The Audience Funnel: Observations of Gesture based interaction with multiple large displays in a City Center

Daniel Michelis[1], Jörg Müller[2]

[1]Anhalt University of Applied Science, Germany
d.michelis@wi.hs-anhalt.de

[2]Deutsche Telekom Laboratories, TU Berlin, Germany
hans-joerg.mueller@telekom.de

Abstract. We present data from observations of Magical Mirrors, a set of four large public displays with gesture based interaction installed in downtown Berlin, Germany. The displays show a mirror image of the environment in front of them and react with optical effects to the gestures of the audience. Observations of audience behaviour revealed recurring behavioural patterns, like glancing at a first display while passing it, moving the arms to cause some effects, then directly approaching one of the following displays and positioning oneself in the center of a display. This was often followed by positioning oneself in the center of the other displays in order to investigate the different effects. From these observations we deduced a framework of interaction with gesture-based public display systems. It describes the phases of passing by a display, viewing & reacting, subtle interaction, direct interaction, multiple interactions and follow-up actions. We collected quantitative data of these behavioural phases by observing 660 passers-by on 2 weekend evenings and show how many passers-by pass the thresholds between these phases. This 'Audience Funnel' should provide a framework to encourange systematic investigation of public display systems and enable comparability between different studies.

1 Introduction

Large-size electronic displays have been visible in public areas for a number of years. As a consequence of their ever-increasing presence, these displays have made dynamic digital media a more and more common feature of public space. Advertising, informational, and entertainment content has already established itself in the public realm, now new content and types of application can be seen cropping up everywhere. The development of new and cheap display technologies that are available as foils, like organic light-emmiting diodes (OLED) or electronic paper, may accelerate this development until finally a 'display wallpaper' may cover much of public space.

After a long phase of only passive display technology, interaction possibilities are now entering these display ecologies. (Multi)Touch technologies are now widely investigated and used in industrial and research settings, and also gesture based interaction gains more attention. From the usage setting, such public display technology is very different from traditional computing technologies. While in a PC setting, a user usually uses a computer in an office for a prolonged time in order to achieve a specific task, usage of public displays is much more opportunistic. Social context and motivation for technology use become the most important factors, and the concept of a task is often not applicable at all. Therefore, it is very important to study the use of such technologies in everyday social settings with passers-by, who have their own goals and context. From these observations, models of audience behaviour can be built that allow to design, implement and evaluate public display technologies. In this paper, we deduce a framework of audience engagement with public displays,

which describes the interaction process as a series of different phases. Usually, but not necessarily, users need to enter one phase before they can proceed to the next one. The boundaries between these phases can be described as thresholds, which need to be passed by the user to interact more closely with the display, or at which the interaction process is aborted.

This paper goes beyond previous studies in two important ways. First, to our knowledge it is the first systematic observation of audience behaviour towards both gesture-based public displays (without dedicated controllers) and multi-display installations in a public setting.

Second, based on these observations, we provide a quantifiable model of audience behaviour that enables the comparision of the performance of different public display systems.

2 Related Work

Frameworks of audience behaviour towards public displays can be categorized into ad-hoc models and models based on observations. Ad-hoc models are designs of public displays that react to their audience, e.g. their distance from the display. They define different audience situations that can be sensed by the display and where the display would show different content depending on the measured situation. In contrast, observation based models are derived from data of actual use of a public display system in a real public setting. A system is designed and deployed in a pubic space, and the audience behaviour towards the display is observed. The behaviour is

then categorized and used to deduce a framework. This framework can then be used to predict, describe and evaluate audience behavior towards these displays.

Ad hoc models: The Hello.Wall

Hello.Wall (2) is an ambient display with a relatively low number of large pixels. It shows light patterns depending on the context. It is accompanied by a hand held device, the View Port, which can be detected by the Hello.Wall via RFID and Wlan. With nobody near the display, the Hello.Wall shows an abstract pattern encoding e.g. the number of people in the building. When a user passes by the display, it shows secret codes notifying the user of certain events. The user can then hold the ViewPort over the wall to obtain more detailed information. Thus, the Hello.Wall proposes a division of space into ambient, notification, and interactive zones.

When nobody is in the range of the sensors of Hello.Wall, the display is in *ambient* mode. It then shows general information that is defined to be shown independent of the presence of a particular person.

In the *notification* zone, the user is identified by Hello.Wall by his hand held device. The display content reacts to user presence and encourages the user to more closely approach the display's surface and to begin interaction with the content.

As the viewer enters the *interaction* zone, he can interact with Hello.Wall to obtain even personal information. Therefore, he needs to hold his View Port over one of the pixels of Hello.Wall, where then personal information is shown on the display of the View Port.

