The Puppeteer Display: Attracting and Actively Shaping the Audience with an Interactive Public Banner Display

Gilbert Beyer, Vincent Binder, Nina Jäger, Andreas Butz
University of Munich (LMU), Munich, Germany, name.surname@ifi.lmu.de

ABSTRACT
We present a wide interactive *banner display* installed at a city sidewalk and the findings from two long-term field studies investigating the opportunities of public displays to *actively shape the audience.* In order to improve parallel usage and dissolve crowds, our wide display subtly directs individual users by visual stimuli and manipulates the audience like a puppeteer, thus reversing the notion of adaptive content being implicitly manipulated by the users.

We first investigated visual signifiers which *attract initial users approaching sideways*, and then others, which *actively influence user positions* and regulate audience constellations. We found that dynamic visual stimuli such as *frames* and *ellipses* are effective (1) to direct users in front of the display, (2) to distribute multiple users along the display, (3) static frames are more effective than moving or interactive ones, and (4) these visual stimuli also work indirectly by inducing social pressure among users.

Author Keywords
Wide Displays; Public Displays; Adaptive Displays; Attention; Interactivity; Framing; Visual Signifiers.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

ACTIVE BEHAVIOR SHAPING
Very wide advertising displays are often found along urban sidewalks and passages, where people are initially passing them sideways. Their classical paper-based variant is called a *banner display* and displays information along the entire pathway. Their future counterparts, digital and *interactive banner displays*, can be installed in the same locations but will require novel interaction concepts to be effective. We identified two opportunities of such displays for *active behavior shaping*, i.e. for manipulating user behavior by visual signifiers to increase the effectiveness of the display:

DIS '14, June 21 - 25 2014, Vancouver, BC, Canada
Copyright 2014 ACM 978-1-4503-2902-6/14/06…$15.00.
http://dx.doi.org/10.1145/2598510.2598575

Figure 1. Wide banner display installed at a city sidewalk. Users need to be *attracted when approaching sideways* (top) and *optimally distributed* along the display (bottom).

First, for attracting initial attention and conveying that they are interactive, banner displays *can use their entire width*, which may make up for the disadvantage that they are not approached frontally and in full view (see Figure 1 top). In fact, attention-grabbing techniques for frontal approach will often not be transferable one-to-one to other trajectories.

Second, interactive banner displays allow *many users to interact in parallel*. This, in turn, will work much better if members of arriving groups distribute along the screen surface, and if active users will rearrange to make room for new arrivers (see Figure 1 bottom). This may not happen by itself, since social interaction with public displays can involve close bystanders [11], people interacting behind each other in several rows [12], or even conflicts between parallel users [14]. We also observed many users standing close together during our initial tests with the wide banner display, which could often be attributed to group affiliation but effectively led to active group members impeding each other and possibly also preventing passive members from taking an active part. If such crowds and clusters could be dissolved and *users distributed more evenly*, the percentage of active users may increase and groups may stay longer altogether, which is desirable for an advertising display. For convenient simultaneous interaction users should also pick *suitable positions*, where they have *sufficient space to* *interact, don't impede each other* and allow new arrivers to *find empty spots*. If such preferable conditions do not emerge spontaneously, the display may become active by itself to increase its effectiveness. We propose that this could be accomplished by the use of *visual signifiers*, with which the display subtly manipulates user positions.

We therefore investigated with a wide interactive public display (1) how visual signifiers have to be designed to *attract the initial attention* and signal that the display is interactive if users are *approaching sideways*, and (2) once people are engaging, if and how visual signifiers and dynamic strategies for displaying them can be effectively used to *moderate and actively direct* users in front of the display. We propose that this process which we call *visual audience moderation* poses the problems of *positioning*, *repositioning* and *distributing* users. If such basic goals can be accomplished by visual techniques, it shall also be possible to shape desired specific audience constellations.

To explore these questions of active behavior management, we designed and installed a wide interactive banner display at a city sidewalk and conducted two long-term field studies. While the first study showed that the key factor for initially attracting attention is interactivity of signifiers regardless of different position and movement strategies, the second study revealed that dynamically employed visual stimuli such as *frames* and *ellipses* are effective (1) to direct engaging users in front of the display, (2) to distribute multiple users along the display, and that (3) static frames are more effective than moving or interactive ones. We also noticed that (4) these visual stimuli also work indirectly by inducing some type of social pressure among users.

During the two studies with the display, we were able to observe strong effects of visual signifiers manipulating users like a *puppeteer*, thus reversing the common notion of public displays just adapting to explicit and implicit user behavior and being manipulated by the audience.

RELATED WORK

Our work builds on and extends prior art in large interactive public displays, with a focus on the following subject areas:

Social Constellations in front of Displays

From Reeves et al. [17] and Dalsgaard et al. [7] we apply the notions of user-spectator roles and their social relations such as performative interaction. For our work, we have to consider *natural attraction cues* such as the honeypot effect and *social reasons* for user positioning as described by Brignull et al. [3]. Using narrow displays, Müller et al. [12] report that users often start interacting right behind existing groups, blocking the way of other passers-by, and Michelis et al. [11] that people initially stand by the first user. With our wide display offering enough space we observed similar *by-standing effects*. Multiple displays can provide spaces for further users [11,19], yet this may imply one user per display and thus separate and static interaction.

Attracting Attention and Signaling Interactivity

Michelis et al. [11] report on the effectiveness of mirror images of the user augmented with different visual effects for attracting attention. To signal interactivity, Müller et al. [12] compare different mirror representations and further cues such as attract loops and call-to-action, Beyer et al. [5] outline the challenges of an unaware initial interaction for shaped displays, Grace et al. [9] investigate a dynamic skeletal representation of users combined with visual cues, and Akpan et al. [1] compare the influences of various places and spatial contexts in this regard. In this work, we explore how a successful initial interaction can be accomplished with a long display, with users *approaching sideways* in a typical urban sidewalk trajectory.

