

Coding Tactile Symbols for Phonemic Communication

Siyanz Zhao, Ali Israr, Frances Lau, Freddy Abnousi
Facebook Inc., 1 Hacker Way, Menlo Park, CA, USA
{siyanz, aliisrar, flau, abnousi}@fb.com




Figure 1. Haptic messages are delivered through the skin when it is not convenient for users, e.g. while, a) holding grocery bags in both hands or b) driving.

ABSTRACT

We present a study to examine one's learning and processing capacity of broadband tactile information, such as that derived from speech. In Study 1, we tested a user's capability to recognize tactile locations and movements on the forearm in the presence of masking stimuli and determined 9 distinguishable tactile symbols. We associated these symbols to 9 phonemes using two approaches, random and articulation associations. Study 2 showed that novice participants can learn both associations. However, performance for retention, construction of words and knowledge transfer to recognize unlearned words was better with articulation association. In study 3, we trained novel participants to directly recognize words before learning phonemes. Our results show that novel users can retain and generalize the knowledge to recognize new words faster when they were directly train on words. Finally, Study 4 examined optimal presentation rate for the tactile symbols without compromising learning and recognition rate.

Author Keywords

Touch; Haptic Interfaces; Haptic Language

ACM Classification Keywords

H.5.2. User Interfaces: Haptic I/O

INTRODUCTION

The human skin is a powerful yet underutilized communication medium. Embedded with a wide variety of receptor systems, our skin informs us about environmental conditions, such as temperature and pressure, and intrinsic properties of objects in contact, such as compliance and roughness [18,43]. Moreover, touch is crucial for interpersonal and social communication. It conveys emotions,

comfort and authority [11,12]. Despite the skin's rich capacity to transfer meaningful information, it has shown limited success to translate language to sighted users [2,5, 20,33]. In this paper, we investigate the skin's ability to convey broadband information, such as the one derived from a continuous stream of speech, and examine artificial means to communicate coded tactile cues to novel users. Our work lays the foundation for using the skin to alleviate visual and auditory sensory overload and to deliver discreet messages to users when it is not convenient for them to engage with other modalities.

Most previous efforts have used vibrations to communicate coded messages, which has low throughput on the skin. Specifically, on the forearm 2-3 pure vibration points could be correctly identified [41]. This recognition rate substantially decreases when presented with an accompanying masking stimulus [36]. Tan and colleagues improved the user's performance by incorporating multidimensional cues (vibrations combined with finger motion [35]). Eagleman and colleagues also improved the transmission by incorporating movement cues between two vibrating points [25]. It is our hypothesis that multidimensional haptic stimuli will be easily differentiated, memorized and retained, therefore, in this paper we examined identification of both vibration points and vibration movements, sandwiched between two masking stimuli, presented on the forearm.

Another consideration factor for communication through the skin is to determine a coding scheme that maps tactile symbols to natural units of language. Previous studies have utilized a variety of schemes to map acoustics, phonetic, alpha-numerals and symbolic features to distinct tactile features [10,14,20,21,25]. For example, the Tadoma method (a speech communication method common in deaf-blind communities [29]) conveys rich facial actions during speech production on a receiver's hand. The success of such scheme relies on relaying information via established neural pathways for optimal tactile (or tactual) presentations,

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5620-6/18/04...\$15.00

<https://doi.org/10.1145/3173574.3173966>

which not only ease the comprehension of coding but also reduces cognitive demands and learning resources [27,32].

Finally, an important consideration for successful tactile communication is the training method. Efficient training leads to good memorization and retain performance with minimal training time. It should also support generalization of the learned knowledge to recognize new words. In this paper, we evaluated two common perceptual learning approaches for haptic language learning [34]. First is the “bottom-up” approach (Study 2), in which users are trained on phonemes and tested to construct words. They are also tested to recognize new (untrained) words to see if they can generalize their learning beyond the trained set. The second training approach is the “top-down” method (Study 3) where participants are trained directly to recognize words before generalizing the training to recognize phonemes.

The contributions of our work are: 1) 9 discriminable haptic patterns (Study 1); 2) a method of meaningfully associating haptic patterns with phonemes using their articulation placements (Study 2); 3) two training approaches, i.e., top-down vs. bottom-up (Study 3); 4) an optimal display speed that maximizes transmission rate and recognition accuracy (Study 4); 5) a range of real-life applications of our work.

RELATED WORK

In 1957, Frank Geldard introduced the concept of tactile literacy [9], suggesting that our skin can transmit language and share the load with the visual and auditory systems. Using five tactors (vibrators for tactile communication), Geldard devised a vibratory communication system that conveyed alpha-numerals characters as vibration codes presented on the back. Since Geldard’s initial success, many researchers have explored the idea of transmitting language through the skin. These systems can be roughly grouped into two categories by the method used: acoustics-based approach and symbol-based approach.

The acoustic-based approach converts speech audio directly to vibrotactile stimulations. It typically uses a custom algorithm to process the incoming speech signal and spatially map them to haptic stimuli, simulating how cochlea decomposes spoken language to different frequencies. For example, Brooks and colleagues used a tactile vocoder that is worn on the ventral forearm [6]. It maps low frequency information to the wrist and high frequency to the elbow. A user study with one subject showed a learning rate of 2.7 words/hour for the first 150 words and a learning rate of 3.9 words/hour to acquire additional 100 more words. The average accuracy in the last 50 sessions was 75.6%. A more recent example of the acoustic approach was by David Eagleman and Scott Novich [7,25,26]. They designed and implemented a vest, VESTVoice, that takes (real-time or recorded) audio and maps it to the back, from low to high frequency. Seven participants were trained using the system for 12 days to recognize 100 unique words. Their performance after the training was 35% – 65%. When a set of new words was

played to the participants, their accuracy was 30%, which demonstrates some transfer of knowledge with their approach. In general, the acoustic approach has fast presentation rate (similar to the delivery speed of spoken words). However, it requires extensive training and does not generalize well for users to recognize untrained words.

The second approach to delivering haptic language is a symbol-based approach. This approach breaks language down to components, e.g., letters, and codes each with an identifiable vibration pattern [23]. The benefit of the approach is its freedom to design the haptic patterns. For instance, haptic icons [37] and tactons [4] combined multiple physical parameters, e.g., frequencies, amplitudes, rhythms and waveforms, to create large sets haptic symbols. With tactons, people’s accuracy rate was 77% for a set of haptic icons containing 84 distinguishable stimuli. Another example of the symbol-based approach is Vibratese [9] which uses 5 vibrators to encode English letters and numbers to unique vibratory patterns. However, Vibratese required extensive training (12 hours) to learn. Another example is Skin Reading [23] that translates 26 letters to different vibrotactile patterns on a user’s hand. It requires 3 one hour-long sessions to learn. But once learnt, the recognition accuracy can reach over 90% accuracy. Comparing to the extended training, *haptic phonemes* designed by Enrique, MacLean, and Chita [8] was a set of 9 haptic symbols associated with 9 arbitrary semantic meanings. With 25 minutes of training, participants reached 80% accuracy. Another system, EdgeVib [20], reduced the training time further to 15 minutes by converting Graffiti characters (for letters and numbers) to spatiotemporal vibration patterns on the wrist. User studies showed an average recognition rate of 87.25% on letters and 83.3% on a short message. While both haptic phonemes and EdgeVib achieved shorter training time, their presentation rates are relatively slow. A haptic phoneme is roughly 2 seconds and a 3-vibration letter in EdgeVib takes about 1900 ms. This translates to ~31 haptic phonemes (letters) per min.

