VSP

Advanced Robotics 0 (2008) 1–22

ADVANCED ROBOTICS
www.brill.nl/ar

Full paper

Human Perception of Pneumatic Tactile Cues

Yeongmi Kim [a,*], **Ian Oakley** [b] **and Jeha Ryu** [a]

[a] HMCI Lab., GIST, Buk-gu, Gwangju, South Korea
[b] ETRI, Yuseong-gu, Daejeon, South Korea

Received 31 July 2007; accepted 21 September 2007

Abstract

The ability to combine tactile, on-the-skin, sensory cues with the kinesthetic ones that widely available force-feedback devices are able to produce is a desirable one, opening the door to the production of more realistic, compelling virtual environments. Pneumatic air-jet displays can be easily mounted on existing force-feedback devices and we believe have the potential to provide useful complimentary tactile information. However, there is little concrete psychophysical data relating to pneumatic displays — a fact that hinders their adoption. This paper addresses this challenge, and presents brief psychophysical studies examining localization rate, the two-point threshold and the stimulus intensity of cues produced by pneumatic air jets. Using insights gained from these studies, we also present a more concrete application-focused investigation where we evaluate the effect of combining tactile and force-feedback cues in a simple object manipulation task in a virtual environment. We show that task completion times are significantly improved with the addition of tactile information, validating our ideas and suggesting this topic warrants further attention.

Keywords

Pneumatic tactile display, psychophysics, tactile array design, virtual reality buttons

1. Introduction

Widely available point-force haptic devices (such as the PHANToM [1]) provide high-quality kinaesthetic cues and are capable of displaying a range of haptic, or touch, information about virtual or simulated environments. Through motors, or other mechanical elements attached to carefully designed exoskeletal structures, they are capable of rendering the shape, compliance, friction and, crudely, the fine grain details and surface texture of computer-generated objects. However, a generally accepted weakness of these devices is that they are incapable of presenting

* To whom correspondence should be addressed. E-mail: kym@gist.ac.kr

DOI:10.1163/156855308X314524

information directly to the surface of the skin. In order to operate a point-force device a user must hold, in one way or another, a tool-like end-effector and experience the forces generated through this mediating implement. Typical end-effectors take the form of pen-like tools [2], graspable spheres [3] or thimbles into which one's finger is placed in order to feel the forces [1]. The physical aspects of this grasping behavior determine the tactile, on-the-skin, experience of the displayed haptic cues.

While this kind of tool-based interaction is sufficient for many situations and indeed in some cases (such as laparoscopic surgical simulations [4]) it is an entirely desirable, it can be a problem in others. One significant example of this is where it is the intention to present a haptic experience as if the user was unencumbered by equipment; to create the illusion that they are exploring the virtual world with their bare fingers and hands, and not through a mediating tool. In this interaction style, essentially enabled by using the thimble-like end-effectors mentioned above, it is desirable that users should experience not only the net reaction forces derived from touching objects in the virtual world, but also the deformations to, and changes to the pressure exerted on, the skin of their finger. The fact that this latter form of sensory cue is not conveyed with current technologies reduces the realism of the experiences they are capable of imparting and may reduce the effectiveness with which users are able to interact with virtual worlds. Certainly, studies examining the effect of gloves (which impair cutaneous perception) on the completion of physical tasks have found them to lead to substantial reductions in performance [5].

A number of previous authors have highlighted this issue in the virtual reality (VR) domain and attempts to address it have focused, perhaps unsurprisingly, on creating small fingertip displays that can be used in conjunction with force-feedback devices. For example, Wagner *et al.* [6] integrated a tactile pin array into their WAM (Whole-Arm Manipulator) force-feedback device. They describe a study using this system that investigates a user's ability to discriminate the compliance of virtual objects and found that the inclusion of tactile information led to increased levels of sensitivity. Debus *et al.* [7] describe a similar investigation of a combined force-feedback and tactile device. They embedded a vibrating element within the handle of their force-feedback hardware, essentially augmenting it with the ability to produce rudimentary tactile cues. They then investigated user performance in a simple tele-operation task when information in different sensory modalities was available to participants. They conclude that optimal performance was attained when users experienced a combination of visual, force-feedback and tactile cues.

In particular, one technology we believe deserves closer attention in this research domain is that of pneumatic air-jet displays. These are devices which produce tactile sensations by blowing jets of air on a user's skin. In this our motivations largely mirror those of Amemiya and Tanaka [8] in their description of a pneumatic display designed to be mounted on the base of the finger pad. Perhaps the most significant reason why air-jet displays are well suited to integration with a force-feedback device is because they do not require a mechanical assembly at the point of contact with the user's skin. Flexible, lightweight tubing can connect the point at which the air is impelled with force with the stimulus display locus under the user's fingertip. Furthermore, with a carefully designed system it is possible that this required tubing could be conveniently routed along the existing armature of the haptic device to allow the construction of a small, lightweight, dense fingertip display. However, this mechanical simplicity has a number of attendant disadvantages. One of these is that the length (and diameter) of the tubing used clearly effects the time between the initial display of a stimulus and it reaching the user. In some situations this kind of latency may be simply unacceptable.

Furthermore, skin stimulation through air-jets has not been studied as extensively as, for example, pressure or vibration. Thus, while we can state with confidence the localization rates and two-point limits for pressure stimuli delivered by pins pressed against the skin, it remains somewhat unclear what psychophysical parameters apply to the perception of streams of air delivered to the skin. While skin deformation definitely occurs, it is unlikely to do so in so explicable a manner as with the application of a directly quantifiable physical stimulus (in the form of a pin pushed upwards with a certain force or to a certain position). If the skin is placed tightly against the surface of the pneumatic display, then it is likely that a corresponding feeling of pressure will result. However, if the skin is not pressed tightly enough, some air will escape, altering the sensation, and if the skin is slightly raised then the sensation will likely differ considerably.

These concerns are echoed by Amemiya and Tanaka [8], who briefly describe a complex study that attempts to determine the two-point threshold for their pneumatic device and involves the variables of body site (thumb or forefinger), delivered air pressure, air-jet nozzle diameter and inter-jet spacing. Unfortunately, due to its complexity, it is hard to draw concrete conclusions from this study, beyond the simple observation that each variable appears to be capable of exerting a substantial effect on perceptual performance. Reflecting the effective absence of literature and the potential we see for such displays, this paper takes up this research challenge. We describe the design of a multi-element pneumatic display and a battery of psychophysical tests intended to serve as a practical guide for stimulus design for fingertip air-jet systems. As with Amemiya and Tanaka [8], the focus of our interest in these studies is not on the dynamics of individual jet performance, but instead on the perceptual practicalities of designing an array consisting of multiple air-jets; the basic questions of how such an array should be laid out on the skin to ensure that it appropriately matches human perceptual abilities. We conclude with an application-focused study integrating force-feedback and pneumatic stimuli, and speculate about the future of this work.