Adaptive Displays and Proxemics

Adaptive Displays such as that by Vogel et al. [20] tailor the content in regard to implicit and explicit cues of users, can offer shared use, and adapt to the positions, orientations or trajectories of users. Klinkhammer et al. [10] assign visually separated, adaptive personal territories for each user around a tabletop. The potentials of proxemic interactions with a large vertical display, and how they can mediate simultaneous interaction, have been illustrated by Ballendat et al. [2], and Wang et al. [21] present a display that keeps track of users' actions and encourages them to move closer, thereby regaining their attention. With our study we want to explore if the notion of displays adapting the content to users can be extended or reversed by displays also adapting and actively shaping the audience.

Very wide Interactive Displays

Several examples for very wide interactive displays in public spaces have been presented such as City Wall [14], the public information display by Grace et al. [9], and the Climate Wall [6] which allows users to grab and drag words along the façade when passing by. Screenfinity [18] rotates, translates and zooms content to enable comfortable reading on a large display while passing it. In our case, we propose a digital and interactive counterpart of classical paper-based banner displays, a scalable solution that can be easily installed at similar sidewalk locations as their ancestors.

Visual Signifiers

One type of signifiers we designed for influencing user positions are *frames*. Pinhanez et al. [15] discuss the role of frames. They projected visual elements on black backgrounds as frameless displays to contextualize them with the surroundings. Beyer et al. [4] show the influence of frames on user positions. They compared a frameless cylindrical display with the same display subdivided by large *visual frames* simulating display bezels and found that users positioned themselves in front of these frames. This supported our hope that visual frames could also be *employed dynamically* to direct users. We will extend on this work by using dynamically positioned, moving and interactive visual frames and further signifiers.

THE BANNER DISPLAY

The interactive banner display we used in our studies is installed in a street window (see Figure 1). Its format is inspired by paper banner displays found along sidewalks.

City Environment

The banner display is installed in the city of Munich in a lively shopping and nightlife district with gastronomy and hotels nearby, providing a steady flow of novice users and ensuring a broad demographic. The vast majority of passers-by approached the banner display on the sidewalk at an angle of 90°. We observed that only few people crossed the street to approach the display frontally.

Hardware and Software

The display consists of 4 frameless and horizontally aligned luminous plasma screens (MPDP). It is 3.75 meters wide and 0.52 meters high, its lower edge is at 1.25m height. While we tried to avoid discontinuities wherever possible, we could not avoid a single discontinuity caused by a window ledge of 5cm width in the middle of the display. To enable interaction for the complete interaction space, two *Microsoft Kinect* sensors were installed 15cm below the display. Sensor ranges overlapped in the center. For providing feedback to users passing sideways, user recognition based on the depth stream of the Kinects was used. For users that had stopped and interacted frontally with the display, the Skeleton stream was also used, exchanging Skeleton data between different processes via memory-mapped files. Applications were written in C#, the WPF framework, the Microsoft Kinect SDK and Emgu CV.

FIELD STUDY 1: ATTRACTING USERS SIDEWAYS

In our first study with the banner display we evaluated how passers-by could be enticed to interact by visual signifiers *when approaching sideways*, i.e., how their attention could be attracted and the interactivity of the display be conveyed at a sidewalk trajectory of 90°. As [18] confirms, the major challenge in sideways orientation of displays is to make passers-by turn their heads away from their movement direction to become aware of the screen. To stimulate an *unaware initial interaction* with our display, we therefore designed different interactive visualization strategies.

Independent Variables

For actively drawing passers'-by attention to the side by visual signifiers, we identified the following variables for displayed visual stimuli in relation to the user (see Table 1):

Signifier Variable	Values
Interactivity	*non-interactive – reactive*
Position	*parallel – ahead – full-screen*
Direction	*following – reversed – orthogonal*
Representation	*user's self – foreign being*

Table 1. Variables for visual stimuli attracting user attention.

First, a visual stimulus for attracting attention can be *non-interactive* without any relation to user behavior, or *reactive* to immediate and passive behavior such as the user's movement in front of the display. Since a stimulus *parallel* to the user at an angle of 90° may be hard to perceive, the horizontal position can be altered, displaying it *running-ahead* in front of the user, or simultaneously at any *full-screen* position. To signal the user a correlation between the stimulus and his or her movement, it may also continuously follow the user, or move at user speed in *reversed direction*, thus doubling the perceived speed when looking at the screen. The stimulus may also emerge *orthogonally* to the user. The signifier representation is strongly content-dependent involving effects of size, shape, color, animation and image schema that cannot be fully explored. Yet, stimuli can represent the users' *self* as explored by [11,12,1], or *foreign beings* that react to user's movement.

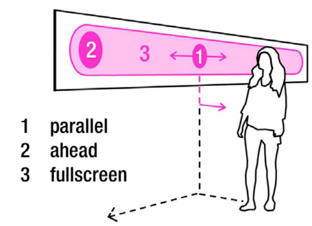

Figure 2. Positions and directions for visual signifiers that attract the attention of passers-by approaching sideways.

Visual Content

To instantiate different conditions from the variables and allow comparisons, we designed several *Underwater World* contents (see Figure 3). All stimuli were part of this visual theme with various types of fish and used constant colors and graphical styles. We chose the Underwater World as (1) it is a *non-abstract* visual metaphor users can easily understand (2) it provides a constant basis for comparisons and is *extensible* and *generalizable* to other contents, (3) the swimming fish *allow implementing the variables* interactivity, position, direction, and (4) fish *make sense* when seen ahead, full-screen or from the side. Besides, the fish had also a building-related meaning, which was not evident to the arbitrary audience: people working behind the large window were often compared to fish in a fish tank.

