Sweep-Shake: Finding Digital Resources in Physical Environments

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 the *Sweep-Shake* system, a novel, low interaction cost approach to supporting the spontaneous discovery of geo-located information. By sweeping a mobile device around their environment, users browse for interesting information related to points of interest. We built a mobile haptic prototype which encourages the user to explore their surroundings to search for location information, helping them discover this by providing directional vibrotactile feedback. Once potential targets are selected, the interaction is extended to offer an hierarchy of information levels with a simple method for filtering and selecting desired types of data for each geo-tagged location. We describe and motivate our approach and present a short field trial to situate our design in a real environment, followed by a more detailed user study that compares it against an equivalent visual-based system.

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, gestures.

1. INTRODUCTION

There is an increasingly rich and large set of geo-tagged information available through the web: photos, Wikipedia articles, blog entries, current event schedules, music and videos. Some of this information is explicitly about the place, for example, a description of an historic building. But this common class of content is just one possibility. People might leave digital content in an area as a form of self-expression – digital graffiti or a beautiful poem in a derelict, prosaic part of town. There is also content that is automatically associated with a location, such as the music listened to by others in an area as they pass through [8]. A number of services provide access to this content using location information provided either by the user (via a desktop map search, for instance) or automatically using GPS or cell-phone base station tower information.

Such location orientated views of the web work well on the large screens of conventional computers, but innovations are needed for effective mobile discovery of this content. Our work is motivated by a desire to provide people with 'heads-up', non-screen-based ways of finding digital content associated with the places they pass through. Small-screen, visually-based approaches require the user to grapple with the digital world when they should be immersed in the physical. The need to divide attention between the world of the small screen and the physical environment might limit the much longed for (at least in the research community) vision of a fusion of the physical and digital worlds.

Consider this example interaction scenario, supported by our system, to search the surrounding geo-located information space by simply pointing a mobile device:

> *Walking around town, Ben holds his mobile media player in hand. The street is full of life and he's enjoying people-watching and checking out the new shop window displays. With his hand at waist level, he loosely points toward a bookshop in the distance; there's no response from his device. So, he sweeps it over to his right, to a music store, over the road: the device vibrates. He pauses, and after another gesture, he gets vibration feedback telling him there is some new music content. Selecting it, he begins to listen to the store's top-ten sales. But, today he's a window-shopper, not a buyer and continues on his way.*

We built an experimental set of apparatus with a range of haptic features to explore options for supporting this interaction. The *Sweep-Shake* system helps the user to feel the presence of information in the space around them by providing vibrotactile feedback as a cue to the direction of geo-tagged content. Both the physical area the information relates to and the amount of content available are indicated by haptic pulses, and by sweeping the device around a location the user can assess the possibilities. The system also further extends the haptic response to help the user explore the information space in more detail, by allowing them to zoom in on and filter the available data into different categories.

To help us evaluate this concept, a comparison system was constructed, using a visual interface instead of haptic feedback. This system was designed to provide the same pointing interaction, but with a map display replacing the vibrotactile pulses, allowing us to compare several usability aspects of the *Sweep-Shake* system against a viable alternative interface.

We begin by situating the approach amongst developments in related fields, considering previous approaches and highlighting the rationale behind our designs. Sections 3 and 4 describe the *Sweep-Shake* prototype and our visual control system in more detail. In Section 5 we describe our aims and motivations behind exploring these systems, and then present an exploratory field study under-

MobileHCI'09 September 15-18, 2009, Bonn, Germany
ACM 978-1-60558-281-8.

taken to help understand the value of the approach in a scenario similar to that involving Ben, above. A second study, looking more closely at the effectiveness of the interactions the *Sweep-Shake* prototype affords against our visual system, is described in Section 6. Finally, we close the article with our conclusions and several pointers to future work.

2. BACKGROUND

Researchers have only relatively recently begun to investigate the potential offered by pointing-based mobile interaction. Initially, Egenhofer [4] proposed several *Spatial Information Appliances* before the necessary technologies were integrated into mobile devices. These included a *Smart Compass*, providing turn-based GPS guidance toward a location, and *Geo-Wands*, that help users to identify geographic targets by pointing toward them. Increasingly, with the proliferation of sensory hardware, prototypes combining gestures and sensor data to create devices such as these are being developed.

Fröhlich *et al.* [5] conducted a 'Wizard-of-Oz' style study to assess the viability of 'point-to-select' against several other methods of interaction, concluding that pointing gestures were 'highly attractive and efficient' forms of location selection. Building upon this work, Simon *et al.* [15] investigated the spatially aware mobile phone, a conceptual device to connect the physical and digital worlds. Their framework used a three-dimensional model of a location in conjunction with knowledge of the user's position in order to create a line-of-sight visualisation from their location.

Rukzio *et al.* [13] studied three techniques (touching, pointing and scanning) for locating smart objects, finding touching and pointing to be the preferred interaction techniques if the user had a line of sight or was close to the target device. Pointing was seen as a quick technique that required some cognitive effort but a low amount of physical effort, especially when objects were not within touching distance. Results from their study also showed both pointing and touching to be intuitive techniques, particularly among older participants who wanted to be able to avoid mobile device input as much as possible. More recently, Strachan and Murray-Smith [16] studied the efficacy of bearing-based target interaction, addressing problems that can arise due to the uncertainty of a user's location and heading. A study of their probabilistic approach toward these issues showed that targets could effectively be selected, even when fairly closely spaced.

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 extends the pointing phase of this process to include information discovery and filtering, helping the user to interact more closely with the location they are exploring.

