Dude, Where's My Car? In-Situ Evaluation of a Tactile Car Finder

Martin Pielot, Wilko Heuten
OFFIS Institute for Information Technology
Oldenburg, Germany
martin.pielot@offis.de
wilko.heuten@offis.de

Stephan Zerhusen, Susanne Boll
University of Oldenburg
Oldenburg, Germany
stephan.zerhusen@googlemail.com
susanne.boll@informatik.uni-oldenburg.de

ABSTRACT
We present an in-the-large, in-situ study of a "car finder" application for mobile phones. The interface is an instance of the magic wand metaphor, i.e. the mobile phone vibrates when it points in the direction of the stored location (e.g. the car). The rationale of using tactile feedback is that visual interfaces may not be ideal when the application is used on the move, since its users may easily become distracted. To investigate if users would use tactile feedback and whether it can lower distraction and is accepted, we published the application for free on Google Play. We collected anonymous usage data between July 2011 and July 2012. We provide evidence that tactile feedback was used about half of the time and can lower the users' distraction.

Author Keywords
Handheld Devices and Mobile Computing; Tactile & Haptic UIs; Research in the Large; Mass Participation Trials

ACM Classification Keywords
H.5.2 User Interfaces: Haptic I/O

BACKGROUND AND MOTIVATION
Thanks to increasingly capable mobile phones, we have got used to querying digital information anywhere and with respect to our location. For instance, we can search for nearby restaurants on a map, use the phone as a magic lens to see if nearby houses are for rent, or keep track of the location of our parked car. Oftentimes, these so-called location-based services are designed to be used on the move (see Figure 1). Typical usage environments, such as sidewalks or pedestrian zones, can be very lively and full of obstacles. Hence, it is important to design the interaction with these applications in a way that users do not become distracted from their surroundings.

Common location-based services often rely on visual feedback, such as maps, to communicate location information. However, interacting with mobile devices on the move leads

NordiCHI '12, October 14-17, 2012 Copenhagen, Denmark

Figure 1. Location-based services, such as car finders, are popular applications for mobile devices. Since they are used on the move, users need to keep track of obstacles. Thus, user interfaces of location-based services should allow to focus on the environment.

to fragmented attention [7]. In a survey by Madden and Rainie [4], one in six (17%) cell-owning adults reported to have physically bumped into another person while being busy with their mobile phone. Thus, we need to investigate how to enable non-distracting interaction when using location-based services.

As a solution to this problem, novel interaction techniques are emerging that use gestures and non-visual feedback to allow eyes-free interaction. For example, as illustrated in Figure 2, users receive vibro-tactile feedback when they point at the location of an object with the mobile device. Fröhlich et al. [2] has suggested the term *magic wand* for such interfaces. Comparing the magic wand paradigm against several other methods, Fröhlich et al. conclude that such pointing gestures are efficient and intuitive forms of interacting with geographic targets. Several other research groups have shown that magic wand interfaces are well-suited to communicate the locations of waypoints and points of interest for navigation and orientation [5, 10, 11, 12, 15, 16]. For example, Robinson et al. [12] found that travellers can reach a destination by simply cueing its location with a vibro-tactile magic wand interface. They report that this technique was well received and argue that it may help overcome the "tyranny of turn-by-turn navigation". In our previous work [8, 10, 9], we provided evidence that this technique can significantly reduce the user's distraction.

However, all of the previous studies share some limitations: all of the results are limited to one geographic area and a rather homogeneous sample of participants. Also, the participants did not navigate out of a need but because they were asked to do so as part of the study. Hence, it is not clear whether the magic wand interaction technique will remain appreciated and beneficial in in-situ use, i.e. the use out of a necessity in daily life.

Therefore, we conducted the first in-the-large, in-situ study of a tactile magic wand interface. We implemented a simple car finder application, the *6th Sense Car Finder*[1] for Android phones that, besides the typical visual arrow, provides vibro-tactile feedback when pointing at the parked car (see Figure 2). We released it to the public via Google Play (formerly known as Android Market) and collected anonymous usage data over a period of 12 months. We found that about half of the users enabled vibration feedback. In addition, we provide evidence that the tactile magic wand can lower the level of distraction.

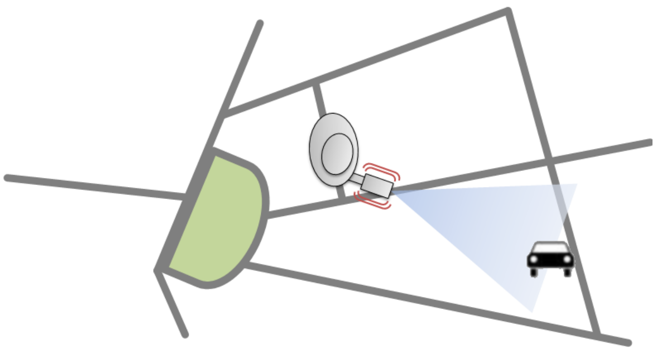

Figure 2. Tactile Magic Wand Interface: the mobile phone vibrates when pointing at the parked car.