While most of the work reviewed so far decomposes language to either words or letters, there is an approach that uses phonemes as building blocks. The benefit of delivering phonemes via haptic displays is faster delivery rate. On average, an English word is composed of 5 letters [3] while it contains ~3.34 phonemes [19]. Delivering less symbols per word can speed up the overall presentation rate. Israr and colleagues [15] explored schemes that display phoneme-based speech input to the hand and the forearm. Instead of decomposing frequencies to different locations, they mapped places of articulation of phonemes to locations. Their studies showed good discrimination in consonants on the forearm. Inspired by Israr and colleague’s phonemic-place of articulation-approach, we custom designed symbols for each phoneme to maximize perceptual differences between them while incorporating place of articulation for mapping.




Figure 2. The haptic device used in the study.

DESIGNING HAPTIC SYMBOLS

Participants and Common Apparatus

All participants are employees of our organization and within the age range of 20-50 years old. None of the participants had prior experience with similar haptic devices. New participants were recruited for each study.

The haptic device, used in all studies, contained 6 voice-coil actuators (model: TEAX13C02, Tectonic Elements, UK). Figure 2 shows the device and the layout of actuators. The layout is similar to that of Braille. The actuators were separated apart beyond the 2-point threshold, which is 40 mm on the forearm [39,40]. The device was mounted on the dorsal side of the left forearm to free the right hand for other tasks. The areas of the actuators that directly contacted the skin were circular, with a diameter size of 1.8 cm. To reduce the contact area to focus the vibrations, we mounted cylindrical rubber tips ($d=1$ cm) to the voice coils.

The actuators were covered with insulated tape and were Velcroed onto a shin guard that conformed around the forearms. The actuators were driven by a custom electronic driver to amplify the audio output from a laptop. An audio interface (MOTU Ultralite mk3) was used to drive the 6 actuators individually through an USB port of a laptop.

Study 1: Recognizable Haptic Patterns on the Forearm

To begin the design of haptic language, it is critical to understand the spatial capability of the display area, i.e., the dorsal part of the forearm. Therefore, we first tested for optimal parametric values to generate vibrotactile signals and determine a set of simple spatiotemporal patterns that were easily distinguishable. We investigated 4 parameters: frequency, duration, location and spatial pattern. As low amplitude stimuli are easily masked by successive stimuli, we opted to use a fixed amplitude in the study.

Stimuli

Study 1 examined 2 frequencies (30 Hz and 250 Hz) and 2 actuator durations ($D = 150$ ms and 400 ms). The frequency values were selected to target two different receptors in the skin: Pacinian corpuscle cells, which respond more actively to high frequency stimuli, and Merkel cells, which are more active with low-frequency signals. A linear onset and offset of 10% of the actuator duration was used at the beginning and the end of the stimulus.




Figure 3. A visual illustration of the AXB paradigm.

Stimuli were either a single-point actuation or two-point sequential actuations. The single-point actuations are perceived as a single “poke” on the skin while the two-point actuations trigger apparent motion between the points, which feels like a point traveling from one actuator to the other. This illusory motion has been used in vibrotactile display to transmit, e.g., spatial, information [1,22,38,42].

The Stimulus Onset Asynchrony (SOA , in milliseconds) between the two actuations was determined by a previous study as a function of duration, i.e., $SOA = 0.17 \times D + 45$. This function yields the smoothest illusory motion in participants (see [44] for the definition of smooth apparent motion). We did not use motion patterns that were in more than one direction because past work indicates that focusing on another direction draws attention away from perceiving a separate direction [17] and lowers the accuracy of correctly detecting the motion. The total number of haptic patterns was 36; 6 single-point locations and 30 two-point movements (6 positions for the start point \times 5 positions for the end point).

Note that the total stimulus duration for two-point actuations was longer than the one-point actuations. For the two-point actuations, the total stimulus duration in the 150 ms condition was 250.5 ms ($D + \text{onset/offset} + SOA$) vs 165 ms for a single-point ($D + \text{onset/offset}$) and in the 400 ms condition, it was 593 ms vs 440 ms. We will refer to the duration of an actuator as the actuator duration and total duration of a stimulus as the stimulus duration.

In language, phonemes do not stand alone and the same applies to haptic phonemes. Therefore, Study 1 used an AXB paradigm, where the participants were asked to identify the stimulus (X) sandwiched in a sequence of three stimuli (AXB) [36]. Both forward and backward maskers (A and B) are randomly selected from the 36 patterns. The actuator duration of all three stimuli were the same in each trial, i.e. 150 ms or 400 ms. The Inter-Stimulus Interval (ISI), i.e., the temporal gap between when a stimulus stops and when the next stimulus starts (Figure 3), was 500 ms.

Procedure

After participants sat comfortably in front of the testing computer, the experimenter explained the task and assisted in placing the display on the forearm. Each factor was actuated in sequence to ensure that all were perceivable at a similar intensity. Adjustments were made if necessary. After the display was worn properly, they went through 5 training trials before starting the main experiment. In each trial of the experiment, participants felt three haptic patterns (AXB) and their task was to identify the middle one (X) by

typing the actuator number (Figure 5a) with the keyboard. For motion patterns they typed the numbers in the order of movement, e.g. number 12 was the motion from location 1 to location 2. Upon submitting their responses, they got correct answer feedback before the next trial started.

In order for the participants to refer to the sheet when responding, they had a sheet with the layout and numbering of the actuators (Figure 5a). Each participant was tested on 720 trials (2 frequencies \times 2 durations \times 36 patterns \times 5 repetitions). The whole study took roughly 2 hours, split into 2 one-hour sessions on two sequential days.

Results and Discussion

Four participants were tested (2 females). A repeated-measures ANOVA was done with duration and frequency as the within factors and the total number of symbols correctly identified as the dependent variable. The test showed a significant effect of duration [$p=0.03$, $\eta^2=0.83$] and frequency [$p=0.02$, $\eta^2=0.87$]. There was no significant interaction between the two variables [$p=0.55$, $\eta^2=0.13$]. Figure 4 plots participants' performance for each duration and frequency. Specifically, the performance was better with high frequency (250 Hz) and long actuator duration (400 ms). The accuracy difference between the 2 actuator durations was less than 5% even though the longer duration was more than twice the length of the shorter one. Considering that using 150 ms as the duration could greatly increase the transmission speed, we presented the stimuli at 150 ms and 250 Hz in subsequent studies.

