Designing Interactive Advertisements for Public Displays

Hasibullah Sahibzada, Eva Hornecker, Florian Echtler, Patrick Tobias Fischer
Bauhaus-Universität Weimar, Bauhausstr. 11, 99423 Weimar, Germany
first.middle.lastname@uni-weimar.de

ABSTRACT
Although public displays are increasingly being deployed in everyday situations, they are still mostly used as auto-active information sources. Adding interactivity can help to attract and engage users. We report on the design and in-the-wild evaluation of an interactive advert for a public display in a tourist information center. We evaluate and compare 3 different variants – non-interactive, interaction using body tracking, and interaction using personal mobile devices – with respect to attracting the attention and interaction from passersby. We further compare these variants with an iterated version of the body tracking system with an extended tracking area. Our findings include an unexpected reluctance of passersby to use their mobile device in public, and the increased interactive area for body interaction resulting in increased engagement and spontaneous multi-user interaction, while removing the so-called 'landing effect'. Based on our findings, we suggest guidelines for interactive adverts on public displays.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords
public display; advertisements; full-body interaction; mobile devices; Kinect; in-the-wild study

INTRODUCTION
Traditional advertisements present content in a static form (posters) or as linear video, passively exposing the consumer to product information [5]. But people tend to ignore adverts and public displays [7, 21, 16]. Interactive adverts hold the potential to increase attractiveness and attention levels [15, 2]. Both body/gesture-based interaction and mobile phone-based input [21, 22, 27, 24] have been suggested as suitable modalities, but it is not clear which is most effective, in particular for a public situation, and how to optimize attraction levels and ease of use.

Our research investigates how interactive elements can increase attraction of adverts, how to increase the attention levels of passersby and how to motivate users to engage in subsequent stages of interaction, from first attention over implicit and explicit interaction to follow-up action. Our interactive advert promotes the Bauhaus walks, a guided tour by students from our university of its historical buildings in the town. The tour is currently taken by around 5000 visitors per year. The core idea of the advert is that users move across a map of the town and trigger short information snippets on top of historic sites. They have a time window in which to explore 5 sites, and then are shown a brief video providing information on when and where to find the walks. We developed 4 system variants in a user-centered design process in collaboration with the Bauhaus walk guides and deployed and evaluated these in our local tourist information center.

CHI 2017, May 06 - 11, 2017, Denver, CO, USA

ISBN 978-1-4503-4655-9/17/05…$15.00
DOI: http://dx.doi.org/10.1145/3025453.3025531

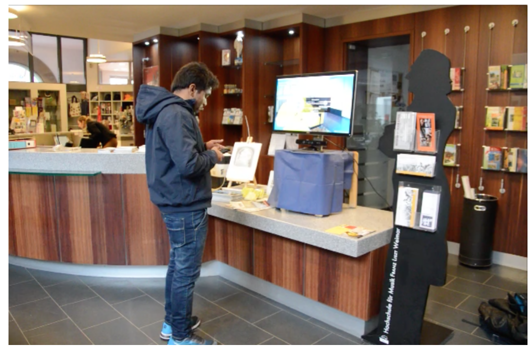

Figure 1. Hardware setup in field study at tourist information center.

We compare a non-interactive advert (a video showing the same information as the other versions) serving as baseline, with a Kinect-based version where users' body movements are translated into their silhouette moving across the map, and a mobile interaction version where users use their phone as controller to navigate across the map. In addition, we developed and deployed a second version of the body interaction system with an extended tracking angle. These 4 versions were evaluated in a field study in the Weimar tourist information office where they were deployed for 5 days each, respectively 3 days for the second body interaction version. We evaluated for commonly used measures of success of adverts and pervasive displays, such as the number of glances received, the number of people taking attention and interacting, the overall length of engagement, and how many people transitioned from each stage of engagement to the next [23, 4, 19, 24, 27].

Our study revealed that for our specific design decisions and application scenario, the mobile phone interaction mode was

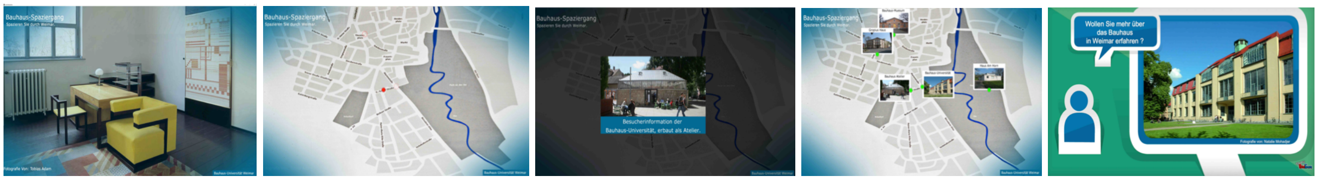

Figure 2. Screenshots of the non-interactive advert with start screen, map, pictures of sites appearing over the map and final information on the walks.

not successful, creating too large a threshold for interaction for our context and user population. We demonstrate the utility of a silhouette display in a multi-user scenario and explore how to increase its effectiveness. Extending the tracking angle of our display removed the landing effect and resulted in higher attention rates. Our observations further show that advert content needs to be tightly integrated into the interaction process in order to be taken notice of and that a match between content and interaction mode supports this.

RELATED WORK

When pervasive computing technology is used for advertisement, this is termed *pervasive advertisement* [18]. Examples include Bluescreen [26], which selects and displays ads on detecting users and allows them to choose and manipulate content by positioning themselves to the display, and the Prospero project [9]. Many auto-active digital displays are already used by the advertisement industry and increasingly, companies experiment with interactive adverts, such as games where players can win a coupon [11], or awareness campaigns [12].

A well known issue is display blindness; people tend to ignore advertisements [7, 21]. Huang et al [16] found that many factors, such as display position, eye-catching content, and size influence this along with viewers' expectations. People often intentionally ignore displays [16, 21] because of information overload, as they expect only uninteresting content. Interactive ads hold the potential to increase attractiveness, attention levels and recall [15, 2]. More research is needed on which interaction modalities are best. Direct feedback on body movement has been found effective for prompting attention and interaction [20, 24]. But more (physically demanding) movement can negatively influence recall from large interactive screens, and non-interactive screens sometimes even outperform interactive scenarios [23]. She et al [27] argue for the effectiveness of smart mobile devices, as they enable personal and discrete interaction, sharing and storing of information, and can be used from a distance. While studies of mobile interaction with public screen adverts have shown that users like this [30, 22], these tend to be proof of concept and to use an experimental study setup, even when run in 'real' locations. The question thus remains whether mobile input is effective in-the-wild.

A common measure in public display research for how much attention a display garners is the number of glances from passersby who turn their head towards the display [14, 16, 21] in relation to how many ignore it. Advertisers are interested in the conversation rate (conation step [27]), that is, how many viewers take the intended follow-up action. With internet ads, this can be automatically tracked via cookies and logfiles, but for public displays it is much harder to discern. A common approach to model effectiveness of advertisement is the *conversion funnel* [4]. This begins with viewer's attention, as higher attention benefits the recall of factual content of ads [15, 25] and is seen as first phase of engagement. Attention then leads to interest and motivation, desire (e.g. to own the product) and finally action (buy it) [13]. The funnel metaphor indicates that at each step less people continue.