Ad hoc models: The interactive public ambient display

The interactive public ambient display (3) uses a touch overlay and a VICON motion tracking system to detect the audience's body posture and enable explicit and implicit interaction. The system adapts to the user in four interaction phases. When no user is nearby, the display in ambient mode enables the user to get a general overview of the information at a quick glance. When a user passes by, the implicit interaction mode depending on the body posture of the user shows whether some urgent notification waits for his attention. When a user approaches the display, he can interact with it using gestures in subtle interaction. When the user steps in front of the display, touching it and covering it with his body, it shows even personal information. The system was tested informally with four users.

If the user is outside of the interaction range of the display, the display is in *ambient display phase*. It forms a central context anchoring all subsequent interaction and gives the user an overview of what kind of information or interactive function the system offers.

If the user enters a certain radius around the display, the display enters the *implicit interaction phase*. It identifies his body position and orientation and infers his openness to receiving information. Depending on this, peripheral notifications are displayed that aim to draw the user closer to the display and enter the next interaction phase.

If the user starts to consciously interaction with the display, it enters the *subtle interaction phase*. More detailed descriptions of the notifications and/or the current state of the available public information are displayed.

If finally the user stands close to the display, a *personal interaction phase* is entered. The user can interact with the display via touch, and his body can help occlude the view of their personal information from others.

Observation based models: The Opinionizer

The Opinionizer consists of a big projection display with a nearby laptop placed on a table. People can type in comments on a topic of interest, which are then shown on the display. Over time, a group discussion and shared content emerges. The system has been deployed at two parties, users were videotaped. Brignull and Rogers (1) divide the process of interaction into three general phases. These phases include peripheral awareness activities, focal awareness activities, and direct interaction:

Peripheral awareness activities introduced as eating, drinking and socializing elsewhere at the party. In general, people are peripherally aware of the display's presence and do not know much about it.

Focal awareness activities are already associated with the display. People are engaging in socializing activities associated with the display, such as talking about, gesturing to and watching the screen being used. They give the display more attention and learn more about it.

Direct interaction activities concern the active engagement with the interaction system. Individuals or groups type in their opinion to the display.

In their analysis Brignull and Rogers not only look at these three kinds of activities but also at the transition zones between them.

Their analysis revealed that the transition zones between different types of activities represent a key bottleneck in public interaction behavior. They state that in particular to cross the threshold from peripheral to focal awareness activities, people need to be especially motivated. One example would be to switch from chatting to someone on the other side of the room to deciding to move within view of the display to have a better look. One effect they observed that helped users cross the threshold was the *Honey Pot effect*. Whenever some users already interacted with the displays, bystanders were much more likely to cross the thresholds to focal awareness and even direct interaction activities. There was a progressive increase in the number of people around the display, increasing a sociable 'buzz' in the area. On closer observation, people started making comments to each other about the display and whether they had already used it.

Observation based models: The CityWall

CityWall (15) is a large Multitouch display that was installed in a street in Helsinki during the summer of 2007. It showed pictures downloaded from Flickr on a time line which could be zoomed and panned, where the pictures could be resized, rotated, and moved with simple one- or twohanded gestures. The display attracted a lot of attention and in only eight days of installation, 1199 persons interacted with the

system. Usage was videotaped and analyzed, and a variety of social configurations of use were identified. The main focus of the observations was not how users approached the display, but rather the social constellations that appear when groups of people interact with the display simultaneously. We summarize the observations regarding users' approach of the display.

Although the photos on the CityWall were continuously moving, in interviews users stated that it was difficult to *notice the display* itself, let alone the fact that it was interactive. In some cases users noticed the display while seeking shelter from rain under the roof that protected the display.

The CityWall was in use 8.8% of its total uptime. However, in 19% of the investigated cases, CityWall was already in use by someone else when a new user entered the display. When somebody was already using the display, the interaction drew a lot of attention from other passers-by. This observation might be similar to the *honeypot effect* already observed in the Opinionizer study.

In many cases when the display was already in use, new users waited a few steps behind the current users for their turn. This was called a *stepwise approach* by the authors.

Because CityWall was 2.5m wide, it was often *used in parallel* by multiple users independently. Conflicts occurred for example when the photos of one user covered the photos viewed by other users.

While in the most cases, users worked independently, in some occasions they cooperated in *teamwork* in order to achieve common goals. In one case for example, two users holding beer cans used their free hands to together scale one photo.

In some cases, the users started to spontaneously follow *emerging goals*. In one case for example, they started to play soccer by building a goal out of two photos and using a third photo as a ball.

Discussion of related work

In the studies presented, we see that substantial work has already been done regarding ad-hoc and observation based models of audience behaviour. However, there seems to be a kind of disconnect between the ad-hoc frameworks and the observation based frameworks of interaction. The ad-hoc frameworks have never been tested with larger numbers of users, and the observation based frameworks have been investigated with systems that do not actually adapt to the user in different phases. The ad-hoc models focus on interaction capabilities of the display and system behavior, utilizing a relatively linear framework of user interaction. By contrast, the observation based models present more complex and non-linear interaction frameworks. More emphasis was placed on the transitions between the phases, and some additional non-obvious behaviors have been observed. For example, the honeypot effect has been observed in both the Opinionizer and CityWall studies. Additionally, unintended interaction behavior like playing soccer on the CityWall has been observed, providing examples of appropriation and emerging goals of the users.