2.1 Non-visual location-based feedback

Moving away from visual feedback, Strachan *et al.* [17] use 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 similar way, to obviate usage of a screen, Holland *et al.* created *AudioGPS* [7], using audio to provide representations of the direction and distance of waypoints, finding the system adequate for navigation tasks to within eight degrees of a target. Jones *et al.* [9] created a similar system to help users navigate through a virtual environment using ambient spatial audio, providing cues as to the direction and distance of a specified target by adjusting the fade and balance of an audio track. Brewster *et al.* [2] used structured audio messages to help users navigate through four levels of an hierarchical menu, finding over 80% accuracy for location identification. Similar to our zooming interface (described in detail in the next section), cues are used to help users discover any currently available navigation possibilities and determine their present location in the menu hierarchy.

Van Erp *et al.* [18] created a vibrotactile belt to help indicate to a user the direction of waypoints via haptic feedback. Their system used eight tactors placed around the user's waist to help indicate the direction of a waypoint: each tactor vibrated to indicate a certain direction, and distance was also encoded into the vibrational feedback. Similarly, Lindeman *et al.* [10] used tactors placed around the user's torso to help them in a building search task, finding that using directional vibrotactile feedback helped improve users' performance and significantly reduced the number of undiscovered areas. 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 they suggest that using haptic feedback for this task can help users to accurately navigate around obstacles without losing track of their path. Luk *et al.* [11] describe a prototype for mobile haptic interaction where piezoelectric actuators are used to provide several tactile sensations ranging from simple buzzes to complex 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 information discovery, providing cues as to the location of several possible points of interest rather than guiding the user to their journey destination.

3. SWEEP-SHAKE SYSTEM

The *Sweep-Shake* system allows users to discover geo-tagged information in the environment around them with simple pointing and sweeping gestures, as illustrated in Figures 1 and 2.

As the user passes their focus near an information hotspot, vibrational feedback is provided to indicate both the direction of, and quantity of information related to, this potential point of interest. If the user is curious, they can 'zoom in' to explore in more detail. Here we offer options for filtering and selecting several data types – for example, text, images, videos and audio content – with a straightforward directional motion. Tilting the device toward the edges of their initial target gives the user a quick overview of the types of location-related information they can access, as shown in Fig. 3.

The system, then, allows the user to seek out digital resources associated with their environment without having to pause to look at a screen during the discovery process. We envisage people using the device both while they are stationary and when walking. The following sections outline how the design is implemented and describe the system in more detail, highlighting its novel interactive features.

3.1 Mobile hardware

We combine several separate devices in our design concepts, allowing us to quickly develop and test prototype interaction meth-

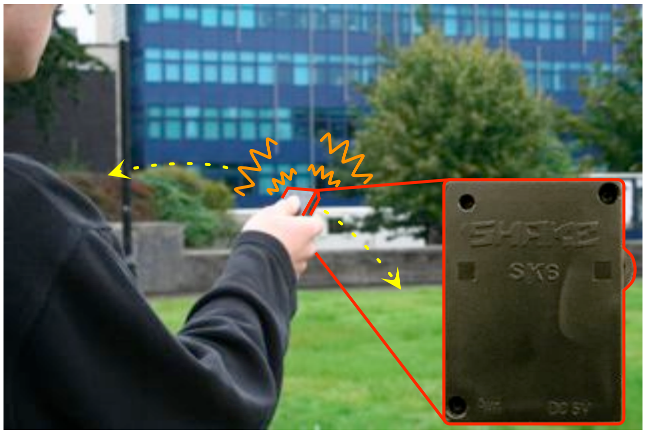

Figure 1: *Sweep-Shake* system in use: the SHAKE device (shown inset) vibrates when the user points toward a location with geo-tagged content.

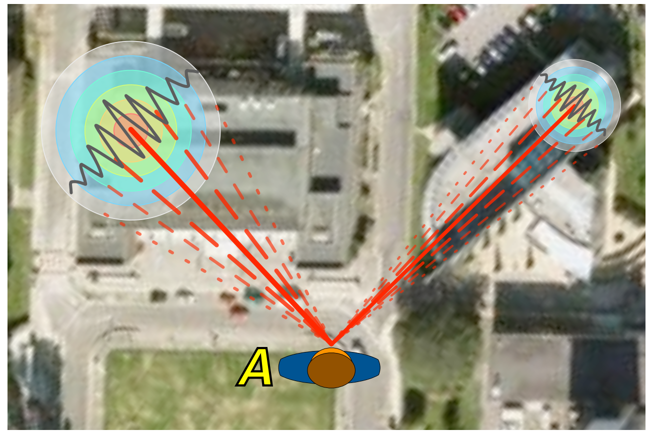

Figure 2: Haptic feedback: the spread of the haptic area helps the user, at point *A*, to determine the size of the target, and the increase in vibration frequency toward the centre of the target guides them to its centre point.

ods. To determine a user's location we use a standard GPS receiver. For real-time orientation and heading (compass) data we use SHAKE sensor packs (see [19]). The SHAKE SK6 is a small Bluetooth device incorporating three-axis accelerometers, magnetometers and angular rate sensors, dual-channel capacitive input sensors and a navigation button (visible in Fig. 1, inset) which allows for user input. Also included is a programmable vibrating motor that we use to provide a range of vibrotactile effects as haptic feedback.

Signals from each of these devices are communicated wirelessly using a Bluetooth serial port profile to a Sony VAIO Ultra-Mobile PC (UMPC), which records the data and and displays point of interest information to users on request.

3.2 Haptic browsing

The *Sweep-Shake* system's browsing mode helps users to explore the spaces around them to search for interesting geo-tagged information, without the need to look at a screen until they have discovered data that interests them. Users hold a SHAKE device in their hand and, whilst moving around their environment, can either point specifically toward locations that they are interested in to detect whether information is present, or sweep the device around the area to search for possible information hot spots using a scanning motion. If they point toward a location that contains geo-tagged information, the device gently vibrates to indicate the presence of information about that target.

Scanning the area around the target will help the user determine the exact direction of the object (see Fig. 2): the vibrotactile feedback increases in magnitude until they are pointing directly at the target. This interaction also offers the user an indication of the quantity of geo-tagged information that is present: targets with little information appear small and take up little of the user's field of view; targets with more information appear larger, indicating that a larger amount of geo-tagged data is present.

