

Evaluating Haptics for Information Discovery While Walking

Simon Robinson, Parisa Eslambolchilar, Matt Jones
Future Interaction Technology Lab
Computer Science Department
Swansea University
Swansea, SA2 8PP, UK
{s.n.w.robinson, p.eslambolchilar, matt.jones}@swansea.ac.uk

ABSTRACT

In this article we describe and evaluate a novel, low interaction cost approach to supporting the spontaneous discovery of geo-tagged information while on the move. Our mobile haptic prototype helps users to explore their environment by providing directional vibrotactile feedback based on the presence of location data. We conducted a study to investigate whether users can find these targets while walking, comparing their performance when using only haptic feedback to that when using an equivalent visual system. The results are encouraging, and here we present our findings, discussing their significance and issues relevant to the design of future systems that combine haptics with location awareness.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: Haptic I/O, Interaction Styles

General Terms

Design, Experimentation, Human Factors

Keywords

Mobile computing, location-aware, haptics

1. INTRODUCTION

Geo-tagged views of the places we visit have become increasingly common in recent years with the proliferation of online mapping and low-cost location-based systems. *Google Maps* shows the user aerial views of a location, whilst overlaid *Panoramio* photos and *Wikipedia* articles provide a sense of context and history about landmarks, interesting information snippets or even just quirky views of nearby objects. These services work well on the large screens of conventional computers, but lack elements of usability in mobile scenarios, where it is tempting to simply shrink the display and provide the same interaction on a smaller screen. While they

are, of course, functional, systems like these require the user to transfer attention from their physical surroundings to an alternative, digital representation of the location, and concentrate on interpreting the links between the two.

Our work is motivated by a desire to remove this need to grapple with the digital world; instead, we aim to provide immersive, 'heads up' ways for users to discover content associated with the places they visit. In this article we present and explore one such approach to mobile information discovery, using our prototype haptic system. This novel interaction prototype allows the user to interact with digital information in their nearby environment without the need to concentrate on a screen to discover geo-located data. Instead, interesting places are discovered by pointing and scanning, and 'felt' using haptic feedback that is directly related to the real-world size and position of the target.

We begin in Section 2 by reviewing previous research into pointing-based location interaction, mobile device feedback techniques and device usage while moving, situating our approach amongst these developments. Section 3 describes the haptic prototype in more detail, highlighting important aspects of its implementation and its novel contributions. In Sections 4-6 we describe a study conducted to evaluate our prototype, assessing its performance against an equivalent visual-based system, presenting our method, results and discussion of their significance. Finally, in Section 7 we conclude the article and present several pointers to future work.

2. BACKGROUND

Our prototype employs physical pointing to enable location interaction, combined with haptic feedback to help users discover information while on the move. In this section we first review previous mobile pointing-based systems, continuing to examine non-visual feedback for location-based systems. Finally, we review previous literature discussing user attention and interaction while moving, looking at the potential pitfalls of using a mobile device when on the move.

2.1 Pointing-based interaction

Researchers have only relatively recently begun to investigate the potential offered by mobile pointing-based interaction, creating many imaginative devices as the sensor technologies required to enable this interaction have reduced in size and cost. In an early contribution to this area, Egenhofer [2] discussed the possibilities for several *Spatial Information Appliances* before the necessary technologies were integrated

© The Author 2009.
Published by the British Computer Society

into mobile devices. These included a *Smart Compass*, providing turn-based GPS guidance toward a location, and *Geo-Wands*, that help users to identify physical geographic targets by pointing toward them.

Fröhlich *et al.* [3] assessed the viability of ‘point-to-select’ against alternative methods of interaction via a ‘Wizard-of-Oz’ style study. The results showed that the relation of location services to physical objects was helpful and important, and suggested that pointing gestures were ‘highly attractive and efficient’ forms of location selection. In a related project, building upon this work, Simon *et al.* [15] investigated the spatially aware mobile phone, a conceptual device to connect the physical and digital worlds. Using a three-dimensional model of a location in conjunction with knowledge of the user’s position, their framework was able to create a line-of-sight visualisation from the user’s location.

Approaches such as these demonstrate useful, usable methods for enabling pointing-based interaction, but they rely on a visual display as the basis of the interaction process: pointing achieves location selection, but discovery of the available information falls back to visual menu navigation. Our design uses the pointing phase of this process for discovery of both the location and quantity of the information available, helping the user to interact more closely with the location they are exploring, and only resorting to a visual display to show the resulting location information.

2.2 Non-visual location-based feedback

Moving away from visual feedback, Strachan *et al.* [17] used location and heading data in conjunction with real-time trajectory prediction to guide a user along a path to a desired target location. By pointing and tilting a device around their environment, the user can browse the route features around them, with both audio and haptic feedback directing them toward their destination. When the user is heading toward the target the audio signal is clear and there is no haptic feedback, but if they move off track the audio is distorted and vibrotactile feedback increases. In a related paper, Strachan and Murray-Smith [16] studied mobile interaction with virtual targets, with both vibrotactile and audio feedback. Their research addressed specifically the problems that can arise due to uncertainty in the user’s location and heading data, and offered a probabilistic approach to this problem. An experiment using their system showed that targets could be selected effectively, even when fairly tightly spaced.