All frameworks have in common that they describe the interaction process as a number of phases, starting with a phase where the user ignores the display and ending with close and sometimes personal interaction. A special emphasis is placed on the thresholds between the phases. Up to now, no quantitative data has been presented of how many users pass the threshold from one phase to the next.

The Magical Mirrors Installation

Magical Mirrors is an installation of four large public displays in downtown Berlin, Germany. The displays show a mirror image of the environment in front of them and apply optical effects reacting to the gestures of the audience (Figure 1).

The Magical Mirror displays are installed in a 90° angle to walking direction behind storefront windows. They are 1,20m times 1,80m in size and installed at 1m height with the camera directly underneath them (Figure 2). The foot walk is about 4m broad and during one hour on a weekend evening, on average 330 passers-by passed the displays. In addition to the Magical Mirror displays, three big projection displays at the edge of the building showed screenshots from interacting users to audiences across the adjacent junction and down the neighboring streets. The displays were active each evening after sunset from February 2006 until June 2007. The displays show a mirror image of the audience together with one of four different visual effects. The *Aura* effect shows a white aura around the boundaries of the audiences' bodies. The *Luminary* effect shows a cloud of numbers around the fastest moving image region (movement center) (e.g. a hand). The *Flexibility* effect shows a

band that follows the movement center. The *Progression* effect consists of a number of flowers which grow towards the movement center.

Observation Method

We conducted observations of the use of Magical Mirrors in two different phases. The first observation phase aimed at exploring the audience behaviour in order to identify recurring behavioural patterns and from these patterns deduce an interaction framework. The second phase aimed at generating quantitative interaction data according to the interaction framework.

The primary objective of the first observation phase was to get as close as possible to the interacting audiences without disturbing their behaviour. A series of electrical cubicles was installed on the other side of the sidewalk from the displays, at about 4m distance. One of the authors conducted observations sitting on one of these cubicles with a bottle of beer, an image that is not uncommon in the streets of Berlin. From this position he was able to do detailed observations of interaction at all four displays, without disturbing the audience.

Observations

During the first observation phase it quickly became apparent that only very few passers-by hurried past the displays without noticing them. As the displays reacted to their body movements, during an average hour all but about ten passers-by at least

glanced at the displays, often turning their heads towards the displays as they passed by.

Approach

As the displays were installed along a sidewalk, almost all passers-by passed the displays sequentially. A very common behaviour was that the passers-by passed by the first display, and noticed the motion on the display caused by their own body out of the corner of their eye. They then turned their head to look at the display, without reducing their walking speed. While turning their head to keep looking at the display, they walked on. When they reached the next display or the one after that, they often carried out conscious movements, like waving their arms and hands, to trigger the display effects. They often reduced their speed or stopped for a brief moment to interact with the display.

Interaction

Interestingly, almost all users showed the same behavior when entering close interaction with the displays. Almost all users tried to position themselves in the middle of the displays, walking back and forth until they were content with their appearance on the display. This process usually took a few seconds and distinctly marked the point when the passer-by started prolonged interaction with the display. This positioning enabled the passer-by to make best use of the screen real estate for

the effects, and was usually followed by exploring the visual effects for some dozen seconds.

Emergent Goals

After some time merely exploring the effects, many of the users seemed to develop emergent goals for their interaction. One user for example after some initial exploration tried to use the *Flexibility* effect to make the band span the entire screen from one side to the other. So he started jumping from one side of the display to the other, while the band on the display followed his movements. Another example of such emergent goals is one user who tried to control the screen using his tongue. So he moved closer to the camera, and moved his tongue to make the band follow it. One user using the *Aura* effect apparently decided to fill the whole screen with his white aura. He suddenly started gesturing wildely with his arms and jumping from side to side in front of the camera, thereby filling the screen with his white aura.

Multiple Interactions

In a significant number of cases, after exploring the interaction with one screen, users went on to also explore the different effects shown on other screens. For example, when they had passed a first screen without stopping and then started interaction at the second or third screen, they went back to the first screen and started interaction there. Many users repeated this until they had explored all four screens.

Groups

Groups often showed a very distinct behaviour from individuals. Pairs of passers-by reacted similar to individuals. They fit together on one screen and while usually one of them initiated the interaction, they often explored the interaction together at the same screen. Groups of three or more members however behaved very differently. While also with groups, usually one of the group members discovered the displays and started interaction, they did not all fit the same screen. During the first observation period, at certain times only one of the four displays was active. In these cases, the remaining group formed a circle and waited for the interacting person to finish. After some time, group pressure built up and the group urged the interacting person to join them again and go on. In other cases however, all four displays were active at the same time. In these cases, after a short time of waiting, quite often the rest of the group approached the free displays and started interacting there themselves.