We further examined and categorized major types of errors in identifying the patterns. Note that there were overlapping trials between the following types of errors. Therefore, the error rates did not add up to 100%.

The most common type of mistake came from confusions between non-diagonal motion and diagonal motion. A non-diagonal motion was between two actuators that was aligned either vertically (e.g., actuator 1 and 3) or horizontally (e.g., actuator 1 and 4). Out of the total 1671 errors, 1284 errors (76.84%) were due to diagonal vs non-diagonal confusions. Specifically, participants were more likely to misperceive diagonal motions as non-diagonal ones: 595 out of 959 diagonal trials (62.04%) were




Figure 4. Recognition rate by actuator duration and frequency. Error bars show 1 standard deviation. Each shape represents a participant.

misidentified versus 689 out of 1431 non-diagonal trials (48.15%). This result aligns with Jones's work [17] where the authors found that participants were more accurate at identifying haptic motions that are in cardinal directions (either up/down or left/right).

The second common error was confusions between single-point and two-point actuations (20.47% of total errors). In addition, directional errors, e.g., an up-down motion being perceived as a down-up motion, were also common (11.49% of total errors). Both types of errors could be a result of the short SOA values. While the SOA-duration equation yields smooth motion, it did not guarantee discriminable patterns. In later studies, we increased the SOA between the two-point actuations by 20 ms ($SOA = 0.17 \times D + 65$). This ensured a strong sense of illusory motion while creating a noticeable distinction from a single-point pattern.

Another 6.5% of errors was a result of misidentifying long and short patterns. This type of error was only applicable to vertical patterns as there was only one distance for the horizontal patterns. A long pattern was created from actuators on the wrist-side to the elbow-side or vice versa while a short pattern as a vertical pattern between two adjacent actuators. A closer look at the data revealed that there was an interaction between the perception of distance and actuation duration. That is, at 150 ms long-distance symbols were likely to be judged as short-distance symbols.

In order to identify patterns that were distinguishable and recognizable by the participants, we applied an algorithm to reduce the 36 patterns to a smaller set. First, the algorithm clustered patterns that were easily confusable with each other into groups based on a stimulus-response confusion matrix (included as supplementary material). We defined a pattern as being easily confusable with another one if they were falsely confused over 12.5% of the trials. To maximize the distinguishability, if two groups shared confusable patterns, we combined the two groups. The algorithm then selected one pattern from each group such that the chosen pattern was least confusable with the remaining patterns outside of the group. The accuracy rate threshold for "least confusable" was 75%.

This algorithm formed 9 confusion groups, which yielded 9 final symbols, shown in Figure 5b. Note that the resulting 9 symbols only contained non-diagonal and short distance motions, consistent with the error analysis from Study 1.

Discriminable Haptic Patterns on the Forearm

To confirm that people could perceive the 9 symbols accurately, we did a small study with 2 novel participants. The patterns were generated using an actuator duration of 150 ms and a frequency of 250 Hz. With the new SOA of 90.5 ms, the stimulus duration for the two-point actuations was 270.5 ms. Participants were shown a visual illustration of the 9 symbols (same as Figure 5b), labelled as pattern 1–9. The same AXB paradigm was used and participants were given correct answer feedback after each trial. Each




Figure 5. (a) The layout of the actuators on the dorsal side of the forearm; (b) The 9 haptic symbols identified from Study 1 and their associations to language using either random pairing or place of articulation.

participant completed 10 repetitions for all nine phonemes (90 trials in total).

The study result showed an accuracy rate of 88%. Using the method introduced in [36], we calculated the information transfer (IT) with the 9 haptic patterns to be 2.59 bits. There was an effect of training. During the second half of the study, i.e., the last 5 repetitions, the participants could accurately recognize 91% of the trials, which gave an IT of 2.75 bits. This short study verified that the 9 selected haptic symbols were distinguishable and recognizable.

PAIRING HAPTIC PATTERNS WITH PHONEMES

Past literature has shown that human memory is built on associations. Meaningful associations can help people recall learnt material [13,31]. Therefore, it is critical to design meaningful associations between the haptic patterns and phonemes. As phonemes are sound articulated by a speaker, an intuitive association is to map where a phoneme is pronounced in the mouth, i.e., place of articulation, to where a haptic symbol is played on the forearm. For example, if a phoneme is pronounced on the front of the mouth (e.g. bilabial /m/) the haptic symbol will locate at the wrist end of the forearm. Similarly, a phoneme placed on the back of the mouth (e.g., velar /k/) is associated with a symbol that is presented at the elbow end of the forearm.

Another association we created between the haptic symbols and phonemes is with duration. In English, there are two main types of phonemes, consonants and vowels. In English, vowels can be long and short vowels but consonants are only short [28]. While the lengths of consonants and vowels vary in spoken English, vowels are generally longer than consonants. In our design, we paired vowels with longer symbols, i.e., the two-point actuations, and consonants with shorter symbols, i.e., single-point actuations.

With this association schema, we chose 5 consonants (/m/, /d/, /dh/, /w/, /k/) and 4 long vowels (/A/, /AI/, /E/, /OO/) and mapped them to the haptic symbols from Study 1. To examine effects of association on haptic learning, we proposed *random associations* between the haptic symbols with phonemes. Both associations are shown in Figure 5b.

We hypothesized that the articulation associations would help users learn the phoneme-symbol associations, or *haptic phonemes*. The boost in performance would carry on to people’s performance on learning to recognize *haptic words*, a sequence of haptic phonemes. Compared to people who used the articulation associations, people who learned the random associations would perform worse at learning the haptic phonemes and haptic words.

LEARNING OF HAPTIC PHONEMES AND WORDS

The remaining three studies focus on participants’ learning performance of haptic phonemes and haptic words. Study 2 examines whether participants can associate haptic symbols with phonemes and apply the knowledge to perceive words. In this “bottom-up” approach, learning is first done in small building units and scaled up to construct words.

In comparison, Study 3 explores a different training method, a “top-down” approach, where participants are directly trained to recognize words and then are tested on individual building units. We used the two approaches to see if there is a difference in learning performance and training time. Study 2 and Study 3 also measure retention of learned haptic phonemes and whether the knowledge can be applied to recognize new haptic words. Study 3, along with Study 4, evaluates optimal delivery rate of haptic phonemes by testing a set of duration and ISI combinations without compromising participant’s recognition performance. The results of these studies inform future design of haptic phonemes and haptic words.

Study 2: Perceiving and Learning Haptic Phonemes and Words

Study 2 addresses the following questions:

- Can people associate phonemes with haptic symbols? (Day 1)
- Can people learn the associations? If yes, what is the performance on learning and recalling the associations, i.e., haptic phonemes? (Day 1)
- How much is retained 24 hours after training? (Day 2)
- Can a meaningful sequence of haptic phonemes, i.e., haptic words, be perceived? (Day 2)
- How well can people learn to recognize haptic words over time? (Day 2)
- Can trained users generalize the haptic learning to perceive words they have not trained on (transfer of knowledge)? (Day 2)
- How do meaningful associations affect the learning performance for the above tasks? (Day 1 & 2)

The training method in Study 2 is analogous to the analytic (or “bottom-up”) approach in perceptual learning [34]. This approach trains people on phonemes before they learn to concatenate phonemes to perceive words. The benefit of this approach is that after establishing performance on recalling phonemes, errors that occur in perceiving a word are a result of misperceiving a sequence of phonemes rather than faults on perceiving individual phonemes.

