions. *ATZ/MTZ Extra (The New BMW 5-Series)* (August). 51–53.
22. Immersion. Vibetonez, 2005.
23. Klatzky, R.L., Lederman, S., Hamilton, C., Grindley, M. Feeling textures through a probe: Effects of probe and surface geometry and exploratory factors. *Perception & Psychophysics*.
24. Leung, R., MacLean, K., et al., Evaluation of Haptically Augmented Touchscreen GUI Elements Under Cognitive Load. in *Proceedings of 9th International Conference on Multimodal Interfaces (ICMI '07)*, (2007).
25. MacLean, K., Enriquez, M., Perceptual Design of Haptic Icons. in *Eurohaptics*, (Dublin, Ireland, 2003).
26. Oulasvirta, A., Tamminen, S., et al. Interaction in 4-second bursts: The fragmented nature of attentional resources in mobile HCI. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI '05)*, *CHI Letters, vol. 7* (1). 919-928.
27. Poupyrev, I., Okabe, M., and Maruyama, S., Haptic Feedback for Pen Computing: Directions and Strategies. in *ACM Extended Abstracts of the ACM Conference on Human Factors in Computing Systems - CHI 2004*, (2004), 1309-1312.
28. Rensink, R.A. Visual Sensing Without Seeing. *Psychological Science*, *15*. 27-32.
29. Tan, H.Z., Charlotte M. Reed, Lorraine A. Delhorne, Nathaniel I. Durlach, and Natasha Wan Temporal Masking of Multidimensional Tactual Stimuli. *Acoustic Society of America*, *114* (6). 3295-3308.
30. Tan, H.Z., Durlach, N.I., Reed, C. M., Rabinowitz, W. M. Information transmission with a multifinger tactual display. *Percept Psychophys*, *61* (6). 993-1008.
31. Tan, Y., Validity of Haptic Cues and Its Effect on Priming Visual Spatial Attention. in *11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, (2003).
32. Tang, A., McLachlan, P., Lowe, K., Saka, C. R. and MacLean, K., Perceiving Ordinal Data Haptically Under Workload. in *Seventh International Conference on Multimodal Interfaces (ICMI 2005*
33. Ternes, D. Building Large Sets of Haptic Icons: rhythm as a design parameter, and between-subjects MDS for evaluation *Computer Science*, University of British Columbia, Vancouver, 2007.
34. Ternes, D., & MacLean, K. E. (in review), Building large sets of haptic icons: rhythm as a design parameter. in *Submitted to ACM CHI '08*, (Florence, Italy, 2008).
35. van Erp, J.B.F., & Spapé, M. M. A., Distilling the Underlying Dimensions of Tactile Melodies. in *Proceedings of Eurohaptics*, (Dublin, Ireland, 2003), 111-120.
36. Wickens, C.D. Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science*, *3* (2). 159-177.