INTCOM 1630
23 November 2007 Disk Used
ARTICLE IN PRESS
No. of Pages 10, Model 5+

Available online at www.sciencedirect.com
ScienceDirect
Interacting with Computers xxx (2007) xxx–xxx
Interacting with Computers
www.elsevier.com/locate/intcom

Did you feel something? Distracter tasks and the recognition of vibrotactile cues

Ian Oakley *, Junseok Park

Q1 *Electronics and Telecommunication Research Institute, 161 Gajeong Dong, Yuseong Gu, Daejeon 305-700, Republic of Korea*

Received 12 July 2007; received in revised form 14 September 2007; accepted 3 November 2007

Abstract

Research on vibrotactile displays for mobile devices has developed and evaluated complex multi-dimensional tactile stimuli with promising results. However, the possibility that user distraction, an inevitable component of mobile interaction, may mask (or obscure) vibrotactile perception has not been thoroughly considered. This omission is addressed here with three studies comparing recognition performance on nine tactile icons between control and distracter conditions. The icons were two dimensional (three body sites against three roughness values) and displayed to the wrist. The distracter tasks were everyday activities: Transcription, mouse-based Data-entry and Walking. The results indicated performance significantly dropped in the distracter condition (by between 5% and 20%) in all studies. Variations in the results suggest different tasks may exert different masking effects. This work indicates that distraction should be considered in the design of vibrotactile cues and that the results reported in lab based studies are unlikely to represent real world performance.

Keywords: Wearable computing; Mobile; Haptic interface; Tactile icons; Tactons

1. Introduction

Simple vibrotactile cues, in the form of a buzz when a message or call is received, are a standard and useful feature of mobile phones, the most widely available type of handheld computer. Their success can be attributed to their firm fit with the usability constraints of one of the most common phone tasks: signifying alerts. They can provide attention grabbing notifications to users engaged in unrelated tasks (which screen-based visual cues cannot) and do this discreetly (which speaker-based audio cues cannot) and without explicitly interrupting users.

A number of authors (Brown et al., 2006; Brown and Kaaresoja, 2006; Chang and O'Sullivan, 2005) have suggested that vibrotactile cues can play a greater role in mobile interfaces. They have focused on increasing the expressiveness of the vibrotactile cues, to enable systems which can convey more than the binary information required to indicate the arrival of a call or message. One key motivation for this work is that multi-modal interfaces may be better suited to many mobile scenarios than the currently dominant point and click graphical interfaces which evolved for desktop computing. Essentially, people look, listen and feel as they move and interpret the world around them through the course of their daily activities. Similarly, a device to which you can look, listen and feel may well be much more useful than one which demands your visual attention to complete even simple tasks.

However, there has little consideration of the other crucial aspect of mobile interaction: the environment (Pirhonen et al., 2002). Work on the design of vibrotactile cues is now approaching the point at which designers or system developers can use it to select a stimulus set and be confident that users will be able to reliably distinguish between its members. However, it still remains unclear whether they can do this out and about, performing tasks in the real

---

* Corresponding author. Tel.: +82 428603975; fax: +82 428605545.
E-mail addresses: ian@etri.re.kr, ian@whereveriam.org (I. Oakley), parkjs@etri.re.kr (J. Park).

0953-5438/$ - see front matter © 2007 Published by Elsevier B.V.
doi:10.1016/j.intcom.2007.11.003

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003

world, or only within a controlled lab environment. It is not certain that the performance levels reported in the literature will be maintained whilst users are engaged in common activities such as walking down a street, holding a conversion, working at a computer or driving a car.

The work in this paper begins to address this issue. It evaluates recognition performance with a two-dimensional set of vibrotactile icons (or tactons) displayed on a wearable wrist display while users are performing three different distracter tasks. Two tasks involve completing work at a computer (one is Transcription, the other mouse-based Data-entry), while the final one is mobile and has the participants walking around. The tasks were chosen to represent common activities and explore different aspects of distraction. For instance, the Transcription and Data-entry tasks are likely to consume considerably more mental resources than the walking task. Similarly, the transcription task involves small, rapid, controlled movements of the hand and wrist on which the tactile display is mounted, while the Data-entry and Walking tasks do not. Exploring these varied scenarios may offer a more nuanced view of the effects of distraction, and this kind of systemic, in-context evaluation offers a significant contribution to the emerging body of work on tactile cues for mobile devices. Awareness of these issues is of direct relevance for anyone designing a mobile application featuring vibrotactile cues, and work on this topic is required before rich vibrotactile output will successfully transition from the lab to the streets.

2. Related work

Researchers have long acknowledged that the skin is a valuable and under used conduit for information, and there is a substantial and growing body of work on vibrotactile perception and display systems (e.g. Cholewiak and Collins, 2003; Cholewiak et al., 2004; Jones et al., 2006). Some of the earliest practical investigations were in the domain of sensory substitution, whose advocates suggest that a visually impaired user can learn to interpret visual information encoded and presented to another sense (Collins, 1970). Wearable vibrotactile display systems have also been developed for specialist applications. For example, Sklar and Sarter (1999) describe a simple two tactor system intended to provide alerts for pilots. Van Erp et al. (2005) investigated the display of navigation cues using a belt based vibrotactile display system worn by soldiers steering speedboats or flying planes. Lindeman et al. (2005) describe a broadly similar belt display coupled with an application providing spatial awareness information designed to support military personal as they explore a hostile area. The idea that vibrotactile displays can encode spatial information has also been investigated for more everyday tasks. In one of the earliest uses of vibrotactile cues for general mobile computing, Tan and Pentland (1997) describe and evaluate a 3 by 3 back mounted array of tactors (a general term for vibrotactile display elements) and how it could be used to present directions to drivers. Evaluations of these systems have yielded promising results, indicating that users can successfully assimilate and understand the novel feedback.