IN-THE-LARGE EVALUATION

The study we conducted follows the novel *in-the-large* [1, 3, 6] method, i.e. an application for mobile phones serves as study apparatus and is made available to the public for free via app stores. We published a car finder application on Google Play in July 2011. We collected anonymous logging data to investigate (1) if users appreciate and use the tactile feedback, and (2) if the benefits of tactile feedback on the user's level of distraction, which were found in previous studies, are evident in in-situ usage, too.

Car Finder App

When publishing the apparatus of a study on app stores, experimenters give up control over how the application is actually used. It becomes difficult to ensure that the application is used in the way that we want to study [3]. To make it less likely that our apparatus is used in unintended ways, it is kept as simple as possible (see Figure 3). Our *6th Sense Car Finder*, the apparatus in this study, simply determines the direction and the distance of the selected location in relation to the user's position and orientation, and prominently visualises it by a big arrow. As long as the device roughly points at the specified place (indicated by green cone) it vibrates in short pulses (100ms). The pause between the pulses varies depending on the distance to the destination (100 - 500ms). We follow the guideline by Szymczak et al. [15] that "*shorter periods (with constant pulse length) [are] mapped to closer distances.*" The vibration continues if the display is turned off. Thus, users can find their car without visual distraction. To allow unbiased conclusions on how users like the tactile feedback, it is randomly enabled/disable on the very first start of the app. Additionally, it can easily be toggled from the main view.

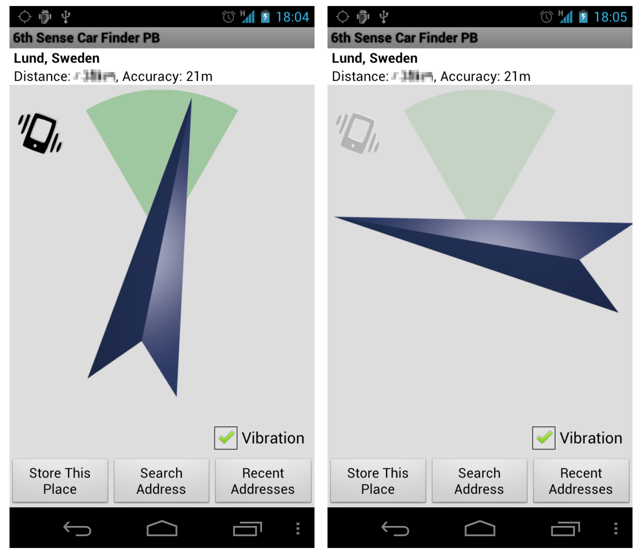

Figure 3. Main view of the Car Finder App: visual cues help the user to quickly grasp how the vibration feedback works.

Study Design

Since we aimed at investigating the effect of tactile feedback on the user's level of distraction, we set up the study as an experiment. The tactile feedback served as independent variable with two levels: disabled = control condition, enabled = experimental condition. We discarded using a between-group design, since this would have meant to force users to either always or never use tactile feedback. Since we heavily advertised the tactile feedback as unique selling point, we chose a quasi-experimental design instead. That is, the application allows users to turn on and off the tactile feedback (switch conditions) at any time.

Data Collection

To investigate how the tactile magic wand is used, we added a logging framework to the car finder. Once per second, this framework takes a snapshot of the current usage context and uploads it to our servers. Each *context snapshot* consists of a number of fields, including time stamp, the quality of the GPS signal, the movement speed, and so on. To infer the acceptance of the tactile feedback we stored in each log entry whether it was enabled or not.

To approximate the user's level of distraction, each context snapshot contains three flags: (1) whether the display was turned on or off, (2) whether the display was covered, and

[1] `https://play.google.com/store/apps/details?id=org.zerhusen.carfinder`

(3) whether the user held the device in a way that allowed to check the display. We used Android's lifecycle[2] events to check whether the display was turned on (application is *resumed*) or off (application is *paused*). To infer whether the display was covered, we used the devices' proximity sensor. In each context snapshot, we logged whether there was something covering the display. Furthermore, from previous study and a number of informal field tests, we knew that users have a strong tendency to slightly tilt the device in the direction of their face when they read content on the move. We therefore logged the tilt of the device for each context snapshot. We considered the user to be likely to check the display when the display was turned on and the device was tilted towards the face, but with a pitch of less than 23° (see Figure 4).