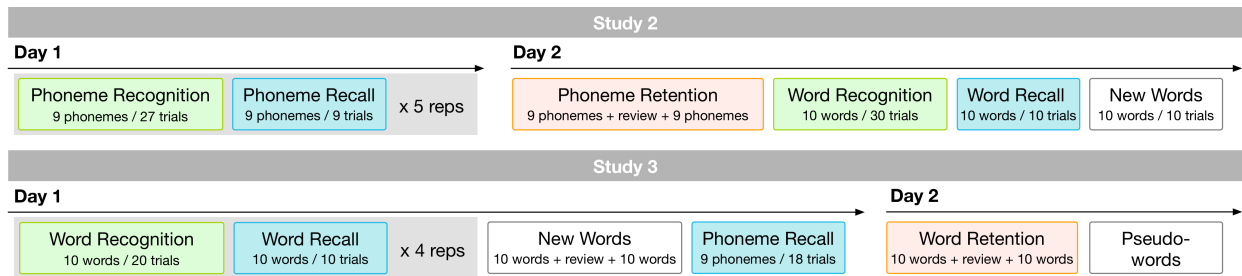


Figure 6. Task timeline for Study 2 and Study 3.

Stimuli

With the 9 phonemes introduced in Study 1, we formed 20 two- and three-phoneme words (Table 1). Other than the phoneme /dh/, the remaining 8 phonemes were counterbalanced so that they had similar number of occurrences in the word list. Then, we divided the 20 words into two lists. The occurrence frequency for each phoneme, excluding /dh/, was maximally balanced in each list.

Word List 1	2-phonemes: <i>they, woo, die, me</i> 3-phonemes: <i>mime, mood, make, weed, cake, wide</i>
Word List 2	2-phonemes: <i>may, key, do, why</i> 3-phonemes: <i>doom, dime, meek, womb, deed, wake</i>

Table 1. The words used in Study 2 and Study 3.

Based on the result from Study 1, the actuator duration in Study 2 was 150 ms with an ISI of 500 ms. With 10% onset and offset, the total duration of a haptic consonant is 180 ms ($150 + 0.2 \times 150$). For a haptic vowel, the total duration is 270.5 ms ($180 + 90.5 [SOA]$). The throughput of the 2-phonemes and 3-phonemes words was shown in Table 2. Two types of associations were tested in Study 2: articulation associations and random associations. The association type was randomly assigned between subjects.

Procedure

Day 1: The Day 1 study used the AXB paradigm as Study 1. Each trial contained 3 phonemes with 500 ms gap in between. The participants were instructed to type the phoneme that corresponded to the middle stimulus.

There were two types of trials: recognition and recall. The recognition trials were similar to the training procedure in perceptual learning studies. Participants received correct answer feedback after typing and submitting their responses. A reference sheet with phoneme-haptic symbol associations (Figure 5b) was available for the participants. In a block of recognition trials, the phonemes were shown in a random order for 3 times (a total of $3 \times 9 = 27$ trials).

For the recall trials, participants could not use the reference sheet nor did they receive correct answer feedback. The task was designed to evaluate how many of the phoneme-symbol associations the participants remembered. In a block of recall trials, each phoneme was randomly presented once (a total of 9 trials). Day 1 contained 5

repetitions of recognition block–recall block (180 trials in total; see Figure 6) and lasted about 45 minutes.

After the participants put on the haptic display, the experimenter triggered the six actuators in sequence to ensure that all were in good contact with the skin. Then, the experimenter presented the participants with the reference sheet that corresponded to the association condition assigned to them. For the participants in the articulation association condition, the experimenter explained to them how place of articulation and articulation length were used to map the haptic symbols to the phonemes.

To familiarize with how the haptic symbols would feel on the skin, the participants went through a training screen where they could repeatedly play each symbol. To demonstrate the presentation speed of the three stimuli, the training screen also had a Play button that randomly played a trial. The participants freely explored the training screen before they started the main experiment. The experimenter logged the time to complete each block.

Day 2: After verifying that all 6 tactors were properly contacting the skin, the participants went through a retention test on the haptic phonemes. Without reviewing the reference sheet, they were presented with the 9 haptic symbols in a random order and in the AXB paradigm. As in Day 1, they were asked to type in the phoneme associated with the middle symbol. After the retention test, they briefly reviewed the reference sheet and repeated the same test to get a measure of how a short review affected the recall of the phoneme-symbol associations (Figure 6).

After the retention test, the participants continued with learning haptic words, using word list 1 (Table 1). The experimenter showed the participants the word list and asked them to read them out loud. This was to confirm that participants could pronounce the words correctly, which

	Actuation duration	ISI	Word duration
Study 2&3	150 ms	500 ms	950.5 / 1721 ms
Study 4	100 ms	200 ms	522 / 924 ms

Table 2. Word duration given different duration and ISI values. In the last column, the first value was for 2-phoneme words (CV) and the second was for 3-phoneme words (CVC).

was critical to the phonetic approach. After the participants read through the list, it was removed from the participants so that they would recognize the words solely based on the symbols they felt on the forearm.

The same recognition and recall trials were used. In each trial after a word was presented, the participants typed the word that they felt. If the sequence of haptic phonemes they recognized did not make up a meaningful English word, they were instructed to type the phonemes in the order that were felt. In the recognition trial block, each word was repeated 3 times (30 trials) and in the recall task block, each word was repeated once (10 trials). The reference sheet and correct answer feedback were provided for the recognition trials but not the recall trials. The participants completed one repetition of recognition–recall block (40 trials).

The last task in the Day 2 session was to recognize a new list of words, i.e., word list 2. Unlike the previous task, participants were not told of the content of word list 2. They typed the words based on what they recognized through the skin. No reference sheet or correct answer feedback was provided. Each word was presented once (10 trials). The experimenter logged the time taken to complete each task. All tasks on Day 2 were completed in 30 minutes.

Results

A total of 11 new participants completed the study (8 female). 5 of the participants were assigned to use the articulation associations and the remaining 6 were assigned to the random associations.

Phonemes: The average time to complete 5 repetitions of phoneme recognition-recall trials was 30.4 minutes (SD=5).

First, we examined the effect of association type on the phoneme-level recognition. As we observed a training effect in the second part of Study 1, we conducted a repeated-measures ANOVA with the repetition block as the within-subject variable, the association type as the between-subject variable, and recognition accuracy, i.e., percentage of trials out of total trials where participants correctly recognized the phonemes, as the dependent variable. The result showed no significant effect of association type on the recognition accuracy [$p=0.32$, $\eta^2=0.11$]. There was also no effect of repetition on the recognition performance




Figure 7. Study 2 phonemes: the average response accuracy for phoneme recognition and phoneme recall tasks. The error bars plot 1 standard deviation.