Honeypot effect

Whenever there was already somebody interacting with the display, it was much more probable that somebody walking down the sidewalk would also start interacting with the displays. This effect may be similar to the honeypot effect observed in the Opinionizer and CityWall studies. Thus we observe that this effect not only exists for keyboard and multitouch interaction, but also gesture-based interactive displays in this public setting.

Direct approach

As described above, at the corner of the building at the first floor, three big projection displays were installed that showed screenshots of the interaction at the Magical Mirrors. At several occasions, we could observe passers-by who saw these displays from the other side of the junction, then crossed the street and directly started interacting with the displays. Additionally, sometimes people directly approached the displays who appeared to already know them.

The 'Audience Funnel' Framework

From our observations of audience behaviour towards Magical Mirrors, we derived a framework of audience interaction. The framework includes different (not necessarily sequential) phases, where at each transition between phases only a certain percentage of the audience can be retained. We call this framework the 'Audience Funnel'. The typical process we observed was that passers-by first glanced at the displays, then moved their arms to cause some reaction of the displays. Then they positioned themselves in the center of the displays to explore the effects. Finally, they approached the other displays and in some cases took photos of the displays and posted them on the web. We distinguish these six different phases of audience interaction and also provide a vocabulary of users in these different phases, which we hope will be useful to describe interaction public display systems in a comparable manner. The phases we identified are *passing by, viewing & reacting, subtle interaction, direct interaction, multiple interaction* and *follow-up action*.

Passing By

Everyone who happens to be present in a certain vicinity of a public display can be called a *passer-by*. The specific area depends on the concrete instance of the public display, and should involve anyone who in principle could see the display. For operationalization, this area usually must be restricted to just the people sufficiently close to the displays that they can be observed in practice. In the case of Magical Mirrors we operationalized this variable as the number of people who entered a 4m radius of the displays (effectively, everybody who was passing the displays at the same side of the street).

Viewing & Reacting

As soon as a passer-by shows any observable reaction to the displays, such as looking at it, smiling or turning his head, he can be considered a *viewer*. The mere fact of somebody shortly glancing at a public display can be very difficult to observe manually. Future eye tracking technology or camera based eye contact sensors might make this observation feasible, and some audience reactions, like craning the head while passing by or smiling, might be easily observable.

Subtle Interaction

As soon as the viewer shows any signs of movement that is intended to cause some reaction by the display, we can call him a *subtle user*. In the case of Magical Mirrors,

this often coincided with briefly pausing in front of the display or by approaching the surface of the screen. Subtle interaction occurs at several meters distance from the display, where the person engaged in the interaction does not occupy any part of the display for herself and allows for the simultaneous interaction of others. We have operationalized this variable to include all viewers who appeared to do some movements with the intention of seeing the reaction of the display. Although this operationalization may seem difficult to actually observe, in practice such actions proved to be quite obvious.

Direct Interaction

As described above, after some initial subtle interactions users usually tried to position themselves in the center of the display (see figure 4). This is a very distinct feature for Magical Mirrors that allows us to distinguish between subtle interaction and direct interaction. Such a user can be called a *direct user*. In the case of Magical Mirrors, this coincided with the user entering a relatively small area of about 1m around the displays. Once within the interaction zone, the user blocks the view of the display from others. A person interacts directly when she actively engages the display for a period of time and enters the interaction zone in front of the displays.

Multiple Interaction

Many users started to interact with the other displays after a phase of direct interaction with one display. Such a user can be called a *multiple user*. Additionally, whenever a person consciously stops the direct interaction by leaving the interaction zone, but then returns to re-engage the display, this is also considered a multiple interaction. We operationalized this variable as any user entering direct interaction with at least one other display (or the same) after having interacted with one display.

Follow-up action

As described, many users conducted follow-up actions after direct or multiple interaction. For example they took photos of themselves or their friends while interacting with the displays and uploaded these to the web.

Data Collection

The primary objective of the second observation phase was to collect detailed quantitative data on interactions with the displays based on the interaction framework deduced from the observations. Therefore one of the authors conducted the observations from a parked car at a distance of 10 to 20m to the four displays, with a clear sight to all of them. From this position, statistics of passers-by in the different interaction phases were kept. Observations were carried out for one-hour periods between 7pm and 9pm on two different weekend evenings on 28.10.2006 and 25.11.2006.

Detailed statistics were kept on the phases of passing by, subtle interaction, direct interaction and multiple interaction. From the point of observation it was not possible to reliably observe whether passers-by were not looking at the displays, and almost all passers-by obviously looked at the displays anyway, thus we decided not to measure the phase of looking or reacting for this particular study. Similarly, we found it very difficult to reliably count the number of people who took various follow-up actions, such as taking photos, and we decided not to measure this phase either.