Figure 4. The tilt of the device is used to determine whether the user is likely to be checking the display.

To avoid ethical issues, no personally identifiable information was collected. In particular, we did not collect device IDs, GPS coordinates, or addresses.

RESULTS

The application was published on the Android Market on July 15, 2011. Until July 01, 2012, it was downloaded 5,844 times. According to the device locale most participants live in the US (36.9%) and Europe ($> 34.0\%$).

Since we were interested in usage on the move, we filtered out all context snapshots where there was no accurate GPS location available and where user was not moving. Applying this filter, we kept 136,123 context snapshots by 765 distinct users, which cover 37.8 hours of use on the move.

In average, a user enabled tactile feedback during 54 % (Mdn = 97 %, SD = 48 %) of the trip time. Due to the large standard deviation and the median being close to 100 % we printed the data for each individual user. We found that almost all users either never (0 %) or always (100 %) enabled tactile feedback. Thus, tactile feedback was used by about half of the participants.

To analyse the effect of tactile feedback on the level of distraction we divided all context snapshots into two groups: the experimental group where the tactile feedback was enabled, and the control group where the tactile feedback was turned off. For each of these groups we analysed, in how many snapshots the display was turned off, the device was covered, and the user was likely to be checking the display (in % of all context snapshot of the respective group). We used t-test to test for significant effects.

[2]`http://developer.android.com/reference/android/app/Activity.html#ActivityLifecycle`

There was a significant effect of tactile feedback on *display turned off* ($t(136,121) = -15.87, p < .001$). In the experimental group, the display was turned off more often (M = 0.41, SD = 0.49) than in the control group (M = 0.36, SD = 0.47). Thus, with tactile feedback enabled, users on the move turned off the display more often.

There was a significant effect of tactile feedback on *display covered* ($t(136,121) = -37.35, p < .001$). In the experimental group, the display covered more often (M = 0.39, SD = 0.48) than in the control group Control (M = 0.28, SD = 0.45). Thus, with tactile feedback enabled, the screen of the device was covered more often, e.g. by putting the device in the pocket.

There was a significant effect of tactile feedback on *checking the display* ($t(136,121) = -37.6, p < .001$). In the experimental group, users held the device less often in a way that allowed them check the display (M = 0.25, SD = 0.43) than in the control group (M = 0.36, SD = 0.48). Thus, with tactile feedback enabled, users were less likely to check the display.

DISCUSSION

We analysed 136,123 context snapshots of 765 distinct users that reflect the same number of seconds of the car finder's use on the move. Despite a few exceptions, about half of the users always used tactile feedback while the other half never used it. When tactile feedback was enabled, users significantly more often turned the display off, covered the display, and held the device less often in a way that easily allows checking its display.

The use of tactile feedback by about half of the users indicates that there is a general acceptance of magic wand interfaces with tactile feedback. Since tactile feedback was enabled/disabled randomly on first start there was no bias towards using it. Also, since it could be toggled from the main screen, users could easily turn the tactile feedback off if they did not want to use it. Hence, we assume that, for most of the time, tactile feedback was only used when the user desired it.

In terms of distraction, tactile feedback had significant positive effects on all three dependent variables (*display turned off*, *display covered*, and *checking the display*). This aligns with findings from previous field studies [8, 10, 9], which show that tactile feedback can reduce travellers' level of distraction. Yet, these variables are only approximations of the level of distraction. Just because the display is turned on, it does not necessarily mean that the user is looking at it and, hence, distracted. However, previous studies [8, 13, 14] unanimously report that travellers are often highly distracted by navigation systems. For example, Rukzio et al. [13] report that participants checked the mobile device's display every 5.8 seconds in average. Hence, if the display is visible to the user, there is a significant likelihood the user will be looking at it. The fewer situations where users can potentially read the display, the less likely they are to be distracted. Consequently, our results indicate that our users were less likely to be distracted with tactile feedback enabled.

In contrast to our previous work [9], where we tested the use of vibration patterns with a similar approach, this study has two advantages. First, in our previous work tactile feedback was enabled on first start-up. This did not allow an unbiased judgment on how well tactile feedback is appreciated by smartphone users. Second, the interface we tested here is much more intuitive. Even despite the bias, the vibration patterns from our previous work were used much less than the car finder's magic wand interface (23.3 % / 47 % of all context snapshots, respectively).