[$p=0.26$, $\eta^2=0.13$]. However, the difference between participants was significant [$p<0.001$, $\eta^2=0.98$].

The same analysis was replicated on the 5 repetitions of the recall trials with recall accuracy, i.e., percentage of correctly recalled trials out of total trials. While there was no significant effect of the association type on the recall accuracy [$p=0.62$, $\eta^2=0.03$], there was a significant difference between repetitions [$p<0.001$, $\eta^2=0.40$]. A significant difference between participants was also observed [$p<0.001$, $\eta^2=0.97$]. Figure 7 plotted people’s accuracy on the recognition and the recall trials for the five repetitions, combining the two association conditions.

On the second day, people’s retention rate was 71% for the first time they did the retention task and 82% for the second time. A repeated-measures ANOVA with association type as the between-subject variable, repetition block as the within-subject variable, and retention rate as the dependent variable showed no significant effect of association type on the retention rate [$p=0.43$, $\eta^2=0.02$]. There was also no significant difference between people’s performance on the two repetitions of the retention tests [$p=0.42$, $\eta^2=0.07$].

Words: The average time to complete the word recognition–recall task was 12.2 minutes (SD=6.6) and was 4.8 minutes (SD=2.6) to recognize the new set of words.

A repeated-measures ANOVA with repetition block as the within-subject variable and the association method as the between-subject variable was done on the word recognition rate. We found a significant effect of the association type [$p=0.01$, $\eta^2=0.52$] and repetition blocks [$p=0.01$, $\eta^2=0.4$], but no interaction between the two [$p=0.2$, $\eta^2=0.36$]. To examine the effect of association on the recall accuracy, an independent samples T-test was done but showed no significant effect of association type [$t(9)=0.81$]. The average word recall accuracy for the articulation association was 74% (SD=19.5%) and for the random association was 63.3% (SD=23.4%) (Figure 8).

New Words: We conducted an independent samples T-test for the effect of association on new words recognition. The test showed a significant difference between people’s performance using the two association types [$t(9)=2.36$, $p=0.04$]. The mean for the articulation associations




Figure 8. Study 2 words: Interaction between association methods and the round of repetition on word recognition rate. The error bars show 1 standard deviation.

($M=46\%$, $SD=16.7\%$) was higher than that of the random associations ($M=26.7\%$, $SD=10.3\%$) (Figure 8). Among the errors in the articulation association condition, 78% were due to false recognition of a single phoneme (i.e., recognized *wake* as *week*) and 7.8% were due to misspelling (i.e., recognized *womb* as *woom*).

Discussion

Study 2 confirmed that participants could learn to associate haptic symbols with phonemes. The type of association used did not significantly affect the learning performance on the phoneme-level, i.e., phonemes recognition rate, recall rate, and retention rate, nor did it significantly affect the word recall rate. Despite of the lack of effect on the phonemes, the articulation association significant helped the participants recognize words (Figure 8) and new words. The lack of association effect on the phoneme-level might be due to the small number of phonemes people had to remember. If the phoneme set were larger, the association type would be more important to the phonemic performance.

In terms of training effect, people’s performance improved over time for the phoneme recall task and the word recognition task. Compare to the phoneme recognition task where no significant training effect was observed, these two tasks were more difficult. This indicates that training is necessary for people to learn the haptic phonemes and the haptic words. The average total training time for the phonemes and words combined was 42.6 minutes.

Study 2 also looked at memory retention of haptic phonemes on the next day. The result indicates that the participants had consolidated the haptics-phoneme associations to short-term memory. The recognition–recall training scheme may have boosted the retention performance as Roediger and Karpicke [30] showed that testing (recall task) after studying (recognition task) leads to better performance in delayed recall.

We also examined transfer of knowledge using our phoneme-haptic symbols approach. While 46% accuracy rate on recognizing the new words is not ideal, most of the errors were due to falsely recognizing a single phoneme. Therefore, there is room for improvement. In day-to-day scenario, context provides priori-probability of what words are more likely to appear. However, in the new word recognition task, the participants had no prior knowledge as to what words they would feel. Preparing their mind for the words, such as showing them the word list beforehand, may improve their performance. We looked at this in Study 3.

In Study 2, the participants learned the phonemes before recognizing words. A different learning method, a synthetic




Figure 9. An illustration of time to recognize a haptic word.

approach or a “top-down” method, is examined in Study 3.

Study 3: Top-Down Training Approach to Learning Haptic Words and Phonemes

Study 3 also spanned across 2 days. The research questions addressed in each day were independent. For Day 1, we evaluated the learning performance with the “top-down” training method. For Day 2, we wanted to, without affecting the recognition rate of the new words, increase the presentation rate of the phonemes in a word by reducing the stimulus duration and the ISI between phonemes. It is important to note that the shortest duration and ISI required to recognize a word that a person has never felt before is longer than the time required to recognize a word the person has trained on. This is due to increased chunk size in memory after training [16].

The concept of chunking is that individual pieces of information can be grouped into larger units [24]. When a person perceives a new word, one needs time to process each phoneme after it is presented. After feeling all three phonemes, he or she concatenates the three phonemes into a word. After that, additional time is needed to interpret the word and retrieve it from memory (Figure 9). The total time to perceive the word is $T_{p1} + T_{p2} + T_{p3} + T_c$. However, after extensive training, the person can perceive a sequence of phonemes as a single chunk without processing each phoneme individually. This reduces T_{p1} , T_{p2} and T_{p3} , leaving the total perception time to approach T_c .

Given the chunking effect, in Study 3 Day 2, we want to measure the shortest duration and ISI values needed for a person to reliably perceive words they haven’t learned before. Therefore, to get a close estimation of the shortest duration and ISI values, we need participants to be familiar with the haptic phonemes but not with haptic words. That is why this task is scheduled on the second day, after having trained on the haptic phonemes. The results of Study 3 Day 2 will inform the presentation speed of haptic phonemes.

Stimuli

The stimuli used in Day 1 are same as Study 2. Each consonant had a stimulus duration of 180 ms and a vowel 270.5 ms. The ISI between each phoneme was 500 ms. The 9 phonemes and both word lists were used.

In Day 2, we tested 3 duration values (50, 100, 150 ms) and 3 ISI values (100, 200, 300 ms). To have a condition comparable to the previous studies, we also added duration of 150 ms and ISI of 500 ms. This gave us 10 combinations of duration–ISI values. To get a reliable measure of accuracy, each combination was repeated 10 times (a total of 100 trials). The duration value in a single trial was the same for all three haptic phonemes.

To avoid the effect of chunking on perception time, we would, ideally, use 100 unique words to avoid repetition. However, the 9 chosen phonemes could not generate 100 different 2- or 3-phoneme English words. To make up for the lack of words, we used pseudo-words that were in the

form of CVC (consonant-vowel-consonant) or VCV (vowel-consonant-vowel). This formation was the closest representation of the words from list 1 and 2. It also ensured that people could pronounce the pseudo-words with the phonetic components. All tasks were timed manually by the experimenter from start to end.