As the user moves around their environment, the GPS device retrieves their location (latitude, longitude). This position data is used to automatically refresh the selection of available points of interest to include all those within 100 metres of the user's current location. This relatively short range restriction was an intentional design decision, made as a result of our previous findings [12] suggesting that specifying distance is a complex task regardless of display, and that closer objects are selected more often than those in the distance. In addition, closer objects occlude those further away – we assume the user has an interest in the points of information in their immediate vicinity, again as a result of findings in [12]. Even with this simplified pointing interaction we are still able to provide users with real-time location information, using a system that is realistic with current mobile technology. There is no need for a complex location model; only geo-tagged information (currently widely available on public websites such as *Wikipedia* and *Flickr*) is required.

As the user browses the information around them they are able to leave the browsing mode and focus more specifically on one particular location, zooming in to view in more detail. Pressing the SHAKE's navigation button whilst vibrational feedback is being felt causes the system to move into a zooming mode, and also causes a distinct 'zoom in' vibrotactile pulse to be generated to indicate to the user that they are focusing on a particular information hotspot.

3.3 Haptic zooming

The *Sweep-Shake* system's zoomed mode is intended to help the user to browse the available information about the place they have selected. However, instead of simply retrieving and displaying all nearby geo-tagged data, any relevant data is first segregated into four distinct clusters based on the types of information that exist at that place.

This novel interaction method provides a simple way for users to find and filter specific types of location-related information. By tilting the SHAKE to four corners around the location of their original target, the user can detect the presence of text, images, video and audio content via vibrotactile feedback (see Fig. 3). At each of the four zones, a vibrotactile pulse is generated to indicate the presence of that type of information, with a different pre-set pulse being generated for each of the information categories. If no pulse is felt then no data of the requested type exists, and the user can simply zoom out and move on.

Whilst pointing to a specific category of information, if the user discovers content of the type they are searching for then they can press the SHAKE's navigation button to retrieve and display it on the UMPC's screen, again feeling a distinct haptic pulse to confirm this action. The data is filtered before being presented: for example if the user pressed the selection button whilst feeling vibration indicating the presence of images at a location, only images would be displayed.

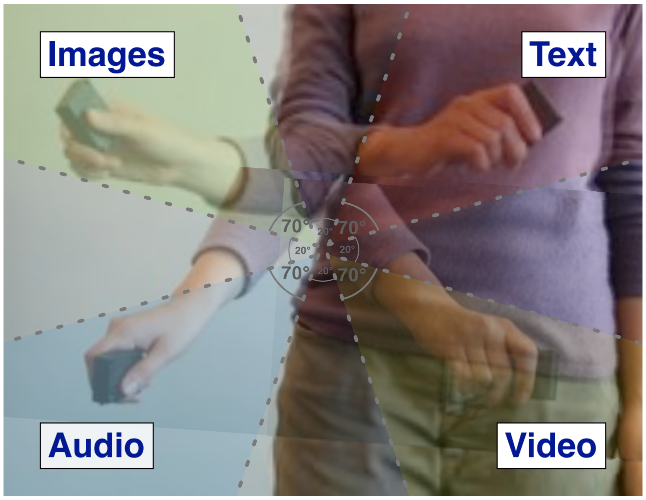

Figure 3: Haptic zooming: pointing toward the four haptic areas triggers feedback.

Figure 4: The visual browsing comparison system, showing the on-screen display (inset). Point *A* has previously been selected, point *B* is currently being pointed at by the user and is activated. Point *C* represents the user at their current position. The two remaining points, *D* and *E,* have not yet been visited.

When the user has finished browsing the information, pressing the navigation button will return to the zoomed mode, again generating a confirmation pulse, and the user can explore the other categories of information that are available. If the user is not interested in retrieving any further information, or there is none available, then pressing the navigation button again will generate a distinct 'zoom out' pulse and return to the original browsing mode.

4. VISUAL MAP SYSTEM

In addition to our haptic *Sweep-Shake* prototype, we constructed a visual-based system to provide a comparison during our user studies. This system is a visual analog of the haptic system in its implementation, to the extent that it was possible. The system differs in that instead of offering vibrotactile feedback, the user is given a visual display of their actions. For this system, users hold a UMPC with a SHAKE firmly attached to the back, ensuring that any movements made by the user whilst holding the display are directly transferred to and recorded by the SHAKE. The UMPC has several control buttons positioned around its screen, and we use one of these to allow the user to control the zooming interaction in an equivalent way to that used in the *Sweep-Shake* system.

Similar to the initial browsing functionality of our haptic design, the visual system provides a browsing mode, presenting the user with an aerial photo of their location with markers overlaid to show potential objects of interest in the same 100m field of view (see Fig. 4). As in our haptic system the user is able to determine the quantity of information at each location, illustrated by the size of the marker: a larger radius indicates more information is present, and the centre point of the object, shown as a cross at the centre of each icon.

As the user moves around their environment, their position and orientation are retrieved from the GPS device and the display is refreshed to show their current location and field of view. As they turn to face potential targets the map display is re-oriented to ensure that the visible display represents the real-world surroundings as seen by the user.

When a potential target is centred in the user's field of view, it is visually highlighted to show that it is available to be explored. Pressing the zoom control button on the UMPC at this stage will switch to the zoomed-in mode, displaying a visual zooming in effect. This zooming effect lasts the same length of time as the pulses given in our haptic system, helping to ensure we are able to compare the systems fairly. In a similar way to the *Sweep-Shake* system, any available information is clustered into type categories and represented by four distinct data icons, each in separate corners of the display. Users can select an information category by touching its icon on the screen of the device, at which point the relevant filtered information is displayed.

When the user has finished browsing the information, pressing the zoom button again will return to the cluster of data types, and pressing a second time will return to the initial aerial photo browsing mode.