As work has started to consider more general purpose mobile computing tasks, there has been focus on less cumbersome output devices. Oakley et al. (2006) evaluated a three by three array of wrist mounted tactors and concluded that the localization rate for wristwatch style, laterally arranged tactors is above 90%. By experimenting systematically with attributes such as frequency, amplitude and waveform Brown et al. (2006) and Enriquez et al. (2006) have both created frameworks for meaningfully conveying multi-dimensional information with short bursts of vibration from a single tactor. Luk et al. (2006) describe an advanced tactile device which can accurately perturb the skin of the finger, and discuss interaction scenarios revolving around embedding this in the side of a handheld computer. They conduct several preliminary lab based studies (on fundamental tasks such as the identification of tactons or directional cues) and conclude that their device has a promising future role in mobile computing scenarios.

One key factor that separates this work from that dealing with specialist interfaces is the lack of evaluation in realistic situations. For example, while Van Erp et al. (2005) tested recognition performance with users steering the speedboats their system is intended to support, there are few contextualized investigations of the complex multi-dimensional cues now appearing in the literature. Whilst the everyday tasks that typical users can be expected to perform are less demanding than controlling a speeding boat, the cues authors are suggesting be used are also much more complex. Systems for specialist applications have tended to rely on a large number of well spaced tactors distributed over a large portion of the body, each of which can emit a single binary cue. In contrast, and for the sake of convenience, the literature on general mobile computing is focusing on relatively small numbers of tactors in close proximity, each of which can render a range of different cues. Relatively little research has examined these kinds of cue in context, and it currently remains unclear whether users will be able to accurately perceive such stimuli when engaged in even mildly distracting tasks.

Some work has appeared examining this issue: Tang et al. (2005) and Chan et al. (2005) both use relatively large, sophisticated stimulators delivering cues to the fingertips (the most sensitive area of the body) and consider the effects of distracter tasks on recognition performance of tactile stimuli. Although they conclude that tactile recognition is unaffected by the presence of a distracter task, arguably their choice of body site and the sophisticated devices they use predispose them to this conclusion, and may not generalize to mobile or wearable scenarios using simple, lightweight and practical tactors delivering stimuli to less sensitive regions of the body.

Underlying this entire research question is attentional theory, something which has received little explicit consideration in the literature. In this discipline, human attention has long been viewed as a process of allocating limited resources (Kahneman, 1973). It suggests that we have insufficient perceptual capacity to simultaneously attend to our entire environment and instead focus on specific parts, at the cost of ignoring others. This process takes place according to a complex logic and that can be, for example, affected by the ability of particular aspects a perceived environment to mask, or obscure the perception, of others (e.g. Lloyd et al., 1999). In multi-sensory systems, this viewpoint is further complicated by multiple-resource theory, which proposes that (among other things) our limited attentional resources differ among our sensory modalities (Wickins, 1991). Consequently, it predicts that cross-modal presentation of stimuli should lead to increased levels of information absorption. While this proposition has been relatively well established in the audio-visual domains, research explicitly including the tactile modality is in its infancy (Sklar and Sarter, 1999). However, given the successes of the now commonplace high amplitude cell phone buzz to indicate an incoming call and of the evaluations of specialist tactile interfaces described above, it seems likely that this theoretical explanation is appropriate in many cases.

What this paper seeks to examine is whether this model holds for the more sophisticated cues, and less intrusive device configurations, now appearing in the literature. However, rather than focusing on an exacting examination of the nature of tactile attentional masking, it adopts a high level approach and seeks to determine if typical, commonplace distracter tasks disrupt the perception of complex tactile stimuli. This applied issue is of direct relevance for system developers or designers seeking to incorporate complex tactile cues into their interfaces, and resolving it may serve as a spur for more purely theoretical work.

3. Method

3.1. Experimental overview

The goal of this work is to investigate the perception of the kind of complex tactons now routinely appearing in the literature using a wearable display and while users are engaged in a range of everyday activities. Three studies were conducted, each completed by a different set of participants and each comparing a control condition in which subjects were idle against a different, commonplace, distracter task: Transcription, mouse-based Data-entry and Walking. Tactile stimuli varied on two dimensions: body site and roughness. Using this range of variables allowed the exploration of whether some aspects of vibration are more resilient to distraction than others, and also whether some tasks are more or less interfering than others. The results provide a window onto real world tacton recognition performance.

3.2. Experimental design and measures

Each of the three experiments had the same structure, composed of four discrete stages always delivered in the same order: training, practice, control and distracter. This structure was adopted to provide participants with maximum experience with the cues before experiencing the distracter task. The training stage lasted up to five minutes in length, and simply involved participants using a GUI (identical to the experimental UI described in Section 3.3) that enabled them to play each of the vibrotactile cues. This enabled participants to become familiar with the experimental stimuli and included an informal check that the magnitude of each stimulus was considerably above threshold. Participants were not able to adjust the magnitude of the cues.

The remaining three stages used a stimulus identification paradigm, similar to that used in much of the work on vibrotactile display (Brown et al., 2006; Oakley et al., 2006). Essentially, this involved a pause, the display of one of the cues, then the presentation of a UI with which users could specify which cue they had just experienced. The practice condition involved 27 trials (each cue, three times) while the control and distracter conditions both contained 54 trials (each cue, six times). The pause before each trial was three seconds long in the practice and control conditions and varied randomly between 10 and 25 s during the distracter condition. The practice stage was used to ensure users were familiar with the experiment; no data was gathered. In the control stage users were at rest, just performing the experimental task. In the distracter stage users performed the experimental task at the same time as a distracter task. In line with the majority of the literature on the recognition of vibrotactile cues (e.g. Cholewiak and Collins, 2003), the experimental measure was error rate. Task completion time was not measured as any increases observed in the distracter conditions would likely reflect only that the participants were busy and not necessarily that their perception was impaired: it is unsuitable as a measure as it would incorporate a fundamental confounding influence.