Study Conditions

Since people passed the display at a consistent distance of about 1m, the distance for the running-ahead stimulus was set to 1m in pretests, in order to position the corresponding stimuli at about 45°. If the user turns and changes direction, this stimulus also turns and smoothly catches up to its position. We offered no further functionality in order to discourage interaction beyond the initial reactions, as following passers-by should be attracted by the visual stimuli alone and not by the honeypot effect [3,12]. The baseline for *Interactivity* was the Underwater World just displaying animated stimuli. We also included mirror images and silhouettes as used by [11,12,1] as additional

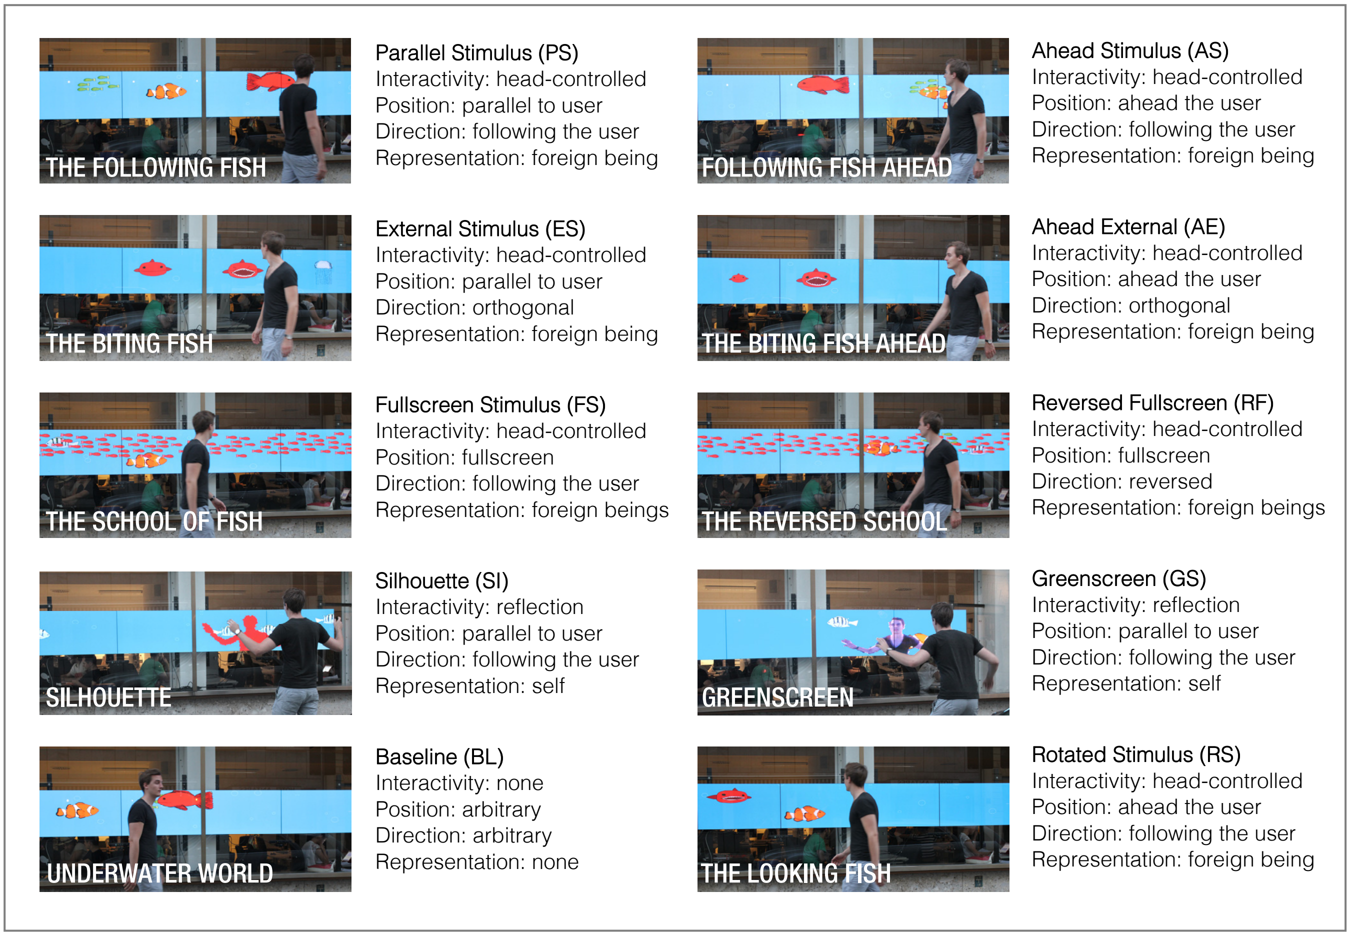

Figure 3. The single conditions of the first study instantiating the independent variables.

Representation baselines, as it remains unclear if such stimuli are also effective when perceived from the side. We further tested a rotated stimulus similar to the perspective-corrected contents of [13,18]. In pretests we discarded combinations of variables that did not prove to be practical such as different positions and directions of mirror images.

Procedure and Data Analysis

To provide optimal lighting conditions for tracking, the first study was conducted in the afternoon and evening hours over a period of 6 weeks. Subjects were just the arbitrary passers-by in the street without any pre-selection. We displayed multiple instances of the conditions, showing each for 45 minutes before choosing the next such that data would grow evenly amongst conditions. 1–2 conditions were shown per day, then the next when certain external conditions (lighting, weather, no parked bikes) were met again. Data was acquired by a field rater supervising the display from an unobtrusive location and by multi-perspective video recordings resulting in 26 hours of analyzed video material (see Figure 4). One camera is directed from the street to correlate visual effects with user reactions. In addition, Kinect position data was logged and 20 semi-structured interviews conducted, but only with users already engaging with the display, as the interviewer had to remain concealed first in order not to attract the attention of approaching interviewees. In the video analysis using *Noldus Observer*, all people passing the screen were counted as single person, pair or compact group and marked as *reacting* if people turned their head looking at the display or *ignoring* otherwise. If people stopped or started to engage with the screen, the time intervals from looking at the display until *stopping* or *aware interaction* were recorded. We then *grouped the conditions* to compare content differing in only one of the independent variables.