2. Pneumatic Display Hardware

2.1. System Architecture

Many different actuators have been adopted to provide tactile sensations for information transmission, entertainment, education and other purposes. Mizukami and

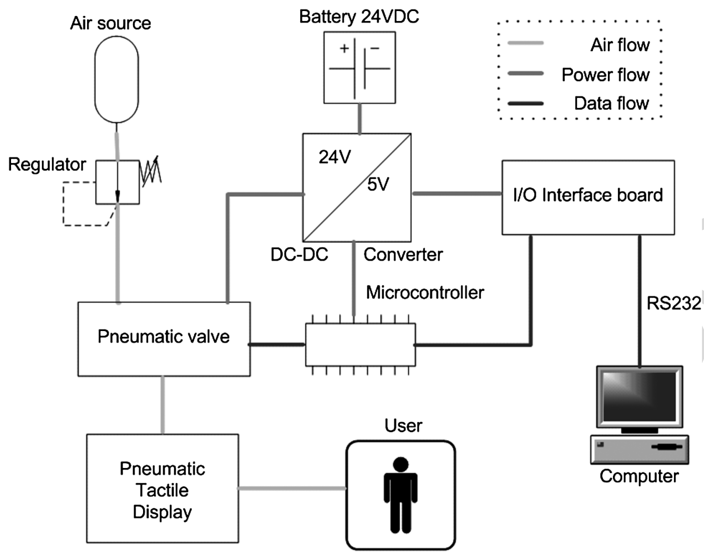

Figure 1. System configuration of the pneumatic tactile display.

Sawada proposed using a compact and low-power shape memory alloy actuator in a form of a string as a sensory aid tactile display for impaired and elderly users [9]. Kim *et al.* [10] describe a pin-array-type tactile display using piezoelectric actuators designed to provide surface sensations. In this work, they attached the tactile display to a PHANToM force-feedback device to render cues generated by a medical simulator focused on palpation tasks. Generally, tactile displays provide broadly similar functionalities: they all can apply a number of cues to the skin in a grid-like arrangement. However, depending on the characteristics of the actuators (and the requirements of the application), the evoked sensations can differ considerably. In our system, relatively simple hardware was deployed to provide air-jet sensations to users. Figure 1 illustrates the overall system architecture of the pneumatic device which consists of an air supply, regulator, I/O interface (and control board), rechargeable battery and pneumatic valves. As shown in Fig. 1, the air supplier provides pressurized air and the regulator keeps the pressure constant. The valves mediate the delivery of this pressurized air to the array according to control signals from an embedded processor (Mexx ATMega 128). A host PC used an RS232 serial connection to pass instructions to this microprocessor. Alongside its serial port, this chip features a total of 56 digital output lines; 25 of these were connected to (and controlled the state of) the binary valves forming the pneumatic array. The host computer sent sequences of bits to the ATMega indicating which of its outputs (and therefore the array jets) should be active. A full cycle of 25 such values occupies 4 bytes and each of these display-sized packets was prefixed with a header byte containing synchronization information. The total packet size was therefore 5 bytes. These packets were sent 500 times a second, corresponding to a total throughput of 2500 bytes or 20 000 bits per second, comfortably within

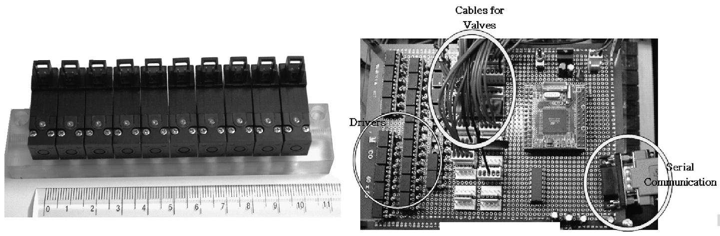

(a) **(b)**

Figure 2. Details of the hardware for the pneumatic system. (a) Size of 10 array valves. (b) Overall I/O interface and control circuit.

the bandwidth of the RS232 serial link. This communications protocol is similar to that adopted by other authors investigating tactile arrays [11]. In all the studies we presented users with air at a constant 1.034 bar. Figure 2 shows the valves used to control the flow of air (Yonwoo Pneumatic YSV10s), and the I/O interface and control circuit. These valves are binary devices (either on or off) with a switching latency of approximately 20 ms and a pressure range of up to 8.0 bar. An AQV214 photo MOS relay was used to control each valve individually.

2.2. *Design of the Pneumatic Display*

For the purposes of these studies, we constructed three prototype pneumatic arrays with different jet sizes, subsequently termed small, medium and large. The medium array was cut into a resin pad, whereas the other two arrays were formed using rapid prototyping tools and plastics. Thin plastic tubes embedded in these bases served as outlets for the air. Each array had the same basic design, and was intended to be mounted on the base of the index finger and held in place by an elastic sheath attached to its sides and tip and covering the top of the finger. It was constructed to snugly fit around the entire finger and this served to mechanically limit the variability of possible placements of a user's fingers against the array surface. The small array and medium arrays consisted of 25 individual air-jets arranged in a 5×5 grid. The large array consisted of nine air-jets arranged in a 3×3 grid, as a larger array would have extended beyond both the physical extent of the fingertip.

The entire device had the same dimensions for each array and the grid of air-jets had the same center point. However, spacing was not uniform. In each row of the array (running across the short axis of the finger pad) the air jets were mounted directly against one another, yielding an inter-stimulus distance of the external diameter of the tubes. However, small structural supports with a width of 0.8 mm were placed between the columns of the array (running down the long axis of the finger pad) resulting in a larger inter-stimulus spacing. The small array used tubing with an external diameter of 1.0 mm and an internal diameter of 0.5 mm; the overall array size was 5×8.2 mm. The medium-sized array featured jets with an external diameter of 2.4 mm and an internal diameter of 1.5 mm. Its total size was 12×15.2 mm. Finally, the large array used tubes with an external diameter of 4.5 mm and an internal diameter of 2.5 mm. As its dimensions were only 3×3 jets, its overall size was

RA
AR:m v 1.22 Prn:26/04/2008; 14:33
ar2565
P.6 (1-22)
by:Daiva p. 6

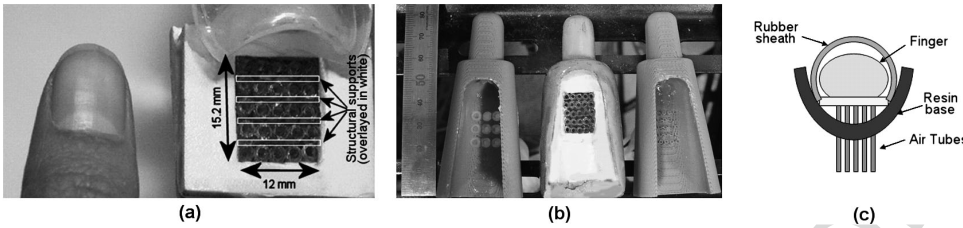

Figure 3. Details of the pneumatic array. (a) Initial prototype of the medium-sized array, overlaid with size information. (b) Final three arrays: large, medium and small arranged from left to right. (c) Schematic of array design.