3.3. Materials

All three studies used VB232 tactors (http://www.tactaid.com), a relatively high quality tactor than can be driven by standard audio output, and a subset of the vibrotactile cues designed, described, evaluated and released by Brown et al. (2006). Nine 500 ms stimuli were used in total. They varied along two dimensions, body site and roughness, with three values on each. The three body sites were located in a band around the wrist 5 cm back from the base of the thumb. One was situated on the left side of the wrist, one on the right and one on the upper, dorsal surface squarely between them. This arrangement is shown in Fig. 1, and was chosen as a wristwatch style arrangement has been examined by other authors (Oakley et al., 2006) and seems

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003

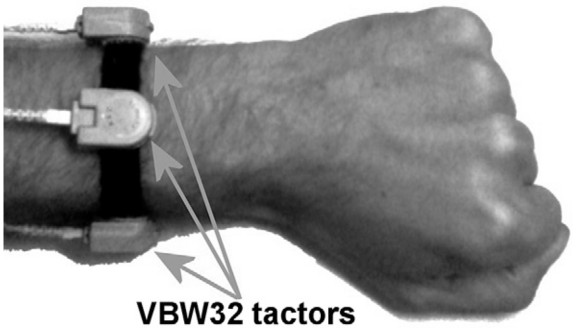

Fig. 1. Experimental device mounted on the wrist of a participant.

a likely candidate design for future development of this concept.

All stimuli were 250 Hz sine waves as this is the frequency at which the VB232 tactors resonant, and to which the human skin is most sensitive (Rogers, 1970; Verrillo, 1963). The effect of roughness was created by amplitude modulation. An unadulterated sine wave was labeled smooth, one modulated with a 30 Hz sine wave rough and one modulated with a 50 Hz sine wave in between these values. These waveforms are illustrated in Fig. 2, but those seeking a full description should refer to Brown et al. (2006). Stimulus amplitude was established through an informal rating process to be significantly above threshold.

Each study featured a different distracter task; details are provided in Section 3.8 below. However, the on-screen interface facilitating user responses remained broadly the same, consisting of a window with nine buttons arranged in a grid, one for each of the possible cues. The axes of the grid corresponded to the two stimulus dimensions. Left to right signified tactor placement, while top to bottom signified smooth to rough. The Transcription and Data-entry studies featured an identical interface, while the Walking study used a variation on this designed for mobile PDA use: large, high contrast buttons in a screen devoid of other UI elements. They are both shown in Fig. 3. Simple icons were used to represent the position and roughness of each cue. For example, the top left icon always corresponded to a smooth cue delivered to the left of the wrist while the bottom right icon indicated a rough cue delivered to the right of the wrist.

3.4. Participants

All participants completed only one study. The Transcription and Data-entry studies had 8 participants, while the Walking study featured 9. All participants bar 1 were right handed. In the Transcription study there were 3 female and 5 male participants with a mean age of 32. The mean participant age in the Data-entry study was 26 and the group composed of 3 female and 5 male participants. The Walking study involved 3 female and 6 male participants with a mean age of 28. Most participants were workers at either our institution or an associated one. The remainder were acquaintances of one of these subjects. They were not financially compensated.

3.5. Procedures

The Transcription and Data-entry studies were conducted in a quiet office with the user seated in front of a desktop PC and wearing enclosing headphones to mask any audio cues that might emanate from the vibrotactile devices. The tactile display was mounted on the wrist of their non-dominant hand and covered with a loose cloth to obscure any visual cues. An experimenter remained with the participants through the training and practice stages, but they were left alone to complete the experimental stages. Keyboard shortcuts for the experimental interface were disabled; participants were required to respond to the experimental stimuli using the mouse. This minimized the possibility that participants would make unintentional responses to the stimuli.

The Walking study had a broadly similar setup except participants completed all stages of the study upright: standing in the first three stages, and walking up and down a quiet (but by no means abandoned) corridor in the distracter stage. The experimental interface was displayed on a PDA held in the dominant hand and responses were entered through thumb taps. They were instructed to keep their non-dominant hand (with the tactors attached) idle in a resting posture at their side for the duration of the study.

3.6. Hardware and software

Two platforms were developed to conduct these experiments. The Transcription and Data-entry studies were

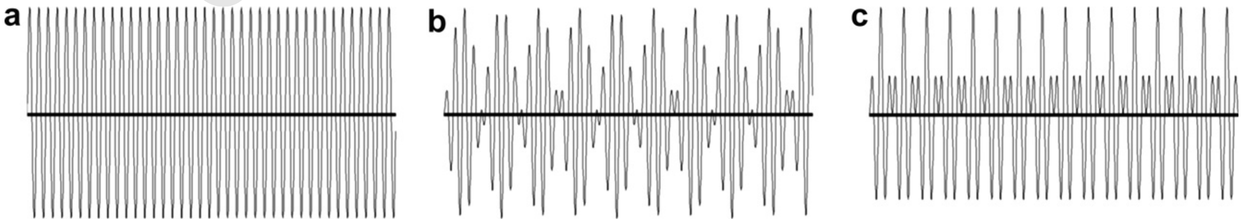

Fig. 2. Comparative illustrations of the three waveforms used to create vibrotactile stimuli: (a) 250 Hz sine wave, (b) 250 Hz sine wave with 30 Hz amplitude modulation, (c) 250 Hz sine wave with 50 Hz amplitude modulation. See Brown et al. (2006) for a full description.