Procedure

Day 1: The first day contained three tasks: word-level recognition–recall, new word recognition, phoneme recognition. In the word-level recognition–recall task, participants went through 4 repetitions of recognition–recall blocks. In each block, they completed the recognition task on the 10 words from word list 1 twice before going through the recall task once. For the new word recognition section, a similar procedure from Study 2 Day 2 was used: without having prior knowledge on word list 2, participants got the 10 words in a random order. Based on results from Study 2, we saw a large proportion of errors that we believed could be avoided had the participants had prior knowledge of the word list. To confirm this, after going through the word list once, participants were briefly shown word list 2 to review. The sheet was taken away after the review and the participants repeated the task for the second time. To measure how the top-down learning method trains the participants on recognizing the phonemes, a short phoneme recall test was delivered at the end of the Day 1 session. Without referring to the phoneme reference sheet, the participants went through a recall task on the phonemes. Each phoneme was randomly presented twice (a total of 18 trials). All tasks on Day 1 took 50 minutes to complete.

Day 2: The Day 2 session started with a retention task on the words from word list 1. Same as in Study 2 Day 2, participants completed the task twice, once before review of the haptic phonemes and words and once after.

After the retention task, the participants were briefed on the concept of pseudo-words and the task. They were asked to recognize and type the three phonemes in each pseudo-word. For phonemes they could not recognize, the participants were instructed to replace it with “-”. The participants completed the 100 trials in 35 minutes.

Results

A total of 9 novel participants (4 females) completed the study. Among the 9 participants, 2 did not finish the Day 2 session. Their data were taken out in the analysis.

Day 1: The average time to complete the word recognition – recall task was 26 minutes (SD = 6.6) and the new word recognition task 6.1 minutes (SD = 1.6). To examine the training effect on people’s word recognition rate, we conducted a repeated-measures ANOVA with repetition block as the within-subject variable and word recognition rate as the dependent variable. Unlike Study 2, it showed no significant effect of training [$p=0.09$, $\eta^2=0.29$]. The mean word recognition rate for all repetitions was 88.6% (SD = 11.0%). For the word recall task, there was a significant

ISI (ms)	Duration (D)		
	50 ms	100 ms	150 ms
100	22.6 (11.1)	21.4 (17.7)	28.6 (21.9)
200	24.3 (29.0)	38.6 (18.6)	37.1 (14.9)
300	32.9 (17.9)	40.0 (18.2)	45.7 (25.1)
500	–	–	44.3 (28.2)

Table 3. Average accuracy (%) for pseudo-words of all durations and ISIs in Study 3 Day 2. Standard deviations are in parentheses.

difference in people’s performance between the 4 times they repeated the task [$p=0.03$, $\eta^2=0.39$]. The recall rate improved from 60% (SD = 28.9%) in the first repetition to 82.9% (SD= 23.6%) in the last repetition. The participants’ difference was also significant [$p<0.001$, $\eta^2=0.90$].

For the phonemes, the average number of phonemes recognized correctly was 90.5% (SD=7.7%). For word list 2, the recognition accuracy rate before knowing the words was 55.7% (SD=29.9%) and for the second time after they were told of the list content, it was 72.8% (SD=23.6%). A repeated-measures ANOVA, with word recognition rate as the dependent variable, was done to check if the accuracy difference between the two times was significant. The result was marginally positive [$p=0.048$, $\eta^2=0.58$].

Day 2: On the second day, people remembered 7 out of the 10 words (SD=0.13) the first time that they did the retain task. The second time, the average number of words remembered was 8.67 (SD=0.19). A repeated-measures ANOVA with words remembered as the dependent variable revealed no significant difference between performance for the two repetitions [$p=0.19$, $\eta^2=0.31$].

The recognition accuracy rate for the pseudo-words was listed in Table 3 for all durations and ISIs. A repeated-measures ANOVA with duration and ISI as the within-subject variables and recognition accuracy as the dependent variable showed a significant effect of ISI [$p=0.004$, $\eta^2=0.6$] but not duration [$p=0.076$, $\eta^2=0.35$] on the recognition performance. Interaction between duration and ISI was not significant [$p=0.62$, $\eta^2=0.10$]. There was a significant difference between subjects [$p<0.01$, $\eta^2=0.86$]. Pair-wise t-tests on the 3 ISI levels, combining the durations, showed a significant difference between ISI of 100 ms and 200 ms [$t(6) = -4.21$, $p<0.01$] and between 100 ms and 300 ms [$t(6)=-4.48$, $p<0.01$]. But there was no significant difference between 200 ms and 300 ms [$t(6)=-1.29$, $p=0.22$]. The t-tests also showed no significant difference between ISI of 500 ms and 300 ms [$t(6)=0.54$, $p=0.6$], nor between 500 ms and 200 ms [$t(6)=1.39$, $p=0.2$].

We also found a significant difference between the CVCs and VCVs [$t(6)=3.62$, $p=0.01$]. The mean recognition rate was 41% (SD=17.2%) and 26% individually (SD=14.5%).

Discussion

With the top-down training approach, the training time to recognize haptic words and phonemes was 26 minutes. Because the designs for Study 2 and 3 differ, we cannot directly compare the results. However, it is worth-noting that the training time with the top-down approach was 16.6 minutes shorter than the bottom-up approach while the word recognition accuracy was similar in both studies.

When recognizing the new words, the participants performed much better after they reviewed the words. This aligns with our hypothesis that preparing their minds by showing them the word list can improve their performance.

The Day 2 study informed us on how short we can design the haptic words. Actuation duration did not significantly affect people's performance. While this suggests we can shorten the duration to be 50 ms, it is too short that people may not focus their attention in time to accurately perceive the signal in a real-life scenario with distractions. For the ISI, 200 ms is the most optimal value. The participant's performance for 200 ms was comparable to 300 and 500 ms and it delivers the fastest transmission rate among the three.

In the next study, we quickly evaluated the new duration and ISI to confirm that naive users can learn to recognize haptic phonemes and words presented at the new speed, using the top-down training approach.

Study 4: Learning with Optimal Presentation Rate

Stimuli and Procedure

Same phonemes and words from Study 3 were used. Based on Study 3, actuation duration was set to 100 ms and the ISI to 200 ms. The procedure was the same as Study 3 Day 1.

Results

4 novel participants (3 female) were tested. The average time to complete the word recognition – recall task was 27.5 minutes (SD=7.2) and 5.5 minutes (SD=0.6) for the new word recognition task. Same ANOVA analyses were done on word list 1 as in Study 3 Day 1, first with recognition rate then with recall rate as the dependent variables. We found no significant effect of training on recognition rate [$p(3)=0.5$, $\eta^2=0.22$]. The mean recognition rate for the 4 repetitions was 76.3% (SD=14.6%). However, the training effect was significant for the recall rate [$p(3)<0.01$, $\eta^2=0.72$]. It was 52.5% (SD=15%) in the first repetition, and 82.5% (SD=8.7%) in the last repetition. For the new words, the recognition rate was 55% (SD=23.8%) the first time the participants felt words list 2 and 87.5% (SD=12.6%) after the participants reviewed the list. Their phoneme recognition accuracy was 73.6% (SD=11%).

