

Haptic Phonemes: Basic Building Blocks of Haptic Communication

Mario Enriquez, Karon MacLean and Christian Chita

Department of Computer Science

University of British Columbia

201-2366 Main Mall

Vancouver, B.C. V6T 1Z4, Canada

{enriquez, maclean, cchita}@cs.ubc.ca

Abstract

A *haptic phoneme* represents the smallest unit of a constructed haptic signal to which a meaning can be assigned. These haptic phonemes can be combined serially or in parallel to form haptic words, or *haptic icons*, which can hold more elaborate meanings for their users. Here, we use phonemes which consist of brief (<2 seconds) haptic stimuli composed of a simple waveform at a constant frequency and amplitude. Building on previous results showing that a set of 12 such haptic stimuli can be perceptually distinguished, here we test learnability and recall of associations for arbitrarily chosen stimulus-meaning pairs. We found that users could consistently recall an arbitrary association between a haptic stimulus and its assigned arbitrary meaning in a 9-phoneme set, during a 45 minute test period following a reinforced learning stage.

Categories and Subject Descriptors

H.5.2 User Interfaces: Haptic I/O

General Terms: Design, Experimentation, Human Factors, Languages

Keywords: Touch, haptic interfaces, haptic icons, tactile language

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 sitting in a theatre, or using a remote control that requires visual inspection while watching TV in a darkened room.

As the intelligent aids we currently use in our daily tasks become more complex, they often entail a proportional increase in 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 periods of time. Interface designers have an increasingly difficult task as they try to simultaneously address the market-driven need for more embedded functionality, while the number and form of the controls are limited by the physical dimensions of the device.

One possible approach to this complex set of design constraints is to divert some of the information flow through the touch sense. A viable implementation would allow for an increased information flow to be conveyed without overloading the visual or auditory senses, and, at the same time, not require additional interface controls: the haptic force feedback can be embedded into the existing set.

With few exceptions, haptic force feedback research has been devoted to direct rendering of virtual environments. Most force feedback devices enable 3D manipulation of graphically displayed models, and tend to be expensive and non-portable. They are used in desktop applications such as training simulators for laparoscopic surgery, sculpting of 3D models and design of mechanical assemblies. The user feels and manipulates the same thing he sees; the graphical image is rendered haptically.

But haptic feedback is also well suited for a radically different contribution, by rendering abstract models or concepts as a new modality for communication. At the lowest level, devices and objects notify users of an event, their identity or their current state or contents. Simplistic versions, such as pager vibrators, have existed for years. However, we argue that this binary or amplitude-graded signal contains far less intelligible information than may be possible with systematic, perceptually guided design; and that in the future, it may support expressive and nuanced communication that qualifies as a new haptic language.

A recurrent research theme for our group is to create ways in which the underutilized haptic sense can absorb more of the demand that modern interfaces place on vision and audition. Our group has used *haptic icons* (brief haptic stimuli associated with meanings) to provide users with information from a given device.




Figure 1. Building haptic words (haptic icons) through concatenation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICMI'06, November 2–4, 2006, Banff, Alberta, Canada.
Copyright 2006 ACM 1-59593-541-X/06/0011...\$5.00.

In order to construct haptic icons, which can also be viewed as haptic *words*, we need to first understand the communication utility of their smallest distinguishable building blocks. *Haptic phonemes* can be constructed of simple waveforms with a fixed frequency and amplitude presented through a haptic display. These phonemes can be assigned meanings which, when combined to create haptic words, can represent increasingly elaborate families of concepts that are related both semantically and haptically.

Overview

The main goal of this project is to explore the degree to which a set of haptic stimuli can convey information alone, without requiring a reinforcing visual image. It is desirable to maximize the perceivable information density of a distinct haptic signal, ideally to rival the level of information content of a graphic icon. Our approach is to train a typical non-expert user to associate an arbitrary meaning with each of a set of haptic phonemes, and then test whether these associations can be consistently remembered.

The results obtained here suggest that subjects are able to learn and later remember meaning-phoneme associations after a brief learning period. Furthermore, the associations are persistent for at least 45 minutes after the learning period.

In the remainder of this paper, we will discuss prior related research and our experimental approach and methodology; then discuss our findings and their implications.