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003

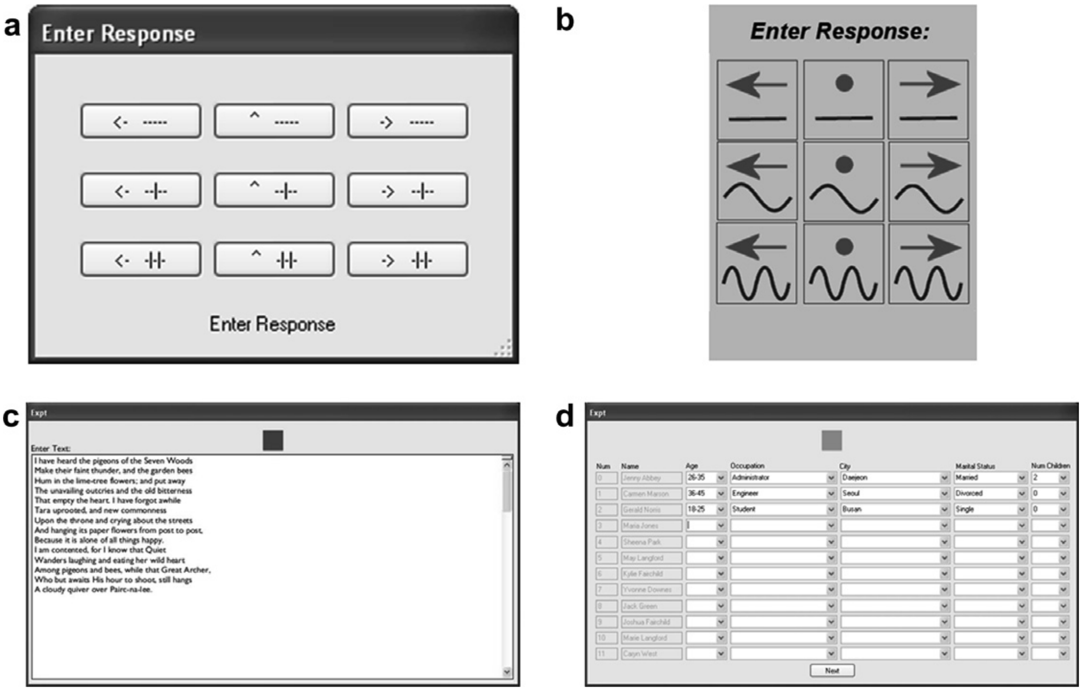

I have heard the pigeons of the Seven Woods
Make their faint thunder, and the garden bees
Hum in the lime-tree flowers; and put away
The unavailing outcries and the old bitterness
That empty the heart. I have forgot awhile
Tara uprooted, and new commonness
Upon the throne and crying about the streets
And hanging its paper flowers from post to post,
Because it is alone of all things happy.
I am contented, for I know that Quiet
Wanders laughing and eating her wild heart
Among pigeons and bees, while that Great Archer,
Who but awaits His hour to shoot, still hangs
A cloudy quiver over Pairc-na-lee.

Num	Name	Age	Occupation	City	Marital Status	Num Children
0	Jenny Abbey	26-35	Administrator	Daejeon	Married	2
1	Carmen Marson	36-45	Engineer	Seoul	Divorced	0
2	Gerald Norris	18-25	Student	Busan	Single	0
3	Maria Jones					
4	Sheena Park					
5	May Langford					
6	Kylie Fairchild					
7	Yvonne Downes					
8	Jack Green					
9	Joshua Fairchild					
10	Marie Langford					
11	Caryn West					

Fig. 3. Experimental interfaces: (a) response screen for Transcription and Data-entry studies, (b) response screen for Walking study, (c) distracter task from Transcription study, (d) distracter task from Data-entry study.

developed using Microsoft's C# and net tools and ran on a standard Microsoft Windows XP computer. The three tactors were controlled using the audio output of the PC's Dolby 7.1 compatible soundcard. To achieve individual control of the tactors DirectSound, a gaming API, was used to create a simple 3D sound scene. By positioning sound sources to the front (left and right simultaneously), to the rear and to either side of the user (again simultaneously) it is possible to output the three separate stereo channels needed to drive the tactors. To resolve variations in the volume of sounds delivered in each of these channels the outputs were routed through small headphone amplifiers, whose volumes were manually adjusted until they were subjectively and informally deemed equivalent.

The Walking study using a system based around three PDAs, two Dell Axim X51vs running PocketPC 2005 and one iPaq hx4700 running PocketPC2003. Applications were developed in C++. Each PDA controlled one tactor through its headphone jack. One PDA acted as a master and used Bluetooth links to control a small audio-capable slave application on the other two devices. The slave applications simply listened for instruction bytes informing them to play one of the nine experimental stimuli. The master PDA issued such instructions, was the site of the main experimental application, and was held in the user's dominant hand throughout the study. The slave PDAs were placed in a backpack for the duration of the experiment. Once again, headphone amplifiers and a subjective equalization procedure were used to ensure that the magnitude of the signals generated by each PDA were equivalent. The amplifiers were also stowed in the backpack. The final weight of the bag was just over 1 kg.

3.7. Hypothesis

The central hypothesis of this work is that recognition rates for tactile icons will decrease when users are engaged in distracter tasks; that distracter tasks will exert a masking effect and that lab based studies do not accurately represent real world performance. Beyond testing the truth and magnitude of this assertion, the three studies reported here also hope to examine it in additional detail. Considering a range of tasks encompassing differing levels of physical and mental activity may reveal if specific kinds of task mask, to a greater or lesser extent, tactile perception. Furthermore, by examining multi-dimensional tactons, it may be possible to ascertain if some parameters are more resilient to this masking than others.