With over 750 users and over 38 hours of use on the move, we report from a multitude of usage data compared to previous work. However, as in all unsupervised studies, we cannot completely nail down whether our users always acted as we assume [3, 9]. Consequently, the internal validity of our findings, i.e. to what extend the observed effects can be attributed to the suspected cause, is lower than in previous work [5, 12, 16]. On the other hand, our results reflect true in-situ usage. This increases the study's external validity, i.e., the ability to generalise findings beyond the scope of the study. At the same time, the data we provide in this paper adds external validity to findings of previous lab and field studies, showing that the benefits of tactile feedback do not only occur in controlled settings but also *in the wild*.

CONCLUSIONS

In this paper we have presented results from the first in-the-wild, in-situ study of a car finder application with a vibro-tactile *magic wand* interface. By analysing anonymous log data from over 750 users, we provide evidence that tactile feedback is well appreciated in a magic wand interface, and that it has a positive effect on distraction when used on the move.

The contribution of this paper is three-fold. First, we have shown that a good share of end-users accept tactile feedback in magic wand interfaces, without providing any special introduction or training. Second, we provide evidence that the benefits of tactile feedback on the user's level of distraction are also present in in-situ use, outside of controlled user studies. Third, by comparing the results to our previous work [9], we show that magic wand interfaces are more appreciated by the average smartphone user than vibration patterns.

On the bottom line, we have shown that the concept of a tactile magic wand can easily be grasped by novice users, is well appreciated, and has positive effects that make worth its implementation. These findings can provide input for reaching a decision, when considering to implement non-visual feedback in customer-grade location-based applications.

ACKNOWLEDGMENTS

The authors are grateful to the European Commission which co-funds the IP HaptiMap (FP7-ICT-224675). We like to thank our colleagues for sharing their ideas with us.

REFERENCES

1. Cramer, H., Rost, M., Belloni, N., Bentley, F., and Chincholle, D. Research in the large. using app stores, markets, and other wide distribution channels in ubicomp research. In *Proc. UbiComp '10 - Adjunct* (2010).
2. Fröhlich, P., Oulasvirta, A., Baldauf, M., and Nurminen, A. On the move, wirelessly connected to the world. *Commun. ACM 54* (2011), 132–138.
3. Henze, N., Pielot, M., Poppinga, B., Schinke, T., and Boll, S. My app is an experiment: Experience from user studies in mobile app stores. *IJMHCI 3*, 4 (2011), 71–91.
4. Madden, M., and Rainie, L. Adults and cell phone distractions. Tech. rep., Pew Research Center, 2010.
5. Magnusson, C., Molina, M., Rassmus-Gröhn, K., and Szymczak, D. Pointing for non-visual orientation and navigation. In *Proc. NordiCHI '10* (2010).
6. McMillan, D., Morrison, A., Brown, O., Hall, M., and Chalmers, M. Further into the wild: Running worldwide trials of mobile systems. In *Pervasive* (2010).
7. Oulasvirta, A., Tamminen, S., Roto, V., and Kuorelahti, J. Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile hci. In *Proc. CHI '05* (2005).
8. Pielot, M., and Boll, S. Tactile Wayfinder: Comparison of tactile waypoint navigation with commercial pedestrian navigation systems. In *Pervasive* (2010).
9. Pielot, M., Poppinga, B., Heuten, W., and Boll, S. Studying tactile navigation systems in-situ in the large. In *In Proc. CHI '12* (2012).
10. Pielot, M., Poppinga, B., Schang, J., Heuten, W., and Boll, S. A tactile compass for eyes-free pedestrian navigation. In *Proc. INTERACT '11* (2011).
11. Robinson, S., Eslambolchilar, P., and Jones, M. Sweep-shake: finding digital resources in physical environments. In *Proc. MobileHCI '09* (2009).
12. Robinson, S., Jones, M., Eslambolchilar, P., Murray-Smith, R., and Lindborg, M. "I Did It My Way": Moving away from the tyranny of turn-by-turn pedestrian navigation. In *Proc. MobileHCI '10* (2010).
13. Rukzio, E., Müller, M., and Hardy, R. Design, implementation and evaluation of a novel public display for pedestrian navigation: the rotating compass. In *Proc. CHI '09* (2009).
14. Rümelin, S., Rukzio, E., and Hardy, R. Naviradar: a novel tactile information display for pedestrian navigation. In *Proc. UIST '11* (2011).
15. Szymczak, D., Magnusson, C., and Rassmus-Gröhn, K. Guiding tourists through haptic interaction: vibration feedback in the lund time machine. In *Proc EuroHaptics '12* (2012).
16. Williamson, J., Robinson, S., Stewart, C., Murray-Smith, R., Jones, M., and Brewster, S. Social gravity: a virtual elastic tether for casual, privacy-preserving pedestrian rendezvous. In *Proc. CHI '10* (2010).