Results

Of the 1866 encounters passing the display, we counted 1469 single persons, 343 pairs and 54 groups. At least one person reacted to the display in 52% of the groups, 35% of the pairs and 27% of the single persons.

General Observations and Interviews

There was a significantly higher attention of pairs and groups ($\chi^2(2)$=23.514, p<0.00001) which we could attribute to other than statistical reasons: usually only one member noticed the visual stimulus first, the other members then took notice by his or her reaction or hint. Pairs were not as

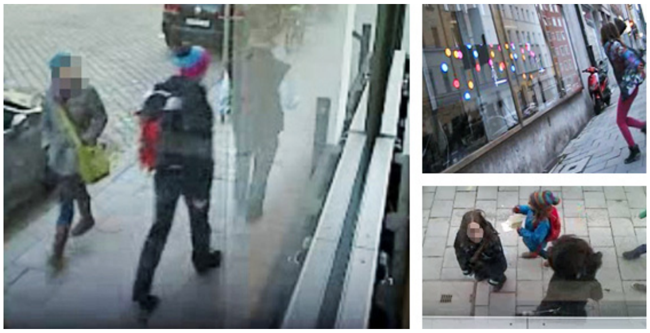

Figure 4. Three of five camera perspectives we used for recording interaction with the banner display.

attentive as groups as they were often in conversation. Especially the younger population stopped to engage with the interactive content. Of those who interacted (~7,5% of those looking at the display), all of the 20 interviewed stated that they had discovered the interactivity by the element moving along with them when passing by.

Video Analysis: Interactivity

Comparing all interactive content variations with the non-interactive baseline revealed that interactivity was correlated with attention to the display: On average 30% of passers-by reacted to the interactive conditions by looking at them, while only ~12% turned their heads towards the baseline content. This difference in attraction efficiency is significant ($\chi^2(1)=10.482$, $p<0.005$).

Video Analysis: Stimuli Conditions

Analyzing the interactive conditions individually revealed that there was no exceptional outlier: comparing the *horizontal position*, the *running-ahead* stimulus (AS) with 31.3% reactions was almost equally effective to the *sideways* stimulus (PS) with 28.7%. Comparing the *direction of movement* between conditions FS and RF, the school of fish was slightly more effective when moving along with the user than moving in *reversed direction* (34.2% to 29.7%). Stimuli moving along with the user (PS and AS combined) with 30% were likewise only slightly more effective than the fish approaching from the side (ES and AE) with 25.7%. Comparable values were obtained for both the greenscreen images and the rotated stimulus with 30% of passers-by reacting. The silhouette with 33% reactions was also a relatively effective stimulus.

Discussion

The first study revealed that for attracting users which are approaching sideways, *interactive stimuli* can significantly increase attention, even if barely within the field of view. The study also showed that different strategies to *position*, *move* or *animate* similarly designed stimuli were in principle equally effective. To our surprise, the *position ahead* did perform only slightly better than the *sideways stimulus,* bringing no essential advantage for attracting attention. We assume that the *school of fish* was most effective, as it was the largest in size and thus the most eye-catching. The stimulus moving in the *same direction* as the user was more effective than the reversed direction, possibly as it is easier for the user to understand that this motion relates to the own movement. All illustrated stimuli were equally effective to mirror images that have been reported to be effective with narrow displays [11,12]. This means that in sidewalk situations designers have more equally suited visual options for attracting attention.

FIELD STUDY 2: VISUAL AUDIENCE MODERATION

After user attention has been caught and users are turning towards the display, we were interested whether visual stimuli can be used to actively influence user positions and regulate audience constellations, thereby *visually moderating the audience*. This process of active behavior management poses the following basic problems:

User Positioning

Visual signifiers in the display must be able to effectively draw arriving users to arbitrary positions they would not have chosen themselves. For this, they must be *wittingly or unwittingly understood* and *immediately accepted by users*.

User Repositioning

When a user already interacts with the display, it might become necessary to reposition him or her in order to free space for new arrivers. Visual signifiers should be able to *make users step aside* and clear the needed space.

Audience Distribution

Visual signifiers should manage to *dissolve crowds* caused by passive group members gathering around an active user, and *distribute users more equally*. This may be achieved by guiding single users towards empty spots, but also by more undirected signals addressing the entire crowd.

If interactive content will succeed to accomplish these basic goals, dynamic visual signifiers might also be used to *shape specific audience constellations*. If, for example, an interactive ball game requires a certain number of players or a minimum distance between them to work properly, it might direct the active users accordingly.

Independent Variables

For actively moderating user positions, we identified the following variables for our visual signifier strategies:

Signifier Variable	Values
Movement	*static – moving*
Interactivity	*non-interactive – interactive*
Representation	*any behavior-effective stimulus*

Table 2. Variables for visual stimuli directing user positions.

For distributing users along the display, visual signifiers can be displayed at various *static* positions, or continuously *move* along the screen surface. Signifiers may just be *non-interactive*, *react* to the current user position, or even *act in relation to previous states* and constellations of multiple users. Preferably, signifier representations should be content-related, but first of all they have to be an *effective positioning stimulus*. Related Work [4] describes unwitting positioning effects in front of large visual frames simulating display bezels, yet it is not clear if visual frames have the same effect when displayed dynamically on the screen.

