



A Design Space for Pervasive Advertising on Public Displays

Florian Alt

A DESIGN SPACE FOR PERVASIVE ADVERTISING ON PUBLIC DISPLAYS

Von der Fakultät für Informatik, Elektrotechnik und
Informationstechnik der Universität Stuttgart zur Erlangung der
Würde eines Doktors der Naturwissenschaften (Dr. rer. nat.)
genehmigte Abhandlung

vorgelegt von

FLORIAN ALT

aus München

Erstgutachter: Prof. Dr. Albrecht Schmidt
Zweitgutachter: Prof. Nigel Davies, PhD

Tag der mündlichen Prüfung: 07.12.2012

Institut für Visualisierung und Interaktive Systeme
der Universität Stuttgart

2013

ABSTRACT

Today, people living in cities see up to 5000 ads per day and many of them are presented on public displays. More and more of these public displays are networked and equipped with various types of sensors, making them part of a global infrastructure that is currently emerging. Such networked and interactive public displays provide the opportunity to create a benefit for society in the form of immersive experiences and relevant content. In this way, they can overcome the display blindness that evolved among passersby over the years. We see two main reasons that prevent this vision from coming true: first, public displays are stuck with traditional advertising as the driving business model, making it difficult for novel, interactive applications to enter the scene. Second, no common ground exists for researchers or advertisers that outline important challenges. The provider view and audience view need to be addressed to make open, interactive display networks, successful.

The main contribution made by this thesis is presenting a design space for advertising on public displays that identifies important challenges – mainly from a human-computer interaction perspective. Solutions to these core challenges are presented and evaluated, using empirical methods commonly applied in HCI.

First, we look at challenges that arise from the shared use of display space. We conducted an observational study of traditional public notice areas that allowed us to identify different stakeholders, to understand their needs and motivations, to unveil current practices used to exercise control over the display, and to understand the interplay between space, stakeholders, and content. We present a set of design implications for open public display networks that we applied when implementing and evaluating a digital public notice area.

Second, we tackle the challenge of making the user interact by taking a closer look at attracting attention, communicating interactivity, and enticing interaction. Attracting attention is crucial for any further action to happen. We present an approach that exploits gaze as a powerful input modality. By adapting content based on gaze, we are able to show a significant increase in attention and an effect on the user's attitude. In order to communicate interactivity, we show that the mirror representation of the user is a powerful interactivity cue. Finally, in order to entice interaction, we show that the user needs to be motivated to interact and to understand how interaction works. Findings from our experiments reveal direct touch and the mobile phone as suitable interaction technologies. In addition, these findings suggest that relevance of content, privacy, and security have a strong influence on user motivation.

Third, this thesis makes a set of contributions towards understanding audience behavior, which is particularly important for advertisers in order to choose appropriate content and to select suitable locations for future advertising displays. Our findings provide an in-depth understanding of the honeypot effect as a powerful interactivity cue. Furthermore, we identify a number of interesting effects (e.g., the landing effect) and explain how developers could design for them.

We envision the results of this thesis to provide a basis for future research and for practitioners to shape future advertisements on public displays in a positive way.

ZUSAMMENFASSUNG

Heutzutage sehen Stadtbewohner täglich bis zu 5000 Werbebotschaften, viele davon auf öffentlichen Bildschirmen. Immer häufiger sind diese Bildschirme vernetzt und mit Sensoren ausgestattet, was zukünftig zu einer globalen Display-Infrastruktur führen wird. Mit vernetzten und interaktiven (Groß-)Bildschirmen werden für den Benutzer wirkungsvolle Erlebnisse mit für ihn relevanten Inhalten möglich. Für die Passanten entsteht ein Mehrwert und sie nehmen gleichzeitig mehr Inhalte von Großbildschirmen wahr. Zwei Hauptgründe sind dafür verantwortlich, dass diese Vision bis heute nur teilweise realisiert werden konnte: Zum einen ist Werbung nach wie vor das vorherrschende Geschäftsmodell, und neue, interaktive Anwendungen haben es schwer, sich gegen traditionelle Werbeeinhalte durchzusetzen. Zum anderen fehlt so etwas wie eine gemeinsame Basis für Wissenschaftler und Werbetreibende, die wichtige Herausforderungen zur erfolgreichen Einführung (aus Sicht der Anbieter als auch der Benutzer) von interaktiven Displays zusammenfasst.

Der Hauptbeitrag der vorliegenden Arbeit besteht in der Entwicklung eines Gestaltungsraums für Werbung auf öffentlichen Bildschirmen und der Identifizierung von damit verbundenen wichtigen Herausforderungen, vor allem im Hinblick auf die Mensch-Computer-Interaktion. Mithilfe gängiger empirischer Methoden aus der Mensch-Computer-Interaktion werden Lösungen für diese zentralen Herausforderungen aufgezeigt und evaluiert.

Zuerst beschäftigt sich diese Arbeit mit den Schwierigkeiten, die bei einer gemeinsamen Nutzung von Bildschirmen entstehen. Dazu wurde eine beobachtende Studie an herkömmlichen schwarzen Brettern durchgeführt. Diese ermöglichte Einblicke in die Bedürfnisse und Beweggründe der verschiedenen Interessengruppen, die Identifizierung der gängigen Praktiken zur Organisation der Displays und ein Verständnis der Wechselwirkung zwischen Raum, Interessengruppe und Inhalt. Aus diesen Beobachtungen heraus wurde eine Reihe von Empfehlungen für die Gestaltung von offenen Bildschirmnetzwerken abgeleitet, welche bei der Implementierung eines digitalen schwarzen Bretts angewendet und in Labor- und Feldstudien evaluiert wurden.

Des Weiteren wurden drei Hauptfragen genauer betrachtet: Wie kann die Aufmerksamkeit des Passanten erregt, Interaktivität kommuniziert und der Nutzer zur Interaktion bewegt werden? Ausschlaggebend für jede weitere Handlung ist es, dass die Aufmerksamkeit erregt wird. Es wird ein Ansatz vorgestellt, der die Blickinformationen als Eingabemodalität nutzt. Die Blickinformationen werden in Echtzeit ausgewertet und dazu verwendet, Inhalte auf der Seite entsprechend

anzupassen. In einer Benutzerstudie konnten gezeigt werden, dass sich dadurch die Aufmerksamkeit signifikant erhöht; zudem konnte ein Effekt auf die Einstellung des Nutzers in Bezug auf die betrachteten Inhalte nachgewiesen werden. In einem zweiten Schritt soll der Benutzer verstehen, dass ein Bildschirm interaktiv ist. Es konnte gezeigt werden, dass sich hierfür das Spiegelbild des Benutzers gut eignet. Um schließlich den Benutzer zur Interaktion anzuregen, muss er motiviert sein und die verwendete Interaktionstechnik verstehen. Die Ergebnisse unserer Studie zeigen, dass die direkte Touch-Interaktion mit dem Bildschirm und die Verwendung des Mobiltelefons als Interaktionsgerät geeignet sind. Einen starken Einfluss auf die Motivation des Benutzers haben zudem die Relevanz der Inhalte, der Datenschutz und die Privatsphäre sowie die Möglichkeit zur sicheren Authentifizierung.

Nicht zuletzt trägt diese Arbeit zum Verständnis des Benutzerverhaltens bei. Dies ist besonders wichtig für Werbetreibende, um die Inhalte und passende Standorte für zukünftige Werbeflächen auszuwählen. Die Ergebnisse ermöglichen ein weitreichendes Verständnis des sogenannten Honeypot-Effekts, der als deutlicher Hinweis auf Interaktivität dient. Zudem wurden einige interessante Effekte gefunden (der Landing-Effekt, Effekte der Gruppeninteraktion), und es wird erklärt, wie Entwickler diese in ihrem Design berücksichtigen können.

Mit dieser Arbeit soll eine gemeinsame Basis für die künftige Forschung und für die beteiligten ausübenden Berufe gelegt werden, damit die Werbung der Zukunft auf öffentlichen Bildschirmen so gestaltet werden kann, dass sie von allen Interessengruppen gewinnbringend eingesetzt und genutzt werden kann.

ACKNOWLEDGMENTS

In the past four years, I had the privilege of meeting and working with a number of outstanding people. They contributed to the research leading to this thesis – not only in a professional but, even more importantly, also in a personal way. This is reflected through the use of the editorial ‘we’. These acknowledgements are meant to point out all those great people I had the opportunity to work with and I sincerely apologize to everybody I missed to mention.

My deepest thanks go to my supervisor **Albrecht Schmidt** whose passion for research has ever been an inspiration for me. His never-ending flow of ideas provided the foundations of my work. Looking at his track record over the last years I keep wondering from where he took the time to provide useful advice whenever I needed it.

He also takes credit for assembling the outstanding working groups I had the pleasure to work with over the last years. I want to particular mention **Alireza Sahami Shirazi** who was not only my office mate over the past years but became one of my best friends as we shared a flat in Düsseldorf. Thanks also to **Bastian Pfleging** (for always coming up with the best offers during our pre and post conference activities), **Paul Holleis** (for teaching me a lot about writing papers and being my tennis partner in Essen), **Dagmar Kern** (for the best Carnival parties in Köln), **Stefan Schneegaß** (for all the soccer-and-BBQ evenings in my garden), **Christian Winkler** (for always letting me win in tennis when we played for correcting exams), and **Thomas Kubitz** (for the sauna nights in Oulu). Further thanks go to my colleagues **Tanja Döring**, **Niels Henze**, **Elba del Carmen Valderrama Bahamondez**, **Anja Mebus**, **Nicole Recksing**, and **Tilman Dingler** as well as to all the undergraduate students that contributed to my work. **Jennifer Huang** and **Grace Tai** did a fantastic job when proof-reading the thesis. I would also like to thank **Julian Mennenöh** (for being the god of data mining and statistics, which he recently proved as the winner of the Data Mining Cup 2012) and **Stefanie Kristes** (for several cinema nights in Düsseldorf and the annual Christmas market visits).

During the four years, a lot of amazing people crossed my way who supported me in various ways. I would like to express my special thanks to **Jörg Müller** for inviting me to spend 6 months of my PhD in Berlin. He is not only a brilliant researcher, but also a great friend who made my time in Berlin even more enjoyable. I also want to thank **Andreas Bulling** for the great collaborations – next time it is definitely up to me to visit you. Further thanks go to **Jonna Häkkinen** and her former team – particularly **Ari-Heikki Sarjanoja**, **Jussi Huhtala**, and

Arto Puikkonen – for hosting us several times in Finland, including unforgettable sauna events and party nights. Particular thanks also go to **Timo Ojala** for all his support during the UbiChallenge 2012. And finally I would like to thank all those senior researchers who, in one way or the other, contributed to the thesis through their useful comments and feedback over a beer or glass of wine after a long conference day, including **Elaine Huang**, **Bo Begole**, **Aaron Quigley**, and **Enrico Rukzio**.

The PD-Net project provided a great basis for ideas and discussions. I want to particularly thank **Nigel Davies** not only for being my external examiner, but for the great evenings at his house, and the ManU jersey. Further thanks go to **Nemanja Memarovic** – aka the paper machine – who became a very good friend and with whom I enjoyed having a lot of discussions and parties (just remembering the open bar night in SF prior to our early morning talk). Further thanks go to **Adrian Friday** (for asking all the questions), **Marc Langheinrich** (for always reminding us to design for privacy and for the roof-top party in Zurich), **Ivan Elhart** (for making sure that everybody abided the UNO rules in Finland), and **Rui Jose** (for the great dinners we had in Portugal).

I want to thank my parents, **Günther Alt** and **Friedlinde Alt** for their unconditional support of whatever I did in life. It is thanks to my father and my brother **Michael Alt** that I first became interested in technology when I was a great deal less large, and thanks to my mother that I nowadays (still) enjoy all those travels around the globe. Sorry for so often being away when you needed me.

Last, but certainly not least, **Susanne Reisenweber** probably had to struggle more with this thesis than I had to. Without your support, deprivations, and love I surely could not have made it. I sincerely apologize for all those nights I spent writing papers, the weeks of endless travels, and having been away from you for more than 7 years. Thank you for always standing by me.

TABLE OF CONTENTS

List of Figures	xvii
------------------------	-------------

List of Tables	xix
-----------------------	------------

I INTRODUCTION AND MOTIVATION	1
--------------------------------------	----------

1 Introduction	3
-----------------------	----------

1.1 Research Questions	5
1.2 Methodology	6
1.2.1 Prototypes	7
1.2.2 Evaluation	7
1.2.3 Ethics	8
1.3 Summary of Research Contributions	9
1.3.1 Design Space	9
1.3.2 Prototypes	10
1.3.3 Conceptual Foundations	11
1.4 Research Context	11
1.5 Thesis Outline	13

II BACKGROUND	21
----------------------	-----------

2 Public Displays	23
--------------------------	-----------

2.1 History of Public Display Research	24
2.2 Commercial Display Systems	28
2.3 Technology	31

2.3.1	Displays	31
2.3.2	Projectors	32
2.3.3	Sensors	34
2.4	Engagement with the Display	36
2.4.1	Attention	36
2.4.2	Motivating Interaction	39
2.5	Interaction Modalities	44
2.5.1	Presence	44
2.5.2	Body Position	44
2.5.3	Body Posture	45
2.5.4	Face Detection and Face Recognition	45
2.5.5	Gaze Direction	46
2.5.6	Speech	46
2.5.7	Gestures	47
2.5.8	Keyboard and Mouse	47
2.5.9	Touch	47
2.6	Metaphors	48
2.6.1	Posters	48
2.6.2	Window	48
2.6.3	Mirror	49
2.6.4	Overlay	49
2.7	Interaction Models	49
2.8	Evaluation Techniques	51
2.8.1	Previous Work	53
2.8.2	Research Questions	53
2.8.3	Approaches to Research	55
2.8.4	Methods and Tools	58
2.8.5	Assessing Validity	60
2.8.6	Guidelines	61
3	Pervasive Advertising	63
3.1	Definitions	65
3.1.1	Marketing	65
3.1.2	Advertising	67
3.1.3	Pervasive Advertising	68

3.2	Objectives	68
3.2.1	Advertising Programs	69
3.2.2	Advertising Performance	71
3.2.3	Targeting	77
3.3	Opportunities	78
3.3.1	User Feedback	78
3.3.2	The Long Tail	79
3.3.3	Engaging Experiences	79
3.3.4	Personalization and Context Adaptivity	80
3.3.5	Audience Measurement	80
3.4	Challenges	82
3.4.1	Calm vs. Engaging Advertising	82
3.4.2	Privacy vs. Personalization	82
3.4.3	Private vs. Public Advertising	83
3.4.4	Local vs. Global Advertising	83
3.4.5	Persuasion vs. Ethics	83

III THE PROVIDER VIEW

4 System Design

4.1	Related Work	91
4.2	Sharing Public Display Space	92
4.2.1	Value Propositions of PNAs for Stakeholders	93
4.2.2	Content Exchange and Access Control	94
4.2.3	Learning from Practices in the Analog World	95
4.3	Study Design	95
4.3.1	Observational Studies	96
4.3.2	Interviews	96
4.4	Data Analysis	98
4.5	Findings	99
4.5.1	Stakeholders and Motivation	100
4.5.2	Displays and Content	101
4.5.3	Managing Content and Supporting Memory	103
4.5.4	The Role of Shared Displays in a Space	104

4.5.5	Interplay of Space, Stakeholders and Content	104
4.5.6	Needs of Content Providers and Display Owners	105
4.6	Design Implications	106
4.6.1	Design for Specific Uses of Notice Boards	106
4.6.2	Respect the Neighborhood Focus of PNAs	107
4.6.3	Support the Emergent Profiling of a PNA	108
4.6.4	Design for Flexibility of Input	108
4.6.5	Support Retrieval of Information	109
4.7	Discussion	110
4.8	Case Study: Car Advertisements	110
4.8.1	Privately-owned Cars: Online Survey	111
4.8.2	Privately-owned Cars: Qualitative Interviews	114
4.8.3	Company-owned Cars: Qualitative Interview	115
4.8.4	Requirements	115
4.8.5	A Platform for Dynamic Contextual Advertising	117
4.8.6	Related Work	118
4.8.7	Summary	118

5 Audience Behavior 121

5.1	Related Work	122
5.1.1	Interactive and Public Displays	123
5.1.2	Designs of Non-Planar Screens	123
5.1.3	Assumptions of Current Designs	124
5.1.4	Audience Behavior in Front of Flat Displays	125
5.1.5	Effect of Display Frame on Audience Behavior	125
5.2	Prototype of a Cylindrical Display	126
5.3	Hypotheses	127
5.4	User Study Design	129
5.4.1	Setup	129
5.4.2	Participants	130
5.4.3	Procedure	130
5.5	Data Analysis	131
5.6	Results	133
5.6.1	Walking	133
5.6.2	Shoulders	134

5.6.3 Time Spent for Interaction 135

5.6.4 Post-hoc Analysis 135

5.6.5 User Experience 136

5.6.6 Semi-Structured Interviews 136

5.7 Discussion 137

5.7.1 Moving Around the Column 137

5.7.2 The Sweet Spot 137

5.7.3 Time Spent for Interaction 138

5.8 Recommendations 138

5.8.1 Design for Walking Interaction 138

5.8.2 Place Columns in the Path of Users 139

5.8.3 Enable Gesture-Based Interaction 139

5.8.4 Use Frameless or Semi-Framed Content 139

5.9 Summary 140

6 Targeting 143

6.1 Related Work 146

6.1.1 Targeted Advertising and Profiling 146

6.1.2 Public Display Networks 147

6.1.3 Interaction Techniques 147

6.2 Opportunities of Pervasive Advertising 148

6.3 Concept and Profiling 151

6.3.1 Creating Profiles 152

6.3.2 Targeting Content Towards the Users 158

6.4 Implementation 160

6.4.1 Requirements 160

6.4.2 Architectural Concept 161

6.4.3 Implementation 162

6.5 Privacy 164

6.6 Discussion 164

IV THE AUDIENCE VIEW 167

7 Attracting Attention 171

7.1	Background and Related Work	173
7.2	Gaze-Based Adaptation of Content	176
7.2.1	Metrics	178
7.2.2	Approach	178
7.2.3	Hypotheses	179
7.2.4	Cognitive Effect	179
7.3	Prototype	180
7.3.1	Eye Tracker	180
7.3.2	UsaProxy	180
7.3.3	Apparatus	182
7.4	Study I: User Attention	185
7.4.1	Methodology	185
7.4.2	Participants	186
7.4.3	Setup and Procedure	186
7.4.4	Data Analysis	186
7.4.5	Results	187
7.5	Study II: User Interest	188
7.5.1	Setup and Procedure	188
7.5.2	Data Analysis	189
7.5.3	Results	189
7.6	Summary and Discussion	189

8 Communicating Interactivity 193

8.1	Related Work	194
8.1.1	Psychological Cues & Interactivity	197
8.1.2	Representation: Recognizing Yourself	197
8.1.3	Abstraction, Biological Motion, Body Schema	199
8.1.4	Perceptual Causality and Animacy	200
8.1.5	Summary	200
8.2	Prototype	201
8.2.1	Software	201
8.2.2	Hardware	201

8.3	Pre Study	201
8.4	Lab Study	202
8.4.1	Objective	203
8.4.2	Conditions	203
8.4.3	Task and Stimulus	204
8.4.4	Apparatus and Design	204
8.4.5	Results	204
8.4.6	Summary	205
8.5	Field Study	206
8.5.1	Deployment	206
8.5.2	Conditions	207
8.5.3	Data Analysis	209
8.5.4	Findings	209
8.6	Summary	219

9 Enticing Interaction 221

9.1	Related Work	223
9.2	Lessons Learned From Traditional PNAs	225
9.3	Research Questions	227
9.3.1	Content	227
9.3.2	Interaction Techniques	228
9.3.3	Privacy	228
9.4	Prototype	229
9.4.1	Digifieds Server	230
9.4.2	Digifieds Display Client	232
9.4.3	Digifieds Mobile Phone Client	234
9.4.4	Digifieds Web Client	235
9.5	Lab Study	236
9.5.1	Interaction Techniques	237
9.5.2	Tasks	239
9.5.3	Setup and Data Collection	241
9.5.4	Procedure	242
9.5.5	Findings from the Lab	242
9.5.6	Implications from the Lab	245
9.6	Deployment in the Wild	247

9.6.1	Evaluation	248
9.6.2	Findings from the Real World	251
9.6.3	Lessons Learned from the Real World	258
9.7	Case Study: Secure Authentication	259
9.7.1	Concept	261
9.7.2	Apparatus	262
9.7.3	Evaluation	263
9.7.4	Discussion	264
9.8	Summary	265
 V CONCLUSION		 269
 10 Conclusion		 271
10.1	Summary of Research Contributions	271
10.1.1	Shared Use of Public Displays	273
10.1.2	Making the User Interact	275
10.1.3	Audience Behavior Around Public Displays	277
10.2	Future Work	279
10.2.1	Towards a Public Display Architecture	279
10.2.2	Key Performance Indicators	282
10.2.3	Payment Models	283
10.2.4	New Technologies and User Experiences	283
10.2.5	Opportunities of Heterogenous Environments	284
10.2.6	Displays as Gateways	285
10.3	Concluding Remarks	285
 VI BIBLIOGRAPHY		 287
 Bibliography		 289
 VII APPENDIX		 333

LIST OF FIGURES

1.1	Thesis Outline	14
2.1	Interaction Models for Public Displays	50
3.1	Stimulus-Organism-Response Paradigm	72
3.2	States and Processes of Consumer Behavior	74
3.3	Three Levels of Involvement	74
4.1	Data Walkthrough	98
4.2	Types of Displays	102
4.3	Retrieving Information	103
4.4	Survey Results: Compensation and Technical Solutions	112
4.5	Survey Results: When Do Users Want to Advertise on their Cars?	112
4.6	Survey Results: What Do People Want to Advertise For?	113
5.1	Prototype of a Cylindrical Display	126
5.2	User Study Setup	129
5.3	Room Layout for the User Study	133
5.4	Typical User Behavior in Front of the Cylindrical Display	134
5.5	Typical User Behavior in Front of the Flat Display	134
6.1	Example for Creating Target Groups	151
6.2	Evolution of a Consumer Profile	153
6.3	Defining Target Groups	157
6.4	Calculating the Individual Fit of an Advertisement.	159
6.5	CAdEt System Architecture	161
6.6	CAdEt Advertiser's Client	163
7.1	Creating Adaptive Content Based on Gaze	178
7.2	Experimental Setup	181
7.3	Layout of the Test Website	183
7.4	Adaptation of Content	184
7.5	Creating Adaptive WebAds	185
8.1	Inadvertent Interaction	198
8.2	Looking Glass Prototype	202

8.3	Results from the Lab Study	205
8.4	Study Location	206
8.5	Representations Used in the Field Study.	208
8.6	Interaction Durations	210
8.7	Understanding Interactivity in the Call-to-Action Condition	211
8.8	Knowledge Building Up Over Time	212
8.9	Landing Effect for a Group	213
8.10	Landing Effect for a Couple	214
8.11	The Honeypot Effect	214
8.12	Interaction in Multiple Rows	216
9.1	Digifieds – Conceptual Architecture	229
9.2	Digifieds – System Architecture	230
9.3	Conceptual Layout of the Display Client	231
9.4	Digifieds Display Client	231
9.5	Digifieds Mobile Client	233
9.6	Ways of Retrieving Content	233
9.7	Digifieds Website	235
9.8	Overview of Interacting with Digifieds	236
9.9	Creating Content for a PNA	237
9.10	Posting Content on a PNA.	238
9.11	Retrieving Content from a PNA.	239
9.12	Study Arrangement	240
9.13	Deployment of Digifieds in the Wild	246
9.14	Display Groups Concept	247
9.15	Observations at the Public Library	249
9.16	Field Trial	250
9.17	Distribution of Digifieds Usage Over the Day	251
9.18	Preferred Content	252
9.19	Comparison of Content	253
9.20	Viewer Interest	254
9.21	Privacy Concerns with Regard to Content	255
9.22	Privacy Concerns with Regard to Technology	255
9.23	Content Posted with Digifieds	257
9.24	Sample Images	261
9.25	Standard 10-Digit Keypad	262
10.1	Conceptual Architectures	280

LIST OF TABLES

1.1	Summary of Research Questions.	6
2.2	Comparison of Different Display and Projector Technologies . . .	33
2.3	Average Number of Participants for Different Methods.	61
4.1	Overview of Study Locations	97
5.1	Overview of Comparisons	132
7.1	Comparison of Random / Gaze-based Image Selection	188
8.1	Measures Coded With Our MatLab Software	210
8.2	Total Number of Interactions	211
9.1	Differences in Task Completion Times	244
10.1	Overview of Contributions to Research Questions.	272
10.2	Mixing Information and Ads: Sample Scenarios	282

I

INTRODUCTION AND MOTIVATION

Chapter 1

Introduction

Computing technologies permeate our everyday life. Many devices provide us with ubiquitous access to the Internet as well as with means for computing and communication, hence changing our world. Currently, mobile phones are the prime computing platform worldwide [159], but further technologies are catching up at a rapidly accelerating pace. Tablet computers are a fast growing market [153] and more and more places are being augmented with large interactive public displays. Several major drivers foster this development: technologies and services are cheaply available, computing devices are easy to use, and they provide value to the user [319]. This thesis looks at public displays as one part in the ubiquitous computing infrastructure that is currently emerging, with a strong focus on human-computer interaction challenges. We believe that networked public displays with large screens may become a novel communication medium that can not only implicitly deliver relevant information to the right person at the right time, but also create value to the user through persistent experiences.

One core issue standing in the way of a wide adoption is the question of who will pay for public display infrastructure in the future. Currently, many public displays are stuck with advertising as a business model. On one hand, displays that have been deployed as a replacement for traditional advertising screens simply adapted existing models. On the other hand, displays showing informative or entertaining content (e.g., InfoScreens in stations), also rely upon advertising-driven models as the passerby cannot easily be charged for using the display. Ongoing projects, such as PD-Net (see Section 1.4) are exploring novel use cases, which may enable new business models, such as pay-per-use or subscriptions. The success of these

use cases will rely upon creating meaningful content that goes beyond the mere adaptation of traditional advertisements. For example, content providers will need to anticipate the fact that in many cases, interaction times for public displays are in general very short and thus make content produced for television unsuitable.

In the last years, sensing technologies are becoming commercially available and an increasing number of public displays are networked. This enables new content, but also provides new forms of advertising. These new forms – today referred to as *Pervasive Advertising* – are likely to be adopted quickly by marketers as they have the potential to add a new quality to interaction, audience measurement, and personalization. First and foremost however, it is clear, that this may worsen the negative perception of public displays as an advertising medium, since marketers may easily use these novel means in a way that ignores the viewer's needs and disregards their privacy. In the worst case, this may ultimately lead to a future as sketched by the motion picture 'Minority Report', where all power is in the hands of the advertisers. They clog users with pervasive spam and manipulate them subconsciously to buy products they do not need. In contrast, this thesis empirically investigates building blocks that can help to realize a more positive vision. By applying methods commonly used in human-computer interaction, it explores how networking and sensing capabilities offer novel opportunities with regard to more attractive content and applications for the viewers, whose role may change from passive observers to active users. In this positive vision, users are more likely to be willing to accept, interact with, and pay for the use of public displays – and be it only with their increased attention.

It is widely understood that advertising is likely to stay as the prevailing business model, but that further use cases are emerging for public displays in the future. To earn acceptance, public displays need to provide a benefit for society, for example by providing ubiquitous access to information, emergency services, and exciting experiences. In many ways, this development is similar to that of the World Wide Web, where commercial as well as philanthropic content co-exists and is mutually beneficial. Truly pervasive advertising will use all available channels to communicate with the user [249]. Pervasive displays are one promising option as the effort to establish initial contact is minimal. We expect that new forms of advertising will combine public displays, mobile devices, and social networking into a persuasive, informative, and engaging media for communication between companies and customers. As a foundation to make this vision come true, it is essential to understand the design space for pervasive advertising on public displays from different views. We take two different perspectives – the provider view and the audience view – and identify the core research questions. For each of the research questions, we contribute and evaluate a potential solution.

1.1 Research Questions

Since the era of static advertising displays is ending and networked displays start exploiting the full potential of pervasive computing technologies, the design space needs to be well understood. This requires fundamental challenges to be tackled that currently stand in the way of public displays becoming valuable for society.

The design space is explored in two major steps. In the initial step, the providers are put into the focus (see Table 1.1). As display space becomes more open, there is an increasing need to understand the different stakeholders, their needs, and the interplay between stakeholders, content, and space (R1). Legal agreements between display and space owners regulated what content was allowed on displays in times when these were still operated by outdoor advertisers only. However, as third parties are granted access to public displays, mechanisms need to be established to distribute control among the stakeholders. Furthermore, the display location, new display shapes, and novel interaction techniques strongly impact on the way the audience behaves around public displays. This makes it necessary to rethink the way content needs to be designed in the future (R2). For example, placing a display in a highly frequented subway station might require designing it for passing interaction, as users are not able to stop in front of the display. Finally, sensing techniques enable the user to be identified and content to be targeted or be personalized. Yet, this needs to be done in a way that respects the users' privacy and makes public displays more interesting to them while at the same time providing a benefit for the advertiser (R3).

In the second step, the user will be put into the focus. Whereas they used to be passive observers for decades, we may see a shift of power towards the user. Finally, they can not only perceive but also be immersed, participate, and respond. Hence, understanding interaction challenges in the field of pervasive public displays is key to addressing the fundamental questions.

Before being able to exercise this newly gained power, some fundamental challenges need to be tackled. First, the user's attention needs to be attracted. The fact that public displays have been used as pure broadcast medium over decades lead to a phenomenon today known as the 'display blindness' [255]. Researchers found, that people are not expecting content worthwhile to be perceived, which creates an inherent need to make users aware of the presence of interesting content (R4). Second, users need to understand that public displays are interactive and much more can be done with them than just perceiving the content (R5). Finally, users need to be enticed to interact with the display (R6), which entails not only motivating them to interact but also providing intuitive interaction techniques.

Table 1.1: Summary of Research Questions.

Research Question	No.	Chapter
I. The Provider View		
How do systems need to be designed to cater to the providers' needs?	(R1)	Chapter 4
How does the audience behave in front of public displays?	(R2)	Chapter 5
How can the user be targeted in a meaningful, privacy-preserving way?	(R3)	Chapter 6
II. The User View		
How can the attention of the user be attracted?	(R4)	Chapter 7
How can interactivity be conveyed in an intuitive way?	(R5)	Chapter 8
How can users be enticed to interact?	(R6)	Chapter 9

Note, that many ways exist as to how a design space can be structured. Based on the presented research questions, we opted for the straight-forward separation between providers and audience. While this may seem artificial – particularly because stakeholders could be distinguished in a more fine-grained way – it helps to structure and present the contribution of this thesis in a comprehensive way.

1.2 Methodology

'Public displays' is a relatively young research field, its inception going back to the 1980s. However, it was only in 2000s when dropping computer hardware prices made the field widely accessible and research is finally becoming truly 'public'. As a consequence, no common ground for the design and evaluation of public display applications exists today, as previous work mainly consists of isolated spots in the design space that present solutions to rather specific problems.

In a similar way, 'Pervasive Advertising' has a very recent history. It was in 2008 that the community gathered for the First Workshop on Pervasive Advertising¹, which evolved to become the prime venue for the community over the last years. Pervasive advertising faces the challenge of being strongly tied to different research disciplines, most importantly computer science and marketing, but also psychology, sociology, and communication science. Consequently, the

¹ Pervasive Advertising Workshop website: <http://www.pervasiveadvertising.org>, last accessed February 13, 2013

community initially faced the challenge of bridging a substantial knowledge gap between the disciplines involved. In an attempt to address this challenge, the community put considerable effort into assembling an edited book that introduces the field from different angles [249]. However, as the field is maturing, research on pervasive advertising is still mainly descriptive, similar to public displays.

Due to the lack of commonly accepted guidelines with regard to design and evaluation, the design space presented in this thesis is the result of a bottom-up approach where a set of small to medium scale projects conducted over four years were used to identify and address fundamental challenges. The results are meant to support the design of future applications concerned with advertising on public displays based on recommendations and guidelines. The research that led to our results follows a user-centered design approach. Evaluation in different stages was rigorously conducted in context and constant user feedback helped to identify shortcomings and to iteratively improve the developed prototypes.

1.2.1 Prototypes

All research prototypes presented in the context of this thesis are a result of a close collaboration with my colleagues as well as external researchers, and many undergraduates and student assistants contributed to the development as a part of their work. Of particular value were collaborations with colleagues in marketing and usable security, whose complementary expertise strongly contributed to the prototype design. All prototypes were developed with the ultimate goal to deploy and evaluate them in the lab or in the field, depending on the research question being addressed. The prototypes deployed in the field required significant effort in order to ensure robustness and comply with any technology or privacy limitations. An overview of the prototypes can be found in Section 1.3.

1.2.2 Evaluation

As the name already indicates, public displays are mainly used in public and semi-public spaces. This often makes research a challenging and time-consuming task. This thesis addresses different types of research questions. Some of them require highly controllable environments, for example the evaluation of novel interaction techniques with regard to user performance. Others require public settings in order to obtain ecologically valid data, for example the investigation of audience behavior, social impact, or privacy concerns.

In cases where public displays need to investigate in the wild, researchers have to cope, among others, with the following issues:

- **Outdoor use:** The operation and evaluation of displays in outdoor environments poses considerable challenges with regard to technology and software, hence requiring higher effort than conducting studies in the lab.
- **Location:** Setting up public displays in the desired location is often challenging, as permission needs to be obtained. Furthermore, users should not be in any risk.
- **Privacy:** The evaluation of public displays often requires cameras – not only for the functionality of the prototype but also for logging and observations. This, however, is an issue in public spaces due to privacy concerns and legal restrictions.

Dealing with all these issues requires a lot of effort and can be very costly. As a consequence, we cooperated with different research partners. Some of them already have the required infrastructure, for example UbiOulu in Oulu, Finland (see Chapter 9). Others had the opportunity to install infrastructure on their premises, for example the Telekom Innovation Laboratories in a Telekom store in Berlin, Germany (see Chapter 8).

Furthermore we developed methods to overcome privacy issues when using cameras, e.g., by using depth videos from the Kinect instead of the video stream. This still allowed assessment of what was happening in front of the displays while at the same time preventing identification of the user (see also Chapter 8). During the course of the thesis, a broad set of tools and methods was used for evaluation, including surveys, interviews, manual and automated observations, and logging.

1.2.3 Ethics

Most of the research presented in this thesis was conducted in the context of the European project PD-Net (see Section 1.4). Regulations required a rigorous ethics process to be followed prior to any experiment. This ensured that personal information collected and processed during the study was treated in accordance with the project's ethical guidelines. First, the process required a worksheet to be completed containing information on the study, the involved project staff, aims and methods, and a detailed assessment of the risks and planned precautions,

including a detailed description how data was being stored and processed. Second, this worksheet was covered by one or more study process templates to be considered during the design and while running the study. Sample documents from the ethics process of the Digifieds deployment can be found in Appendix II.

1.3 Summary of Research Contributions

This thesis makes three main contributions to the field of pervasive advertising with a focus on pervasive displays and human-computer interaction: first, we present a design space; second, we report on the development of a set of prototypes; third, we provide conceptual foundations for interactive, pervasive advertising.

1.3.1 Design Space

Exploring the design space, we identified the following fundamental challenges:

- **System design:** Future applications for public displays will require many stakeholders to be involved, including, among others, space owners, display owners, content providers, advertisers, and users. Their interplay, needs, and motivations need to be understood for future systems in order to provide necessary functionality.
- **Audience behavior:** An understanding of how the audience behaves around interactive public displays is crucial to choose suitable locations and content.
- **Targeting:** Advertisers can exploit the full potential of public displays as they allow identification and targeting of the user. This also provides benefits for the users as they are provided with more interesting content.
- **Interaction:** We envision future advertising applications for public displays to exploit different interaction techniques to draw in the user. In order to achieve this goal, the steps required to make the user interact, need to be understood.

1.3.2 Prototypes

This work contributes a set of prototypes used to explore the design space and provide concrete solutions to major challenges. All prototypes were evaluated in lab and/or field studies.

Prototype	Description	Chapter
CarAds	We built a contextual advertising system to use with cars where previously defined advertising campaigns could be shown based on the car's context (e.g., projection of location-based ads on the rear wind shield). The concept was evaluated in the context of a focus group and interviews.	Chapter 4
CAdEt	CAdEt is a context-aware advertising platform that uses sensing techniques to identify users in the vicinity of public displays. This way, information on the user's shopping behavior can be obtained and merged with their self-stated interest to draw a more comprehensive picture of the user. Advertisers can create campaigns with a target group profile that is then used to present suitable content to the user in front of the display.	Chapter 6
AdaptiveAds	Using an eye tracker, we built a web-based prototype that allows a user's attention to be increased by adapting content based on their gaze behavior. We use an HTTP proxy to embed the required code and adaptive content into the website. The approach was evaluated in the lab.	Chapter 7
Looking Glass	Looking Glass is an interactive ball game that uses different representations of the user to communicate interactivity. Representations are a mirror image, a silhouette, an avatar, and an abstract representation. The prototype was deployed in a public setting in Berlin for three months in summer 2011.	Chapter 8
Digifieds	Digifieds is a digital public notice area that allows passersby, event organizers and also professional advertisers to post classified ads. The prototype allowed suitable interaction techniques for public displays to be investigated, potential privacy concerns of the users to be understood, and envisioned content to be identified. Digifieds was deployed in Oulu in the context of the UbiChallenge in July 2011 and is still active as of the time of submission of this thesis.	Chapter 9
GazeAuth	We built a prototype of a gaze-based authentication system that allows the user to securely authenticate in front of a public display. Instead of a PIN pad, we use images with which users can define passwords, consisting of a series of password points. This makes the approach significantly more secure. The approach was evaluated during a lab study.	Chapter 9

1.3.3 Conceptual Foundations

From the results of our studies, we derive a number of lessons learned, guidelines, and design recommendation. We envision both the results and the implications to be useful for the future design of public display applications. The most important contributions are:

- **Design implications for the shared use of public display space:** Our observations allowed us to identify the needs of important stakeholders. We can understand the current processes and structures that help fulfill their goals. We propose a set of principles and concrete ideas as to how they could be applied.
- **Understanding how the audience behaves around public displays:** Results of our studies on audience behavior revealed a set of interesting effects, such as the sweet spot or the landing zone, and provided an in-depth understanding of the honeypot effect and interaction within and between groups. This allows researchers to choose appropriate locations and content for their displays.
- **Making the user interact:** Results from our three studies looking at the different stages of making users interact, provide useful solutions to tackling the most immanent challenges. Specifically, we show that (1) analyzing gaze can be used to increase attention without knowledge about a user's interests, (2) a mirror representation of the passerby can be used to convey interactivity and significantly increase the number of interactions, (3) that privacy, security, and interesting content are crucial to motivate users to interact, and (4) that direct touch and mobile phones are suitable interaction techniques in terms of maximizing usability and preserving privacy.

1.4 Research Context

The research leading to this thesis was carried out over the course of four years at the *University of Duisburg-Essen* (User Interface Engineering Group) and at the *University of Stuttgart* (Group for Human Computer Interaction). Several cooperations with experts from the field in the context of various projects resulted in publications that contributed to this thesis.

PD-Net

A major part of the research was conducted in the context of the European project PD-Net. PD-Net “aims to lay the scientific foundations for a new form of communications medium with the same potential impact on society as radio, television and the Internet”. An overview of the objectives, project partners, and publications can be found on the project website². Collaboration and discussion with all project partners involved provided invaluable input to this work. Of particular importance was the collaboration with the Ubiquitous Computing Group of the *University of Lugano* led by Professor Marc Langheinrich. Joint work with his research team (Nemanja Memarovic and Ivan Elhart) led to numerous project ideas and publications, also beyond the scope of this thesis (e.g., [11; 16; 200; 229; 230; 231; 232]).

UbiChallenge 2011

One part of this work was conducted in the context of the First International Ubiquitous City Challenge (UbiChallenge 2011), hosted by the University of Oulu (Professor Timo Ojala) and the City of Oulu. During the challenge, we had the opportunity to deploy and evaluate the Digifieds prototype on public displays in a real-world setting in the City of Oulu, Finland over the summer of 2011 [10].

Telekom Innovation Laboratories, TU Berlin

During the four years, two joint projects were conducted together with Professor Dr. Jörg Müller from the Telekom Innovation Laboratories at the TU Berlin. In December 2009, we conducted a study on audience behavior around large cylindrical displays, including partners from the LMU Munich (Gilbert Beyer) and Fraunhofer FIRST in Berlin. Results of the study have been published at the CHI 2011 conference [35].

During a six month stay in 2011 with the group of Professor Dr. Jörg Müller, we worked towards a solution for communicating the interactivity of public displays. The concept, development of a prototype and its evaluation in the lab and field, were published in an award-winning paper at the CHI 2012 conference [254]. Furthermore, an edited book on Pervasive Advertising was published in the Springer HCI Series [249]. The book is to a large extent based on contributions of the community to the Pervasive Advertising Workshop Series we organize since 2009.

² PD-Net Project website: <http://pd-net.org>, last accessed February 13, 2013

University of Cambridge

Together with Dr. Andreas Bulling, an expert in eye tracking and usable security, we worked towards a novel, gaze-based authentication technique for public spaces. The approach has been published at the CHI 2012 conference [55].

University of Duisburg-Essen

Together with Dr. Moritz Balz from the *Ruhr Institute for Software Technology*, we taught a practical course on pervasive computing, which led to the development of the first version of the CADEt advertising platform. His expertise in software engineering provided valuable input with regard to implementation and architecture.

Due to the strong relation of this thesis to the field of marketing, we closely collaborated with Julian Mennenöh and Stefanie Kristes from the *Chair for Marketing and Retail* throughout the entire thesis. Their view from the marketing perspective was particularly helpful when we developed our prototypes. Furthermore, Julian Mennenöh's expertise in data analysis and statistics helped us to understand many of the marketing-related effects observed during our studies.

Joint work resulted in an award-winning paper at the AmI 2010 conference [7] and further publications, also beyond the scope of this thesis (e.g., [17; 18; 233]).

University of Zurich

In the context of the ethnographic study presented in Chapter 4 we visited Professor Elaine Huang and Gunnar Harboe at the University of Zurich. To analyze our data, we created a data workflow and extracted findings that lead to a publication in Pervasive 2011 [11].