In research on public displays, various models of the user interaction process can be found. Besides of Vogel's [28] zones of interaction model (ambient display, implicit, subtle and personal interaction) and Brignull and Roger's [6] interaction phase model (from peripheral over focal awareness to participation and direct interaction) the most relevant for our work is Müller et al's [19] *audience funnel* model. It relates to the *conversion funnel*, as its first phase is attention, where passersby notice the display, react to it, and become motivated to interact. The audience funnel model assumes a reactive display that draws people in via subtle indirect interactions (e.g. reacting to users' proximity to the display) before they engage in direct (e.g. touch) interaction. Once a user interacts, others might join, and finally they leave and might engage in follow-up actions. Effectiveness of public displays is measured via several factors, including how many passersby glance at the display [14, 16, 21], how many started interacting [20, 14], and for how long they engaged.

In our analysis, we build on the conversion funnel model, and adapt this to the behavior we can reliably observe, taking inspiration from the audience funnel model. This concerns indicators of interest and viewers' involvement in terms of them interacting with the advert and its content.

SYSTEM DESCRIPTION

Our system was installed in the local tourist information center. Its main components in the final iteration are a 32 inch screen, situated at eye level (~1.60 m) on a counter next to the central info desk (see figure 1), and three Kinect cameras arranged below the screen to detect and track people. Both components are connected to a Windows 10 computer and a local WiFi network used in the mobile condition (described below).

The non-interactive version of the advert (see figure 2) shows the start screen (without a call-to-action), followed by a map of Weimar with blinking dots on selected points of interest. Pictures of these locations are successively shown as miniatures on the map, followed by a short video giving information on the walks and how to participate.

Figure 3. The body interaction initial attention grabbing phase with the 'come closer to play' prompt then transitions to the map game interface with instructions while scaling down the silhouette. Users need to walk around in front of the screen to move their silhouette around the map.

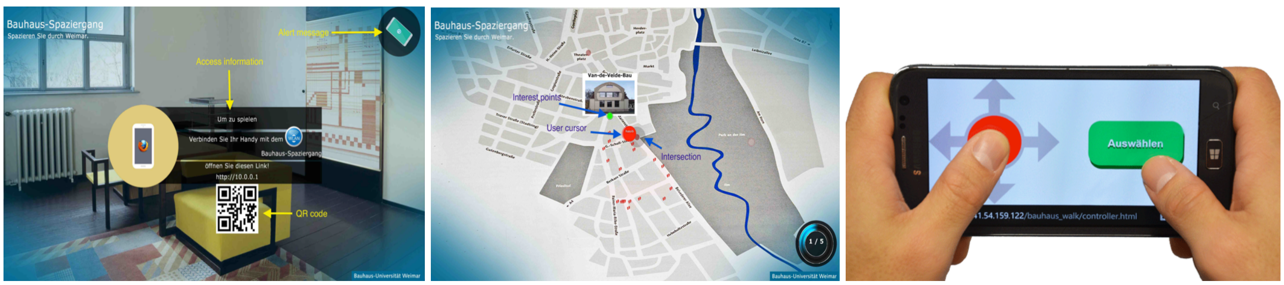

Figure 4. Initial Screen for the Mobile Interaction Modality (here, silhouette feedback is provided), Map Interface and Mobile Controller Interface.

For the interactive versions, the same pictures are revealed on the map (with name tag and short description) when their location is hit by the user. For each round of interaction, 5 locations are selected at random (as incentive for further exploration) and indicated by colored dots on the map. Once all locations are explored, the final information video is shown.

Body Interaction

The initial iteration of the body interaction variant used only a single, front-facing Kinect camera. When people pass by the screen, their silhouette is shown as an overlay on the start screen (a strategy found effective in previous work [20, 24, 30]) and the *call-to-action* prompt says 'to play come closer' (see figure 3 left). In addition, in the top right corner an alert icon appears if the user moves out of camera range. The system can detect up to 7 people and project their silhouettes, each in a different colour.

While people are far away or in the non-central area of the tracking range, this corresponds to the 'implicit interaction' phase of the audience funnel model [19]. Once they enter the central area and face the display (so they can read the call-to-action), they enter the subtle interaction phase. If they are near the screen for more than 3 seconds, the system transitions to the map interface and scales the silhouettes down to fit with the map, while fading the map in over the previous image (figure 3). Users can move around in front of the screen to navigate their silhouettes on the map, and a point of interest is activated by contact with any part of a silhouette. Interaction finishes once all five points of interest are explored or a 40 second timer runs out. To emphasize the gaming aspect, a progress icon in one corner displays this information (n/5 and a timer running down). Finally, the information video advert is shown.

Mobile Interaction

The mobile variant also uses the silhouette representation on the start screen to attract attention and then guides the user to use their phone as controller. While technologies such as Bluetooth and NFC may be utilized to attract attention, these are not yet supported by a majority of devices (or might not be switched on). We thus sought to maximize the reach of this attention-grabbing phase using the silhouette representation.

In this variant, the call-to-action instructs the user to connect to the local WiFi network and open a specific website shown as a QR code (see figure 4 left). Once the user has connected via the QR code (or typed in the IP address) and turned their phone into landscape mode as instructed, the interface transitions, representing the users' cursors with differently coloured circles on the map. The interface on the mobile is a simple HTML5 page which displays a cursor that can be moved around to navigate the map and a select button. The mobile interface uses WebSockets to ensure low latency for interaction and is based on the MMM framework by Weissker et al. [29].

Body Interaction with Wide Angle of Tracking

Based on observations from the first phase of in-situ evaluation and an initial analysis of data, an improved version of the body interaction system was developed that aimed to increase the attention and interaction levels. We had observed that people often walked past the screen due to its positioning alongside a path of travel (cf. figure 7) and were often distracted by the book desk opposite (turning their head away) or looked straight onward. Consequently, when the screen lay still in their field

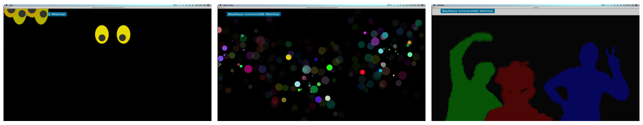

Figure 5. Following eyes, fireworks animation and silhouette representation for the initial 'attention grabbing' study.

of view, they had not yet entered the Kinect's tracking area (and thus received no visual feedback).

The main impetus of our design iteration was thus to increase the angle of view. We accomplished this by using two additional Kinect cameras facing at an angle outwards. As we relied on the standard Kinect data interpretation toolkit (Microsoft SDK), data from each camera is interpreted individually, i.e. not merged into a single tracking space. As it would have been confusing to have silhouettes switch colour when passing a tracking boundary, we used only a single colour for all silhouettes. All other interactions with the system were the same as in the initial body interaction variant.

THE SYSTEM DESIGN PROCESS

Our system design is based on an iterative user-centered design process and insights from the literature. The idea for the advertisement was developed in a co-design process with guides from the Bauhaus Walks and the system design was iterated with repeated user tests of prototypes.