13.5×15.1 mm. These three sizes were selected as the internal diameters increase linearly from 0.5 through 1.5 to 2.5 mm. It should be noted, however, that there are considerable variations in the external diameters and, correspondingly, in the ratio between these two distances. Figure 3 illustrates the three arrays.

3. Psychophysical Experimental Overview

In order to gain an understanding of the kinds of tactile cue that can be effectively displayed pneumatically, we engaged in a battery of brief tests intended to assess basic psychophysical properties: performance in localization, two-point threshold and magnitude summation tasks. The fixed physical structure of the three pneumatic arrays restricts the design of these studies: it is unsuitable for classic psychophysical methodologies such as the method of constant stimuli. Such methodologies require a more flexible and higher resolution set of stimuli, something challenging to achieve with air-jet displays, which are based on rigid tubing. Reflecting this, the approach adopted in this paper is to conduct a number of studies on each array and contrast the performance among these. In this way it is hoped to gain practical insights to inform future array design: what sizes and spacing of air-jets are best suited for this purpose. This approach can be likened to that taken by Tan *et al.* [12] in their informative description of the role human factors play in the design of force-reflecting interfaces.

3.1. Experiments and Participants

Three groups of subjects were used in these studies, subsequently termed groups A, B and C. These groups were used for purely logistical reasons — it afforded the flexibility (with respect to timetabling and participant fatigue) to conduct a wide range of studies rapidly. Each participant group featured eight members. Groups A and B were composed of four men and four women, while group C had six men and two women. The average age of the participants in each group was 22, 29 and 27 years, respectively. Group A completed four brief experiments, two investigating localization performance and two examining the two-point threshold of the mid-sized array. Group B completed a study looking at the perception of stimulus intensity. Finally, group C completed five experiments — three investigating localization performance and two examining the two-point threshold of small and large arrays.

3.2. Localization Studies

The five localization studies shared a similar interface and procedure, and could be differentiated by the size of the air-jets and the distance between the stimuli sites used in each. The intention was to contrast user performance with these different stimulus sizes and levels of inter-stimulus spacing. Two studies were conducting using the small array, two the medium and one the large. Each study was based on the presentation of nine stimuli (each an individual air-jet) arranged in a square pattern. Each of these cues was always presented a total of 20 times (leading to a total of 180 trials for each subject in each experiment) and in a random order.

For each of the three arrays, one study was conducted using the nine jets in the central 3×3 square (subsequently termed the dense studies). In the second studies using the small and medium arrays, eight of the stimuli were positioned on the array rim (the four corners and the four mid-points between them) and the ninth remained in the center (subsequently termed the spread studies). This is illustrated in Fig. 4 for the mid-sized array. For the small array, the horizontal and vertical distances between stimuli centers in the dense study were 1.0 and 1.8 mm, respectively. In the spread study these values were 2.0 and 3.6 mm. For the small and medium arrays these figures were 2.4 and 3.2 mm for the dense study, and 4.8 and 6.4 mm for the spread study. Finally, the study on the large array used horizontal and vertical spacing of 4.5 and 5.3 mm, respectively.

In all studies, each trial commenced with a screen instructing participants to press a key to begin. Upon completion of this action there was a 1-s pause, followed by 500 ms of stimulus presentation. Participants then had to press a key on the numeric keypad (the square arrangement typically situated on the right of a keyboard) to indicate the location of the displayed stimuli. The bottom left of this keyboard

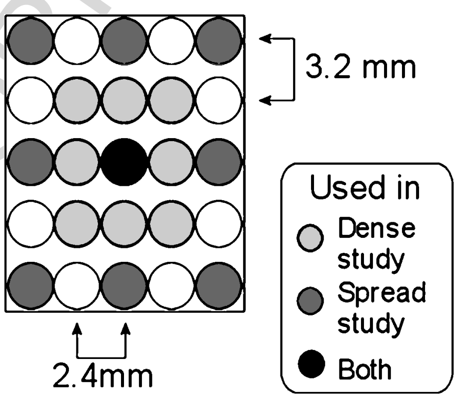

Figure 4. Diagram of the pneumatic array showing stimuli used in localization studies. Marked sizes refer to the medium-sized array.

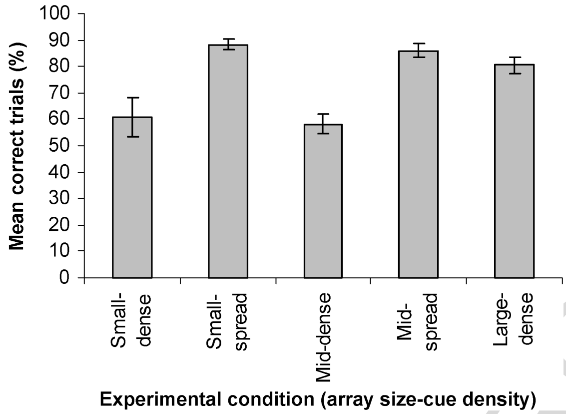

Figure 5. Mean percentage of correct trials in localization studies, shown by study.

(the key marked with the number 1) corresponded to a stimuli on the bottom left of the fingertip array. Similarly, the number 9, at the top right of the keypad, indicated a stimuli at the top right of the fingertip. The graphical interface to the study matched the spatial layout of the numeric keypad (and featured appropriate numbering), and after each trial graphical highlighting took place to indicate both the user's response and the correct answer. After this stage, a new trial began. Prior to each experimental session participants completed a practice session which was identical in structure, but consisted of only half the number of trials.

3.2.1. Localization Results

The localization rates from each study are presented in Fig. 5. A 2×2 ANOVA on the mean recognition rate from the small and medium arrays (array size *versus* cue density) showed an effect of cue density ($F(1, 15) = 41.14$, $P < 0.001$) but not array size ($F(1, 15) = 0.32$, $P = 0.57$), nor an interaction between these factors ($F(1, 1) = 0.001$, $P = 0.97$). To compare the performance of the large array, a one-way ANOVA on the mean recognition rate from the three dense conditions was also conducted, revealing a significant effect ($F(2, 23) = 6.09$, $P < 0.01$). *Post-hoc* t-tests showed differences between the large array and both the small and medium arrays (respectively, $t(15) = 2.54$, $P < 0.05$ and $t(15) = 4.68$, $P < 0.001$). In order to compare performance across the finger pad we conducted a brief analysis of the mean recognition rate by jet position. This data is shown in Fig. 6. This analysis did not reveal a significant effect ($F(8, 4) = 0.605$, $P = 0.76$).