5. EVALUATING THE SYSTEMS

We performed studies aimed at exploring the extent to which haptic feedback can be used for exploratory navigation of the physical environment. We had two research questions:

I Browsing: How effectively can people identify digital resource targets using haptic feedback?

II Zooming: How effectively can people filter resource types using haptic feedback?

To help understand these questions we carried out two studies. The first study – an exploratory field trial – placed the *Sweep-Shake* system in a realistic environment and explored its usability possibilities and suitability for point interaction when moving, providing qualitative feedback about its value to potential users.

The second study investigated the use of the *Sweep-Shake* system's vibrotactile feedback as a cue for target discovery, compared to the visual map system. This study provided us with data to help assess the relative performance of the haptic system when used to perform a stationary task in a more controlled environment than that of our first study.

The following sections describe the two studies carried out, detailing our methods, findings and discussion of the results we collected.

6. EXPERIMENT 1: EXPLORATORY FIELD TRIAL

We performed a field trial to enable us to examine the usage of the *Sweep-Shake* system in a realistic scenario, and to help identify aspects of its design that could be refined. This campus-based study made full use of the system's potential for finding information whilst mobile, using live information about several locations.

6.1 Method

Four participants aged from 18 to 35 were recruited for a 45-minute field study. Three participants were students, one was a member of university staff; two participants were male and two female. One participant had previous experience of accelerometer-based interaction methods with the Nintendo Wii games console; the remainder reported no prior experience of this style of interaction.

Before the study commenced, we placed information hot-spots over five buildings located on the university campus. These target points all contained at least five separate items of information about the location, ranging from web page text, interior and exterior images to videos and audio content of recent and upcoming events. The system was in range of at least three of these points from any point on the route through the campus, giving participants a choice of targets to select at all times wherever they moved.

At the start of each study session, each participant was met individually and introduced to the equipment and its purpose, followed by a walkthrough of a short example usage scenario. Participants were then first introduced to the visual system to illustrate the interaction methods employed, then given a demonstration of the *Sweep-Shake* system and, as a form of training, were asked to practise both target selection and zooming in turn to mark a number of targets from a window in our laboratory.

Each participant was then taken outside to the far edge of the university campus, and was asked to walk through the area, exploring the space around them to find any information that might be of interest to them. As the participants moved through the campus, the researcher observed from a distance, but did not interact with the participant. When participants reached the opposite side of the campus, they were asked to provide verbal feedback resulting from their experience of using the system.

6.2 Findings

All participants found and zoomed in to each of the five targets available. Here we discuss behavioural observations and subjective feedback that arose from participants' experiences of the system.

6.2.1 Participant behaviours

Each participant chose to begin the study by standing still and using the device to systematically scan around the starting location for potential points of interest. From this point onward, participants slowly walked through the university campus, at first stopping whenever they felt a vibration, but toward the end of the session, continuing to walk and instead interpreting the feedback as they moved.

All participants viewed each category of information at least once, but after this first interaction displayed interest primarily for text and photos, with only one participant viewing further video and audio content after they had initially discovered its presence.

Two participants held the SHAKE by their waist as they walked, moving it around to scan constantly, and then temporarily walking more slowly when they felt a vibration.

6.2.2 Verbal feedback

All participants commented that the system was fun to use, and that they enjoyed pointing toward locations to feel the presence of related information. Each participant also stated that once they had been able to understand the system from their interaction with the initial point they discovered, they were much more confident in its usage. One participant felt that the interaction was like 'playing a game to catch the points', and that had the vibrotactile effect been more stable their performance would have been better.

Two participants recalled specific experiences of being lost in foreign cities where they would have liked to have been able to find information around them to help get a sense of their location. Two of the participants suggested a 'guide me' mode, similar to that of a standard GPS device but instead using haptic feedback to indicate the general direction of the target location. Finally, two participants thought that finding important information quickly would be hard with this system, and that the vibrotactile feedback could be more effective if it were smoother in its responses. These participants had felt the haptic feedback to be either off or on, rather than a steady increase in intensity toward the centre of the target.

6.3 Discussion

Encouragingly, all participants were able to discover the targets while moving, after only basic training. Participants offered positive comments about the ability to feel and explore real geo-tagged data, and enjoyed interacting with their surroundings in this way.

Participants seemed to be uninterested in the audiovisual content that was available, instead preferring to perform a quick scan of static text and images. Although this finding could be due to the small amount of content available for this study, it could also suggest that users are reluctant to commit to watching or listening to these types of media whilst on the move. Future extensions of this system could add methods to allow users to take a copy of interesting location information with them on their onward journey, possibly improving the user interaction with these types of content.

The 'guide me' mode suggested by participants is an obvious extension of the *Sweep-Shake* system, with participants requesting the direction of a location beacon rather than turn-by-turn guidance. Visual or audio-based versions of these types of systems are commonly used in tracking beacons, though possible issues can arise when obstacles have to be surmounted and routes are indirect (*cf.* [9]).

It was interesting that two participants held the device by their waist and used it as a 'ping': a background cue to let them know when interesting locations were discovered. As suggested in our initial scenario, this interaction method allows the user to concentrate on the physical environment, only switching to the digital when potentially interesting information is available. This interaction is much more of a negotiation between user and target rather than a simple information request and delivery, and future work in this area could lead to modeless information discovery, with the system's state being automatically changed dependent on the user's behaviour toward potential points of interest.

7. EXPERIMENT 2: HOW EFFECTIVE IS THE SWEEP-SHAKE SYSTEM?

In addition to our exploratory study, a controlled study was performed to enable us to evaluate and compare the *Sweep-Shake* prototype against the visual system. Our main interest was to explore how effectively haptic feedback allows users to find, select and zoom in to potential information points in the space around them with regards to accuracy and time taken.