2. Building Haptic Words

In our definition, haptic phonemes represent the smallest recombinant module of a physical haptic stimulus. They must satisfy the following requirements:

- **Differentiable:** All phonemes must be distinguishable from one another when presented either alone or in any used haptic word (icon) combinations.
- **Identifiable:** Once a meaning has been associated to a stimulus to create a phoneme, it must be easy to remember.
- **Learnable:** The associations between meanings and stimuli should be intuitive and easy to learn.

While a haptic phoneme cannot be broken into smaller recombinant units, it does have multiple minimal dimensions which can be uniquely parameterized. For example, a single phoneme must have both a waveform (a specification of the temporal path its signal takes) and a frequency (the rate at which that path is traversed). We can leverage this by associating sub-meanings with the different dimensions – e.g. frequency might indicate urgency of a cell phone call, whereas waveform could specify identity.

Haptic words (or icons) can be built from haptic phonemes using two approaches (Figure 2):

1. **Concatenation:** Phonemes are combined serially to create a word (number and complexity of required haptic syllables must be determined via user experiments), following an analogy with English word construction.
2. **Superposition:** Phonemes are combined in parallel to create a word of the same length as the longest original phoneme, following a musical chord analogy.

In the work described here, we consider phonemes in isolation.

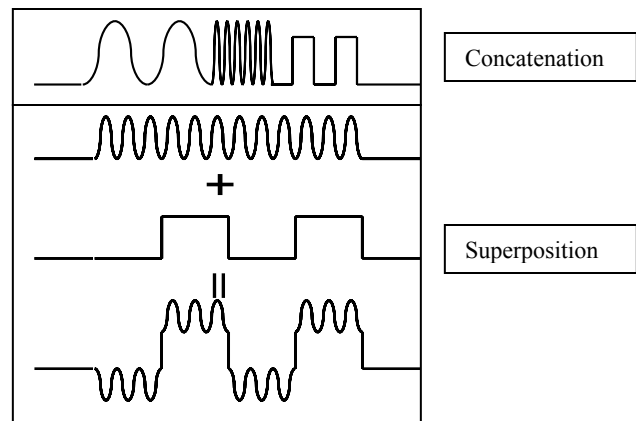


Figure 2. Using concatenation and superposition to create haptic words from haptic phonemes

3. Related Work

Some examples of information-rich haptic signaling are already appearing in commercial products, primarily in the realm of system-person communication. Cell phone manufacturers like Nokia are experimenting with distinctive vibration patterns [1, 2] hypothesizing that just as a ring-tone can be associated with a certain caller, a customized vibration can transmit the same meaning with less intrusiveness.

In 2001, BMW was the first to introduce haptic feedback on the automobile market with the iDrive™ in its 7-series vehicles (now found in other models as well). The iDrive is a force-feedback knob designed to help users access secondary vehicle functions such as audio and climate-control systems [3]. It varies the knob's feel (via programmed compliance and damping) to create a range of detent sensations, with different sensations mapped to different control functions. It was initially greeted with skepticism because of its strangeness, poor usability, and the significant learning required. However, some design iterations combined with BMW's tenacity in holding the radical concept on the market long enough for user experience to develop has met with appreciation of its functionality [4]. However, iDrive uses haptics to provide contextual cues (detents, stops, etc.) and not to communicate concepts or meanings. Perhaps the design of more information-rich cues would help make the device be more user-friendly.

Chan et al. [5, 6] developed a set of 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 designed haptic icons could be learned to a high degree of accuracy in under three minutes and remained identifiable even under significant cognitive workload. However, the haptic icons used in this study were purposefully designed to be intuitive and the rationale behind the associations used was explained to the users prior to the test.

A comprehensive understanding of our tactile psychophysical capabilities is emerging through the work of researchers such as Klatzky and Lederman, who have documented our exquisite sensitivity to texture felt through a probe [7]. Tan et al. [8] measured information transfer rates of 2-3 bits/second for vibrotactile stimuli independent of duration: appreciable content

can be conveyed through this channel. Further, there has been some work in using haptic stimuli to orient attention in another sensory modality; for example, by using taps on the back to direct gaze [9].

Another series of experiments [10] address the question of how we perceive haptic signals by using an efficient version of Multidimensional Scaling (MDS) to extract perceptual axes for complex haptic icons. Results show that a set of icons constructed by varying the frequency, magnitude and shape of 2-sec, time-invariant wave shapes map to two perceptual axes, which differ depending on the signals' frequency range; and suggest that expressive capability is maximized in one frequency subspace (7-25 Hz).

Tang et al. [11] tackle the problem of visual information overload by exploring how haptic feedback can be used as another means for information transmission. Their experiment shows that people can perceive and accurately process haptically rendered ordinal data while under cognitive workload.