Pre-Study

We compared 3 simple visualizations (figure 5) that react to people's proximity (Kinect-based tracking) to determine the best attention-grabber for our advert. Whenever a person is detected, (1) a pair of eyeballs appears on-screen and follows the person, (2) coloured animated fireworks appear at their rough location, (3) a coloured silhouette appears (similar to [20], who found this to be more effective as representation of the users' body than an abstract avatar or video image). For our study, we utilized a 14" screen installed above eye level in the entrance area of our university canteen that shows local adverts. As baseline we used a static advert shown by default.

Over four days, for 2 hours each, the visualisations were shown and one person counted glances of passersby. Only 7.6 % looked at the traditional advert (9 out of 118 passersby) whereas the silhouette attracted 15.82 % glances (from 139 passersby). The following eyes were less effective (12.98 %) followed by fireworks (10.1 %). A X^2 test revealed that the silhouette attracts significantly more glances than the non-interactive display (X^2, (1, N=257)=4.046, $p < .05$ (p=.04)). No significant difference was found for the other visualisations. This confirmed the effectiveness of a silhouette display for attracting attention compared to more abstract representations.

Focus Groups and Iterative Development

We ran two focus group sessions with Bauhaus Walk guides, one to determine the target group and develop initial ideas, and one for feedback on two alternative designs that had emerged (moving Bauhaus-design chess pieces and walking across a map). The map idea was considered more suitable for the advert content of a guided walk and easier to understand. Participants further suggested that the advert should allow for multi-user interaction (also considered important in the literature on interactive adverts [27]).

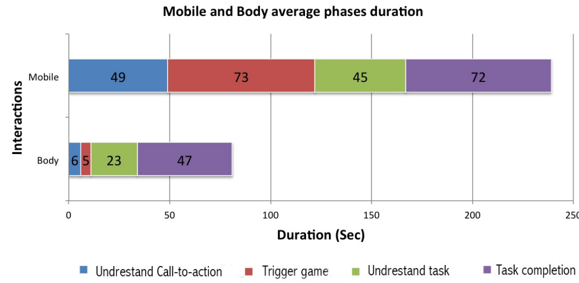

Figure 6. Comparison of body and mobile interaction duration for each phase in the high fidelity prototype evaluation study.

An early paper prototype was tested (for mobile and silhouette-body) with 5 participants in WOz manner, with a researcher moving a paper body silhouette across a display in response to participant actions. This resulted in design changes and simplifications. This test indicated difficulties for mobile interaction, e.g. understanding that one needs to scan a QR code to get started, and to turn the phone into landscape mode. It further appeared not to be as 'natural' to move around a map (on a 2nd display) by navigating with the finger across a screen.

A functional prototype of both interface versions was tested with 12 participants using a think-aloud approach. We investigated how long each phase of interaction took (understanding the 'call-to-action' screen, getting started, understanding the map navigation task, and the interaction itself) (figure 6). This user test revealed that mobile interaction overall took about 3 times as long, mostly due to understanding what to do and managing to start the game. All phases of interaction took much longer, including map navigation (on average 72 seconds in the mobile modality versus 47 with body interaction). Participants took 6 seconds to react to the call-to-action in the body interaction modality, but 49 seconds for mobile interaction. Comments from participants indicated that seeing their own body represented on the map provided a strong clue of 'walking' that made it easy to grasp the interaction method. Some participants suggested that movement interaction might be too embarrassing in public, and that mobile phone interaction would be better suited, even though most preferred the body interaction system. This test resulted in further iteration of the system design, in particular, simplifying the start phase of the mobile version.

IN-SITU EVALUATION

The final system was evaluated in-the-wild, in the Weimar tourist information center. As baseline condition, the non-interactive advert consisting of a video provided the same information as could be gained from the interactive version. The non-interactive, body interaction and mobile version were deployed for five days each in three consecutive weeks. The improved, wide-angle version of the body interaction system was deployed for three days, a few weeks later.

Figure 7. Setup in our tourist information center, with the display at the right hand side next to the information counter.

All four versions of the advert system were deployed at the exact same location and without any changes to the surroundings. Figure 7 shows the setup of our installation in the tourist center. The building entrance is just visible at the top right end of the photo. Visitors first pass by the counter, which is shaped as a half circle. Two tables with books surround the counter. Our display was placed on a desk in the adjacent space to the counter, which opens up to a larger open area (at bottom of photo). The 32 inch screen was set up at eye level (~1.60 m). Visitor levels were fairly consistent over the entire time period. From our observations, around 60 % of passersby were elderly and 15 % children.

Data Collection and Measures

In the context of our project, it was logistically not feasible to measure the actual conversion rate, which would have required questioning all participants of the Bauhaus walks in the weeks following the in-situ deployments and identifying which system condition they had been exposed to. Therefore, we limit our study to the initial phases of the conversion funnel, i.e. attention (or attraction and awareness [27]) and further engagement with the advert.

A researcher spent 2 hours of in-situ observation on each day of observation and manually counted people glancing at the display (turning head towards the display for up to 3 seconds) vs. the number of people who ignore it.

In addition, we analysed the Kinect depth data recorded over the entire deployment. Due to space and processing limitations, only a 2D colored silhouette image was stored once per second. These images were analysed alongside logfiles (which provide accurate time stamps). This data was analysed by one of the authors and three people trained to count our measures of engagement. We measured engagement as defined by attention paid (stopping in front of the display for at least 3 seconds) and overall engagement duration, based on the Kinect data. We further measured the number and duration of interactions, also based on Kinect data, as well as how many people viewed the final video advert (the first step of conation [27]). From the Kinect silhouette images we further identified incidences of the honeypot effect (where interacting people attract bystanders who later also want to interact, creating a social buzz) [6] and landing effects [20], where a person passes by the display and, having noticed the interactive response after passing, walks back to explore this.

Furthermore, a researcher took notes of passersby behaviour for 2 hours per day. In addition, we interviewed a subset of people after they interacted with the display or watched the non-interactive video version. These interviews were audio recorded and transcribed.

Due to our design iteration, there are two phases of data collection. The first phase covers one week each for the non-interactive, mobile and body-interaction prototype. This enables us to compare our measures over an identical number of hours and days across these 3 conditions. The second phase of data collection is based on 3 days of deployment of the second wide-angle body interaction prototype. Besides of using an identical setup, we were lucky for the weather to be very similar on these days.

FINDINGS

In the following, we first focus on quantitative measures related to the conversion funnel, i.e. glances, attention and interaction, and the number of honeypot effect and landing incidences. Then we discuss further findings.

We ran statistical tests based on the manual counting of passersby in the Kinect image data to verify that the number of overall passersby for each interface condition was similar, in order to compare between conditions for interface effects. In the first study phase, 1031 people passed the non-interactive prototype, 995 went past the body interaction system and 852 passed the mobile interaction version (this is much higher than the average numbers of participants for lab and field studies reported in Alt et al's survey [3]). An ANOVA revealed no significant difference (F2,5)=0.8873, p>0.05 (p=0.437). We calculated the η^2 effect size index for this to be 0.18 (18 % difference in number of passersby). For the second stage of analysis, data from three days of deployment (the same weekdays) each was considered, resulting in lower overall numbers (629 for non-interactive, 597 for body interaction, and 679 for the iterated body interaction version). An ANOVA revealed no significant difference ((F2,3)=0.1449, p=0.868).