1.5 Thesis Outline

The thesis consists of ten chapters, which are distributed into five parts. The *Background* part contains an in-depth introduction to public displays and pervasive advertising. It is followed by the two main parts of the thesis. The *Provider View* part tackles provider-relevant challenges, including system design, audience behavior and targeting. The *Audience View* part focuses on the user and addresses challenges related to interaction. At the end of the thesis, the *Conclusion* contains

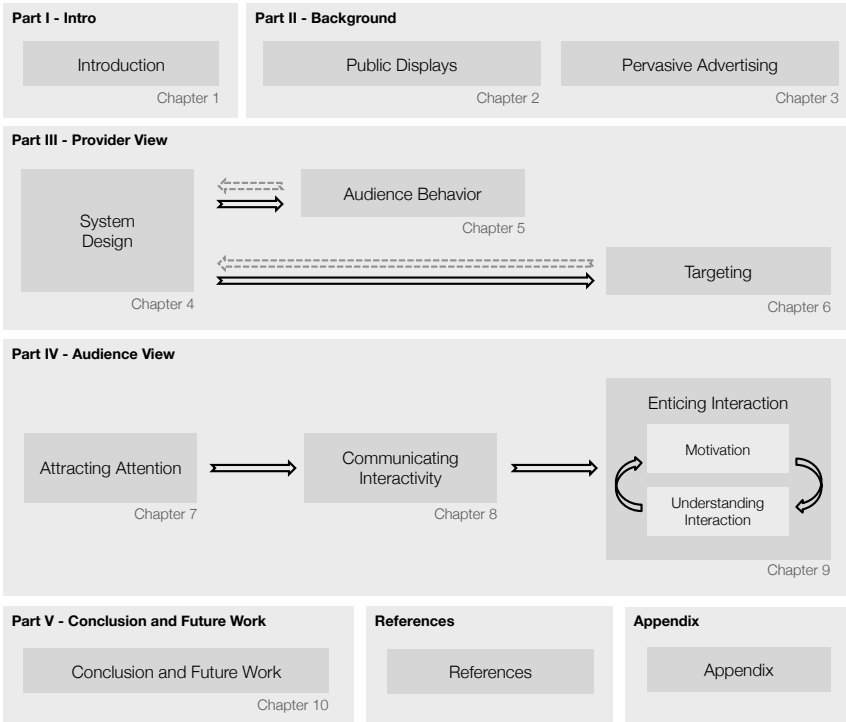


Figure 1.1: Thesis Outline.

a summary of the research contributions and future work. Related work is described within the background chapters (Chapter 2 – Public Displays and Chapter 3 – Pervasive Advertising) and is also integrated into the other chapters. A visual outline of the thesis is depicted in Figure 1.1.

Part II: Background

Chapter 2 – Public Displays: This chapter provides an in-depth introduction to public displays. It starts with a history of public display research, including pointers to the most important work in this field before presenting selected commercial public display systems and deployments. The following section offers a brief overview of display, projector, and sensor technologies.

Then, a closer look is taken at the different dimension encompassing people engaging with public displays. We look particularly at attention (models of attention, attracting attention, quantifying attention) and motivation, pointing out important concepts from psychology. After this, a comprehensive overview of interaction modalities is presented. Furthermore, useful examples for their application are provided. We also look at metaphors used by public display applications as well as at different interaction models. The remainder of the chapter looks at the evaluation of public displays. This section is based on an extensive literature review and our own experience of the past four years.

Chapter 3 – Pervasive Advertising: To bridge the knowledge gap between readers from different fields, this chapter provides an introduction to the field of pervasive advertising. It starts with the definition of important terms, including marketing, advertising, and pervasive advertising. Thereafter, a section introduces advertising objectives. It takes a closer look at how advertising programs are created and provides an in-depth introduction on measuring advertising performance. In this context, state-of-the-art measures and models are presented. The chapter concludes with outlining opportunities and challenges in pervasive advertising.

Part III: The Provider View

Chapter 4 – System Design: Large public displays have become a regular conceptual element in many public spaces, shops and businesses, where they advertise products or highlight upcoming events. In this chapter, we are interested in exploring how these isolated display solutions can be interconnected to form a single large network of public displays, thus supporting novel forms of sharing access to display real estate. In order to explore the feasibility of this vision, we investigated today’s practices surrounding shared notice areas, i.e. such as shop windows or notice boards – places where customers and visitors can put up event posters and classifieds. In particular, we looked at the content posted to such areas, the means for sharing it (i.e., forms of content control), and the reason for providing the shared notice area. Based on two-week long photo logs and a number of in-depth interviews with providers of such notice areas, we present a systematic assessment of factors that inhibit or promote the shared use of public display space. These ultimately led to a set of concrete design implication for providing future digital versions of such public notice areas in the form of networked public displays.

The remainder of the chapter presents a case study looking at cars as future advertising space. With new display technologies, we expect static displays or uniformly painted surfaces (e.g., onto car doors or the sides of vans and trucks) to be replaced with embedded dynamic displays. Furthermore, we see an opportunity for advertisements to be placed on non-commercial cars: results of our online survey with 187 drivers show that more than half of them have an interest in displaying advertising on their cars under two conditions: (1) they will receive financial compensation, and (2) there will be a means for them to influence the type of advertisements shown.

Chapter 5 – Audience Behavior: In the future, interactivity will create an emerging need to understand how the audience behaves around public displays, as this behavior may have a strong influence on which content is appropriate and where such displays can be ideally placed. For example, placing interactive displays next to a street may be dangerous or deploying them in a highly frequented subway station may lead people constantly bump into each other. This problem becomes even more pronounced, as displays can be soon manufactured in arbitrary shapes. As one example of such non-planar displays, we explore cylindrical displays as a possible form of such novel public displays in this chapter. We present a prototype and report on a user study, comparing the influence of the display shape on user behavior and user experience between flat and cylindrical displays. Based on the results, we derive recommendations for designing content as well as interaction and where such displays should be optimally placed.

Chapter 6 – Targeting: One of the major opportunities in pervasive advertising is the ability to target content to the passerby, hence making content more interesting to them. At the same time, this is on one hand, often perceived as an invasion of privacy since data about the user is collected; on the other hand drawing a comprehensive picture of the users is challenging as only their current shopping behavior can be taken into account, but not their personal interests. This chapter presents an approach that considers both information about the user’s self-assessed interest as well as about implicitly gathered information on their behavior, leading to what we call *adaptive profiles*. In conjunction with campaign profiles created by the advertiser for the presented content, a best match can be calculated, even for multiple users, and thus appropriate content can be presented. This chapter provides an in-depth introduction to the approach and presents a prototype implementation.

Part IV: The Audience View

Chapter 7 – Attracting Attention: This chapter addresses the fundamental problem of attracting the user’s attention. As a result of people not expecting interesting content on public displays, users build up a so-called display blindness. Consequently, ways need to be found to increase the attention towards content on public displays. The approach presented in this chapter analyzes a user’s gaze behavior on-the-fly and uses this information to adapt content elements on the screen accordingly. As eye trackers are still difficult to use in front of public displays, we built a web-based prototype and tested the approach in the lab using a standard PC. In an experiment with 12 users, we compared randomly updated content (baseline) with content chosen based on their gaze behavior. We assessed how attention changes the user’s attitude according to the Elaboration Likelihood Model (ELM) and explain these changes.

Chapter 8 – Communicating Interactivity: One major challenge of interactive public displays is, that it is not obvious to the passerby that they are interactive. We designed an interactive installation that uses visual feedback to the incidental movements of passersby to communicate its interactivity. An initial lab study reveals, that mirrored user silhouettes and images are more effective than avatar-like representations. Furthermore, it takes time to notice the interactivity (approximately 1.2 s). In a field study, three displays were installed during three weeks in shop windows, and data about 502 interaction sessions were collected. Our observations show: (1) Significantly more passersby interact when immediately shown the mirrored user image (+90%) or silhouette (+47%) compared to a traditional attract sequence with call-to-action. (2) Passersby often notice interactivity late and have to walk back to interact (the *landing effect*). (3) If somebody is already interacting, others begin interaction behind the ones already interacting, forming multiple rows (the *honeypot effect*). Based on the findings, we derive a set of design recommendations that can be used to build public display applications and shop windows that more effectively communicate interactivity to passersby.

Chapter 9 – Enticing Interaction: To entice people to interact, two major challenges need to be overcome: users need to understand how the interaction works and, they need to be motivated to interact. This chapter first looks at suitable and easy-to-understand interaction techniques before more closely investigating factors that influence a user’s motivation to interact. As a use case, we chose public notice areas (e.g., bulletin boards), which are in their

traditional (paper-based) form very popular and frequently used. Based on an observational study on traditional public notice areas we developed Digifieds, a digital public notice area that allows users to create, post, and retrieve classified ads from public displays. A lab study identified suitable interaction techniques but also revealed considerable privacy concerns, since sensitive data is input to the display. To further investigate these issues, we deployed Digifieds in the context of the UbiChallenge and conducted further field and long-term studies. In the remainder of the chapter, a case study is presented that introduces image-based authentication utilizing gaze as a promising future authentication method.

Part V: Conclusion

Chapter 10 – Conclusion: The conclusion summarizes the contributions made in this thesis. Furthermore, it identifies and discusses potential areas for future work.

II

BACKGROUND

Chapter 2

Public Displays

Public displays have the potential to become the next ubicomp technology to have a major impact on society. They are increasingly saturating our everyday life as computing hardware prices continue to drop and can be found not only in public spaces but also in retail stores, restaurants, bars, and urban administration buildings. Unlike other ubicomp technologies, advertising has established itself as the major business model for public displays. However, advertising does not provide a direct benefit for the viewer, leading to viewers developing a so-called display blindness [255] where they tend to simply ignore public displays, because they do not expect them to be personally relevant or interesting.

Since public displays are becoming interactive, new opportunities and challenges emerge both for the viewer and the provider. Viewers can benefit from ubiquitous information access and engaging user experiences whereas providers can increase the impact by sensing who their audience is and tailoring the content. In order to be compelling to the viewer, content that goes beyond advertising is needed [12]. This, however, requires rethinking the way content on and interaction with public displays needs to be designed and how interactive capabilities can be exploited.

The following chapter provides an in-depth introduction to public displays. It starts with a history of display research, including early art installations, situated and collaborative displays, ambient displays, and novel application domains such as entertainment and advertising. This is followed by a presentation of selected commercial systems and installations. Then, an overview of both display and sensor technologies is provided. The third part looks into engagement

with interactive public displays, introducing the basic concepts of attention and motivation. Next, a comprehensive overview of interaction modalities is presented. These modalities enable implicit and explicit interaction with public displays, ranging from presence sensing via gaze to gestures and touch-based interaction. Afterwards, different metaphors to present content on interactive public displays are explained. Then, an overview of different interaction models is provided. The final part of the chapter looks into the evaluation of public displays. Based on an extensive literature review, it introduces research questions, study types, methods and tools, discusses different approaches to assess validity, and provides a set of guidelines to help inform the design of public display studies.

This chapter is based on the following publications:

- J. Müller, F. Alt, D. Michelis, and A. Schmidt. Requirements and Design Space for Interactive Public Displays. In *Proceedings of the International Conference on Multimedia* (Firenze, Italy), MM'10, pages 1285–1294, New York, NY, USA, 2010. ACM
- F. Alt, S. Schneegaß, A. Schmidt, J. Müller, and N. Memarovic. How to Evaluate Public Displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays* (Porto, Portugal), PerDis'12, pages 171–176, New York, NY, USA, 2012. ACM
- J. Müller, F. Alt, and D. Michelis. Introduction to Pervasive Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*. Springer Limited London, 2011

2.1 History of Public Display Research

The history of research on digital public displays goes back to the 1980s when first artists and later researchers started to experiment with interconnecting remote places through video shown on digital displays and audio. In November 1980, Kit Galloway and Sherrie Rabinowitz created *Hole-in-Space*, an installation connecting sidewalk-facing windows in New York and Los Angeles [115]. This concept of long-lived video (and audio) links that could be used for casual and informal interaction was scientifically explored further by various research institutions operating offices across different cities, e.g., the *MediaSpaces* project [42; 123] at Xerox PARC or the *VideoWindow* [107] at Bellcore Labs.

It was in the early nineties when Marc Weiser and his colleagues developed several examples of display devices and illustrated how displays of different sizes could be embedded into a working environment to support users with solving different tasks [378]. This not only included portable displays for personal use ('tabs' and 'pads') but also large-scale wall-mounted displays ('boards') as situated public or semi-public displays. He envisioned these displays to be networked and hence to collaboratively support users in remote locations. It was also at that time when the idea to implicitly target content to users based on their interests was born. Weiser and Brown propose deploying information displays in the periphery so that users can decide what they want to look at [379].

Inspired by Weiser's ideas, the following decade saw public display systems emerge that delivered content in a situated manner (e.g., the flexible ubiquitous monitor project – *FLUMP* [106]), as well as several approaches to create wearable displays (e.g., the *Thinking Tags* [47], the *Meme Tags* [47], the *WearBoy* [209] and the *BubbleBadge* [101]). At the same time, more and more ambient displays emerged such as the Waterlamp and the Pinwheels used in Ishii's and Ullmer's *ambientRoom* [155], the *OfficePlant #1* [44], and the *Information Percolator* [139]. Greenberg's *DynamicPhotos* [125] aimed at raising awareness and fostering direct interaction between distant people. Rodenstein [304] used windows and foil to provide short-term weather forecasts through projected images and animations.

As display prices started to drop significantly at the beginning of the new millennium, it not only became affordable to install displays in public locations but projections and flatter displays also enabled the displays to be installed on walls and in narrow corridors. For example, Pinhanez's *Everywhere Display* integrated a steerable projector and a camera allowing touch-enabled user interaction on arbitrary surfaces [285]. The early 2000s saw a lot of research on so-called door displays, drawing from the existing practice of leaving notes on doors. The *Dynamic Door Display* allowed both public and private information to be displayed [259]. Similarly, *Outcast* [222] showed personal information and let users leave text messages. *Hermes* enabled visitors to leave pictures, as well as text or stylus-based scribbles [65]. The *Room Wizzard* showed information on meeting room availability and allowed rooms to be reserved ad-hoc [268].

Following research on ambient and abstract information representation in the late 1990s, researchers started to use (flat-panel) displays instead of physical artifacts to convey information. *WebAware* displayed hits on the group's website [297; 343]. The *Hello.Wall* [290; 354] was the first wall-sized ambient display, consisting of LED clusters and conveyed information via light patterns. Interaction was enabled using so-called ViewPorts via RFID.

At the same time, raising awareness and encouraging social interaction through public displays moved to the focus of researchers. *GroupCast* [222] aimed to serve as a conversation starter by displaying information that the people standing in front of it were likely to be interested in. The *Interactive Wall Map* [219] encouraged people to talk about things they experienced in certain places. While these two systems were only deployed in offices, *Ticket2Talk* [224] and the *Opinionizer* [52] fostered conversations at larger gatherings such as conferences. In contrast to the aforementioned system, which aimed at immediate and direct interaction between individuals, the *Notification Collage* [126] and the *Aware Community Portals* [314] tried to create awareness of community activities. In the *Plasma Poster Network*, users were able to post content via email and a web interface to large plasma displays [71]. The *semi-public displays* provided a small set of applications to raise awareness of group activities [148], for example an attendance panel that indicates colleagues that are planning on attending an event. Further projects that looked into increasing community awareness include the *MessyBoard* [103] and the *Community Wall* [124; 346].

Several years later, raising awareness and encouraging social interaction moved again into the focus of researchers, e.g., with Chew et al.'s *Sparks* [68], a conceptual peripheral display system that could be used to support and mediate conversations in a conference setting. *AwareMedia* [27] was a display designed for a hospital environment that provides an overview of staff assignments and locations, as well as surgical interventions.

Further use cases for public displays included interactive kiosks for information access and shared work surfaces. In 2004, Vogel and Balakrishnan presented a semi-public display prototype that allowed information on a display to be accessed through gesture-based interaction [372]. *BlueBoard* [312] provided access to personal information and allowed content to be exchanged between users. *IM here* [150] allowed content to be entered via a web interface and could, for example, be used to notify the members of a meeting that the meeting was about to start. *Dynamo* [161] allowed media content to be shared on a large public display surface.

Towards the end of the 2000s, the mobile phone became an increasingly prevalent device, making it the interaction device of choice for public displays. It allows not only people in the vicinity of public displays to be detected, e.g., based on a Bluetooth MAC address but also enabled them to actively interact with the display. *C-Blink* allowed information to be transferred from the mobile phone to the display encoded in hue differences between subsequent frames of blinking mobile phone displays, which were recognized by a camera integrated into the

display [242]. In a similar way, Sahami et al. used the flashlight of a phone and a display camera to create a novel pointing device [335]. Ballagas et al. detected movements of the phone and analyzed the optical flows from the phone's camera to enable pointing on large displays [25]. Rohs and Gfeller used visual codes that let the user specify which display to interact with [307]. The *TxtBoard*, developed by Kenton O'Hara et al., allowed family members to show messages sent via SMS to the display in a semi-public manner [266]. Further research focused on using Bluetooth as an interaction modality. Cheverst et al. looked at the user acceptance of this interaction technique [65]. Davies et al. explored how the Bluetooth device name could be used to implicitly control the content shown on public displays [83]. Finally, the mobile phone was used to exchange content with the display. *iCapture* used visual codes to select and download content [240], whereas *Digifieds* worked based on QR codes and alphanumeric codes [10].

During that time, researchers also started looking into how interaction with public displays could be enabled in a secure privacy-preserving way. Berger et al. proposed using the mobile phone as a complementary display on which sensitive information would only be displayed to the user [34]. Sharp et al. presented a system that censors sensitive information on the display and shows the uncensored version only on the private display (e.g., a mobile phone) [332]. Furthermore, they presented an architecture for increasing web browsing security on public terminals by hiding confidential information from the terminal [331].

A new emerging application domain at that time was using public displays for entertainment. *Jukola* allowed patrons in a bar to select the music via PDA devices and a touch-enabled public display [267]. *Fancy a Schmink* was a multi-player game where a public display was used to visualize social networks formed by players [300]. *MobiLenin* let diners influence the music choice of a restaurant. A large display was used to display the results of the voting [315]. The *CityWall* in Helsinki let users upload photos to a back-projected, multi-touch enabled display in a shop window to foster 'active spectatorship' [281].

Another application domain for public displays that has emerged is navigation. *GAUDI* (Grid of AUtonomous DIspays) enabled public displays to show personalized navigation information [185]. Stahl et al. presented navigation as a possible scenario of their content scheduling framework for public displays [349]. Rukzio et al. presented the *Rotating Compass*, which combines a public display and feedback via vibrations on the mobile phone [309].

With the focus moving towards the optimization of content selection in the second half of the 2000s, advertising finally moved into the focus of researchers. Müller and Krüger suggested a methodology for creating models based on identifying

users via Bluetooth, gaze recognition, and personal profiles [252]. The *Prospero* project [77] shifted power to the user, making them not only the recipient of content on public displays but also allowing them to shape what is presented on the display as an active contributor.

With further decreasing hardware prices for displays, research-centered display networks emerged. Lancaster University deployed the *e-Campus* system [353]. In 2009, the city and university of Oulu deployed 12 so-called *Ubi-hotspots* [272], which were open to researchers in the following years to deploy their applications in the context of the UbiChallenge [270]. Finalists of the first UbiChallenge in 2011 included (1) *FunSquare* [228], an application that aimed to stimulate social triangulation among the passersby of a display by presenting self-generative, locally relevant content, (2) *CLIO* [303], an application that allowed memories of a city to be shared among the citizens, (3) *RunWithUs* [121], a service designed to encourage people to exercise regularly by sharing information on the display network, and (4) *Digifieds* [10], a digital public notice area that allowed citizens to post and retrieve classified ads from public displays.

From 2009 to 2013, the European FET-Open project *PD-Net* aimed to lay the scientific foundations for public displays as a potential communication medium of the future. In the context of the project, the consortiums initiated and organized the *First International Symposium on Pervasive Displays*³ (PerDis'12) in Porto, Portugal.

2.2 Commercial Display Systems

Apart from prototypes that were developed and deployed in the context of research projects, public displays have long been commercially used. The Motogram – later nicknamed the “zipper” – was deployed on the facade of the New York Times building in 1928 [53; 63]. In the 1960s, so-called split-flap displays were installed as arrival and departure boards at airports as well as in train stations and as electronic scoreboards in stadiums [265; 385]. The first large video-enabled color displays was deployed during the olympic games in Montreal in 1976 [6]. With the advent of projectors and cheap flat-panel displays, an increasing number of displays has been deployed in public spaces [82; 279].

The following table provides an overview of selected commercial display systems:

³ Symposium on Pervasive Displays website: <http://www.pervasivedisplays.org>, last accessed February 13, 2013

Name	Description
PrintSign	Digital signage display with stand; allows “printing” content onto screen Link: http://www.ambientweb.co.uk/sectors/smartsigns/printsign.htm
Sony Ziris	Digital signage suite (Ziris Create, Ziris View, Ziris Manage); supports content creation and scheduling, content transfer, content play-out and monitoring of software and hardware Link: http://pro.sony.com/bbsc/ssr/cat-digitalsignage/resource.solutions.bbsscms-assets-cat-digsignagedev-solutions-Ziris.shtml
Planar Systems / Cool Sign	Solution for managing digital signage networks (CoolSign Manager, CoolSign Network Controller); can be used on Windows PC’s by installing the CoolSign Player software; used, e.g., at Chicago O’Hare airport Link: http://www.planardigitalsignage.com
Dynamax / PointOfView	PointOfView NG2 Player can be used as standalone screensaver (includes scheduler and web-based administration interface); PointOfView NG2 Enterprise Server for managing display networks Link: http://www.dynamaxworld.com/Digital-Signage-Out-Of-Home.html
3M Digital Signage	Digital signage software, mainly hosted by 3M on Internet-based servers (including user interface for administration, scheduling and content distribution), play-out software for single and multi-display deployments Link: http://www.3mdigitalsignage.com/
Netpresenter	Content delivery within organisations (e.g., distributing announcements to employees); content as subscribeable channels (slide sequence); Netpresenter Editor or Powerpoint extension to create content Link: http://www.netpresenter.com/solution/digital-signage
AdSpace Network	Operator of advertising screens used in shopping malls (US); revenue created by selling airtime for advertisements to customers Link: http://www.adspacenetworks.com/
InfoScreen	Operating public displays (mainly projections) and selling advertising airtime (Germany, Austria); schedules a mixture of editorial content (news, trivia, short documentaries), no audio Link: http://www.stroeerdigital.de/infoscreen
Blinkenlights	Non-commercial, artistic installation in Berlin; shows animated light patterns on an eight-story building; user-generated content (patterns) Link: http://www.blinkenlights.net/
BBC Big Screens	Large outdoor screens showing a mixture of programmed content (BBC TV feeds, local communities, artists, live-broadcasts), supports audio Link: http://www.bbc.co.uk/bigscreens/
E Ink	Display technology developed at the MIT; uses arrangement of tiny capsules filled with differently colored and charged particles; high contrast, wide viewing angle, sunlight-visible, flexible, thin; commercially used by Neolux, Midori Mark, and Hamburg city train [31]

Name	Description
Visual Planet	Manufacturer of touch-sensitive foil to convert regular LCD displays or back-projected displays into touch-enabled interactive displays Link: http://www.visualplanet.biz/
SeeSaw Networks	Commercial service working as advertising aggregator; management of advertising space, provides advertising base Example: Life Pattern Marketing (advertisements placed on displays in locations a person is likely to visit during the course of a day) Link: http://www.seesawnetworks.com/
Screenfeed	Manages delivery of licensed content to any public display (content channels, Flash); allows regions of the screens to be filled with new content Link: http://www.screenfeed.com/
Taxi TV	Indoor displays in taxis, showing local news, weather, traffic and advertisements; geo-targeted advertising Link: http://www.verifonemedia.com/vnet/taxi.aspx
ClearChannel	World's largest outdoor advertiser; advertising on different types of digital displays (Digital Outdoor Network, Spectacolor, Transit Shelters, Airport) Link: http://clearchanneloutdoor.com/
Kinoton	Cylindrical LED screens; Litefast MediaPlayer Software allows HD videos and images to be rendered on displays Link: http://www.kinoton.de/en/products-solutions/360-display-systems.html
Dynascan	360 degree indoor and outdoor displays (using spinning LED technology); CMS supporting DVD, Blu-ray, TV feeds and live video Link: http://dynascanusa.com/ds360.html
AdWalker	Mobile advertising display (walking monitors) embedded into a padded sports vest (17 inch LCD screens) Link: http://www.adwalkers.com/
Clo Winebar	Bar in New York that allows drinks to be ordered using touch screens Link: http://walyou.com/bars-and-restaurants-themes-geeks/
Inamo Restaurant	Restaurant in London providing an interactive ordering system; illustrated food and drinks menus are projected on the table surface; allows even taxis to be ordered Link: http://www.inamo-restaurant.com/pc/
Tesco Subway Store	Virtual Subway Store in Seoul that allows passersby to shop with their mobile phones Link: http://www.designboom.com/technology/tesco-virtual-supermarket-in-a-subway-station/

2.3 Technology

Today, a multitude of displays and sensor technologies are available, each of whose properties make them more or less suitable for indoor or outdoor environments. In order for researchers and practitioners to understand which technologies to use when and where, this section briefly introduces various techniques and discusses their advantages and disadvantages.

2.3.1 Displays

Whereas public displays have existed for centuries, e.g., in the form of antique columns still to be found in places like the Hathor temple in Egypt or Trajan's Column in Rome, the age of digital displays began in the late 19th century. The earliest version of the cathode ray tubes (CRT) was invented by the German physicist Ferdinand Braun in 1897. Commercial use of display technologies goes back to 1922, when monochrome CRTs entered the market. Colored CRTs have been available since 1954. More than half a century later, in the early 2000s, liquid crystal displays (LCD) started to gradually replace them. Liquid crystals were discovered in 1888 [344], but it was not until 2008 that the sales of LCD screens surpassed the sales of CRT units. The principle of plasma television was first described by Kalman Tihanyi in 1936, 30 years before the first monochrome video display was co-invented by the University of Illinois and PLATO Computer Systems. In 2010, 18 million units were shipped globally.

Today, there are four main technologies used in displays:

- The *cathode ray tube (CRT)* consists of a vacuum tube that contains three electronic guns (red, green, blue phosphor dots). The electrons are accelerated and deflected by internal or external means in order to create images in the form of light emitted from a fluorescent screen.
- In contrast, *liquid crystal displays (LCD)* do not emit light but use the light-modulating properties of liquid crystals. LCDs have widely replaced CRTs in many applications, since they are available in a wider variety of screen sizes, do not suffer from burn-ins, and are more energy efficient.
- *Plasma Display Panels (PDP)* use small cells containing electronically charged iodized gases, which form plasma when a voltage is applied across the cell. They are bright, can be produced in fairly large sizes (up to 3.8

meters diagonally), and, in contrast to LCD displays, do not suffer from degradation at high viewing angles. However, they consume significantly more power than CRTs or LCDs.

- The *Organic Light-Emitting Diode (OLED)* consists of light-emitting diodes, in which the emissive electroluminescent layer is a film of organic compounds. They emit light in response to an electronic current. OLED displays can use either passive-matrix (PMOLED) or active-matrix (AMOLED) addressing schemes. The difference is that AMOLED requires a thin-film transistor backplane for switching on and off individual pixels, yet it allows for higher resolutions and larger display sizes. Although OLEDs are expected to be available at lower costs in the future and offer wider viewing angles, improved brightness, better power efficiency and fast response times, manufacturing them is still extremely expensive and their current reflectance of around 80% results in poor readability in bright ambient light (e.g., outdoors).

2.3.2 Projectors

Today, we also see projections being used as public displays. The history of the projector goes back to the early 15th century, when Johannes de Fontana sketched a monk holding a lantern with a small translucent window that projected an image onto a surface. From then on, the projector continued to be developed and improved. One important milestone was the invention of limelight by Sir Thomas Drummond and Michael Faraday, from which the first projectors for theaters emerged. It was not until more than a century later, in the mid 1990s, that the first projectors became available for the mass market, which were produced by Texas Instruments using DLP technology. In 1996, Focus Systems presented the first multimedia ultra portable wireless projector, and around the same time, the first LCD projectors emerged. State-of-the-art projectors are based on 3-D projection technology and use lasers as lighting sources.

There are two main technologies currently used:

- *Digital Light Processing (DLP)* projectors create an image through small mirrors on a semi-conductor chip, called Digital Micromirror Device (DMD), where each mirror represents a pixel in the projected image. Light sources used in DLP projectors are usually replaceable metal halide lamps, high-power LEDs, or lasers.

Table 2.2: Comparison of Different Display and Projector Technologies.

	CRT	LCD	Plasma	OLED	Projection	
					LCD	DLP
Weight	--	+	-	++	++	++
Temperature	-	+	-	-	-	-
Power Consumption	-	+	-	-	-	-
Visual Angle	+	-	++	++	-	-
Response Times	++	--	++	++	-	+
Screen Burn-ins	yes	no	yes	yes	yes	no
Dead Pixels	no	yes	yes	no	yes	yes
Screen Size	7-40"	any size	>32"	any size	illumination dependent	

- *Liquid Crystal Display (LCD)* projectors use prisms or dichroic filters in order to separate light to three polysilicon panels (red, green, blue). Individual pixels can be opened and closed in order to allow light to pass or to block it. As a result, different colors and shades can be produced by a combination of different pixels.

Discussion

Table 2.2 provides a brief overview of the advantages and disadvantages of the presented technologies. Most digital signage uses LCD screens today. The major reason is that high ambient light and glare potentially wash out the plasma's image. In contrast, LCDs are prone to very low temperatures. In very cold weather, low response times make them almost unusable for dynamic content, such as videos or animations. Projectors are mainly used indoors, because the brighter the ambient light is, the greater the illumination and the lower the contrast of the projectors needs to be. Furthermore, projectors generally need a different setup to avoid shadows of users standing close to the projection surface (e.g., back projection or high position).

2.3.3 Sensors

The increasing prevalence of smart environments requires the integration of more and more sensors for obtaining information on the environment. The information collected is then processed further and used to modify the environment through the use of different types of actuators. Sensors can either be integrated into the infrastructure (e.g., allowing information to be gathered on weather conditions, traffic congestion, etc.) or personal devices. Mobile phones, for example, now come with many integrated sensors (e.g., GPS, cameras, microphones, accelerometers, and digital compasses) that enable the user to collect individualized data. The following section provides an overview of sensor and actuator technologies.

A range of *optical sensors* (from motion detectors to cameras) is available, which makes it possible to collect very simple (motion-related) but also very complex (human-behavioral) information. The (semi-)automatic analysis of camera images is called computer vision. Today, cameras are so inexpensive (they cost only fractions of a dollar to produce) that they can be integrated into virtually any device. With systems on a chip (SoC), processing power and storage can be directly integrated into the camera. Hence, integrated systems can be built that preserve user privacy, for example, by only outputting the number of detected faces and not the faces themselves. In order to analyze the three-dimensional composition of a scene, stereo cameras are traditionally used. In addition to the normal camera image, a depth map is calculated that provides the distance of all objects in the camera's visual field. Stereovision relies on good features (e.g., textures) detected in the image and verified in both camera images (similar to the human visual system), and requires considerable processing power.

Recently, depth cameras have become available. Depth cameras generate depth maps by illuminating the scene with infrared lighting. Two general technologies prevail: (1) Time-of-flight cameras (e.g., SwissRanger 4000 by Mesa⁴) use technologies such as modulated light sources in combination with phase detectors to measure how long light takes to travel from the camera to the object and back to the camera, and (2) structured light cameras (e.g., Microsoft Xbox Kinect⁵) project a light pattern onto a scene, after which a vision system calculates depth information from the pattern's distortion relative to the objects in the scene. Using depth images, some operations such as background subtraction are much easier to perform than with normal camera images. The recent price decline in depth cameras has also led to a significant number of applications, e.g., gesture control.

⁴ Mesa website: <http://www.mesa-imaging.ch/prodview4k.php>, last accessed February 13, 2013

⁵ Kinect website: <http://www.xbox.com/en-US/kinect>, last accessed February 13, 2013

Audio sensors such as microphones can either provide low-level information requiring only minimal processing (e.g., noise level, base frequency) or high-level information (e.g., speech recognition). Microphone arrays can be used to determine the location of sound sources.

Today, *location sensors* can be used to obtain information on position, collocation, and proximity of users both outdoors (e.g., GPS, GSM, and WiFi) and indoors (e.g., Optitrack, Vicon). Approaches often vary highly in granularity. Indoors, location sensors are typically embedded in the environment, as with the Active Badge system [377]. In the context of advertising, location sensing can be used for tracking the path customers take through the aisles of a supermarket or for mobile displays.

To obtain information on direction, orientation, inclination, motion, or acceleration of a device, many mobile phones now come with *accelerometers* and/or *gyroscopes*. Whereas accelerometers measure proper (relative) acceleration of a device, gyroscopes measure orientation and rotation (using the principles of conservation of angular momentum). Together, these sensors make it possible to accurately recognize movement within a three-dimensional space. Accordingly, different types of contexts can be detected such as whether it is stationary on a table or moving in a car. Acceleration is of particular interest when it comes to analyzing usage patterns.

With the advent of the iPhone, there has also been a proliferation of devices using *touch sensors*. In addition to smartphones, more and more displays and tabletops are equipped with support for (multi-)touch input. Different technologies are used to create touch surfaces. Resistive touch screens use two flexible sheets coated with resistive material, which can register the precise location of a touch by detecting where the sheets are pressed together. For capacitive sensing, a small voltage is applied to a conductive layer, creating an electrostatic field. When a conductor such as the human hand comes close to or touches the surface, a capacitor is formed, and the change in capacitance can be measured from the corners of the panel. Optical touch technologies such as FTIR use light sensors or cameras along with computer vision to detect fingers and objects on and above surfaces. State-of-the-art technologies also include PixelSense⁶ (e.g., Microsoft Surface 2.0), which uses IR sensors integrated into a LCD display to detect activity on top of a surface without using a camera. We refer to Schmidt et al. for further information on sensor technologies [320].

⁶ PixelSense website: <http://www.microsoft.com/surface/en/us/pixelsense.aspx>, last accessed February 13, 2013

2.4 Engagement with the Display

The following section looks at how people behave in front of displays and discusses how this behavior can be quantified. Müller et al. [250] present a model of the different interaction phases, the so-called audience funnel (see also Section 2.7). This model is based on the one presented by Brignull and Rogers [52] and focuses on audience behavior that is readily observable by an outside observer. People must exceed a threshold to transition from one interaction phase to the next. For each pair of interaction phases, a conversion rate can be calculated of how many people are observed passing from one phase to the next, and different displays can be compared using these rates.

In the first interaction phase, people are merely *passing by*. In the second phase, they are *looking at the display* or reacting to it, e.g., by smiling or turning their head. *Subtle interaction* is only available when users can interact with the display through gestures or movement, e.g., when they wave a hand to see the effect on the display. *Direct interaction* occurs when users engage with a display in more depth, often positioning themselves in the center in front of it. People may *engage multiple times* with a display system, either when multiple displays are available or if they walk away and come back after a break. Finally, people can take *follow-up actions*, like taking a photo of themselves in front of the display.

Thresholds exist between the phases that need to be overcome in order to make users proceed to the next phase. For example, not all passersby will look at a display, and not all who look at it will engage in subtle or direct interaction. We propose that the major key to overcome the first threshold (looking at the display) is to attract the attention of passersby. In order to overcome the second threshold (subtle interaction), the curiosity or interest of onlookers should be raised, and in order to overcome the other thresholds (direct interaction, follow-up actions), people must be motivated. All of these thresholds may be increased due to the nature of the interaction in public. Thus, adequate measures must be taken in order to mediate these issues and lower the thresholds.

2.4.1 Attention

Human-computer interaction often assumes that the user is aware of the computer in the first place. This is not necessarily the case for public displays. In contrast to other computing technologies, public displays are not owned by their primary users (the audience). They are installed in public contexts, where they compete

for audience attention with various other stimuli (e.g., other signs, traffic, and people). There has been discussion on how much attention ubiquitous computers should attract. On the one hand, it has been argued that if the environment is filled with ubiquitous computers, they should remain calm and slide effortlessly between the center and periphery of attention [379]. On the other hand, it has been argued that they should engage people more actively in what they do [305]. If public displays fail to attract sufficient audience attention, however, they may not be used at all.

Models of Attention

Generally, the information processing power of the human brain is limited, and at any point in time, a surplus of sensory input arrives at the brain that can not be processed in detail. Attention denotes the process in which the human brain decides which of the numerous sensory inputs to apply the most computational power to. Visual attention is often modeled with a ‘spotlight’ metaphor, in which a certain region of the visual field is selected for more detailed processing. This spotlight often coincides with the fovea, but can change in location and diameter. In general, attention is influenced both by bottom-up processes (external stimuli like a suddenly appearing message) and top-down processes (internal stimuli like the goal of a user looking for a letter in a certain color).

Itti and Baldi [158] present a computational model for bottom-up attention. The sensory input image is split into representations of colors, intensity, and orientations (in the human brain, specialized neurons exist for these representations). From the representations, various feature maps are computed, which are then normalized and combined into conspicuity maps. These conspicuity maps are combined into a single saliency map. In a winner-take-all process, the most salient region is selected and attended to and consequently inhibited. This repeats itself with the next attended region, which will be a different one (inhibition of return). This bottom-up model only takes into consideration the mere sensory input to the brain. Yet, this process is complemented by top-down processes, in which the focus of attention is influenced by the current task, knowledge, and cues. Hamker [130] present an extended model combining bottom-up and top-down processes. In particular, internal goals are modeled to influence the attention process.

In addition to these neuro-computational models, applied models were postulated, in particular to inform human-computer interaction design. Weiser and Brown [379] propose a model of the center and periphery of attention, where users could only centrally attend to one thing at a time but could monitor multiple

things simultaneously in the periphery of their attention. In their proposal for Calm Computing, Weiser and Brown suggest that devices should be designed so that they effortlessly slide back and forth between the center and periphery of attention. They believe that users could thereby attend to more things simultaneously in the periphery of their attention and then take control of them by re-centering them in the center of their attention.

(Not) Attracting Attention

Two general models of what attracts (visual) attention that have been proposed are behavioral urgency and Bayesian surprise. Change blindness can be used in order not to attract attention, and specifically for public displays, the *honeypot effect* has been shown to strongly attract attention.

Franconeri and Simons [112] hypothesize that attention is captured by stimuli that indicate the potential need for immediate action. It has been found that the abrupt appearance of new objects [166] and certain types of luminance contrast changes [97] capture attention. In addition, moving (towards the observer) and looming stimuli have been found to capture attention [112]. Since all of these stimulus properties hint at the potential need for immediate action (e.g., an animal approaching), behavioral urgency may be a useful model to predict how much attention a stimulus will attract.

Itti et al. [156] propose a model of Bayesian surprise for bottom-up visual attention, which measures the difference between posterior and prior beliefs about the world. This is different from Shannon's concept of information, since instead of relying on objective probabilities, it considers only subjective beliefs. They implemented a model of low-level visual attention based on Bayesian surprise to predict eye movement traces of subjects watching videos. The model performs better than other models predicting attention based on high entropy, contrast, novelty, or motion.

Change blindness is an effect that shows how the attention-attracting effect of changes can be avoided. In certain circumstances, people have surprising difficulty in observing seemingly obvious major changes in their visual fields, e.g., road lines changing from solid to dashed or a big wall slowly changing color. Visual effects that cause change blindness include blanking an image, changing perspective, displaying 'mud splashes' while changing the image, changing information slowly, changing information during eye blinks or saccades, or changing information while occluded (e.g., by another person). Intille [154] proposes using change blindness to minimize the attention a display attracts while updating content.

The honeypot effect [52] is described by Brignull and Rogeres in the context of the Opinionizer public display while it was shown at a party. Whenever a crowd of people had already gathered around the display, this crowd seemed to attract a lot of attention, and other people were much more likely to also gather around the display. Similar effects can be observed with the CityWall display [281], as well as with the Magical Mirrors installation [236]. Although attention plays a role for any multimedia system, it plays a crucial role for multimedia on public displays because of the strong competition for audience attention in public spaces.

Quantifying Attention

Quantifying attention is a challenging task, particularly for public displays. Marketing research so far has mainly focused on settings in which high attention occurs, since it is assumed that high attention equates to high recall, which equates to more effective advertising [137]. However, attention towards public displays usually occurs in the order of just seconds [147]. As a consequence, in order to understand the effectiveness of advertisements, research on limited attention spans is crucial. One of the early examples is Dennis's et al.'s work on mall atmospherics [89]. Their work shows that digital signage can convey information at the time and in the place where customers are in the mood to shop⁷ and that it has an effective and entertaining component. In this way, digital signage can support perceptions and emotions in a positive way.

Furthermore, attention is not dependent on a viewer's distance to the screen – in fact, people standing in front of a screen can still easily ignore it [255]. Consequently, attention might be less important than elaboration (or generally involvement), which can be stipulated by increasing the relevance through adapting or contextually personalizing content. In Section 3.2.2, an in-depth discussion on suitable advertising measures is provided.

2.4.2 Motivating Interaction

Traditional paper-based public displays have served as read-only media (e.g., posters and billboards). With interactive displays, users need to be motivated to make use of them and need to be given an incentive for using them. Typically, people do not go out in order to look at public displays. Instead, they tend to come across a public display (e.g., while waiting at a bus stop) and become

⁷ According to Jugger, 70% of all buying decisions are made at the point of purchase [170].

motivated by external factors to look at it. The entry of interactive displays into public space is part of a greater trend: the spread of computer usage from the workplace into public life. While task-oriented theories simply regard the 'how' of an activity and not the 'why', they leave questions concerning underlying motivations unanswered [336].

Malone presents a distinction between tools and toys to differentiate systems that have an external goal from those that are used for their own sake. Tools are task-oriented. They are designed to achieve goals 'that are already present in the external task'. Toys either need to provide a specific goal or enable the user to create their own, emergent goal. A tool should be easy to use, while a toy needs to provide a challenge that motivates the user [213].

In spite of its increasing significance in human-computer interaction, motivation has only been investigated in an isolated manner. There is still a significant need for advancement in understanding the motivation behind the user's activity [212]. Particularly little is known about how the design of public displays will invite interaction [1].

In his studies, Michelis identifies the following building blocks for motivating interaction in public space [236]. His list of motivating factors is based on the work of Thomas Malone, who investigated motivating principles for designing traditional human-computer interaction [212].

Challenge and Control

The first motivating factors, challenge and control, are based on the notion that the ability to master an interaction while still being challenged will increase the motivation to carry out this interaction. Flow [79] is a state of mind where the user is fully immersed in an activity while feeling energized and focused. Simply said, flow can be achieved in a channel between too little challenge (leading to boredom) and too much challenge (leading to anxiety). In human-computer interaction, people strive for an optimal level of competency that allows them to master the challenges presented by the application [50].

The Magical Mirrors study revealed that viewing the consequence of one's own interactive behavior was the most important element for making the challenge a motivating factor. With such displays, the users were motivated to explore and master the interactive functions of the display [236]. Another factor that played an important role was the presence of an emergent goal to the interaction, in which a distinction between set and emerging goals could be made. Emerging goals arose from the interaction of the individual with the Magical Mirrors displays.

Since emerging goals have a strong motivating effect, interactive environments should not only provide a set of goals but also allow for the design of one's own emerging goals [50]. Moreover, the intrinsic motivating challenge of an activity appears to increase if clear and direct feedback is given following one's own behavior as well as the attainment of the goal [212].