Holland *et al.* created AudioGPS [5], which used audio to provide representations of the direction and distance of waypoints. Their informal study found the system adequate for navigation tasks to within eight degrees of a target, but there were some issues with latency under field conditions. In a similar system, Jones *et al.* [7] used ambient spatial audio to help guide users around a virtual environment, providing cues as to the direction and distance of a specified target by adjusting the fade and balance of an audio track.

Shin and Lim [14] used a vibrotactile jacket in conjunction with an ultrasound sensor array to provide obstacle detection and feedback for visually impaired users, and suggested that the use of haptic feedback for this task could help users to navigate around obstacles reasonably accurately and without losing track of their path. Similarly, Van Erp *et al.* [18] used eight tactors placed around the user’s waist to help indicate to them the direction of waypoints. Each tactor vibrated to indicate movement in a certain direction, and dis-

tance was also encoded into the vibrational feedback. Luk *et al.* [8] describe a prototype for mobile haptic interaction where piezoelectric actuators are used to provide several tactile sensations ranging from simple buzzes to complex moving patterns. Particularly interesting is the application of this technique to a navigation task, where the device ‘strokes’ the user’s thumb to prompt them to move forward.

These systems demonstrate the use of audio and haptic feedback for navigation tasks, often guiding the user to a specific pre-set location. Our approach, however, uses haptic feedback for the discovery of all available information targets in the area around the user, providing them with cues as to the location of possible points of interest rather than pointers toward their journey destination.

2.3 Using while moving

Several researchers have shown that using a device while moving can have a detrimental effect on user performance and mobility. Mustonen *et al.* [9] studied users reading text on a device screen while walking, finding that participants’ performance decreased while moving. Similarly, Barnard *et al.* [1] compared user performance when sitting and while walking, finding that moving significantly increased the time taken for word search and reading comprehension tasks. Pascoe *et al.* [11] looked at usage of a mobile device for field-workers, and suggested that the introduction of minimum attention user interfaces and the addition of context awareness could help to improve the effectiveness of mobile interaction in this environment. Oulasvirta *et al.* [10] investigated attention to the screen of a mobile device in several different environments, ranging from a laboratory to a busy street, finding large differences in user mobility between these situations.

As demonstrated in the papers here, managing attention to a device while moving can be a complex undertaking. In the next section we describe our prototype system, which attempts to provide an alternate interface to the discovery of geo-tagged data, allowing a more fluid, simplified interaction flow.

3. HAPTIC SYSTEM DESIGN

The haptic system allows the user to discover geo-tagged information in the environment around them by making pointing and sweeping movements with a mobile device, as illustrated in Fig. 1. In the following sections we describe in more detail its usage, implementation and the useful, novel interactive features it provides.

3.1 Haptic browsing

We designed a framework to allow haptic discovery of geo-located data, helping users to discover the location of interesting information in the area around them. When using our prototype, as the user moves around their environment their position (latitude, longitude) is used to refresh a selection of available points of interest near their location. The system is designed for one handed interaction with minimum attention required, and allows the user to search for geo-tagged information – for example: text, images, videos, audio content – without the need to look at a screen until they have discovered a target that interests them. Holding the device in their hand, the user can point and scan around the points of interest nearby, feeling gentle vibrotactile feedback when they move their focus around each target.



Figure 1: The haptic system in use: the SHAKE device (shown inset) vibrates when the user points toward a location with geo-tagged content available. Pressing the SHAKE's navigation button (shown at A) selects the target.

The feedback provided to the user helps them to discover both the direction of and quantity of information available about a point (see Fig. 2). Quantity is indicated by the spread of the vibration: popular targets with lots of data appear larger and take up more of the user's scanning range. To help the user determine the direction of the actual target, the vibrational feedback increases in intensity toward its centre.

3.2 Implementation

Our prototype system consists of a Sony VAIO Ultra-Mobile PC (UMPC), connected over Bluetooth to a Sensing Hardware for Kinesthetic Expression (SHAKE) SK6 sensor pack. The SHAKE (described in [19]) is a small integrated sensor device that contains a variety of compact sensing hardware. For our application we use its tri-axis accelerometers, gyroscopes and magnetometers for orientation, its internal motor to provide vibrotactile feedback, and the navigation button on its side to allow users to select targets. The SHAKE device is attached firmly to the back of the UMPC so that any movements the user makes whilst holding it are directly recorded. All data received from the SHAKE, in addition to timings, button presses made by the user and logs of when vibrotactile feedback pulses were sent and their intensities are recorded by the UMPC at all times.

Providing vibrotactile feedback on a device that is also used to gather sensor data introduces problems with sensor interference. In order to address this, we first low-pass filter the accelerometer, magnetometer and gyroscope data, then combine these to obtain a stabilised orientation matrix which provides a heading value that fluctuates less than using the magnetometer data alone would allow. This technique does, of course, introduce a small lag, but we feel this is acceptable to counter the more significant usability problems that would arise if we left the data unfiltered.