Allen et al. present an exploratory study addressed at understanding the feasibility, with respect to ease of learning and usability, of efficient, eyes-free music playlist navigation based on symbolic haptic renderings of key song parameters. In this study, users were able to learn haptic mappings for music parameters to usable accuracy with 4 minutes of training. These results indicate promise for the approach and support for continued effort in both improving the rendering scheme and implementing a haptic playlist system [12].

4. Approach

We have previously tested subjects' ability to learn and remember semantic, intuitive associations and found encouraging results [5]. What kind of performance is obtainable with the harder case of deliberately *arbitrary* stimulus-meaning associations?

To test this proposition, we needed a mechanism for learning the associations as well as a test with a performance measure indicating how well the user is able to perceive and process the signals and relate them to their respective meaning. We took a three-stage approach. In the **self-guided learning stage**, users were first presented with a graphic interface that showcased a set of phonemes and allowed them to learn associations between the phonemes and a set of previously determined arbitrary meanings with no semantic relation to the haptic stimuli. This was followed by a computer-guided **enforced learning stage**, where users were asked to recognize and categorize a series of randomly presented haptic phonemes into boxes labeled with their respective meanings, while receiving reinforcement feedback about errors. Finally in the **test stage**, users performed the enforced learning task but without reinforcement.

5. Phoneme Creation

In order to ensure that phoneme stimuli met our design specifications for discriminability and identifiability, we followed a series of steps described here.

Prior research [10] shows that a person can clearly distinguish a set of 10 or more haptic signals when delivered through a haptic knob. Here we used a similar haptic knob to deliver a set of 9 haptic stimuli by combining 3 levels of waveform and frequency. For simplicity, we used periodic waveforms with fixed frequencies; each signal had a duration of 2 seconds.

We required a set of haptic stimuli which was (a) perceptually well-distributed in frequency and waveform and (b) perceptually equalized for amplitude (all signals "feel" the same amplitude). Good perceptual spread means that users perceive key stimulus attributes as varying linearly [10]. This entailed determining a perceptual transformation that simultaneously modulated both dimensions (Section 5); for the waveform, we also had to ascertain a transformation path.

Finally, phonemes were created by assigning arbitrary meanings to these stimuli.

Perceptual Transform of Waveform

For this experiment, we decided to use simple waveforms varying from a triangle to a square wave; prior work had revealed that these haptic signals are perceptually distinctive, whereas a triangle and a smooth sine wave, for example, are indistinguishable [13]. We explored several waveform "morph" functions in search of one which varied linearly between these endpoints in a user's perception, rather than feeling similar along most of its range. The most obvious path (a linear interpolation) failed to elicit a range of intermediate sensations.




Figure 3. Perceptually intermediate waveforms from triangle to square

The transform described in Figure 3 best satisfied our criteria. A pilot user study revealed that the main factor that influenced the perceived sensation for this intermediate waveform is the height of the vertical component, with shorter vertical components feeling more like a triangle and longer ones more similar to a square waveform.

Multidimensional Distribution

To achieve uniform perceptual spacing of stimuli composed of simultaneously varying parameters, we chose a large initial set and then refined a final subset using a previously developed Multidimensional Scaling Analysis (MDS) tool [10, 14].

Our initial set contained 25 stimuli composed of 5 waveforms (triangle, square and 3 intermediates) displayed at 5 different frequencies (3, 7, 13, 18 and 21 Hz). The MDS analysis of user comparison data results in a mapping of these stimuli according to their perceived relationships (Figure 4). We then selected a subset of 9 stimuli which maximized perceptual spread and differentiability according to this map. The selected stimuli were composed of Triangle, Morph*0.4 (vertical line 40% of total amplitude) and Square waves presented at 7, 10 and 18 Hz.