The results of the Magical Mirrors study support the importance of emergent goals for motivation [236]. In order to turn an interaction into a challenge, the behavioral outcome should, however, be somewhat uncertain and the end result should be unknown. The motivating effect of control is based primarily on recognizing a cause and effect relationship and the freedom of choice in performing the interaction. For motivating users, the perception of control is more important than actual control and can even have a motivating effect in the absence of actual control [5].

Curiosity and Exploration

As one of the most important foundations for intrinsically motivating behavior, curiosity is evoked through novel stimuli that present something unclear, incomplete, or uncertain. This causes users to search around for possible explanations within their environments, and their behavior is motivated by a desire to avoid potential insecurities.

Curiosity is described as a precursor to explorative behavior, through which people make accessible previously unavailable information about their environment. Explorative behaviors allows people to utilize exploration as a means to avoid insecurities [88]. Specific explorations are attempts to reduce the degree of incongruity and therefore the level of stimulation. However, if the stimulation falls below an optimal level, the individual is motivated to make further explorations in order to re-establish the optimum.

Curiosity appears to belong to the most important characteristics of intrinsically motivating environments. In order to stimulate curiosity and to influence motivation, the interaction should not be designed in a way that is either too complex or too trivial. Interactive elements should be novel and surprising but not incomprehensible. On the basis of his or her prior experiences, the user should have initial expectations for how the interaction proceeds, but these should only be partially met [212]. Thus, in reactive environments, a motivating optimum of complexity is also fostered through the interplay of surprising and constructive interaction. The desired behavior for the interaction can initially be activated by surprising elements and maintained through constructive elements. In contrast to

perceptible changes that appeal to people's sensory curiosity, cognitive curiosity relates to anticipated changes. People are motivated in this way to optimize their cognitive structures [244]. To increase motivation through curiosity, it appears at first sufficient to convey to the individual a sense of incompleteness, discrepancy or dissipation and to present, through the interaction, the chance to abate these sensations. However, during the interaction, it should be made especially clear how completeness can be attained [212].

Choice

Choice as a motivating factor is based on the observation that the motivation for a behavior appears to increase if, in the process, people can select between alternatives in behavior. This choice enables them to control their behavior and to make active decisions regarding behavior for the individual situation. Preferable are those alternatives that best correspond to one's own preferences and through which not only the behavior itself but also its effects can be controlled [160].

With an increase in the number of possible choices, the likelihood increases that a feeling suited to the individual can be found. Even with very trivial choices or ones that only exist in the imagination of the individual, a motivating effect was clearly proven [50; 78; 160]. Given that the mere presence of choice appears to promote intrinsic motivation, it can therefore be established that the sensation of autonomy and control increases as a result. The greater the number of choices perceived, the stronger one's own autonomy and control appears to be.

On the other hand, it was demonstrated that a number of alternatives that exceed an optimal level [160], as well as the absence of choice and opportunities for control [176], lead in various ways to a reduction in intrinsic motivation. In summary, the offer or presence of interaction alternatives can be a strong motivating factor and encourages the performing and maintaining of specific interactions [364]. This could also be seen in the Magical Mirrors study [236].

Fantasy and Metaphor

In general, imaginary settings also appear to have a motivating effect on behavior. In these fantasy settings, the constraints of reality are switched off so that one imagines possessing new abilities. In interacting with computers, one of the initial user reactions is oftentimes the inspiration of fantasy. The extent to which interactive environments inspire fantasy determines their attractiveness and generates interest in the reception of the interaction [277]. The use of metaphors allows for operationalizing fantasy concepts [41].

By employing metaphors, fantasy elements can be directly integrated into human-computer interaction. Since they usually refer to physical systems, metaphors can help the user to comprehend the interaction prior to actual use, motivating them toward the reception of the interaction [177]. In the Magical Mirrors study, the metaphor used was the distortion mirror known from annual fairs and amusement parks [239]. Since the interaction bears resemblance to already-known situations, it can be grasped more easily and utilized more efficiently. By doing so, metaphors do not need to reproduce the world realistically, since the abstract, conceptual, or symbolic representation can prove equally effective as live images [187]. The significance of metaphors in human-computer interaction is supported by a series of research projects. If new forms of interaction are linked to familiar traditions, it appears easier for users to carry over already-established behaviors.

Collaboration

In contrast to the first motivating factors, collaboration is based on the interaction with other human beings. A condition for its motivating effect is the opportunity that the individual can influence the interaction of other people [50]. This also appears to apply when multiple individuals engage in communal activities via the use of computers. With the linking of computers via the Internet, human-computer interaction was also expanded around a social component [91].

In addition to social interaction over the Internet, the use of interactive public displays increasingly plays an important role in collective interaction located in one place [204]. The motivation to collaborate is increased, for example, through functionalities that make the effects of one's own behavior visible. With a view toward cooperation and competition, differences can be ascertained between the social value orientations individualistic, cooperative and competitive. While people with a cooperative orientation also hold the preferences of others to be important, people with a competitive orientation seek to maximize their own preferences in relationship to the preferences of others. In this case, collaboration is especially motivating if individual behavior is recognized by others [374]. If the efforts and effectiveness of one's own behavior are recognized and valued, people are motivated to repeat this behavior again. If the collaboration is continued, the probability of sustained recognition is even greater.

The visibility of one's own behavior is also one of the most important foundations for recognition [213]. The degree to which collaboration has a motivating effect is influenced by the personal experience of the individual and can strongly vary according to each particular situation. Alongside individual orientation, cultural differences also play a role [168].

2.5 Interaction Modalities

State-of-the-art sensing technologies enable a wide range of modalities for implicit and explicit interaction to be supported in front of public displays (cf., Section 2.3.3). The following section provides a comprehensive overview of these modalities and presents selected examples.

2.5.1 Presence

A wide variety of sensors allow the presence of an audience to be sensed or even be identified in the vicinity of a display, for example based on cameras, microphones, Bluetooth and RFID scanners, or pressure sensors. Presence sensing can be used to trigger implicit interaction, often with the aim of getting the user involved into interaction with the display.

Hello.Wall [354] is an ambient display that reacts to people as they pass by. The installation uses RFID-based ViewPorts carried by users for identifying them in the wall's proximity and triggers the emission of information via light patterns on the wall. Other systems that use RFID for identifying the users are BlueBoard [312] and AutoSpeakerID [224].

Mahato et al. use the Bluetooth-friendly name of mobile phones to encode users' interests in order to implicitly personalize the environment [210]. In a similar way, Davies et al. use the Bluetooth-friendly name as a unidirectional channel to control the environment, e.g., by encoding a search query whose result is then shown on a nearby display [83]. Alt et al. use Bluetooth for implicitly generating profiles that could then be used in pervasive advertising environments [7].

The Thinking Tags [46; 48] and Meme Tags [47] use infrared to identify communication partners and exchange information. Infrared badges are used to identify display users in the Interactive Wall Map [219] and in GroupCast [222].

2.5.2 Body Position

Cameras or pressure sensors in the floor can be used to identify presence and exact body position of a person in front of a display. Knowing the position allows the system to provide finer-grained interaction and to deliver content more precisely, e.g., by displaying or updating content close or in relation to the user's position.

Beyer et al. [35] use a camera to determine a user's position in front of their cylindrical screen and then encourage the user to interact with content that follows the user as they move around the column. The Community Wall uses a grid of infrared-based movement sensors to track users' positions in real time [124].

2.5.3 Body Posture

Body posture can be used to estimate the direction a user is facing. In combination with proximity sensing, the path a user takes while approaching the display can be determined. Technical solutions for measuring body posture include motion tracking, 3-D cameras, and low-frequency waves [384].

Vogel and Balakrishnan present a public ambient interactive display [372]. This system uses cameras to determine the transition between implicit and explicit interaction based on which interaction phases are defined. Similarly, Sawhney et al. use cameras in order to find out whether people approaching the display were looking, glancing, or simply passing by [314].

2.5.4 Face Detection and Face Recognition

Today, a variety of software and hardware is available for analyzing video and extracting information about the users based on their faces. In the simplest case, a system can detect faces, which allows the number of people in front of displays to be counted. Sophisticated systems can extract more fine-grained information such as age, gender, or even the mood of a person, allowing for more specific user targeting. Finally, face recognition allows users to be identified, thus enabling personalization of the content shown on the display.

Fraunhofer's SHORE includes means for detecting whether a user's mood is happy, sad, surprised, or angry [191]. The eMir system classifies facial expression in order to encourage interaction with a public display [100]. Grasso et al. use face detection based on real-time video analysis in order to identify the user [124]. Commercial face detection solutions include Samsung PROM⁸, which provides statistical information on the number, gender, and age group of users.

⁸ Samsung PROM website: <http://www.samsunglfd.com/solution/feature.do?modelCd=Samsung%20PROM>, last accessed February 13, 2013

Face recognition systems are widely available now. Airports support immigration procedures [198], commercial solutions are offered, e.g., from NEC⁹, and mobile apps such as Recognizr¹⁰ enable users in a camera's visual field to be recognized. The latest Android phones such as Samsung's Galaxy Nexus can also use facial recognition to unlock the phone [43].

2.5.5 Gaze Direction

Camera-based technologies such as eye tracking allow different aspects of a user's gaze behavior to be precisely assessed. For example, knowledge about the user's gaze direction can be used to measure whether they looked at a public display. While coarsely-grained gaze information is mainly useful for audience measurement, knowing the exact gaze position can enable both explicit interaction and implicit adaptation of the content (see Chapter 7).

A rough estimation of the user's gaze direction can already be achieved with a simple camera [217; 258]. Similarly, specialized devices that are able to detect eye contact can be used to determine whether users looked at a certain object, e.g., the Xuuk EyeBox2¹¹. Mubin et al. use the EyeBox2 to adapt product lighting in a smart shop window based on the user's gaze direction [245]. ReflectiveSigns [251] uses gaze detection to learn about audience content preferences.

2.5.6 Speech

Microphones in the vicinity of a display can not only be used to sense keywords of ongoing conversations (e.g., allowing advertisements to be targeted) but also enables an estimation of the number of people close by. Based on this information, content can be adapted implicitly or voice commands can be used to let users explicitly control the content on a public display. One example of using microphones is LaughingLily, an ambient display that is able to detect the level of activity in a meeting situation [20].

⁹ NeoFace Facial Recognition: <http://www.nec.com.au/solutions-services/neoface-facial-recognition.html>, last accessed February 13, 2013

¹⁰ Recognizr website: http://www.readwriteweb.com/archives/recognizr_facial_recognition_coming_to_android_phones.php, last accessed February 13, 2013

¹¹ Xuuk Eyebox2 website: <https://www.xuuk.com/Products.aspx>, last accessed February 13, 2013

2.5.7 Gestures

Gestures have been subject to research for many years. While several technologies enable gesture interaction (accelerometers, touch sensors, mouse, and gaze-tracking), cameras are the most popular among public displays. Hand gestures are used for indirect explicit interaction, e.g., for manipulating objects or controlling the screen.

The Pendle [371] is a gesture-based wearable device that integrates environment-controlled implicit interaction and user-controlled explicit interaction. Vogel and Balakrishnan present an overview of gesture-based interaction techniques for public displays [372].

2.5.8 Keyboard and Mouse

The aforementioned modalities are often not understood (at least not easily) at first glance, especially when it comes to explicit interaction. In contrast, a standard *keyboard or mouse* provides an easy means for enabling interaction with a public display.

The Opinionizer [52] looks at how people socialize around public displays. They deliberately used a keyboard as an interaction device to avoid any obstacles for using the system. In Prospero [77], a keyboard was used as one possible form of authentication with username and password.

2.5.9 Touch

Although touch interfaces have been available for many years, their popularity increased with the advent of the iPhone and other mobile multi-touch devices. In public displays, touch sensors enable direct interaction. Users can explicitly interact with the screen by manipulating objects.

CityWall [281] allows multiple users to simultaneously interact with a large touch-enabled display. Boring et al. present the Touch Projector [45], a system that allows users to interact with remote screens through a live video image on their mobile phone.

2.6 Metaphors

Understanding how users intuitively perceive the world around a display is essential for the design of interactive display applications. Hence, we present the results of an analysis of existing public display applications. We identify four prevailing mental models used as interaction metaphors.

2.6.1 Posters

By definition, a poster is a piece of printed paper (including text and graphics), which can be attached to walls or vertical surfaces. Many digital public displays follow this metaphor, adapting content that was originally designed for their analog counterparts. However, more and more digital public displays are enhanced with sensing capabilities, allowing people in the vicinity to implicitly or explicitly interact with the content. One example is CityWall [281]. The screen is deployed in a shop window and enables multi-person interaction. The research focuses on phenomena arising from public deployment, e.g., parallel interaction, conflict management, and gestures.

Interactive content on public displays following the poster metaphor aligns well with what users currently expect from public displays, which is being digital counterparts of traditional posters. However, this leads to the tendency for displays to be ignored by users due to the association with traditional advertising posters [255]. Hence, approaches following the poster model face the challenge that they need to put a special focus on grabbing a user's attention.

2.6.2 Window

Following the window metaphor, this mental model creates the illusion of a link to a remote and often virtual location. In contrast to the poster, windows enable two-way interaction: users can look inside and parties on the remote side can look outside.

As displays were first used in the 1980s to (virtually) connect physically separate spaces, many of them used the window metaphor. One of the most prominent installations was *Hole-in-Space*, that used the window metaphor to create a link between two display windows [115]. Later, more long-lived installations based on full-duplex audio and video connections were used in the context

of ‘media spaces’ to connect research offices in different cities (e.g., Xerox PARC) [42; 123]. Similarly, the VideoWindow [107] was deployed at the Bellcore Labs and connected remote conversation partners via near life-sized displays, leading to almost face-to-face-like interactions. Such a window metaphor can be extended to other modalities like punching. Remote Impact [246] allows two remote players to enter the same interaction space.

2.6.3 Mirror

Mirrors in the real world are objects with a reflective surface. Several research projects follow the metaphor of a mirror to encourage interaction. For example, the Magical Mirrors present its audience a mirror image of themselves augmented with optical effects like a ribbon following their hands [236]. Other installations embed users within a different context, e.g., a scene at the beach or on top of a mountain. Schönböck et al. show that making users a part of the display has a strong potential to catch their attention as they pass by [323]. Müller et al. use a mirror metaphor to convey interactivity in a quick and easily understandable way [254].

2.6.4 Overlay

Finally, projectors enable the creation of overlays, which allow content to be displayed within another context. In contrast to the aforementioned models, overlays are frameless in that they can seamlessly integrate with the environment. Pinhanez present the Everywhere Displays Projector [285], an LCD projector that allows images to be projected onto different surfaces of an environment. One application, the Jumping Frog, presents a frog on any surface in the environment. If somebody tries to touch the frog, it ‘jumps’ to another surface nearby.

2.7 Interaction Models

The following section looks at interaction models that describe the behavior of the audience in front of public displays. Such models can be useful during the design process, for the setup, and to measure the effectiveness of a public display application.




Figure 2.1: Interaction Models for Public Displays: (a) Three Zones Model (Streitz et al. [354]); (b) Extended Model (Vogel and Balakrishnan [372]); (c) Public Interaction Flow Model (Brignull and Rogers [52]); (d) The Audience Funnel (Michelis and Müller [238]).

In the context of Hello.Wall, Streitz et al. [290; 354] present an interaction model based on three zones: an ambient zone, a notification zone, and a cell interaction zone (Figure 2.1a). Vogel et al. [372] refine this model by separating the interaction zone into a subtle and a personal interaction phase and by generalizing the notion of the notification zone into an implicit interaction phase (Figure 2.1b). As a neutral state, they define a so-called ambient display to be an anchor point for subsequent interactions. This model is strongly geared towards information presentation and is therefore mainly suited to model single-user interaction.

Brignull and Rogers [52] present a model of public interaction flow. Deploying a system called Opinionizer at a book launch party and at a welcome party allowed them to explore how groups socialize around public displays, how they walk up to them and how they change roles, i.e., change from an onlooker to an active participant and vice versa. In their first step, three activity phases were explored: the peripheral awareness activity, the focal awareness activity, and the direct interaction activity. Understanding these activities allowed the

conceptual framework of analyzing public interaction, depicted in Figure 2.1c, to be developed. The strength of the framework can be seen in that it (1) supports multi-user interaction, and (2) takes people moving between the different activities into consideration. In contrast, it disregards both implicit interaction and explicit interaction from a distance (probably due to the use of keyboard and mouse as input devices). The study seems to be also limited by the fact that many people (at least briefly) knew each other (which might be an explanation for the strong observed effect of social embarrassment), preventing triangulation effects between strangers from being modeled.

Michelis et al. [238; 250] present the audience funnel, which builds upon the *Public Interaction Flow Model* and focuses on observable audience behavior (Figure 2.1d). It consists of several interaction phases, for each of which a conversion rate can be calculated as the user moves to another phase. The phases are: passing by, viewing and reacting, subtle interaction, direct interaction, multiple interaction, and follow-up actions. Between the different phases, certain thresholds exist that need to be overcome in order to enter the next phase. To overcome the first threshold and transition from ‘passing by’ to ‘viewing and reacting’, the passerby’s attention must be captured. To overcome the second threshold and move on to ‘subtle interaction’, the onlooker’s curiosity must be piqued. Subsequent thresholds can be overcome by motivation. The strength of this model is that it can be used to calculate conversion rates and thus provide a measure of success for public display content or applications. This is especially interesting to advertisers. However, the model does not consider social interaction and describes a rather linear process where people are unlikely to move back and forth between the different stages.

This overview of interaction models shows that no comprehensive model exists. Interaction is strongly depending on a number of different aspects, e.g., the supported interaction techniques, the content, the audience, and the envisioned type of interaction.

2.8 Evaluation Techniques

Interactive public displays are leaving the labs and are being deployed in many places. They already permeate public spaces, shop windows, malls, workspaces, and public institutions, and they are equipped with sensors such as cameras, which enable presence and motion sensing. At the same time, new (consumer) devices

and software enter the market (e.g., the Microsoft Kinect or Leap Motion¹²), providing opportunities for researchers to create novel interaction techniques. Hence, there is an emerging need for both practitioners and researchers to understand how to best evaluate public displays with regard to effectiveness, audience behavior, user experience and acceptance, and social as well as privacy impacts.

Since no commonly accepted guidelines currently exist for how (applications for) public displays should be designed, the evaluation of these is both crucial and challenging.

- **Display deployments are often opportunistic.** As new infrastructure or real estate is created, the premises are often augmented with public displays with only little knowledge of the audience.
- Simulations of the environment in which a display is deployed are difficult, because there are **no (dynamic) models** yet (e.g., of the stream of visitors passing through a pedestrian area). As a result, evaluation has to be conducted in context (in both the real world and in the lab).
- There is **not one single goal** that public displays (or their content) try to achieve. Ads most likely strive to maximize public attention, interactive games may want to create an engaging experience, informative applications such as a public transport schedule may aim at maximizing usability, and some displays may be deployed to show warnings to passersby or support the fast evacuation of a building.
- **Measuring the effectiveness** of a display is difficult. Unlike in the Internet, it is often not possible to monitor user interaction, but sensors might allow richer information based on the interaction to be extracted in the future. Using sensors, however, might raise privacy concerns (e.g., when using a camera in public space) and thereby restrict the means for evaluation.

In order to tackle these challenges, this section provides guidelines for evaluating public displays. Our work is grounded in a comprehensive literature survey. Based on this, we identify common study types, paradigms, and methods, including their respective advantages and disadvantages. We discuss and validate them, ultimately deriving guidelines that can help researchers and practitioners choose the best evaluation method for their public display.

¹² Leap Motion website: <http://leapmotion.com/>, last accessed February 13, 2013

2.8.1 Previous Work

As of January 2012, 522 papers can be found in the ACM Digital Library concerning public displays or digital signage. Most of these papers evaluate their own concepts and deployments. Even now, more than 30 years after Hole-in-Space, neither design guidelines that cover a broad spectrum of systems and applications nor generally accepted evaluation guidelines exist for public displays. However, several ideas have surfaced in recent years.

Cheverst et al. report on challenges of evaluating situated displays deployed in a community setting [67]. Storz et al. published lessons learned from the deployment of the eCampus public display network [353], which provides useful information for informing the design of public display (networks) but only little information with regard to evaluation. Mankoff et al. look at the evaluation of ambient displays, focusing mainly on effectiveness and usability [214]. Starting from Nielsen's usability heuristics, they create a modified set to be used for the evaluation. Finally, Matthews et al. use activity theory to evaluate peripheral displays [218]. They identify an initial set of evaluation metrics (appeal, learnability, awareness, effects of breakdowns, and distraction) that vary depending on the importance of the display but did not focus on evaluation methods.

Though many research papers provide useful lessons learned or recommendations based on their findings [147; 353], most previous work focuses on a rather specific application domain (e.g., community/situated displays and ambient displays), draws conclusions only from findings of their deployment(s), or treats evaluation only on the side or on a high level. To overcome these limitations, we base our findings on a comprehensive literature review, identifying research questions, research types, research approaches, and methods used in public display research.

2.8.2 Research Questions

In the following section, we describe the most popular questions that researchers tried to answer during their evaluations. Note that many projects tackled subquestions to these questions (e.g., numbers of glances or interactions as subquestions to display effectiveness).

- **Audience behavior:** A major focus is how the audience behaves around a display. Prior work identified effects such as the honeypot [52; 254] (interacting users attract more users), the sweet spot [35] (a preferred position in front of the screen), or the landing zone [254] (people only realize

that a display is interactive after passing by). Though mostly conducted in the real world, we also found examples of lab evaluations [35]. Audience behavior can be assessed by observations [104] and log data [306].

- **User experience:** User experience describes the overall experience of interacting with a display. This is important in public display research, since good user experience may lead to a higher motivation to use the application and possibly draw the user in for as long as possible. Researchers examined different interaction techniques and their effect on user experience based on (standardized) questionnaires, e.g., interactions mediated through a mobile device [25] or direct touch [296].
- **User acceptance:** Often used in early stages of the development process, user acceptance investigates users' motives and incentives for interacting with a display. It can be assessed qualitatively based on subjective feedback, e.g., in focus groups to collect the target group's view and concerns [67] or quantitatively based on questionnaires [235].
- **User performance:** Effectiveness from a user's perspective is often measured when evaluating novel interaction techniques, e.g., based on a camera, mobile phones [25] or direct touch [76]. User performance can be quantified by measuring task completion times and error rates [25; 76; 321].
- **Display effectiveness:** Interesting from an economic perspective, several studies aim at measuring the effectiveness of public displays (e.g., how many people passed by a display [251], how many looked at it [147; 255], and how many started interacting [254]).
- **Privacy:** Some projects aim to understand the users' privacy concerns. Alt et al. looked at how mobile phone can overcome privacy issues [10]. Shoemaker and Inkpen explore an interaction technique that allows private information to be shown on a shared display [337].
- **Social impact:** Social impact is subject to a great deal of research. Researchers have evaluated how displays could foster social interaction [228], how users engage in social interaction [254], which types of communities form around displays [11; 67], and which social effects occur [51; 281].

In Appendix I, a summary of research projects classified by research questions can be found. We distinguish between evaluations that were conducted prior to creating a prototype by asking users or running an ethnographic study and those that evaluated a prototype in the lab, field, or in the context of a deployment.

2.8.3 Approaches to Research

The following section provides an overview of different study types and paradigms we found throughout the literature review. They will be briefly explained, followed by a discussion of their advantages and disadvantages. In our own research, we adhere to the notions presented by Nielsen [260] and Lazar et al. [202].

Study Types

We categorize related work according to the three principle types of research: descriptive, relational, and experimental research.

Descriptive Research

Descriptive research aims at merely describing what is going on in a certain situation. This description can be qualitative (e.g., observations [255], interviews [11], and focus groups [14]) or quantitative (e.g., photo logs [11]). In descriptive research the variables do not need to vary, e.g., multiple prototypes for comparison are not needed. It is striking that the vast majority of public display research includes descriptive methods. Good examples are the CityWall [281] and Worlds of Information [164]. In both studies, a single prototype is deployed and user behavior around the display is measured, analyzed, and described (with, e.g., observations and questionnaires). One major benefit of descriptive research is that no hypotheses need to be tested, and therefore, the hypotheses do not need to be derived from a general theory. It is especially suited for a research field like public displays that is in an early phase and does not possess general theories yet. However, descriptive studies of single prototypes create isolated spots in the design space of public displays with no relation to other studies. This makes it difficult to compare results and designs, and ultimately to understand the structure of the entire design space. Hence, in the long run, the progress of public display research may be hindered if it continues to focus purely on descriptive studies.

Relational Research

Relational research aims at showing that two or more variables covariate, i.e., that there is a relation between two or more factors. Relational research does not, however, indicate causality, i.e., it is unknown which of the variables causes the other to change or whether both depend on a third, unknown variable. Relational studies are rare in public display research, in particular because not many relationships between different dependent variables are considered to be interesting. Exceptions include ReflectiveSigns [251], where it is shown that the time people spend looking at display content does not correlate with people's stated interest.

Experimental Research

Experimental research aims at determining causality, i.e., that one variable directly influences another variable. Experiments possess the following characteristics: they are based on hypotheses, there are at least two conditions, the dependent variables are measured quantitatively and analyzed through statistical significance tests, they are designed to remove biases, and they are replicable [202]. Experiments aim to (fail to) refute hypotheses, and these hypotheses are usually derived from theories. Therefore, entire theories can be refuted by experiments. Experiments can be conducted in the lab (more control) or in the field (higher ecological validity). Whereas a lot of experiments have been conducted in the lab (e.g., in order to evaluate user performance with regard to a novel interaction technique [25; 45; 76; 179]), real-world experiments are rare in public display research, partially because no coherent theories of public displays exist. Another reason is that multiple variations of a prototype would need to be developed, making such experiments in the real world particularly time-consuming. One example is Looking Glass [254], where the influence of different interactivity cues on how many people interact with the displays were tested.

Research Phases

Research methods can be used in different phases during a project. In the beginning, there is usually a phase of *requirements analysis*, and there is no prototype yet [11; 14]. Typical methods used during requirements analysis encompass ethnography and techniques to question users such as focus groups, interviews, or questionnaires. After the first prototypes are developed, a phase of *formative studies* usually follows. These are intended to give direction to the design process and find properties and problems of the current prototype. Formative techniques include deployment-based research, lab studies, and techniques to ask users [10] (see below). When the final prototype exists, *summative studies* come into play. They usually try to make some conclusion about the final prototype, e.g., by comparing it to a baseline in a lab [35; 254] or field study [254].

Paradigms

We identify five evaluation paradigms used either to inform the design of a prototype (ethnography, asking users) or to evaluate a prototype (lab study, field study, deployment-based research).

Ethnography

In ethnographic studies, usually certain (social) settings are investigated without intervention, e.g., without deploying a prototype. Ethnographic studies have been used to inform the design of public display systems. Alt et al. conducted an ethnographic study to assess the motivation and intentions of stakeholders as well as social impact [11]. Huang et al. investigate current practices around public displays [147]. The advantage of ethnographic studies is that they provide valuable information that could be used to enhance the design of a public display (system). The disadvantage is that they often require a lot of effort.

Asking Users

Similar to ethnography, users can be questioned using interviews, questionnaires and focus groups. Yet, hypothetic questions usually lead to poor answers (e.g., “*What functionality would you expect from a future version of this app?*”). Therefore, users are often also prompted with a prototype, and these methods are combined with lab or field studies, ethnographies, or deployments.

Lab Study

Lab studies aim at evaluating a system within a controlled environment. Lab studies can be descriptive, relational, or experimental. During the lab study both qualitative data (e.g., interviews [254] and observations [35]) and quantitative data (e.g., task completion time and error rates [25; 45]) can be collected. The advantage of lab studies is that external influences (such as other passersby, environmental conditions) can be minimized and (sensitive) equipment for proper measurements (such as cameras and sensors) that would be difficult to deploy in public can be used [35]. The disadvantage of lab studies is that they may provide only low ecological validity and that the dynamics of the real world are excluded.

Field Study

In contrast to lab studies, field studies aim at evaluating a system or technique in a (semi-) public setting. In contrast to deployment-based research, they are rather short (days to months) and focused on a single research question. Similar to lab studies, they may be descriptive, relational, or experimental. Data collection in the field is often cumbersome and time-consuming, since automation may be difficult due to privacy issues (e.g., recording video in a public space). The advantage, however, is that a high ecologic validity of the data can be assumed. Furthermore, there are aspects such as effectiveness [251], social effects [228], audience behavior [254], and privacy implications [10] that are almost impossible to measure in the lab. The disadvantage of field studies is that they are usually

complex due to the high number of potential influences and require a tremendous effort in preparation (finding a suitable place, legal issues, etc.). Traditional methods are observations and logging.

Deployment-based Research

Deployment-based research is a kind of action research that introduces technology (e.g., public displays) into a social setting (e.g., a city), to address some research questions derived from theories [10; 254]. User feedback and involvement are then obtained, and in an iterative process, the deployment is improved. At the same time, this data is used to build and refine the theories, which in turn generate new research questions that can be addressed through changes in the deployment. In contrast to field studies, deployments are integrated into the everyday life of their users, and in contrast to ethnography, researchers do not intervene by deploying a prototype.

There is a continuum from cultural probes over technology probes to deployments [151]. *Cultural probes* support users with things like cameras to document their lives, while *technology probes* introduce small prototypes in order to understand a given domain, sometimes without the scientific rigor introduced in experiments. Only *deployments*, however, really become permanent useful artifacts in everyday life. Deployments enable researchers to investigate longitudinal effects of use that cannot be investigated with other means. They are also the only method that can really get rid of the novelty factor, which influences other kinds of studies. On the other hand, the maintenance of such deployments binds considerable resources. Examples of deployment-based research are Hermes [66], the Wray display [67], eCampus [353], and UBIOulu [272].

2.8.4 Methods and Tools

Interviews

Interviews are often semi-structured, i.e., the interviewer follows some pre-defined guidelines but would dig deeper if they discover interesting findings. Interviews can also be conducted in context (e.g., shortly after the subject used the system). Interviews are a powerful method for understanding the user's views (e.g., concerns, problems, and opinions [10]).

Questionnaires

Questionnaires are a useful method for the quantitative evaluation of public displays. Questionnaires can be *standardized*, thus allowing different systems to be assessed and compared with regard to, e.g., usability (System Usability Scale [26]), user experience (AttrakDiff¹³), or task load (NasaTLX [134]). On the other hand, questionnaires can be *customized* and used to ask the user about their personal views. Questionnaires have been used in many of the reviewed projects in order to assess, e.g., user experience [35], user performance and the users' views [25].

Focus Groups

Focus groups are used in early stages of the design process (usually as soon as an early prototype exists), to discuss it with people of the potential target group. They are run with 5-8 people in sessions lasting about 1-2 hours, including a demonstration of the system and hands-on trials, followed by a discussion. The discussion is led by one of the researchers based on guidelines and tries to answer important research questions. Alt et al. use a focus group to assess the users' views on contextual mobile displays [14]. Cheverst et al. use focus groups to discuss multiple system designs with different degrees of interaction [67]. The advantage is that feedback (e.g., on potential issues) can be provided in very early stages of the design process. On the other hand, opinion leaders in the group may prevent some people from stating their (contrasting) views.

Observations

Observations are most powerful when it comes to (post-hoc) analysis of audience behavior [35], as well as effectiveness [254] and social impact [164]. In general, two forms of observations can be distinguished: automated and manual. During *automated observations* users are observed by cameras installed in fixed locations (potentially filming both the screen and the viewer) [35]. The video footage can be analyzed post-hoc using computer vision methods such as shape or movement detection, eye recognition, or manual annotation and coding [35]. When conducting *manual observations*, data is gathered by observers, e.g., by taking field notes or pictures and videos from both the subject and the display [254]. In this case, the observers usually hide in a location from which both the screen and the interacting persons can be seen [10]. The advantage of the method is that users behave most naturally if they are not aware of being under surveillance,

¹³ AttrakDiff website: <http://www.attrakdiff.de/en/Home/>, last accessed February 13, 2013

making the findings highly ecologically valid. On the downside, video-based observation may compromise the subjects' privacy, and it may be very difficult to make conclusions on why users behaved in a certain way. Therefore, observations are often combined with (post-hoc) interviews.

Logging

User interaction can also be logged (e.g., time to perform a task and number of clicks) during the study and analyzed post-hoc. This method is particularly helpful when conducted over a longer period of time. Means for logging include all types of sensors that allow motion [35], eye gaze [324], presence [254], or user interaction [10] to be assessed. Logging is used in many observations, e.g., to assess trajectories [35], time of day [10] and type of content [324] with which interaction occurred. The advantage is that a lot of data can be gathered with literally no effort. The disadvantage is that logging often collects personal data.

2.8.5 Assessing Validity

Our literature review reveals a broad set of study types and methods that have been used by researchers during their evaluations. We find that most of the projects used various methods in parallel, showing the complex nature of public display evaluations. In the past four years, we have carried out more than 30 studies involving public displays in the context of various (large-scale) research projects, which allowed us to validate the methods described above. In the following section, we provide a high-level summary of both our experiences and findings from previous work.

Internal, External & Ecological Validity

Most studies can be criticized as not exercising sufficient control over confounding variables (internal validity), not generalizing findings to cover other settings and situations (external validity), or not testing a realistic situation (ecological validity). Internal, external, and ecological validity usually cannot be achieved at the same time. Instead, studies must often sacrifice one or two of them to improve the third. In HCI, internal validity is often prioritized above the other two, leading to highly controlled lab studies with rather low ecological validity. In contrast, public displays are, by nature, a very social phenomenon. Behavior in the public space may be very different than what is expected from the lab [147]. Hence, ecological validity is often prioritized above internal and external validity.

Table 2.3: Average Number of Participants for Different Methods.

	Field Study			Lab Study		
	avg.	med.	sd.	avg.	med.	sd.
Interview	26.9	15.0	29.5	16.0	12.0	2.9
Observation	35.5	36.0	12.7	14.3	6.0	5.9
Questionnaire	38.4	32.0	37.4	15.0	12.5	5.0

(Note that only within-subject designs were considered.)

Study Size

An often occurring question is what sample size is required in order to be able to draw well-founded conclusions. Although there is certainly no ultimate answer, we analyzed the number of participants from more than 80 studies (cf. Appendix I). Table 2.3 provides a rule of thumb as to which sample sizes may be most appropriate for which method.

2.8.6 Guidelines

First of all, researchers should start with clear research questions and decide whether to run a descriptive, relational, or experimental study. We derived the following set of guidelines to help researchers with the design of their public display studies.

Choose to focus on internal, external, or ecological validity. Oftentimes, control, generalizability, and realism cannot all be achieved at the same time. It is important to make clear which kind of validity will be at focus and which validity will be partially sacrificed. In public displays, ecological validity is typically more valued, but it has to be clear how internal and external validity are reduced, and measures can also be taken to improve these (e.g., randomization to decrease the influence of confounding variables).

Consider the impact of the content. Public display research is not possible without content, but the impact of content and other factors on usage is indistinguishable [353]. Thus, every study is at risk of producing results that are only valid for the particular kind of content tested. Testing different contents might help.

Understand the users. Public displays may have vastly different users at different locations and different times. In [254], for example, school children in the morning behaved very differently from drunk people at night. Furthermore, it is generally advisable to practice triangulation and combine different methods towards a common research question. Experiments can often be accompanied by descriptive research, and quantitative by qualitative measures. The deployment in public space also introduces ethical issues, and anonymization of any required data will usually be necessary.

Check for common problems. One of the most common problems with public displays is that they do not receive a lot of attention. This happens because passersby do not expect them to be interactive [254] nor do they expect to find interesting content on the displays [255]. Engaging people that actively promote displays might help to raise audience awareness of a display's interactive capabilities [271]. Creating customized content that reflects the users' interests may also help but is expensive [353]. A possible solution could be the use of autopoiesic content, i.e., self-generated content [228]. However, the effect of display blindness might be caused by other factors imposed by the environment. For example, at one location in the Looking Glass study [254], very few people looked at the display, because their heads were turned in the opposite direction to look for oncoming traffic as they approached. In another example, Storz et al. [353] point out technical challenges. Their initial deployment was much shorter than envisioned, because their hardware was not built for a setting that had a high volume of diesel fumes, which in turn caused the projectors to shut down automatically. Besides influencing the performance of the hardware, physical conditions and the setting can also influence the user acceptance. Previous research shows that placing a display in a leisure-oriented environment such as a swimming hall can attract more attention than in a business-oriented environment such as a municipal service center [271]. Hence, understanding the environment and its challenges before the deployment is crucial.

Chapter 3

Pervasive Advertising

The increasing affordability of powerful mobile devices, combined with cloud-based data storage and advances in wireless communications, have made it easier than ever for people to access information and obtain a wide range of products and services anytime, anywhere [319]. However, a key obstacle to the deployment of ubiquitous computing systems in public spaces is the question of who will pay for them. Two online payment schemes prevail. Users either buy a product, such as an app or song, or they subscribe to a service, such as one that tailors content to passersby. While Internet access itself is typically a paid service, most Web-based services – e-mail, social networks, news, and so on – as well as many apps, are free to users and thus rely on advertising.

Just as it pays for other forms of media, including free newspapers, radio, and TV, we believe that advertising – not direct sales or service contracts – will also underwrite the future ubicomp infrastructure. We envision this infrastructure to start as display networks and later expanding to encompass entire interactive spaces. However, for advertising-based public display networks to become truly pervasive, using all available communication channels [249], they must also provide a tangible social benefit and be engaging without being obtrusive. In many ways, public display networks must parallel the World Wide Web, blending commercial with informative content.

Advertising performs a vital function in society by conveying information about products and services, which benefits both producers and consumers. It is hard to imagine a modern capitalistic economy without it. At the same time, people

often respond negatively to advertising, especially crass efforts at manipulation. Those living in large cities see up to 5,000 ads per day [351], and many regard the proliferation of signs and billboards as a form of urban blight. Sao Paulo, Brazil, even went so far as to ban all public advertising in 2007 [57].

The perceived advertising glut is a product of market distortion. Ads are intended to consume our attention, but advertisers cannot be charged for accessing this scarce resource, so the market equilibrium shifts toward consuming maximal attention. To 'price in' this negative externality, city governments forbid certain kinds of ads or require companies to purchase a license to advertise, using the proceeds to operate public infrastructure such as bus stops and toilets. These regulations must often be amended to address challenges posed by new technologies, such as digital billboards.

Pervasive Advertising will affect everyone's life – whether we like it or not. The full potential of pervasive advertising has not yet unfolded, but we are currently at a crossroads, where a decision will be made that leaves us in a better or worse off position. One direction might lead us into a world where we are clogged with pervasive spam, we are being spied on and being (subconsciously) manipulated to buy things we do not need. It was only in 2010, when NEC deployed the 'Next Generation Digital Signage Solution' [363] in airports and malls in Japan. It is a new advertising system that allows content to be tailored to the passerby by determining a person's gender and age.

However, we still have the choice to take the future into a more beneficial direction, creating a world in which pervasive computing unfolds its full, positive potential. In this world, we are provided with information anytime, anywhere as we need them and advertisements create inspiring experiences. Advertisements will strike a balance between being calm when we do not need them, and being engaging and inspiring when we want to participate. Our privacy is well protected and we can inspect, change, and delete any data about us that is available to advertisers. While ads might still persuade us to do things that are in our interests, unethical persuasion is avoided, persuasion strategies are overt, and we have the ability to express our opinion of them.

Pervasive Advertising and Pervasive Commerce will, in the near future, allow the customer to be reached anywhere and at anytime [308]. Public Displays (or Digital Signage) is likely to play a key role as it has the potential to influence the customer throughout the entire purchase decision process. Chandon et al. showed, that 75% of all buying decisions are made at the point-of-sale [62]. Consequently, public displays can unfold their full potential in retail environments as they inform the customer at the right time and at the right spot. At the same time, public

displays achieve recall rates (15-30%) similar to those of print (10–20%) and TV (20–30%) advertisements [361], hence making them a suitable means not only for product promotion and brand building but also for interactive information exchange and subconscious priming [341].

One major issue we are observing today is that there seems to be a considerable knowledge gap in terms of advertisements of the future between different disciplines – mainly marketing, computer science, psychology, and communication science. Whereas advertising experts are largely unaware of pervasive computing technologies and how these may influence and change their business field, computer scientists and experts know only very little about advertising as they may have experienced it, so far, only as a consumer. This background chapter aims at closing the most fundamental gaps by shedding light onto the core principles of the respective fields. We provide a definition of the most important terms, before looking into the objectives of (pervasive) advertising. We also look at advertising performance and discuss how measures may change in the future. Finally, opportunities and challenges in pervasive advertising are being presented.

This chapter is based on the following publications:

- J. Müller, F. Alt, and D. Michelis. Introduction to Pervasive Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*. Springer Limited London, 2011
- F. Alt, J. Müller, and A. Schmidt. Advertising on Public Display Networks. *IEEE Computer*, 45(5):50–56, 2012

3.1 Definitions

3.1.1 Marketing

The terms advertising and marketing are often used synonymously. Marketing is confused with advertising and selling techniques. This is partly due to the fact that the concepts, strategies, and instruments of marketing are often not visible to the typical consumer. Marketing plays an encompassing role and is integrated into a company's entire value creation process.

Definition: Marketing

“Marketing is the activity, set of institutions, and processes for creating, communicating, delivering, and exchanging offerings that have value for customers, clients, partners, and society at large.” [19]

This definition shows the broad sweep of the marketing discipline. It also describes three important marketing trends. First, marketing serves more than just the purposes of a given business. It also includes the general activities and institutions beyond the scope of traditional organizations. This conceptual extension takes into account that marketing is no longer conducted by businesses alone, but also by agencies, self-organized groups, or even individuals. Moreover, the classical understanding of marketing is no longer restricted to functional aspects. The conceptual scope is much broader. Marketing also includes non-functional pursuits. For example, activities that are not necessarily associated with a corresponding output. The second important change is the relationship between companies and customers. Since the second half of the twentieth century, marketing is no longer limited to the one-way value delivery from companies to customers. Today, buyers and sellers enter into long-term relationships that focus on exchange rather than one-sided interaction. Customers have become active participants in the company-client relationship through the intensive use of digital communication technologies. The third significant change concerns the recipients of marketing activities since these are no longer restricted to customers per se, but also now include partners and ‘society at large’. Marketing is increasingly viewed as an exchange between companies, customers, and community groups, where all involved parties see each other as equal partners and adapt their expression accordingly.

Marketing is based on the *marketing mix* – a set of marketing instruments that companies use to reach their customers directly. The classical instruments of the marketing mix fall into four areas: product, price, place, and promotion [182].

Product Product-related instruments are all activities and procedures that take into account the needs and requirements used to design current and future products. In this sense, the product policy encompasses the maintenance of successful products and the planning and realization of product innovations.

Price The pricing policy is based on the decision related to the type and extent of compensation that customers pay to use the company’s services. Price instruments not only refer to the actual price but also discounts, surcharges, or price timing.