When the filtered compass heading intersects a nearby point of interest, the device generates vibrational feedback based on the tangential distance from the centre of the target, and increases in intensity as the distance decreases. When providing feedback, closer targets occlude those further away –

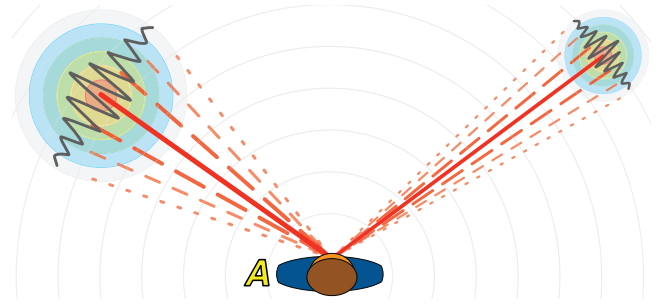


Figure 2: Haptic feedback: the spread of the haptic area helps the user, at point A, to determine the size of the target, and the increase in vibration frequency toward the middle of the target guides them to its centre point.

we assume the user has an interest in the points of information in their immediate vicinity.

3.3 Proposed benefits of the system

The novel interaction method we present in our haptic prototype provides a simple way for users to seek out digital resources whilst moving through their environment, without having to pause to look at a device's screen during the discovery process. We envisage people using the device whilst stationary, but believe it offers users the most benefits while walking and moving around their surroundings: the user can view, experience and interact with the physical features of their environment while simultaneously exploring the digital accompaniments to the world around them. This type of interaction with the system is the area we are most interested in, and in the next section we describe a user study undertaken to investigate usage in a similar scenario, evaluating the prototype's performance against an alternative, visual-based system.

4. EXPERIMENT

We performed a lab-based user study to study the system's usability when the user is walking. In this section we first describe an equivalent visual-based device that was constructed to allow comparisons between our prototype and a viable alternative system, then detail the method used to evaluate and compare these systems.

4.1 Visual comparison system

To enable a fair comparison between the haptic system and a viable alternative we constructed a second prototype that is a visual analog of the haptic system, to the extent that it was possible. Where the haptic system provides vibrotactile feedback, this device shows icons as a visual representation of the information targets that are available for discovery in the area around the user (see Fig. 3). Users hold the same hardware as in the haptic system, but on the screen they see a radar-like display that rotates as they turn, updating to ensure that the display shows the targets that are currently in the user's field of view. Feedback for target interaction is achieved by changing the colour of the target icons when the user points directly toward them, and the centre of each target is indicated by a cross (see points B and C, Fig. 3). In addition, the quantity of information available is indicated



Figure 3: The visual system in use for the study, showing inset the display the user sees. The user is represented on the screen by the icon at A and a highlighted target is shown at point B. The icon used for unselected targets is shown at C.

by the size of the icon: larger icons indicate more geo-tagged information is present.

4.2 Study design

Our main interest was in evaluating whether using haptic feedback could help improve the user experience whilst on the move, comparing their performance against an equivalent visual system. We measured user performance over several factors, specifically: whether users could find targets, the affect on their normal walking speed, the time taken to select targets, any false positives generated, and their perceptions of the systems.

4.2.1 Participants

Twenty participants aged between 18 and 55 were recruited for a 15-minute lab study. Thirteen participants were undergraduate students, two were postgraduates and five were members of university staff in areas unrelated to HCI. Ten participants were male and ten female. Six participants had previous experience of sensor-based interaction from using the Nintendo Wii games console.

4.2.2 Conditions

Prior to the study, participants were allocated between two conditions: the haptic system and the visual comparison system. When allocating participants, those with prior sensor-based interaction experience were randomly divided so that both systems had an equal number, then the remainder of the




Figure 4: Study layout: grey circles indicate positions of cones; arcs T_1 and T_2 are target zones. User walks from point A to end and repeats from B.

participants were allocated randomly to give ten participants per system.

4.2.3 Measures

To quantify system performance we used the Percentage of Preferred Walking Speed (PPWS) measure [12] as an indication of the system's affect on a user's normal behaviour, but used a different method for recording this to mitigate concerns about its inaccuracy. PPWS is measured by recording a participant's average preferred walking speed before the experiment, then expressing walking speeds recorded during the study as a percentage of this value. During an initial pilot study between our systems, however, we found the commonly used pedometer-based measurement of PPWS to be unreliable for shorter distances. When using the system in our pilot study, participants took more steps than they had when their pace was measured beforehand. When using this data to calculate percentage of normal walking speed this would show a speed improvement, when in actual fact this is not the case.

To solve this accuracy problem, we opted to place a tape measure at the side of the walking circuit, and referred to this to measure the distance walked by participants, instead of using a pedometer. Measurements were taken from the front of the participant's foot to the nearest 10cm. In addition, to correct for the curved walking course, at the end of the study all recorded distances were scaled up in line with a previously-calculated average maximum circuit length of 10.2m, obtained prior to the study by measuring by three participants' average distance walked around the circuit.

In addition to the PPWS measure we were interested in the time taken to select targets, which could indicate whether certain target positions took longer to select than others. We also recorded the number of false positives, defined as when the participant pressed the selection button when not actually pointing toward a target, allowing us to highlight interaction issues or cases of feedback confusion.