3.2.2. Localization Discussion

The two-way analysis of the data from the small and medium arrays shows a clear effect of cue density, and none of array: the larger distances between stimuli in the medium array did not afford any advantages in the localization task. However, the large array appears to cross a critical threshold, as its dense arrangement offered significant increases in recognition rate when compared to the other two arrays. On the surface, this is a somewhat surprising result as all recognitions rates are

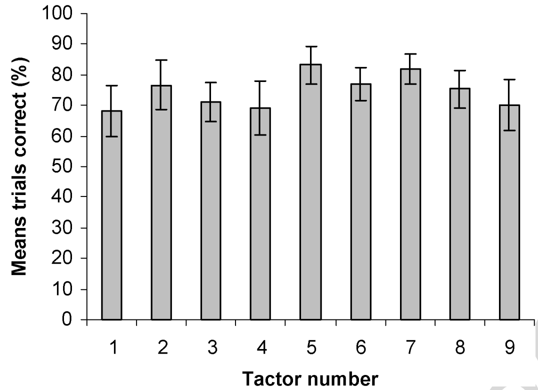

Figure 6. Mean percentage of correct trials in localization studies, shown by tactor.

considerably above chance (approx. 11%) and no particular jet location led to divergent performance. In this situation, a consistent increase in performance with array size would be a typical psychophysical relationship — an observation that suggests some other factor is involved.

Spatial acuity on the finger pad is a complex issue. For pressure stimuli produced by pin pricks distances of between 2 and 3 mm [13] are commonly reported to be practical limits. This is generally based on the receptive field size of the mechanoreceptors in the fingertip: between 2 and 3 mm for both Meissner (which detect 'flutter' and are therefore likely to be activated by air-jets to the skin) and Merkel receptors. However, figures of as low as 0.15 mm have been recorded using sophisticated experimental methodologies [14] — a result attributed to using the overlapping receptive fields of multiple receptor cells to triangulate a more precise position. In the dense studies with both the small and medium arrays conducted here (featuring stimulus sites with centers situated between 1 and 3.2 mm apart) the error rate hovered around the 60% mark; users experienced considerable difficulty with the task. These figures indicate that inter-cue spacing of in excess of 3 mm will be required to facilitate accurate localization. Indeed, the relatively high accuracies found in the medium and large arrays use a spacing of between 4.5 and 6.4 mm, suggesting that, practically, air-jet cues may be considerably harder to localize than pressure cues and require correspondingly higher separation distances.

However, the results from the spread study using the small array (with inter-center distances of 2 and 3.6 mm across and along the finger) are significantly higher than those from the dense study on the medium array (2.4 and 3.2 mm spacing), suggesting another factor is at play. The increase in the inner diameter of the tubes (from 0.5 to 1.5 mm) is a likely candidate. Essentially, an increase in cue diameter causes a corresponding drop in tactile acuity — smaller cues are more readily localized. This is an extremely likely hypothesis, based on the assumption that as a smaller cue activates fewer mechanoreceptors, it is therefore easier to identify its spatial position.

In conclusion, these three studies show clear advantages, in terms of the spatial resolution with which cues can be detected, for using air stimuli of the smallest possible dimensions. The 0.5-mm cues deployed in the small array could be reliably localized when arranged in a 3×3 grid only 6×10.8 mm. This is comparable to the spatial resolution results recorded using pressure cues [13], suggesting air-jets may be a viable alternative to these technologies. Larger stimuli require proportionally larger spacing to achieve equivalent performance levels and should be avoided if high localization performance is required.

3.3. Two-Point Studies

Three studies were conducted to gauge the two-point threshold for air-jet stimuli, one on each array. The experiments shared a similar task and procedure. In each, participants tapped a key to begin a trial and after a 1-s pause were presented with a stimuli which they then had to judge as either consisting of one or two separate jets of air. Feedback was given regarding their response, but not as to its correctness. Both studies were preceded by practice sessions with half the duration of the experimental sessions.

These studies took advantage of the physical constraints of our array design to present points of varying degrees of separation. Each used four stimuli, one of which was generated by a single air-jet, the remaining three being generated by a pair of jets. The pairs of stimuli were all adjacent, but differed in the directionality of this adjacency. As the array design is not uniformly spaced along its axes (as illustrated in Fig. 3 and due to the presence of structural supports between its rows, but not its columns), it features different inter-jet spacing between horizontally, vertically and diagonally adjacent jets. This discrepancy was use to generate stimuli pairs at different inter-stimuli spacing, as shown in Fig. 7 for the mid-sized array. The distances between jet centers for pairs in these three orientations were 1,

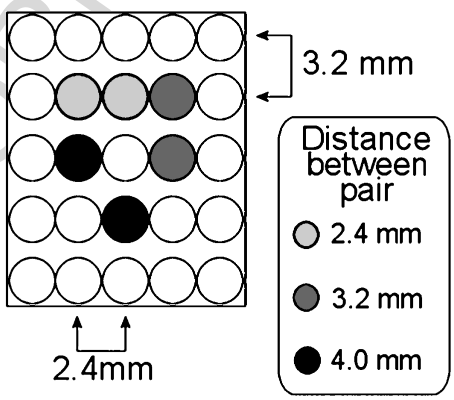

Figure 7. Diagram of the pneumatic array showing a sample set of stimuli used in two-point threshold studies. Marked sizes refer to the medium-sized array.

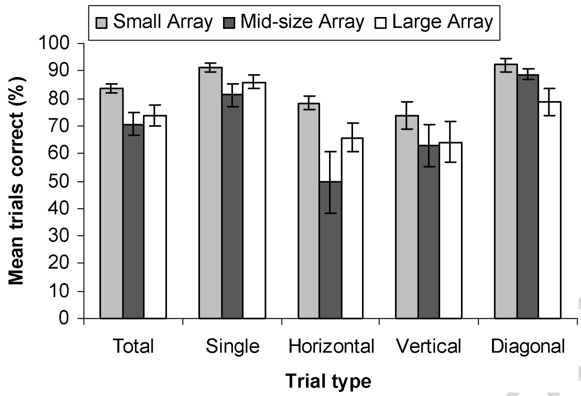

Figure 8. Mean percentage correct trials from two-point threshold studies. Includes overall mean and mean by each stimulus type.

1.8 and 2.06 mm for the small array, 2.4, 3.2 and 4.0 mm for the mid-sized array, and 4.5, 5.3 and 6.95 mm for the large array.

To control for possible response biases in these studies, each of the three stimuli pairs was presented 30 times, while the single-jet stimulus was presented a total of 90 times. This equalized the number of times participants were exposed to individual and pair stimuli, and led to a total of 180 trials. All stimuli were presented on the central 3×3 portion of each array and randomized for the physical limits of this configuration. This led to nine possible locations on which to display the single jet — six for each of the horizontal or vertical pairs and eight for the diagonal pair.

3.3.1. Two-Point Results

Figure 8 illustrate the accuracy with which participants correctly judged the stimulus in each trial as being composed of one or two points in all three experiments. It includes the total mean accuracy rate and a breakdown of this into the mean from each of the four types of cue. A two-way ANOVA analyzing cue type against array size revealed significant trends in both main effects: cue-type ($F(3, 23) = 14.46$, $P < 0.001$) and array size ($F(2, 23) = 6.51$, $P < 0.005$). No interaction was found ($F(3, 2) = 1.41$, $P = 0.219$). *Post-hoc* t-tests revealed that the small array offered increased accuracy when compared to the other two arrays (both at $P < 0.005$). No difference was found between the mid-sized and large arrays ($P = 0.28$). Cues composed of a single jet and those composed of a pair of diagonally arranged jets offered significant improvements in accuracy compared to cues formed by pairs of either horizontal or vertical jets (all at $P < 0.001$). No other differences were observed.