Place Instruments of distribution consist of all measures designed to sell products and services as well as to logistically organize their distribution. Sales measures directly or indirectly target the purchase process. Distribution measures ensure the availability of the product for the customer.

Promotion Promotion includes all market communication activities of the company. This includes defining communication objectives and target groups, selecting communication channels, and determining the size of the communication budget. Promotional instruments can be divided into ‘above-the-line’ and ‘below-the-line’. Traditional advertising in newspaper, television, radio, cinema, and out-door advertising are considered ‘above-the-line’ measures; all other new communication tools are considered ‘below-the-line’ (e.g., sales promotions, sponsoring, product placement, event marketing, and online communication). Advertising is by far the most important communication tool in the marketing mix [226].

Marketing is hence a medium or long-term strategy for a product line or brand to be established in the market.

3.1.2 Advertising

As shown, advertising is one of four areas of the marketing mix. It is defined traditionally as a mass communication process designed to change the recipient’s attitudes and behavior [226]. Kotler and Keller have formulated a standard definition of advertising that should also define the use of the term in this thesis:

Definition: Advertising

“Advertising is any paid form of non-personal presentation and promotion of ideas, goods, or services by an identified sponsor. Advertisers include not only business firms but also charitable, nonprofit, and government agencies.” [182]

The general goal of advertising is to transmit information to a specific group of recipients in order to achieve the desired effect. Accordingly, the task of advertising is to systematically plan, design, coordinate, and control all communicational activities of an organization with respect to relevant recipient groups in order to contribute to the marketing objectives.

3.1.3 Pervasive Advertising

After decades of development in the laboratories, pervasive computing technologies are finally in a position to reshape our world. Analogous to developments on the Internet, it is our belief that advertising will be the business model that drives pervasive computing. As illustrated above, the goal of advertising is to impart information, evoke emotions, and trigger actions. The properties of pervasive computing (automation, interactivity, ubiquity) make it a powerful tool for achieving these goals. These properties have the potential to change advertising in six main ways: symmetric communication, long tail, experiences, personalization, audience measurement, and automated persuasion.

As stated previously, advertising is defined as any paid form of non-personal presentation and promotion of ideas, goods, or services by an identified sponsor. Pervasive computing environments are saturated with computing and communication capabilities, yet these features are integrated so seamlessly for the user that it becomes ‘the technology that disappears’ [379]. Based on these definitions, we define pervasive advertising as:

Definition: Pervasive Advertising

“Pervasive advertising is the use of pervasive computing technologies for advertising purposes.” [249]

3.2 Objectives

This section first provides an overview of traditional advertising objectives. It then shows how to reach these goals by developing advertising programs. Third, it looks at advertising performance encompassing state-of-the-art measures and explains how advances in pervasive advertising can be used to draw a more elaborate picture of the consumer and cognitive processes that may ultimately lead to purchase decisions. The section concludes with an introduction to targeting.

Advertising objectives have a direct means-end relationship to overall marketing and business objectives. The achievement of advertising objectives therefore contributes to the fulfillment of higher corporate goals [226]. Advertising objectives should be defined by content, scope, time, and target segment. They can be differentiated into cognitive, emotional, and conative goals as the following categorization of advertising objectives shows [316]:

Cognitive Objectives If the corresponding need is minor, the essential information about the product or service should be communicated. A need is considered minor, whenever the recipient is aware that the need is currently present and that it can be satisfied by available offers. In this case, it is often sufficient to communicate the key features of the product.

Emotional Objectives Emotional advertising aims to link a product or service with specific emotions that lead to a differentiation from other offerings. Emotions are used primarily in mature markets with technically and functionally interchangeable products.

Conative Objectives Conative objectives are related to actual consumer action. In this case, advertising aims to stimulate potential buyers to buy, order, use, or take action in any form.

3.2.1 Advertising Programs

In order to reach the advertising objectives, advertising programs need to be developed. This development consists of five steps: setting objectives, establishing a budget, choosing and creating the message, selecting the media-channels, and, finally, evaluating the results. These five steps are known as the *5Ms* of advertising: Mission, Money, Message, Media, and Measurement [182].

Mission

Defining the advertising objective is determined by the target market and the positioning established by the organization's marketing strategy. The objective or advertising goal is the desired result of the communication process between advertiser and recipient within a given timeframe. According to the advertiser's needs and aims the objectives can be classified as described above, into cognitive, emotional, and conative advertising goals.

Money

Defining the budget is generally determined by the product's life cycle, the existing consumer base and market share, competition, buying frequency, and the substitutability of the offer. Of course, the budget varies widely with the media and technologies used.

Message

Developing the advertising message and positioning the advertisement is a creative and an analytical task. Using market research, advertisers learn about how, when, and where their target audience will most likely perceive their message. Based on this knowledge, they decide the position, frequency, and other aspects of the advertising message. The actual creation of the message, that is to say, the design, layout, logo, etc. is the creative part of the development process. The latter is just as important for the advertisement's effectiveness: "The ad's impact depends not only on what is said, but often more importantly, on how it says it" [182]. In addition, as postulated by McLuhan, the medium influences the perception of the message [225]. This is of particular interest when new technologies are used to communicate with potential customers.

Media

Media types vary in aspects such as reach (percentage of the target market exposed to the medium), frequency (e.g., frequency of message display), or impact (e.g., persuasiveness of the medium). In principle, advertisers try to find the best balance between reach, frequency, and impact and the corresponding costs. Among the variety of advertising modes and techniques we would like to highlight out-of-home and point-of-purchase advertising. Both are highly relevant to the field of pervasive advertising.

Out-of-home advertising refers to a range of advertising methods designed to reach people in their everyday environments. Most of these environments are public or semi-public areas in which regular activities such as working, shopping, or traveling take place. Typical environments are shopping malls, airports, train stations, or city centers. Frequent types of out-of-home advertising are described by Stalder in [350].

The term *point-of-purchase advertising* describes ways to communicate with potential customers during the actual act of purchase. Next to classical in-store TV advertising, other forms of in-store advertising include ads on shopping carts, aisles and shelves, in-store demonstrations, or coupon machines. A significant number of consumer purchase decisions take place at the point-of-purchase. Point-of-purchase advertising has a strong potential to remind consumers of certain offers while making their final decisions as well as stimulating spontaneous purchases.

Measurement

The final step is the evaluation of the results. Has the ad been communicated effectively? In order to measure the communication effect, an advertising objective needs to be set in advance. An advertising objective is, according to Kotler “a specific communication task and achievement level to be accomplished with a specific audience in a specific period of time” [182].

Different objectives and different advertising media allow different measuring techniques. Whereas measuring the click-through-rate of an online banner is relatively simple, measuring the communication-effect of a newspaper advert is still very complex. This differentiation should be considered when developing new advertising techniques.

3.2.2 Advertising Performance

In order to express the efficiency of an advertisement or an entire campaign, marketers are interested in conversion rates and the return on investment. The *conversion rate* is defined as “the percentage of visitors who take a desired action” [274] (e.g., buying a product). In the Internet, online stores traditionally calculate visitor to customer conversion. With techniques for more fine-grained tracking of a user, a so-called conversion funnel emerged where online marketers measure a multitude of conversions, e.g., from users of a search engine to visitors of a website, from visitors of a website to contacts, subscribers of a newsletter, or users interested in downloads, etc. In this context, the click-through rate is an important metric, which is defined as the ratio of clicks onto an advertisement to the overall number of page impressions. The *Return on Investment (ROI)*¹⁴ is a concept used for optimizing the budget spent on advertising in support of the advertising strategy [206]. Improving the ROI is an important way to increase marketing effectiveness as well as revenue, profit, and market share.

In order to calculate conversion rates and return on investment in online marketing, the following measures are commonly used:

- **The CPM (Cost per Mille) or CPT (Cost per Thousand)** charges advertisers based on the exposure of a message to a specific audience. ‘Per Mille’ means per thousand impressions of the message. In this model it is assumed that each page served to the browser is being viewed by the user.

¹⁴ Often referred to as Return on Marketing Investment (ROMI).

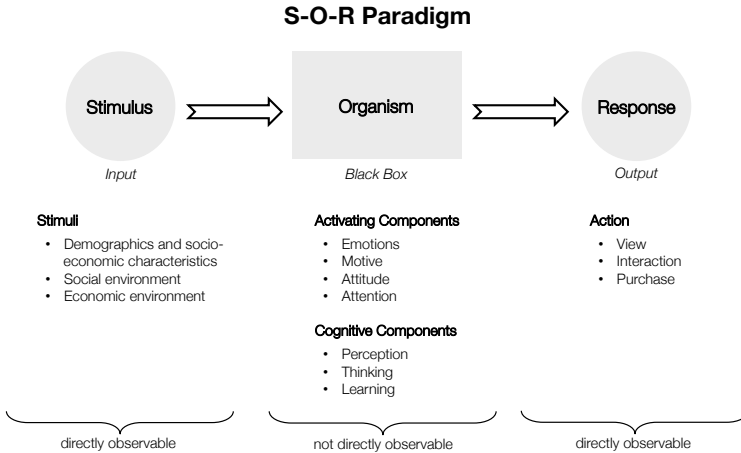


Figure 3.1: Stimulus-Organism-Response Paradigm (adapted from [186; 197]).

- The **CPC (Cost per Click)** charges advertisers for each click a user makes on a listing and hence redirects him to the advertiser's website. In contrast to the variant CPV (Cost per View), where advertisers pay for each unique user view of an advertisement (e.g., a pop-up), it is irrelevant whether or not the visitor makes it to the target website.
- The **CPA (Cost per Action)** is a performance-based model where the advertiser only pays as a user completes a transaction, hence shifting the risk to the publisher. Two variants exist:
 - In the **CPL (Cost per Lead)** model, advertisers pay for any action that potentially leads to a sale (e.g., completing a form or signing up for a newsletter).
 - The **CPS (Cost per Sale)** charges the advertiser every time a sale is made.

With the use of pervasive computing technologies, the effectiveness and efficiency of advertisements can reach a new quality and measures can be applied also to the real world. Classical research on advertising effectiveness is classified based on the so-called *S-O-R paradigm* [21] (Figure 3.1), which today forms the basis of many comprehensive and partial models [186; 197]. In contrast to

the classic Stimulus-Response paradigm that considers cognitive processes as a 'Black Box', the neo-behavioral S-O-R- paradigm considers hypothetical constructs, such as perception, emotion, attitude, and motivation as intervening variables. Understanding these hypothetical constructs allows different states and processes of consumer behavior to be identified (Figure 3.2) [365]:

- **Knowledge / Cognition:** Knowledge describes the degree of awareness about an objects' properties and relations that can be altered through information acquisition and processing.
- **Involvement / Arousal:** Arousal is the basis of each behavior. It can be stimulated by the time of day, internal processes (e.g., thinking) or external stimuli (e.g., listening to music). The level of arousal impacts the intensity of the processes and the degree of cognition. Involvement, as a specific form of activation, is particularly important when describing consumer behavior as it reflects the willingness of the consumer to acquire and process information.
- **Feelings / Emotions:** Feelings and emotions are states of inner arousal with a positive or negative direction. Feelings are (more or less consciously) perceived emotions. Emotions can be triggered by external stimuli or inner, neuronal processes.
- **Motives / Needs:** Motives can be activated through feelings of lack or external stimuli. Each behavior can be explained through motives.
- **Attitude:** Attitude describes an inner mindset of a person (positive or negative) towards an object (persons, situations, items) that is rather consistent over time. This strongly influences the person's buying behavior. Attitudes are learned during the socialization process either through self experience or by adopting the experience of others.
- **Values / Measures:** A value is the basis for ethical action. The set of all consistent values and measures forms a consumer's value system and has a (direct or indirect) influence on the buying behavior.
- **Lifestyle / Personality:** The lifestyle is an observable behavior pattern which is formed through personal and societal values. The lifestyle is often used as the basis for segmentation.

These states and processes are difficult to measure, particularly in an automated way. As a result, current advertising metrics try to abstract from these processes.

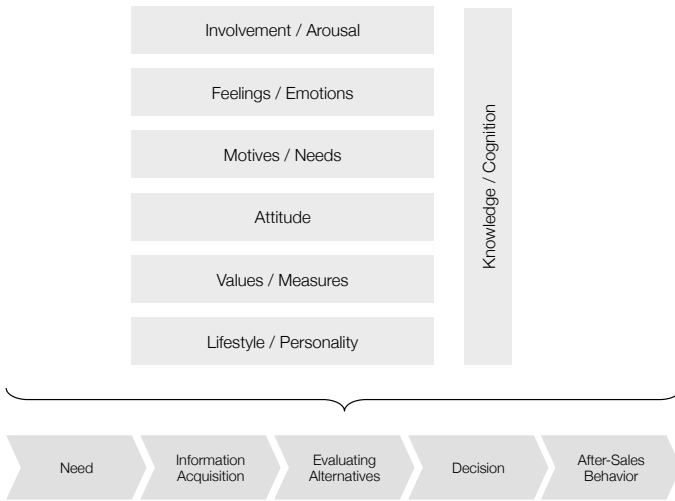


Figure 3.2: States and Processes of Consumer Behavior (adapted from Trommsdorff [365]).



Figure 3.3: Involvement comprises of three levels (from low to high): attention, comprehension, and elaboration. The higher the level, the more attentional capacity is allocated to analyze the message which results in more durable cognitive and attitudinal effects (adapted from Greenwald [127]).

For example, ‘contact’ as one of most important metrics in advertising today, is an abstraction of attention. Attention – “a process in the attempt to select information for reception” [172; 334] – is the first level of ‘involvement’ (Figure 3.3), a term that goes back to the 1940s [333] and has been adapted to advertising by Krugman in 1965 [190]. Involvement is a person’s perceived relevance of an object based on inner needs, values, and interests [386] and describes the strength of the relationship between consumers and an object, such as a product or a service. It is widely believed that the effectiveness of an advertising message is being moderated by the involvement of the audience [127]. As a consequence, involvement is used in modern marketing to categorize buying decision processes [197]:

High Involvement Purchases are important for the consumers and are strongly related to their individual personality and self-assessment. Due to high risks, they are usually based on complex decision processes and require a lot of time for choosing among product alternatives (e.g., luxury articles).

Low Involvement Purchases are based on decision processes only to a certain extent as they are less important and of lower risk, and often follow specific behavioral patterns (e.g., buying generic products such as sugar or toilet paper).

Marketing and psychology researchers are highly interested in how involvement can be achieved and which consequences emerge [29; 291]. To better expelling the processes, dual processing theories are widely used. These usually distinguish a conscious and a subconscious processing path [264]. The most referenced theory is the so-called *Elaboration Likelihood Model* (ELM) [283]. The ELM explains the changes in attitude through advertising. It differentiates between two routes to persuasion or elaboration: the central and the peripheral route. Only the central route leads to a (positive or negative) deep and persistent change in attitude. The peripheral route leads to a rather weak and less persistent change in attitude, which seems intuitively less desirable for an advertiser. The decision of which route to take depends on two factors: the motivation to process and the ability to process. The first (motivation to process) depends on the personal relevance of the message for the recipients as well as on their personal responsibility or involvement towards the message and their need for cognition. The second (ability to process) is dependent on things like extent of distraction, the frequency of message repetition, the existing knowledge and the free cognitive capacities, etc.

When looking at measuring involvement, technology has come only so far as to enable attention to be measured in an automated way (or, as already mentioned, actions to which attention can be abstracted, e.g., page impressions). Comprehension and elaboration are difficult to assess and expensive methods such as interviews and questionnaires are required. In order to measure the efficiency, the major challenge to be overcome is the operationalization of these constructs. Examples are recall and recognition or number of 'intentions to recommend'. A comprehensive overview is provided by Trommsdorff [365].

Advances in pervasive computing will make it possible to more comprehensively assess the states and processes, or their manifestations that are observable. Two examples are emotions and engagement (as an observable form of involvement).

Example I: Emotions

Advertising has always used emotions to influence a user's perception of a brand or product (emotional objectives). Primary emotions are interest, happiness, surprise, grief, anger, honesty, disregard, fear, shame, and feeling of guilt [162]. With novel technologies that allow the users' emotion to be observed and assessed in real time, a new quality can be added to advertisements. On one hand, content can be tailored towards the user based on their current emotion [100]. On the other hand, implicit feedback of the user on whether or not they liked an advertisement can be gathered. Different technologies can be used to assess emotions. Küblbeck et al. presented SHORE, a camera-based system to analyze facial expressions and then determine emotions [192]. Further technologies include psycho-biological instruments (e.g., electro-encephalography), skin conductance, or gaze measurement (width of pupil). However, the problem with these approaches is that the direction of the emotion is usually difficult to be assessed.

Example II: Engagement

Novel, interactive systems have the potential to engage the users (e.g., making them interact with an interactive game on a public display). Engagement can be seen as a result of involvement and measuring engagement brings researchers and marketers closer to comprehensively assessing involvement. Different ways of quantifying engagement exist for different media. Online advertisers defined engagement as the proportion of impressions that consumers engaged with (e.g., by mouse-touch) and for how long [75]. This measure draws from the fact that users tend to follow the mouse cursor movement with their eyes. Hence, an estimate of the share of impressions that users actually saw can be provided. The problem is that users can obviously see an ad without a mouse-touch and, if the dwell is too low, users are likely not to have perceived the advertisement. Also mobile advertising is looking into engagement and refers to it as the 'duration of sight' [330]. Capturing action and reaction between the first impression and the last 'moment of sight' allows to be understood how the user interacts and can make mobile ads more effective. Engagement on public displays could be quantified through the level of interactivity. For example, using expressive body gestures to control a game could have a different impact than simply touching a button – not only on the users themselves, but also on other passersby (cf., the honeypot effect [254]).

3.2.3 Targeting

An important strategy in advertising is market segmentation [345]. The aim of market segmentation is to split a heterogeneous market into homogeneous partial markets (target groups) based on different segmentation criteria. The STP model as presented by Kotler [182; 183] identifies three steps: segmentation, targeting and positioning.

Segmentation

Segmentation is the process of separating the market into clearly definable groups of buyers in a way such that needs of this group can be satisfied with a particular marketing mix. The resulting groups (segments) need to be homogeneous with regard to the expectations towards a product or service, the buying habits, media consumption, etc. since every buyer has their own needs; each theoretically represents a separate market. However, developing marketing programs for each single customer is in general not feasible, especially as companies are facing large numbers of customers who only buy small amounts of products. Based on the desired granularity of the segmentation, customers can be grouped into a continuum of markets ranging from mass markets (null segmentation) via target group markets and niche markets to micro markets.

Targeting

In the second step, the so-called targeting, the company needs to come up with a measure of attractiveness for the segments and then decide to target one or more segments. Evaluation criteria include so-called company independent criteria, for example size and potential of the segment, level of competition, potential for growth, and risks. In contrast, company dependent criteria include investments, production capacities, rentability, know-how, brand loyalty of customers and competitive advantages. For an overview of methods used for evaluating and choosing segments (e.g., checklists, profiles, scoring) we refer to [113].

Positioning

Finally, during positioning, suitable marketing instruments (see Section 3.2) are selected in order to reach the marketing objectives. The decision of which instrument needs to be made for every segment and strongly depends on whether an instrument is applicable to separate customers or to segments only. This is usually a technical question. For example, using face detection in front of a

public display to find out whether the person standing in front of the display is male or female allows marketing instruments on segment level to be applied (such as a simple advertising application that shows cars to men and jewelry to women). However, if the very customer standing in front of the display should be addressed via direct advertising, face recognition or other technologies that enable identification need to be put in place. A comprehensive overview of the segment specific use of marketing instruments can be found in [113].

3.3 Opportunities

Pervasive computing technology will change advertising. The analog world offers useful hints with regard to new opportunities [11], but new technologies require rethinking current practices [30]. Major opportunities lie in the shift of power between stakeholders, the long tail, more engaging experiences, personalization, audience measurement, and persuasion.

3.3.1 User Feedback

Classical advertising follows a mass media approach in which a small number of advertisers distribute their advertisements to the masses. This unidirectional communication model produces an asymmetrical distribution of power. All the power is concentrated in the hands of advertisers who decide which ads to show when and where. At best the audience has the option to ignore, protest against, or vandalize the resulting ads. For some people, such an asymmetrical distribution of power creates a feeling of being at the mercy of advertisers.

Since pervasive computing is interactive it offers the opportunity to transfer a significant degree of power to the audience. This fundamentally alters the unidirectional communication model by allowing the audience to communicate opinions directly to advertisers and others. Companies should treat customers as equals. This can benefit both consumers and companies, since a closer bond is created, and because it allows companies to learn from their customers much faster. Practical examples include the ability of the audience to choose the content they like or to submit their own content. Also, social media will foster communication within communities [348; 359; 375]. Eventually, this may lead to a democratization of ads and the look of public spaces [366].

3.3.2 The Long Tail

By definition, pervasive computing is highly automated and many things that required individual attention in classical advertising will also become automated. This significantly lowers the cost of and effort needed to produce individual advertising campaigns. Starting a new campaign may be as easy as filling out a few fields on a website and may cost only a few cents. This price decline enables very small companies and even individuals to launch their own tiny, local campaigns. It is important to remember that not only big companies are interested in advertising as the communication of sponsored messages. Even a small restaurant or market stall must advertise, just like anyone who wants to sell an old bicycle. Even someone who wants to surprise his wife at an airport or make a birthday surprise to a friend might be interested in displaying something in public. Some examples of how to accomplish this are presented in [167].

3.3.3 Engaging Experiences

Pervasive computing offers powerful media that respond to all senses. Large, bright displays that surround us create powerful visual impressions, but it is also possible to appeal to our hearing, what we feel, haptics, and even our sense of smell. As Norman [262] explains, there are three levels to interactive computer systems. The lowest level is *visceral*, or in other words, what initial visual impression the technology makes. This can be described as ‘the first impression’. The second level is *behavioral*, which is related to the look and feel, and described as ‘how it feels’. The third level is *reflective*, meaning, e.g., what we think others think about us when we use it.

Since most traditional ads are not interactive, they do not go beyond the first level. Pervasive advertising, however, needs to properly address all three levels. It has a look and feel, and also makes us reflect when we interact with it. These properties make pervasive advertising a much more powerful tool. Since pervasive computing is all around us, engaging experiences can follow us and surprise us wherever we go. Furthermore, since pervasive advertising is digital, it is very easy to create new experiences all of the time. Together, this helps creating a wow-effect possible over and over again. Analog posters will look relatively pale compared to the intense and memorable experiences that can be created with pervasive advertising. Examples of what this can look like are presented by van Waart et al. [366]. By having fun, users will be more likely to pay attention to content, and perhaps draw the attention of onlookers.

3.3.4 Personalization and Context Adaptivity

Personalization and context adaptivity are at the core of pervasive computing and provide natural powerful tools for advertising. In personalization / user modeling, computers learn the preferences and behavior models of groups or individuals. This fits naturally with the target groups as a core concept of marketing. It is important to remember that in marketing, target groups are used in different ways: for the development of the product as well as for the placing of the advertisements. The dilemma of advertising traditionally was that the properties defining the target group have to be measurable, and target groups have to be accessible based on these criteria. This restricts them basically to demographics and other easily assessable criteria. Pervasive computing allows many more things to be quantified, thus making it feasible to develop target groups based on measurable demographic criteria or actual behavior. Pervasive computing allows measuring all kinds of things in real time, building user profiles, and adapting advertisements, thus making far more finely tuned target groups accessible [278].

In addition to personalization, adaptation to the context is more fine-tuned due to automation and better sensors. Traditionally, a huge effort was required to post different ads, for example, depending on the weather. Using pervasive advertising, however, things such as advertising ice cream when the sun is shining and umbrellas when it is raining become minor. It can be assumed that when advertisements are much better adapted to the context, for example, when showing products in the hometown context of the audience, they are more effective. How this can be achieved is described by Strohbach et al. [355] and Bauer and Spiekermann [28].

3.3.5 Audience Measurement

Audience measurement has always been an integral part of advertising, mainly because ‘if you can’t measure it, you can’t improve it’ (Lord Kelvin). Any advertising campaign is driven by goals, and goals can only be set according to factors that are measurable. Limited measurement capabilities also limit the scope of what can be achieved. Traditional audience measurement methods include diaries that are, however, subjective, error-prone, and labor-intensive. Later, more reliable technologies such as the Personal People Meter¹⁵ or software apps integrated with cable television were invented that monitor and record whatever program people are viewing.

¹⁵ Arbitron’s Portable People Meter: http://www.arbitron.com/portable_people_meters/home.htm, last accessed February 13, 2013

Pervasive computing provides powerful sensors for measuring the actual behavior of people. The immense opportunities this provides have been demonstrated by the Web. Because clicks are easily measurable, information on Internet usage could be collected by measuring page hits (e.g., comScore¹⁶, Wakoopa¹⁷, and Hitwise¹⁸). At the same time whole new business models and paradigms have emerged. Ads are often paid for on a click-through basis. Campaign success is completely transparent using tools such as Google Analytics¹⁹, enabling advertisers to optimize and cancel campaigns based on live data. Even the fully automated optimization of advertising campaigns is possible. Google Website Optimizer²⁰, for example, can automatically run statistical tests on user behavior in different versions of campaigns and optimize the campaign accordingly.

Pervasive advertising makes it possible to apply this entire approach to the real world. User behavior, for example, whether people looked at an advertisement, can be easily measured using computer vision and face detection and is already commercially available (e.g., TruMedia²¹, Quividi²², stickyPiXEL²³, and Cognition²⁴). Tracking when audiences interact with the ad is basically free, and even things such as eye tracking may soon be ubiquitous. This will allow advertisers to set goals that are far more detailed (e.g., 50% of bald men between the ages of forty and sixty should have read the first sentence in this text block advertising hair implants), and optimize their campaigns in a rapid loop. If it were possible, for example, to track how many people frown and turn away after seeing a specific part of an ad, advertisers can determine which aspects of an advertisement the audience prefers and adapt their campaigns accordingly.

¹⁶ comScore website: <http://www.comscore.com/>, last accessed February 13, 2013

¹⁷ Wakoopa website: <http://wakoopa.com/>, last accessed February 13, 2013

¹⁸ Experian website: <http://www.experian.com/hitwise/>, last accessed February 13, 2013

¹⁹ Google Analytics: <http://www.google.com/analytics/>, last accessed February 13, 2013

²⁰ Google Website Optimizer: www.google.com/websiteoptimizer, last accessed February 13, 2013

²¹ TruMedia website: <http://www.tru-media.com/>, last accessed February 13, 2013

²² Quividi website: <http://www.quividi.com/>, last accessed February 13, 2013

²³ Sticky Pixel website: <http://www.stickypixel.com/>, last accessed February 13, 2013

²⁴ Cognition by Intel: <http://intel.cognitionvision.com/>, last accessed February 13, 2013

3.4 Challenges

3.4.1 Calm vs. Engaging Advertising

That technology should be calm and require minimal attention has been a core feature of pervasive computing from its very inception. Weiser and Brown [379] proposed that, when computers saturate the surrounding environment, calm computers will be most effective. Key to this is the effortless sliding of information between the center and periphery of our field of attention. This conforms with the concept of context-adaptive displays where the requirements and wishes of the users are obtained from data using various sensors, and then the content of the screen is being ‘magically’ adapted. It became clear over the years that predicting or even quantifying what users want through observation alone is very difficult or even impossible. In response to Roger’s observation of these facts, she proposed the seemingly oppositional paradigm of engaging computing: computers should provide great experiences and engage users more in how they currently behave [305].

It is our belief that pervasive advertising on public displays should be both calm and engaging. Although this might seem like a contradiction, it is not. Calm advertising means that advertisements should be easy to ignore. Engaging advertising means that ads should provide engaging experiences when the person is actively engaging with them. Both can be achieved at the same time. A pervasive ad could appear as calm, mildly flowing water when nobody engages with it, and then convert to an engaging mini-game once somebody pays attention.

3.4.2 Privacy vs. Personalization

Privacy has been an important topic for pervasive and context-aware computing from the beginning. Yet, there is immanent tension as personalization is not possible without knowing the user. In pervasive advertising, there is a huge incentive for advertisers to collect as much user data as possible. Thus, it is crucial that user privacy is protected. This can happen either through industry self-regulation, government mandated regulation, or both. The degree to which privacy is protected and guaranteed will determine both how relevant and thus successful advertisements are and whether users trust advertisers. To gain such trust requires effort but trust can also quickly be lost; guaranteeing user privacy is one of the foremost challenges facing pervasive advertising.

3.4.3 Private vs. Public Advertising

One major challenge lies in whether an advertisement is being presented to the target audience in a public or a private way. Whereas outdoor advertising is clearly public, many ad channels are private or semi-private. TV spots are usually presented to small co-located groups, and Web ads as well as mobile ads are presented to one person only. The decision whether or not to present an advertisement in a public or a private way becomes essential with personalization, as advertisements could affect the passerby's privacy (e.g., a targeted advertising for an expensive car can reveal that the person in front of the display is wealthy). Consequently, it should be ensured that personalized advertising on public displays is not deterministic and that the choice of an advertisement can not be tracked back to the user.

3.4.4 Local vs. Global Advertising

Public displays are deployed in a specific physical environment, e.g., in a historical railway station or in front of a post-modern office building. Advertisers face the challenging task of blending in (or contrasting) the same advertising content with different environments, which is not the case for mobile devices. A similar problem exists on the Web, where web designers struggle to match dynamic ads with the corporate layout (color, font, etc.) of a page, which can be partly overcome using CSS. As the owner of the place is usually the stakeholder making the final decision as to what can be displayed, we believe that this may become a major challenge in the future.

3.4.5 Persuasion vs. Ethics

Public display designers can apply persuasive technologies [108] to create more effective systems. Since persuasion is an integral part of advertising, however, designers must be wary of crossing the line of manipulation. We believe that the use of deception, coercion, operant conditioning, or surveillance to exploit vulnerable consumers is unethical, and that a system's success will partly depend on the choices the designer makes.

III

THE PROVIDER VIEW

Outline

The following part of this thesis is concerned with the design space from the view of the display and advertisement providers. It addresses three core challenges that these stakeholders face when they deploy public displays. First, though currently large outdoor advertisers dominate the advertising landscape, we see a trend that more than ever, small businesses are entering the scene and that even users take the role of the advertiser. This makes it necessary to rethink the way advertising systems are designed. Second, interactivity and the ability to create arbitrary-shaped public displays require the understanding of how the audience behaves around such displays in order to appropriately design the content and place these displays. Third, pervasive advertising offers novel opportunities to target the customer with more interesting content. However, this should be done in a privacy-preserving and participatory way to increase user acceptance.

This part of the thesis takes a closer look at these three steps.

- **Chapter 4 – System Design** To influence the design of future, shared advertising displays, this chapter presents an ethnographic study that investigates today's practices surrounding shared public notice areas (PNAs). We look at the content posted to such areas, the means for sharing (i.e., forms of content control), and the reason for providing the PNA.
- **Chapter 5 – Audience Behavior** This chapter presents a prototype and reports on a user study, comparing the influence of the display shape on user behavior and user experience between flat and cylindrical displays.
- **Chapter 6 – Targeting** We present an approach for adapting content towards the users without affecting their privacy. The approach does not only consider personal interest, but also actual shopping behavior and aims at providing more interesting content.

Chapter 4

System Design

Large digital displays are rapidly permeating public spaces. The availability of suitable technologies for outdoor use and sinking prices for large display hardware has led to a transformation from paper-based to digital signage. Urban landscapes are augmented with digital signage solutions by large digital-out-of-home (DOOH) advertisers replacing more and more traditional billboards. Apart from reducing the cost of updating content, these displays allow animations and/or interlacing news content to be added, which increase the visibility and attractiveness for passersby. However, so far these digital displays are not globally networked and access is typically restricted to their owners.

We envision that in the future, these individual displays and isolated display solutions could be interconnected through the Internet. Hence, a canvas across urban space can be provided that allows any type of content to be distributed onto this display landscape, not only from large advertisers but also from neighboring shops, local residents, and visitors. The vision to make public displays more attractive by blending advertising and non-advertising content could be realized in this urban space [12]. Technically, the challenge of such a vision is to create a suitable middleware that supports the remote exchange and programming of arbitrary content onto arbitrary displays, as well as suitable interfaces to interact with such systems. A far greater challenge lies in the design and deployment of suitable control tools that can support the stakeholders' understanding of how these displays ought to be used. Without suitable incentives and means for staying in control, display owners might be reluctant to grant access to their displays and relinquish their control over what is being shown on their in-store displays.

Our central research question is: how can we build digital public display networks that can go beyond today's isolated advertisement displays and instead, provide an open platform for posting and displaying third party content (viewers, customers, charities, etc.)? Yet instead of tackling the technical challenge of such a vision directly, we begin our investigation with trying to understand the social and economical drivers to support this vision: What would motivate display owners to allow others access to their displays? And how would control interfaces and incentive structures have to look like in order to support the widespread uptake of such systems?

We decided to base our research in today's practices surrounding the precursors to our vision: shop-windows, notice boards, and wall hangers, where customers, community members, and visitors can use pen & paper and pins & tape to put up their messages, notices, posters, and classified ads (cf., Figure 4.2). For the purpose of this thesis, we collectively call such boards, walls, and hangers 'Public Notice Areas', or PNAs. We began by observing the use of PNAs in 29 locations in two different countries, using photo logs to document their use and change over 4 weeks. We then held in-depth interviews with most of the people responsible for the observed PNAs, in order to understand their current practices for controlling access to the PNAs, elicit their motivations for offering PNAs, and identify concerns they might have when relinquishing control in a fully digital public display network. The results from the observational studies and interviews are analyzed in a quantitative and qualitative fashion, using a 'data walkthrough' analysis that we developed to give all team members an in-depth view of the data collected by other team members.

In the remainder of the chapter a case study is presented that looks at the users' view on context-sensitive car advertising. The motivation stems from the fact that in the future literally every surface can be a display. Cars are ubiquitous and offer large and often highly visible surfaces that can be used as advertising space. Until now, advertising in this domain has focused on commercial vehicles, and advertisements have been painted on and were therefore static, with the exception of car-mounted displays that offer dynamic content. With new display technologies, we expect static displays or uniformly-painted surfaces (e.g., onto car doors or the sides of vans and trucks) to be replaced with embedded dynamic displays. We also see an opportunity for advertisements to be placed on non-commercial cars. Results of our online survey with 187 drivers show that more than half of them have an interest in displaying advertising on their cars under two conditions: (1) they will receive financial compensation, and (2) there will be a means for them to influence the type of advertisements shown.

This chapter is based on the following publications:

- F. Alt, N. Memarovic, I. Elhart, D. Bial, A. Schmidt, M. Langheinrich, G. Harboe, E. Huang, and M. P. Scipioni. Designing Shared Public Display Networks: Implications from Today's Paper-based Notice Areas. In *Proceedings of the 9th International Conference on Pervasive Computing* (San Francisco, CA, USA), Pervasive'11, pages 258–275, Berlin, Heidelberg, 2011. Springer-Verlag
- F. Alt, C. Evers, and A. Schmidt. Users' View on Context-Sensitive Car Advertisements. In *Proceedings of the 7th International Conference on Pervasive Computing* (Nara, Japan), Pervasive'09, pages 9–16, Berlin, Heidelberg, 2009. Springer-Verlag

4.1 Related Work

Public displays have been subject to research for many years. Many traditional displays are being replaced with digital counterparts due to the falling prices of displays. Many projects have looked at the technical requirements for networking digital displays, mostly within and across offices (e.g., [3; 71]) but also in public space (e.g., [256; 353]). Also, novel interaction methods have been studied, both in terms of user behaviors (e.g., [52; 238]) and interface technology (e.g., [105; 125]). Although technical and architectural suggestions could be drawn from these studies, our initial work in this domain focuses on understanding the design implications from existing practices of posting on PNAs and latent motivations for offering and maintaining PNAs.

There have been several studies of this kind that looked at current practices around publicly available notice boards and displays. Taylor et al. [358] look at community notice boards in a rural village to inform the design of a digital version. Churchill et al. [72] look at community notice boards in an urban area and in their own workspace to inform the design of their Plasma Poster Network, a system that enhances the chance to encounter interaction and awareness of different workgroups' activities. Huang et al. [147] conducted a field study to analyze various paper and digital displays and their actual placement, as well as how much people actually look at them. Based on their findings, they provide design recommendations for increasing the visibility of displays and for better matching between people's behavior and the displays' content. Although there

is some overlap between previous studies and our research, our primary goal is to come up with general implications that inform the design of a public display network system that go beyond display visibility and office space and suburban area settings.

The related works show that public displays have a large potential to foster communities. Redhead and Brereton [295] present a qualitative analysis of local community interaction among its members. One of their main findings is that public displays could increase the perception of unity as well as communities' interests. Some of their findings, e.g., suggestions for content and features of such displays, support the perspective of the community and delivery of local messages. A report of their findings on the usage of a digital community notice board is available in [296].

Studies exploring the impact of digital notice boards on communities have been conducted in several settings. As mentioned previously, Taylor and Cheverst [360] look into notice board practices in a rural village and informed the design of the Wray Display, a community photo sharing display aimed at understanding how digital displays can help to support communities in suburban areas. Churchill et al.'s Plasma Poster Network [71] looks into how displaying social media impacts relationships among co-located colleagues in their workplace. The CoCollage [223] aims at cultivating community in a cafe by showing posters and quotes and hence enhancing awareness, interaction, and relationship among people. Of particular interest is Huang et al.'s [149] finding that people spend less time learning about system capabilities when it is not supporting current use practices. The users' desire and interest to use novel systems needs to be taken into account [238]. This emphasizes the necessity of embedding existing routines in novel system to support its use. Note that in our view, this entails not only catering to users' needs, but also reflecting on PNA providers' motivations.

4.2 Sharing Public Display Space

Traditional public displays are a very common way of communication and they are ubiquitous in our environment. They scale from post-it notes on an office door telling people who stop by 'back in 5 minutes' to graffiti on a train making a political statement. This form of communication is very effective and observers will not be even aware of reading the signs in many cases, but they still do. Such public displays are an example of invisible technologies that allow transparent use, as Weiser suggests [378].

In our work we are particularly interested in public displays for information dissemination and for one-to-many communication. Spaces for these forms of communication and publicly sharing information can be found in many places, such as grocery stores, cafes, and restaurants but also in city administrations, public libraries, universities, and schools. Such places provide space that is visible, accessible, and frequented by people. Examples are notice boards and walls on which people are allowed to attach posters, and windows as well as doors where people can hang up flyers or notes. There is a huge variety of such PNAs, and many types of content can be found. Generally it can be seen that these displays have a function in their environment, and that the form factor of the display and the types of contents shared are influenced by the location, the owner, and the expected audience. In contrast to other forms of communication these displays support the following properties: (1) dissemination of content that is mainly locally relevant, (2) addressing of the receivers by selection of space, and (3) forcing information and content on people that pass through a certain space.

Traditionally, posting information in PNAs also had the function of personal communication from one individual to many receivers. However, this function of public displays has lost importance with the digital social networks and the World Wide Web. Popular forms of content include sales, housing, job and service offers, events, promotions, lost and found, and advertisements, all of which are at least showing one of the characteristics above. In order to derive an understanding of how to create digital displays that provide new flexibility and cost-effectiveness and at the same time retain the qualities of the analog PNAs, we investigate several issues further.

4.2.1 Value Propositions of PNAs for Stakeholders

In the optimal case, PNAs provide value to all stakeholders, including the people owning the space, people providing content, and people observing the content. First, it is important to identify stakeholders for a PNA installation, and their motivations. In many cases there is an interplay between interests, incentives, and value propositions.

Consider the following scenario of a notice board for classified ads in a supermarket. The content provider (e.g., a customer who wants to sell a bike) has the opportunity to reach people in the local community. The observer (e.g., another customer looking for a used bike to buy) via PNA becomes aware of the product they are looking for. The supermarket provides customers with a further reason to visit the store. Values have to be seen in the greater context of the PNA and

its place of use. Here issues such as exercising control over what content can be placed, by defining access to the display and implementing means for restricting content placement to a certain group play a role. The following questions help to identify these issues:

- Who is allowed to post?
- Who decides what content is appropriate?
- What content can be posted?
- What is the motivation for the owner to allow posting?
- What is the motivation for content providers to post?
- What is the motivation for observers to look at the content?

4.2.2 Content Exchange and Access Control

The utility of public displays depends on the fact that displayed information is useful. In general, content is posted to the display and removed or overwritten after a certain time. The mechanisms used both for posting and removing content as well as for enforcing usage policies are essential success factors. Often it is desirable that content can be created ad-hoc, without specific knowledge or tools (e.g., writing with a pen provided at the board onto pre-printed cards). Yet, the simpler the means for content provision, the more likely that spam or inappropriate content will be posted.

Control mechanisms for content creation and posting restrict the use to different degrees, ranging from having the notice board in a public area (where users posting content may be observed) to explicit approval of content. We suggest the following questions to identify and structure these mechanisms:

- How can content be placed onto the display?
- How flexible/easy is content creation?
- What tools are required to create content?
- What is the process for content approval?
- How is content removed or overwritten?

- How is access control implemented for content providers and viewers?
- How are viewers supported to help them remember content?

4.2.3 Learning from Practices in the Analog World

Many different types of PNAs are in use in different places. This multitude evolved over a long time and many of their properties fulfill a certain need. Similarly, many different kinds of content are publicly posted. Here, too, a long tradition exists for creating and designing content, ranging from artistic expression to minimalistic presentations. In this chapter we aim to identify these rich characters of different displays, content types, and related practices. For deriving design implications for digital displays, understanding practices and the rationale behind these is very valuable. In particular we are interested in the communication aspects that such displays facilitate.

A further important aspect of public displays is that they have a potential function for the community. By posting information in publicly accessible space but in a specific location, a clear addressing to the local community is made. Here it seems interesting to uncover functions that displays have that go beyond the communication of individuals or groups. To learn about the practices we ask the following questions:

- What practices have been established around sharing on a display?
- What are the reasons for these practices?

4.3 Study Design

To answer our research questions we ran a two-week field study during summer of 2010 in four different cities in Switzerland and Germany. The field study involved observational studies (photo logs) and subsequent interviews with the people responsible for the observed PNA, i.e., shop-owners and personnel.

4.3.1 Observational Studies

We aimed at observing a wide variety of locations, displays, and audiences. We looked for any kind of institutions, stores, and restaurants/eateries that displayed public notice boards. Due to the labor-intensive nature of the work, we opted for a convenience sampling of the observed sites (places were located along our work routes), allowing us to regularly visit these places over the course of four weeks. The observed locations were within the local neighborhoods surrounding the universities and central stations of Lugano (Switzerland), Essen, Düsseldorf and Munich (all Germany). An overview of the locations can be found in Table 4.1.

After choosing suitable locations, we identified the persons in charge of the PNA we wanted to observe. We introduced ourselves, explained the purpose of the study and asked for permission to take pictures of the PNA. We provided a written description of the study and explained that all data collected would be used for scientific purpose only. While most people immediately agreed to permit the study and even showed interest in the results, some of them first had to check with central management and asked us to report to the management every time we returned. In two locations we were refused permission to conduct the study, as the management felt that this would strongly intrude their customers' privacy.