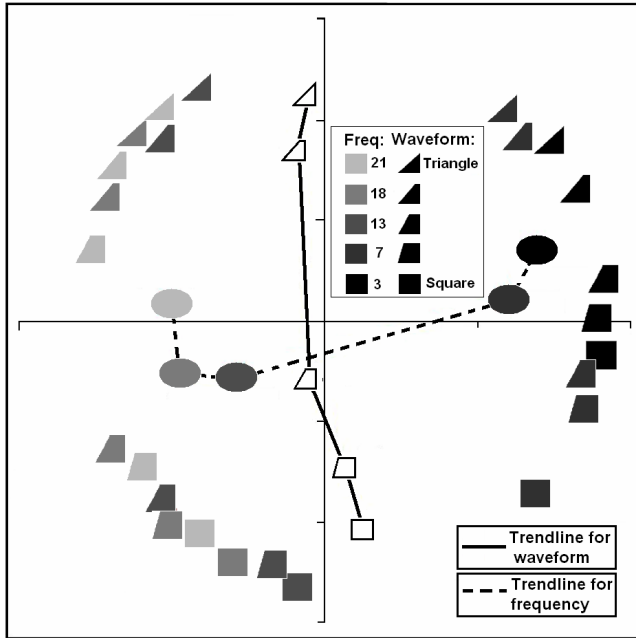


Figure 4. MDS perceptual distribution of initial (25-element) stimulus set. Each point in the horizontal trend is the average value of the respective points for that frequency, and increases roughly from right to left. Likewise, the vertical trend shows average values for each waveform setting: 1 = triangle; 2 = 0.25_morph; 3 = 0.5_morph; 4 = 0.75_morph; 5 = square

Perceptual Equalization of Final Stimulus Set

Finally, we ran a short user study using Parameter Estimation by Sequential Testing (PEST) [15] on our chosen set, attain perceptual equality of the frequency and waveform combinations used. Table 1 shows the resulting relative amplitudes..

Table 1 Relative amplitudes for perceptual equalization of stimuli (2 sec duration), with assigned nonsense meanings

	7 Hz <i>Grass</i>	10Hz <i>Flower</i>	18Hz <i>Tree</i>
Triangle <i>Blueberry</i>	0.62	0.55	0.45
Morph (0.4) <i>Strawberry</i>	0.62	0.55	0.45
Square <i>Orange</i>	0.62	0.55	0.45

Family-Based Meaning Association

The main purpose of this experiment was to test the effectiveness of truly arbitrary associations between haptic phonemes and meanings. In order to ensure that the haptic phoneme – concept associations remained abstract, we chose nonsense meanings that were carefully screened to have no semantic connection to the stimuli themselves. However, we exploited the 2-dimensional property of the set and grouped meanings by what we will call *families* (associated with particular frequencies) and *functions* (associated with waveforms), allowing users to “chunk” groups of items and thus assemble larger sets with less learning effort [16].

The selected family (frequency-related) meanings were types of plants, while functions (waveform-related) were fruits; all 6 items were selected to suggest unique colors, shapes and textures (Table 1). Thus, a given stimulus was associated a dual meaning: an 18 Hz triangle wave could be regarded as either a Tree or as a Blueberry.

6. Methods

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

6.1 Physical Setup and Instructions

Phonemes were displayed on a direct-drive actuated knob, shown in Figure 5. The knob was rubber-covered brass with an outer diameter of 10.5 mm and a length of 16.5 mm. The 1mm thick rubber coating prevented slipping and allowed a better grip of the knob while minimizing compliance. The knob was mounted directly on the shaft of a 20W Maxon DC motor (RE025). This 24-volt motor has a stall torque of 240 mNm with a position frequency roll-off at 200 Hz. No position feedback was required since all signals were delivered open-loop.

Participants were seated at a table so that they could comfortably rest their hand and hold an actuated knob (Figure 5) which was placed on a raised platform matched by a padded armrest. 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 17” 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.



Figure 5 Participants held a haptically enabled knob through which haptic phonemes were displayed

6.2 Protocol

Following instruction, the experiment had three phases: self-guided learning, enforced learning and testing.

Self-Guided Learning Phase

Using the GUI shown in Figure 6a, participants repeatedly selected different Family-Function combinations and felt the corresponding haptic stimulus (consisting of the corresponding waveform at the specified frequency displayed by the haptic knob). Participants could spend up to 5 minutes in this phase and for uniformity, they were not allowed to return to the self-guided learning interface after this point.

Enforced Learning Phase

Throughout this stage, participants interacted with Identification and Reinforcement views of the Enforced Learning GUI (Figure 6b). Participants were presented with 4 randomly selected phonemes (presented as draggable icon tiles) and three destination boxes labelled with either Families (Grass, Flower and Tree) or Functions (Blueberry, Strawberry, and Orange). A left mouse-click on a phoneme tile triggered playback of the corresponding haptic stimulus. With a right mouse-click, the tile could be dragged into the chosen box.