3.8. Study descriptions

3.8.1. Study 1 – Transcription

The distraction task in this study involved transcribing a set of printed poems into a window on the computer screen. The poems were in a document holder situated adjacent to the computer monitor. For right handed users the document was placed on the right of the screen, for left handed users, the left of the screen. The UI also featured an adaptive speed monitor in the form of a colored square above the text entry window. The color of this square was based on a rolling average of the participant's typing

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003

speed, capped to a minimum of 40 key presses per minute. Essentially, it was green when they exceeded their rolling average and red when they failed to do so. It was intended to encourage participants to direct their full attention to the distracter task. Fig. 3 includes a screen shot of this interface.

This task involved both mental and physical aspects. The acts of reading, remembering and typing the text consumed mental resources. Moving the hands and arms to type and occasionally turn the pages of the printed text physically occupied the body, and in particular those parts of it directly engaged by the wearable tactile device.

3.8.2. Study 2 – Data-entry

As with the previous study, this experiment involved entering data from a printed sheet into the computer; the setup was largely similar. In this case, the data took the form of a table listing statistics about people: ID number, name, age, occupation, city of residence, marital status and number of children. An on-screen application mirrored these fields. ID number and name were automatically filled in and participants were required to enter the remainder of the data. Each data item had a fixed number of items (for instance there were six age ranges, and 12 possible cities of residence) and all data-entry took place using drop down list boxes. Participants were required to use the mouse in their dominant hand to do this. As with the previous study, an adaptive speed monitor was displayed. This was based on the rate at which participants altered list box selections and capped at a minimum of 12 per minute. Fig. 3 includes screen shot of this interface.

This task was designed to mimic the mental distraction of the Transcription task, but omit the physical distraction of moving the wrist on which the wearable tactile display was mounted. Participants were requested to keep their non-dominant arm still for the duration of the study.

3.8.3. Study 3 – Walking

The distraction task in this study was simply walking up and down a corridor. Participants were instructed to leave their non-dominant hand (with the wearable device attached) idle for the duration of the study. This task involved little to no mental distraction; participants were able to perform the task nearly autonomously, and could be observed focusing closely on the PDA. It involved a degree of physical distraction in that the bodies of the participants were in motion, but it is important to note that their arms relatively remained relatively still throughout.

3.9. Results

The data from all three experiments are shown in Fig. 4. The mean data for both control and distracter conditions in each experiment are presented, followed by the percentage of trials in which participants responded correctly on the individual stimulus dimensions of body site and roughness. These mean percentage correct data were analyzed using a single ANOVA along similar lines: two conditions (control and distracter, within subjects), by two stimuli components (body site and roughness, within subjects) by three experiments (Transcription, Data-entry and Walking, between subjects). The results revealed effects of condition ($F(1,12) = 9.93, p < 0.01$) and stimuli component ($F(1,12) = 210.6$, $p < 0.001$) but not experiment ($F(2,12) = 1.989$, $p = 0.143$). There were no significant interactions.

As the experimental design used in these studies always places the distracter condition after the control condition, a statistical analysis to determine the presence of any bias (positive practice or negative fatigue or habituation) that this might result in was also conducted. This was achieved by checking for correlations between the trial order and the mean correctness of the response generated by all subjects in both experimental conditions. These raw data are shown in Fig. 5, and two-tailed Pearson's product-moment tests showed a significant positive link (indicating a practice

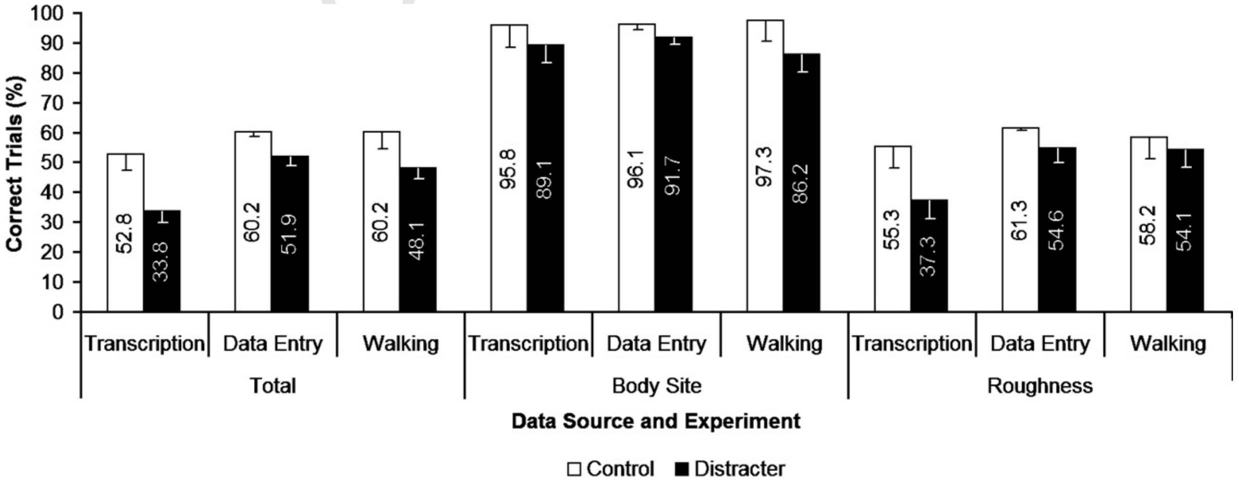

Fig. 4. Percentage correct trials recorded in each study (Transcription, Data-entry and Walking). Data are divided to show total recognition rate and also recognition rates for each stimulus dimension (body site and roughness). Error bars show standard error.