Comparison of Glance Counts

Based on our manual in-situ count (2 hours on 5 days each), 71.17% (274 of 385 passersby) ignored the non-interactive advert display (cp. figure 8), and only 28.83% glanced at it, whereas 34% glanced at the mobile interaction version (from 237) and 41.4% glanced at the body interaction version (106 of 256 passersby). A X^2 test reveals a difference between conditions (X^2(2, N=878)=10.863, p < .05 (p=0.4376)).

Figure 8. Display Blindness at the non-interactive version.

A pair-wise comparison reveals a significant difference between the body interaction and the non-interactive version (X^2(N=641)=10.863, p < .05 (p=.00437)), but none between mobile and non-interactive or body interaction version.

We furthermore compared the number of glances at the non-interactive and the body interaction adverts with glances manually counted over three days deployment of the wide-angle body interaction version. From 225 passersby for this iterated version, 51,11% (115) glanced at the display. The number of glances thus increased by almost 10% compared to the first body interaction version This difference is statistically significant (X^2(1, N=481)=4.5413, p < .05 (p=.033086)). Moreover, the wide-angle body system increased attention in terms of glances highly significantly compared to the non-interactive version (X^2(1, N=610)=30.2247, p < .001 (p=.0)).

Comparison of Engagement and Interaction Measures

Our measures of engagement and interaction are related to the conversion funnel discussed earlier. Here, we aim to study the conversion rate between each level of engagement with the display (engage/attention, interaction, viewing video). Since the manual glance count was only done for 2 hours each day, we here rely on a different measure indicating attention being paid, that is people stopping in front of the display for at least 3 seconds. This information was extracted from the Kinect data. Further measures include overall engagement duration, the number of interactions (based on Kinect data) and duration of the different interaction phases, and how many people viewed the final video advert.

We first compare five days each of non-interactive, mobile and body interaction. Engagement for the non-interactive advert varied between 5 to 100 seconds, on average 34 seconds. Only 79 from 1031 passersby (7.66%) were categorized as engaged (standing in front of the screen for more than 3 seconds).

With the body interaction system, people spent between 10 and 200 seconds in different stages of engagement, on average 42 seconds. Some left in the middle of the interaction and some people only stared at the screen without triggering the map game, and a few repeated the interaction. On average it took 20 seconds to react to the call for action and to trigger the map. Map interaction lasted on average 18 seconds. Most users left when the advert video started, while those that remained appeared to only wait in order to play again. Among 995 passersby, 115 were categorized as engaged, that is 11.56%. Of these, 51 interacted with the system (5.1% of passerby), and 34 persisted to view the video advert (3.4% of passerby), the last stage of the conversion funnel considered here.

For the mobile interaction modality, engagement duration (22 seconds) was only half as long as for the body interaction version (42 seconds), and fared even worse than the non-interactive advert (34 seconds). People did react to the attention-grabbing start screen, seeing their silhouette on the screen, and engaged with it. Based on Kinect images, 77 from 852 passersby were identified as engaged, that is 9%. But they then did not proceed to use their mobile phone. Only two people pulled out their phone and took a photo of the QR code, but did not proceed further. Overall, none of the passersby interacted with the mobile version, and thus none saw the video advert. In terms of the audience funnel, none entered the direct interaction stage.

A one-way ANOVA indicates a significant difference in the number of engaged passersby between the three conditions ((F2,5)=11,20, p <.05 (p=.002)). A post-hoc Tukey test (critical value Q of 5.0430) shows a strongly significant difference between non-interactive and body interaction system (Q=5.6337, p=.0047509) and between body and mobile interaction version (Q=5.9467, p=.0032197), but none between mobile and non-interactive system. The η^2 effect size is 0.8, that is 80% of variance is accounted for by conditions.

We also compared tracking-based data from 3 days each of the non-interactive and the body interaction system with the iterated body interaction version with its wider angle of tracking. There were 45 engaged passersby over three days for the non-interactive system, 69 for body interaction and 104 for the improved body interaction system (of which 61 interacted with the map game and 27 viewed the final video). An ANOVA reveals a significant difference in engagement between conditions ((F2,3)=20.4154, p < 0.05 (p=0.0021)). A post-hoc Tukey's HSD test reveals no difference between non-interactive and original body interaction version (possibly because of the smaller numbers based on 3 days instead of 5), and a strong significant difference between the non-interactive version and the iterated body version (critical value Q of 6.325 and found Q of 8.9627, p<.01, p=.001744) and a significant difference between the original body version and the iterated body interaction version (critical value Q of 4.3341 and found Q of 5.3169, p<.05, p=.0218582). The revised body interaction version thus increased the number of engaged passersby. An effect size η^2 of .89 was calculated, that is 89% of variance is accounted for by conditions. Furthermore, a larger proportion (58.6%) of engaged users continued on to play the map game compared to the initial body interaction system (43x%), although the difference is not statistically significant.

Regarding the conversion funnel, we saw little overall conversion to the final step of viewing the video advert. From observation and the behaviour observable in the Kinect image data, it is clear that most people that stayed for the video did so because they wanted to interact again (and were not genuinely interested in the video). Of 51 interacting users (body interaction version), 17 ignored the advert video entirely by leaving the display, standing at one side, or turning around to talk with others until the video was over to then start over the game. The video was viewed by similar numbers of people in both body interaction conditions, 3.97% in the wide-angle setup

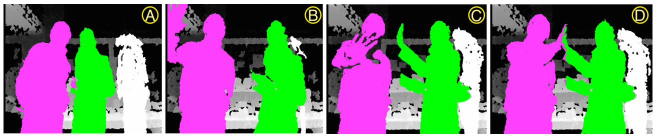

Figure 9. Playing with the silhouettes.

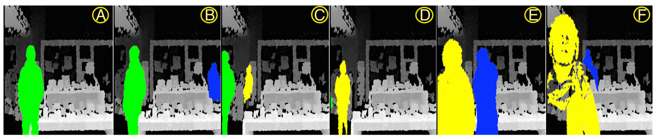

Figure 10. Honeypot effect (mobile version) – the yellow person is attracted to the display after having seen the green person's silhouette.

and 3.41% in the initial setup with a narrow tracking angle. This indicates that video content cannot be simply added at the end, but needs to be integrated into the interactive sequences.

Further Findings and Observations

Well known behaviours from the literature on public displays include the landing effect [20] and the honeypot effect [6]. In the following, we describe our findings for these behaviours as well as other interactions.

From observation, passersby behaviour around the non-interactive advert was calm and passive, with little curiosity towards the display and advert content. Most only viewed short fragments of the advert video. Very often the display was neglected (display blindness effect) even though people stood right in front of it, e.g. reading a leaflet, turning their back to it or talking with their group members (cf. figure 8).