This phase had 20 randomized trials: 10 with boxes labelled by Family, and 10 by Function; a trial consisted of sorting 4 tiles, for a total of 80. Participants were allowed to feel a given tile's stimulus up to 15 times before placing it. To discourage participants from grouping tiles by compared similarity (as opposed to 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 error feedback and required to re-sort incorrectly identified tiles (Reinforcement view).

The vertical layout of the destination boxes matched the order in which the buttons were labelled in the self-guided learning phase (Figure 6a). Our goal was to benchmark learnability of associations for a the case of arbitrary associations. It was acceptable for this learning to be demonstrated with either a spatial or semantic mapping, both of which are arbitrary.

Test Phase

All participants proceeded to the test phase regardless of performance in the 20-trial enforced learning phase. We chose to measure performance following a uniform learning period, rather than the learning time required to achieve a specified performance level, in complement to the approach taken by [5].

The test phase consisted of 20 randomized repetitions of the 9 phonemes tiles, with 10 repetitions sorted into Family boxes and 10 into Function boxes. 180 sort items were randomly allocated to 36 trials with 5 tiles each. To minimize fatigue, a 5-minute break was enforced after Trial 18.

A test trial was identical to the identification step of an enforced learning trial, with the exception that participants sorted 5 tiles instead of 4 (Figure 6c) and were allowed to display a trial 20 rather than 15 times. We used 5-tile trials here (rather than 4 as during learning) to avoid participant re-use of systematic strategies (other than association knowledge) developed during learning. Following the sort, participants *did not* receive feedback on their performance.

Measures

We collected several measures including number and identity of phonemes placed in each destination box by experiment condition, and enforced learning and test phase durations.