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003

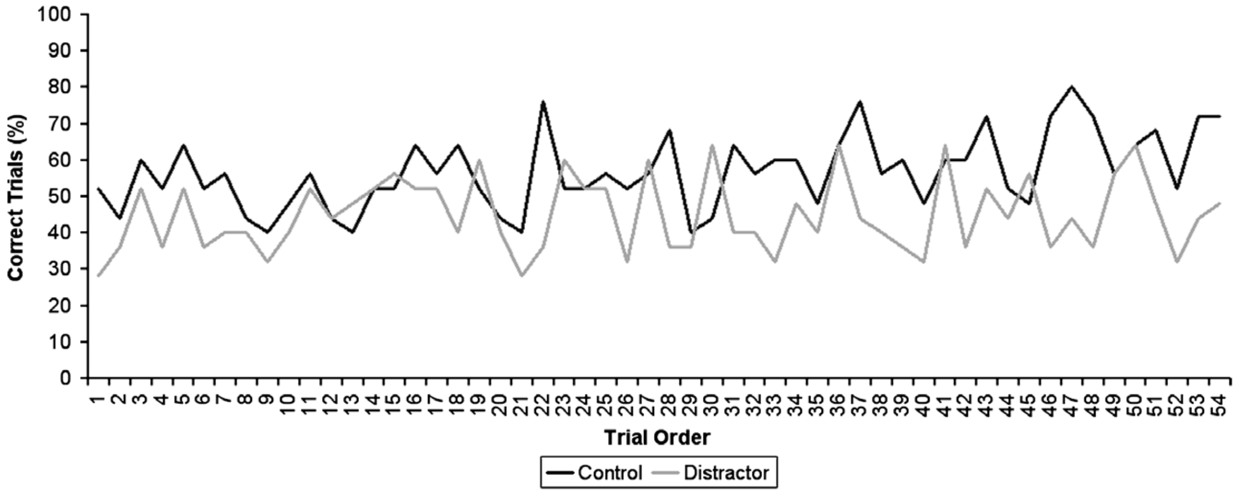

Fig. 5. Percentage correct trials shown by trial order for all subjects. Data from control and distracter conditions plotted separately.

effect) in the control condition ($r(24) = 0.481$, $p < 0.001$) but revealed no change in the distracter condition ($r(24) = 0.121$, $p = 0.385$).

3.10. Discussion

The main result of this work is clear: distraction, in the form of being engaged in other tasks, can mask the perception of vibrotactile cues. Despite the presence of a practice effect in the control condition, and the absence of a fatigue effect in the subsequent distracter condition, a significant reduction in performance of between 5% and 20% was observed across both stimulus dimensions and all three experiments. Had a balanced experimental design been adopted, it seems likely that this difference would have been greater in magnitude. These rates are high enough to have a substantial impact on the usability and usefulness of an interactive system, and taken together indicate that the results reported in lab based studies (e.g. Brown et al., 2006) are not likely to be representative of real world performance: all three distracter tasks have a negative impact on recognition rate. This suggests that the detrimental masking effects observed in this study are likely to appear in any real world deployment of complex vibrotactile cues.

This conclusion underlines the importance of conducting studies on mobile user interfaces in context (Pirhonen et al., 2002). Furthermore, given the relatively simple nature of the tasks studied here, and the quiet, stable environment in which they took place, it is entirely possible that a true real world study (conducted, for example on users reading whilst riding on public transport) will reveal much stronger masking effects. In such an environment, it may be that vibrotactile cues are rendered relatively inexpressive, simply overwhelmed by environmental stimuli. Alternatively, if users are continually alert for detailed vibrotactile cues, environmental stimuli may become increasingly distracting. For example, a user keyed into a wide range of tactile messages may detect many false positives: tactile mirages arising from natural vibrations caused by garments rubbing against one another or the erratic buzzing of a vehicle in which they are traveling. Investigating such situations to establish the veracity of these suggestions is an obvious next step for this work.

It is worth discussing two potential confounds may influence these conclusions. The first of these is that the user response paradigm (the selection of an on-screen button to indicate a particular cue) is the same for each of the studies, and therefore the differences observed among the conditions may be due to some peculiarity of this process. For example, it is possible that the distracter tasks interfered not with vibrotactile perception, but instead with ability to map these to the appropriate button selection action. However, although this is a possible alternative explanation, it is also true that in a real application scenario, a response based on pressing a button (or similar UI element) is a highly likely interaction model. Therefore, the practical difference between these two accounts may well be minimal. Nevertheless, clarifying this point by conducting an alternate version of these studies based on a radically different method for capturing user responses, perhaps by recording spoken utterances, would be a worthwhile activity.

The second confound relates to the perception of the tactile cues. Although an informal process to ensure they were significantly above threshold took place, and all stimulus levels remained the same for all subjects in each study, additional effort could have been expended to specify them. The experimental setup did not allow participants to indicate a failure to perceive a cue, potentially mixing such responses with those in which a cue was detected. However, the approach adopted here reflects several key observations. Firstly, there are few established procedures for establishing the perceptual magnitude of tactile cues, and few uniform standards between different display devices. Consequently, subjective determinations of magnitude are commonly adopted (e.g. Brown et al., 2006). It has also been observed that there are few differences in the ability to localize a vibrotactile stimulus by varying the actuators used (Jones et al., 2005) or the amplitude or frequency (Cholewiak and Collins, 2003) of the cues (it is unclear if this robustness holds for the detection of other stimulus attributes such as roughness). The relatively high localization performance in all conditions suggests participants had little trouble detecting the presence of a cue and goes some way towards validating this approach. It is further reinforced by comparing the data from the studies conducted here to those in the literature.