After we were granted permission, we visited each location on consecutive working days over the course of roughly four weeks, each time taking several pictures of all postings. Pictures were mainly taken in the morning (on the way to work), during lunch break, and in the late afternoon / evening (on the way back home). We tried to make sure that pictures were taken at comparable times of the day. In total, four researchers were involved in the study, each one being assigned a fixed set of locations. Due to scheduling constraints it was not in all cases possible to take pictures on consecutive days. However, we made sure that for each location at least 10 picture sets were taken within maximum four weeks.

4.3.2 Interviews

After our observational study we conducted interviews with people in charge of managing the displays. Those were not necessarily the display owners, but also store managers or regular staff. With the interviews we aim at understanding a range of issues surrounding PNAs: the shops' motivation for having a PNA; the practices for adding, editing, and removing content; any restrictions as to what customers were allowed to post; any problems with the displays; and whether people could imagine substituting the 'analog' display with a digital version.

Table 4.1: Overview of Study Locations.

ID	Name	Description	Obs. Int.	Type	Curated
E1	Turn Headshop (Rack)	Retail	x x	ED	-
E2	Turn Headshop (Door)	Retail	x x	ED	-
E3	Diocese (Office)	Church	x x	ID	x
E4	Diocese (Entrance)	Church	x x	ID	x
E5	Supermarket	Retail	x x	SCD	-
E6	Supermarket	Retail	x x	UCD	-
E7	University Cafeteria	Public Bldg./Gov.	- x	UCD	x
D1	City Administration	Public Bldg./Gov.	x x	ID	x
D2	Adult Education Center	Public Bldg./Gov.	x x	UCD	-
D3	Public Library	Public Bldg./Gov.	x x	ED	x
D4	Child Services	Public Bldg./Gov.	x x	ED	-
M1	Supermarket	Retail	- x	SCD	-
M2	Supermarket	Retail	- x	SCD	-
M3	Supermarket	Retail	- x	SCD	-
M4	Supermarket	Retail	- x	SCD	-
L1	Supermarket	Retail	x x	SCD	-
L2	University	Public Bldg./Gov.	x x	UCD	-
L3	Bakery	Service	x x	UCD	x
L4	Church	Church	x -	ID	x
L5	Supermarket	Retail	x x	SCD	-
L6	Cafe	Service	x -	ED	-
L7	Hairdresser	Service	x -	ED	-
L8	Bar	Service	x x	ED	-
L9	Cafe	Service	x x	ED	x
L10	Pharmacy	Retail	x -	ED	-
L11	Bookstore	Retail	x x	ED	-
L12	Red Cross	Public Bldg./Gov.	x x	ID	x
L13	Laundry	Service	x -	ED	-
L14	Church	Church	x -	ID	x

Abbreviations: Obs: Observation, Int: Interview, Curated: Display is curated
 SCD / UCD: Scaffolded / Unscaffolded Classifieds Display,
 ID: Information Display, ED: Event Display




Figure 4.1: Data Walkthrough: (a) Photos, field notes, and transcribed interviews were printed and affixed to walls. (b) Analysis and annotation of material by researchers. (c) Affinity diagram to identify themes in the data.

We conducted interviews in the locations that our observational study covered. However, for two locations we were not able to get hold of a person in charge. We also included additional locations with similar PNAs to gather further information. For the interviews we returned to the location and tried to identify the person (currently) in charge of the display, asking her or him to answer a set of 10 questions. We offered to return at a convenient time in case people were too busy to talk to us. As interviews took place during business hours, the interviewee’s time was in general limited. Consequently, we limited our interviews to maximum 10 minutes. We either audio-recorded the interviews with a voice recorder (in case people felt comfortable to do so) and transcribed them later or took handwritten notes during the interview. It should be noted that the interviews were limited to the parties who owned and ‘administered’ the displays for reasons of accessibility. Information we gained about the perception and use of these displays by passersby and other stakeholders in the interviews was conveyed to us by display administrators and therefore may reflect their particular interpretation of phenomena regarding the displays. For the purpose of this study, we relied primarily on our observations to gain insight into the practices and needs of passersby and other stakeholders to complement the more direct inquiry into the practices of display owners and administrators.

4.4 Data Analysis

We conducted an extensive qualitative analysis of the photographs and interviews collected. Because of the large volume of data generated by the study and the distributed fashion in which the data was collected, we designed a ‘data

walkthrough' analysis method. The goal of this method was to help team members become highly familiar with data collected by other members of the team, and provide a view on the data that would provide both a comprehensive overview of all of the data from all sites, as well as an individual detailed view of each site.

The team printed out photos from each day of data collection for each site and affixed these photos (in total 298) to walls and whiteboards in a single room. The photos were placed in chronological order, grouped by site. Interview transcripts and field notes were affixed along with the corresponding photographs. Five members of the research team then proceeded to 'walk through' the data, analyzing the photos, interviews, and notes, and writing observations on individual sticky notes that were used to annotate the data collections (Figure 4.1b). As the team discovered patterns and higher-level observations, these were written on a separate whiteboard.

After this exploratory phase, the team then used an affinity diagram to identify themes in the data and associate them with our research questions (Figure 4.1c):

- Who are the stakeholders?
- What 'characters' of displays arise and what determines them?
- What are current practices for sharing display space?
- What is the role of 'posting displays' in a space?
- How do the space, stakeholders, and content interplay?
- What are the needs of people who are posting and display owners and how do the displays satisfy them?

This was achieved by taking all of the observations and categorizing them to derive the general findings. The identification of findings was done as a group, and each observation was discussed as to how it might fit with other observations.

4.5 Findings

Based on the data collected and analyzed, we report in this section the findings, with a focus on current practices. We first identify stakeholders, describe characteristics of displays and content, and uncover the motivation for shared public displays. During the data analysis we did not discover any obvious differences in the data gathered in both countries.

4.5.1 Stakeholders and Motivation

The data from observations and interviews provide a clear indication of a number of diverse stakeholders involved in operating and using public notice boards and shared public displays. On a highest level we can discriminate three different groups: display providers and managers, content providers, and viewers.

Display Providers and Managers

Based on interviews (L1, L2, L5, E3, E4, E5, E6, D1, M1, M2, M3, M4) we discovered that the decision to install a PNA is taken on a higher management level, e.g., in store chains and public authority institutions, and hence each branch or store will have a PNA as standard inventory. Also public and ecclesiastic institutions see information dissemination as a part of their mission and use them to distribute important information about their current activities. In locations where venue and shop owners are running the place (L3, L8, L9, L11, E3, E4, E5, M1, M2, M3, M4), interviews revealed that the decision to have a display and how to use it is in the hand of one person.

The motivation to provide public displays is manifold: retail and service have them to increase customer satisfaction (interviews at M1, M3, M4, E6, D2), public authorities and ecclesiastic institutions mainly used them to disseminate information on their current activities (observations at L4, L14, interviews at E7, D1, D2, D3, D4), and some of them (interviews at E3, E4, E7, D1, D2, D3, D4) feel the need to have a space for third party content as long as it fits within the institution's scope and does not harm their reputation.

In interviews, we found that some venues (L3, E3, E4, M3) have a dedicated person in charge of the content approval, i.e., a notice board manager whereas in other places it is less formalized. In public and ecclesiastic institutions there is typically a dedicated manager, whereas smaller venues are more likely to distribute this role throughout the staff, i.e., each staff member can act as manager.

Content Providers

We see two distinct groups of content providers: classifieds providers and third party advertisers. Both groups seek to distribute information to the target audience. People living in the vicinity of the venue or its frequent visitors can be seen as classifieds providers, seeking for 'matchmaking' opportunities, e.g., students exchanging books, people offering / looking for housing, or selling furniture. The content often defines how long one can expect a poster to remain on the board.

While classified providers are mostly individuals, third party advertisers are usually affiliations: church, government, business, musicians, non-profit, or other. All of them have a common goal of reaching a large audience and advertise in the vicinity of the target community's physical center, e.g., music events have multiple posters at music-oriented bars and universities, church-related events appear within its parish's locality, and even third party ads on government public display are topically focused. Interestingly enough, some of the venues take on the role of the third party advertiser and try to blend in with the rest when advertising their own events (e.g., L8).

There is an inherent tension between display owners and content providers as both rely upon each other (e.g., a PNA without content is not interesting and a person providing content cannot do so without space). The best way to minimize the risk of conflicts is to create a shared understanding about venues' board expectations (e.g., it is clear what content is expected on a certain display). It seems that this is quite common for PNAs as there were very few reports on abuse of the displays (interviews at L1, L5, L8, L9, L11, E1, E2, D1, D2, D3, D4).

Viewers

The motivations for viewing content ranges from clear information needs (e.g., someone looks for a place to rent) to accidental reading (e.g., waiting at the bus shelter and reading the posts in lack of any other occupation). Viewers are typically related to the location (e.g., they work or live close by) and may act at some time as viewer and at some other time as content provider. Many PNAs are located near high-traffic areas with guaranteed waiting time, e.g., next to printers or copiers, whereas other locations use them for decoration, e.g., bars. In cases where people are waiting, it is very likely that they browse through the PNA's content. Claypool et al. provide supporting evidence [73].

4.5.2 Displays and Content

During our observations and data analysis we discovered a number of different display types that are targeted to specific types of content. In the following we discuss typical groups that are commonly in use for PNAs (Figure 4.2).

Scaffolded Classifieds Display Our observations indicate that retail stores and supermarkets favor a well-organized arrangement of their PNA (e.g., L1, L5). These areas are highly scaffolded with preprinted cards provided at the




Figure 4.2: Types of Displays: (a) Scaffolded Classifieds Display (b) Non-Scaffolded Classifieds Display (c) Information Display (d) Event Display

display, which can be filled in and inserted into several rows of slots. Their content is in general informal and hand-written and sometimes includes tear-aways (e.g., name and telephone number). Content creation is very fast (in the order of a minute). Typically, content providers are asked to provide a date to later remove outdated content. Content not fitting the scaffold is attached next to the board. The content usually has a high turnover.

Un scaffolded Classifieds Display These displays are characterized by the absence of prescribed structure leading to flexible and ad-hoc posting. Typically they are not well organized. Content in any form can be placed at any position, even if it fully or partially occludes other content. Interviews showed that for most of these displays (e.g., L2, L8), there is no particular person in charge to check and remove posters placed in improper place or with unwanted content. These displays reflect the self-service nature of the postings. Content posted on such displays is in general similar to the aforementioned displays, with less structured layout, mixed sizes of posts, more colorful and more event-related. We discovered such displays at university, the adult education center, and a grocery store (L2, L8, D2).

Information Display As part of their information duty, many institutions, churches and libraries provide curated PNAs (observations at L4, L12, L14). They are characterized by formal, mostly professional content, including ads and events. In general they have a smaller number of postings compared to the above-mentioned types. Content is thematically focused (even if from third parties) and often applies to a larger vicinity. There often is an approval process through existing/formalized organization networks. These PNAs typically have a means for prohibiting unauthorized postings (e.g., by having a glass front pane, see Figure 4.2c).



Figure 4.3: Retrieving Information: (a) Flyers (E1), (b) Posts with takeaway tabs (E6).

Event Display Observations showed that bars and retailers offer event-focused PNAs (E1, L6, L8, L9). They are characterized by professional ads (posters, flyers), are thematically tightly scoped (e.g., techno events) and contain mostly third party content. The content is usually colorful, sometimes chaotic (depending on whether scaffolding is provided) and often provides some form of urban aesthetic.

4.5.3 Managing Content and Supporting Memory

Different mechanisms exist that help viewers to remember content they have seen. We came across 3 types of practices with regard to supporting memory: (1) information that is meant to stay on the board and where viewers are expected to remember essential information, (2) content that is completely taken away, and (3) content that offers parts with contact details that can be taken away. To encourage others to take information with them, current practices include providing multiple copies (e.g., flyers) or posts with takeaway tabs (Figure 4.3). If viewers have to remember content, it is important that this is as easy as possible, e.g., by providing an easy-to-remember URL. With the wide availability of mobile-phone cameras we also see a further practice conceptually combining (1) and (3). As people take photos of public display content, they take the information with them and at the same time leave the content for others.

We found different practices of post management with regard to cleaning displays.

- **Expiration Date:** The most popular practice we discovered for cleaning content is the removal of outdated content. This can be done either with a posted expiration date (e.g., after 30 days of posting) or with an implicit expiration date (e.g., the date of an event).
- **Complete Cleanup:** Another practice is entirely erasing PNAs based on a regular schedule, e.g., monthly. We found this practice for PNAs that contained too much content for manually selecting stale content (L2).
- **Curated Content:** Especially in municipal and official institutions we found that displays are often curated. Content is usually submitted at the reception or sent via mail, and has to be approved prior to being published. We observed that curated displays are in general very tidy and posts are not attached above each other so that they obscure other content.

4.5.4 The Role of Shared Displays in a Space

One function of displays is that they tend to create a central location for community activity (observations at E1, E2, E7). Such communities may be geographical, cultural, and also religious. Even though platforms such as Ebay or Craig's list exist, means for locally exchanging goods are still of interest. Heavy items are especially popular, since shipping costs can outweigh their price sometimes. Notice boards offer a convenient way of offering items to a local community where it is likely to find interested persons who are able to fetch items personally. Additionally, we found content that is relevant to the expected community visiting the location. E.g., in a music bar (observation at L8), PNAs mainly included promotion for music events (parties, live shows, concerts, etc.).

Further, boards are often used as community support tools. An example is the adult education center, where the PNA is mainly used to exchange study-related material (e.g., people offered their course material to people in lower courses).

4.5.5 Interplay of Space, Stakeholders and Content

We found that PNAs are often placed in high traffic areas (e.g., next to the main entrance of a location or in spots with waiting times) and that this had a strong impact on turnover in content. The same appeared to be true for communication and information hotspots. PNAs were often placed next to opinion boxes, store hours, space for prospects, and content from the display owner itself.

Most interesting though, is that posts are very local in nature. We found that content with no relation to a certain location (e.g., products, movies) appeared only rarely. In contrast, content seemed to be locally highly relevant. An example are posts on babysitting, as these people would usually not cover too large distances. This finding is also supported by the fact that information on local events is often posted in multiple locations in close proximity (e.g., in various stores in the pedestrian area). Similarly, boards often seem to express the identity of a venue. During our interviews we found that certain types of content is removed (e.g., political content, certain music events) if the venue could not identify with the content.

4.5.6 Needs of Content Providers and Display Owners

We observed that PNA owners often have an *agenda*. This in general correlates with their motivation for having a PNA (e.g., information obligation, dissemination of culture). Interestingly, persons in charge of managing the PNA sometimes serve as a gatekeeper, hence supporting certain events in an altruistic way as they feel that competing events are already sufficiently well publicized. As an example, the public library D4 refuses to announce Mardi Gras events but instead favors independent theaters.

Further, many places provide *support for posting*. Such support includes tables for writing posts as well as scaffolding in the form of structured cards and pens.

In addition, we found that *flexibility of content creation* is supported in many places. Whereas especially scaffolded displays provide standard cards, most of the display owners allow in parallel placement for customized postings. Thus, homemade posts, professionally printed or colored to make them more visible and eye-catching, can also be found. Another example we observed was that it seemed to be OK for most display owners to extend the provided display space in order to fit more content, e.g., by attaching content next to the designated posting area.

We also found many places that allow *multiple copies* of the same item of content to be posted. This was observed for different types of PNAs. On unscaffolded displays posters are freely tiled up next to each other, making a larger area more distinguishable from other parts of the board. This practice was observed also with scaffolded displays. Multiple posts of the same ad/post appear on different PNA locations and sometimes with small or no differences that are not easy spotted: in cases of non-professional handwritten posts, multiple posts are hard

to recognize. Only by reading them carefully and comparing contact details can these be noticed. Obviously, having multiple copies of a single post increases the chances of the post being seen. The motivation for the content providers to put up multiple copies of the post is that they are afraid that their post is going to be removed for some reason.

Finally, we found that there are often difficulties of indicating a *venue's expectations* of the board. When asking for problems with PNAs, several display owners reported on discussions with people in order to explain what content is or is not allowed. Content that is removed is mainly political content, offending or provocative content, competing content, and content that does not fit the agenda of the display owner. There are very few examples where we found explicit notices that certain types of content are not allowed or have to be approved. One example was at the adult education center, where a notice stated that teaching content is to be approved first.

4.6 Design Implications

Our observations identified important stakeholders and some of their needs, as well as current processes and structures that help them to fulfill their goals, or that emerge as a result of their joint activities. Naturally traditional displays are not technically networked. However we could observe a connectedness on a conceptual level. Connectedness was exhibited by enforcing the same policies across a set of displays or by having the same design and structure across all displays. Our design implications assume that the stakeholders' needs will still have to be met within an open digital public display network, but that the concrete processes and features do not need to match existing solutions one-to-one. We propose five broad principles, and offer specific ideas on how to apply them.

4.6.1 Design for Specific Uses of Notice Boards

One of the clearest patterns in the observations was the great variety of notice boards and surfaces, depending on the purpose they were meant to fulfill, particularly the type of content they present, and on the preferences of the stakeholders. This indicates that no single design may be an appropriate replacement for all the current uses of notice boards. A system that displays large numbers of classified ads should optimally look different and work differently from one that seeks to

create awareness of local events. The notice board styles, posting form factors, and content described above offer a starting point for developers to target their systems; designers should complement this general description with specific knowledge and understanding of their particular users. A digital notice board designed for a chain of coffee shops might for example draw on elements from unscaffolded classifieds and event displays and would consider the type of clientele, the activities that take place in the shop, and the chain's design aesthetic and brand image.

The design should take into account that users interact with more than one display. In a department store they may see one at the entrance, one in the elevator, and one where they try on some shoes. The design should capitalize on the fact that users will be exposed to a network of displays. The displays should be in a desired order if the technologies used allow for it, e.g., creating a story within the display content across the physical display network.

4.6.2 Respect the Neighborhood Focus of PNAs

Of the thousands of postings we observed across the 29 sites, every one, with almost no exception, was related to the local area or to the community that used the space. The vast majority of these were classifieds relevant to a limited area or notifications for local events. Across the very different styles of notice boards, it is clear that the neighborhood is the audience for and the source of the postings. If a digital notice board is to play the same role, it should be based on postings and ads that have a clear connection to the place or neighborhood, not on centralized advertising campaigns. Especially for networked display systems, where technically there is no limitation for the distribution we recommend to design posting procedures in a way that supports locality and to restrict the content to a certain neighborhood. Such mechanisms are most likely a property of the system architecture as well as a part of the actual design.

Note that this is not simply a matter of geographical restriction, but of community identity, even more though when using display networks. The handwritten nature and tear away tabs of many classified postings create an indirect physical connection between advertiser and reader, which may lead to greater intimacy and trust. Furthermore, the aesthetic of a concert poster communicates both the intended audience and is used by that audience to provide a shared group identity. A digital notice board might attempt to capture the direct intimacy of handwritten notes by allowing posters to record short video messages as part of postings.

4.6.3 Support the Emergent Profiling of a PNA

The different styles of notice boards we observed reflect, among other things, different agendas on the part of the owners of the space. Some are considered decorative or as a way to express the identity or to support the image of the place. Some are appreciated as providing another activity for customers to engage in, increasing the importance of the venue to the community and potentially even attracting customers. Perhaps the most important agenda, however, is in disseminating information the notice board owners have an interest in or sympathy with, such as a library posting notices for classical concerts or an adult learning center showing notices for trading course books.

Board owners actively use their control to promote all of these priorities. At the same time, we found that in many cases they have difficulty articulating their agendas, and when asked, many were unable to give a good explanation for why they even offered a PNA. Even more importantly, they rely on third-party posters for most of the content. These factors mean that the actual profile of the PNA is usually not pre-defined, but emerges from the interplay of the interests of the board owners and users.

This goes to the core of reasons why venue owners are willing to offer free advertising space and must be taken very seriously. In a globally networked digital system, the owner of the venue may no longer be the owner of the board and may have a much less direct control of the content that is posted. A digital replacement must give the venue owners overall control over the board profile, and designers must recognize that this is usually not a choice that is made only once, but a day-to-day activity. Features could be provided to the managers or owners of each venue allowing them to easily choose which postings to allow or remove (and perhaps learning and automating their patterns over time) and to oversee the overall appearance of the board.

4.6.4 Design for Flexibility of Input

Across all locations we found an impressive variety of posts, comprising hand-written notes in various sizes and colors on a set of different materials, printed notes enhanced with images or maps, and professionally designed advertisements. We recommend that digital notice boards provide means for preserving this flexibility, supporting:

- **Ad-hoc posters:** For people (coincidentally) approaching the display a mechanism has to be provided, which allows for on-site creation of notes (e.g., pre-defined templates, standard input devices such as mouse, keyboard, or touch).
- **Sophisticated posters:** People preparing content in advance need to be given means to easily transfer it to the board locally, e.g., via a scanner, a USB stick or Bluetooth and in networked settings remotely, e.g., via a web interface.
- **Professional posters:** In a similar way, people distributing professionally designed content (e.g., flyers) need to be granted easy access to the board.

As a result of the flexibility of current approaches we see that the entry barrier is kept low, hence attracting a large number of posters. We believe that for globally networked display systems the success heavily depends upon speed and ease of use of the content creation mechanisms.

4.6.5 Support Retrieval of Information

The opportunity to retrieve information is a crucial prerequisite for the success of classifieds as well as for event promotions. In the design of this mechanism two basic options exist: providing a pointer for the user to retrieve (e.g., a URL or a phone number) or to provide a copy of the content to retrieve (e.g., a flyer). Traditional systems use a combination where a pointer is retrieved (e.g., by tearing off a piece of paper) and this also contains a minimal summary of the content. In digitally networked systems, many mechanisms exist to support users equipped with mobile devices in taking content with them, including users taking photos of the content, to sending this information via SMS or Bluetooth, by providing a QR code, or simply by printing it out. Whereas the primary motivation is to preserve information such as a date, a name, an address, or a telephone number, additional implicit information is embedded with the takeaway information: the number of missing takeaways, e.g., indicates high interest. Similar information could be provided by displaying the number of poster downloads, or even by restricting the number of possible downloads on digital displays.

4.7 Discussion

Traditional displays are all around us and interaction with them is a common activity. The communication characteristics we can find for traditional displays are in many ways complementary to the Internet. The importance of using public displays to address a larger number of people with general content has declined with the advent of digital social networks and the World Wide Web. However public displays play a major role in addressing groups of people that can be found in a certain location. We believe that understanding practices in the use of traditional analog displays can provide valuable insights for future generations of globally networked public displays.

By analyzing the usage of a large number of public displays situated in different contexts and based on follow-up interviews with stakeholders we collected and described various practices related to public displays. In particular we were interested in the various stakeholders and their motivation in the use of public displays. Additionally we investigated how content can be provided and removed and how access control is implemented. Based on our reflection and understanding of the data, we suggested a set of design implications for digital display systems.

The central design recommendation is to take the context of a potential display into account: the passersby, the neighborhood and community in which the display is situated, and the display owners' expectations with regard to content. Beyond this flexible content creation, content posting, and content control are central to allow a broad set of people using it. To increase the effectiveness of posted content it is important that viewers can take information with them.

As a result of our findings we implemented a research prototype (see Chapter 9) that allowed the identified challenges to be investigated in the digital world.

4.8 Case Study: Car Advertisements

Apart from PNAs, further spaces exist that could be shared for advertising in the future – for example cars. Cars are ubiquitous in today's societies and, given their size and shape, offer potential space for placing highly-visible ads. Vehicles are used and seen by people, and can therefore lead to many points of visual contacts between people and cars. For example, a pedestrian looking at cars driving by, a driver or passenger driving behind another car and looking at its rear, and a person walking by a row of parked cars and glancing at them.

Until now, very few privately-owned cars have become advertising spaces and commercial vehicles have mostly advertisements that promote the services and products of the company that own them. Advertising on vehicles has remained mostly static, in the sense that the ads do not adapt to their immediate situation or context. The first approaches to offer dynamic advertising content involved mounting electronic displays on cars where the content might be location dependent, e.g., taxis in the Boston area²⁵. With advances in display technologies, we anticipate that there will be new ways for attaching or embedding electronic displays into car surfaces, such as car doors or the sides of vans and trucks. This will create new opportunities for advertising. We argue that context will play a major role for efficiency and acceptance of advertising displays.

This section explores the potential of cars as dynamic and contextual advertising space. In the following, a summary of the results of an online survey with 187 drivers and their attitudes towards car-based advertising are presented. Furthermore, we describe the requirements, concept, and design space for a contextual advertising system for vehicles.

To understand users' expectations, motivations, and constraints for providing their car as a platform for advertising, we conducted an online survey and follow-up interviews with car owners and a car fleet manager.

4.8.1 Privately-owned Cars: Online Survey

In our survey, we were particularly interested in (1) when users want to make their car available for advertising, (2) product categories users would be willing to show advertisements for, (3) the acceptance of different technological solutions, (4) possible forms of compensation for showing advertisements on private vehicles, and (5) parts of the car where user's would prefer ads to appear on.

The survey was completed by 187 persons (130 males) with an average age of 27.4 years. The majority of participants drive compact-sized vehicles (128) and medium-sized vehicles (41). Recruitment for the survey was done via email.

First, we asked for the *preferred compensation schema* (Figure 4.4, left). We suggested different types of compensation, including monthly allowance, discount on car purchase, benefits from third parties (coupons for cinema/concerts, etc.),

²⁵ Taxis with roof mounted signs that change their content according to the GPS position, see <http://www.clearchanneltaximedia.com/products/taxi-tops-digital-smart.asp>, last accessed February 13, 2013




Figure 4.4: The participant preferred a monthly allowance and discounts on the car purchase as compensation (left). Easy to detached foil and fixed labels were favored as technical solutions (right). (5-Point Likert scale, 1=not interesting at all, 5=very interesting)




Figure 4.5: Most users would agree to always advertise Tech on their cars.

and coupons for free fuel. On a 5-point Likert scale (1=not interesting at all, 5=very interesting), 155 of 187 (82.9%) participants were interested (gave a rating of 3 or higher) in a monthly allowance, 139 (74.3%) in a discount upon purchasing the car and 137 (73.3%) in coupons for free fuel. Only 50 (26.7%) subjects were interested in receiving coupons from third parties as compensation.

Second, we asked *when users would like to provide their cars as advertising space*. Users were asked to choose from several options: at any time, while driving, while I am not close to my car, and while I cannot see my car. We found that 105 subjects (56.1%) did not care when ads are displayed, while 37 subjects (19.8%) preferred not to be close to their car while it was showing ads, 24 subjects (12.8%) did not want to see their car while it was advertising (Figure 4.5).




Figure 4.6: Users do not like to advertise for controversial categories; advertisements on culture, music, sports, and clothing are preferred.

Third, we evaluated the *effect of the advertised products* on the attitude of car owners towards advertisements. We provided users with a selection of different categories, including music, clothing, fast food, beverages, tobacco, firearms, politics, culture, eroticism, and sport events. We found that users preferred not to show advertisements on their cars for controversial categories (Figure 4.6). These included tobacco (124 subjects, 66.1%), eroticism (144 subjects, 77.0%), firearms (151 subjects, 80.8%), and politics (116 subjects, 62.3%). On the other hand, they strongly preferred showing advertisements on culture (173 subjects, 92.5%), music (167 subjects, 89.3%), sport events (161 subjects, 86.1%), and clothing (151 subjects, 80.7%).

Fourth, we looked into the *acceptance of different technical solutions* (Figure 4.4, right). Using a 5-point Likert scale (1=not acceptable at all, 5= totally acceptable), we found that the most popular solution was easy-to-detach foil (92.0% with a rating of 3 or better), followed by fixed labels (55.6%). Rather unpopular were painting (43.9%) and roof-mounted equipment (19.8%).

Finally, we wanted to know which *parts of the car participants preferred to have the advertisements appear on*. Participants used a 5-point Likert scale to evaluate the following locations: rear/trunk, rear windshield, rear side windows, sides/doors, roof, front lid. The result shows that most participants favor

the sides/doors (83.9% with a rating of 3 or better), followed by the vehicle's rear/trunk (78.6%), roof (77.0%), rear side windows (67.4%), rear windshield (60.9%) and front lid (60.4%).

The survey revealed the following findings with regard to privately-owned cars:

1. Incentives are essential for convincing users to provide their cars as advertising space. Compensation for placing ads on cars in preferred order is: monthly allowance, discounts upon car purchase, and coupons for free fuel.
2. More than half of the participants would not care if they were close to their car while ads were being shown. However, one third preferred not to have the ads shown when they were close by or could see the ads themselves.
3. Users want to stay in control of the products advertised, especially those falling under controversial categories, such as eroticism, firearms, and tobacco.
4. For technical solutions, easy-to-detach foil was rated by far the most preferred option while roof-mounted equipment was rather unpopular.
5. The side doors, trunk and roof are the favorite areas on the car for advertisement placement.

4.8.2 Privately-owned Cars: Qualitative Interviews

In order to gain further understanding of the car owners' views on car advertising, we conducted four follow-up interviews with people who participated in the online survey. The subjects were two students (each 21 years old), a contractor (50), and a real estate administrator (45).

First, we asked them about concrete *incentives and values*. All subjects would be interested in a monthly allowance (80-120 euros monthly), two of them would also like to receive coupons for free fuel (80-100 euros monthly). Second, we asked about *how to select the products to advertise for on their cars*. Two favored a category list, where they could select groups of products, and the other two favored positive lists for a more fine-grained selection. All of them wanted to be able to update the list at any time. Third, we asked about *privacy concerns regarding contextual and location information* that might be transmitted about their car. Three of the subjects were not concerned at all, adding that they were participating in payback programs anyway. However, all four subjects stated that their GPS data should not be provided to the advertisers.

4.8.3 Company-owned Cars: Qualitative Interview

To complement the view of privately-owned cars with that of large companies that own cars for business use, we conducted a telephone interview with a fleet manager of a pharmaceutical company, who is responsible for a fleet of approximately 600 cars with a replacement rate of 200 cars/year. During the interview, we discovered issues and concerns regarding the hypothetical introduction of car advertisements in the fleet. One of the manager's primary concerns was the additional administrative workload. This additional workload is partly due to the need to equip cars with the required technical gear and partly due to the necessity to define products that should not appear as ads on the fleet's vehicles. The manager also felt that a monthly allowance of about 10% of the monthly leasing rate of the cars would be an attractive incentive. A final and interesting finding was with regard to the acceptance among the employees using the cars. Since cars are a part of an employee's compensation, as well as a status symbol, the manager felt that it would be essential (especially among sales personnel) to allow the employees to have a strong influence on the decision of whether or not to allow advertising on their cars.

4.8.4 Requirements

Based on the findings from our survey and interviews we derived the following requirements for context-aware car advertisements.

Means for expressing preferences. One central requirement is that car owners have the power and means to specify what advertisements are shown on their car. It is also essential that the system is easy to use. Potential solutions include giving the user a list of products and brands that they can accept or deny (not individually but as a whole). Asking the user to provide preferences for each product may lead to a lengthy and tedious process and asking for each product category may be too general to accurately specify their preferences. Therefore, we suggest creating a preferences editor based on a combination of product groups and brands. This approach scales well to large numbers of advertisements and appears to satisfy the users' concerns.

Simple technical solution. The technical solution for advertising must be very simple to apply and remove. A major change to the appearance of the car (e.g., display system on the roof) is also not welcomed. The required effort by the car owner for administrating, configuring, and maintaining such a display surface

must be minimal. One potential technical solution is a self-contained display unit consisting of a magnetic foil (to make it easily attachable to the car door) with a bendable display (e.g., eInk display) on top and built-in communication and context-sensing. This unit would potentially be powered by harvesting energy.

Defining context. In the case of contextualized advertisements on vehicles, the advertising message would be broadcasted to all people who focus directly or indirectly on the car. Thus, the audience is not limited to a specific user but rather to a group defined by the context. We identified the following information:

- **Spatial:** Location is an important factor when dealing with context-aware ads. Using knowledge of the geographic location of the display, ads can be targeted at the audience expected in that area. It may also be possible to estimate the potential view context (e.g., distance of the observer).
- **Demographic:** From spatial information, demographical data can be derived, e.g., data on the people living, working, or staying in the area.
- **Temporal:** For the advertiser, knowing when ads are shown is important. The temporal context can either be date- and/or time-based or even further abstracted (e.g., during rush hour, lunch break, after work).
- **Weather:** For several product groups (e.g., cloth, drinks, food), the current weather conditions, temperatures and forecasts may be relevant to decide which ads to display.
- **History:** Past information on displayed advertisements can facilitate the decision-making process when other context-information is unavailable.

Gathering context. There are two steps to gather the context: (1) collecting objective data from the context using sensors embedded in the display, and (2) complementing the sensor data with external information. For the first step, location data via a GPS sensor is an obvious option but additional sensors that provide information about proximity (e.g., WLAN or Bluetooth scan) or weather conditions are also possible. Other types of context information (e.g., traffic, weather, demographics) may be gathered from external sources, e.g., of the Internet or from a database.

Matching context and content. One of the challenges of context-aware advertising systems is matching the ad content to context (e.g., time, location, etc.). In a traditional advertising business, the marketing group decides when, where, and under what circumstances to display the content. With context-aware systems

designed to make the decisions about content selection automatically, context information would have to be available in real time. Instead, we recommend giving the advertisers the freedom to specify the conditions in which an ad is shown. Then, the current context of a car and its display(s), the car owner's preferences, and the conditions specified by the advertiser could simply be matched to the context data.

4.8.5 A Platform for Dynamic Contextual Advertising

As a proof-of-concept, we implemented a prototype platform that supports contextual advertising on dynamic displays. The system is set up based on a client-server architecture. In order not to rely upon the car's internal computational resources, we implemented a thin client, requiring as few computations and low data storage as possible.

The two core functionalities of our car advertising system are (1) the matching of the current context in which the car is situated with the available advertisements, and (2) the smart and efficient exchange of information between multiple cars and the ad server. In the existing prototype, we follow an exact matching approach. Our implementation supports a push-pull communication model: cars send an update of their current context to the server periodically. At the same time, the server itself constantly calculates the best available ad campaign for every client's car. The best-fitting advertisement is then pushed to the client. This is technically realized by dynamically modifying the web page used by the client to display the content.

The front-end of our system consists of two web interfaces – one for the advertiser and one for the car owner. The advertiser's interface provides a way of specifying both the content of the advertising itself, as well as additional settings that define the context in which an ad has to be shown. The decision algorithm needs to know where, when, and under what circumstances to display the ads. The web interfaces were implemented using PHP and currently support the input of locations, date, and time, as well as weather and temperature conditions. The advertiser can select the desired area(s) on a map that is realized by the Google Maps API. The advertiser's interface also supports the upload of pre-designed content (e.g., images). In the car owner's interface, the people providing their cars as advertising space can specify their preferences and tell the system what ads they will and will not allow.

The client side includes three different agents for gathering the context: (1) a GPS agent reporting the location of the client, (2) a time agent providing the time, and (3) a weather agent pulling current weather information and weather forecasts. The server keeps a record of which ads were shown where and on which car. This information could be used as a basis for the car owner's compensation and incentives could be provided in order to change their parking behavior (e.g., parking in an highly frequented area will provide a higher revenue than parking in a backyard). For advertisers, the information can be used to assess how successful their campaigns are.

4.8.6 Related Work

Examples of mobile applications that support the selection of advertisements based on context are SMMART [196], Ad-me [146] and SmartRotuaari [269]. The first two approaches use Bluetooth, the third one uses WLAN in order to determine the location of a device. All systems support the concept of making decisions about the to-be-displayed content based on context information.

When it comes to distributing ads on mobile clients, two approaches prevail: push (sending information automatically to the client) and pull (sending the information only upon the client's request). Both approaches are discussed by Finin et al. [294] and Varshney [367]. Another interesting approach was taken by Castro and Shimakawa [87]. They try to migrate from a push to a pull model, where users can decide if they wanted to receive more ads after an initial notification.

Finally, research on context-aware advertisements for public displays is closely related to our work, since both types of advertising media try to attract large audiences with different personal backgrounds and interests. Work that focuses on maximizing the exposure of ads based on context is presented by Karam et al. [173] and Kern et al. [178]. Both approaches take the interests of the people in the vicinity as context in deciding which ads can be tailored to the target group.

4.8.7 Summary

Our survey shows that cars have the potential to become an interesting, dynamic, and context-aware advertising space in the future. However, it is essential that users stay in control of their vehicles, have an appropriate means for specifying

their preferences and are provided with an agreed-upon compensation model. We recommend that location, demographics, time, weather, and history be considered as important contextual parameters. As a proof-of-concept, we implemented a client/server-based platform that provides a web interface for advertisers and car owners to enter content and preferences. On the client side, several context agents provide real-time data about the current context of the advertising medium. On the server side, a matcher and content creator assemble and prepare web pages for rendering.

Several challenges arise while trying to provide dynamic car advertisements. Although static deployments, such as labels or printed advertisements, are widely available, robust technological solutions for easily setting up context-sensitive advertisements on cars are still under development. National traffic regulations can also lead to additional challenges. For example, according to a telephone interview we conducted with a representative of the German Technical Monitoring Association (TÜV), the use of light-emitting or light-reflecting materials is only allowed after passing a strict approval procedure, due to their high risk of distracting other drivers on the road.

Chapter 5

Audience Behavior

As public displays are being widely deployed there is an emergent need to understand how the audience behaves in front of these displays in order to appropriately chose their location and the content to be shown. This need will become even more pronounced as interactivity leads to considerable changes in audience behavior. People will be attracted and spend more time with public displays, probably interacting for several minutes. As a consequence they may be so immersed that they do not realize anymore what is happening around them. This creates the need to place displays in a way such that users are not put at any risk (e.g., if a display is deployed in a shop window facing the sidewalk, people interacting might not realize that they are just a step away from the street with cars frequently passing by). Furthermore, if the space where the display is setup is highly frequented, an interactive game may lead to that passersby constantly bump into people who have started to play. As space and display owners understand the way the audience behaves they can use this information to refine their placement of the display or the type of content shown.

This chapter focuses on audience behavior with regard to different display shapes. Due to the deployed display technologies, today the majority of displays are flat, rectangular, and framed. However, we learned from history, that there are many successful forms of non-planar displays. One popular form are columns. Freestanding columns have the benefit of high visibility due to their concise and elevated shape, and can also provide more screen real estate on the same floor space. In addition, columns were extensively available inside buildings for structural reasons. Famous examples for ancient cylindrical displays are

Trajan's column in Rome or columns in the Hathor temple in Egypt. Even today the most popular form of non-planar displays are cylindrical screens, such as cylindrical bulletin-boards, inflatable columns used at events, or street furniture columns used for cultural information, public announcements, and ads (Morris or Litfaß columns). With advances in technologies, bendable displays will allow nearly any surface to be turned into a display and hence allow displays of almost any arbitrary shape and size to be created for no additional costs compared to flat displays.

We opted to investigate cylindrical displays as one possible form of novel, arbitrary-shaped public displays. Hence we present a prototype of an interactive cylindrical display and report on a user study comparing user behavior in front of classical, flat displays and cylindrical displays showing the same content (we use the terms cylindrical display and column interchangeably). This chapter makes the following contributions. We present an interactive cylindrical display, which reacts to passersby by adapting the content according to their movements. Based on a lab study we show that people in front of cylindrical displays tend to move more and explore the content from a wider range of standing positions. We explain how this influences the way content should be designed. Finally, we discuss advantages and disadvantages of flat and cylindrical displays based on the findings of our study.

This chapter is based on the following publication:

- G. Beyer, F. Alt, J. Müller, A. Schmidt, K. Isakovic, S. Klose, M. Schiewe, and I. Hausen. Audience Behavior Around Large Interactive Cylindrical Screens. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada), CHI'11, pages 1021–1030, New York, NY, USA, 2011. ACM

5.1 Related Work

In recent years, a variety of prototypes of non-planar displays have been developed. Most public deployments, however, still use planar, framed displays and their designs are based on implicit assumptions that may not hold for non-planar displays. While a number of studies about audience behavior for planar displays exist, their results may not generalize to non-planar displays. Finally, the specific case of cylindrical displays introduces the new category of semi-framed displays.

5.1.1 Interactive and Public Displays

Inspired by early work such as Media Spaces [42] and the Digital Desk [380], a variety of interactive and public displays have been proposed. The Plasma Poster Network [71; 72] is a network of touch enabled public displays installed in hallways supporting office communication. Similar to our approach, many of such displays use computer vision to enable interaction from a distance. The seminal work of Myron Krueger explores vision-based interaction with large displays where the displays react to the users' movements [189]. In ReflectiveSigns [251], a network of displays reacts implicitly to the audience's viewing behavior. The Intelligent Kiosk [70] is an early example of an interactive Avatar on a public display that reacts to passing people. Similarly, the AmiQuin [235] is a virtual mannequin in a shop window that reacts to the audiences' body movements. Finally, Malik propose vision-based interaction techniques with multiple fingers for distant displays [211].

5.1.2 Designs of Non-Planar Screens

Organic user interfaces [143] have been proposed as computer interfaces that use non-planar displays for input and output. Lin at al. present the i-m-Tube, an interactive tubular display [207]. The first version uses a single projector and a convex mirror to create a back-projected cylindrical display, which supports multi-touch. Benko at al. present Sphere [33], a multi-touch enabled spherical display. They discuss unique properties of spherical displays and multi-touch interaction techniques for such displays. Some of the properties for spheres are similar to properties of cylindrical displays (for instance the user can not see the whole display at any time), while others are not (like a smooth transition between vertical and horizontal surfaces). Lee at al. [203] present foldable displays like fans, newspapers, scrolls, and umbrellas. Their prototypes are implemented using projection, infrared markers, and the Nintendo Wii remote camera. They discuss affordances of different shapes and possible interaction techniques. Volumetric displays appear to show content inside the display volume. Grossmann at al. focus on such displays' unique properties and show how they can be used for collaborative applications [128; 129].

Regarding cylindrical displays, a number of commercial designs exist. Many commercial cylindrical displays use rotating LED's (Kinton²⁶, Dynascan²⁷), where the drive system has to be well adjusted to get a jitter-free image. Other technical solutions include static mechanical designs including rings of LED modules (Barco²⁸) and projection-based setups (Pro Display²⁹). In [32], Benko outlines challenges when designing gestural interactions with non-flat surface computing interfaces derived from the development of three prototypes (a gesture-enabled flat display, a sphere, and a dome). Challenges identified include walk-up-and-use functionality, linking heterogeneous devices, usability from multiple directions, and compelling applications.

5.1.3 Assumptions of Current Designs

Most current deployments of interactive public displays use planar, framed displays enabling interaction either through touch or body gestures. The CityWall [281], for example, is a large multi-touch display installed in downtown Helsinki that supported browsing photo collections. Worlds of Information [164] is an extension of the same system that includes touch interaction with 3D spheres of photos. Magical Mirrors [238] is a deployment in downtown Berlin where passersby could see their own mirror image on a display and interact with virtual objects through body movements. Such designs for flat displays usually start from a number of implicit assumptions that need to be questioned for non-flat displays. Many designs assume that: (1) people stop walking before they interact, (2) users can perceive the content of the entire screen at any time, (3) users can see what other people do when interacting with the display, (4) shoulders are usually parallel to the display, (5) the position centrally in front of the display is preferred, and (6) content is not distorted. For example, touch displays are obviously difficult to use while walking, and even for Magical Mirrors a walking user would quickly leave the camera view. For the CityWall, users could easily scale photos so big that they hide the view for others, which would be a problem if users cannot observe the effects of their actions on other parts of the screen.