Finally, we used a questionnaire based on the NASA TLX instrument [4], rating six aspects of participants' perceived performance – mental demand, physical demand, time pressure, performance, effort and frustration – from 1 (low demand) to 7 (high demand). We also recorded any comments given by participants, and noted interesting participant behaviours.

4.2.4 Tasks

We based our experiment on procedures developed by [12] and refined by [13], aiming to simulate a realistic environ-

ment, but maintain control over the variance that can occur in studies such as these. During the study, participants walked around a circuit in a university corridor (see Fig. 4), negotiating obstacles and using the system at the same time to select targets. Participants completed three primary tasks during the study:

- Walk three 9m lengths of the corridor to measure their average preferred walking speed.
- Find and select each of 30 targets at opposite ends of the circuit:
 - Stand at point *A*, touch the device screen and start walking toward point *B*. While walking, attempt to select a target in zone T_1 .
 - Repeat from point *B* to *A*, selecting a target in zone T_2 .
 - Continue to repeat these steps until all 30 targets have been selected.
- Complete a TLX questionnaire to rate their perceptions of the system.

The targets used were identical in size, and appeared 12° wide on each system. Targets were equally spread over zones T_1 and T_2 at the same positions in each zone, but participants could only see one target at once. The order of the targets was random, but the order was the same for each participant and system.

4.2.5 Procedure

At the start of each study session, participants were met individually and introduced to the system they would be using and its purpose via a verbal walkthrough of a typical usage scenario. Participants were then given a demonstration of the system to select example targets in the lab, and asked to use the system themselves to select three sample points as a form of device training.

Each participant was then taken to the corridor used for the experiment and, before starting the study tasks, was asked to walk three 9m lengths of the corridor while the researcher timed them, allowing us to calculate their average preferred walking speed. Participants then used the system while walking along the corridor circuit, attempting to select targets at the same time.

Participants started at point *A* and were asked to simultaneously start walking and touch the device screen to begin target selection. Once walking, participants followed the line of the course, weaving around the cones while trying to select the current target. When the participant successfully selected the target, they stopped walking immediately so the researcher could note down the distance they had walked. Participants were told they should not stop walking until they had either selected the target or reached the far side of the course. If the participant reached the end of the course without selecting the target, they stopped and selected the target from the end of the circuit, but the researcher noted this event. After the participant successfully selected a target, they walked to the opposite end of the corridor (point *B*) and repeated the procedure, continuing to repeat this until all targets had been selected, walking 30 lengths of the corridor in total.

Measurement (units)	Haptic	Visual
Distance walked per target (m)	3.5 (2.5)	4.3 (2.8)
Time to select each target (secs)	6.7 (4.8)	7.2 (4.5)
No. false positives per target	0.5 (1.2)	1.2 (1.7)
Original walking speed (m/sec)	1.5 (0.1)	1.4 (0.2)
PPWS (% of original speed)	37.7 (19.6)	43.6 (18.1)

Table 1: Means and standard deviations (in brackets) of each of the measures recorded.

After selecting all targets, participants completed a TLX questionnaire and offered any verbal feedback resulting from their usage of the system. Finally, at the end of the study all participants were rewarded with a bookstore gift voucher as a token of our appreciation.

5. FINDINGS

All participants found and selected each of the 30 targets placed around the study area. Each participant also completed the TLX questionnaire and, in the majority of cases, offered verbal feedback about the system they had used. In the following sections we present our results, comparing data from the haptic system against the visual alternative.

5.1 Outlying results

During the study, two participants (one haptic, one visual, both female students in the 18-25 age group, and with no prior accelerometer experience) were confused about undertaking the tasks they were set, and struggled to make sense of and complete the experiment. When initially analysing recorded data, the PPWS measurements from these participants were found to be significantly set apart from the rest (over two standard deviations from the mean). Consequently, we discarded all the data from these participants before continuing with analysis. In addition, 20 cases (3.7% of the remaining 540 selected targets) where participants reached the end of the circuit without successfully selecting the target were removed from the data set as any PPWS measurements taken from these results would falsely show a speed decrease, and other measurements would be inaccurate as participants were not moving when they eventually selected the target.

No further outlying results were found, and we conducted detailed analysis on the data from the remaining nine participants and 260 targets selected per system (520 targets total).

5.2 Participant performance

Table 1 shows the mean and standard deviation of each measurement for each system over all targets. Participants using the haptic system have achieved nearly 38% of their normal walking speed, while participants using the visual system were able to walk slightly faster, at around 44% of their preferred speed.

Statistical analysis using ANOVA shows a significantly lower PPWS is attained when using the haptic system than when using the visual system ($F = 10.25, p = 0.001$), but the haptic system has generated significantly fewer false positives than the visual ($F = 29.5, p < 0.001$). No significant difference was found between the two systems in the time taken to select targets ($F = 0.93, p = 0.335$).