Q2

The results also show body site was more easily identified than roughness, an effect which has previously been noted. In particular, the tactons used in this work were proposed and studied by Brown et al. (2006) who reported mean recognition rates of 97% for body site and 54% for roughness (when three roughness stimuli were used). The overall mean 96% and 58% recognition rates achieved in the three control conditions here are broadly consistent with these figures. However, despite the similarities between the results, the procedures used in this work differed from those adopted by Brown in several important ways. Brown's body sites were spaced along the length of the ventral forearm (near wrist, near elbow and in between) and all stimuli were presented for 2000 ms. The closer body sites (around the wrist) and shorter stimulus presentation times (500 ms) used in this work do not appear to have influenced recognition rates, suggesting there may be little advantage to more widely spaced tactors, or longer stimulus events. In the absence of distraction tasks, short bursts of vibration emitted from a wrist watch style display appear to be as easy to perceive as much longer stimuli coming from points distributed over the entire forearm. This observation does not represent an entirely concrete conclusion, instead serving to demonstrate that the practical question of how to optimally arrange and stimulate a wearable array of tactors remains currently unanswered (Cholewiak and Collins, 2003). However, it remains compelling, suggesting that further attention to a wrist-watch style display is warranted, and that short vibrations may be as effective as longer ones. This is encouraging evidence supporting the future deployment of vibrotactile devices and cues, as it is likely that the kind of increments in user convenience such a display represents will be required for the widespread adoption of such systems. As Pierce et al. (1999) point out in relation to virtual reality display peripherals, users can be reluctant to don elaborate or cumbersome equipment simply to interact with computer systems.

Previous research on wearable vibrotactile displays has highlighted the role of anatomical reference points in improving localization. This refers to the fact that stimuli delivered to easily identifiable body sites – such as the wrist or elbow on a display positioned up the length of the arm (Cholewiak and Collins, 2003) or center of the spine or stomach on a display around the torso (Cholewiak et al., 2004) – are more accurately recognized than those at less easily specified sites. In a direct comparison between arm and torso based tactile displays, Jones et al. (2006) conclude that the torso is a more suitable body site, arguably because of the presence of a more readily identifiable set of bodily landmarks. However, the studies in this paper support Oakley et al. (2006) in their suggestion that anatomical landmarks in the form of the cardinal points around the wrist (top, bottom and sides) are easily recognized, and offer a level of performance similar to that which can be observed in torso mounted displays with a much higher inter-tactor spacing. The 96% recognition rate for localization in the control conditions of the studies reported here is compelling high. However, it remains true that the issue of anatomical reference points in vibrotactile perception is one which is not fully explained and deserves further, closer attention.

One of the objectives of this work was to establish if recognition performance varied among the different distracter tasks, or if the stimulus parameters were affected differently. Such information would cast light on the kinds of cues, and the kinds of tasks that might be best suited for vibrotactile display. From the point of view of attentional theory, this objective can be expressed as seeking to determine, at a high level, the relative masking abilities of distracter tasks, and whether certain stimulus parameters offer more or less resistance to this. However, the main analysis did not uncover any such difference, suggesting distraction exerted its masking effect uniformly, and the nature of the tasks and cues used had no influence on performance.

While this may be the case, an examination of the raw data leads to one result that stands out in this respect: roughness recognition in the Transcription study descends from 55% to 37% (little over chance) and causes a corresponding drop in total recognition rate. Given the magnitude of this drop a further brief analysis of this data was conducted. *T*-tests comparing the data from the distracter condition of the Transcription study to that of the Data-entry and Walking studies both showed significant differences (respectively, $t(14) = 2.47$, $p < 0.05$ and $t(15) = 2.39$, $p < 0.05$), and although these results are not sufficient to overturn the fact that no interaction between the factors of experiment and stimuli component was recorded in the main analysis, they do suggest that different distracter tasks may exert different effects on performance. Further as the body site data does not appear to show this effect, it may also be that different stimulus parameters are affected differently. A more powerful study designed to tease apart these factors would provide useful insights. For example, as this effect is observed in the Transcription task, the only one which includes motions of the forearm, one compelling possible explanation is that it is these movements disrupted the perception of the roughness of the cues. The presence of such an interfering, masking, link between local motor activity and vibrotactile perception would have wide reaching implications for stimulus design and be of considerable importance to anyone seeking to deploy a vibrotactile

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003

interface. It may also serve to explain the general prominence of torso mounted displays in the literature (e.g. Van Erp et al., 2005; Lindeman et al., 2005; Jones et al., 2006); the torso is not simply larger and in possession of more anatomical landmarks than other body areas, but also relatively inflexible and therefore not subject to localized movements which can mask displayed cues.

4. Conclusions and future work

The studies described in the paper suggest that distraction, through a process of attentional masking, negatively influences tacton recognition performance. Given that much recent work on tactons has focused on mobile or wearable scenarios where distraction is inevitable, this paper concludes that it is a factor that researchers and system designers can not afford to ignore. To do so would compromise both the usability and effectiveness of the interfaces they create. However, the work presented in this paper only represents an initial effort to explore this issue, and many questions remain unresolved.

Future work on this topic includes additional studies to determine whether different kinds of task exert different effects on the recognition of tactons; to explore the precise properties of the masking behavior observed here. In particular, tasks which involve movement of the body part hosting the tactile display seem likely to more detrimentally affect performance. Furthermore, although this paper is concerned with the effects of distraction on performance, future work should consider how to design tactons to be resilient to the effects of distraction. Numerous strategies suggest themselves, the simplest being repeated stimuli presentation.

However, this seems inelegant (not to mention potentially annoying) and does not address the fundamental problem. If a cue was difficult or impossible to perceive on its initial presentation, it may well be the case that this remains true in subsequent ones. Environmental noise, for example, may well overwhelm a vibrotactile message regardless of how frequently it is presented. An alternative and potentially more promising approach involves developing an interaction model which is based on cues directly related to the task at hand, rather than unrelated, as those studied here. Williamson et al. (2007) provide an example of how this might be achieved. In their system, vibrations are delivered in response to (and in part based on) rich motion input. By presenting cues only when a user is attending to them, and also varying them based on the parameters of user input, it may be possible to convey information more reliably. A similar concept was also explored by Sekiguchi et al. (2005). Fundamentally, this idea is grounded in the work of Lederman and Klatzky (1993) which suggests that haptic perception is an active process of exploration and that performance is greatly reduced in situations where users are merely passively exposed to cues. Indeed, establishing whether there is a distinction in tacton recognition performance analogous to that between the high levels of acuity observed when actively exploring a physical object versus the relatively poor performance found when passively experiencing contact would be of considerable interest, and may open the door to much more effective tacton design strategies.