Conditions and Content

We designed five signifier strategies to position, reposition and distribute users. They all involved a simple ball game as background content, in which users can bounce shapes falling from above with a visual representation of their tracked skeleton. The game did not prefer any orientation or position in front of the screen. Conditions only differed in the behavior of the visual signifiers such as movement, interactivity and representation. The baseline was just the ball game without any visual signifiers. For our study we designed two basic visual signifier representations: *frames* and *ellipses*. The color, size and aspect ratio of the frames were partially predetermined by the other content. They had to provide enough space for the visual user representations, but still fit on the physical screen. The exact size, aspect ratio and bezel thickness were determined in *iterative pre-tests*. We noticed that a suitable shape of the signifier was crucial, as frames too small or too large could be misunderstood or disregarded. The visual frames had no additional functionality. The single signifier strategies are:

Static Frames (SF)

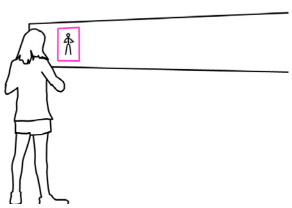

The most basic signifier concept were *static frames*. In this condition, two static visual frames were displayed in addition to the game (see Figure 1). Frame distance was chosen to minimize any mutual interference between players. Absolute frame positions were altered by imperceptibly moving them (crossing the screen in 7 min.). This slow movement was essential for randomizing the frame positions thus minimizing the influence of any external variables. This condition was designed to find out whether users would position themselves in front of the frames.

Static Ellipses (SE)

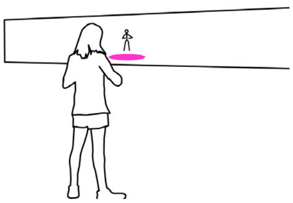

As an alternative to the visual frames we designed *static ellipses,* to investigate if there are any further visual signifiers inducing positioning effects as reported by [4]. They were also white and had an aspect ratio of 5:2. The ellipses also provided sufficient space for one user representation and appeared in the lower part of the screen to convey the impression of *spotlights* on a floor. Just as the static frames, the ellipses moved only with a very slow speed to randomize positions.

Figure 5. The *Static Ellipses* condition.

Moving Frames (MF)

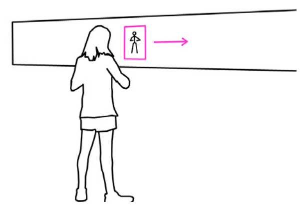

In contrast to the static frames, with the *moving frames* a single frame was moved in horizontal direction along the screen surface at a *perceptible speed*. It took 90 seconds for the frame to traverse the complete screen once, which corresponds to a moving speed of 0.15 km/h at which users could easily adapt their position to the frame by occasionally stepping aside. This was designed to check whether users would be willing to follow the frame and adapt their position repeatedly.

Dynamic Frames (DF)

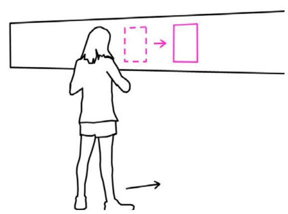

In contrast to the previous concepts, *dynamic frames* were only created upon user interaction and reacted to the user's position: When a user started to interact with the screen, he received an individual frame, which after half a second moved sideways by one frame width. This combines the positioning stimuli of the static and moving frames: It was designed to test if users would recognize a frame that was *individually created* for them when they had approached, and consequently perform a side step towards the displaced frame.

Multiple Dynamic Frames (MDF)

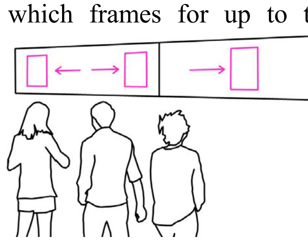

This is an extension of the dynamic frames concept, in which frames for up to three simultaneous users were coordinated with each other. The first approaching passer-by received a frame where he or she stopped. If a second user approached and stopped near the first user, the frame split into two frames moving apart until their positions provided enough space for both to comfortably interact in parallel. Hence the first user received a stimulus to reposition himself and make room for the second user, which might not have happened automatically. At the same time the second user received a stimulus to move to the resulting free space and occupy it. If a third person approached, a third frame appeared in a distant space. This condition was used to verify, if dynamic frames can be used to direct multiple users simultaneously and optimally distribute them across the screen surface.

Procedure and Data Analysis

The second field study was conducted in the evening hours over a period of five weeks. The procedure and data collection largely corresponds to the first study. 28 hours of video were analyzed, 47 log files collected and 20 semi-structured interviews conducted. In the developed *coding scheme*, each stopping passer-by was recorded as numbered subject and described by the following behaviors:

Positioning Behavior

For each user the state *framed* was set if he positioned fully within the boundaries of a stimulus, *boundary* if he was not fully within but still tangent to it, and *off frame* if he stood apart from signifiers or was in range of one by pure chance.

Repositioning Behavior

This behavior kept count of *how many times* a user repositioned himself in front of a signifier, e.g. when he left the signifier position to observe another player, and then returned to resume interaction in front of the signifier.

Distribution Behavior

To quantify crowding effects and their resolution, the state *crowded* was set if pairs or groups crowded together and interfered with each other, or *dispersed* if all users owned at least one arm's length personal distance to the neighbor.

User Activities

User activities such as *interacting* or *observing* were scored to *correlate* active or passive engagement of users to their positioning behavior, and to reveal *social constellations*, i.e. the combined behaviors of two or more persons.

Results

Within 28 hours of analyzed video the behavior of 304 passers-by who stopped and engaged with the display was scored. Of these, 267 interacted elaborately with the screen, and 37 were mere observers only watching other players.

Behavior in front of the Baseline

Overall 61 passers-by engaged with the baseline, with 1 passerby only watching others. To relate the baseline to the other conditions, we analyzed the *repositioning behavior* and *distribution* of users: Only one third of the active users (20 out of 60) repositioned, e.g., to gain more space while interacting, while the rest remained quite *static*. In 54% of the cases in which multiple users interacted in front of the wide screen, they crowded together instead of using the available free space. Such crowding of groups and pairs could be attributed to *bystanders* watching the initial user from close positions behind or next to the user, partly in orthogonal orientation to the display. If a passive member joined the interaction, then often without detaching from the crowd, and ignoring other users, which resulted in users interfering with and occluding each other (see Figure 6).

Figure 6. Typical by-standing and crowding effects.