3.3.2. Two-Point Discussion

Accuracy in this experimental task increased with the two most distinguishable cues. The single jet and the diagonally separated pair were correctly identified more than 85% of the time in all studies. The other two pairs were identified with less than 70% accuracy — a figure made more dramatic by the fact that chance in this

RA
AR:m v 1.22 Prn:26/04/2008; 14:33 ar2565 P.12 (1-22)
by:Daiva p. 12

task is at 50%; participants clearly experienced some difficulty. The fact that this occurred in all three studies, despite the fact they used different stimulus sizes and spacing, suggests that larger jet sizes yield higher two-point thresholds. This is confirmed by the fact that performance was also significantly better using the small array. Together these findings strengthen the conclusions from the localization study that smaller jets offer better performance in fundamental tasks based on uniquely identifying them on the skin. Indeed, in the literature, the two-point threshold for pressure stimuli generated by pin-pricks to the skin has been reported to be as low as 1–2 mm [13]. This figure is broadly consistent with the data from the small array (although the two larger arrays offer considerably poorer performance), once again suggesting that appropriately configured air-jet displays can provide cues with the same resolution and accuracy as the more commonly studied pressure displays.

In conclusion, participants performed significantly better with the small array than the other two, despite the fact this features the smallest distances between pairs of stimuli. This strongly supports the superiority of smaller cues over larger. However, it is worth noting that the ability to distinguish pairs of stimuli or perform accurate localization may not always be desirable. In some instances, such as when rendering contact with an object, delivering a sensation to an area of the skin, rather than to a set of discrete points, may be appropriate. This issue is examined in the next study reported in this paper.

3.4. Stimulus Intensity

Unlike many forms of tactile display, one current technological limitation of most air-jet displays is that they are unable to present stimuli of different intensities. This is largely due to the fact that most valves (which are used to control the airflow), and especially those that are small, reasonably priced and quiet, are binary. They are either off or on and consequently so are the stimuli that they can create. However, delivering cues that have the capability to grow or shrink in magnitude remains an attractive goal. One way this can be achieved with a display composed of binary elements is by changing the size of the skin area to which a stimulus is applied by changing the number of array elements used to display it. At threshold levels, a relationship between the size of skin contactor and the perceived magnitude of a displayed vibrotactile cue has been established by Verrillo [15], more recently investigated by Oakley *et al.* [11] and similar effects are well documented in other sensory modalities, such as vision [16].

Reflecting this work, we conducted a study intending to gauge whether participants can effectively discriminate between stimuli composed of different numbers of simultaneously active air-jets. In fact, given our previous localization and two-point threshold results, it is reasonable to suggest that by activating groups of adjacent air-jets, we can in fact stimulate different sized areas of the skin. Based on this perceptual assumption we conducted a study in which participants experienced a pair of stimuli, each created by activating between one and nine air-jets, and then had to judge which of these two was of greater magnitude or if they were the identical. This study used the medium-sized array and the air-jets were drawn from its central 3×3 portion. Both localization and two-point performance using this array configuration was observed to be relatively poor in the previous studies, suggesting cues in this arrangement are sufficiently dense as to approach spatial thresholds. Each of the nine possible magnitudes was twice compared against all cues, leading to a total of 162 trials (nine magnitudes $\times$ nine magnitudes $\times$ two presentations). Each stimulus was always produced by a single arrangement of one or more adjacent tactors (unlike in the studies on two-point localization which explicitly varied which jets stimuli were displayed on).

In this study, each trial consisted of a subject depressing a key to begin, a 1000-ms pause, followed by a 500-ms stimulus presentation, another 1000-ms pause and a final 500-ms stimulus presentation. Participants were required to respond by pressing keys corresponding to whether they thought the first stimulus was greater than the second, that they were the same or that the second was greater than the first. An on-screen interface reinforced these instructions and highlighted user responses, but not their correctness. Immediately prior to the experimental session, participants completed an 81-trial practice session.

3.5. Stimulus Intensity Results

The data recorded in this study consisted of pair-wise judgments comparing the perceived intensity of each individual stimulus against the full set of stimuli. This data was analyzed by tallying the number of times a given stimulus was rated as being greater, the same or less than another; as perceived intensity increases or decreases, these three counts should similarly increase or decrease. These data are shown in Fig. 9. We confined our analyses to the number of times each stimulus was rated as greater than another. An ANOVA indicated the variations in this statistic attained significance ($F(8,7) = 129.65$, $P < 0.001$), while *post-hoc* t-tests incorporating

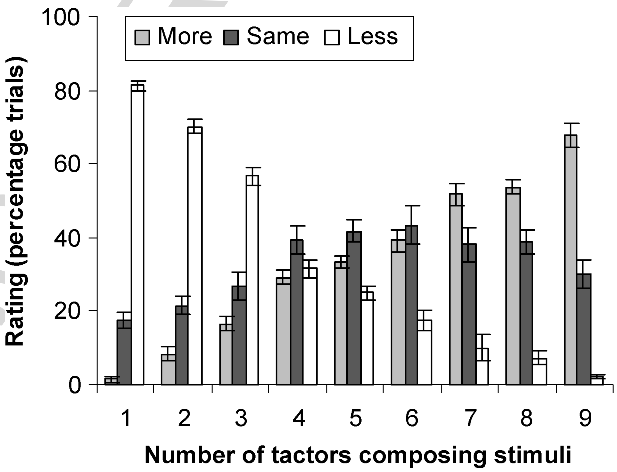

Figure 9. Ratings data from the intensity study.

Table 1.
P values from t-tests in the intensity study comparing how often each stimuli is rated as more intense than the other (italics indicate significant results)

Number of air jets in second stimuli	Number of air jets in first stimuli							
	2	3	4	5	6	7	8	9
1	0.08	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*
2		0.46	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*
3			*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*
4				1.0	0.64	*0.01*	*0.01*	*0.01*
5					1.0	*0.01*	*0.01*	*0.01*
6						0.07	0.14	0.01
7							1.0	0.12
8								0.20

Bonferroni confidence interval adjustments revealed a large number of significant differences, summarized in Table 1. These results strongly suggest that by varying the size of a pneumatically created tactile stimulus, we can effectively vary its perceived magnitude. This finding is made especially valuable in light of the fact that the individual jets within our array are only capable of producing a stimulus of a single magnitude. By demonstrating the viability of an alternative mechanism to vary the intensity of a cue, we open the door to the production of a whole new range of dynamically changing stimuli.