Acknowledgements

This work was supported by the IT R&D program of the Korean Ministry of Information and Communications (MIC) and Institute for Information Technology Advancement (IITA) (2005-S-065-03, Development of Wearable Personal Station). Yeongmi Kim and Jeha Ryu donated their time and facilities so that one of the experiments could be conducted: we are grateful for this. The comments of several of the anonymous reviewers have also made a significant contribution, both improving the quality of the paper and broadening its scope. Finally, we extend our thanks to all our experimental participants.

References

Brown, L.M., Brewster, S.A., Purchase, H.C., 2006. Multidimensional tactons for non-visual information presentation in mobile devices. In: Proceedings of Mobile HCI 2006. ACM Press, pp. 231–238.

Brown, L.M., Kaaresoja, T., 2006. Feel who's talking: using tactons for mobile phone alerts. In: CHI'06 Extended Abstracts on Human Factors in Computing Systems. ACM Press, pp. 604–609.

Chan, A., MacLean, K.E., McGrenere, J., 2005. Learning and identifying haptic icons under workload. In: Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC 2005). IEEE Computer Society, pp. 432–439.

Chang, A., O'Sullivan, C., 2005. Audio-haptic feedback in mobile phones. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'05). ACM Press, pp. 1264–1267.

Cholewiak, R.W., Collins, A.A., 2003. Vibrotactile localization on the arm: effects of place, space and age. Perception and Psychophysics 65, 1058–1077.

Cholewiak, R.W., Brill, C.J., Schwab, A., 2004. Vibrotactile localization on the abdomen: effects of place and space. Perception and Psychophysics 66, 970–987.

Collins, C.C., 1970. Tactile television – mechanical and electrical image projection. IEEE Transactions on Man Machine Systems 11, 65–71.

Enriquez, M., MacLean, K.E., Chita, C., 2006. Haptic phonemes: basic building blocks of haptic communication. In: Proceedings of the 8th International Conference on Multimodal Interfaces (ICMI'06). ACM Press, pp. 302–309.

Jones, L., Lockyer, B., Piateski, E., 2006. Tactile display and vibrotactile pattern recognition on the torso. Advanced Robotics 20, 1359–1374.

Lederman, S.J., Klatzky, R.L., 1993. Extracting object properties through haptic exploration. Acta Psychologica 84, 29–40.

Lindeman, R.W., Sibert, J.L., Mendez-Mendez, E., Patil, S., Phifer, D., 2005. Effectiveness of directional vibrotactile cuing on a building-clearing task. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'05). ACM Press, pp. 271–280.

Lloyd, D.M., Bolanowski Jr., S.J., Howard, L., McGlone, F., 1999. Mechanisms of attention in touch. Somatosensory and Motor Research 16, 3–10.

Luk, J., Pasquero, J., Little, S., MacLean, K., Hayward, V., Levesque, V., 2006. Haptics as a solution for mobile interaction challenges: initial design using a handheld tactile display prototype. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'06). ACM Press, pp. 171–180.

Kahneman, D., 1973. Attention and Effort. Prentice-Hall, Englewood Cliffs, NJ.

Oakley, I., Kim, Y., Ryu, J., 2006. Determining the feasibility of forearm mounted vibrotactile displays. In: Proceedings of the 14th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'06). IEEE Computer Society, pp. 27–34.

Pierce, J.S., Pausch, R., Sturgill, C.B., Christiansen, K.D., 1999. Designing a successful HMD-based experience. Presence 8, 469–473.

Pirhonen, A., Brewster, S., Holguin, C., 2002. Gestural and audio metaphors as a means of control for mobile devices. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'02). ACM Press, pp. 291–298.

Rogers, C.H., 1970. Choice of stimulator frequency for tactile arrays. IEEE Transactions on Man Machine Systems 11, 5–11.

Sekiguchi, Y., Hirota, K., Hirose, M., 2005. The design and implementation of ubiquitous haptic device. In: Proceedings of WorldHaptics 2005 (WHC'05). IEEE Computer Society, pp. 527–528.

Sklar, A.E., Sarter, N.B., 1999. Good vibrations: tactile feedback in support of attention allocation and human-automation coordination in event-driven domains. Human Factors 41, 543–552.

Tan, H.Z., Pentland, A., 1997. Tactual displays for wearable computing. Personal Technologies 1, 225–230.

Tang, A., McLachlan, P., Lowe, K., Saka, C.R., MacLean, K., 2005. Perceiving ordinal data haptically under workload. In: Proceedings of the 7th International Conference on Multimodal Interfaces (ICMI'05). ACM Press, pp. 317–324.

Van Erp, J.B., Van Veen, H.A., Jansen, C., Dobbins, T., 2005. Waypoint navigation with a vibrotactile waist belt. ACM Transactions on Applied Perception 2, 106–117.

Verrillo, R.T., 1963. Effect of contactor area on vibrotactile threshold. Journal of the Acoustical Society of America 35, 1962–1966.

Wickins, C.D., 1991. Processing resources and attention. In: Damos, D. (Ed.), Multiple-task Performance. Taylor and Francis, London, pp. 3–34.

Williamson, J., Murray-Smith, R., Hughes, S., 2007. Shoogle: multimodal excitatory interaction on mobile devices. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'07). ACM Press, pp. 121–124.

Please cite this article in press as: Oakley, I., Park, J., Did you feel something? Distracter ..., Interact. Comput. (2007), doi:10.1016/j.intcom.2007.11.003