7.1 Method

In order to support our interest in the effectiveness and efficiency of the browsing and zooming actions, the *Sweep-Shake* and visual systems were modified to focus on target selection rather than information retrieval. To enable this, no geo-tagged data was created for the systems to ensure that users' behaviours were not affected by the quality of the data retrieved. In addition, to allow us to concentrate on participants' ability to find specific locations, each target became un-selectable after it had been selected once. In the *Sweep-Shake* system, once targets had been selected, pointing toward them no-longer triggered vibration feedback; in the visual system this was achieved by changing the icon used for the target once it had been selected (see point *A*, Fig. 4).

7.1.1 Participants

Thirty-two participants aged from 18 to 65 were recruited for a half-hour study. 14 participants were university staff members, 18 were students; 16 participants were male and 16 female. Nine participants had previous experience of accelerometer-based interaction methods with the Nintendo Wii games console; the remainder reported no prior experience of this style of interaction. None of the participants had taken part in the first study.

7.1.2 Conditions

Participants were equally split between two study conditions: the *Sweep-Shake* system and the visual system. Randomly, four of the the nine participants with prior accelerometer experience were allocated to the *Sweep-Shake* system and the remainder to the visual system. The 23 remaining participants were randomly allocated between the two systems, giving 16 participants per system. Before the study began, two sets of six distinct pre-set targets were created within a 100 metre radius of the participants' location during the study. Half of these targets were used for an initial training session, the remainder were used for the tasks participants were asked to complete. Each of these second set of targets was also randomly allocated between one and four sub-targets, with 15 created in total over all six targets. These sets of targets were identical for each participant over the entire study.

7.1.3 Tasks

Participants completed the following tasks during the study:

- Find and select each of the six targets in the area around them (see Fig. 5). When a target was found:
  - Zoom in to the target
  - Find and select each available sub-target
- Complete a questionnaire to rate their usage of the system
- Give any verbal feedback resulting from their usage of the system

7.1.4 Measures

Five measurements were automatically recorded from each participant whilst they browsed around their environment. These readings, and our reasons for recording them were:

B_1 The number of targets successfully found and selected, recorded to allow us to measure the effectiveness of the targeting feedback and interaction methods that each system employed.

B_2 The time taken to select each target, which helps to indicate whether participants were able to find targets within a reasonable amount of time, and whether some targets could have been harder to find than others. For the first target this is recorded from the system start time; for subsequent targets this is measured from the time the participant zoomed out from the previous target.

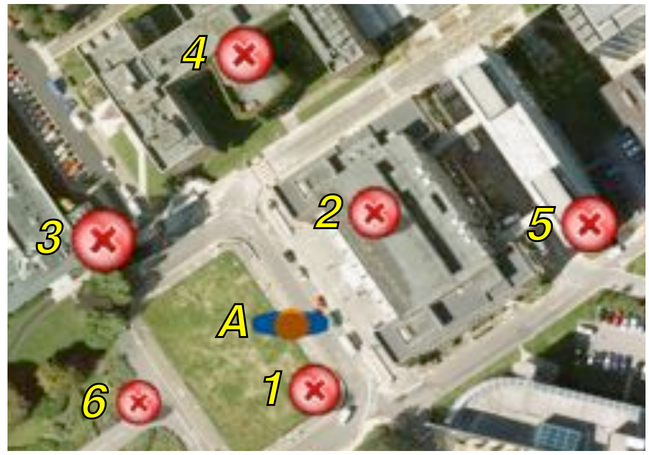

Figure 5: Targets used for the study. The participant, standing at point *A*, points toward and attempts to zoom in to each target in turn in any order they wish.

B_3 The number of target activations without selections (i.e. when a user pointed at a target but failed to press the button to select it), which gives a measure of the ability of the system to successfully communicate to the user the presence of a target and allow them to select it.

B_4 The delay between the participant activating a target and the selection of the target, which could highlight any differences between the participants' ability to respond to visual or haptic feedback cues.

B_5 The number of false positives (i.e. when a participant pressed the button to select a target but was not actually pointing toward any of the targets), allowing us to highlight any areas of unclear feedback or interaction confusion.

When zoomed in on individual targets, different automatic measurements were taken:

Z_1 The number of sub-targets successfully found and selected from each target, allowing us to measure the effectiveness of the feedback given when zoomed in upon a location.

Z_2 The time taken to find each of up to four sub-targets for each target, showing whether participants would be able to find information categories within a reasonable amount of time. For the first sub-target this is recorded from when the participant zoomed in; for subsequent sub-targets this is from the time the previous sub-target was selected.

Z_3 The number of times participants needed to zoom in on each target to select all of its sub-targets, which helps to show if participants could find all the information about a target without repeated interactions.

Z_4 The time spent feeling the haptic feedback for each sub-target, giving an indication of whether the feedback given was clear and quickly interpretable (not applicable to the visual system).

Z_5 The number of times the participant touched or pointed to each sub-target, indicating whether participants would be able to get an overview of the information available about each target without the need for multiple interactions.

Participants were also asked to complete a questionnaire based on the six factors of the NASA Task Load instrument (TLX) [6], examining their perception of the costs involved in using the system. They were asked to rate the mental, physical and temporal demand imposed, their success in performing the selecting task, the overall effort needed and their frustration with the system. Each of these dimensions was rated on a scale of 1 (negative, e.g. high frustration, low performance) to 7 (positive, e.g. low mental demand, high performance). In addition, each participant was asked to rate, on the same scale of 1-7, specific aspects of the prototype's usage and usability. The features rated were: their overall ability to identify the actual targets they were marking, how fast they felt they were able to mark the targets, and the perceived usefulness of the system for finding directions, points of interest, urgent information or for simply filling time.

7.1.5 Procedure

At the start of each study session, each participant was met individually and introduced to the study and its purpose, followed by a discussion of a short usage scenario. Participants who had been allocated to the visual system were then given a short usage demonstration and, as a form of training, practised both target selection and zooming by marking a number of targets from a window in our laboratory. Participants allocated to the *Sweep-Shake* system were first introduced to the visual prototype to help illustrate to them the interaction methods. These participants were then given a demonstration of the *Sweep-Shake* system and were asked to perform the same training exercise as participants using the visual system.