However, this was different for the body interaction version. Here people were frequently playing with their own silhouette and engaging in various ways. Figure 9 shows two people who play with their silhouettes and the colour effect. From the observations during the field study, we conclude that the simple prompt to 'come closer to play' was well understood once users had realized the display was interactive and were close enough to read. Similarly, the metaphor of walking around a map appeared to worked well. The silhouette display triggers the landing effect as desired, and also attracts bystanders.

Out of 51 who played the game in the body interaction mode, four repeated the interactive game twice and two people thrice. Interestingly, this number was reduced for the second system version with a wide angle of tracking. Here, overall numbers increased (61), but only one person played twice. The numbers are not big enough though for statistical comparison.

Honeypot Effect

The mobile version was the least effective in triggering a honeypot effect, with only two occurrences. The non-interactive video advert resulted in seven and the body interaction version in 15 incidences. An ANOVA (($F_{2,5}$)=12.29, p=.001) and subsequent post-hoc Tukey HSD test shows that the body interaction version created significantly more honeypot incidences than the non-interactive version (Q=4.2762, p=.0264780) and strongly significantly more than the mobile version (Q=6.9488,

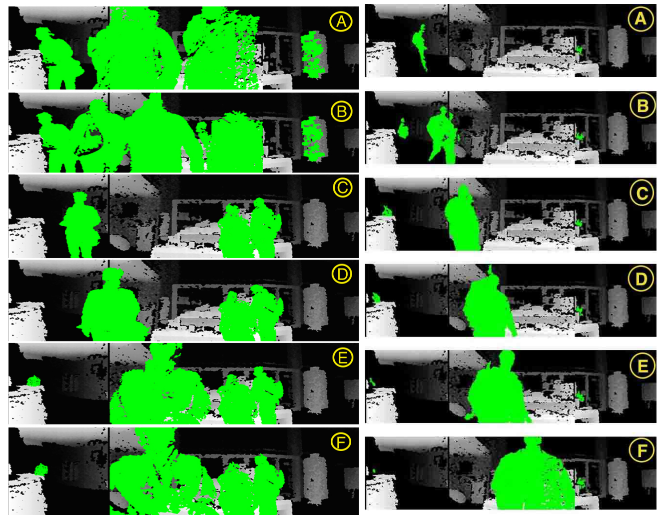

Figure 11. (Left:) Honeypot effect in the iterated body version. Two men interact for a while (frame A). A women is busy at the help desk at the left, but often looks toward the engaged pair, apparently curious. The two men leave (frame B). The woman at the left is alone and watches herself on-screen (frame C). She approaches (frame D), comes closer and starts interacting (frames E, F). (Right:) The wider tracking area means users notice the interactivity much earlier due to early visual feedback

p=.0010053). The effect size η^2 is 0.47 for non-interactive versus body interaction conditions and 0.64 for for body versus mobile interaction.

The number of incidences for three days of deployment remained identical for the original and the iterated body interaction system (10 each, 5 for the non-interactive version). An ANOVA revealed no significant difference between conditions.

Figure 11 (left) shows a honeypot effect occurring from the iterated body interaction version where two men interact and attract a woman that interacts once they have left. Figure 10 exemplifies a honeypot effect from the mobile condition.

Landing Effect

Interestingly, during analysis of the Kinect image data from the non-interactive advert (it was used for data collection), we found incidences of the landing effect, where passersby went past the screen and turned back as they reached the other side, beginning to look at the display. These were not frequent, but over five days, four incidences occurred. This was also the case for the honeypot effect (seven incidences). These effects thus are not limited to interactive or reactive systems. For the mobile version, four landing effects were found, and with body interaction, a total of twelve.

An ANOVA (($F_{2,5}$)=7.5294, $p < .05$ (p=.0076)) and subsequent post-hoc Tukey's HSD test revealed that the body interaction advert led to a significantly higher number of landing effect incidences than either the non-interactive (Q = 4.7527, p=.0144554) or mobile interaction version (Q=5.9467, p=.014454) (with a critical Q value of 3.7711). The effect size η^2 for non-interactive versus body interaction conditions was 0.66; for body versus mobile interaction, this also was 0.66.

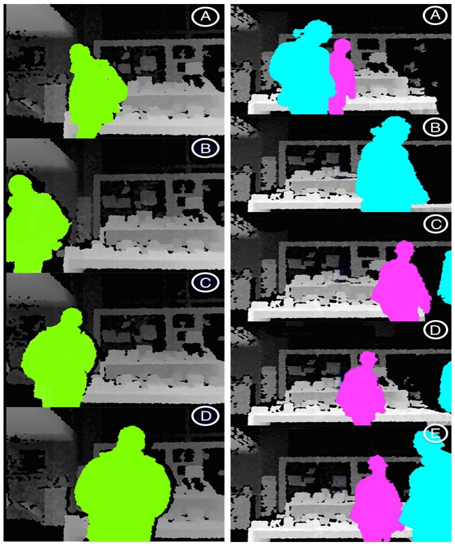

Figure 12. Landing Effect (left: non-interactive, right: body interaction)

Figure 12 shows two examples of the landing effect from the non-interactive and the initial body interaction setup (left, a person coming from the right turns back just after approaching the left end of the screen; right, two people coming from the left stop and return just as they reach the right hand side). We were initially surprised to find a much lower number of landing effects for the second version of the body interaction system with its wider angle of tracking. Over three days, there were six incidences, compared to seven with the original body interaction version. With data from three days (identical weekdays from the field study for non-interactive and original body interaction), an ANOVA was run, which revealed no significant difference.

But this should be interpreted positively, indicating that our design iteration was effective. With the revised version, people notice the interactivity much earlier, before they have passed through the middle of the interaction zone, and thus do not need to turn back. Figure 11 (right) exemplifies this, showing a passersby who notices himself as he enters the focal area of the screen (frame D, E) and then begins to interact. The landing effect thus could be considered a side-effect of a public display's spatial setup and its interaction zone size in relation to human walking speed. It takes people around 1.2 seconds to react to their silhouette when walking past it [20]. For a setup where people approach the screen from the front, there is more time and focal attention. But for a setup such as ours, where people walk past the screen, the display reaction can easily go unnoticed.

Other Observations

From our impressions during observation, users showed the most fun and joy in both body interaction conditions. Passersby were attracted quickly and showed signs of curiosity, waving their hands or moving about explicitly to learn about the interactivity. People also frequently called their friends over to jointly play with the coloured silhouettes. Figure 15 (left) shows an example of a lady calling her friend, who joins

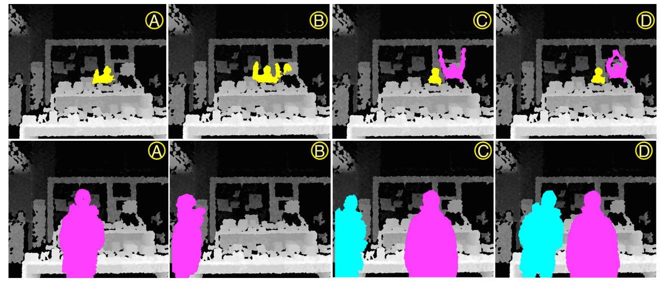

Figure 13. (Top): Noticing interactivity from afar. (Bottom): Calling each other: in frame (A) a person is engaged with the display and then moves out. In frame (C), the person calls a friend to join. In frame (D), both are engaged.

her in the iterated version of the body interaction system and figure 13 (bottom) shows a similar incidence from the initial body interaction variant.