It is possible to combine the targets for the two routes along the circuit and analyse them together, due to the fact

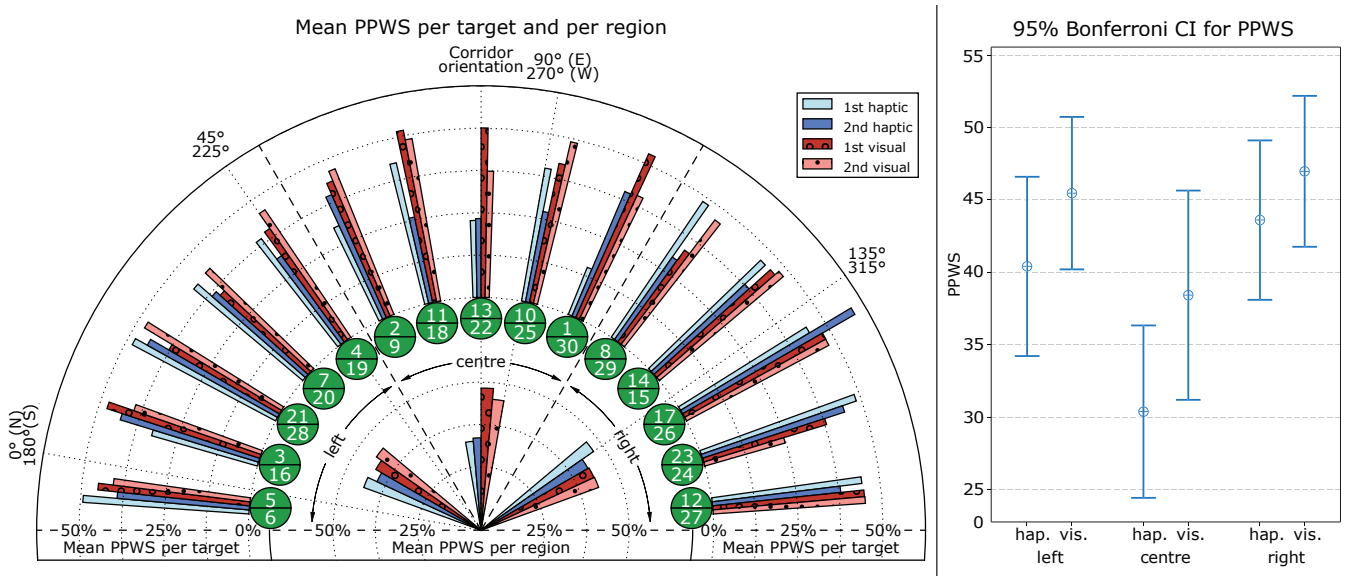


Figure 5: Left: regionally-separated targets, showing the location of each target and the order in which they were selected (circular icons), and the mean PPWS for each target and region per system. Right: Interval plot showing PPWS for each target region per system.




Figure 6: Spread of false positives, showing the number of false positives recorded in the area around each target.

that the targets are mirrored at each end of the course. After pooling the data from these targets, we are also able to segment target positions into separate regions in order to highlight any interesting results that may have arisen due to the target locations. We have chosen to split the targets into left, centre and right of the user’s position, and these results are detailed in the next section.

5.2.1 Regionally-separated targets

Fig. 5 illustrates the mean PPWS per target and per region. Targets in zone T_2 are shown in line with those in zone T_1 to allow visualisation of differences in performance at each target location. Each circular icon shows the order in which participants visited the targets, outer bars show the mean P-PWS per target and inner bars combine these to give a mean value per region.

Interestingly, when segmented into these sections, targets to the left and right of the participant appear to have produced similar PPWS results between systems. ANOVA analysis supports this, with targets to the left and right show-

Behaviour	Haptic %	Visual %
Directly to target	33	49
Probing around target	34	24
Expanding sweep	18	14
Sweeping several areas	6	12
Touching target edges	4	1
Other / unclassified	5	0

Table 2: Percentage spread of participants’ target searching behaviours.

ing no significant difference between systems ($F = 2.82, p = 0.095$ and $F = 1.45, p = 0.23$ respectively), while targets in the centre affect the participant’s PPWS significantly more ($F = 5.52, p = 0.02$): participants using the visual system were able to maintain a higher PPWS for targets in the centre of their field of view.

This pattern is not shown when considering the false positives generated per region (left: $F = 10.05, p = 0.002$; centre:

Evaluating Haptics for Information Discovery While Walking

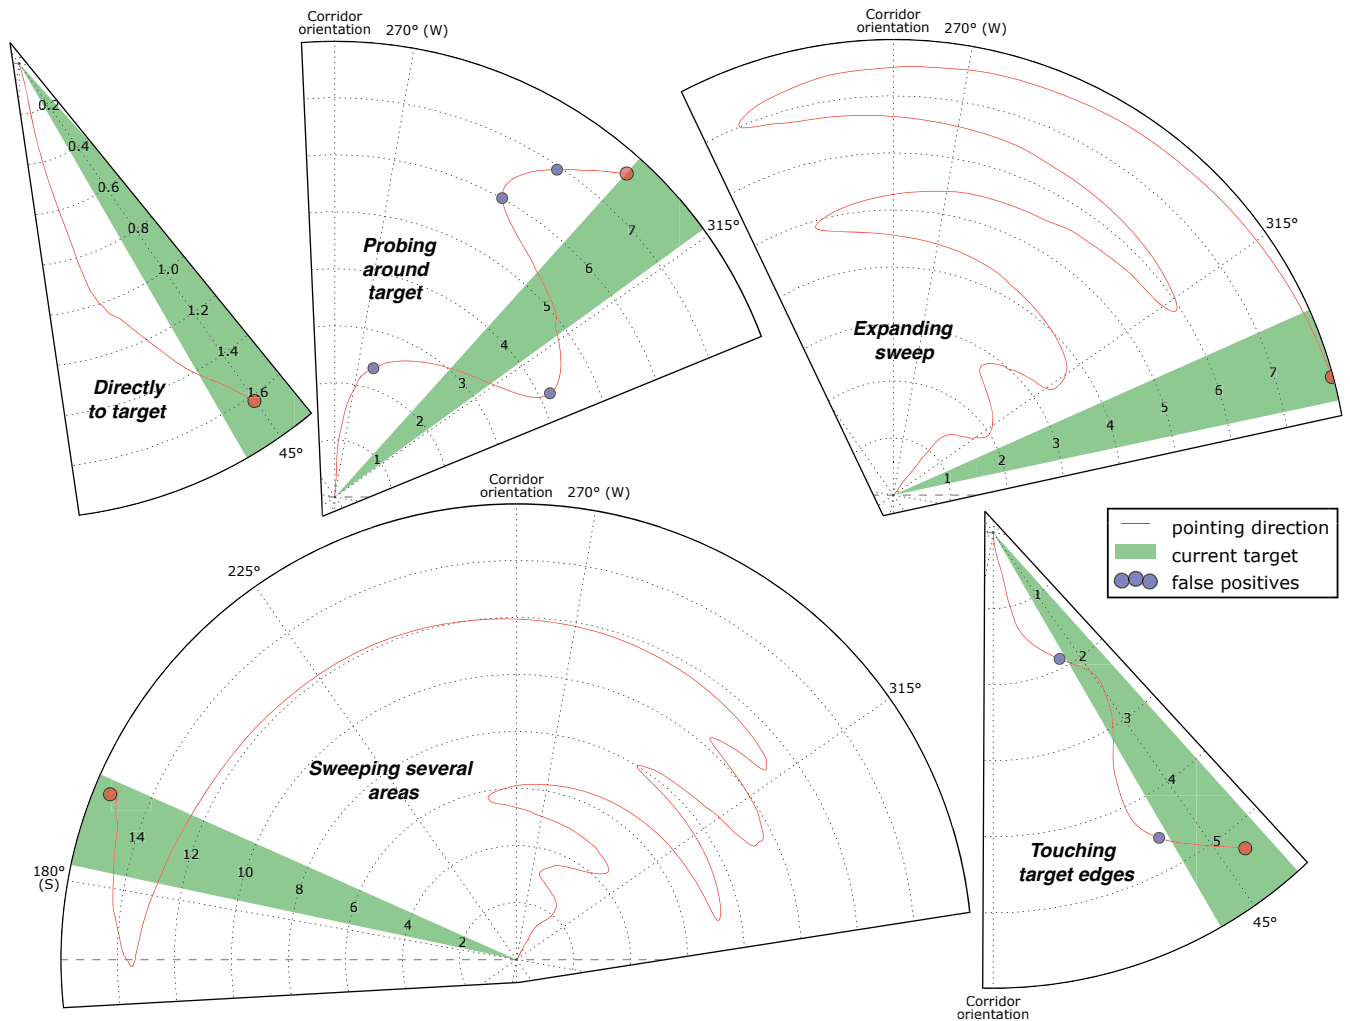


Figure 7: Representative target searching behaviours. The direction the participant was pointing is shown over time (seconds) from the tip to the outside edge of each arc, with the current target area highlighted.

$F = 7.05, p = 0.009$; right: $F = 13.49, p < 0.001$): the haptic system generates significantly fewer false positives regardless of the target region. Delving deeper into this interesting result, we can compare the direction the participant was pointing when they triggered the false positive with the area displayed on the visual system's screen at the time. This analysis shows that 75% of false positives recorded on the visual system were found to have been recorded when the target was visible on the screen but not selectable (i.e. not directly in front of the participant). Fig. 6 illustrates the spread of false positives for each system, shown in relation to the target position, highlighting the denser concentration of false positives recorded on the visual system in the area where the target was visible.

No significant differences were found when considering the time taken to select each target for the left and centre regions ($F = 0.11, p = 0.746$ and $F = 0.11, p = 0.741$ respectively), but a significant difference was found for the right region ($F = 4.21, p = 0.042$): the haptic system took significantly less time to select targets located to the right of each participant.

5.3 Target searching behaviours

Although heading data was recorded for each target, it is difficult to automatically analyse the participants' scanning and searching behaviours due to the variance between participants and the range of target locations. Instead, we have opted to categorise these behaviours into six different types of searching behaviour. Table 2 gives the percentages of behaviours of each type, and Fig. 7 illustrates examples of each type of behaviour.

5.4 Participant subjective ratings

Fig. 8 shows the spread of each of the TLX ratings given by participants. ANOVA on participants' TLX responses shows a significant difference between those of participants using the haptic and those using the visual system for their perception of the time pressure ($F = 6.12, p = 0.027$): participants felt the haptic system put them under less pressure to complete the task quickly. No further significant differences were found when considering any of the remaining TLX ratings ($p > 0.05$).