Following this short training session, each participant was then taken outside to a fixed location in an open space on our university campus – every participant in the study stood in the same location. As an additional training exercise participants were informed that there were six targets in the area around them, and were then asked to use the browsing capabilities of the system they had been allocated to locate and select each of the first set of six targets, with no zooming interaction.

Participants were then informed that there were a further six targets to explore, and proceeded to use the system to select and zoom in to each of the second set of six targets, finding and selecting each available sub-target. While participants completed this task, the researcher observed their behaviours and methods used in finding and selecting targets. After completing this task, participants completed the questionnaire to rate their usage of the system. Finally, participants were asked for verbal feedback resulting from their usage of the system, and this was recorded by the researcher.

7.2 Findings

All participants attempted to find (and believed they had succeeded in finding) each of the six targets presented to them. Each participant also completed the rating questionnaire and offered several verbal comments about the system they had been asked to use. We discuss our findings below, comparing results from the *Sweep-Shake* system against the visual map alternative.

7.3 Objective results

Table 1 shows the mean and standard deviation of each measurement for the browsing modes of each system. Six participants using the *Sweep-Shake* system found all targets, with the remaining ten participants missing between one and four of the six targets. Thirteen participants using the visual system found all targets, and the remaining three found all except one. 177 false positives were recorded over all participants: 91 when using the *Sweep-Shake* system and 86 when using the visual system.

Measurement	Sweep-Shake	Visual
Number of targets found	4.5 (1.5)	5.8 (0.4)
Time to select each target (seconds)	16.5 (22.3)	8.8 (5.6)
Activations w/o selections (per target)	9.3 (6.0)	0.7 (0.5)
Delay after target activation (seconds)	1.7 (1.4)	1.9 (1.6)
False positives (per target)	0.9 (1.1)	0.9 (0.6)
Total time taken (seconds)	105.2 (32.3)	81.7 (26.4)

Table 1: Mean and standard deviation (in brackets) of each of the measures recorded when browsing, and the total time taken to complete the task.

Measurement	Sweep-Shake	Visual
Time to find each sub-target (seconds)	3.0 (4.5)	2.1 (2.3)
Zooms required (per target)	1.0 (0.2)	1.0 (0.4)
Interaction time taken (seconds)	8.4 (11.1)	*n/a*
Number of interactions (per sub-target)	3.5 (3.7)	1.0 (0.2)

Table 2: Mean and standard deviation (in brackets) of each of the measures recorded when zoomed in.

Table 2 shows the mean and standard deviation of each measurement for the zooming modes of each system. Participants using the *Sweep-Shake* system found 67.5% of the available sub-targets on average, while participants using the visual system found 98.8%.

We conducted further analysis on the results recorded using GLM ANOVA both between targets and between systems. The following sections list statistically significant results, in addition to figures showing the 95% confidence interval for the mean of each value.

7.3.1 Time to find targets

Fig. 6 shows the spread of times taken to find each target. No significant difference was found between targets when using the *Sweep-Shake* or visual systems individually ($p > 0.05$), but when comparing systems there is a significant time difference between the *Sweep-Shake* and visual ($f = 10.32, p = 0.002$): participants using the visual system took less time to select targets.

7.3.2 Activations without selections

Fig. 7 illustrates the spread of results for each target over both systems. Significant differences were found between targets when using both the *Sweep-Shake* ($f = 5.91, p < 0.001$) and visual ($f = 8.75, p < 0.001$) systems, possibly indicating that some targets have been harder to activate regardless of the system used, or that participants have attempted to explore their environment before focusing in on one target. When comparing systems there is a significant difference ($f = 37.67, p < 0.001$): the visual system causes less activations without selections.

7.3.3 False positives

Fig. 8 illustrates the number of false positives recorded for each target. A significant difference is evident between targets when using the *Sweep-Shake* ($f = 6.69, p < 0.001$) and visual ($f = 3.72, p = 0.004$) systems, showing that some targets have been harder to select regardless of the system used. No such significant difference is evident when comparing systems ($f = 0.16, p > 0.05$): the number of false positives produced was not caused by one particular system.

7.3.4 Time to find sub-targets

When comparing systems for the time to select sub-targets, the visual system is significantly faster ($f = 8.21, p = 0.004$).