The enhanced version with 3 Kinects covering a wider angle allowed for more people to interact simultaneously than the first version. People standing towards the side of the screen who would not come to the fore were still able to interact. Figure 15 (right) illustrates this kind of bystander interaction.

This resulted in more group interaction, with far more people (almost two thirds) interacting in groups (versus on their own) with the enhanced version (41 from 61 people over three days, in a total of 19 groups) compared to the first version where around half interacted in groups (27 from 51 interacting users over five days, in overall 11 groups). The size of groups remained similar, with most people interacting in pairs, and occasionally with three people, maximally four.

For both versions of the body interaction system, recall in subsequent interviews with people who had interacted was good, as they all remembered that the advert was for a walk or tour and related to the Bauhaus.

DISCUSSION

Our findings confirm that a silhouette display mirroring users' movements is effective in attracting people to a public display (cf. [20, 30, 24]). It was very effective for the body interaction system and also worked as attention grabber for the mobile version (but did not support further transitioning to mobile phone interaction). Moreover, in our pre-study in the canteen, from three interactive feedbacks (silhouette vs. following eyes vs. fireworks) only the silhouette received significantly more attention in terms of glances than the traditional static advert.

Different to Müller's study, our study investigates three modalities, and also an extended version of the body-tracking setup.

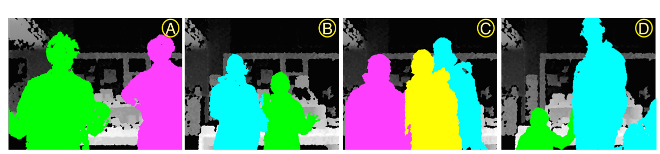

Figure 14. Group interactions with the body interaction version (4 different groups).

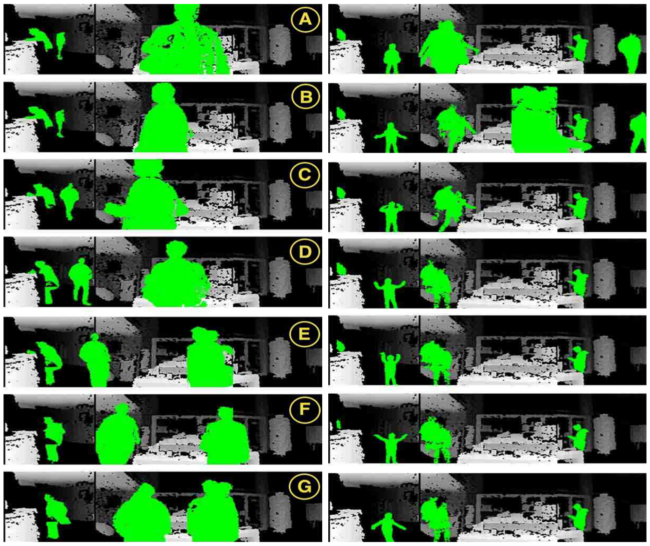

Figure 15. (Left) Calling others. A lady engages with the screen for a while (frame A). She turns and calls her friend, who is further away, looking at books (B). Her friend approaches. The lady is back at the screen (D), makes space for her friend (F), who finally also starts interacting (frame G). (Right) The enlarged tracking area increases opportunities for bystander interaction. A girl stands with her parents at the information desk. She recognizes herself on-screen, waving her hand to confirm, and starts to play with her silhouette, but remains at a distance.

Moreover, our interactive advert consists of different phases including a final video, and we study conversion rates between these phases. Previous studies (e.g. [21]) have largely focused on a younger audience due to their setting on university premises; however, given demographic developments, we consider our older user group to be a strength of our study.

While some authors argue vehemently for the effectiveness of mobile phones to interact with adverts [27] and this has been explored in experimental studies with proof-of-concept applications [22], our in-the-wild study revealed that people hesitated to use their phones. While this may be due to the usage context (cf. [1]) it provides evidence against She's claims [27] that mobile devices are ideal for the attraction stage and provide conveniences of personal and discrete interaction. In addition, if passersby are unwilling to pull out their phones in the first place, then other means of mobile interaction (e.g. iBeacons, QR codes or NFC) will also be ineffective.

Our mobile variant also reveals a dilemma for system designers - while a reactive silhouette display might be most effective for attracting initial attention, it did not motivate users sufficiently to then switch interaction mode to mobile interaction (which might improve retention due to being less physically demanding (cf. [23])). The prompt to pull out a phone might clash with expectations. It remains an open question whether other forms of inciting people to pull out their mobile may be more successful (e.g. a non-verbal prompt by a mirrored silhouette of the user that morphs into a figure holding a mobile).

In comparison to other in-the-wild studies our results for the body interaction system are encouraging. Parra et al [24] had 10.84% reacting to their mirror image in a busy train station, only at most 3% interacting (5.1% in our case), and only a small fraction watched a final video.

Isssues with Mobile Interaction

We found that our implementation of mobile phone interaction had too high a threshold, especially for an older audience beyond teenagers and university students. We interviewed some people that had been looking at the display after they left the area. This revealed that many consider their phone to be private and do not like to use it in public, and in general were sceptical of using it in this context. Some did not own a smartphone. This is surely influenced by the user group in our setting, with primarily middle-aged and older visitors attending the tourist office. Uptake levels in other settings with younger visitors might be higher. Our pre-study of the hi-fi prototype showed that getting started (and understanding what to do) takes far longer and is more difficult to convey than with the body interaction variant. Even though we acknowledge that the mobile interaction process could be improved (no dedicated WiFi network, no prompt for a name), the fact that only 2 people went as far as pulling out their phones and tried to scan the QR tag shows that the initial hurdle was too big.

Effects of Wide-Angle Body Interaction

The increase of attention for the revised body interaction advert can be traced to several factors. With wide angle tracking, people can see themselves on the screen when approaching from all sides and have a chance to do so when walking past the entire area around the screen. This increases exposure time to the reactive image, in which people's attention is caught and they can begin to understand how to interact (cf. [24]). Other studies found that it takes around 1.2 seconds to understand interactivity with a silhouette display on a large screen [20].

This, interestingly, removes the landing effect. It only occurs when the time span for passing along the interactive zone is just within the range of recognition time (1.2 seconds) so that people still have time to react, but are already past the central zone. Thus the landing effect is frequently an artifact of the size of the interactive zone and of visual flow/path.

Contrary to expectations, the wide-angle body version did not result in more honeypot effects. This is possibly because people in a tourist center do not stay around for long and thus the number of people who can get involved in a honeypot is limited. Furthermore, people may be hesitant to approach when strangers are already interacting with the system, and may prefer to wait until the system is free.