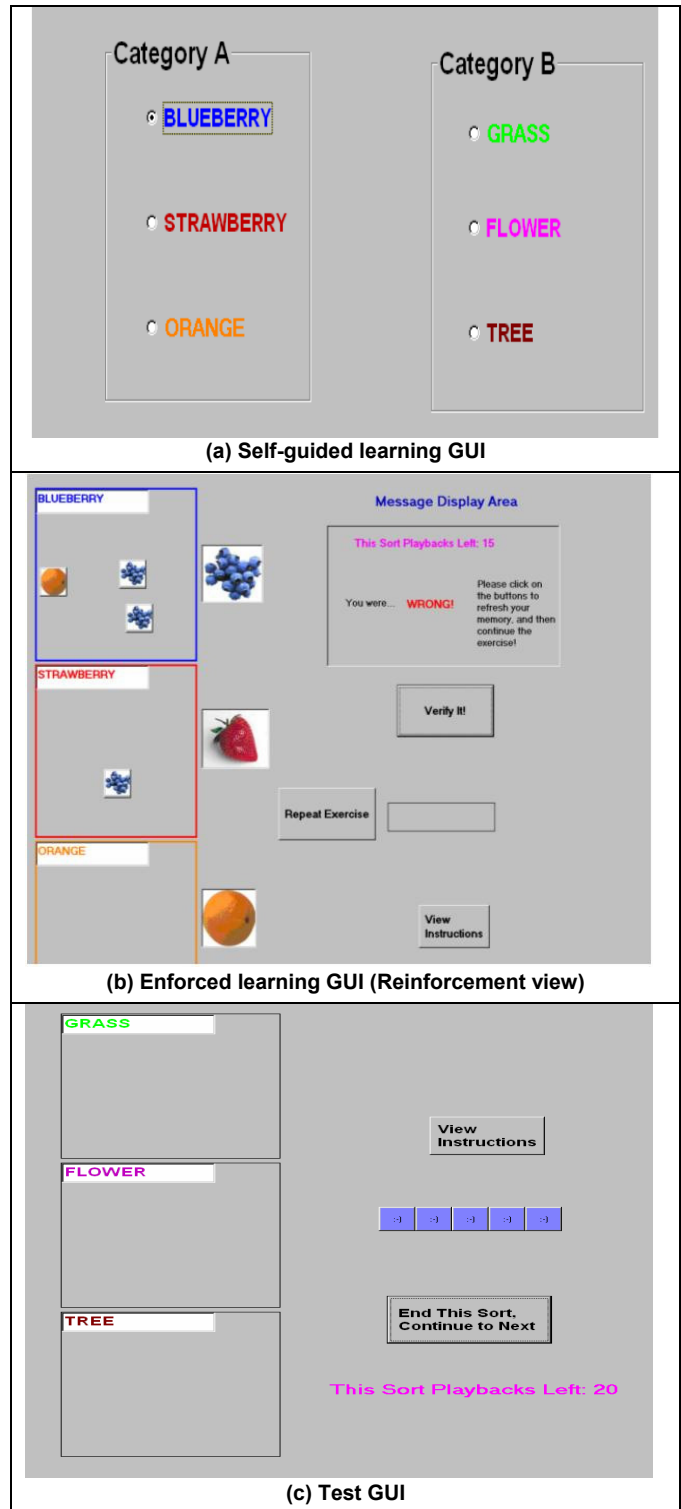


Figure 6. Graphic interfaces for the three experiment phases. During Identification (b), tiles are initially displayed as in (c); then at reinforcement, feedback text and tile graphics are displayed.

7. Results

12 participants (6 male, 6 female; age range from 20-45 with median 27) took part in this study; most were university graduate students in Computer Science. Most had no experience with haptic displays, and the rest had moderate experience. Each was paid \$10 for an approximately 1.5-hour session.

7.1 Enforced Learning

The enforced learning phase had an average duration of 19.3 minutes (stdev. 4.8, min 10.6 and max 28.2 minutes). On average, participants made 19.4 out of a possible 80 errors (76% correct, where chance would be 33%; stdev 6.17, min 12, max 33 errors). During this phase, participants made on average 76% correct responses. Individual performance during this phase is shown in Figure 7.




Figure 7. Individual performance during Enforced Learning: % correct responses

7.2 Average Identification Performance

Figure 8 summarizes the correct and erroneous phoneme identifications made in the test phase, for waveform (Function) and frequency (Family). Correct identification averaged across all three waveforms was 22.0/30 (73%) where chance performance would be 33%. Identification of meanings assigned by frequency was slightly better, with, an average for all frequencies of 24.2/30 (81% correct, chance=33%).

Figure 8 also illustrates confusion patterns. For example, subjects on average confused the triangle waveform 8.1 /30 times (27%) with the morphed waveform, but only 0.66 (2%) with the square; 7Hz stimuli were confused 6.2 (21%) times with 10Hz, but only 1.3 (4.3%) times with 18Hz. While the ordinal structure of the stimuli tends to drive miss-associations towards the center categories (morph and 10 Hz), it is useful to consider asymmetry within each plot and the two plots with each other: Triangle-to-Morph errors are made more often than Square-to-Morph, and there are more confusions overall for waveform than frequency.

7.3 Individuals Analysis

Figure 9 shows the aggregate identification performance for each subject alone. Most participants apparently formed strong, *correct* haptic stimulus-meaning associations. A small number appeared to form consistent *incorrect* associations. For example, Participant 1 consistently identified the morph stimulus at 7Hz for morph at 10Hz (70% of the time) and the square stimulus at 7Hz for square at 10Hz (70%). Participant 7 consistently identified the triangle stimuli for morph (80% of the time). From this we conclude that

both Participants 1 and 7 are a representative example of the category of participants who consistently learned a certain association, albeit the wrong one.

Others learned some mappings well, but had difficulties with others. For example, Participant 4 achieved near-perfect scores for the triangle (90% accuracy) and morph (87% accuracy) stimuli but exhibited poor results for the square waveform stimuli (56% accuracy).




Figure 8. Average number of placements made in Test phase. Correct placements are enumerated by the bars marked '*'. Each figure represents average behavior across subjects of 90 item sorts per subject, where 30 correct placements into each of the three destination boxes would be a perfect performance.




Figure 9. Individual identification performance (average of correct responses for each destination box)

The best-performing participants (4, 5 and 10) made mistakes in only one of the frequency or waveform associations.

8. Discussion

Overall, the results of this experiment show that with training of about 25 minutes, participants demonstrated learning of correct associations for the difficult case of arbitrary associations with a performance rate of 73 or 81% for waveform and frequency respectively (chance 33%). An individual's analysis showed that some of the participants were consistently making the same mistakes, suggesting that incorrect but consistent associations were learned. We view these results as promising, particularly in consideration of the intentionally non-intuitive associations and brief training period. We already know that more intuitive mappings can lead to even shorter learning periods (Chan et al observed a 3-minute average for performance-gated learning of a set of similar size and structure but metaphorically mapped [5]).

8.1 Validation of Methodology

Adequacy of Learning and Learning Protocol

During enforced learning, participants made on average 76% correct responses, which in turn resulted in comparable and reasonable test phase performance (average 77%). Test phase performance indicates that learning worked overall, and its similarity to enforced learning phase performance implies that learning had plateaued.