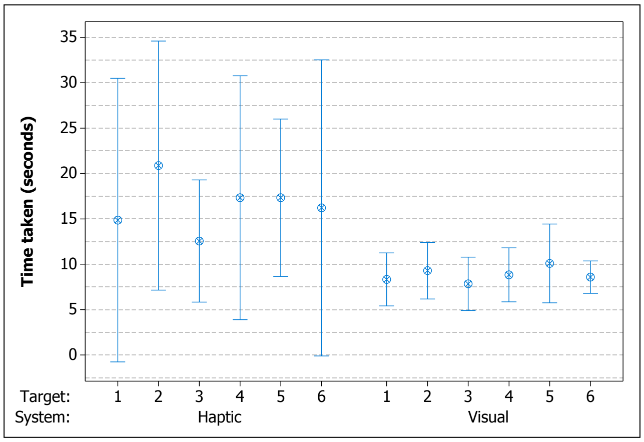

Figure 6: Time taken to select each target

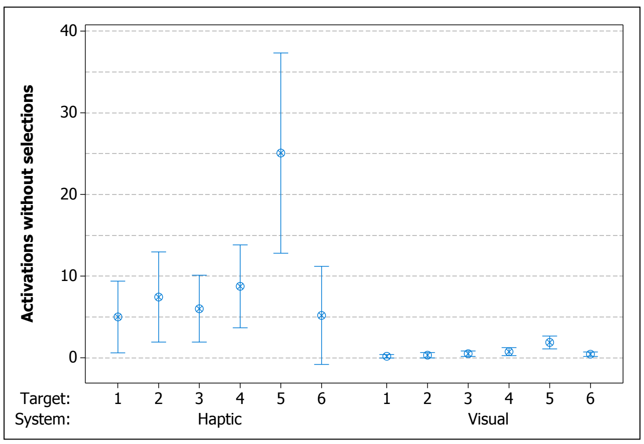

Figure 7: Activations without selections

7.4 Subjective ratings

Fig. 9 (top) shows the mean of each of the TLX ratings given by participants. Significant differences were found between the two systems when considering performance ($f = 11.67, p = 0.002$) and frustration ($f = 6.02, p = 0.02$): the visual system was rated significantly better on these aspects. Fig. 9 (bottom) shows the ratings given for system usage and usefulness. Significant differences are evident for identification accuracy ($f = 19.35, p < 0.001$) and speed of marking ($f = 12.45, p = 0.001$), with participants rating the *Sweep-Shake* system lower on these aspects. No further significant differences were found when considering any of the remaining subjective ratings.

7.5 Verbal feedback

Four participants of the 16 using the *Sweep-Shake* system said they would like to be able to use the device in their everyday lives, with two of these noting that using the system would get much easier with practise. Three participants specifically stated that using the *Sweep-Shake* system was fun, and five others said that the act of pointing at a location was very easy and they could get used to it quickly. One user commented that the system was 'much more helpful than my GPS for finding places', but that it would be harder when the information space was cluttered, a statement echoed by three others.

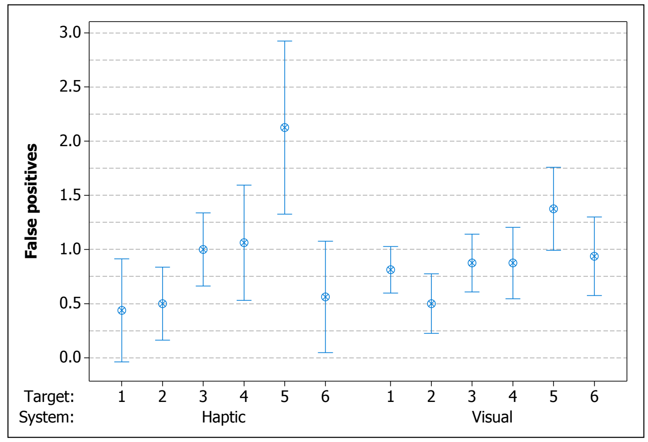

Figure 8: False positives

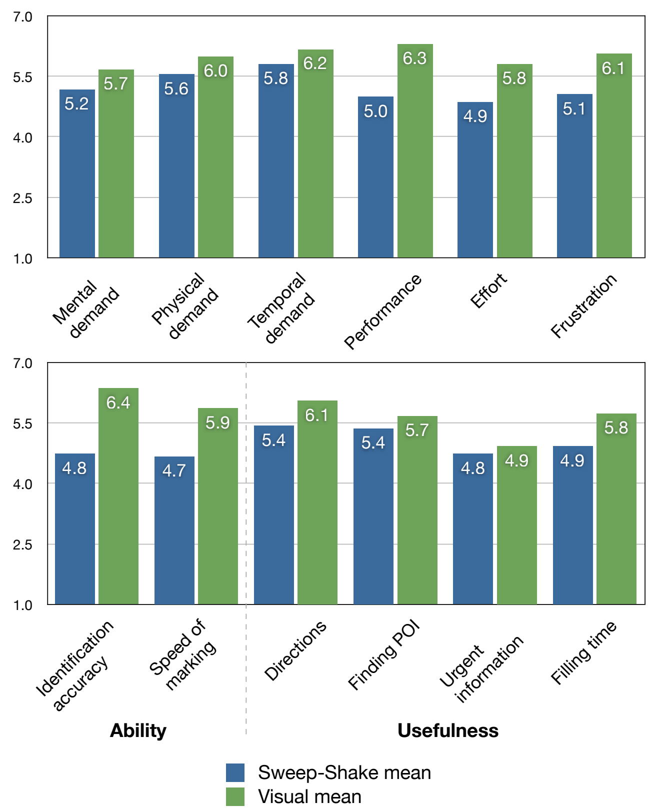

Figure 9: Subjective system usage ratings. Top: participants' TLX scores, bottom: ratings for system usage and usefulness of the system for specific tasks.

Four participants found the different vibration pulses hard to distinguish from each other, and another worried about the social acceptability of pointing to objects. A further three participants noted difficulty in remembering which mode the system was in when attempting to zoom in to targets. Six participants said they appreciated the general idea of the device, and specifically mentioned the same 'guide me' mode suggested by participants in our field study, indicating that this is a common request.

Four participants of the 16 using the visual map system commented that they found it easy to use, and that they really liked being able to see the buildings around them from above. One participant commented that they preferred the zoomed in mode for its more comprehensive view of the surroundings at that particular location. Another participant stated that they found the discovery of each target easy primarily because they were familiar with the test area – had the study been conducted in a new location then they felt they would have had less success in completing the task.

Three participants using the visual system found it to be fiddly and awkward, and struggled to use the re-orienting display. In addition, nine participants mentioned problems in seeing the display due to reflections on the screen of the UMPC.

7.6 Discussion

It was encouraging that participants were able to find information and select targets with only vibrotactile feedback to aid them. In addition, the ability of participants using the *Sweep-Shake* system to explore the available information when zoomed in on targets suggests potential for more successful target discovery after further user training and exposure. Indeed, it is important to consider these results in light of the low degree of familiarity with haptic interfaces. Users have very little (if any) experience with vibrotactile interaction, but extensive experience with GUI-based systems. In addition, the the *Sweep-Shake* system offers much lower resolution feedback than the visual interface, but participants have still been able to use it to successfully discover and select targets.