Figure 8: Spread of TLX responses for each system.

5.5 Observed behaviours and verbal feedback

All participants initially used their right hand to hold the device, and used their left hand to tap the screen to begin the task. After finding a small number of targets, five participants using the visual system began to use both hands to hold and steady the device, while all haptic participants continued to use only one hand to sweep and select.

The majority of participants using the haptic system tended to look almost straight ahead, appearing quite focused on the vibrational feedback they were searching for. Most visual participants appeared from observation to struggle to weave between the cones while looking at the screen, though only three mentioned this in post-study comments.

One participant using the haptic system commented: “if you stop looking down at it it’s much easier. I’d like to be able to use this in real life – it would be very helpful,” while a second stated that “if you’re not looking at it you’re concentrating much more on the feeling rather than the screen,” and “once you get the hang of it it’s really easy.” Another participant stated that they “didn’t need to concentrate [very much],” but also that it was “easy to go past too far.” and have to backtrack to find the target.

Participants using the visual system offered comments more aligned with tourist applications: “[this] would be really helpful for my travelling instead of going to tourist information,” though several criticisms of the system were also raised: “[I have to] catch them as they go past,” and “the delay [lag] makes it much harder.”

6. DISCUSSION

Analysis of the study results shows, encouragingly, that haptic feedback can help users to find and select targets while still able to maintain walking at around 38% of their normal speed on average. Participants using the visual system achieved a slightly higher percentage of their normal walking speed – around 44% on average. When looking more

closely at the results, we can see that for targets to the left and right of the user there is no significant difference in the speed achieved, and only for targets in the centre does the visual system enable a significantly faster walking speed.

Participants have slowed down their walking in all cases, as might be expected. However, there is no great gain in visual performance that could be predicted, and participants found all the targets using the haptic system. Although a comparison of targets overall shows that the visual system allows a significantly faster walking speed, it is important to consider this result in light of user interface familiarity. Users have very little (if any) experience with vibrotactile feedback for interaction, but most will have extensive experience with visual, GUI-based systems. Furthermore, the similarity between haptic and visual results when targets are separated into regions should be seen as a positive point: with very little training, when interacting with targets to either side of their location, there is no evidence that the familiar, visual-based system provided the users with any additional benefit. In addition, while the visual system requires concentration on a screen, the haptic allows the user to focus on their surroundings, as observed during the study.

For targets directly in front of the user, the visual system has allowed participants to walk significantly faster than the haptic. This could be due to the ease of this task when targets are directly ahead – participants started the task and were able to see and select the target almost immediately, whereas participants using the haptic system may have automatically gone straight into the sweeping motion, after which more searching is needed to get back to the initial starting position. For areas to the left and right of the user, the haptic interface shows promise, but the design evidently needs improvements for targets in the area directly in front of the user.

Unexpectedly, the visual system has generated more false positives than the haptic. Referring to the visual data – see Fig. 6 – and the 75% of false positive locations that occurred when the target was visible, this appears to be due to participants predicting target positions before they are selectable. When using the haptic system, fewer false positives are generated, at the expense of time to select targets in some cases. This result highlights the difference in feedback resolution between the two systems, and the trade-offs that result from the differing response levels. The visual display allows users to see a target’s location before they are able to select it, but this can result in unnecessary button presses. The haptic system only notifies the user of a target’s presence when it is selectable, resulting in fewer false positives, but a slower, more exploratory discovery process.

When rating the systems the participants using the haptic prototype rated their perception of time pressure significantly better than those using the visual system, with no significant differences for any other ratings. Verbal feedback, too, seems to have indicated user appreciation of the haptic system, with several positive comments about its ease of use. Participants using the visual system offered several similar comments, but some found the interface difficult to use. Interestingly, only one of the haptic participants commented that they had difficulty using the system, despite the fact that the systems are essentially the same in their interaction method, and only the feedback differs. One explanation for this result is that participants using the visual system were able to see their errors, such as occasions where they missed targets and had to backtrack, or when they came close to se-

lecting but pressed the selection button at the wrong time. Participants using the haptic system, however, could only detect these events on occasions where they skipped over the target and had to revisit it – pointing near to the edge of the target would not provide feedback.

7. CONCLUSIONS AND FUTURE WORK

In this paper we have explored and discussed the viability of using haptic feedback to help provide awareness of information targets in a user's surroundings. The results are promising, and illustrate the potential success of future haptic location-based systems.

Haptic feedback has allowed users to find and select virtual targets in their physical environment while they are moving. Encouragingly, in two thirds of cases performance using haptic feedback has not been significantly different from that using an equivalent visual system. Crucially, in the visual case, exploration of the information space requires a 'heads-down' interaction style which appeared to be distracting to the users. Haptic feedback allowed both interaction with the environment and discovery of information in a 'heads-up' manner.

The study undertaken investigated usage of the system in a realistic but controlled environment and, as such, it is important to consider some of the limitations of these findings. We attempted to simulate users' interaction with their physical surroundings by asking users to walk around a simple obstacle circuit while completing the study tasks. In reality, however, obstacle avoidance is not the only element of interaction with the physical world – interaction and communication with other people, crossing roads or simply observing the scenery are just a few of the many things people do while moving around their environment. However, we believe that these additional tasks support the need for non-visual interactive devices that allow usage in parallel with users' everyday tasks.