Discussion

With the shorter duration and ISI values, the participants' accuracy on word recognition and recall were comparable to Study 3 (Table 4). So was the performance for the new words recognition task. The phonemes recognition rate went down from 90.6% (SD=7.7%) in Study 3 to 73.6% in

	Stud 3 Day 1	Study 4
Word Recog	88.6 (11)	76.3 (14.6)
Word Recall	82.9 (23.6)	82.5 (8.7)
New Words	72.8 (25.6)	87.5 (12.6)
Phonemes	90.6 (7.7)	73.6 (11)

Table 4. Comparison between Study 3 Day 1 and Study 4 results. Word recognition accuracy was the average of all trials. Word recall and new words accuracy were from the last repetition. Standard deviations are in parentheses.

Study 4 (SD=11%). Overall, we think the new duration and ISI values could be used to present haptic words.

GENERAL DISCUSSION AND FUTURE WORK

This paper presents a method to communicate language components on the skin, using vibrotactile patterns. We associated haptic symbols with location of articulation of phonemes. User studies showed that meaningful phoneme-symbol associations helped participants remember haptic words and phonemes. Users recognized 20 haptic words and 9 haptic phonemes after 26 minutes of training and they could still strongly remember the words and the phonemes 24 hours after the training.

The training order also affected learning performance. Users can directly learn to recognize haptic words and, in the process, they acquire the knowledge on the individual phonemes. This top-down training is time efficient and gives the same learning result.

Our work has many applications. It can be used to alleviate visual and auditory sensory overload and provide a means to discreetly deliver information to users in a non-intrusive manner. For example, an augmented reality device could use tactile coding to deliver information about physical surroundings or to deliver notifications, without being visually disrupted. Another application space of our work is to deliver information to users when they are unable to read content, e.g., running or carrying groceries in both hands. More broadly, our work can be extended to a complete input/output system that can not only receive tactile messages but also send tactile information to other devices.

Future directions of our work will focus on expanding the 9 phonemes, which is a limitation to the current work. To generate more haptic symbols, we will explore using multiple locations on the arm, such as the dorsal side of the forearm, upper forearm, and/or bilateral arms. We will also increase the physical parameters used to make the symbols, such as various waveform and curvilinear haptic patterns, e.g., a loop around the arm. Given that working memory is limited, it is crucial to examine how increase in phonemes and words will affect people's learning performance with our training method. Another direction of future work is to consider if people can recognize haptic words in situations where their attention was not fully devoted to the system.

ACKNOWLEDGMENTS

We thank Prasad Chakka, Robert Turcott, and Wahyudinata Setiawan for valuable technical support and guidance.

REFERENCES

1. David S Alles. 1970. Information Transmission by Phantom Sensations. *IEEE Transactions on Man-Machine Systems* 11, 1: 85–91. <https://doi.org/10.1109/TMMS.1970.299967>
2. Gabriel Arnold, Charles Spence, and Malika Auvray. 2016. Taking someone else's spatial perspective: Natural stance or effortful decentring? *Cognition* 148: 27–33. <https://doi.org/10.1016/j.cognition.2015.12.006>
3. Vladimir Bochkarev, Anna Shevlyakova, and Valery Solovyev. 2012. Average word length dynamics as indicator of cultural changes in society. *Social Evolution and History* 14: 153–175.
4. Stephen Brewster and Lorna M. Brown. 2004. Tactons: structured tactile messages for non-visual information display. *Proceedings of the fifth conference on Australasian user interface - Volume 28*, 15–23.
5. P. L. Brooks and B. J. Frost. 1983. Evaluation of a tactile vocoder for word recognition. *The Journal of the Acoustical Society of America* 74, 1: 34–39. <https://doi.org/10.1121/1.389685>
6. P L Brooks, B J Frost, J L Mason, and K Chung. 1985. Acquisition of a 250-word vocabulary through a tactile vocoder. *The Journal of the Acoustical Society of America* 77, 4: 1576–9. Retrieved August 14, 2017 from <http://www.ncbi.nlm.nih.gov/pubmed/3157716>
7. David Eagleman. 2014. Plenary talks: A vibrotactile sensory substitution device for the deaf and profoundly hearing impaired. In *2014 IEEE Haptics Symposium (HAPTICS)*, xvii–xvii. <https://doi.org/10.1109/HAPTICS.2014.6775419>
8. Mario Enriquez, Karon MacLean, and Christian Chita. 2006. Haptic phonemes: basic building blocks of haptic communication. In *Proceedings of the 8th international conference on Multimodal interfaces - ICMI '06*, 302. <https://doi.org/10.1145/1180995.1181053>
9. Frank A. Geldard. 1957. Adventures in tactile literacy. *American Psychologist* 12, 3: 115–124. <https://doi.org/10.1037/h0040416>
10. KS Hale and KM Stanney. 2014. Handbook of virtual environments: Design, implementation, and applications. *Human Factors and Ergonomics Series*: 1232. Retrieved from <http://books.google.com/books?id=6GnFavZHZZcC>
11. Matthew J. Hertenstein, Rachel Holmes, Margaret McCullough, and Dacher Keltner. 2009. The communication of emotion via touch. *Emotion* 9, 4: 566–573. <https://doi.org/10.1037/a0016108>
12. Matthew J. Hertenstein, Dacher Keltner, Betsy App, Brittany A. Bulleit, and Ariane R. Jaskolka. 2006. Touch communicates distinct emotions. *Emotion* 6, 3: 528–533. <https://doi.org/10.1037/1528-3542.6.3.528>
13. Thomas S. Hyde and James J. Jenkins. 1969. Differential effects of incidental tasks on the organization of recall of a list of highly associated words. *Journal of Experimental Psychology* 82, 3: 472–481. <https://doi.org/10.1037/h0028372>
14. Ali Israr, S.-C. Kim, Jan Stec, and Ivan Poupyrev. 2012. Surround haptics: Tactile feedback for immersive gaming experiences. *Proceedings of the 2012 ACM annual conference extended abstracts on Human Factors in Computing Systems Extended Abstracts - CHI EA '12*: 1087–1090. <https://doi.org/10.1145/2212776.2212392>
15. Ali Israr, Peter H. Meckl, Charlotte M. Reed, and Hong Z. Tan. 2009. Controller design and consonantal contrast coding using a multi-finger tactual display. *The Journal of the Acoustical Society of America* 125, 6: 3925–3935. <https://doi.org/10.1121/1.3124771>
16. Gary Jones. 2012. Why Chunking Should be Considered as an Explanation for Developmental Change before Short-Term Memory Capacity and Processing Speed. *Frontiers in psychology* 3: 167. <https://doi.org/10.3389/fpsyg.2012.00167>
17. Lynette A Jones, Jacquelyn Kunkel, and Erin Piatieski. 2009. Vibrotactile Pattern Recognition on the Arm and Back. *Perception* 38, 1: 52–68. <https://doi.org/10.1068/p5914>
18. Roberta L. Klatzky and Susan J. Lederman. 1999. Tactile roughness perception with a rigid link interposed between skin and surface. *Perception & Psychophysics* 61, 4: 591–607. <https://doi.org/10.3758/BF03205532>
19. Lori F. Lamel, Robert H. Kassel, and Stephanie Seneff. 1989. Speech Database Development: Design and Analysis of the Acoustic-Phonetic Corpus. *Speech Input/Output Assessment and Speech Databases 2*: 161–170.
20. Yi-Chi Liao, Yi-Ling Chen, Jo-Yu Lo, Rong-Hao Liang, Liwei Chan, and Bing-Yu Chen. 2016. EdgeVib. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology - UIST '16*, 595–601. <https://doi.org/10.1145/2984511.2984522>
21. J. G. Linvill and J. C. Bliss. 1966. A direct translation reading aid for the blind. *Proceedings of the IEEE* 54, 1: 40–51. <https://doi.org/10.1109/PROC.1966.4572>
22. Granit Luzhnica, Sebastian Stein, Eduardo Veas, Viktoria Pammer, John Williamson, and Roderick Murray Smith. 2017. Personalising vibrotactile displays through perceptual sensitivity adjustment. In