As might be predicted, participants using the visual system were able to discover and select targets more accurately and faster than those using the *Sweep-Shake* system (9.3 versus 16.5 seconds). This difference in speed and accuracy, although clearly a negative aspect for important, time critical information access using the *Sweep-Shake* system, can also be seen as a positive point. Participants are stimulated to explore their environment whilst also taking in the visual scene in front of them, rather than concentrating on the digital representation on their mobile screen. This stimulated exploration is perhaps illustrated by the number of activations without selections in our results – participants seem to have skimmed over the available targets before deciding to go back and explore them in more detail later.

Participants had difficulty in selecting some targets when using both the *Sweep-Shake* and visual systems – target five in particular has caused a larger number of false positives and activations without selections. One possible explanation for this is its distance from the user. Targets further away take up less of the user's field of view so may be harder to select, suggesting a need for our systems to compensate for distance when determining the initial target size. Alternatively, the selection of this target could have been affected by its proximity to target two, but this seems unlikely as targets became un-selectable after being selected once.

When asked to rate the systems, the visual display is subjectively rated only marginally higher than the *Sweep-Shake* system. In their verbal feedback, too, participants seem to rate their usage of the *Sweep-Shake* system similar to that of the visual system, with several comments about ease of use and the natural feel of pointing to locations. Participants using the visual system gave positive comments about the map display giving them a view of their surroundings from above, a finding that was also highlighted in our earlier work [12].

Participants using the *Sweep-Shake* system seem to have had difficulty identifying the mode they were in at times, and false positives and verbal comments resulted. Clearly this is an area where further haptic feedback development could offer usability improvements. While it is relatively easy for users to assess mode changes in visual systems, a richer set of haptic forms may be needed to clearly communicate shifts in system state. Surprisingly, this issue was not raised in our field study, suggesting that the presence of real location data helps to alleviate difficulties in determining the system state.

It was interesting that several participants mentioned the use of haptics applied to a familiar GPS navigation task, suggesting that they prefer wayfinding to be a background task, available on request, rather than its common implementation of pushing all available information whenever it could be useful. Previous work (e.g. [17]) has investigated the use of feedback when the user leaves a pre-defined path, but this result suggests that perhaps users would prefer to receive these prompts only when desired, rather than on every deviation from the ideal route.

8. CONCLUSIONS AND FUTURE WORK

This paper has discussed using haptic feedback to provide information awareness and content filtering in a physical environment. The results are encouraging, and illustrate the potential success of haptic location-based interaction: approaches such as these may allow people to investigate their physical and associated digital worlds in an engaging, 'heads-up' way. We believe it is vital to provide such mechanisms if mobile access to (and creation of) place-based information is to reach its full potential.

Clearly, haptic feedback has allowed users to find and filter virtual targets in their physical environment. Although the feedback provided is not as effective as a visual representation of the same scene, the visual display requires full-screen, 'heads-down' engagement with the device, which could be particularly distracting while on the move. We are currently conducting a further study to investigate refinements to the *Sweep-Shake* system that could enhance the user's experience over that of a visual system while walking.

Further work is needed to improve and clarify the haptic cues that are provided, particularly so that users can sense the change in level from target to sub-targets. One possibility would be to combine haptic feedback with other techniques, such as audio. Indeed, previous investigation has highlighted advantages of this approach: Chang *et al*. [3] found that basic audio-haptic couplings enhanced the user experience, while Ahmaniemi *et al*. [1] evaluated dynamic audiotactile feedback for gestural interaction, finding that the combination of haptic and audio feedback improved the accuracy achieved by participants when attempting to perceive the difference between vibrotactile feedback textures.

These techniques could help alleviate problems encountered in our design, allowing for richer vibrotactile feedback, coupled with audio to better inform the user of its meaning. In addition, while we have presented a one-level hierarchy of information (target to sub-target), with improvements to the current basic haptic and gesture recognition technologies, we can envisage a richer structuring and navigation of content, based on the framework presented in this article.

9. 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.

10. REFERENCES

[1] T. Ahmaniemi, V. Lantz, and J. Marila. Dynamic audiotactile feedback in gesture interaction. In *MobileHCI '08: Proceedings of the 10th international conference on Human*

computer interaction with mobile devices and services, pages 339–342, New York, NY, USA, 2008. ACM.

[2] S. A. Brewster, V. Räty, and A. Kortekangas. Earcons as a method of providing navigational cues in a menu hierarchy. In *HCI '96: Proceedings of HCI on People and Computers XI*, pages 169–183, London, UK, 1996. Springer-Verlag.

[3] A. Chang and C. O'Sullivan. Audio-haptic feedback in mobile phones. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1264–1267, New York, NY, USA, 2005. ACM.

[4] M. Egenhofer. Spatial information appliances: A next generation of geographic information systems. *First Brazilian Workshop on GeoInformatics*, October 1999.

[5] 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.

[6] 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.

[7] 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.

[8] M. Jones and S. Jones. The music is the message. *interactions*, 13(4):24–27, 2006.

[9] 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.

[10] R. W. Lindeman, J. L. Sibert, E. Mendez-Mendez, S. Patil, and D. Phifer. Effectiveness of directional vibrotactile cuing on a building-clearing task. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 271–280, New York, NY, USA, 2005. ACM.

[11] 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.

[12] S. Robinson, P. Eslambolchilar, and M. Jones. Point-to-geoblog: gestures and sensors to support user generated content creation. In *MobileHCI '08: Proceedings of the 10th international conference on Human computer interaction with mobile devices and services*, pages 197–206, New York, NY, USA, 2008. ACM.

[13] E. Rukzio, K. Leichtenstern, V. Callaghan, P. Holleis, A. Schmidt, and J. Chin. An experimental comparison of physical mobile interaction techniques: Touching, pointing and scanning. In *UbiComp '06: Ubiquitous Computing*, volume Volume 4206/2006, pages 87–104. Springer Berlin / Heidelberg, 2006.

[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. Shoogle: 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.