3.6. General Discussion

The motivation underlying these psychophysical tests was to attempt to determine the perceptual qualities of the stimuli we can produce with our pneumatic display. Our overall goal is to integrate this tactile device as a fingertip display mounted on the armature of a force-feedback device, such as the PHANToM [1]. A thorough understanding of the nature of the cues we can deliver will be important in ensuring that we can create constructive, rather than destructive, pairings of tactile pneumatic cues and kinesthetic force-feedback cues. Although not a complete account, the studies we report provide valuable insights as to how this might be done. We have concrete data regarding localization, the two-point threshold and the ability of participants to discriminate the size (or intensity) of presented stimuli. To demonstrate how this psychophysical knowledge can be applied to a practical scenario, in the next section we describe a simple study investigating the utility of combining pneumatic tactile and force-feedback cues to a simple object manipulation task in a virtual environment.

4. Applied Study — VR BUTTON

Haptically rendered VR buttons are three-dimensional virtual objects that can (through the medium of a force-feedback device) be pushed and respond to this action by moving in a manner similar to that of a button in the real world. After an initial resistance is overcome, the button moves back freely for a short distance before coming to a hard stop. This dynamic behavior creates a click-like sensation — a potentially valuable piece of feedback. While it is possible to produce such buttons relatively easily using commercially available software toolkits [17], they are not a common feature of virtual environments. Raymaekers and Coninx [18] report one reason for this. They describe a study of a selection task in which participants activated VR buttons by either pushing them (as detailed above) or by using a dedicated physical switch (like a mouse button) positioned on the tip of the end-effecter of a force-feedback device. This latter condition led to a 50% reduction in task completion times, strongly suggesting that requiring VR buttons be pushed in order to be activated is a far from optimal interaction technique.

Despite this negative evidence, creating virtual buttons that behave like real ones remains a relatively desirable goal. If the intention is to immerse users in an environment, or realistically display some complex virtual object to them, then imbuing objects with real-world dynamics is a requirement. Although the use of a dedicated hardware button may be more efficient, it is far from a physically realistic approach and seems likely to destroy any illusions of the realism of the environment. Furthermore, in the kind of situations we are considering, in which a user interacts by placing their entire finger in an enclosed haptic display, it is unclear where a physical switch could be mounted.

We suggest that combining pneumatic tactile and force-feedback cues relating to button pushing may increase user performance. As inspiration for this, we cite Bicchi *et al.*'s [19] work on tactile flow. One of the fundamental demonstrations of this concept involves improving the accuracy of a feeling of contact by displaying the same net force to a greater or lesser area of skin. By appropriately adjusting the amount of skin being subjected to a stimulus, Bicchi *et al.* were able to create a more realistic percept. Similarly, we may be able to create a more realistic, and potentially easier to use, version of a haptic button by using our pneumatic display to deliver appropriately changing stimuli to the surface of the skin as the button moves through the process of being pressed.

4.1. Experimental Task

The experimental task in this study was the selection of cube-shaped buttons in a VR environment through the mechanism of moving to the button's surfaces and pushing against them. There were nine buttons arrayed horizontally on a flat base at the bottom of the virtual environment and participants interacted with them by pushing down on their top surfaces. The buttons moved in response to the user's efforts and a button press event occurred after a certain travel distance. Three conditions were considered. In the first, No-Air, the only feedback presented to users

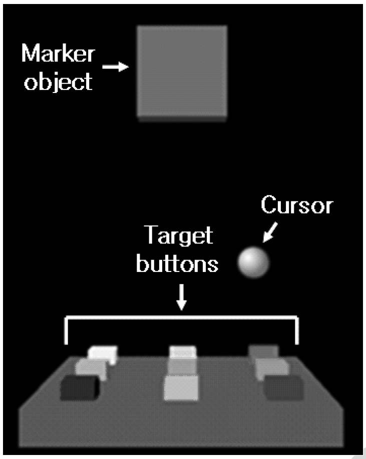

Figure 10. Experimental interface used in the button study. The nine buttons are arranged on the grey object at the base of the screen, the large red object at the top is the marker and the yellow spheres is the cursor.

was in the form of kinesthetic cues from the point force haptic device as it rendered contact with the dynamically moving virtual buttons. In the second, Binary-Air, this feedback was accompanied by the uniform activation of the entire pneumatic array on contact with the button. The third, Multi-Air, also featured pneumatic feedback, but varied the magnitude of this according to the depth of the button's travel — the further the travel, the more active jets.

The graphical interface to the study is shown in Fig. 10. Each trial consisted of a single button press and color was used to designate the current target — each of the nine buttons had a different color (the three primary colors, the three combinations of pairs of primaries, and black, white and orange). The trials commenced with the participant briefly touching a large marker object at the top of the virtual environment. In response to this, it changed color (from a neutral grey) to match one of the buttons, indicating which should be pressed. After a successful button press, both the marker and the target button became grey until the start of the next trial.

4.2. Experimental Design and Measures

This study used a repeated measures design — all participants completed all three feedback conditions. It was fully balanced design with six order conditions. Each feedback condition was composed of 90 trials (10 for each button), delivered in a random order and preceded by an identically structured half-length practice session. There were 12 subjects, six female and six male, with a mean age of 27. All were right-handed and none reported any significant impairment in either their haptic sensory–motor system or their color vision. Two subjects completed each order condition. One subject was a member of Group C from the previous studies, while the remaining 11 only participated in this study.

Aspects of both task completion time and error rate were measured. Time was decomposed into three stages: Touch time, Click time and Release time. Touch time encompassed the period from the start of a trial (touching the marker object) until contact was made with the target button. In cases where the target button was touched (and then released) multiple times during a trial, Touch time was set to be the last of these occurrences. Click time spanned from the Touch time until the button was successfully selected and Release time spanned from this selection until the user moved off the button. Error measurement consisted of recording the number of button touches that occurred prior to a successful selection on the basis that such behavior would indicate instability and uncertainty in the button-targeting procedure.

4.3. Hardware

This experiment was conducted using a modified version of the mid-sized pneumatic array featuring jets with an external diameter of 2.4 mm (and an internal diameter of 1.5 mm) and a PHANToM force-feedback device [1]. The modifications to the array were simply to add a short finger-like protrusion to its tip, which was designed to fit snugly inside the standard thimble attachment of the PHANToM. Thus, we were able to augment the normal thimble mode of interaction with the PHANToM (which involves placing a finger inside the thimble) with our pneumatic array with little effort. Problems that might emerge with this *ad-hoc* solution are that the additional weight it adds to the tip of the PHANToM is unbalanced and the fact that the user's finger ends up positioned approximately 2 cm in front of the device's fulcrum — the point at which it is designed to accurately apply the forces it generates. These issues are unresolved at the current time, but we believe sufficiently minor that a meaningful empirical investigation can be undertaken with this hardware configuration. We regard this device as a proof-of-concept prototype, and suggest that considerable additional development would be required to create a robust, reliable and generally applicable device with the features our preliminary system possesses. The combination of modified pneumatic and PHANToM force-feedback device is shown in Fig. 11.