²⁶ Kinton 360° Display Systems: <http://www.kinton.de/en/products-solutions/360-display-systems.html>, last accessed February 13, 2013

²⁷ Dynascan 360° Displays: <http://www.dynascanusa.com/ds360.html>, last accessed February 13, 2013

²⁸ Barco website: <http://www.barco.com/en/>, last accessed February 13, 2013

²⁹ Pro Display: <http://www.prodisplay.com/dual-projection-screens-digital-signage.html>, last accessed February 13, 2013

Similarly, users start playing ‘soccer’ by throwing photos around, which would be very different if one could not observe the actions of other users. Vogel and Balakrishnan [372] use a parallel shoulder position as an indicator of user interest, but it is not clear if this is valid for non-planar displays. Also, the distortion of content could render many current interfaces unusable on non-planar displays.

5.1.4 Audience Behavior in Front of Flat Displays

Any observations of audience behavior towards public displays have been conducted using planar displays, and it is unclear whether existing findings generalize towards non-planar displays. Scott et al. [325] explore how people use regions on round tabletops, which share some properties with cylindrical displays, but where the display itself is flat. Huang et al. [147] show that for existing planar displays, very few people look at them or stop in front of them. This may also be related to the fact that many flat displays are installed either on walls, orthogonal to walking direction, or above eye height. Cylindrical screens in contrast are often installed directly in the way of passersby. Müller et al. [255] show that what people expect on public displays depends on the immediate surroundings. As columns are usually freestanding, expectations may differ. Brignull and Rogers [52] and Peltonen et al. [281] observed that users are waiting for their turn to access the display. For cylindrical displays, there is no central position where the user could ‘own’ the display, but all positions around the display are equal. Therefore, turn taking could be different.

5.1.5 Effect of Display Frame on Audience Behavior

A major difference between planar and non-planar screens may result from their frame. In this context Manovich [215] presents a theory of the imprisonment on the viewer’s body by the screen apparatus on one hand and the requirements of the image perspective on the other hand. In cinema, e.g., the body of the viewer is confined to a seat and the head is aligned to forward view, hence providing the best viewpoint. The same is true in classical arts (e.g., a painting in a museum) where the viewers seem to position themselves centrally at some distance in front of the screen. We show that this effect does not appear in front of non-planar screens. Another interesting conception from Manovich is the description of the screen’s frame as a clearly defined rectangle, constituting a ‘viewing regime’. Anything outside the frame can be ignored by the viewer, while they immersed




Figure 5.1: Prototype of a Cylindrical Display: The display consists of an acrylic rear-projection screen, 4 foil mirrors, and 8 standard projectors with 1024x768 pixels. The flower content was designed for passing-by interaction.

themselves into the content inside the frame. Pinhanez and Podlaseck [286] discuss advantages and disadvantages of frameless displays, also claiming the significance of the frame to serve as a reference for the viewer to orient inside the scene and position them accordingly. Cylindrical displays introduce a new category in between framed and frameless displays. In this case of semi-framed displays, a frame is provided on the top and bottom, but not left or right.

5.2 Prototype of a Cylindrical Display

For our research we use a prototype of an interactive cylindrical display. Though our prototype is projector-based, we envision that future versions will be based on bendable display foil. In the following, the hardware and software setup of this display is described.

The prototype of the cylindrical display consists of a cluster of eight standard projectors, four foil mirrors, ten PCs, and a rear-projection screen. It has a height of 2.2 meters and a diameter of 1.5 meters. The 4:1 projection screen is 1.1 meters high, has a diameter of 1.3 meters and a resolution of 2048 to 512 pixels. Each cluster element projects onto a mirror reflecting the projection onto about one quarter of the screen. For a viewer independent image blending we use a special rear projection screen with a low gain factor. To detect position and movement of the passersby around the column, a sensor interface is installed above the column consisting of a high-resolution camera and a 185° fisheye lens. The hardware setup is depicted in Figure 5.1.

The raw projection on the cylindrical screen is heavily distorted due to the curvature of the screen and the projection angle. To correct these distortions we use the calibration technology described in [39]. This software is also used to resolve the blending function for the real time correction between overlapping regions of adjacent projectors. The visualization software for the displayed contents on the column screen is based on a distributed rendering system. To enable user interaction, we implemented a motion tracking software, using OpenCV. The motion tracker uses frame differencing to detect motion, and calculates the angle, speed, and pixel distance of moving blobs from the column. The Kalman filter is used to smooth these trajectories.

We developed a VRML-based application framework that enables us to display the same kinds of interactive applications within the coordinate systems of flat as well as round screen shapes.

5.3 Hypotheses

As no studies regarding user behavior towards cylindrical screens exist, it is important to understand how people move and behave around them. This knowledge can then be used as a basis to develop applications that exploit the properties of the new format and investigate more elaborate topics such as multi-user interaction. As there are situations where only a single user is interacting, we decided to concentrate on single-user interaction, leaving multiple users for future work.

Based on informal observations of colleagues and visitors we pose three general hypotheses that characterize behavior around cylindrical displays. First, we assume that users walk more when interacting with cylindrical than with flat displays. If true, this is an important property, since many flat displays are

designed for people standing in front of them. Designs that work for people standing, with rather high complexity and small fonts, may not work for people walking. Second, we assume that while users seem to have their shoulders parallel to flat displays, they would have their shoulders in a certain angle to columns. This would be an important property for gesture-based interaction: while users can use both arms equally for flat displays, one arm would be turned away from a column, making symmetric gestures difficult. Further, it would be difficult to move any arm against the direction the user is facing, so an entirely different gesture set would need to be designed. Finally, we hypothesize that due to the more active engagement users would spend more time interacting with columns.

Hypothesis 1 – Users walk more when interacting with the column.

- 1a: Users walk longer distances when interacting with the column.
- 1b: Users spend more time walking when interacting with the column.
- 1c: The position of users has a higher variance when interacting with the column.

Hypothesis 2 – Users position themselves with shoulders parallel to the flat display but not to the column.

- 2a: The users' shoulder position is less often parallel to the display when interacting with the column while walking.
- 2b: The users' shoulder position is less often parallel to the display when interacting with the column while standing.

Hypothesis 3 – Users spend more time overall interacting with the column.

In addition, we had several hypotheses for the viewing behavior of participants. We hypothesized that participants would look more often at the cylindrical display, but for shorter bursts. We also hypothesized that participants would look at the left half of the column when walking clockwise and the right half when walking counterclockwise. While these hypotheses were formed before the design of the user study, we additionally conducted a post-hoc analysis of the data to explore further observations we made.




Figure 5.2: User Study Setup: room with the cylindrical display (a), room with the flat, rectangular display (b).

5.4 User Study Design

To test these hypotheses we conducted a study comparing single users' behavior in front of interactive flat and cylindrical displays. The study was conducted at our lab over the course of two days. In the following we describe the study design and setup and report on the recruiting process as well as the study procedure.

We opted for a lab study due to the following reasons. (1) For the anticipated measurements, a highly controllable environment was required for statistical data analysis. This would have been difficult to achieve in public due to the high amount of external influence and fragile, technical equipment. (2) We used cameras to assess the users' behavior during the study. This would have been a major issue in public due to privacy reasons. To create an authentic scenario where people behave in a semi-natural way, we created a situation where participants were free to visit different rooms containing various exhibits. Thus, we created a controllable, yet still realistic scenario, in which we did neither influence people's behavior nor were they aware what we were measuring.

5.4.1 Setup

For the study we prepared 4 rooms at our lab each of which contained a prototype. Two of the rooms contained 'fake' prototypes, which were functional (one was an interactive flat screen with content that reacted to the viewers' head movements and facial expressions, the other one was a non-interactive dome projection and showed a movie) – however their only purpose was to create a more realistic situation and to distract from the displays under investigation.

To compare the flat and the cylindrical display we designed two similar rooms, one containing the cylindrical display, and one the flat display (Figure 5.2). We designed the two rooms as similar as possible while still preserving the situation in which a flat or cylindrical display would normally be deployed. The distance between door and displays was equal and both displays are approached from the same angle. While the doors would open in different directions, both displays would visually appear at the same position in the room.

The flat display is equipped with one standard projector and uses the same rear-projection material as the cylindrical screen. For the size of the flat display we had to choose between a display that has the same size as the visible display area of the column when a user is standing in front of it, and the total size of the cylindrical screen ‘unrolled’. Being in the situation to choose whether to deploy a flat or cylindrical display one would usually have to choose between two displays that take the same floor real estate and be in the first situation. For this practical reason, we decided to test a flat display with the same size as the visible area of the column first and leave the second situation for future work.

The flat display uses a webcam as a sensor. For user tracking we setup four webcams in each room, which detected the users’ position in the room. From the webcams we created a synchronized, time-stamped video file using OpenCV. In order to analyze and manually annotate the video files we attached markers to the floor hence creating a grid of $60 \times 60 \text{ cm}^2$ squares.

5.4.2 Participants

We aimed for a diverse round of participants to represent people in public spaces. People were recruited from bulletins in the surrounding neighborhoods (train stations, shopping malls, university campus, etc.) and mailing lists in the days prior to the study. In total, 15 people participated in the study (10 males, 5 females). The average age was 32.7 years; participants were students, artists, taxi drivers, office workers, and technicians.

5.4.3 Procedure

The participants were invited to the labs and asked to report to the doorman, who led them to the laboratory. We started with an initial briefing. First, we had the participants fill out a demographic questionnaire (asking for age, gender,

nationality, profession, and, if applicable, their major). Second, we explained to them the setting of the study. They were told that there were four exhibition rooms containing different project works. We did not tell them which or how many expositions were in each room. Instead, we asked them to simply walk through the rooms and spend as much time there as they wanted. In order to minimize compliance effects, we stressed to them that there was no minimum and no maximum time we expected them to spend in each room. We told them to fill out a one-page questionnaire available in each room after they finished their visit. After that, we wanted them to move on to the next room until they visited all four rooms. Furthermore, we told them that several cameras were setup in the rooms, needed for the purpose of the study (we did not tell the participants that some prototypes were interactive).

After the initial briefing we guided the participants to the first room and explained to them that arrows attached to the walls and floors would lead the way through and to further rooms. Though we surveyed the people in the different rooms (using webcams), we did not interrupt or talk to them during the study. Since we used a within-subject design (all participants saw all expositions) the order of the exhibits was counter-balanced among the subjects. Whereas the ‘fake’ prototypes were always seen in the same order, we switched positions for the rooms containing the flat/cylindrical screens. After the participants finished the fourth questionnaire, we conducted a semi-structured interview with them. All interviews were audio-recorded. Finally we debriefed the participants and explained to them the purpose of the study. Participants received a compensation of 10 euros.

5.5 Data Analysis

Complete time-stamped and synchronized videos from four different perspectives of all user behavior were recorded in both rooms. All user positions were transcribed. A coding schema was devised containing codes for walking, shoulder angle, and head direction. The complete video recordings were coded by two independent raters using the Mangold INTERACT coding software. In order to ensure inter-rater reliability, one video was coded by both raters and Cohen’s Kappa was computed for all codes. Inter-rater reliability was satisfactory for walking (Kappa=.61) and shoulder (Kappa=.67) codes (we considered Kappa \geq .61 substantial agreement [199]). Unfortunately, reliability was not satisfactory for head direction codes, such that we had to drop all hypotheses related to viewing

Table 5.1: All comparisons with Wilcoxon signed-rank test with paired samples.

Measure	Column		Flat	
	mean	std. dev.	mean	std. dev.
Distance walked (m)**	47.3	24	21.2	13.7
Time spent standing (%)*	44.9	25.9	62.8	16.6
Mean duration of stops (s)**	3.5	2.3	9.9	13.1
Max duration of stops (s)*	12.9	11.9	38.7	52.7
Total time spent (s)**	97.7	53.4	172.8	138.6
Time spent with shoulders parallel (%)**	41.5	21.3	69.5	17.2
Time spent with shoulders parallel while walking (%)**	22.1	10.6	46.3	16.1
Time spent with shoulders parallel while standing (%)	70.0	26.1	82.0	18.6
Stops per minute (1/min)	8.3	4.3	6.8	3.3
Mean distance from display (m)	1.5	.23	1.7	.41
Variance in location (Rows)**	5.6	1.4	2.3	.96
Variance in location (Columns)**	3.7	1.4	.93	.57

(*:p<.05, **:p<.01, ***:p<.001)

behavior. In addition to testing the hypotheses, a post-hoc analysis of various variables of user behavior (e.g., duration of stops) was conducted. Because we could not ensure normal distribution for all dependent variables, all hypotheses were tested using the non-parametric Wilcoxon signed-rank test with paired samples using a level of significance of p<.05. The AttrakDiff questionnaire was evaluated using the associated software, and semi-structured interviews were partially transcribed to cover non-repetitive user statements.

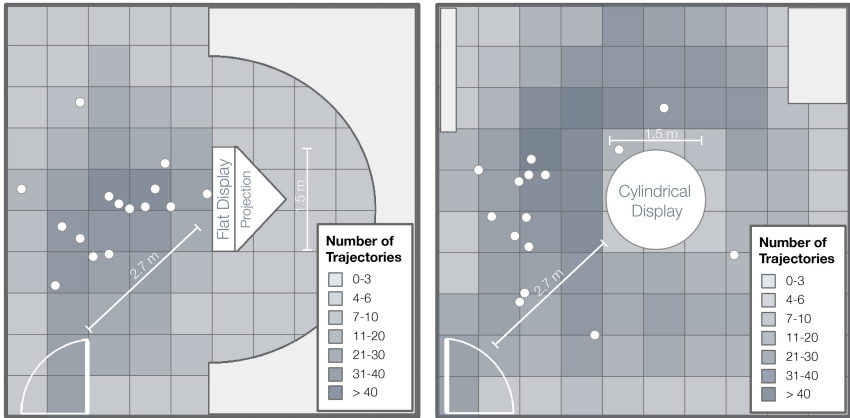


Figure 5.3: Room Layout for the User Study: room with flat display (left), room with cylindrical display (right). For the data analysis we separated the rooms into a grid consisting of $60 \times 60 \text{ cm}^2$ squares. The heat map shows the participants' trajectories of each square. White dots indicate the participants' first standpoint after entering the room.

5.6 Results

In the following, we report on the results of our data analysis. An overview of the results is presented in Table 5.1.

5.6.1 Walking

All hypotheses regarding the walking behavior of the participants (1a, 1b, and 1c) are supported by the data. Participants cover significantly more meters with the column screen (on average 47.3 m) than with the flat screen (21.2 m) ($p < .003$). With the column, participants spend most of their time walking (only 44.9% standing), while with the flat display, they spend most of their time standing (62.8%) ($p < .027$). Also, they cover many more different locations with the column. The variance of location in rows is more than twice as high (5.6) for the column than for the flat screen (2.3) ($p < .001$), and the variance in columns is more than three times as high (3.7) for the column than for the flat screen (.93) ($p < .001$).




Figure 5.4: Typical User Behavior in Front of the Cylindrical Display: (a) The user interacts immediately. (b) The user keeps walking around the display. (c) The shoulder position is mainly orthogonal throughout the interaction phase.




Figure 5.5: Typical User Behavior in Front of the Flat Display: (a) The user approaches the display. (b) The user stops in front of the display. (c) The user maintains a parallel shoulder position, even during interaction.

5.6.2 Shoulders

The hypotheses of the angle of shoulders towards the display are partially supported. Overall, participants spend most of their time with their shoulders parallel to the flat display (69.5%) (Figures 5.4 and 5.5).

In contrast, with the column participants spend most of their time with shoulders not parallel (only 41.5% parallel) ($p < .001$) (Figures 5.4). If this data is divided by whether participants are walking or not, the picture is more diverse. When walking, participants spend 46% of their time shoulders parallel to the flat display, while this is the case for the column only 22% of the time ($p < .001$). When standing, there is a trend towards participants having their shoulders parallel

to the flat display more often (81%) than to the column (69%), however, this difference is not significant ($p < .069$). While hypothesis 2a is supported by the data, hypothesis 2b is not (compare also Table 5.1).

5.6.3 Time Spent for Interaction

Hypothesis 3 assumed that participants would spend more time with the column than with the flat display. This hypothesis is not at all supported by the data. Indeed, participants spend almost twice as much time with the flat display (2:53 min) than with the column (1:38 min). A post-hoc analysis shows that this is significant ($p < .002$).

5.6.4 Post-hoc Analysis

When the participants stop for the first time in front of the flat display, most participants stop in a very narrow area in front of it, while the locations of the first stops are more diverse for the column (Figure 5.3). The same figure also depicts a heat map of where participants move. This nicely shows the ‘sweet spot’ area in front of the flat display, which is much more spread out for the column. Also, participants spend slightly more time until they first stop with the column (10.4 s) than with the flat display (6.7 s), and then continue walking sooner (after 5.7 s) than with the flat display (22.5 s). These differences however are not significant. Participants stop slightly more often (per minute) with the column (8.3 stops) than with the flat display (6.8 stops), although this difference is not significant. However, if participants stop, they stop for more than twice as long (9.9 s) with the flat display than with the column (3.5 s), and this difference is significant ($p < .006$). The longest time participants spend without moving is even more extreme. While with the column display, participants spend on average no more than 12.9 s on a single location, participants spend up to 38.7 s in a single location with the flat display. Additionally, participants seem to approach the column a little bit closer (avg. distance 1.5 m) than the flat display (1.7 m), however, these differences are not significant.

5.6.5 User Experience

To evaluate the user experience, all participant filled in an AttrakDiff questionnaire [135] for both displays. The results reveal that from a pragmatic as well as from a hedonic point of view, the participants consider the column to be of higher quality. Considering the attractiveness, the column is rated very attractive whereas the flat display is rated medium attractive. All results from the questionnaire are not significant, which might have been caused by our rather small sample size.

5.6.6 Semi-Structured Interviews

The participants do not seem to feel bothered by the experimental setup:

“As soon as you start playing around with the media, you forget that you are being observed.” (P2)

Interestingly, many participants experience the flat display like ‘television’ (e.g., P3, P4, P7), while they experience the column as being more ‘dynamic’: ‘

“With the column you are more actively engaged. The flat display had a television feeling. The round was more dynamic. TV is more static.” (P3)

They also feel more passive in front of the flat display:

“You stand in front of it like in front of a TV. With the column, you are more actively engaged. It was more like a TV situation; you walk in front of it and think, what happens now? With the column, you immediately had the feeling, you do more with it.” (P7)

Some participants consider the column as more comfortable:

“It was more spatial. [...] It is better if you can move around a fixed point than back and forth.” (P1)

“For me, round is more comfortable than square. It was more organic.” (P3)

Note that all interviews were held in German and have been translated post-hoc.

5.7 Discussion

Our study reveals that single user behavior differs significantly between planar and cylindrical displays. Due to the setup of the study, there are certain limitations with regards to the generalizability of the results.

1. We only evaluated single user interaction. Hence no conclusion can be drawn towards behavior in multi-user scenarios.
2. Unlike in the lab, for real world deployments there are many factors influencing movement patterns, such as streets, buildings, or pedestrian traffic. However, as columns are often deployed in small free places, spatially similar to our lab setting, we believe that the observed movement patterns generalize to such situations.
3. With regard to the content, we believe that the findings are valid for most similar interactive content. Yet, very different content (e.g., a fixed framed video), unsuitable for the column, might lead to different results.

In the following we outline and discuss major findings.

5.7.1 Moving Around the Column

Compared to a flat, rectangular display, a column does not provide any boundaries left or right. This lack of borders indeed seems to have the effect of making viewers move freely around the column (Figures 5.3 and 5.4). Participants spend most of their time walking and covered significant distances, looking at the column from various locations and stopping quite often, but only for relatively short times, on average only 3.5 seconds. This leads to much more diverse body postures as opposed to a flat display. Only about 30% of the time is spent with shoulders parallel to the display, and even when participants stop, the shoulders are not parallel 30% of the time.

5.7.2 The Sweet Spot

The data shows that for flat displays, there is a relatively small area in front of the display where participants get themselves in a frontal position ('the sweet spot').

Though this looks like an obvious finding, we are not aware of any prior studies that support this assumption with data. This area is positioned centrally in front of the display, about 1.5 meters away from it. Participants seem to approach this area quickly after entering the room, and stop in this position with their shoulders parallel to the display, facing the display frontally (Figures 5.3 and 5.5). From this position, they can see the entire screen from the best perspective, while the entire frame was still in the visual field. These observations are well aligned with the theory from Manovich [215]. Participants stop in the same position on average for up to 39 seconds, with one participant not moving for more than three minutes. When they move again, almost half of the time they still have their shoulders parallel to the display, moving back- and forwards, or sideways without turning. Overall they spend most of their time standing and do not cover great distances.

5.7.3 Time Spent for Interaction

As opposed to the hypothesis, participants spend significantly more time with the flat display than with the column, almost twice as long. This is an interesting finding, which we believe is worthwhile to be investigated in more detail. As people stand still to identify with the screen image and ignore the real world, they might easily spend more time than if moving around, exploring also the physical space.

5.8 Recommendations

Based on these results we derive recommendations for designing content on planar and cylindrical displays.

5.8.1 Design for Walking Interaction

Cylindrical displays are most suitable to keep people in motion, even if they do not frontally face the display. Unlike flat displays where motion is often simulated in an unnatural way (e.g., ‘running’ in the same spot in front of a flat display using Xbox Kinect Sports track-and-field), cylindrical displays provide a way to support real motion. Thus, columns and their content should be designed for

walking, and be set up at sites where they can raise the attention of viewers by implicit interaction and encourage them to move on. A way to achieve this is to keep content simple, such that it can be perceived while passing by. Another good way is content that moves with the audience as they move around the column, adapting to the location of the viewer rather than requiring the viewer to adapt his location to the display. We propose that flat displays, because of their sweet spot, may be more suited for waiting situations and longer dwell times, and may support more complex content.

5.8.2 Place Columns in the Path of Users

Cylindrical displays are suitable to be located within the path of the users and can better be used while passing. For flat displays people tend to take a fixed standing position for interacting, probably blocking the way for others. In contrast, columns allow users to avoid bumping into other people by constantly moving around the display. As a result, content on cylindrical displays should adapt to passersby and could attract them towards non-crowded areas around the display.

5.8.3 Enable Gesture-Based Interaction

Previous research shows that three interaction techniques are highly suitable for the use on public displays: touch, gestures/body position, and mobile phones. Whereas touch and mobile phone based interaction is applicable for flat displays where people stop in front, the motion-fostering character of cylindrical displays makes them more suitable for gesture-like interaction. We observed that almost 60% of the time, the viewers' shoulders are not parallel to the display. Consequently they can not use both hands equally well for interaction as one arm is constantly turned away from the display. It may be best to enable interaction with one arm, and not requiring movement of the arm against walking direction. Another difficulty is the question of which arm is geared towards the display depends on walking direction and not on dominance of the hand.

5.8.4 Use Frameless or Semi-Framed Content

Digital columns are semi-framed. As a result they do not have a 'sweet spot' any more, hence posing another difficulty since designers and audience can not

use left or right frames to orient themselves anymore. Using framed content, like videos or images, on a column, of course can create a virtual frame (and thus probably another ‘sweet spot’), but that may not put the cylindrical shape to its best use. As explained in [286], frames serve to create a distance between the content and the surface (in this case, the column), while frameless content integrates smoothly with the surface. In order to integrate well with the column and exploit the frameless nature, frameless content is ideal.

5.9 Summary

Non-planar, digital displays have the potential to change the experience of public displays in urban spaces. Presenting content on such displays, creating an engaging user experience, and exploiting the new properties of these screens requires us to rethink the way we design content and applications. Observing users’ behavior is a necessary first step in developing guidelines for designing interactive columns.

In this chapter, we presented the same content on a cylindrical and a planar display and observed different user behavior. We found that users move more when interacting with cylindrical displays. The sweet spot of flat displays, where users tend to position themselves, does not exist for cylindrical displays. In our study, participants spend significantly more time interacting with the flat display than with the cylindrical display. These findings can influence the design of cylindrical displays in the following way. Content for cylindrical displays should be designed for audiences in motion. Furthermore, they are suitable to be placed in the path of the users. Gestural interaction should be possible, ideally using one hand only, because users may not have both shoulders parallel to the display. As they are only semi-framed, frameless content is best suited for cylindrical displays. If a design goal is for people to spend little time engaging with the display (e.g., public display in a busy area) then cylindrical screens may be more appropriate than flat ones.

We see five important limitations of this study and areas of future work. First, for this study we decided to compare cylindrical displays to flat displays that cover the same floor real estate. A follow-up study should compare both displays to a flat display of the size of the cylindrical display unrolled. Second, multi-user scenarios are an obvious application area for cylindrical displays. Because there is no sweet spot, multiple users can approach the display on equal grounds and do not need to take turns. For this reason users may also feel less observed

as if standing in an exclusively occupied spot. Third, while our experimental setup enabled the detailed investigation of motion behavior under lab conditions, behavior in public places might differ. A field study of the prototype will be the next step in investigating the properties of cylindrical displays. Fourth, in this study we compared display formats using one specific interactive application. It would also be interesting to evaluate user behavior with different applications. Fifth, we did not investigate gaze behavior or recall of content on the displays. It would be interesting to investigate whether the increased movement when interacting with columns leads to increased recall of content.

Chapter 6

Targeting

Measuring the effectiveness of public displays is at the focus of researchers, display owners, and marketers alike. Knowledge about how many people looked at a display, or more specifically (a certain region of) the content, and how it is perceived can be used not only as a measure for impact based on which advertisers could be charged, but it allows the content to be enhanced in real time, thus making it more interesting to the viewer. Today, technologies exist that enable pure audience measurement (e.g., simply counting the number of people that pass a display, see also the audience funnel [238]) and more recently, technologies such as age and gender detection based on analyzing the passerby's face provide more fine-grained information. However, this chapter is more concerned with how information obtained from audience measurement technologies can be used to better shape the image of the customer, and how this information can be fed back to the advertising system in order to create more powerful advertising.

In former decades, the customers' needs and interests were unknown until they were encountered at a customer touch point (e.g., in a store), so that marketers had no means to reach potential customers except to maximize the exposure of advertisements. Although good sales personnel tried to learn as much of the customers' needs and interests as soon as they entered a store (e.g., based on physical appearance, gender, age group), this information was not available for marketers. In recent years, technologies evolved which allowed to not only shape the picture of a customer on-the-fly but also to observe their (shopping) behavior over a longer period of time and make this information available for marketers. In the real world, the most prominent example (still) is shopper loyalty cards,

which provide means for keeping track of the items a customer purchased, hence allowing potentially interesting products to be selected. However, this method has several drawbacks. First, the user's cooperation is required to track what they buy, e.g., by presenting a shopper loyalty card upon checkout. Second, tracking can only be done on a coarse granularity since only the purchase can be identified but not the browsed items. Third, the provision of targeted advertisements is restricted to direct mail, as the user cannot be identified on-the-fly.

With the advent of the Internet, several of these issues could be tackled by assessing a user's click stream in online stores. Information could be aggregated implicitly so that no cooperation with the customer was required. Furthermore, it was possible to not only analyze what a customer bought but also what they were looking at or searching for. Finally, content can either be adapted in real time as well as be distributed through other channels, such as email or direct mailing.

As researchers and practitioners aim at applying these techniques to public displays, the major obstacle is the identification of the user. Users in public space are in general anonymous and technologies such as face recognition are, though available, difficult to deploy due to severe privacy issues. In recent years, various pervasive computing technologies emerged. Most notable are mobile devices, but also a variety of location and tracking tools based on different types of sensors (e.g., NFC readers, eye trackers, etc.) have permeated the advertising space, creating new opportunities and challenges for out-of-home advertisers:

- Audience sensing in the vicinity of a customer touch point (e.g., a public display) enables real-time audience measurements as well as adaptation of the content based on the users' interests.
- Ubiquitous technologies, such as mobile phones, are taking implicit and explicit user interaction with advertising content into consideration. Hence, a way of actively choosing interesting content is provided and at the same time, it is possible to learn about a user's potential needs and interests.
- Networked advertising environments allow such systems to not only be deployed for single stores, but also on a large scale across arbitrary stakeholders (e.g., store owners, display providers, advertisers).
- Even in cases where no identification of the user is possible, sensors can be used on-the-fly to detect what they are currently looking at and probably interested in, for example by analyzing their gaze as they look at different pieces of content on a public display.

In this chapter we look into how existing technologies can be used to provide adaptive (advertising) content on public displays via large scale advertising environments with the potential to increase exposure and uptake by passersby while at the same time preserve their privacy. We present CAdeT (*Context-aware Advertising Environment*), a platform where an advertising network of arbitrary size can be created. Our current setup allows scanners at points-of-sale and public displays to be integrated, hence providing means for adapting advertisements towards the users' interests and also letting users interact with the system.

First, we present the approach of *adaptive user profiles*, which allows users to setup a profile based on a self-assessment. This profile is then automatically enhanced with aggregated data about the 'real' shopping behavior. On one hand, the users' subjective interests are taken into account. This is especially interesting for advertisers since it reflects the users' intention and provides a strong indication whether or not a user is interested in a specific product. On the other hand, implicit observations might hint on interests not mentioned explicitly or reveal changes in behavior over time. Hence, a more detailed image of the user's interest can be drawn and advertisements can be targeted more precisely rather than simply statically assigning users to target groups (see Section 3.2.3 for an introduction to targeting).

Second, we present a *prototype system*, which uses pervasive technologies such as Bluetooth to create profiles in an anonymous and unobtrusive way. Linking profile information to the MAC address of a Bluetooth enabled device, such as a mobile phone, allows the identification of a user without the need to store any personal data such as name, gender, or address. Furthermore, a user may easily stop using the system by switching off the Bluetooth functionality of his phone.

This chapter is based on the following publications:

- F. Alt, M. Balz, S. Kristes, A. S. Shirazi, J. Mennenöh, A. Schmidt, H. Schröder, and M. Gödicke. Adaptive User Profiles in Pervasive Advertising Environments. In *Proceedings of the European Conference on Ambient Intelligence* (Salzburg, Austria), Aml'09, pages 276–286, Berlin, Heidelberg, 2009. Springer-Verlag
- J. Mennenöh, S. Kristes, F. Alt, A. S. Shirazi, A. Schmidt, and H. Schröder. Customer Touchpoints im stationären Einzelhandel – Potenzial von Pervasive Computing. *Marketing Review St .Gallen*, 27(2):37–42, 2010

6.1 Related Work

This section draws upon research from different areas, such as pervasive advertising, public display networks, interaction, and profiling / recommendation systems. In the following relevant work from these fields is being presented.

6.1.1 Targeted Advertising and Profiling

Previous research investigated how pervasive technologies can aid advertisers to deliver suitable advertisements to costumers in appropriate ways at the appropriate moment in time. Sharifi et al. show how public electronic displays could improve the efficiency of advertising systems if the display is aware of the identity and interests of the audience [329]. Ranganathan and Campbell examine the delivery of ads to selected costumers using pervasive technologies [293]. BluScreen is an intelligent public display that selects and displays ads in response to users detected in the vicinity through the presence of their Bluetooth enabled devices [280]. Karam et al. use the BluScreen system and focus on how content for the public display can be selected based on the audience's preferences and areas of interest [173]. The advertisement selection tries to maximize the exposure of as many advertisements as possible to as wide of an audience as possible without any prior knowledge and need of any user action. Müller et al. detect and analyze audience attention in order to adapt digital signs [255]. Thawani et al. propose an architecture for context-aware real time advertisements in a live broadcast stream [362].

One important area in pervasive advertising is profiling and recommender systems for targeting the customer. Davies et al. show how to generate profiles in pervasive environments based on user interaction [83]. Further research focuses on how available technologies can be used to provide more suitable and relevant information based on profiles. Bilchev and Marston introduce the concept of a distributed user profile constructed based on different profiles that can be used for personalized online advertising [40]. Shannon et al. present a platform to extract information from a Facebook profile for targeting personalized ads [328]. AdROSA is a system for automatic web banner personalization, which integrates web usage and content mining techniques [175].

Recommender systems have been subject to various research projects. Lawrence et al. present a personalized product recommender for supermarkets using PDAs [201]. Adomavicius et al. introduce a multidimensional approach for

incorporating contextual information into recommender systems [2]. Amazon introduced item-to-item collaborative filtering [208], a hybrid of content-based and collaborative filtering.

In the context of pervasive advertising, many projects looked into how profiles could be created and assessed for enabling targeted content. However, to the best of our knowledge no approach has been presented yet, which allows an aggregated profile based on a user's self-assessed profile and implicitly gathered information from the system to be generated. In contrast to existing systems we derive data from different types of interaction and show how this information can be used to integrate both information into a meaningful profile.

6.1.2 Public Display Networks

Public display networks have been mainly deployed for research purposes. However, none of them are commercially used for advertising. eCampus is a network of public displays installed at the campus of Lancaster University [352]. The Plasma Poster Network is a network of digital bulletin boards displaying community-generated content for informal networking [72]. Greenberg and Rounding present the Notification Collage, a groupware system that enables collocated and distributed colleagues to post and share information [126].

Display owners manage and own most of the current public display networks. With displays becoming cheaper and more widely available we envision that networks will be shared among different stakeholders. With our platforms we aim at laying the foundations to easily integrate displays from multiple owners and grant access to different stakeholders.

6.1.3 Interaction Techniques

Currently, commercial public display installations show mainly mere adaptations of analog content and do not provide means for interaction. In the future, however, we believe that the impact of advertisement on public displays may benefit from means for interaction. In this context the interaction can be intended (explicit) or not (implicit). With explicit interaction, the user has the goal to interact with the environment, in implicit interaction, even though the action performed by the user is not primarily aimed to interact with the environment, it is interpreted as interaction. The assumption in this situation is that the environment has a

certain understanding of the user's behavior [317]. The advantage of implicit interaction is that no extra effort is required from users. For both implicit and explicit interaction with interactive displays various modalities can be used.

With regards to explicit interaction on public displays, Sahami et al. use the flashlight on the mobile phone [335]. Furthermore, Ballagas et al. explore interaction using the camera of a mobile phone [24]. An overview on phone-based interaction is available from Ballagas et al. [25]. Vogel and Balakrishnan present an interactive public ambient display [372] that uses explicit (e.g., hand gesture and touch) implicit (e.g., body orientation, position) interaction techniques.

One main goal of implicit interaction in pervasive advertising is to gather additional information about the user. Analyzing a user's click-stream on the WWW is a common implicit tracking technique. However, tracking interaction with public displays requires additional sensing. Schmidt et al. describe how Bluetooth can be used to identify users in the vicinity of a public display [318]. An overview on how sensing can be used to obtain information about the potential viewers is presented by Schmidt et al. [318]. Müller et al. discuss mobile phone interaction techniques with public displays [220]. Reitberger et al. [302] and Meschtscherjakov et al. [234] looked into interaction with ambient displays.

6.2 Opportunities of Pervasive Advertising

In traditional advertising a one-way communication from the advertiser to the customer prevails (*push communication*). In this case, the advertiser is mostly unaware of the audience exposed to their messages. Current practices to estimate the audience include the German G value (cf., Chapter 7). With pervasive computing technologies a new communication channel between the advertiser and the customer is available, which allows the customer to become an active part of the communication process (*pull communication*) [54].

From a customer's view, they now have the opportunity to share information on personal interests with the advertiser, e.g., by using their mobile phone to explicitly input information about their interests to an advertising system. Furthermore, they can also interact with the advertiser through the system to obtain further benefits, e.g., coupons or additional information. The advertiser benefits from this interaction because they can explicitly and implicitly gather information on the customers' interests. For example, such a system can recognize certain customers as they enter the vicinity of a display or as they (implicitly) interact with

content (similar to generating a click stream in the WWW). A system can then recognize the content as being of interest to the user. As a result, the advertising system can now use both knowledge about the user in the vicinity of a display and information about the target group of an advertisement to select appropriate content. It can be presumed that a message with high relevance is potentially more valuable to the customer.

Besides these benefits for the customer we envision further benefits in terms of effectiveness and efficiency for the advertiser. For illustration of the first part we draw upon the Elaboration-Likelihood-Model [283] (see also Section 3.2.2). Within the ELM, the first determining factor is the motivation to process. Information on both the user's interest allows the advertiser to display the ads according to the customers' targets and to those customers with a probably higher motivation to process due to a higher personal involvement and relevance of the message. We therefore presume an increased effectiveness of the advertisement, as the probability is higher that the elaboration will follow the central route and therefore positive attitude changes will be more persistent.

In terms of efficiency today, large parts of the advertising budgets are spent on scattershot ads where the location is used as the sole parameter to reach a target group (e.g., the location of a display in close proximity of a school and the knowledge that lots of pupils pass the display on the way to the bus stop can be used to advertise toys or candies). Yet, scattering losses are inevitable since people who are not a part of the target group are exposed to the advertisements. Now with information that the right group is targeted, it can be assumed that a campaign's efficiency will increase.

This additional information can also add a new qualitative dimension to billing. Since showing an advertisement to the target group rather than to people that do not fit the target group is more likely to generate revenue, display providers could adjust what they charge advertisers based on the exposure to the desired target group. Hence, exposure time could either be maximized and the overall cost for the advertiser be justified, or exposure time could be minimized without scarifying the envisioned impact. With our system we envision to enable a new business model for outdoor advertising, where advertisers are charged based on the actual exposure of their content towards the audience.

The implications can be summarized as follows:

- Customers benefit by being exposed to advertisements fitting their interests better and by getting the opportunity to become an active part of the communication process.

- Advertisers can potentially increase their campaign effectiveness and efficiency. This entails increasing the probability of elaboration on the central route by reaching customers with a higher overall motivation to process. Furthermore, advertisers can reduce scattering losses through better targeting and by better measures for evaluation and billing.

The following scenarios aim at clarifying the potential of pervasive advertising in different real-life settings using different technologies.

Scenario 1 – Subway Station A public display is deployed at a subway station. People waiting for the train can connect to it with their Bluetooth enabled phones. The use range is limited to persons carrying phones that can operate software requiring Bluetooth, but advertisers can assume to interact with people having a technological background and who can afford advanced hardware. Since no other assumptions of the user base can be made, Bluetooth devices are the lowest common denominator. Users are identified by the unique Bluetooth MAC address and, thus, can be detected each time they pass by. This bears the potential for follow-up ads or repeated exposure. If the user starts interacting with the system, e.g., by filling out a questionnaire about their interests via a mobile application, the ads displayed could be matched to their interests or the interests of all users in the vicinity.

Scenario 2 – University The system is deployed at a university campus. Students at this university are used to paying with a student ID card in cafeterias that are equipped with an RFID chip and a barcode denoting the matriculation number. The system will read either of them; the RFID chip can even be used to detect users in the vicinity of a terminal. Hence, the usage is limited to students of this university; however, the students can connect to the system by just using the ID cards. Beside displaying advertisements on campus as stated above in the subway station scenario, displays can be used as information terminals allowing students to interact with the display privately. Displayed information and ads can be adapted to the student's interest and visited courses, e.g., new course material is automatically downloaded on the student's campus desktop. This individual desktop could be accessible not only on campus but online at any other place.




Figure 6.1: Example for Creating Target Groups.

Scenario 3 – Grocery Store The system is deployed at a grocery store. The retailer has implemented a loyalty program that is actively in use. Some customers use RFID loyalty cards and some use their Bluetooth-enabled mobile phone to participate in the loyalty program. Like in the above-mentioned scenarios, ads can be adapted to the customer. The retailer can use the customers' individual shopping history to point out interesting promotions to them on the display or trigger additional sales in the aisle ahead. Similar to the university scenario, touch screen displays can allow the customer to interact privately, e.g., to receive coupons, view additional product information, or to comment on items. Besides understanding the shopping behavior by analyzing the shopping carts, the retailer can gain further shopper insights, e.g., by tracking the shopping route in-store and the time customers stay in certain areas of the store via Bluetooth or RFID. This information can be used to improve ad sequences or the store layout.

6.3 Concept and Profiling

In order to more precisely target content, information about the customer is required. Such information is gathered and stored based on profiles. Profiles represent the user's attitude towards multiple areas of interest (categories) in a

multi-dimensional space R^N . Figure 6.1 depicts a (very simple) two-dimensional space based on which consumers could be assessed using the dimension attitude towards sports and music. This allows advertisers to use multiple criteria isolated or in certain combinations to define the target group for their ads.

In general, profiles may be created in two different ways. First, users can explicitly share information with the advertiser. Many online stores explicitly ask users about their areas of interest when setting up their personal profile or when signing up for shopper loyalty programs. Second, data can be derived from the user's shopping behavior. While in online shops this can be done implicitly, e.g., by analyzing a user's click-stream, the user's cooperation is required in stores, e.g., by presenting their shopping loyalty card upon checkout.

Based on the profiles it is possible to form target groups – either customized or based on cluster analysis. Advertisers might then select one or more target groups for their campaigns or define own target groups (an introduction to targeting can be found in Section 3.2.2; current practices as used by Amazon are described in [208]). The advantage of a cluster analysis is that the system assures a minimum of registered users in each group and shows which kinds of users really exist. It prevents advertisers from selecting extremely differentiated groups, which do not contain a large enough number of users. Yet, the groups might not fit all possible campaigns and custom target groups. To tackle these shortcomings our approach supports also customized target groups.

6.3.1 Creating Profiles

Our approach generates profiles, which consist of ratings on a scale from 1 to 10 in several categories. These depend on the advertising environment and may be adjusted to the needs of the provider. In our lab setup, we specified 45 categories, including for example music, events, and places. They are used to define different types of profiles: the consumer profile, the campaign profiles, and the target group profile for a specific campaign.

Figure 6.2 depicts a simplified example on how a consumer profile evolves based on the interaction with different campaigns. The campaign (bottom) targets people with a moderate interest in *Opera* (grey area, value 0 to 5). The bold black line represents the impact of the campaign (e.g., how strongly the categories change if the user is sensed at the opera) on their profile. As a result, the user's original value for the category *Opera* (top, left) increases from 1 to 6 (top, right) during several visits. In the following the profiling is explained in detail.