Clear extensions of this work are present in the discovery of information for visually impaired users. While research has already been published in this area (see Section 2.2), we believe our interface offers improvements over several systems developed previously. For our prototype, users hold a mobile device which provides feedback for multiple target points; previous systems required a waist belt or jacket, and provided only navigational information for one destination.

Further work could investigate alternative scenarios, such as the discovery of information for later perusal. Previous work (e.g. [6]) has investigated this area, but mainstream applications of this 'delayed interaction' are lacking. An adaptation of the haptic system to delayed information browsing could allow users to explore areas using only a small screen, or indeed no screen at all, requiring only the SHAKE and a small GPS device, rather than the current hardware.

8. ACKNOWLEDGEMENTS

This work is funded by EPSRC project EP/E042171/1, undertaken in collaboration with colleagues at Glasgow University: Roderick Murray-Smith, Stephen Brewster, John Williamson and Craig Stewart. Stephen Hughes provided the SHAKE devices used in our studies.

9. REFERENCES

- [1] L. Barnard, J. S. Yi, J. A. Jacko, and A. Sears. Capturing the effects of context on human performance in mobile

- computing systems. *Personal Ubiquitous Computing*, 11(2):81–96, 2007.
- [2] M. Egenhofer. Spatial information appliances: A next generation of geographic information systems. *First Brazilian Workshop on GeoInformatics*, October 1999.
- [3] P. Fröhlich, R. Simon, L. Baillie, and H. Anegg. Comparing conceptual designs for mobile access to geo-spatial information. In *MobileHCI '06: Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, pages 109–112, New York, NY, USA, 2006. ACM.
- [4] S. G. Hart and L. E. Staveland. Development of NASA-TLX (task load index): Results of empirical and theoretical research. In *Human mental workload*, pages 139–183, 1988.
- [5] S. Holland, D. R. Morse, and H. Gedenryd. AudioGPS: Spatial audio navigation with a minimal attention interface. *Personal Ubiquitous Computing*, 6(4):253–259, 2002.
- [6] M. Jones, G. Buchanan, T.-C. Cheng, and P. Jain. Changing the pace of search: supporting "background" information seeking. *Journal of the American Society for Information Science and Technology*, 57(6):838–842, 2006.
- [7] M. Jones, S. Jones, G. Bradley, N. Warren, D. Bainbridge, and G. Holmes. ONTRACK: dynamically adapting music playback to support navigation. *Personal and Ubiquitous Computing*, 12(7):513–525, October 2008.
- [8] J. Luk, J. Pasquero, S. Little, K. MacLean, V. Levesque, and V. Hayward. A role for haptics in mobile interaction: initial design using a handheld tactile display prototype. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 171–180, New York, NY, USA, 2006. ACM.
- [9] T. Mustonen, M. Olkkonen, and J. Hakkinen. Examining mobile phone text legibility while walking. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 1243–1246, New York, NY, USA, 2004. ACM.
- [10] A. Oulasvirta, S. Tamminen, V. Roto, and J. Kuorelahti. Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile HCI. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 919–928, New York, NY, USA, 2005. ACM.
- [11] J. Pascoe, N. Ryan, and D. Morse. Using while moving: HCI issues in fieldwork environments. *ACM Transactions on Computer-Human Interaction*, 7(3):417–437, 2000.
- [12] H. Petrie, V. Johnson, S. Furner, and T. Strothotte. Design lifecycles and wearable computers for users with disabilities. In *First workshop on HCI for mobile devices*, Glasgow, UK, 1998. Glasgow University.
- [13] A. Pirhonen, S. Brewster, and C. Holguin. Gestural and audio metaphors as a means of control for mobile devices. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 291–298, New York, NY, USA, 2002. ACM.
- [14] B.-S. Shin and C.-S. Lim. Obstacle detection and avoidance system for visually impaired people. In *Haptic and Audio Interaction Design*, volume 4813/2007

- of *Lecture Notes in Computer Science*, pages 78–85. Springer Berlin / Heidelberg, 2007.
- [15] R. Simon, H. Kunczier, and H. Anegg. Towards orientation-aware location based mobile services. In *Location Based Services and TeleCartography*, pages 279–290. Springer Berlin Heidelberg, 2007.
- [16] S. Strachan and R. Murray-Smith. Bearing-based selection in mobile spatial interaction. *Personal and Ubiquitous Computing*, 13(4), 2009.
- [17] S. Strachan, J. Williamson, and R. Murray-Smith. Show me the way to monte carlo: density-based trajectory navigation. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1245–1248, New York, NY, USA, 2007. ACM.
- [18] J. B. F. Van Erp, H. A. H. C. Van Veen, C. Jansen, and T. Dobbins. Waypoint navigation with a vibrotactile waist belt. *ACM Transactions on Applied Perception*, 2(2):106–117, 2005.
- [19] J. Williamson, R. Murray-Smith, and S. Hughes. Shoogles: excitatory multimodal interaction on mobile devices. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 121–124, New York, NY, USA, 2007. ACM.