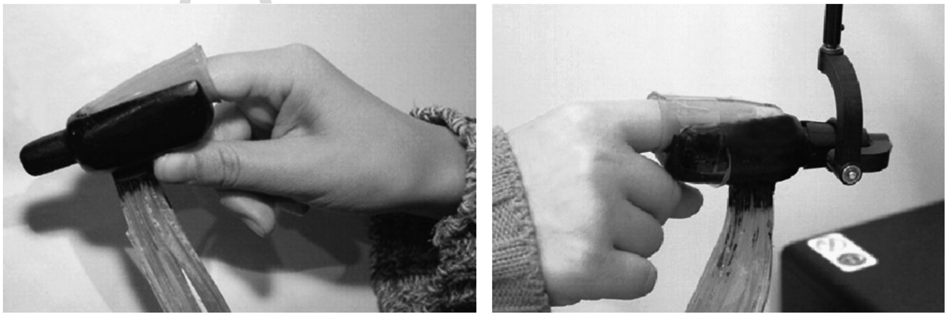

Figure 11. Pneumatic array housing for easy attachment to the PHANToM force-feedback device.

4.4. Materials

The movement (and kinesthetic feel) of the virtual buttons used in this study was generated by a simple model consisting of a pair of virtual springs. The first of these was a relatively strong spring active during the initial portion of the button's movement, the second a weak spring active during the latter portion. Beyond the range of both these springs, the button's motion simply stopped. This profile results in the following click-like behavior. An initial effort is required to push a button, which (after moving a short distance) drops away rapidly until it reaches a hard stop. This is characteristic of other virtual buttons reported in the literature [18].

Two different pneumatic cues sets were used in this study. The first, the Binary-Air condition, involved simultaneous activation of all 25 jets in the array on contact with a button. The second, used in the Multi-Air condition, adopted a more complex paradigm in which the distance the button had moved mediated the number of active air jets. From contact and for the first third of the button's travel only the central air-jet was activated; for the second third, the central nine air-jets were on; for the final third, all 25 air-jets were on. This manipulation leverages the result reported in the intensity study suggesting that perceived stimulus magnitude increases when additional active jets are active [20]. By offering more intense tactile stimuli as the button is pushed towards the limits of its travel, it may be possible to more realistically mimic the process of pressing against a button in the real world.

4.5. Results

The results from the three measures of task time are shown in Fig. 12. A repeated measure ANOVA was used to analyze each of these results. Significant effects of Touch time ($F(2, 11) = 27.642$, $P < 0.001$), Click time ($F(2, 11) = 4.385$, $P < 0.025$) and Release time ($F(2, 11) = 20.882$, $P < 0.001$) were found. *Post-hoc t*-tests incorporating Bonferroni confidence interval adjustments showed that the Binary-Air and Multi-Air conditions yielded improved Touch times (both

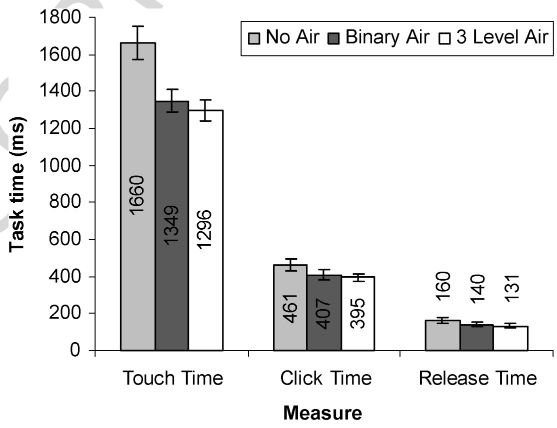

Figure 12. Mean task completion time in the button study.

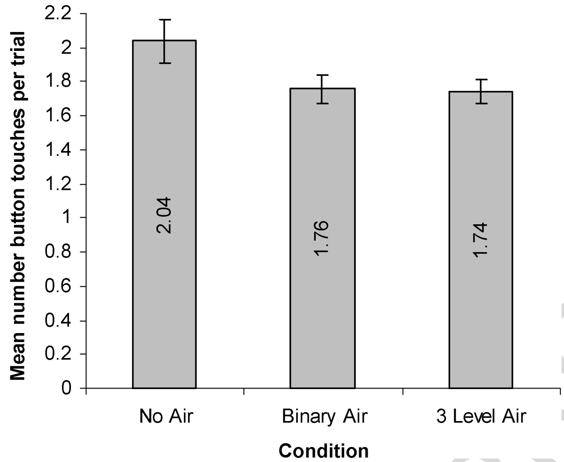

Figure 13. Mean number of button touches per trial in the button study.

$P < 0.001$) and Release times (both $P < 0.001$), but not Click times ($P = 0.294$, $P = 0.054$). Multi-Air offered an improvement over Binary-Air in Release time ($P < 0.05$) but not Touch time ($P = 0.093$) or Click time ($P = 1$). The mean number of button touches in each trial is shown in Fig. 13. A repeated measures ANOVA showed significant differences ($F(2, 11) = 8.578$, $P = 0.002$) and post-hoc t-tests bore these out in the form of significant comparisons between No-Air and both of the other two conditions (both $P < 0.05$).

4.6. Discussion

This study provides simple, but compelling evidence that the combination of proprioceptive, force-feedback cues and cutaneous pneumatic ones can be constructive and yield concrete performance improvements. The addition of even simple binary tactile cues decreased total task completion times by about 20% in a targeting task. Given the prevalence of such fundamental operations in computer use, it seems likely that this improvement would translate into substantial qualitative and quantitative improvements in more complex and realistic scenarios. The vast majority of complex interactive computer systems are based on users making a large number of targeting and selection operations, and improving performance in such atomic tasks is an important method with which to enhance overall system usability.

Although both Touch and Release times were improved (by approximately the same ratio) with the addition of pneumatic cues, the fact that the Release time was relatively brief means that the bulk of the temporal improvement occurred during the Touch time. This results from this stage, covering the time from the start of a trial until the last time the target button was touched prior to a successful selection, were also reflected in the error data. The No-Air condition showed a significant increase in the number of times a button was touched before a successful selection was made and it seems likely that it was this factor that caused the slower Touch times. Participants were able to use the tactile pneumatic cues to more easily determine when they had made contact with a button, which in turn resulted in a more rapid selection. This same factor may also have caused the reductions in the Release time — the tactile cues may have given subjects sufficient confidence that they were in contact with the button that they performed the act of pressing it relatively autonomously and rapidly, in a gestural manner. In contrast, lacking this feedback in the No-Air condition, they made have been spending more time explicitly monitoring their actions, resulting in slower performance.

The only difference detected between Binary-Air and Multi-Air was that the latter offered an improvement in the Release time. This may indicate that the gradually increasing feedback allowed subjects to more reliably detect when they had completed the button press, which in turn let then finish the interaction more rapidly. Such an explanation appeals to common sense — the added value of a gradually escalating cue should appear not immediately after it is presented, but after sufficient time (or distance) has passed for its rise to be observed. However, such a conclusion may go beyond what can be reasonably support by the data — the difference between the Release times in these two conditions is only 9 ms and although this is approximately 7% of the total, it remains an extremely short period of time.