General Observations: The Puppeteer Display

From the beginning, we observed strong positioning effects in front of the visual signifiers: passers-by did align themselves to the center of the *frames* and *ellipses* as soon as they were starting to interact. Also, users did not only align themselves once, but actually became attached to the visual stimuli that were slowly moving across the screen (see Figure 7). From the other side of the street, the display directing its users from one side to another and back, or guiding passive users to empty spots in front of the display, appeared like a puppeteer to us, manipulating the positions of puppets by strings. It was striking how immediately and willingly users accepted this manipulation of their behavior.

Figure 7. Users being subtly directed by visual signifiers.

Positioning in front of the Stimuli

Of the entire audience, 243 engaged with one of the five *signifier conditions*, out of which 207 people interacted and 36 only observed others. Of the passers-by who interacted, about 70% at least once deliberately moved to the center of the offered visual stimuli. Table 3 shows the share of actors aligning to a signifier for each individual condition:

SF	SE	MF	DF	MDF
78,8 %	87,8 %	66,7 %	65,8 %	56,5 %

Table 3. Share of active passers-by aligning to the signifiers.

To our surprise, correlating positions and activities reveals that the strongest position stimulus during interaction is generated by the *Static Ellipses.* With 83% of interaction time in front of the ellipses, they are even significantly more effective than the second best *Static Frames* with 58% of time spent in the central position (see Figure 8). Yet, the time users spent in front of the signifiers also reveals a varying effectiveness of the conditions to hold users: in contrast to the static signifiers, the majority of the interaction time was clearly spent off-frame for the *Moving Frames* and *Dynamic Frames*. People aligned to the *Multiple Dynamic Frames* for about half of the time. The boundary position was only passed shortly. Passive spectators usually preferred off-frame positions.

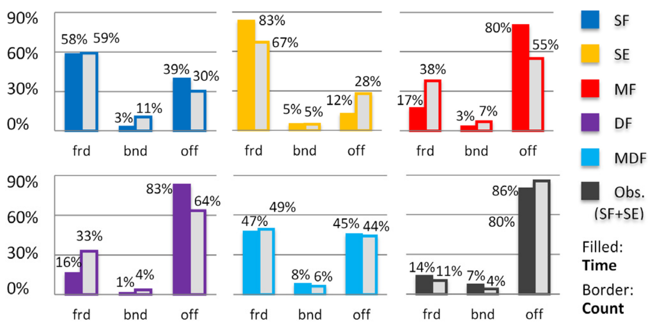

Figure 8. Positioning time and count when *interacting* with (grey: *observing*) the visual signifiers. Note that the signifiers make up only for a small fraction of the screen real estate.

Repositioning in front of the Stimuli

In contrast to the baseline, passers-by were rarely static in their behavior while interacting with the *signifier conditions*. Often, they repositioned themselves more than once within the same continuous interaction. The *Static Ellipses* generated the most repositioning, with 36 out of 41 interacting persons (88%) repositioning at least once and 13 or 32% repositioning a second time. One person was even coming back a fifth time to the ellipse signifier (see Fig. 9). Chi-squared tests showed that with all *signifier conditions* significantly more users repositioned a first time than with the *baseline* (all p<0.005, *MDF*: p<0.05). Looking at the mean durations after the first and second repositioning, interaction times decrease for the *Static Frames* and *Static Ellipses*, but increase for the *Moving Frames* and *Dynamic Frames* (see Figure 9 right). We observed that once users accepted the *moving stimuli* as reference for their positioning, they stayed longer in front of these signifiers the second time.

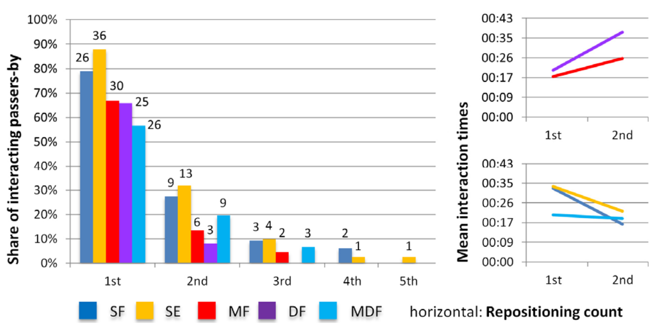

Figure 9. Repositioning count of actors (left) and mean interaction times after the first and second repositioning (right).

Distribution in front of the Stimuli

For all framing conditions, if pairs or groups were interacting, a distribution of users across the display was observed more often than crowding situations. Chi-squared tests showed that the distribution is significantly higher than in the baseline for all conditions but the *Dynamic Frames* (all p<0.005, *DF*: p>0.05). The best results were achieved with the *Multiple Dynamic Frames*, with a comfortable distribution of passers-by in 96% of cases. Also the *Moving Frames* and the *Static Ellipses* performed well with 88% and 86% of pairs and groups distributed (see Figure 10).

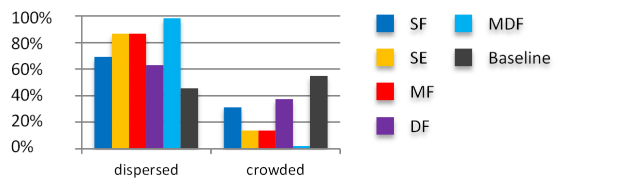

Figure 10. Distribution of two or more passers-by across the signifier conditions and the baseline on a percentage basis.

Social Interaction

If more than one user engaged with the *baseline* condition, the most frequent behavior observed were active users initially being watched by their partners and group members. Such initial *by-standing behavior* sometimes followed by self-contained interaction without adapting one's own position was much more often the reason for close distances than *interaction between users* (e.g. joining hands, embracing or boxing the partner). While close bystanders could also be observed with the *signifier conditions*, active users were detaching from the initial actor, coming back only when the other user had to show something. When interacting persons were observed by others, they tended to position themselves even more eagerly in front of the stimulus (see Figure 11).

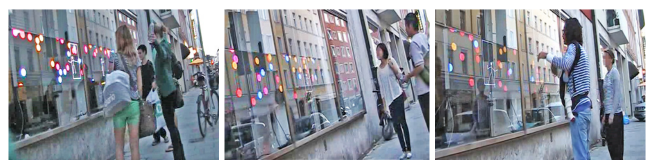

Figure 11. Performing in front of signifiers: The active user demonstrates to the bystanders how to position accordingly.