Nevertheless, the design of the learning phases might be improved. It is possible that some of the wrong associations made by the participants could have been avoided with a 'smarter' training method. In particular, we believe performance could be improved by allowing participants to return to the self-guided learning interface when consistent mistakes in the enforced learning phase have been detected (Figure 9). In addition, a slightly longer self-guided learning phase (here it was limited to 5 minutes) might have allowed a shorter and more effective enforced learning phase.

Conservatism of Results

The reported results could be considered to be conservative (for the abstract association case) in that they do not take into consideration these incorrectly associated stimulus-meaning pairs,

The experiment was designed to test for properly associated stimulus characteristic-meaning pairs. However, we did not test for identification of both meanings associated to any stimulus simultaneously. We believe that if we had done this, it might have made the associations easier to remember.

8.2 Influence of Stimulus and Individual Individuals Performance

During the enforced learning phase, participants made on average 76% correct responses, which in turn resulted in reasonable test phase performance, producing an average of 23.1 (77%) correct responses (stdev = 8.0; max =29.7; min =18).

Participants seemed to experience increased difficulty in isolating the mid-range stimuli, but did significantly better for the extremes of the set. This comes as no surprise since we expected the anchoring effect to facilitate the identification of stimuli whose characteristics fall into either ends of their respective scale [17].

There were no participants who performed generally poorly throughout the experiment. Rather, we distinguish three categories of performance:

1. Did well on all categories (6/12)

2. Did well overall, but consistently learned the wrong association for some categories (4/12)

3. Did well on some categories, but performed poorly on others (2/12)

Most participants (Types 1 and 2) were able to construct a strong and consistent mental binding, despite the fact that the associations presented to them were deliberately designed to hold no intuitive meaning. The predominance of this ability leads us to believe that the majority of users will be able to learn the desired stimuli-arbitrary meaning associations.

With regards to Type 3 participants, we hypothesize that in the same manner that 10% of the population is color-blind; there may be a naturally-occurring difficulty in learning haptic stimuli associations which we might call "haptic-numb". Such a deficit could occur either in perceptual inability to make distinctions, or cognitive difficulty in making associations. To distinguish these, we would need to pair an association test with a stimulus distinctiveness test for every individual.

At the other end of the scale, some participants might have a natural talent in memorizing haptic signals and represent the higher proficiency we observed in this study. These participants probably have a natural ability *both* to make haptic distinctions and form cognitive associations.

Extending the haptic language: Family and Function

Giving meanings to different characteristics of a haptic phoneme allows us to increase the information density of a haptic language. In this experiment, we utilized haptic phonemes created by varying two characteristics (frequency and waveform) amongst 3 possible values. We hypothesized that this approach would allow more information-rich haptic signals to be created.

To actually test this premise, a direct comparison of flat and family-based sets must be made. In either scenario, the experiment paradigm used here was not designed to ascertain maximal learnable set size, but instead measured performance for a given set size under a given learning procedure. Further work will be required to determine how large a number of distinct haptic phonemes can be used in a set while still remaining easy to learn and effective.

Enduring Associations

We measured the number of times that participants correctly sorted the family + function haptic representations into their corresponding family or function. This was taken as an indication of how difficult it is for a person to learn these arbitrary associations. We were also interested in measuring the change in ability to respond correctly to the signals that occurs during 45 minutes of testing.

We found that overall, participants were able to correctly remember the associations throughout a 45 minute test phase after the enforced learning phase ($p=0.028$). This suggests that the associations were stored in a longer-lived form than short term memory. It is reasonable to expect that without further demand, these associations would soon degrade following the end of the experiment; further work will be required to establish their longevity given regular use and reinforcement.

8.3 Design Implications

Test results show that after a brief (5 + 20 minute) training period, participants are able to learn and remember arbitrary haptic

phoneme - meaning associations with an average performance close to 80%; a benchmark which we view to be conservative with respect to a real application which would probably benefit from a more intuitive mappings and ongoing learning reinforcement. This suggests that a similar set of haptic phonemes could be successfully utilized as a means for a device to communicate simple information to its user.