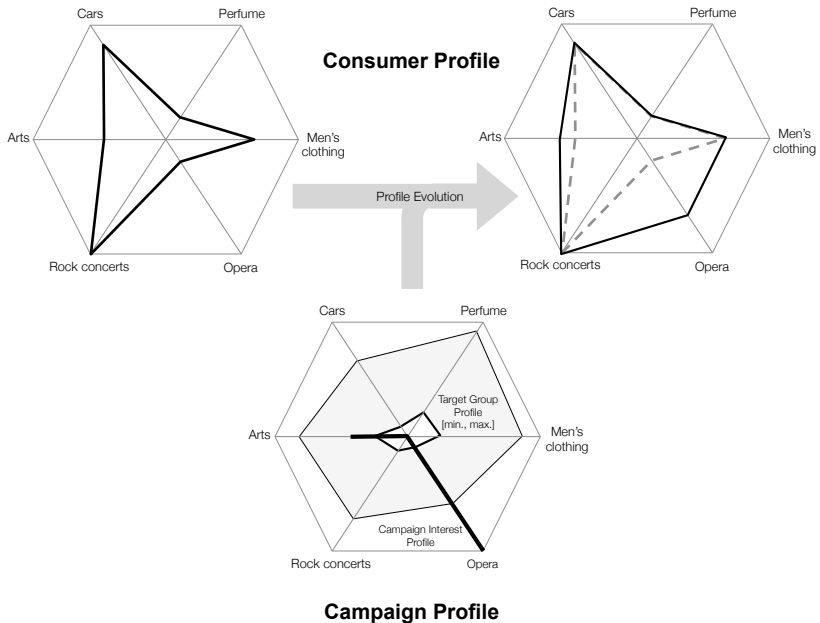


Figure 6.2: By visiting the opera, the consumer profile value increases, since the campaign profile specifies a strong impact on the category *Opera*.

Consumer Profiles

The consumer profile reflects the user's areas of interest and consists of two components. The first component is a self-assessment, allowing users to specify their interest in the given categories ranging from 'not interested' (0) to 'very interested' (10). By providing this explicit profile to the system, a user can benefit in a way such that the ads they receive are more appealing than random ones since they match their interests.

The second component tries to take the actual shopping behavior into account. As every user interaction with the system is captured, it is now possible to enrich the user's self-assessment using their behavior. The system currently tracks the following types of interaction: point of sale visits (A), event/ad watching on a display (B), rating of an event (C), and redemption of coupons (D). This set of interactions is not conclusive and further types of interaction may be considered.

Assuming that a user's visit of a sport store (A) signals their interest for the category 'sport', this store visit should positively affect the user's profile in the assigned categories. Similarly this is true for event/ad watching (B) and the selection of coupons (D) offered by the system as a motivation to participate. We have not yet integrated the event rating (C) into the system, because we doubt if, or how the rating can be translated into effects on the users' interests. However, the impact of a certain type of interaction depends on measures such as the resulting revenue. Since we did not yet deploy and evaluate the system in the real world and do not have any data, we cannot provide evidence that our model reflects the influence of the different types of interactions on the consumer profile correctly. For example, can we really derive a higher interest from more frequent store visits? Or can we conclude that somebody's interest is lower because they forgot to use their coupon on various occasions?

The consumer profile used for targeting content is calculated in multiple steps.

(1) Calculating Individual Activity

For each user, we gather data that shows which action took place in which category. Our system aggregates the different kinds of actions to get one value per user per category indicating how active a user behaves with regard to a certain category (dimension). The aggregation can be displayed as a relative measure, which is the sum of actions per action type:

$$sa_{cat,u} = \sum_{at=0}^{AT} z(a)_{cat,at,u}$$

$sa_{cat,u}$: the normalized activity

at : type of action

AT : number of action types

u : user

cat : category (area of interest, dimension)

$z(a)_{cat,at,u}$: normalized number of actions per action type, user and category

Depending on the weighted impact of the action, normalization is required since scaling is necessary. For example, the overall number of redeemed coupons will, on average, be much lower than the number of point-of-sale encounters. However, this does not indicate that coupons are less important. Hence, the weight parameter can be used to 'equalize' the different distributions in relation to their means and standard deviations (z-transformation).

$$z(a)_{cat,at,u} = z_{cat,at,u}(a_{cat,at,u}) = \frac{a_{cat,at,u} - \bar{a}_{cat,at}}{s_{a_{cat,at}}}$$

$a_{cat,at,u}$: the number of actions per action type and category

$\bar{a}_{cat,at}$: mean number of actions per action type and category

$s_{a_{cat,at}}$: standard deviation of the number of actions per action type and category

Once data is available, such as a positive consumer reaction to an advertisement or the revenue generated by the user, the importance of different action types can be estimated using multiple linear regression analysis.

$$Y_{cat,u} = \beta_{cat,0} + \sum_{at=0}^{AT} [\beta_{cat,at} * z(a)_{cat,at,u}] + e_{cat,u}$$

$$\bar{Y}_{cat,u} = \beta_{cat,0} + \sum_{at=0}^{AT} \beta_{cat,at} * z(a)_{cat,at,u} \sum_{u=0}^U [\bar{Y}_{cat,u} - Y_{cat,u}]^2 = \sum_{u=0}^U e_{cat,u}^2 \rightarrow \min!$$

$Y_{cat,u}$: target measure (e.g. revenue) per user and category

$\bar{Y}_{cat,u}$: estimated target measure per user and category

$e_{cat,u}$: residual value

$\beta_{cat,at}$: relative ‘importance’ of action type per category

Subsequently, the parameters ($\beta_{cat,at}$) would indicate the ‘importance’ in relation to the influence of the action type on business relevant target measures such as the revenue. In the current state, we still have to integrate a measure to track the real impact, e.g., the actual sales possibly triggered by an advertisement. The result could be the weighted normalized action type (individual activity) per category into account.

$$wsa_{cat,u} = \sum_{at=0}^{AT} \beta_{cat,at} * z(a)_{cat,at,u}$$

The measure *sa* (*wsa*) reflects how much the individual activity differs from the overall mean activity of all users per category. E.g., a value of 1 (–1) would indicate that the activity lies 1 standard deviation above (below) average within a certain time period. Users with a high individual activity might be of special value to advertisers since it can be assumed that their probability to generate revenue is highest among those users.

(2) Adjusting the Consumer Profile

Since for each user the self-assessment of interests is available, the consumer profile can now be enriched using the activity measure. An obvious approach would be to linearly increase or decrease the user's self-assessment per category based on whether his activity is above or below average. However, this might lead to ratings beyond the scale. Hence, the self-assessment per category should be normalized as well. Furthermore, the rating of users' interests might have different means in each category. Thus, it seems to be more reasonable to evaluate the relative values instead of the absolute values.

$$z(sas)_{cat,u} = z_{cat,u}(sas_{cat,u}) = \frac{sas_{cat,u} - \overline{sas}_{cat}}{s_{sas_{cat}}}$$

$sas_{cat,u}$: self-assessment per category and user
 \overline{sas}_{cat} : mean self-assessment per category
 $s_{sas_{cat}}$: standard deviation of self-assessment per category
 $z(sas)_{cat,u}$: normalized self-assessment per category and user

The formula uses the self-assessment profile as an input and adjusts the value based on the user's activity intensity. We introduce the alternation rate w ($0 \leq w \leq 1$), which influences how strong the effect of a differing normalized activity on the self-assessment is. If w is 0, the profile will not be altered. If w is 1, the profile will completely adopt the activity profile. Hence, we recommend the interval of $[0.2; 0.5]$ for w , if we assume that only one alternation rate is used for all changes.

$$z(sas)_{cat,u,t+1} = (1 - w) * z(sas)_{cat,u,t_0} + w * sa_{cat,u,t_0}$$

$z(sas)_{cat,u,t+1}$: updated normalized self-assessment per category and user in $t+1$
 $z(sas)_{cat,u,t}$: normalized self-assessment per category and user in t
 $sa_{cat,u,t}$: normalized activity in t

Campaign Profiles

In the previous section we introduced the alternation rate w , which is used for all changes made to the user's profiles. To take into account that many campaigns might be more or less about sports or more or less about music, we introduce campaign profiles. Similar to consumer profiles, campaign profiles consist of one value in each category which can be determined by the advertiser upon creating




Figure 6.3: Defining Target Groups: Advertisers can segment customers based on their interests – in this example sport and music).

the campaign. Its use is to characterize ads, coupons, and points-of-sale belonging to the campaign. This characterization is required to determine which category values of a user are to be adapted. We can use this information to specify the parameter w category- and campaign-specific:

$$w_{cat,cam} = \left[\frac{c\rho_{cat,cam} - c\rho_{cat,cam,min}}{c\rho_{cat,cam,max} - c\rho_{cat,cam,min}} \right] * w_{max}$$

$w_{cat,cam}$: category- and campaign-specific alternation rate

$c\rho_{cat,cam}$: campaign profile (category specific)

$c\rho_{cat,cam,min}$: minimum of campaign profile scale (e.g., 0.0)

$c\rho_{cat,cam,max}$: maximum of campaign profile scale (e.g., 10.0)

w_{max} : maximum pre-defined alternation rate (e.g., 0.3)

For example, the campaign of a sport retailer will have high values in the sport category. A consumer's visit to a point-of-sale belonging to the campaign will subsequently have an effect on the rating of his category sport. Since the campaign has a value zero in cultural events, a point of sale visit will not alter this category. By using campaign profiles, a more differentiated adjustment can be achieved.

Target Group Profiles

While the consumer profiles are calculated as a blend of self-assessment and actual user interaction with the system, advertisers can determine target group profiles using minimum and maximum values for each category. Figure 6.3 shows how advertisers can, based on the two-dimensional example from Figure 6.1, select suitable target groups for their campaigns. In addition, advertisers need to define a campaign profile for each of their campaigns, which only hold one value per category – like consumer profiles. This target group profile is then used to match the campaigns with the user’s interests. Since advertisers can select from a range of values, they may define precisely which users they want to target the ads to and which not. For example, if a sports retailer wants to target only consumers with a high interest in sport, they might define a campaign with a target group profile ranging from values 8 to 10.

$tg_{cat,cam,max}$: category-specific maximum target-value
 $tg_{cat,cam,min}$: category-specific minimum target-value

An interesting idea from the marketing point of view is to identify users that stated a high interest in a certain category during the self-assessment, but do not show any activity with the system in this category. By addressing those users with an advertisement in that category, chances might be higher to trigger activity since general interest can be assumed. This case shows that having both a self-assessed user profile and an adapted user profile is very valuable for advertisers.

6.3.2 Targeting Content Towards the Users

Based on the target group, ads to be displayed are chosen. We distinguish two cases when it comes to matching ads and users. If only a limited number of users are in the display’s vicinity, advertisements are selected by showing the best fit for only one single user taking turns in order of arrival. The main reason is that finding a best match for a small group of users might be difficult if interests among them differ fundamentally. Hence, a more promising strategy is to perform a ‘safe’ match on the single users. With an increasing number of users at the display, we use matching algorithms to get an overall sufficient fit.

Both cases require a measure indicating how well an ad fits (a group of) users. To measure the psychological distance between a campaign and a user’s self-assessment we use the general distance measure L_r , which we adapt to the proposed model. The distance measure L_r should meet two requirements:

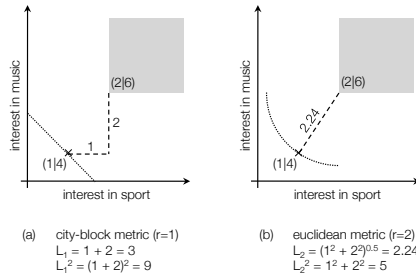


Figure 6.4: Calculating the Individual Fit of an Advertisement.

- If the user's category value lies within the defined boundaries of the target group profile, the distance should be 0. Thus, the distance is measured between the individual value and the boundary if the value lies beyond. Else, the distance is 0.
- It should be considered that a larger distance might indicate an aversion towards a certain category. Thus, many small deviations over all categories should not be assessed as 'worse' as a few large deviations of the same absolute number.

$$L_{r,cam,u} = \left(\sum_{cat=0}^{Cat} (D_{cat,cam,u}^r) \right)^{\frac{1}{r}}$$

$L_{r,cam,u}$ is the overall mental distance between a user's self-assessment and the nearest target group profile boundary. It sums up the category-specific distances. The parameter r is used to alter the distance type.

$D_{cat,cam,u}$ are the category-specific distances per user, category and campaign. It equals 0 if the self-assessment lies within the defined target group $[tgP_{cat,cam,min}; tgP_{cat,cam,max}]$.

Depending on r , different types of distances can be calculated: the city-block metric ($r=1$) or the euclidean metric ($r=2$) (Figure 6.4).

Based on the requirements, we suggest using the squared euclidean metric, which sums up the squared distances for each dimension. This metric takes into account that larger distances might be disproportionately worse than smaller distances. Hence users' objections towards specific categories are well considered and the probability of targeting ads towards 'wrong' users is minimized.

6.4 Implementation

Based on the concepts explained so far we built the CAdEt prototype which aims at being applicable to the scenarios introduced in Section 6.2. This prototype is functional, hence enabling us to realize the presented approach. At the same time, it is based on conceptual considerations that we assume to be applicable for the class of systems for pervasive advertising as proposed by the CAdEt approach. In this section we introduce the system architecture and implementation.

6.4.1 Requirements

The system design is influenced by several requirements arising from the scenarios:

1. **Connection mechanisms:** The scenarios mention Bluetooth and RFID, however, additional technologies might also be supported in the future. The system architecture must provide means to abstract from such technologies.
2. **User authentication:** Depending on the connection mechanisms, different credentials are available for user identification. In the case of Bluetooth, only MAC addresses are used. However, in the third scenario, further technologies are used for authentication (RFID-enabled cards). The CAdEt concept as explained in Section 6.3 might require in addition, more sophisticated user interfaces such as web sites that enable profile adaption by the user, which would require a username/password-based login. Thus, the system must administrate users independent from the specific identification mechanisms. This will also allow extending the system with other authentication procedures.
3. **Distributed system:** The system is inherently distributed since all scenarios embrace components at different locations, which are able to communicate and at the same time access shared data. The system is also loosely-coupled since customers can connect and disconnect ad-hoc, thus demanding a fail-safe communication between the components.
4. **Administration:** The focus of CAdEt's concept is on management of categories according to the concepts explained in Section 6.3. The system must thus provide means to manage, adjust, and query categories accordingly, including relations to ads and user profiles.




Figure 6.5: System Architecture: A central server stores advertisements / campaigns, profiles, and user interactions with the system; at the display, Bluetooth is used to identify the user and present targeted advertisements; an advertiser client allows to create advertising campaigns.

6.4.2 Architectural Concept

This leads to a system architecture as depicted in Figure 10.1 whose components reflect CAdEt's distributed structure and technology abstraction. The central server is used to persistently store ads, events, visitors' profiles, and the users' interactions with the system. Based on this information, the system can calculate a user's profile consisting of personally pre-defined preferences and interactions with the system at any time. This profile is matched against campaigns defined by the advertisers via a web client. The content can be queried by any client, e.g., a public display.

The campaigns generated by the advertisers are presented to the consumers by public displays. Each public display is equipped with a Bluetooth scanner. Therefore, the system can 'identify' users in the vicinity of the display and select/update content according to the profile related to the scanned Bluetooth MAC address.

The advertiser's client (Figure 6.6) allows advertisers to specify campaigns and select user groups to target. The users of this client have full access to their campaigns, but – due to privacy concerns – no access to customer data. Finally, an administration client provides means to export the gathered data in order to create reports. These can be used by advertisers to evaluate the success of campaigns as well as by the providers of the system for accounting purposes. The administrators have access to all data in the system and can provide advertisers with appropriate information about users. With these components, the requirements are fulfilled:

- Connection technologies are provided by on-site devices such as public displays. They communicate with the central server in a unified way and thus encapsulate the heterogeneity regarding connection mechanisms.
- Login and user authentication are provided by the central server, so that different on-site devices can provide user credentials that are matched to existing profiles by means of the database.
- By decoupling data management and on-site access, the system fulfills the requirements with respect to distribution and enables loose coupling, as on-site devices can connect and disconnect to the central server; at the same time, they manage the connection of customers.
- Categories are managed within server-side logic and database. Customers and administrators can access and adjust them with sophisticated UIs.

6.4.3 Implementation

The described prototype has been implemented in Java. The central server is a Java Enterprise application that delegates data access, session management, and provision of connections to the underlying platform. The advertiser's client is a rich client application. This is crucial since editing of campaign data does not only involve a tight interaction with the data store, but also sophisticated graphical editing. In contrast, the administration client is mostly used for creating reports and is thus provided as a web application.

We implemented a mobile phone client in JavaME, which can be run on Symbian phones. The mobile client can be downloaded via Bluetooth at the public displays. We use this application for two purposes: first, users can create their self-assessed personal profile. Second, the user can remotely control the display, e.g., browsing

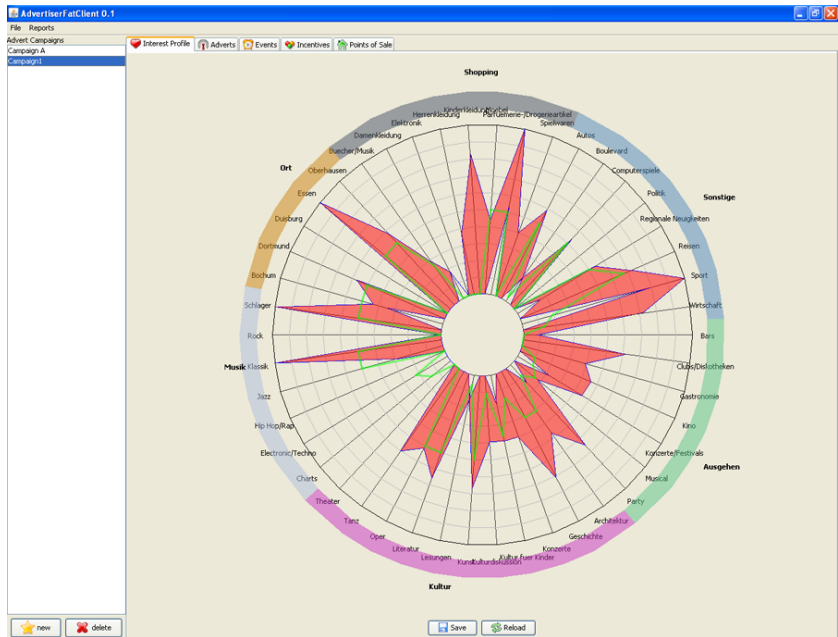


Figure 6.6: Advertiser's Client: Advertisers can specify a *campaign profile* (green line) defining the influence of the campaign on a user's profile, as well as the *target group profile* (red area) defining which user group the campaign is targeted to.

for information on a specific event by using the mobile phone. We also use Bluetooth scanners in points-of-sale to detect the (areas of a) store a user visited. This reveals important information on the users' interests.

The user can enter personal preferences into the system by means of the mobile phone application. This can be done anonymously as users are not required to provide any personal data. Furthermore, users can interact with the advertising board using the mobile phone.

Connections between on-site devices and the central server are made with Java Remote Method Invocation (RMI). This is very efficient and feasible since these communication channels are used only inside the system. Connection to Bluetooth and RFID are provided by on-site devices, which either use built-in connection technology like RFID readers or Java software to create connections, e.g., to Bluetooth devices.

6.5 Privacy

When it comes to storing and assessing information about a user's shopping behavior, this often raises privacy concerns, which might have a strong impact on acceptance and uptake. Consumers are afraid of data being passed on to other companies or advertisers, hence resulting in an ever-growing number of ads they are confronted with. However, the success of shopper loyalty programs shows that most of the users are willing to give away their information in return for incentives [49; 347], even though it is not obvious who has access to the data.

With regards to user acceptance, Kollmann's acceptance model distinguishes different phases of the acceptance process and therefore different types of acceptance [181]. Initially, in the attitude-forming phase, it is essential to point out a system's characteristics in terms of privacy and other factors in order to influence the attitude toward the innovation positively. Once the potential user is aware of the innovation, they can form a positive attitude. As a potential user engages with the system it is important that they experience it positively. They will evaluate these experiences and rethink their attitude and eventually form a decision on whether or not to become a permanent user. Yet, this is not a one-time decision, as once a user started to use the system, they will still carefully reevaluate the decision in terms of privacy, costs, benefit, etc. For a comprehensive overview on privacy recommendations for DOOH advertising we refer to Geiger [120].

Our approach contributes to user acceptance as it makes it possible to target consumers without the need to store personal data such as name, age, gender, or address by linking a profile to an ID such as a Bluetooth MAC address only. Thus, the user stays anonymous and at the same time may deliberately decide whether or not to participate. Yet, it is essential for the success of the system to convince the user of its advantages, such as ads, which are personally interesting for them and make them aware of the fact that they are virtually in control of the system at any time. Furthermore, we believe that the uptake and success of such a system can be significantly increased if the platform providers have high credibility.

6.6 Discussion

This chapter introduced CAdeT, a pervasive advertising environment, which allows the efficiency of targeted public display advertisements to be studied. First, we looked at how both the users' personal interests as well as their shopping

behavior can be taken into account in order to precisely target advertisements to them. Therefore, the concept of adaptive user profiles was introduced. We showed how such a profile could be generated based on a user's self-assessment, the user's individual activity, and the activity intensity compared to other users of the system. Second, a system setup was proposed which allows anonymous and unobtrusive profiling based on Bluetooth scanners followed by an explanation of the implementation and functionality of our system.

With our implementation we provide means for realizing the presented scenarios. First, large scale networks, which gather and make information on the user's interest available in different locations (e.g., at the subway station) can be realized. Hence, content on public displays can be targeted towards the people in the vicinity. Second, the modular architecture of the system makes it possible to later integrate further technologies, such as RFID cards.

Our advertising environment has so far been deployed in a lab setting only and not been evaluated with real-world data. A large-scale user study should look at the following aspects: (1) accuracy and robustness of the profiling algorithm, (2) performance and accuracy of the matching algorithm, (3) acceptance and concerns among user, (4) a usability evaluation of the different types of interactions supported by the system, and (5) multi-user situations.

IV

THE AUDIENCE VIEW

Outline

As public displays and their (advertising) content become interactive, engaging experiences can be created for the user. However, in order for the user to interact, three steps are required. First, the attention of the user needs to be attracted. Second, the interactivity of the public display needs to be communicated. Third, the user needs to be enticed to interact by motivation and providing easy-to-understand interaction techniques.

This part of the thesis takes a closer look at these three steps.

- **Chapter 7 – Attracting Attention** We present a prototype that allows content to be adapted based on the user's gaze behavior. We show that in this way the attention of the viewer towards the adapted content can be increased. Furthermore we discovered an impact on attitude, according to the peripheral route of the Elaboration Likelihood Model.
- **Chapter 8 – Communicating Interactivity** In order to understand how interactivity can be communicated in an efficient way, we present *Looking Glass*, a public display application that utilizes the user representation to make them aware of a display's interactive capabilities. The approach is first studied in the lab before deploying and evaluating it in a public setting in Berlin, Germany.
- **Chapter 9 – Enticing Interaction** To investigate what motivates users to interact and how interaction can be designed in an easy-to-understand way, we present *Digifieds*, a digital bulletin board. Findings from a lab study and a deployment in a public setting in Oulu, Finland reveal suitable interaction techniques and motivating factors.

Chapter 7

Attracting Attention

Attracting the attention of the passerby is the first and most important challenge for digital public displays. If users do not see a display, no further interaction will occur. Before the advent of real-time audience measurement technology it was not possible to detect how many people actually paid attention to a display. As a result, digital out-of-home (DOOH) advertisers created models that allowed the number of passersby and hence potential onlookers to be estimated. One example is the Fraunhofer *frequency atlas* for Germany³⁰, which provides this estimate for streets, public transport infrastructure, and pedestrian areas. The frequency atlas considers different factors, including points of interest (restaurants, recreational facilities), population density and structure, traffic, street categories, and socio-demographic as well as socio-economic data. Furthermore, it considers the contact opportunity for advertising media – the so-called *G value*. This metric is calculated, among other factors, based on the distance to trajectories, occlusion, lighting conditions, distraction through other advertisements, height and orientation of the display, and the complexity of the surrounding environment. For decades, models like the frequency atlas influenced payment models for outdoor advertising media, such as public displays. However, outdoor advertisers are progressively adapting real time technologies that provide feedback on contact (i.e., the number of passersby looking at a display).

³⁰ Frequenzatlas: <http://www.iais.fraunhofer.de/frequenzatlas.html>, last accessed February 13, 2013

As public displays are deploying sensing technology, they allow users to implicitly interact with them through their mere presence. It is finally possible to actively influence both the passerby's degree of attention and potentially their attitude, e.g., through adapting the content. The basic assumption is that users might be more attracted by content that reflects their interest. Traditionally, content on public displays changed frequently and this rate increased even with the digitization of displays. As a result, passersby do not know what type of content to expect but learned over the years that this content is in general not interesting for them and eventually ignore it. This prompted advertisers to design content with visual attributes that potentially attract the passerby's attention. However, more profound cognitive changes (attitude, engagement) might occur as displays become interactive and provide information or services that are of value to users.

In this chapter, we look at the effect of content that adapts to the user. We are particularly interested in the potential of this approach to increase the overall attention as well as to influence the user's attitude. Besides conspicuity of content and colorful content, Müller et al. identified interesting content as one of the most important dimensions that influence whether or not people pay attention to public displays [255]. Our idea is to assess the user's gaze behavior, e.g., how long they look at specific content and then use this information to adapt the content.

It is currently difficult to automatically assess audience behavior – and more specifically attention – in front of public displays. Several camera-based approaches have been presented that try to recognize attention, including eye contact sensors [245], the analysis of head poses [342], and (mobile) eye trackers [324]. Whereas the first of the two approaches lacks accuracy (attention can only be measured on display but not on content level), wearing mobile eye trackers is not feasible in commercial deployments. As a solution we suggest to integrate the eye tracker with the display. One major challenge is how to 'automate' the calibration of these eye trackers in public settings, especially as interaction times are short and it is unlikely that users are willing to spend time on calibrating the system prior to using it. With the increasing accuracy of eye trackers, it will be possible in the future to calibrate the eye tracker on-the-fly, e.g., by displaying a bouncing ball on the display and assuming that the user's eyes will follow it.

To investigate the potential of our approach – using gaze data and adapting content in order to increase attention – we built a web-based prototype. As automatic calibration is still in its infancy, we used a stationary eye tracker and conducted a lab evaluation. As technologies such as Web 2.0 evolved, the Internet became a platform providing a wide variety of applications, services, and content such as news, ads, weather information, and stock trends. Such content elements are

often displayed alongside the main content of web pages and are an important part of many sites' business model. We envision a similar layout of public display content in the future, where information and ads are mixed [12]. Currently, info screens blend different types of content in a time- or space-multiplexed manner (cf., Section 10.2.1). If an eye tracker was available at the display, this advertising content could be adapted and replaced with suitable content in real-time based on the user's gaze as they (implicitly) interact with the content.

Our prototype system allows a web page's content elements to be adapted on-the-fly based on the user's attention (dwell time, number of fixations, and duration of fixations). It assesses the user's attention by analyzing their gaze using a Tobii X120 eye tracker and then feeds back the information to the browser by using an HTTP proxy (Figure 7.5). We additionally embedded JavaScript in each website, which processes the gaze data obtained from the proxy and subsequently adapts the content on the pages. In a lab study, the users surfed on the Amazon website. Using our system, we adapted image elements on the page and measured whether attention would increase for adapted content compared to randomly chosen content. We were able to show an increase of attention towards these elements. There is also evidence for the peripheral route of the Elaboration Likelihood Model [283] implicating an influence on the user's attitude.

This chapter is based on the following publication:

- F. Alt, A. S. Shirazi, A. Schmidt, and J. Mennenöh. Increasing the User's Attention on the Web: Using Implicit Interaction based on Gaze Behavior to Tailor Content. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction (Copenhagen, Denmark), NordiCHI'12*, New York, NY, USA, 2012. ACM

7.1 Background and Related Work

Today, a variety of stationary high-precision eye trackers are commercially available. These are still expensive, but as a mass-market emerges we envision prices to drop significantly and the integration with public displays to become feasible. Current systems support a large freedom of head movement and allow users to behave naturally. With data rates of up to 120 Hz, trackers provide real time information about the gaze point. Currently, eye tracking systems are mainly found

in labs due to their high price tag. However, a lot of research has been focused on self-made *low-cost eye trackers*. Hanson et al. report on the potential of eye trackers built upon components off-the-shelf (COTS) [131]. Li et al. developed the openEyes system [205], an open hardware design with several open-source software tools for eye trackers. They show how the performance of those low-cost trackers could be incrementally improved by replacing parts of the systems with higher quality components³¹. Current developments suggest, that eye tracking systems might be included in standard computers at little or no extra cost in the future as more applications become available. Modern laptops are already widely equipped with devices such as webcams or fingerprint readers, which were also specialized and expensive decades ago.

In general, two categories of eye tracking systems can be distinguished [95] – diagnostic and interactive systems. *Diagnostic systems* focus on the offline analysis of data gathered from user interaction to evaluate usability or user behavior. In marketing, gaze data are particularly interesting, as they provide insights into visual, cognitive, and attentive aspects of human performance, as well as how consumers disperse visual attention over difficult forms of advertising [95]. Further research on diagnostic systems includes tools focusing on web reading. Beymer and Russell present the WebGazeAnalyzer [36], a tool that supports the collection, analysis, and re-analysis of gaze data. The WebEyeMapper [298], analyzes eye tracking data and maps it onto objects on a web page, making it easy for researchers to comprehend the gaze data collected. Diagnostic systems are also used to explore user behavior on the WWW. A prominent example from the field of advertising is the work of Burke on the usefulness of banner ads and banner blindness [58] .

In contrast, *interactive systems* focus on how eye trackers can be used to control applications in scenarios where people are not able to use other input devices, e.g., a surgeon during operations. Chen et al. investigate the correlation between eye and mouse movement [64]. Sibert and Jacob compare a novel gaze-based object selection technique with conventional selection by mouse [340]. Farid et al. investigate how eye gaze can be used to control computer displays (e.g., navigating within large images or multiple video streams) [102]. More specific to the WWW, tracking gaze interaction on web pages requires a real time analysis of the intersection of eye gaze and the DOM bounding boxes of the viewed web page, as done by Reeder et al. [298]. Whereas the previous examples were mainly used to explicitly control a system, researchers also looked into how eye trackers

³¹ See also the homepage for the COGAIN Network of Excellence and the COGAIN Association www.cogain.org, last accessed February 13, 2013

could be used in an implicit way. The MAGIC pointing technique [387] is one of the early projects using gaze information for positioning the mouse cursor and assisting interaction. Another example is Santella et al.'s application for photo cropping based on a user's gaze, which implicitly identifies regions-of-interest based on fixation data [313]. Buscher use eye tacking as a data source for realizing attention-based feedback on subdocument level [60]. As a use case they examine personalized, context-based query expansion / ranking. Biedert et al. discuss using gaze data for creating responsive text [38].

Furthermore, eye trackers have been used to build *attentive user interfaces* (AUIs) [369]. These UIs try to manage the attention of the users through input channels beyond conventional, explicit channels and bring the right information at the right time to the user. Computer vision and other technologies can potentially be used as input channels for AUIs, e.g., the user's presence, body posture, head direction, etc. However, the user's gaze information is a particularly rich resource. EASE, described by Wang et al. [376], uses the gaze data to assist Chinese text entry. Qvarfordt and Zhai [292] developed an interactive tourist system, which senses the users' interests based on eye-gaze patterns and manages data output accordingly. Drewes and Schmidt [94] enhanced the MAGIC system with a touch-sensitive mouse to ease the pointing task in graphical UIs. EyeWindows, presented in [109], is an attentive windowing technique that uses eye tracking for focus window selection. Selker discuss the complexity of interfaces designs based on simple observations of eye behavior [326]. For a framework on increasing user attention through AUIs we refer to [370]. In the context of the WWW, previous work on AUIs has mainly focused on revealing the user's interest, e.g., [4; 22; 73]. Until now analyzing reading behavior provides relatively reliable results and is studied mainly in the context of implicit search queries [81; 96; 111; 132]. In contrast, inferring the relevance of images based on gaze has turned out to be challenging and has only been successful under controlled lab conditions [180].

Despite considerable effort to create attentive UIs based on the user's gaze data, approaches so far mainly focus on enhancing search queries. Little is known about how gaze-behavior towards forms of content other than text can be exploited. To the best of our knowledge this project is the first approach to use real-time gaze data towards (ad) images with the aim to increase attention.

7.2 Gaze-Based Adaptation of Content

When observing users in front of screens, it becomes clear that they are interacting with more than just the provided input devices (e.g., keyboard, mouse, touch display, camera). They are often standing or sitting seemingly motionless while their eyes actively scan a section of the screen. Eye trackers could retrieve this information and allow it to be used in a variety of application domains, both offline and in real-time. The most important ones include UIs for people with disabilities, market research, advertising, and usability. For marketing research and usability engineering, eye trackers are typically used offline, meaning that experiments are carried out prior to analyzing gaze paths and behavior. In UIs for disabled persons, eye trackers can enable gaze-based, real-time interaction. Eye trackers have already been integrated into eye-typing interfaces, wheelchair steering systems, or remote controls.

Gaze data can be used for creating adaptive interfaces. Oppermann defines the term *adaptation* as a process in which a system adapts its behavior to users based on previously acquired and processed information about the user [273]. We see a large potential in adapting content to the attention or even the interest of the user.

The adaptation can be realized in two ways, by tailoring content based on (1) previously gathered knowledge, such as a profile (offline) and (2) actual user behavior (real-time). The first approach (*offline*) relies upon the identification of the user and on having access to the previously assessed data. Such data is provided either explicitly by the user (e.g., a questionnaire on his interests) or implicitly by collecting data on the user's behavior (e.g., encounters at customer touch points, such as purchases). The second approach (*real-time*) assumes that the user's interest can be predicted from the currently accessed content (such as Google adSense) in addition to knowledge about the current behavior. If the user is, e.g., in an online shop looking for a specific item or if they write an email requesting information on a product, providing ads for similar products is then both feasible and sensible. However, when the user looks at more generic content, e.g., a news site, the link to a specific user's interest is less clear. For example, if the user reads about a plane crash on a news site, is it appropriate to advertise for cheap flights from the similar airline?

The advantage of the second approach is that the user's interest is determined in real-time. This approach eliminates the need for storing personal profiles. In our study, this approach has the potential to increase the attention towards content and offers new opportunities for structuring or selecting information. The following scenario points out some of these opportunities.

Scenario – Adaptive Advertisements It is early evening as Paul arrives at a hotel in Stuttgart for a business meeting starting the following morning. After checking in, he decides to go out for dinner. On the square in front of the hotel he passes a public display, which shows an interactive map of the city. As he approaches the display his eyes catch the ad for a Thai restaurant that is displayed alongside the map, next to ads for a sports store, a supermarket, and a large department store. As he is not really up for Thai food at the moment, he looks at the map and tries to identify a close-by area where he could find further dining options. At the same time, the eye tracker at the display recognized that the Thai restaurant received most of Paul's attention among the ads and brings up three more restaurants next to the map, including an Italian Cucina, a German Bar, and a French Bistro. Paul is a big pizza fan and immediately touches the advertisement of the Italian restaurant to get further information on directions and the menu. He also receives a coupon that saves him 3 euros on the pizza. He takes a picture of the coupon with his mobile phone and heads for the restaurant.

The scenario can be captured by a real-time analysis of the user's gaze input. With gaze information it becomes possible to obtain the required information on-the-fly and adapt content instantly and more appropriately. Whereas based on a click or touch attention (or even interest) can be assumed for one page only (the page clicked), gaze input can consider any page element. The drawback of the gaze method is that the intention or motivation for the user is less clear. Clicking or touching a content element is very likely to occur out of interest, but eye contact might be the result of subconsciously scanning the page and may not be related to the user's interest. Therefore, perception or even interest for gaze is much more difficult to determine, as (implicit) interaction only occurs for very short bursts. Hence, the paper focuses on adaptation based on the user's attention rather than on perception or interest. Mello-Thoms et al. [227] and Hauland [136] shows that the dwell time could be used to compare attention distributed between targets. Furthermore, the Elaboration Likelihood Model (ELM) is used to explain cognitive changes.

In the remainder of this section we first look at suitable metrics that can be used to assess user attention. Then we formally present our approach and present the research hypotheses. Finally, we introduce the ELM and discuss how it can be used to evaluate and explain cognitive effects.




Figure 7.1: Creating Adaptive Content Based on Gaze.

7.2.1 Metrics

To analyze the user's gaze behavior, Poole and Ball [287] identified different eye-movement metrics. In the following, we discuss those of interest to our research.

Dwell time (also referred to as gaze or fixation cluster) The dwell time can be used as a measure for attention among a set of targets [136; 227]. Yet, no conclusion on the interest towards the target can be made.

Number of fixations per target The number of fixations for a target can be used as a measure for its importance. Hence, a higher number of fixations indicates that a target (e.g., a page element) is more noticeable or important to the viewer than others [288].

Fixation duration The average duration of fixations per item can be used as an indicator for engagement or for the difficulty a user faces when extracting information [171].

7.2.2 Approach

As users look at different elements, their eyes are constantly moving, even if they are not consciously aware of it. With eye tracking, the user's focus and the content at this position can be determined. This information can be used to derive the attention for a given element. Our general approach for gaze-based adaptation is as follows (see also Figure 7.1). First, a set of content areas potentially attractive to the user needs to be identified. These content areas can be texts, images, videos, or animations and may be located in any part of the screen. Second, an eye tracking system is used to detect the user's gaze behavior towards these content

areas. This information should be monitored in real-time and transferred to a server-side application. Third, the acquired gaze information is used to calculate the user's attention towards each content area. Fourth, an attention measure is used to determine elements to be adapted. Steps 2–4 are constantly repeated.

7.2.3 Hypotheses

Based on our approach, we aimed to answer two hypotheses:

Hypothesis 1 – Adapting content towards the users' gaze behavior increases their attention.

Today, click streams serve as both a measure for the success of web pages and as input data for adaptive content. With eye tracking we envision to find a means for an on-the-fly assessment of what attracts the user's attention and use it in real time to build attentive systems.

Hypothesis 2 – The gaze behavior correlates with the user's interest.

We assume that users tend to look at content elements, which are potentially interesting to them. Finding proof for this assumption would allow the user's interest based on gaze information to be determined.

7.2.4 Cognitive Effect

In order to understand the effect of our approach on the user, we draw upon the widely accepted Elaboration-Likelihood-Model (ELM) [283], which tries to explain the changes in a user's attitude (cf., Section 3.2.2). As previously mentioned, the model differentiates between two routes leading to persuasion or elaboration: the central and the peripheral route. Which route is taken depends on the user and the situation he or she is in. The ELM distinguishes between two factors: the motivation to process and the ability to process. If both factors are true the central route is taken, where conscious information attainment and processing takes place. As a result, the average fixation duration as a measure of engagement would increase significantly. If the peripheral route is taken, unconscious information acquisition takes place. Advertisers should repeat peripheral cues to achieve persuasion [284]. Thus, a change in attitude can be achieved through a significant increase in the number of fixations.

Through an evaluation, we aim to find out whether gaze-based adaptation of content could support (a) the central route and thus lead to higher *engagement* (this would be the case for an increase in the average fixation duration) or (b) the peripheral route and thus lead to a *change in attitude* (increase in the number of fixations). Finding either effect would make the approach highly interesting for advertisers as a strong influence on the user's elaboration could be assumed.

7.3 Prototype

The basic architecture of our system is depicted in Figure 7.2. It is based on a standard web architecture consisting of a web client (browser) and a web server (content provider). An eye tracker implicitly tracks the gaze-based user interaction on a web page and stores the gaze path in a database. In order to realize real-time interaction on a web page, we use an HTTP proxy to (1) insert application code into the website which handles the adaptation, and (2) to read and process the recorded gaze data to update the appropriate content and to trigger the update in the client. Using a proxy has two advantages. First, it allows the system to work with arbitrary web pages, since the proxy can easily insert the required code on any requested web page. Second, no user-sided software installation is required.

7.3.1 Eye Tracker

We use a Tobii X120 eye tracker to extract coordinates of the eye's focal point. The Tobii X120 is a bifocal eye tracker that uses two integrated infrared cameras to monitor gaze behavior. It is table-mounted and supports data rates of up to 120 Hz. We implemented the tracking software to time-stamp and record the entire gaze path in a database at 60 Hz. Hence, the gaze behavior can be monitored with regard to number of fixations, duration of fixations, and dwell time.

7.3.2 UsaProxy

In interactive systems it is necessary to process the recorded gaze data in real-time. We use the HTTP proxy UsaProxy [23]. It allows JavaScript code required for processing the collected data to be inserted on-the-fly on arbitrary web pages. The analysis of the data can be done either on the server or the client side.

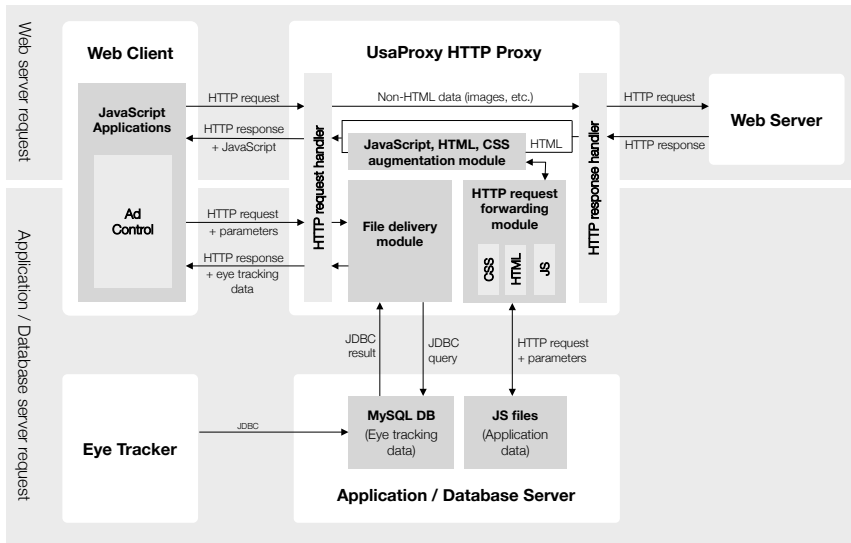


Figure 7.2: Components of the Experimental Setup: UsaProxy HTTP proxy, eye tracker, application/database server, web server.

The UsaProxy's initial task is to embed JavaScript code in any page sent from the web server to the client in response to a standard HTTP request. This makes it possible to realize the embedding of content on the client instead of the proxy. In order to add the script, UsaProxy monitors all HTTP requests that pass through it. The modifications to the original HTML content are small: a `<script>` tag is added inside the document's `<head>`, and its `src` attribute references the required JavaScript. As JavaScript code can be inserted into pages, data from external sources can also be loaded explicitly on to the page. Therefore, the UsaProxy was slightly modified to support XMLHttpRequests to external data sources. The script, which runs inside the browser as a result of the above modifications, uses further XMLHttpRequests to access the content on the external database server. XMLHttpRequest objects provide a convenient way of downloading the data. One problem is that modern browsers require requests to be sent to the same server that supplied the original web page. We were able to circumvent this same origin policy in the following way: the JavaScript component makes a special request to the same server that the web page was requested from. The requested URL is special since it appears to access the 'local' directory `/usaproxy1010/httprequest/` on the server. However, this directory name is intercepted by UsaProxy and redirected to the database server, which then

answers the query and returns the required data. In this way, we are then able to implement a method to handle AJAX requests. The requests contain information about the object the user is looking at. UsaProxy retrieves and sends the current coordinates of the eye's focal point to the browser.