Design Recommendations

From our study, a number of design recommendations emerge. It should be noted that these are based on one study in a particular context (cf. effects of space and place [1] and audience). Limitations of our study are discussed in a following section.

Positioning and path matter. Where people don't walk towards the display but tangential to it and are focused on what is in front of them, they might not notice system feedback as easily.

Extending the interactive area resulted in prolonged exposure time to the system. This increased noticing by 10% and resulted in an increase in engagement and interaction.

Mirrored silhouette feedback works well: Our study confirms findings from previous work [20] on the effectiveness of a silhouette display as visual feedback to attract passersby attention and interest, and provoking playful interaction.

Multi-user interaction. She et al [27] mention multi-user interaction as research challenge for interactive adverts. In our case, simultaneous interaction worked well, people even explored the map together on occasion and often called a partner to come closer and share the interaction. Sometimes bystanders would interact from a distance. We recommend to enable people to collaborate in a simple, non-conflicting way.

Avoid media breaks and interaction modality switches. The video advert at the end of the game constitutes a media break, and wasn't successful. The switch from an attention grabber via silhouette feedback to phone input in the mobile condition did not work, and rather appeared to confuse people. If such a switch is inevitable, it needs to be more fluid in order to influence expectations.

Integrate advert content into the interaction. We found that the final video broke off engagement (similar to [24]). People left or just waited for the end of the video to replay the game, mostly ignoring the video. This means advert content should be integrated into the interaction, or the video needs to be very attractive and provide surprises or some kind of gratification (e.g. different video depending on how well you did, integrating your silhouette into the ad etc.) .

Content and interaction modality fit. The content greatly influences usage of public displays [23]. We received positive feedback in our early evaluations of prototypes on bodily interaction with the map, the interaction (moving about) being a direct match, without much metaphorical translation, to the task of exploring a map. Moreover, the advert content of a guided walk was considered a good fit. Such direct translation will not always be as evident, but this should be considered when generating and deciding between design alternatives.

Mobile phone input needs to be extremely low threshold and only works in contexts where people are apt to use their mobile: Mobile interaction was revealed as problematic, especially at a site where people do not have extended time on their hands, and populated mostly by an older age group that is not comfortable and/or familiar with QR codes. This implies that phone interaction needs to be as instantaneous as body interaction, to utilize apps that people have already installed as part of longer-term brand engagement, with a long-term narrative, or to provide added benefit (e.g. storing information, immediate purchases, cp. [27, 22]). It might require situations/locations with a younger audience, and where people people spend considerably more time, might be bored, and might take their phones out to kill time.

Limitations of Our Study

Many factors are known to influence attention and user engagement, from display location and size [16], orientation [17] to the environmental context [8, 10, 1]. A display in a cafe or train station will have different outcomes than one in a library or workplace. In our case, we tested effectiveness for a tourist information center. Furthermore, we only tested one design for an interactive advert.

Counting of glances was done manually in observation and might have missed some incidences. For the analysis of depth-images of passersby we were unable to determine whether one person might have passed the display twice, creating an error margin for our calculation of numbers of passersby.

The type of engagement in the three conditions is different. In the non-interactive mode, engagement means the user watches the advert, but does not need to do anything. In the mobile mode, engagement encompasses the initial silhouette feedback phase and the actual mobile phone use for the map game – but people did not go beyond the first phase of interacting with the silhouette. In the body-interaction mode, people did interact with the map game. An open question is whether a different attraction mechanism for mobile interaction might result in improved levels of engagements. The iterated wide-angle version of the body interaction system did not utilize multiple colours of silhouettes. This may have reduced the attractiveness of the display and resulted in less group play. It is likely that a multi-colour silhouette display in the wide-angle body version would be even more successful. This system was furthermore only tested for three days, resulting in less statistical power.

CONCLUSION

We presented a multi-part study on interactive advertisements. Our initial analysis confirmed previous findings that a silhouette is effective in attracting users. Based on this data, we designed an interactive advert experience with focus groups, Wizard-of-Oz and high-fidelity prototypes. Several variants of this interactive system were deployed in a public location. Our observations allow us to conclude that in our scenario, a simple prompt to 'come closer to play' in conjunction with the users' silhouette was well understood, as was the metaphor of walking around a map. On the other hand, having initially engaged people switch to use their mobile phone to play was not well received. We found that although prior research [23] indicates that movement-based interaction can reduce recall compared to other interaction modes that require less physical effort and some researchers [27] argue that it has severe limitations compared to mobile interaction with adverts, it received the most attention and active engagement, constituting a trade-off that designers of interactive adverts need to be aware of. As the silhouette representation attracts bystanders, the final iteration of our setup used an increased tracking area to attract users even earlier when passing by the display. This iteration was successful in increasing engagement levels further and attracting more group interaction, but interestingly, due to enhanced exposure time, removed the landing effect. Finally, we provided a number of design recommendations for the design of interactive advertisements.

ACKNOWLEDGMENTS

We thank all volunteers who helped with the design process as well as video and data analysis, all visitors who participated in the study and the very helpful tourist center staff.