Correlating user activities and positions revealed that in the majority of *social constellations* at least one member was interacting in front of one of the stimuli. Table 4 illustrates the social interaction of pairs across all conditions:

Pair Constellations	Duration	Count	Mean time
P1 interacts off-frame/ P2 watches	00:01:22	7	00:00:12
P1 interacts boundary/ P2 watches	00:02:03	7	00:00:18
P1 interacts framed/ P2 watches	**00:07:33**	**38**	**00:00:12**
P1 interacts off-frame/ P2 interacts off-frame	00:20:25	47	00:00:26
P1 interacts boundary/ P2 interacts boundary	00:00:15	1	00:00:15
P1 interacts framed/ P2 interacts framed	**00:11:51**	**35**	**00:00:20**
P1 interacts framed/ P2 interacts off-frame	**00:14:39**	**49**	**00:00:18**

Table 4. Social interaction of pairs: In 122 out of 184 cases or 66% of constellations (rows 3,6,7 combined) at least one of the partners is interacting in front of a visual signifier.

Interviews: Awareness of the Signifiers

Of the 20 interviewees of the second study 10 had interacted with the *Static Frames*, 3 with the *Static Ellipses* and 7 with the *Multiple Dynamic Frames*. All could recall the basic elements displayed on the screen. In particular, all had recognized the frame or ellipse signifier, and of the 16 that had positioned themselves in front of a signifier 15 answered they did so because of the signifier.

Discussion

Quantitative analysis showed that a significant majority of passers-by were influenced in their position by the visual signifiers. Potential reasons for the found behavioral patterns and varying effectiveness of the stimuli may be:

Effect of the Visual Frames

The visual rectangular frame as used in this study turned out to be effective to regulate the positions of users. Users interpreted the virtual frames correctly from the beginning and instinctively accepted the prescribed stimulus instead of choosing the position themselves. This behavior may result from a conditioning of humans for real physical frames such as picture frames or television screens. This behavior did not occur unconsciously, almost all participants of the interview stated to have positioned themselves as a consequence of the frame. This *conscious perception of the frames* is different from the observation of [4] with large static frames. The reason may be that the frames in our case were smaller and did not coincide with the real physical frame, and thus can clearly be identified as a component of the virtual content having certain functionality.

Effect of the Ellipses

The *Static Ellipses*, designed to test if any other stimuli than frames would perform to influence passer-by positions, surprisingly *even outperformed* the *Static Frames*. While ellipses resembling solid floor areas or platforms are not as common as rectangular frames in contemporary digital content, our understanding of spatial perspective may play a role: The most important plane of reference in 2D and 3D space is the ground plane, having a particular relevance for human orientation [8]. This role of the floor may at least partially have caused the preferred positioning of one's own screen representation on a solid platform within virtual space. Yet, if such an instinctive stimulus plays a stronger role than the understanding of this signifier as functional element still has to be evaluated. Individual content-related associations, e.g. interpreting the ellipses as the only illuminated position on a stage, cannot be ruled out either.

Static vs. Dynamic Stimulus

With the slowly moving *Static Frames* and *Ellipses* the *interaction time* inside the stimulus was longer than outside, showing that users do not only align themselves once to these stimuli, but also can become attached to them (as puppets to the string). This could not be observed for the *Moved Frames* and *Dynamic Frames* moving away quickly from the user. Yet, there was no such significant difference between static and moving concepts in regard to the number of first positionings in front of the stimuli. In other words, users did not position themselves less often in front of the *strongly moving signifiers*, but they were simply *not willing to follow them*, showing reluctance to reposition themselves repeatedly. When such a stimulus pulled away, they might have recognized that interaction was also possible outside of the frame, i.e. that it had no actual functionality.

Distribution of Passers-by

Members of a pair or group did often not have enough space to interact without impeding each other in front of the baseline. The signifier concepts can resolve such crowding effects: except for the regionally limited *Dynamic Frames*, all visual stimuli generated a *significantly* better distribution of users during interaction than the baseline, and for all significantly increased reposition numbers were obtained. This means that visual stimuli can not only be used as a tool to control single users, but also to *equalize and dissolve groups,* thus influencing the dynamic of public interaction. The very dynamic and clearly inviting *Multiple Dynamic Frames* with 96% performed best in this situation.

Social and Performative Interaction

If interacting persons were observed they tended to position themselves even more eagerly in front of the stimulus. This behavior can be attributed to a so-called *performative interaction* [7]: The interacting user takes a presenting role, demonstrating interaction to others. This situation creates a kind of social pressure on the actor to not look ridiculous in front of the audience, but instead to act in an exemplary manner and show others how to interact appropriately.

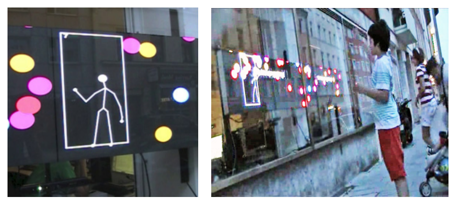

Figure 12. Improving the positioning stimulus by additional functionality: frame with a game counter in the top left corner.

DESIGNING FURTHER SIGNIFIERS

In our study the *frames* and *ellipses* performed well as positioning stimuli. Other good designs may exist, but not all visual signifiers may be equally effective. Signifiers could be designed *content-related*, using visuals which are meaningful in the context of the content. For example, on a field of flowers, sunbeams may be more fitting than frames. Then, the *effectiveness of a visual signifier* as positioning stimulus should be tested first and refined iteratively before use. *Static signifiers* should be preferred as we found strong motion significantly reduces the willingness of users to stay in front of the signifiers. Future work could investigate if the performance of moving frames can be improved by augmenting them with *additional functionality* such as a game counter (see Figure 12). If a user perceives such an obvious advantage of standing in front of the signifier, this may increase the positioning stimulus. Also, *additional visual cues* could improve the effectiveness of the tested positioning stimuli, such as the exogenous and endogenous cues for shifting the attention of users to spatial spotlights explored by Posner et al. [16]. Comparably, if current users in front of the screen have to be repositioned, instead of moving the frames arrows pointing to the frames or flashing frames could be used to first direct the attention to them.