During our observation period on two weekend evenings we could observe 660 passers-by. The results are presented in figure 14. Approximately 640-650 passers-by looked at the displays. 218 entered subtle interaction with the displays, for example by moving their hands and arms while passing by. 207 users entered direct interaction, usually after conducting subtle interaction before. 144 users conducted multiple interactions. Finally, we estimate that approximately 2-5 users took follow-up actions like taking a photo of the displays and posting it on the web. The relative percentages of how many users passed the threshold between the different interaction phases are presented in figure 11. From this data we can see that the threshold between passing by and subtle interaction is strong. About one third of passers-by seemed to make some body movements to consciously cause some reaction of the displays. That is, a large number of passers-by did not interact at all but once passers-by entered subtle interaction there are likely to proceed further along the audience funnel. From 218 users who entered subtle interaction only 9 turned away without direct interaction. Further down the audience funnel another 63 users stopped interacting so that we observed 144 passers-by entering multiple interactions.

Discussion

We compare the observations we made with the use of Magical Mirrors as well as the framework we derived from these observations with related work. Subsequently, we discuss how the Audience Funnel generalizes to other display systems that may not be gesture based or even not interactive.

Some of our observations are similar to the observations in the Opinionizer and CityWall deployments. For example, the honeypot effect we observed is very similar to the same effects in both the Opinionizer and CityWall deployments. Also the emergent goal setting and appropriation we observed in Magical Mirrors has been observed in a very similar way in the CityWall deployment.

In addition to these known effects, we could also observe some new audience behavior schemes, which are caused by the facts that the displays showed a mirror image of the audience and that multiple displays were provided.

One new observation is that passers-by tended to pass by one display before stopping at the next or the one thereafter and initiating interaction there. We can call this behavior *progressive approach*. Because passers-by tend to walk past the displays pretty quickly, multiple displays give passers-by a “landing zone” where they can slow down before they stop in front of one of the displays. While our observations could only give a first indication of how this might work, it is worth further research of how series of displays can be used to progressively capture the audience’s attention.

Another new observation is what we can call the *position in the middle* effect. When public displays show a mirror image of the audience, there seems to be a natural behavior of audiences to position themselves in the middle of the screen. This effect can be used to clearly distinguish between subtle and direct interaction. Difficulties could occur when gesture-based displays offer also (multi)touch capability. While users would on the one side want to position themselves in the middle of the display to start gesture based interaction, this position could be too far away to actually touch the display, thereby increasing the threshold from gesture-based to touch interaction.

The Audience Funnel framework stands in the line of observation-based frameworks and is different in purpose from ad-hoc frameworks presented in related work. It builds on the Opinionizer framework but focuses on quantifiability of the conversions between different phases. Therefore it describes behaviors that are in principle observable, even if some of the behaviors, like looking, are difficult to observe with current technology. The phases of the Audience Funnel are compared to the Opinionizer framework in table 1. While the phase of Peripheral Awareness in the Opinionizer framework describes people who are near the display, but engaged in other activities, we chose the term passing by to describe people who could in principle see the display but do currently not look at it. In the Opinionizer framework the phase focal awareness describes people who engage in focused activities around the display. This overlaps partially with the phase of viewing and reacting of the Audience Funnel. We believe that by focusing on readily observable behaviour it is easier to obtain quantitative data on people in the different phases. As the phase of

Subtle Interaction is specific to gesture based displays, it is not described in the Opionionizer framework. The phase of direct interaction is supposedly identical to the Opinionizer framework. The phase of Multiple Interaction is more relevant for multi display settings and therefore not described in the Opinionizer framework. Similarly, the process of people taking photos may be more relevant in public settings and was not included in the Opinionizer framework probably for this reason.

Table 1: Comparison of the phases of the Audience Funnel and Opinionizer framework.

Audience Funnel	Passing-By	Viewing & Reacting	Subtle Interaction	Direct Interaction	Multiple Interaction	Follow-Up Action
Opinionizer	Peripheral Awareness	Focal Awareness	NA	Direct Interaction	NA	NA

An overview of how the Audience Funnel generalizes to other kinds of displays is presented in figure 2. For each specific display setting the phases need to be operationalized in a different way. Some phases may not be applicable for a specific type of display, or not be measurable in a reliable way in a specific setting. We propose that they still enable the quantitative comparison between different systems. The phase of passing by may be operationalized by passers-by entering a certain area around the display that may depend on the specific setting. This phase should be applicable to any display setting. The phase of viewing & reacting may be operationalized depending on the available sensors, like eye contact sensors, or may be skipped if no reliable sensors are available. It should also be applicable to any kind of public displays. The phase of subtle interaction is specific to displays that respond to the audiences gestures and may be skipped for those who do not. The phase of direct interaction should be applicable to any interactive displays, while the

operationalization depends on the specific way of interaction supported. Multiple interaction is applicable to any interactive public display, but is most relevant when multiple displays are available. Finally, follow-up actions are always available, but the specific kind of follow-up actions one is interested in may depend on the purpose of the public display or the actual audience behaviour. For example, a public display may be intended to enable people to sign up for test drives with a new car, but actual predominant audience behaviour might be that people take videos of the display. Which follow-up actions are the most interesting to measure needs to be determined on a case-by-case basis.