Confusion patterns suggest that some people might need more time to learn correct haptic phoneme-meaning associations. Results show that most participants formed strong mental bindings for all the phonemes with the correctly (or incorrectly) learned associations; and the remainder could make good associations in most cases. Furthermore, it should be noted that most mistakes in identifying the haptic phonemes occurred when participants mistook either end of a characteristic's scale with that of the middle value. Perhaps selecting phoneme sets with only two values per dimension (avoiding mid-values), would result in sufficiently improved identification performance to justify additional dimensions to increase set size. For example, would people do better with 3 dimensions with 2 values on each ($2^3=8$) rather than 2 dimensions with 3 values ($3^2=9$)?

9. Conclusions and Future Work

Inspired in part by human ability to parse non-intuitive graphical icons, the findings presented here suggest that haptic signals can be a robust way to communicate meaningful information to a user: arbitrary associations between haptic phonemes and meanings can be learned to a usable performance level after a 25 minute training period and remembered consistently for a relatively long period of time (45 minutes).

This opens the possibility for the design of interfaces that take advantage of the sense of touch as a communications channel. If arbitrary associations can be learned, this simplifies the task of designing haptic interfaces intended to communicate meaningful information to their users. More intuitive associations, when available, can probably be learned more easily still. Further, we hypothesize that with regular and pervasive reinforced exposure, larger set sizes could be learned to better accuracy. More work is needed to establish this.

In addition to the small experiments proposed in Discussion, most of which are targeted at increasing usable phoneme set sizes, we need to explore ways of increasing the amount of information that can be encoded in a single haptic information module. For example, the set of 9 haptic phonemes tested in this project could theoretically be used (through concatenation or superposition) to create a larger set of haptic words (or icons) which could convey more complex meanings and perhaps open the way for the development of a far-reaching haptic language. How much complexity can be perceptually and cognitively decoded from tactile messages?

.Another important consideration, given the likelihood of multitasking / time-and-safety-critical working environments, is the robustness of haptic icons or phonemes to workload. For example, could users still appropriately utilize the haptic phonemes from this study in a real world situation, such as in an automobile navigation aid or a cell phone identify feature used while walking down a busy urban street? Methodologies for exploring these questions is being developed (e.g. Chan 2005,

McLachlan 2005) but the general concept of designing interfaces for high workload is one with an open future.

10. References

1. Immersion, *Vibetonez*. 2005.
2. 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*. 2006. Montréal, Québec, Canada.
3. Haller, R., *The Display and Control Concept iDrive - Quick Access to All Driving and Comfort Functions*. ATZ/MTZ Extra (The New BMW 5-Series), 2003(August): p. 51-53.
4. Day, J., *Can BMW's iDrive Pass Its Road Test Now?* Electronic Design, 2004.
5. Chan, A., MacLean, K., McGrenere, J. *Learning and Identifying Haptic Icons Under Workload*. in *World Haptics*. 2005. Pisa, Italy.
6. Chan, A., *Designing Haptic Icons to Support an Urgency-Based Turn-Taking Protocol*, in *Computer Science*. 2004, University of British Columbia: Vancouver, BC.
7. 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, 2003.
8. Tan, H.Z., Durlach, N.I., Reed, C. M., Rabinowitz, W. M., *Information transmission with a multifinger tactual display*. Percept Psychophys, 1999. **61**(6): p. 993-1008.
9. Tan, H.Z., Gray, R., Young, J. J., Traylor, R., *A Haptic Back Display for Attentional and Directional Cueing*. Haptics-e, 2003.
10. MacLean, K., Enriquez, M. *Perceptual Design of Haptic Icons*. in *Eurohaptics*. 2003. Dublin, Ireland.
11. 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)*. 2005. Trento, Italy.
12. 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*. 2005. trento, italy.
13. Enriquez, M., *A study of haptic icons*, ed. C. University of British and S. Dept. of Computer. 2002, Vancouver: University of British Columbia. xi, 119 leaves.
14. Pasquero, J., Luk, J., Little, S., MacLean, K. *Perceptual Analysis of Haptic Icons: an Investigation into the Validity of Cluster Sorted MDS*. in *Haptic Symposium 2006*. 2006. Washington, DC.
15. Taylor, M.M., Creelman, C. D., *PEST: Efficient Estimates on Probability Functions*. Journal of the Acoustical Society of America, 1967. **41**(4A): p. 782-787.
16. Miller, G., *The magical number seven, plus or minus two: Some limits on our capability for processing information*. Psychological Review, 1956. **63**.
17. Paulhus, D.L., *Measurement and control of response bias*, in *Measures of personality and social psychological attitudes*, P.R.S. J.P. Robinson, & L.S. Wrightsman, Editor. 1991, Academic Press: New York Academic Press.