Finally, it would be interesting to investigate if actively created audience constellations can *trigger engagement* or improve the performance with interactive games.

CONCLUSION

Very wide displays in public space are usually approached sideways, but also offer sufficient space for multiple simultaneous users. In this paper, we investigated how an interactive *banner display* deployed in the wild can attract attention when approached sideways, and showed how *visual moderation* can direct users to arbitrary positions in front of the display and thus *actively shape the audience.* In a field study, we compared different static, moving and interactive signifiers regarding their effectiveness to direct users. We found that ellipses generated the most effective positioning stimulus, while static frames were also very effective. Moving and strongly dynamic signifiers can also lead to more distribution of the audience along the display, but users were not willing to repeatedly follow moving frames. In pairs or groups of users, social pressure might push one user into a performer role in front of the signifiers. During the study, in our impression the display directed its users like a puppeteer. Based on these observations, we believe that actively directing users will increase the effectiveness of future interactive public displays.

REFERENCES

1. Akpan, I., Marshall, P., Bird, J., and Harrison, D. Exploring the Effects of Space and Place on Engagement with an Interactive Installation. In *Proc. ACM CHI '13*, ACM Press (2013), 2213–2222.
2. Ballendat, T., Marquardt, N., and Greenberg, S. Proxemic Interaction: Designing for a Proximity and Orientation-Aware Environment. In *Proc. ACM ITS '10*, ACM Press (2010), 121–130.
3. Brignull, H., and Rogers, Y. Enticing People to Interact with Large Public Displays in Public Spaces. IFIP INTERACT 2003, 17–24.
4. Beyer, G., Köttner, F., Schiewe, M., Haulsen, I., and Butz, A. Squaring the Circle: How Framing Influences User Behavior around a Seamless Cylindrical Display. In *Proc. ACM CHI '13*, ACM Press (2013), 1729–1738.
5. Beyer, G., and Butz, A. Communicating the Interactivity of differently shaped Displays. CHI 2013 Workshop on Displays Take New Shape: An Agenda for Interactive Surfaces, April 28, 2013, Paris, France.
6. Dalsgaard, P., and Halskov, K. Designing Urban Media Façades: Cases and Challenges. In *Proc. ACM CHI '10*, ACM Press (2010), 2277–2286.
7. Dalsgaard, P., and Hansen, L. K. Performing Perception – Staging Aesthetics of Interaction. *ACM TOCHI 15*, 3, ACM Press (2008), 1–33.
8. Gibson. *The Senses Considered as Perceptual Systems*. Greenwood Pub Group Inc., Cornell University, 1966.
9. Grace, K., Wasinger, R., Ackad, C., Collins, A., Dawson, O., Gluga, R., Kay, J., and Tomitsch, M. Conveying Interactivity at an Interactive Public Information Display. In *Proc. ACM PerDis '13*, ACM Press (2013), 19–24.
10. Klinkhammer, D., Nitsche, M., Specht, M., and Reiterer, H. Adaptive Personal Territories for Co-located Tabletop Interaction in a Museum Setting. In *Proc. ACM ITS '11*, ACM Press (2011), 107–110.
11. Michelis, D., and Resatsch, F. Unlocking the Interactive Capabilities of Large Outdoor Displays. *Information Display 23*, 3 (2007), 24–29.
12. Müller, J., Walter, R., Bailly, G., Nischt, M., and Alt, F. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In *Proc. ACM CHI '12*, ACM Press (2012), 297–306.
13. Nacenta, M. A., Sakurai, S., Yamaguchi, T., Miki, Y., Itoh, Y., Kitamura, Y., Subramanian, S., and Gutwin, C. E-conic: A Perspective-Aware Interface for Multi-Display Environments. In *Proc. ACM UIST '07*, ACM Press (2007), 279–288.
14. Peltonen, P., Kurvinen, E., Salovaara, A., Jacucci, G., Ilmonen, T., Evans, J., Oulasvirta, A., and Saarikko, P. It's Mine, Don't Touch!: Interactions at a Large Multi-Touch Display in a City Centre. In *Proc. ACM CHI '08*, ACM Press (2008), 1285–1294.
15. Pinhanez, C., and Podlaseck, M. To Frame or Not to Frame: The Role and Design of Frameless Displays in Ubiquitous Applications. In *Proc. Ubiquitous Computing*, Springer-Verlag (2005), 340–357.
16. Posner, M., Snyder, C., and Davidson, B. Attention and the Detection of Signals. *Journal of Experimental Psychology: General 109*, 2 (June 1980), 160–174.
17. Reeves, S., Benford, S., O'Malley, C., and Fraser, M. Designing the Spectator Experience. In *Proc. ACM CHI '05*, ACM Press (2005), 741–750.
18. Schmidt, C., Müller, J., and Bailly, G. Screenfinity: Extending the Perception Area of Content on Very Large Public Displays. In *Proc. ACM CHI '13*, ACM Press (2013), 1719–1728.
19. Ten Koppel, M., Bailly, G., Müller, J., and Walter, R. Chained Displays: Configurations of Public Displays can be used to influence Actor-, Audience-, and Passer-By Behavior. In *Proc. ACM CHI '12*, ACM Press (2012), 317–326.
20. Vogel, D., and Balakrishnan, R. Interactive public ambient displays: Transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proc. ACM UIST '04*, ACM Press (2004), 137–146.
21. Wang, M., Boring, S., and Greenberg, S. Proxemic Peddler: A Public Advertising Display that Captures and Preserves the Attention of a Passerby. In *Proc. ACM PerDis '12,* ACM Press (2012), 3:1–3:6.