Table 2: Generalizability of the Audience Funnel for other kinds of displays.

	Passing-By	Viewing & Reacting	Subtle Interaction	Direct Interaction	Multiple Interaction	Follow-Up Action
Magical Mirrors	✓	✓	✓	✓	✓	✓
(Multi)Touch Display	✓	✓	×	✓	✓	✓
Non-Interactive Display	✓	✓	×	×	×	✓

From the data for Magical Mirrors, we see that the Audience Funnel can be very useful when aiming at improving public display systems. For example, if the goal would be to increase the number of multiple interactions, it becomes obvious that there are two major ways the system can be improved. First, one could work at increasing the number of passers-by. This would probably involve installing the displays at a different location where more people pass by or changing the environment so that people take different paths, thus passing by the displays. Second, for Magical Mirrors, a major threshold is between the phases of passing-by and subtle interaction, where the conversion rate is 33%. By systematically changing the design

of the system and observing whether the conversion rate is improved, more people pass down the Audience Funnel. It seems obvious that it would be much more useful to work on one of these points than, e.g. trying to improve the conversion between subtle interaction and direct interaction, where only very few people drop out.

Conclusion

In this paper we presented data from observations of Magical Mirrors. To our knowledge this is the first public deployment of a multi-display system with gesture based interaction. We describe approach towards these displays, interaction behavior (the position in the middle effect), emergent goals, multiple interactions, group behavior, the honey-pot effect and direct approach. From these observations we derive the audience funnel framework that describes audience interaction with public displays. This framework describes the phases of passing by a display, viewing & reacting, subtle interaction, direct interaction, multiple interactions and follow-up actions. We present data from observations of 660 passers-by to quantify how many people cross the thresholds between the different phases. For the specific case of Magical Mirrors, we show that a major threshold is initiating subtle interaction, while the thresholds to direct interaction and multiple interaction are already very low. The audience funnel framework advances the field of public display research in two important ways. Regarding increasing the number of interactions for single installations, it guides the focus towards improving the thresholds where it matters most. Work on improving the system can be spent at the high thresholds, where many

users drop out, instead of working on the low thresholds, where few users drop out anyway. In addition, the audience funnel framework makes obvious the thresholds that come earlier in the funnel should be improved first, because otherwise the audience may already be dropped out before they reach the later thresholds. Second, the audience funnel provides a metric for quantitatively comparing different public display systems, in order to evaluate which approaches are most effective for lowering the different thresholds.

References

(1) Brignull, H., Rogers, Y. (2003), Enticing People to Interact with Large Public Displays in Public Spaces, INTERACT'03

(2) Streitz, N. A., Röcker, C., Prante, Th., Stenzel, R., van Alphen, D. (2003), Situated Interaction with Ambient Information: Facilitating Awareness and Communication in Ubiquitous Work Environments. In: Tenth International Conference on Human-Computer Interaction (HCI International 2003), June 22-27, 2003.

(3) Vogel, D., Balakrishnan, R. (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 (Santa Fe, NM, USA, October 24 - 27, 2004). UIST '04. ACM, New York, NY, 137-146

(4) Michelis, D. (2009), Interaktive Großbildschirme im öffentlichen Raum: Nutzungsmotive und Gestaltungsregeln, Gabler, 2009

(5) Slatta, T. Urban screens: Towards the convergence of architecture and audiovisual media, First Monday Special Issue, Nr. 4, 2006

(6) Cp. Fleisch, E., Mattern, F. (2005), Das Internet der Dinge – Ubiquitous Computing und RFID in der Praxis, Berlin.

(7) Weiser, M. (1991), The Computer for the 21st Century, Scientific American, 265(3), pp. 94-104.

(8) Adams, R., Russel, C. (2007), "Lessons from Ambient Intelligence Prototypes for Universal Access and the User Experience", in: Stephanidis, C., Pieper, M., (Ed.) ERCIM UI4ALL Ws 2006, LNCS 4397, Berlin, Heidelberg, pp. 229-243.

(9) Agamanolis, S. (2004), "Designing Displays for Human Connectedness", in: O'Hara, K., Perry, M., Churchill, E. (Ed.), Public and Situated Displays: Social and

Interactional Aspects of Shared Display Technologies (Cooperative Work, 2), Kluwer Academic Publishers, Norwell, MA, USA.

(10) Ferscha, A., Resmerita, S., Holzmann, C. (2007), "Human Computer Confluence", in: Stephanidis, C., Pieper, M., (Ed.) ERCIM UI4ALL Ws 2006, LNCS 4397, Berlin, Heidelberg, pp.14-27.

(11) Leikas, J., Stromberg, H., Ikonen, V., Suomela, R., Heinila, J. (2006), "Multi-User Mobile Applications and a Public Display: Novel Ways for Social Interaction", in: Proceedings of the Fourth Annual IEEE international Conference on Pervasive Computing and Communications, Vol. 0, IEEE Computer Society, Washington, pp. 66-70.