5. Conclusions and Future Work

In this paper we described the motivations and design of a system that enables the combination of force-feedback and pneumatic tactile cues. Through a series of psychophysical studies, we developed a basic understanding of the perceptual characteristics of pneumatic stimuli. We gained insights into localization rates, the two-point threshold and showed magnitude summation below spatial thresholds. Taking this information on board, we describe a final study which combines proprioceptive and pneumatic cues, and demonstrates even very simple cues can lead to performance improvements and potentially improve the usability of virtual environments.

Development of a better integrated and more robust hardware platform is an imperative for this research; our current device is no more than a proof-of-concept prototype. Additional future work will involve extending the idea of combining pneumatic tactile and kinesthetic cues to more challenging scenarios. Possibilities we are considering include further fundamental studies using more rigorous procedures and more sensitive metrics and assessment protocols, and the enhancement of more complex virtual environments, either featuring larger and more diverse control elements, or those which attempt to achieve greater levels of realism, e.g., in the entertainment industry. We believe that the scope of this technology extends to both gaming and narrative entertainment systems, potentially enriching many kinds of digital content and eventually reaching away from the desktop and into the living room.

Acknowledgements

This research was supported by the Republic of Korea's Ministry of Information and Communication through the Realistic Broadcasting IT Research Center at the Gwangju Institute of Science and Technology.

References

1. T. H. Massie, Initial haptic explorations with the PHANToM: virtual touch through pointer interaction, *Masters Thesis*, Massachusetts Institute of Technology, Boston, MA (1996).
2. R. Q. Van der Linde, P. Lammertse, E. Frederiksen and B. Ruiter, The HapticMaster, a new high-performance haptic interface, in: *Proc. EuroHaptics*, Edinburgh, pp. **00–00** (2002).
3. J. Murayama, L. Bougrila, Y. L. Luo, K. Akahane, S. Hasegawa, B. Hirsbrunner and M. Sato, SPIDAR G&G: a new two-handed haptic interface for bimanual VR interaction, in: *Proc. EuroHaptics*, Munchen, pp. **00–00** (2004).
4. R. Webster, R. S. Haluck, G. Zoppetti, A. Benson, J. Boyd, N. Charles, J. Reeser and S. Sampson, A haptic surgical simulator for laparoscopic cholecystectomy using real-time deformable organs, in: *Proc. IASTED Int. Conf. of Biomedical Engineering*, **LOCATION?**, pp. **00–00** (2003).
5. J. Chodack and P. Spampinato, Spacesuit glove thermal micrometeoroid garment protection *versus* human factors design parameters, in: *Proc. SAE Conf. on Space Station and Advanced EVA*, **LOCATION?**, pp. **00–00** (1991).
6. C. R. Wagner, O. Clatz, D. P. Perrin, R. L. Feller, H. Delingette, N. Ayache and R. D. Howe, Integrating tactile and force feedback with finite element models, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, Barcelona, pp. **00–00** (2005).
7. T. Debus, T. Becker, P. Dupont, T. Jang and R. Howe, Multichannel vibrotactile display for sensory substitution during teleoperation, in: *Proc. SPIE Int. Symp. on Intelligent Systems and Advanced Manufacturing*, Boston, MA, pp. **00–00** (2001).
8. K. Amemiya and Y. Tanaka, Portable tactile display using air jet, in: *Proc. Int. Conf. on Artificial Reality and Telexistence,* Tokyo, pp. **00–00** (1999).
9. Y. Mizukami and H. Sawada, Tactile information transmission by apparent movement phenomenon using shape-memory alloy device, *Int. J. Disabil. Hum. Dev.* **5**, 277–284 (2006).
10. S. Kim, K. Kyung, J. Park and D. Kwon, Real-time area-based haptic rendering and the augmented tactile device for a palpation simulator, *Adv. Robotics* **21**, 961–981 (2007).
11. I. Oakley, Y. Kim, J. Lee and J. Ryu, Determining the feasibility of forearm mounted vibrotactile displays, in: *Proc. IEEE Haptics Symp.*, **LOCATION?**, VA, pp. 27–34 (2006).
12. H. Z. Tan, M. A. Srinivasan, B. Eberman and B. Cheng, Human factors for the design of force-reflecting haptic interfaces, in: *Proc. ASME Dynamic Systems and Control Division Conf.*, **LOCATION?**, pp. **00–00** (1994).
13. K. O. Johnson and J. R. Phillips, Tactile spatial resolution — I. Two point discrimination, gap detection, grating resolution and letter recognition, *J. Neurophysiol.* **46**, 1177–1191 (1981).
14. J. M. Loomis, An investigation of tactile hyperacuity, *Sensory Process.* **3**, 289–302 (1979).
15. R. T. Verrillo, Effect of contactor area on the vibrotactile threshold, *J. Acoust. Soc. Am.* **35**, 1962–1996 (1963).
16. R. Sekuler and R. Blake, *Perception*. McGraw-Hill College, NEW York (2002).
17. Reachin Technologies AB, January (2006). Available at http://www.reachin.se/
18. C. Raymaekers and K. Coninx, Menu interactions in a desktop haptic environment, in: *Proc. EuroHaptics*, **LOCATION?**, pp. **00–00** (2001).

19. A. Bicchi, E. P. Scilingo and D. De Rossi, Haptic discrimination of softness in teleoperation: the role of the contact area spread rate, *IEEE Trans. Robotics Automat.* **16**, **00–00** (2000).
20. Y. Kim, I. Oakley and J. Ryu, Combining point force haptic and pneumatic tactile displays, in: *Proc. EuroHaptics*, Paris, pp. 309–316 (2006).

About the Authors

Yeongmi Kim received her BS degree in Control and Instrumentation Engineering at Korea Maritime University, South Korea, in 2004, and MS degree in the Mechatronics Department, Gwangju Institute of Science and Technology (GIST), South Korea, in 2006. She is a PhD student in the Mechatronics Department, GIST. Her research interests are psychophysics, tactile display design, haptic content authoring and haptic interaction in broadcasting.

Ian Oakley holds a PhD in Human Computer Interaction (2003) and a Joint Honours (First Class) in Computing Science and Psychology (1998), both from the University of Glasgow, UK. He currently works as a Researcher at the Electronics and Telecommunications Research Centre in Daejeon, South Korea and has held previous posts at the Gwangju Institute of Science and Technology (South Korea) and MIT Media Lab Europe (Dublin, Ireland). He is a member of the ACM Special Interest Group on Computer Human Interaction. His research interests focus on multi-sensory interfaces.

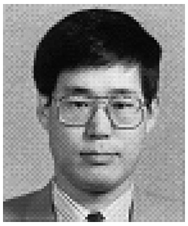

Jeha Ryu received his BS degree at Seoul National University, South Korea, in 1982, MS degree at the Korea Advanced Institute of Science and Technology, South Korea, in 1984, and PhD degree at the University of Iowa, USA, in 1991. He has been a Professor in the Mechatronics Department, Gwangju Institute of Science and Technology, since 1994. His research interests include haptics, haptic rendering, haptic interaction control, design and control of haptic devices, as well as haptic-included modeling and broadcasting.