- Proceedings of the 2017 ACM International Symposium on Wearable Computers - ISWC '17*, 66–73. <https://doi.org/10.1145/3123021.3123029>
23. Granit Luzhnica, Eduardo Veas, and Viktoria Pammer. 2016. Skin Reading: Encoding Text in a 6-Channel Haptic Display. In *Proceedings of the 2016 ACM International Symposium on Wearable Computers*, 148–155. <https://doi.org/10.1145/2971763.2971769>
 24. George A Miller. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review* 63, 2: 81.
 25. Scott D. Novich and David M. Eagleman. 2015. Using space and time to encode vibrotactile information: toward an estimate of the skin’s achievable throughput. *Experimental Brain Research* 233, 10: 2777–2788. <https://doi.org/10.1007/s00221-015-4346-1>
 26. Scott David Novich. 2015. Sound-to-Touch Sensory Substitution and Beyond. Rice University. Retrieved August 21, 2017 from <https://scholarship.rice.edu/handle/1911/88379>
 27. Valeria Occelli, Charles Spence, and Massimiliano Zampini. 2009. Compatibility effects between sound frequency and tactile elevation. *NeuroReport* 20, 8: 793–797. <https://doi.org/10.1097/WNR.0b013e32832b8069>
 28. Marc van Oostendorp, Colin J Ewen, Elizabeth Hume, and Keren Rice. 2011. The Blackwell companion to phonology. In *Sign syllables*. 1299–1344. <https://doi.org/10.1017/CBO9781107415324.004>
 29. Charlotte M. Reed, Nathaniel I. Durlach, Louis D. Braida, and Martin C. Schultz. 1989. Analytic Study of the Tadoma Method. *Journal of Speech Language and Hearing Research* 32, 4: 921. <https://doi.org/10.1044/jshr.3204.921>
 30. Henry L. Roediger and Jeffrey D. Karpicke. 2006. Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science* 17, 3: 249–255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>
 31. John G. Seamon and Pauline Murray. 1976. Depth of processing in recall and recognition memory: Differential effects of stimulus meaningfulness and serial position. *Journal of Experimental Psychology: Human Learning & Memory* 2, 6: 680–687. <https://doi.org/10.1037/0278-7393.2.6.680>
 32. Charles Spence. 2011. Crossmodal correspondences: A tutorial review. *Attention, Perception, & Psychophysics* 73, 4: 971–995. <https://doi.org/10.3758/s13414-010-0073-7>
 33. Charles Spence. 2014. The skin as a medium for sensory substitution. *Multisensory research* 27, 5–6: 293–312. <https://doi.org/10.1163/22134808-00002452>
 34. Robert Sweetow and Catherine V Palmer. 2005. Efficacy of Individual Auditory Training in Adults: A Systematic Review of the Evidence. *J Am Acad Audiol* 16: 494–504. Retrieved August 10, 2017
 35. Hong Z. Tan, Nathaniel I. Durlach, Charlotte M. Reed, and William M. Rabinowitz. 1999. Information transmission with a multifinger tactual display. *Perception & Psychophysics* 61, 6: 993–1008. <https://doi.org/10.3758/BF03207608>
 36. Hong Z Tan, Charlotte M Reed, and Nathaniel I Durlach. 2010. Optimum Information Transfer Rates for Communication through Haptic and Other Sensory Modalities. *IEEE Transactions on Haptics* 3, 2: 98–108. <https://doi.org/10.1109/TOH.2009.46>
 37. David Ternes and Karon E. MacLean. 2008. Designing Large Sets of Haptic Icons with Rhythm. In *Haptics: Perception, Devices and Scenarios*. Springer Berlin Heidelberg, Berlin, Heidelberg, 199–208. https://doi.org/10.1007/978-3-540-69057-3_24
 38. Shintaro Ueda, Masafumi Uchida, Akio Nozawa, and Hideto Ide. 2008. A tactile display using phantom sensation with apparent movement together. *Electronics and Communications in Japan, Part II: Electronics (English translation of Denshi Tsushin Gakkai Ronbunshi)* 29–38. <https://doi.org/10.1002/ecj.10000>
 39. E.H. Weber. 1996. De subtilitate tactus. *H. Ross & D. Murray, E. H. Weber on the tactile senses*: 21–128.
 40. S Weinstein. 1968. Intensive and extensive aspects of tactile sensitivity as a function of body part, sex and laterality. In *First International Symposium on Skin Senses*, 195–222.
 41. Elaine Y. Wong, Ali Israr, and Marcia K. O’Malley. 2010. Discrimination of consonant articulation location by tactile stimulation of the forearm. In *2010 IEEE Haptics Symposium*, 47–54. <https://doi.org/10.1109/HAPTIC.2010.5444681>
 42. GH Yang, D Ryu, S Park, and S Kang. 2012. Sensory saltation and phantom sensation for vibrotactile display of spatial and directional information. *Presence: Teleoperators and Virtual Environments* 21, 2: 192–202. https://doi.org/10.1162/PRES_a_00099
 43. T. Yoshioka, S. J. Bensmaïa, J. C. Craig, and S. S. Hsiao. 2007. Texture perception through direct and indirect touch: An analysis of perceptual space for tactile textures in two modes of exploration. *Somatosensory & Motor Research* 24, 1–2: 53–70. <https://doi.org/10.1080/08990220701318163>
 44. Siyan Zhao, Ali Israr, and Roberta Klatzky. 2015. Intermanual apparent tactile motion on handheld tablets. *IEEE World Haptics Conference, WHC 2015*: 241–247. <https://doi.org/10.1109/WHC.2015.7177720>