REFERENCES

1. Imeh Akpan, Paul Marshall, Jon Bird, and Daniel Harrison. 2013. Exploring the Effects of Space and Place on Engagement with an Interactive Installation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2213–2222. DOI: http://dx.doi.org/10.1145/2470654.2481306
2. Florian Alt, Stefan Schneegass, Michael Girgis, and Albrecht Schmidt. 2013. Cognitive Effects of Interactive Public Display Applications. In *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays (PerDis '13)*. ACM, New York, NY, USA, 13–18. DOI: http://dx.doi.org/10.1145/2491568.2491572
3. Florian Alt, Stefan Schneegaß, Albrecht Schmidt, Jörg Müller, and Nemanja Memarovic. 2012. How to Evaluate Public Displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays (PerDis '12)*. ACM, New York, NY, USA, Article 17, 6 pages. DOI: http://dx.doi.org/10.1145/2307798.2307815
4. Abraham Bagherjeiran, Andrew O. Hatch, and Adwait Ratnaparkhi. 2010. Ranking for the Conversion Funnel. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '10)*. ACM, New York, NY, USA, 146–153. DOI: http://dx.doi.org/10.1145/1835449.1835476
5. A. Bezjian-Avery, B. Calder, and D. Iacobucci. 1998. New media interactive advertising vs. traditional advertising. *Journal of Advertising Research* 38, 4 (1998), 23–32. https://www.scopus.com/inward/record.uri?eid=2-s2.0-0002101166&partnerID=40&md5=ad6a4696560529a4e9f71f460bf7ede8
6. Harry Brignull and Yvonne Rogers. 2003. Enticing People to Interact with Large Public Displays in Public Spaces. In *Human-Computer Interaction INTERACT '03: IFIP TC13 International Conference on Human-Computer Interaction, 1st-5th September 2003, Zurich, Switzerland*, Matthias Rauterberg, Marino Menozzi, and Janet Wesson (Eds.). IOS Press.
7. Moira Burke, Anthony Hornof, Erik Nilsen, and Nicholas Gorman. 2005. High-cost Banner Blindness: Ads Increase Perceived Workload, Hinder Visual Search, and Are Forgotten. *ACM Trans. Comput.-Hum. Interact.* 12, 4 (Dec. 2005), 423–445. DOI: http://dx.doi.org/10.1145/1121112.1121116
8. Elizabeth F. Churchill, Les Nelson, and Laurent Denoue. 2003. *Multimedia Fliers: Information Sharing With Digital Community Bulletin Boards*. Springer Netherlands, Dordrecht, 97–117. DOI: http://dx.doi.org/10.1007/978-94-017-0115-0_6
9. Ben Congleton. 2007. Prospero - A 'Visual Commons' Framework for Community-Aware Public Displays. (2007).
10. Nicholas S. Dalton, Emily Collins, and Paul Marshall. 2015. Display Blindness?: Looking Again at the Visibility of Situated Displays Using Eye-tracking. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3889–3898. DOI: http://dx.doi.org/10.1145/2702123.2702150
11. ddbsthlm. 2011. Mc Donalds Pick 'n' Play. Video. (30 May 2011). Retrieved September 8th, 2016 from https://www.youtube.com/watch?v=7u0ij9D5S4Y.
12. Brendan Fitzgibbons. 2014. Powerful billboard lets you stop child abuse. (2 December 2014). Retrieved September 8th, 2016 from https://www.dramafever.com/news/powerful-billboard-lets-you-stop-child-abuse-/.
13. Anthony G Greenwald and Clark Leavitt. 1984. Audience Involvement in Advertising: Four Levels. *Journal of Consumer Research* 11, 1 (1984), 581–92. http://EconPapers.repec.org/RePEc:oup:jconrs:v:11:y:1984:i:1:p:581-92
14. John Hardy, Enrico Rukzio, and Nigel Davies. 2011. Real World Responses to Interactive Gesture Based Public Displays. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia (MUM '11)*. ACM, New York, NY, USA, 33–39. DOI: http://dx.doi.org/10.1145/2107596.2107600
15. Robert Heath, David Brandt, and Agnes Nairn. 2006. Brand relationships:strengthened by emotion, weakened by attention. *Journal of Advertising Research* 46, 4 (December 2006), 410–419. http://opus.bath.ac.uk/11844/
16. Elaine M. Huang, Anna Koster, and Jan Borchers. 2008. *Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays?* Springer Berlin Heidelberg, Berlin, Heidelberg, 228–243. DOI:http://dx.doi.org/10.1007/978-3-540-79576-6_14
17. Junko Ichino, Kazuo Isoda, Tetsuya Ueda, and Reimi Satoh. 2016. Effects of the Display Angle on Social Behaviors of the People Around the Display: A Field Study at a Museum. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 26–37. DOI: http://dx.doi.org/10.1145/2818048.2819938
18. Jörg Müller, Florian Alt, and Daniel Michelis. 2011. Pervasive advertising. In *Pervasive Advertising*. Springer London, 1–29.
19. Jörg Müller, Florian Alt, Daniel Michelis, and Albrecht Schmidt. 2010. Requirements and Design Space for Interactive Public Displays. In *Proceedings of the 18th ACM International Conference on Multimedia (MM '10)*. ACM, New York, NY, USA, 1285–1294. DOI: http://dx.doi.org/10.1145/1873951.1874203
20. Jörg Müller, Robert Walter, Gilles Bailly, Michael Nischt, and Florian Alt. 2012. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In *Proceedings of the SIGCHI Conference on Human Factors in*

Computing Systems (CHI '12). ACM, New York, NY, USA, 297–306. DOI: http://dx.doi.org/10.1145/2207676.2207718

21. Jörg Müller, Dennis Wilmsmann, Juliane Exeler, Markus Buzeck, Albrecht Schmidt, Tim Jay, and Antonio Krüger. 2009. Display Blindness: The Effect of Expectations on Attention Towards Digital Signage. In *Proceedings of the 7th International Conference on Pervasive Computing (Pervasive '09)*. Springer-Verlag, Berlin, Heidelberg, 1–8. DOI:http://dx.doi.org/10.1007/978-3-642-01516-8_1
22. Masafumi Muta, Soh Masuko, Keiji Shinzato, and Adiyan Mujibiya. 2015. Interactive Study of WallSHOP: Multiuser Connectivity Between Public Digital Advertising and Private Devices for Personalized Shopping. In *Proceedings of the 4th International Symposium on Pervasive Displays (PerDis '15)*. ACM, New York, NY, USA, 187–193. DOI: http://dx.doi.org/10.1145/2757710.2757732
23. Philipp Panhey, Tanja Döring, Stefan Schneegass, Dirk Wenig, and Florian Alt. 2015. What People Really Remember: Understanding Cognitive Effects When Interacting with Large Displays. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, New York, NY, USA, 103–106. DOI:http://dx.doi.org/10.1145/2817721.2817732
24. Gonzalo Parra, Joris Klerkx, and Erik Duval. 2014. Understanding Engagement with Interactive Public Displays: An Awareness Campaign in the Wild. In *Proceedings of The International Symposium on Pervasive Displays (PerDis '14)*. ACM, New York, NY, USA, Article 180, 6 pages. DOI: http://dx.doi.org/10.1145/2611009.2611020
25. Josh Schwartz. 2013. How Engaged Time Affects Brand Recall. (20 May 2013). Retrieved September 12th, 2016 from http://blog.chartbeat.com/2013/05/20/how-engaged-time-affects-brand-recall/.
26. Matthew Sharifi, Terry Payne, and Esther David. 2006. Public Display Advertising Based on Bluetooth Device Presence. In *Mobile Interaction with the Real World (MIRW 2006) in conjunction with the 8th International Conference on Human Computer Interaction with Mobile Devices and Services*. Espoo, Finland. http://www.hcilab.org/events/mirw2006/pdf/mirw2006_sharifi.pdf
27. James She, Jon Crowcroft, Hao Fu, and Flora Li. 2014. Convergence of Interactive Displays with Smart Mobile Devices for Effective Advertising: A Survey. *ACM Trans. Multimedia Comput. Commun. Appl.* 10, 2, Article 17 (Feb. 2014), 16 pages. DOI: http://dx.doi.org/10.1145/2557450
28. Daniel Vogel and Ravin Balakrishnan. 2004. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (UIST '04)*. ACM, New York, NY, USA, 137–146. DOI: http://dx.doi.org/10.1145/1029632.1029656
29. Tim Weissker, Andreas Berst, Johannes Hartmann, and Florian Echtler. 2016. The Massive Mobile Multiuser Framework: Enabling Ad-hoc Realtime Interaction on Public Displays with Mobile Devices. In *Proceedings of the 5th ACM International Symposium on Pervasive Displays (PerDis '16)*. ACM, New York, NY, USA, 168–174. DOI: http://dx.doi.org/10.1145/2914920.2915004
30. Matthias Wölfel and Luigi Bucchino. 2014. Increasing Customers' Attention Using Implicit and Explicit Interaction in Urban Advertisement. In *Proceedings of the 16th International Conference on Multimodal Interaction (ICMI '14)*. ACM, New York, NY, USA, 204–207. DOI: http://dx.doi.org/10.1145/2663204.2663268