(12) Grasso, A., Muehlenbrock, M., Roulland, F., Snowdon, D. (2004), "Supporting Communities of Practice With Large Screen Displays", in: O'Hara, K., Perry, M., Churchill, E. (Hrsg.), Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies (Cooperative Work, 2), Kluwer Academic Publishers, Norwell, MA, USA, pp. 261-282.

(13) Cp.Kortuem, G., Kray, C. (2005), "HCI issues of dispersed public displays", in: Workshop on Distributed Display Environments at CHI 2005.

(14) Cp. Laurel, B. (2003), Design Research – Methods and Perspectives, Cambridge, Mass.

(15) Peltonen, P., Kurvinen, E., Salovaara, A., Jacucci, G., Ilmonen, T., Evans, J., Oulasvirta, A., Saarikko, P. (2008), "It's Mine, Don't Touch!: interactions at a large multi-touch display in a city centre", in: CHI '08: Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems.

(16) Cao, X., Massimi, M., Balakrishnan, R. (2008), "Flashlight Jigsaw: An Exploratory Study of an Ad-Hoc Multi-Player Game on Public Displays", in: CSCW '08:

Proceedings of the 2008 ACM Conference on Computer-Supported Collaborative Work, pp. 77-86.

Figure 1: Passers-by interacting with Magical Mirrors: The effects shown are in this order: Progression, Flexibility, Luminary, and Aura.

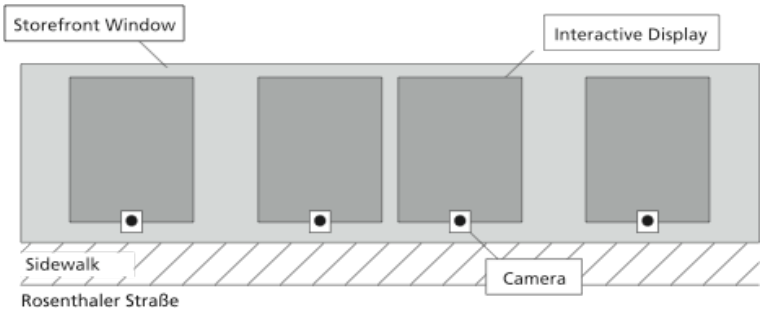

Figure 2: Experimental Set-up with four displays behind a storefront window

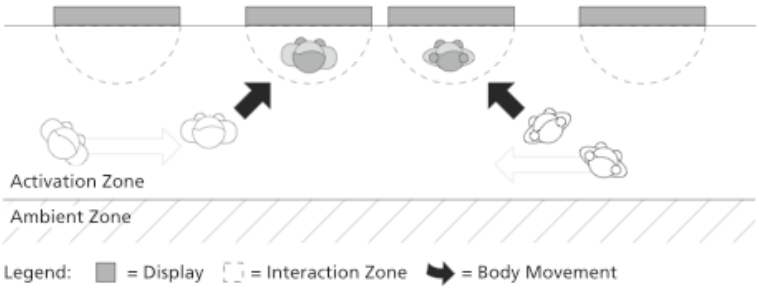

Figure 3: Observed approaching activities of passers-by

Figure 4: Users positioning oneself in the center of a display

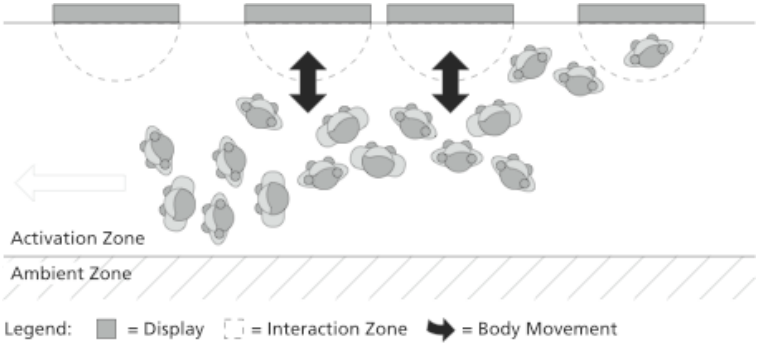

Figure 5: Group dynamics with four active displays

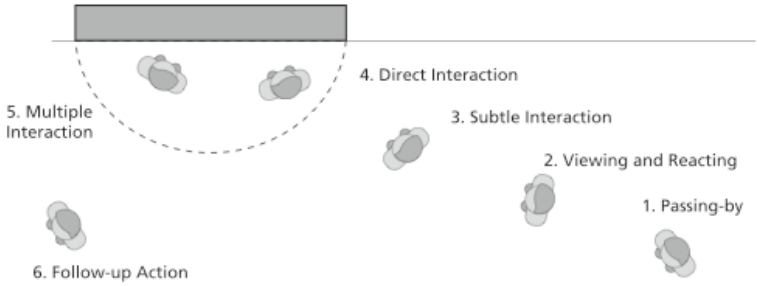

Figure 6: The Audience Funnel

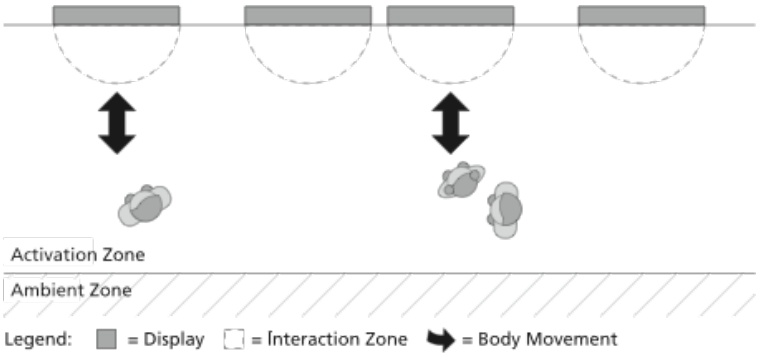

Figure 7: Subtle Interaction outside the Interactive Zone

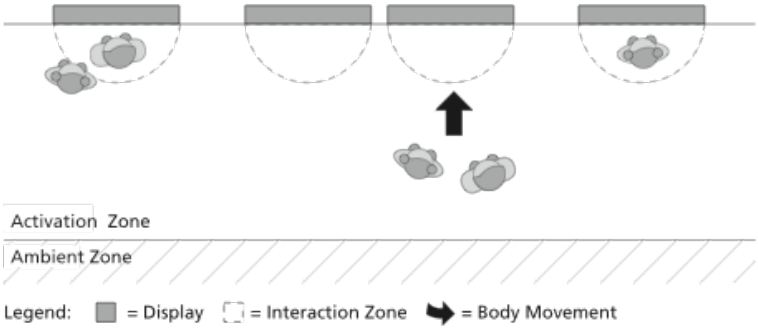

Figure 8: Direct Interaction

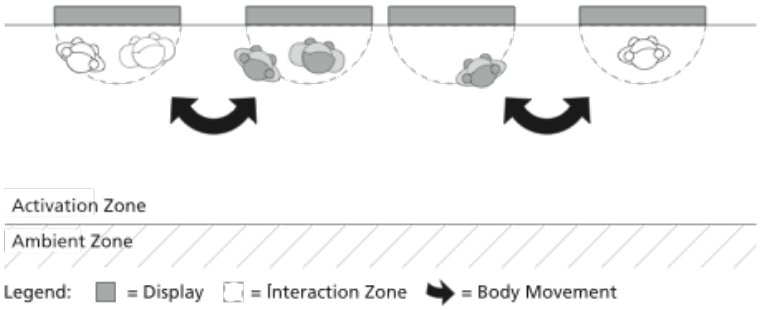

Figure 9: Multiple Interaction.

Figure 10: Taking pictures of oneself as a common follow-up action

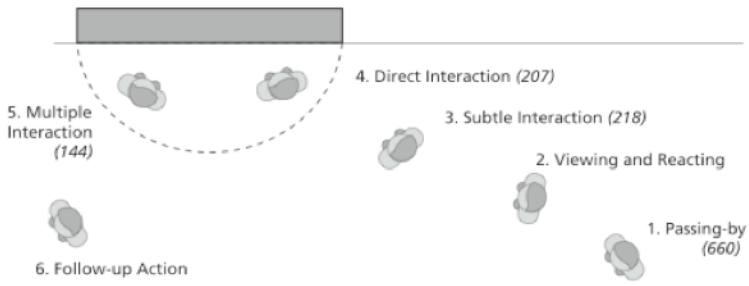

Figure 11: The Audience Funnel for Magical Mirrors

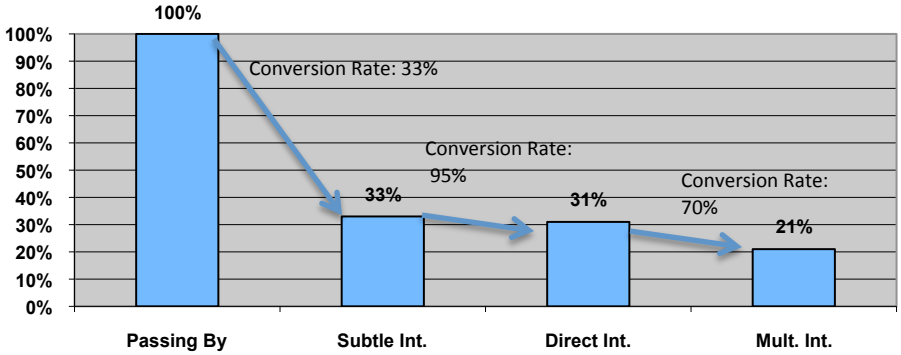

Figure 12:Percentages passing interaction thresholds for Magical Mirrors, together with the associated conversion rates.