7.3.3 Apparatus

The Amazon corporate website served as the test environment, hence, allowing realistic tasks to be performed in a familiar setting, such as searching for a product (e.g., looking up the current price for the iPod touch). We use the UsaProxy to embed additional ad elements into the web page, which then updates based on the user's gaze behavior. Our changes do not alter the look and feel or the URL of the Amazon website. In general, any website with arbitrary content elements is suitable. The sole information the system needs is the position and size of the elements to correctly associate them with the user's gaze. This setup allows us to determine the dwell time, the number, and the duration of fixations on each page element. As suggested by Poole and Ball [287] the dwell time is used as the most appropriate means to compare attention between targets in order to decide which elements to adapt.

Layout of the Test Web Site

Figure 7.3 shows the layout of the test website with three distinct areas.

1. We insert a *task area* (a), allowing arbitrary tasks to be presented to the subjects. This area remains unchanged during navigation on the page. Once a task is solved, the answer can be entered into the text field provided. Clicking on the 'Next Question' button triggers the system to randomly draw the next task from a database.
2. The *advertisement area* (b) is inserted using JavaScript. In this area, we show different advertising elements. To maximize exposure of these elements they are inserted on the left side of the page since this area is most likely to be perceived by the users [289]. Note that this is consistent through all conditions in the study. Since this area is integrated into the corporate Amazon design, it is not obvious for participants that it contains third party content.
3. The *main content area* (c) shows the original Amazon website. It is fully functional and users can freely navigate around the page.

How much is the MP3-Player "Ipod Touch 8GB" on Amazon?

amazon.de

Hallo! Melden Sie sich an, um persönliche Empfehlungen zu erhalten. Neukunde? Bitte hier starten

Mein Amazon | Angebote | Wunschzettel | Gutscheine | Geschenke

Student? 1 Jahr Versandkosten sparen | Mein Konto | Hilfe | Impressum

Alle Kategorien ansehen

Suche: Alle Kategorien

Handys & Zubehör

Fernseher

Digitalkameras

Heimkino

Alle Elektronik-Produkte

Begrüßen Sie die neue Kindle Fire-Familie

Kindle Fire 159€

Kindle Fire HD 199€

Der neue Kindle eReader Nur 79€

amazon Prime

Jede Lieferung am nächsten Tag

Die Amazon.de VISA Karte Jetzt 30 € sparen und dauerhaft Bonuspunkte sammeln

Für Sie dokumentiert (Was ist das?)

Sie haben keine kürzlich angesehenen Artikel. Sobald Sie sich Produkttiteln oder Suchergebnisse angesehen haben, finden Sie diese Seiten zu Ihrer Information hier aufgeführt.

Einkauf fortsetzen: Topseller

Wohnen geht die Liebe, wenn sie...
Eckart von Hirschhausen
★★★★★ (15)
Taschenbuch
EUR 14,99

Greggs Tagebuch 7 -
Dürren geläutert!
Jeff Kinney
Gebundene Ausgabe
EUR 12,99

Verschlung: Der vierte Fall
fig.:
Jussi Adler-Olsen
★★★★★ (70)
Gebundene Ausgabe
EUR 19,99

Taberna Libraria: Die
Magische.
Dana S. Elliott
★★★★★ (41)
Kindle Edition
EUR 3,99

amazon.de

China | Frankreich | Großbritannien | Italien | Japan | Kanada | Österreich | Spanien | USA

Entdecken Sie: AbeBooks | Audible | Book Depository | Amazon BuyVIP | IMDI | Javari | LOVEFILM | MYHABIT | Shoptop

Figure 7.3: Layout of the Test Website: (a) the task area, (b) the main content area, (c) the advertisement area.

Learning and Measurement Phase

To realize the adaptation of the images we define a learning phase and a measurement phase (Figure 7.4). These are continuously swapped while the user is surfing the web. In the *learning phase*, the system collects the user's gaze data hence making it possible to identify the area that received the most attention by calculating the dwell time. For the subsequent *measurement phase* the system updates the ad elements on the next page. Note, that in order to not bias the data gathered, the images are not changed while browsing a page, since this would most likely generate additional attention. Hence, image updates are seamlessly integrated upon reloading a page as a result of clicking a link.

In order to compare the effect of adapting content, the system provides two modes: in the *random mode* (Figure 7.4, top) the advertising images are updated randomly, in the *gaze mode* (Figure 7.4, bottom), the images are updated based on the collected gaze data.

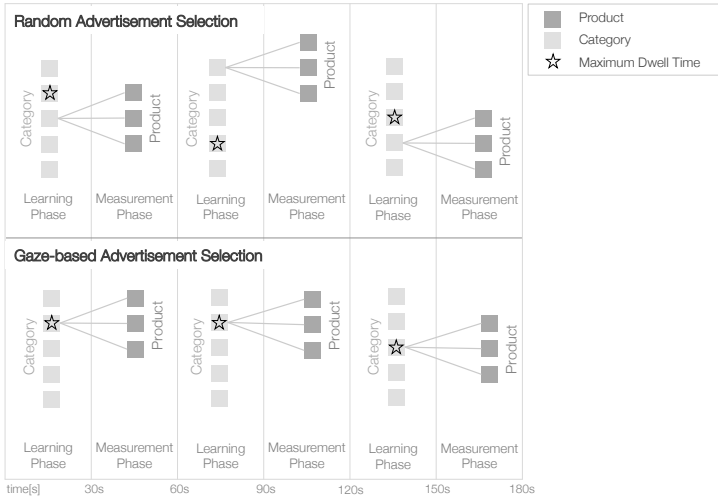


Figure 7.4: Adaptation of Content: (1) in random mode, product images are selected independent of the gaze behavior. (2) in gaze-based mode, product images are adapted based on the user’s attention.

Adaptation of Content

As adaptive content for the study we prepared a set of five images, each representing a certain product category (e.g., cars, mobile phones, games, etc.) and for each of the five categories a set of three images of related products (e.g., BMW, Mercedes, and Seat for the category ‘car’). This resulted in a total number of five categories and 15 product images. As we planned for a within subject design we created a second image set to avoid learning effects.

As mentioned, we use rotating pairs of learning and measurement phases. In order to measure the effect of adapting content to attention, we show the subjects five different category images in the learning phase and measured the dwell time for each image. After 30 seconds, the system triggers a change event so that the five category images were replaced by three product images of the same category upon reloading the page. Based on the software mode the three product images are either chosen randomly or based on the gaze data (by using the category which received the highest sum in dwell time). This procedure is repeated continuously.




Figure 7.5: Adaptive WebAds: As the user interacts with the web, their gaze behavior is implicitly monitored, fed back to the system, and the content being adapted automatically.

7.4 Study I: User Attention

7.4.1 Methodology

We discussed different alternatives with regard to evaluation of the approach. A long-term field study would have allowed a large set of data to be collected and analyzed and the long-term effect on the user and his behavior to be assessed. However, the following challenges prevented us from running the study in the field. First, the eye tracker needs to be re-calibrated as the user changes their position (e.g., every time they sit down on the chair) which would have put an unbearable burden on the user and potentially led to biased data. Second, to collect enough data, we would first need to identify the favorite websites of the user. Then we would need to integrate the adaptable content with these websites to maintain their look and feel to avoid revealing the investigated objects.

As a consequence, the potential of the approach in a controlled lab environment is evaluated. This allowed precise data on the gaze behavior of the user to be gathered and the effects of the approach (changes in dwell time, number of fixations, fixation duration) to be measured. As a baseline condition, randomly selected content is used as this most closely reflects the way image-based ads are currently presented in the World Wide Web. We used a within-subject design where all users would be shown both the random and the gaze-based content (independent variables). The dependent variables were dwell time, the number of fixations, and the duration of fixations.

Furthermore, we had the users fill out a questionnaire after the experiment to assess recognition for the products. We also conducted semi-structured interviews with a focus on user acceptance.

7.4.2 Participants

The participants were recruited via bulletin boards in the neighborhoods surrounding the university, from mailing lists, and lectures. In total, 15 participants (avg. age: 26.5 years) were selected. None of them ever participated in an eye-tracking study before and all have previously used Amazon.

7.4.3 Setup and Procedure

The first study was conducted in an office in the lab. Users did not engage in conversations during the experiment with others around them or turn their heads away from the eye tracker. A table-mounted eye tracker in front of a 22" TFT monitor (Figure 7.5) allowed the users to behave more freely, as they were not restricted to a fixed (head) position.

All of the participants were briefed in the same room where the eye tracker was set up and calibrated. We explained to them about the area for task descriptions on top of the Amazon website. Then, we asked them to solve the provided tasks (e.g., looking up the price for an iPad) and enter their answers into the text field for the next 20 minutes. The participants were informed about the collection of gaze data but we neither revealed the objective of the study nor the additional ads. The users were instructed not to leave the Amazon website. After 10 minutes, the system mode was switched remotely between random and gaze-based mode. To avoid any learning effects, half of the participants started with the gaze-based mode, the other half with the random mode. After the experiment the participants filled out a questionnaire and engaged in semi-structured interviews.

7.4.4 Data Analysis

During the study we recorded the participants' (time-stamped) gaze path, resulting in a total of 722,689 data points. During the analysis we found that for 3 participants the tracker had not recorded gaze-data properly, probably due to their

bright eyes or glasses. Hence, we had to exclude their datasets. We used the software package 'IBM SPSS Modeler' to calculate the dwell time, number of fixations, and average fixation duration per ad element.

7.4.5 Results

To identify the impact of different images on the users' attention and attitude, we used Student's t-Test because of its good performance with sample sizes below $n=30$ [118]. There are 24 'observations' from 12 participants. Thus, we used a paired sample t-test. Since we have the directed assumption of increasing attention, the testing is one-tailed [117].

At first we compared the dwell time between random and gaze-based images. We found an increase of 191.6 ms per participant for the gaze-based adaptation, yet this was not significant (Table 7.1c). When digging further, however, we discovered that there was a significant increase in the number of fixations for the gaze-based images (Table 7.1a). This is a strong indicator that the participants considered these images to be more noticeable and hence paid more attention.

Remarkably, while the number of fixations increased significantly (Table 7.1a), the average fixation duration decreased non-significantly (Table 7.1b). Thus, it is likely to assume that the fixation duration is about the same in the parent population. According to the assumptions of the ELM this indicates an unconscious information attainment in the peripheral route (see Sections 7.2.1 and 7.2.4). An increase in the average fixation duration would indicate more elaborated and conscious information processing in the central route.

Note, that since the average dwell time increases systematically with the number of fixations and the average fixation duration, it seems that the (non significant) decreasing fixation duration lowers the increase of the dwell time (Table 7.1c) in this specific sample. According to the t-test, we would expect the fixation duration to be equal in the parent population.

For recognition we tested whether the participants could recognize randomly chosen or gaze-based product images better on a 5-point Likert scale (1=I definitely did not see the ad, 5=I definitely saw the ad). We found an increase for images shown in the gaze-based condition ($M=3.20$, $SD=0.92$) compared to the test condition ($M=2.72$, $SD=0.51$), yet this effect was not significant, $t(11)=1.358$, $p=0.10$ (one-tailed).

Table 7.1: Comparison of Random / Gaze-based Image Selection: The number of fixations increased significantly ($p = 0.015$).

Measure (mean per user)	random (control)	gaze-based (test)	overall	T-test (p. sample; df=11)	Sig. (one-tailed)
(a) # fixations	24.00	32.17	28.08	2.478	$p = 0.015$
(b) avg. fixation duration	52.20	44.48	48.34	0.550	n.s.
(c) dwell-time	1238.14	1429.75	1333.94	0.511	n.s.

Based on the semi-structured interview, only one participant realized that the Amazon web page was customized. Additionally, none of our participants realized that we adapted content based on gaze information. The quantitative and qualitative data as well as the user feedback show that probably no or very little conscious information processing took place during the study.

7.5 Study II: User Interest

In the second study we aimed to explore the effect of user interest on the gaze behavior. We collected data on a user's gaze towards certain product categories while surfing and later had the users fill out a questionnaire on self-perceived interest towards the categories. We expected a positive correlation (H2).

7.5.1 Setup and Procedure

We used the same setup as in the first study. We recruited participants from mailing lists, Facebook, and lectures. We had a total of 28 participants (avg. age 26.6 years). None participated in an eye-tracking study before (including the previously presented study) and all were familiar with the Amazon website.

The procedure was also similar to the first study. The participants were briefed, the eye tracker was calibrated, and the participants were asked to solve tasks on the Amazon website for 20 minutes. For the purpose of the study, we made sure that each product image was shown equally often and for the same amount of time. The content was *not* adapted to the user's gaze behavior in this study. After

finishing the tasks, we assessed the participants' interest (5-Point Likert scale, 1=not interesting at all, 5=very interesting) by asking them to rate 20 different product categories, five of which we could later associate with the images used in the study. The reason why the user had to rate more categories than used in the study was to make the topic of the study less obvious.

7.5.2 Data Analysis

During the study, we gathered 1,478,617 data points. We had to exclude the data of two participants due to problems with the eye tracker. To find any correlation between the user's gaze behavior and their interests, we analyzed whether the dwell time, the number of fixation, or the average time for fixation increased for the images associated with categories the participant had rated interesting.

7.5.3 Results

There is no significant correlation between the user's interests and gaze behavior. The results in detail are as follows: (1) For products which participants stated a high interest for, the number of fixations decreased ($r=-0.211$, n.s.). (2) Similarly, the dwell time decreased ($r=-0.144$, n.s.). (3) The average duration per fixation increased ($r=+0.136$; n.s.). All findings are non-significant and are likely to be the result of a random effect. This is, especially from a marketing perspective, interesting as it suggest that users may not be aware of their 'real' interests.

7.6 Summary and Discussion

This chapter showed how eye tracking could be used to adapt content towards users based on analyzing implicit gaze behavior. During the evaluation we found that (1) attention significantly increases for content that is adapted based on the user's gaze behavior, (2) there is a potential influence on the user's attitude, and (3) there is no correlation between interest and gaze behavior. Hence, hypothesis 1 was confirmed, whereas hypothesis 2 was refuted.

The observations during the study and the analysis demonstrate that implicit gaze interaction is a powerful modality for creating new user experiences. Without additional effort for the user, content can be tailored to increase their attention,

since the natural gaze movement is a rich resource for information about what they pays attention to. At the same time, the increase in the number of fixations shows that the approach has the potential to affect the users' attitude (cf., the peripheral route of the ELM). In market research and usability studies, eye tracking and offline analysis is best practice and commonly used. We demonstrated that with current technologies, it is possible to benefit from this information in real time. Our research has explored how this information can be used to adapt the user interface in real time and by these means, create an effective feedback loop. We observed that these findings have a major potential for the design and implementation of attentive UIs by creating novel and engaging ways for interacting with information systems.

So far, our research mainly focuses on image-based advertising scenarios. The results show that adapting web ads based on an on-the-fly analysis of gaze behavior is feasible and effective. We demonstrated that if ads are adapted based on the dwell time extracted from gaze in real-time, it is possible to redirect the users' attention towards more specific types of ads. We showed that detailing images the users looked at could help capture the user's attention. Hence, the users' gaze behavior can help provide reactive ads, where the system takes into account the users' attention. We envision, that systems with a potentially better user experience can eventually lead to a more positive perception of advertisements.

We believe, the fact that adaptive content draws attraction independent of the interest should be highly interesting for advertisers. It implicates that new products and services could be effectively advertised if there is a link to objects users are attentive to. An sample application could construct a longer chain of images leading from the user's initial attention towards an object the advertiser would like the user to look at, e.g., they looks initially at a car, then the next picture is a car in front of a house, and the following image is the house itself. We expect that such associative multi-step links may have a higher probability for the user to look at, however we have not comprehensively assessed this. Furthermore, we think that the approach described in this chapter can also be used complementary to traditional ways for targeting advertising, such as profiling.

At the same time, such a technology poses a risk that systems acquire information about the users, which they would rather keep privately for themselves. Our approach supports a non-individualistic customization and protects the user's privacy since it does not require any data to be collected and stored about the user. We believe that this is a strong advantage over other approaches. Even though the service provider can get information about the interests of the user during their interactions with the website, no user profile is generated or stored in the process.

During the study, the participants were asked about their concerns regarding privacy when their gaze data is used to determine ads they are shown. While most of the participants were interested in this new approach, some of them stated that they would turn off this feature if this option was provided. Most of the users did not like having information about their preferences in the hands of the website owners (or in our case the UsaProxy operator). These concerns should be investigated when deploying a system relying upon the user's approval for gathering gaze data.

Our study focused on the specific use case of image based ads on web pages. We chose this narrow focus as it is economically very important and, by narrowing the experiments, we hope to increase the reproducibility. So far, we have no quantitative evidence that these findings are valid for other application areas such as images with content other than advertising or non image-based content. We expect, however, that similar effects hold for other media, such as public displays. One limitation may be that interaction times in front of public displays are in general shorter than in the World Wide Web. Our approach collected 30 seconds of data, but shorter measurement phases may still return appropriate results. The approach may be particularly useful in cases where users are performing search tasks, e.g., browsing through classified ads.

Chapter 8

Communicating Interactivity

In traditional advertising, the major challenge public displays face, is how to attract attention. As displays are becoming interactive two further challenges emerge. First, displays need to communicate that they are interactive and second, they need to entice people to interact. Unlike privately owned devices, such as mobile phones or PCs, people simply do not know or expect that public displays are interactive – an effect that has been amplified by many displays having been used for static ads from their very advent. If public displays cannot communicate their interactivity, they will be hardly used and not fulfill their purpose. We believe that these issues will become even more apparent in the future as current LCD technology for public displays are likely to be replaced by technologies that more closely resemble traditional paper (e.g., e-paper [138]). As a consequence, passersby might not notice that a surface is digital, unless the content is constantly moving.

Relatively little is known about understanding interactivity, which is at the focus of this chapter. Previous solutions involve calls-to-action and attract loops [193]. A *call-to-action*, such as a ‘Touch to start’ label, can be effective. However, text or symbols are language and culture dependent and complex to understand subconsciously. *Attract loops*, such as a video of a person interacting, may create an atmosphere of an arcade game and be complex to understand as well. In the following we investigate how feedback to the passerby’s incidental movements (e.g., a mirror image) can be used to communicate the interactivity of a display. As humans are very efficient at recognizing human motion [80] as well as their own mirror image [241], this technique benefits from these perceptual mechanisms.

After discussing psychological foundations, we report and discuss the results of a lab and a field study. In the lab study we show that a real-time video image or silhouette of the user are equally effective for recognizing interactivity. Avatar-like and more abstract representations are less effective. In the field study we deployed and tested three displays in a shop window during three weeks.

Our observations show:

1. Significantly more passersby interact when they are immediately shown the mirrored user image (90% more) or silhouette (47% more) compared to a traditional attract sequence with call-to-action.
2. Passersby often recognize interactivity after they already passed. Hence they have to walk back – we call this the *landing effect*.
3. Often passersby notice interactivity because of somebody else already interacting. They position themselves in a way that allows them to see both the user and the display. This gives them an opportunity to understand the interactivity. If they start interacting themselves, they typically do so behind the person interacting hence forming multiple rows.

The reported observations can be useful for designers of public displays who want to communicate interactivity to passersby, and more generally, for any designer of devices where users do not know in advance that the device is interactive.

This chapter is based on the following publication:

- J. Müller, R. Walter, G. Bailly, M. Nischt, and F. Alt. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In *Proceedings of the 2012 ACM Conference on Human Factors in Computing Systems* (Austin, TX, USA), CHI'12, pages 297–306, New York, NY, USA, 2012. ACM

8.1 Related Work

Attracting *attention* with public displays and kiosks is not easy [147; 169; 250], and is described as the ‘first click problem’ [169]. Huang et al. observed passersby’s attention towards (non-interactive) public displays and show that

most displays receive little attention [147]. One solution is to use stimuli to attract attention [147; 250]. However, this is challenging in public space. Moving stimuli attract attention, but do not guarantee that the user looks, because there are many objects competing for the attention of the passerby [147]. Another solution suggests using physical objects. For instance Ju and Sirkin [169] show that a physical attract loop (animatronic hand) is twice as effective as a virtual attract loop (virtual hand). While physical objects seem to attract more attention than digital content, they are less flexible and more difficult to update with new content. An overview of the role of attention and motivation as requirements for public displays is provided in Müller et al. [250]. Among others, they report that looming stimuli (moving towards the observer) and motion onset can be used to attract attention.

A literature review identified six commonly applied techniques for *communicating interactivity* of both public displays and tabletops:

1. A *call-to-action* [193], often a simple text label such as ‘touch screen to begin’ was used in [169; 193; 216].
2. An *attract sequence* is originally described as a slideshow [193]. Some multi-touch installations used constantly moving objects [141; 281]. Arcade machines also use a video that either explains the interaction or shows a user performing the interaction. A similar technique is to use constantly moving content, e.g., [281].
3. Nearby *analog signage*, either with a simple call-to-action or a more complex manual, has been used in many deployments, e.g., [193; 216; 281].
4. The *honeypot effect* [52] describes the effect of people being attracted by persons already interacting with a device. Brignull et al. observed this effect and divide the people around the display into the phases peripheral attention, focused attention, and interacting. Further observations of the honeypot effect are reported in [216; 281; 238]. For the CityWall [281], for example, it was observed that people most often notice the wall when someone is interacting with it (in 19% of the cases). In other cases, some passersby had difficulty to notice that the display was interactive.
5. *Persons inviting passersby to interact* can be either users who have already noticed the interactivity and now motivate their friends [216; 281], or researchers standing next to the device inviting users and explaining the interaction [164]. Students are employed as so-called UbiGuides in Oulu, Finland, in order to motivate people to use the displays [271].

6. *Prior knowledge* that a device is interactive can be used if users pass by the same device multiple times, or if they are familiar with the device (e.g., the Microsoft Surface [216]).

After people notice interactivity, *immediate usability* is important. The term was introduced in the context of Shneiderman's CHI photo kiosk [193]. Shneiderman derives three recommendations: (1) Implement an attract sequence tailored for the audience; clearly indicate how to end the attract sequence and begin using the system (e.g., using a call-to-action such as 'touch screen to begin'). (2) Support zero-trial learning. Users should be able to use the interface after observing others or using it themselves for a brief period of time (15-60 s). (3) Encourage users to immediately interact with the content. Users who were not immediately successful would often simply abandon the device. Marshall et al. observed that even a delay of a few seconds after touching an interactive tabletop is problematic [216]. Users are likely to give up and think that the device is not interactive or broken.

Perceived affordances [261] are derived from Gibson's concept of affordances, which are properties of an organism's environment that have a certain relation to the body and skills of the organism. These properties make certain actions affordable. While affordances exist independently of their perception, it is important how users can perceive them. More recently, Norman proposed the more general concept of signifiers [263]. Signifiers may be any information in the environment, which indicate that a certain action is possible or appropriate. This is especially interesting in the context of public displays. For example, smears on a screen may indicate that it is a touch screen.

Several researchers have proposed to use a *shadow or mirror image*³² of users of large displays to indicate and support interaction. They have been used in the context of artistic installations [188], pointing tasks on large displays [338], and interaction above a tabletop [140]. Michelis and Meckel [237] deployed public displays showing a camera image of what was happening in front of the screen. Images are augmented with digital effects guided by motion, like clouds of numbers or growing flowers. The focus of this study centers on the motivation to interact rather than noticing interactivity. Thus, no different user representations were compared and no baseline like call to action was tested. While these works explored various aspects of shadow and mirror metaphors, their application and properties to communicate interactivity of displays were not explored.

³² The main difference between the shadow and mirror metaphors is that when the user is further away from the display, the mirror image gets smaller, while the shadow gets larger. A mirror metaphor also allows to show a detailed RGB image of the user, while this seems unnatural for a shadow.

Attract sequence and call-to-action are practical solutions to communicate interactivity. In the following we explore the representations through mirror images as an alternative.

8.1.1 Psychological Cues & Interactivity

When it comes to noticing interactivity, several psychological concepts provide useful hints as to how such an interactive system should be designed. Figure 8.1 shows that for a certain interaction it is possible to compare whether the manipulation is intentional (or not) and whether the user noticed the effect (or not). Dix et al. discuss a continuum of intentionality between explicit and incidental interaction [93]. Explicit interaction refers to the case where users intentionally manipulate an interactive system. Incidental interaction refers to situations where the interaction is neither intended nor is the effect noticed afterwards. An example is when a user enters a room and the temperature is adjusted automatically, without the user noticing.

A similar concept is implicit interaction [317], which describes situations where the user interacts without being aware of the interaction. When users are aware of the fact that they are interacting, implicit and incidental interaction turn into explicit interaction. We use the term *inadvertent interaction* to describe the situation where users manipulate a device incidentally, but become aware of the effect and thus learns that the device is interactive. When users perceive that a device reacts to their incidental movements, this reaction can be perceived in three different ways. It can be perceived as (1) a *representation* of the user (e.g., a mirror image), (2) an effect *caused* by the user, or (3) an *animate* being or thing reacting to the user. Powerful perceptual mechanisms exist for all of these perceptions. While the focus of this chapter is on the *representation* of the user, a short review of psychological foundations for perceptions of causality and animacy is provided.

8.1.2 Representation: Recognizing Yourself

There are two ways how you could potentially recognize yourself in a mirror: appearance matching and kinesthetic-visual matching [241]. *Appearance matching* is based on a comparison of the image seen in a mirror with the knowledge of how you look like. *Kinesthetic-visual matching* is based on the correlation between the own motion and the visual feedback in the mirror. The question whether an

	Noticed Effect	Unnoticed Effect
Intentional Manipulation	Explicit Interaction	Incomprehensible Interaction
Unintentional Manipulation	Inadvertent Interaction	Incidental / Implicit Interaction

Figure 8.1: While incidental / implicit interaction assumes that the user does not notice the effect, we can distinguish the case where the user inadvertently interacts and then sees the effect.

organism can recognize itself in a mirror has been a topic of investigation since the early work of Gallup [116]. They learned that only humans, chimpanzees, and orang-utans show this behavior. Humans can recognize themselves already in the first months of life [241]. For recognizing somebody else's reflection in a mirror, visual-visual matching can be used if we can see both the person and the reflection. This is presumably easier than kinesthetic-visual matching, which is learned early on in childhood.

When users control a representation of themselves on a display (e.g., a mouse pointer or a mirror image), they need to understand that they are in control. This is similar to the questions in psychology of how humans perceive which part of the world is one's own body (ownership) and controlled by oneself (agency) [165]. It is assumed that information from the own action generation mechanism (intentions) are compared in a tight loop to proprioceptive, tactile, and visual feedback to iteratively correct motion. For example, the perceived position of the pointer is iteratively compared to the intended position (Fitts's Law). The same sources of information (intentions and feedback) are tested for congruence and correlation to determine which parts of the world we control.

From this we learn:

- Visual feedback can override proprioceptive feedback, such that people feel agency for parts of the world which are not actually their own body. As a consequence, people might forget about their real surroundings when immersed in the virtual representation – an assumption that was confirmed during the field study reported later on, where some users were so engaged in the game that they accidentally hit their neighbors. People assume more often that they control something that they do not actually control than vice versa (over-attribution).

- People can experience a continuum between more and less agency, depending on the correlation (amount of noise and delay) [165]. It is important to minimize noise and delay to improve the perception of agency.

In order to leverage these mechanisms, interactive systems could use a representation of the user's presence and appearance like a mirror image (recognizing one's own visual appearance). An abstract representation of the user's motion might, however, also be sufficient (kinaesthetic-visual matching). Unfortunately, we are not aware of any studies comparing different user representations with regard to how efficiently users can recognize themselves.

8.1.3 Abstraction, Biological Motion, Body Schema

Humans can use not only appearance matching, but also kinesthetic-visual matching, to recognize their mirror image. Hence, it is possible to abstract user representation and still have users recognize themselves. This gives the designer of a device much more artistic freedom when designing the user representation. Fortunately, humans have direct perception of the motion of humans and animals from minimal information. It was shown that a video of a dynamic array of point lights (at skeletal joints) is sufficient to see the presence of a walker [80]. For gender recognition, the upper body joints are more relevant (70% accuracy), and adding more points in addition to shoulders, elbows, wrists, and hips does not improve accuracy [184]. From static images of point lights without motion however, not even the presence of a human can be seen. For this section it is especially interesting that we can recognize ourselves and friends, and that we are more effective in recognizing ourselves (43% accuracy) than our friends (36%, 16.7% chance), despite the fact that we see our friends walking more often [80]. This is explained by the fact that both executed and perceived motion are represented in isomorphic representations (the body schema) and can easily be translated into each other.

A system could use minimal representations similar to point light displays to represent users, but it is very important that the representation is dynamic. Upper body parts like wrists and torso might be most effective. In order to use the body schema for representation, however, the feedback needs to directly match to the movements of specific body parts (e.g., head or hand). More abstract feedback that cannot directly be matched to body parts (e.g., averages of the movements of multiple body parts) often needs more time to be recognized [383].

8.1.4 Perceptual Causality and Animacy

Besides for self-recognition, humans also have perceptual mechanisms for causality and animacy. This is easily demonstrated by 2-D movies of simple moving geometric shapes [322]. If one object ‘hits’ a second object, and this second object is ‘pushed’ away, humans have a strong impression that the first object caused the motion of the second. If there is more than a 50-100 ms delay between the two events, this perception starts to disappear. Objects that start from a resting position, change direction to avoid collision, or have directed movement towards a goal can appear to be ‘alive’ [322]. Perceptual causality and animacy can be used to communicate interactivity, and in these cases, known cues causing these perceptions should be used (e.g., collision). In particular, causality can be combined with mirror representations. Since interaction with mirror representations alone is not motivating enough, physics simulations provide motivating interaction and increase the perception of interactivity.

8.1.5 Summary

This chapter focuses on the representation of the user as a cue to interactivity, because such a user representation is a very general tool to support multiple interaction techniques. From these psychological foundations, we learn the following:

1. There are efficient perceptual mechanisms that support this self-recognition.
2. It is unclear how recognition of oneself degrades when the representation is abstracted.
3. It seems crucial that the correlation between the user’s movement and feedback is high (low noise and delay).
4. In order to use the efficient body-schema representation, the feedback should be directly matchable to a certain body part.
5. User representations can be combined with perceptual causality (or animacy) to strengthen the perception of interactivity and provide a more interesting application.

8.2 Prototype

To explore how the user representation can communicate the interactivity of displays, we conducted a series of three user studies. We developed a prototype that was successively refined based on the study results. During these studies the focus was on noticing interactivity rather than attention or motivation. We relied on the motion of the user representation to capture attention and on a very simple ball game (Figure 8.2) to motivate users. More elaborate attention grabbing or motivating techniques would probably increase the total number and duration of interactions.

8.2.1 Software

We use the 3-D rendering capabilities of *OpenGL* to display the user's mirrored image or silhouette and other virtual objects. For detection of users, we rely on the *OpenNI* framework, which provides unique IDs and pixel masks to separate them from the background. The mirrored user representations are directly embedded into the scene to the lower half of the screen (Figure 8.2) and interacts with other virtual objects (balls). We use the *Bullet* physics library to simulate the behavior of these objects constrained to a 2-D plane. Since the simulation is optimized for rigid bodies, we approximate the user's shape with small objects along their contour, which are continuously tracked between frames. We record the depth image stream and user activities for later analysis.

8.2.2 Hardware

The system was deployed on large portrait-oriented public display LCD screens of different dimensions ranging from 40" to 65". The Microsoft Kinect sensor was employed to detect passersby and users.

8.3 Pre Study

In order to see if and how passersby are interacting with a public display, we conducted a pre-study. Our prototype showed the silhouette of the passerby on a 46" portrait LCD monitor. Passersby could interact with a virtual ball using



Figure 8.2: Looking Glass Prototype: The user representation is embedded into the scene.

simulated physics. The display was installed for three days around lunchtime in front of a university cafeteria. Users were observed from a hidden position and interviewed on opportunity basis. Inter-rater reliability was satisfactory (Cohen's Kappa=0.61) [199]. We observed 832 passersby, of which 456 (54.8%) looked at the display, 171 (20.6%) interacted with the display, and 141 (16.9%) stopped walking to interact. People played for 2 to 182 seconds ($\mu=26$ s), and most of them stated that they left due to time constraints. Interestingly, most persons interacted in groups – most single passersby rather hurried past the display.

There are two important conclusions from this study. First, a large percentage of all passersby interacted (in a university setting), so the design is very promising for our purpose. Second, almost no passerby interacted alone. As our design supports only single users, this posed problems as mostly groups of 2-5 users tried to interact simultaneously. Also, almost all passersby stopped moving before interacting, while we expected more interacting while passing by the display.

8.4 Lab Study

The pre-study was followed by a controlled lab study where we removed the attraction and motivation criteria. Hence, we could measure the time required to recognize whether the test application was in an interactive or non-interactive (video playback) mode. Furthermore, the study included the influence of the user representations, for which we evaluated multiple levels of abstractions.

8.4.1 Objective

The objective of this study was to determine the impact of the abstraction of the representation of the user on how quickly users can notice that a display is interactive. We compared the user representations *mirror*, *silhouette*, *avatar*, and *abstract*. In this study, we only focused on noticing interactivity. Participants were asked to pay attention to the display and decide whether the display reacted to their movements or not. No additional virtual objects, that would potentially have biased the motivation of the participants, were shown on the screen. This lab study setup provided a baseline of how quickly users can decide whether a display is interactive under optimal conditions using the different representations. The lab design provided a high degree of control, while at the same time providing a lower degree of ecological validity. To counterbalance, the study was followed by a field study, which offers high ecological validity but less control.

8.4.2 Conditions

We used the following conditions in the lab study:

1. **Mirror Image:** an interactive colored image of the user.
2. **Silhouette:** a white filled silhouette of the user.
3. **Avatar:** a 2-D avatar including head, torso, and hands.
4. **Abstract:** just the head of the user, with abstract eyes and mouth.

All of these conditions can be directly matched to body parts by the user (see Section 8.1.1). For the expected interaction distance at the shop windows the camera could not capture both feet and head of the user. Based on the studies of point-light displays that show that upper body parts are most relevant, we decided to position the camera in a way such that these parts were visible. Based on the same studies, we expect the gain in speed and accuracy from adding feet to the avatar to be low. Related work on stimulus-response compatibility [383] indicates that stimuli that can be directly matched to body parts are more effective than those which cannot. Therefore, we decided for the abstract condition to directly represent the head of the user (instead of, e.g., an average of multiple body parts). All four of these interactive conditions were also presented as non-interactive conditions. In this case, a video of another user interacting with the display was

started as soon as the user stepped in front of the display. These non-interactive conditions should simulate situations where either just a random video was shown on a display, or a different user (e.g., standing behind the participant) would interact with the display.

8.4.3 Task and Stimulus

Users were asked to walk back and forth past the display while following a line on the ground placed at a constant 2 m from the display. On the display, one of the 4x2 different conditions was shown. Users carried a device (Logitech Presenter) and were asked to click on the left button when they believed the display to react to their movements, and the right button when they believed the display did not react to their movements. Users were asked to be as fast and accurate as possible. Time was measured from the moment when they entered the field of view (FOV) of the camera (and thus appeared on the screen in the interactive conditions) until they pressed a button.

8.4.4 Apparatus and Design

An 82" portrait LCD display was used to present the content. The representation of the user was created using a Microsoft Kinect camera and software using *OpenNI*, *NITE*, and *Processing*.

A within subjects design was used with $n=16$ participants recruited from a pool of non-computer scientists. Variables measured were time and accuracy. These 4x2 conditions were repeated in 10 blocks. The order was counterbalanced using a latin square within the participants, and randomized between the participants.

8.4.5 Results

The *selection time* is measured as the time from when the first stimulus appeared (as the user entered the camera's FOV) to the time when the user made a choice. An ANOVA revealed a significant effect for *representation* on selection time ($F_{3,45} = 80.76, p < .0001$). It also revealed a *representation * interactivity* interaction effect on selection time ($F_{3,45} = 6.75, p < .0001$). A post-hoc Tukey test showed that Mirror (1.2 s) and Silhouette (1.6 s) are significantly faster than

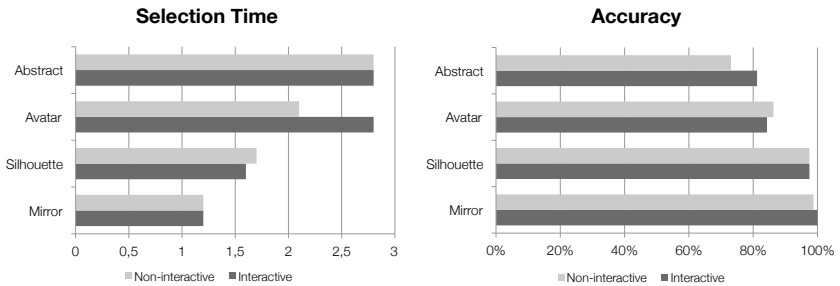


Figure 8.3: Results from the Lab Study: For ‘Mirror’ and ‘Silhouette’, the selection time was shortest (left) and accuracy was highest (right) in the interactive condition.

Avatar (2.8 s) and Abstract (2.8 s) in the interactive condition (Figure 8.3, left). In the non-interactive condition, Mirror (1.2 s) is significantly faster than Silhouette (1.7 s) and Avatar (2.1 s) which is significantly faster than Abstract (2.8 s).

An ANOVA further revealed a significant effect for representation on *accuracy* ($F_{3,45} = 43.09, p < .0001$). It also revealed a representation * interactivity interaction effect on accuracy ($F_{3,45} = 5.84, p < .0001$). A post-hoc Tukey test shows that Mirror (100%) and Silhouette (97.5%) are significantly more accurate than Abstract (84.3%) and Avatar (81.2%) in the interactive condition (Figure 8.3, right). In the non-interactive condition, Mirror (98.8%) and Silhouette (97.5%) are significantly more accurate than Avatar (86.3%) which is significantly more accurate than Abstract (73.1%). Finally, the ANOVA revealed a significant effect for *block id* on accuracy ($F_{9,135} = 5.84, p < .0001$). A post-hoc Tukey test shows that users are less accurate in the first block (74.2%) than in the other blocks (mean: 91.6%).

8.4.6 Summary

From this experiment we learn the following:

1. The Mirror and Silhouette representation are similarly efficient, but both more efficient than the Avatar and Abstract representation.
2. It takes considerable time to distinguish the interactive and the non-interactive conditions even in an optimal environment (1.2 s vs. 1.6 s).



Figure 8.4: Study Location: Displays were finally installed in three shop windows (b, e, f).

The fact that the Silhouette is efficient is valuable, because it provides much more artistic freedom for the designer of a display. While the lab study provided control, ecological validity was low. Therefore, we decided to compare the two most promising representations, Mirror and Silhouette, to a combination of two common traditional techniques, call-to-action and attract loop, and a purely causal technique in a field study.

8.5 Field Study

The objective of this study is to explore how users would notice interactivity and interact with public displays using different user representations ‘in the wild’. We compared the two most effective user representations, Mirror and Silhouette, to the most common strategy in industry, call-to-action combined with an attract loop, and a merely causal condition without user representation. This comparison looks at the ability to attract users with the display as well as their general effect on the social situation in an urban place. A field study was chosen in order to maximize ecological validity, sacrificing the control of the lab.

8.5.1 Deployment

Three displays were deployed for three weeks in shop windows of a store in the city center of Berlin (Figure 8.5). Windows on one side of the store (D, E, F) were close to a well-frequented sidewalk, windows on the other side (A, B) were near a subway entrance. To decide which windows to install the displays, we observed 200 passersby of the street-facing side of the store (C, D, E, F) during

a weekday afternoon until night. The observations showed that there are large differences in how many passersby look into each shop window. The percentages are: Main door C (6%), small window D (12%), small window (13%), small bright window (19%), large window E (29%), small window (16%), large window (29%), second door F (large and bright, 33%). Even 66% looked into the window who walked from right to left making the second door the first window in their path. It seems that the large and bright windows attract more attention, especially if surrounding windows differ. Also for people walking from right to left, we noticed a large percentage (17%) looking straight, and did not notice the last window. Apparently, they looked down a road at the crossing.

For the deployment, we used three LCD monitors in portrait format (65", 46", 46"). Cameras (Microsoft Kinect) were installed below the monitors. For the first week of deployment we moved the displays between the windows A, B, C, D, E, F (Figure 8.5). While window B had the advantage that people could play relatively disturbance free from passersby, windows E and F had a larger number of passersby and attracted most views. Therefore we decided to install the 65" display in window B, and two 46" displays in windows E and F. For the background image we tried different artistic contents, but could not observe a large influence of the content on behavior. The final content was an ad for the store and was created by a professional advertising agency.

8.5.2 Conditions

In our study we tested two variables:

1. User representation (Conditions: Mirror, Silhouette, No-Representation)
2. Interactivity cue (Conditions: inadvertent interaction, attract sequence with call-to-action)

Regarding the application, we opted for a very simple ball game. Ten balls are displayed on the screen, and users can play with them (kick them) using the contour of their representation. The whole game takes place in the 2-D plane of the user representation. In the *Mirror* condition, the user's image from the color camera is extracted from the background and shown on the display. In the *Silhouette* condition, the silhouette of the user is shown on the display, and in the *No-Representation* condition, just the balls are rendered, but no user representation is shown (but interaction is similar to other conditions). In the



Figure 8.5: User Representations: Mirror (a), Silhouette (b), and No-Representation (c). All three representations were tested in an ‘attract loop with call-to-action’ as well as in an ‘inadvertent interaction’ version. In this figure, the corresponding attract loops (a video of somebody stepping close to the display and starting to interact) are shown. In the inadvertent interaction condition, the person in front of the display is shown in the same representation, just without the call-to-action (‘Step Close to Play’).

inadvertent interaction condition, when nobody is in front of the screen, just the background image and balls are shown. The interaction starts as soon as users entered the FOV of the camera. In the *attract sequence with call-to-action* condition, a video of a person demonstrating the interaction is shown together with a label ‘Step Close to Play’ (Table 7.1). The video shows a person in the corresponding visualization (Mirror, Silhouette, and No-Representation) stepping close to the camera and then playing with the balls.

When the user enters the FOV of the camera with a closer distance (1 m), the screen switches to interaction mode. The user is represented in the corresponding visualization and can play with the balls. Conditions were counterbalanced and automatically switched every 30 minutes. This was done to minimize the influence of time of day on the results.

8.5.3 Data Analysis

We collected both qualitative and quantitative data. Qualitative data was gathered from observations, semi-structured interviews, and manual video recording. As quantitative data, complete interaction logs (from NITE person tracking) and videos from the depth camera were kept for each display. For anonymity reasons we only recorded the (anonymous) depth image (Figure 8.10).

Qualitative data was collected daily for three weeks. Since displays worked best in the late afternoon and evening when most of the interaction occurs, at least two researchers were present during these times. Additional observations were conducted as needed. Observations were conducted from inconspicuous positions like from the other side of the street or near the subway entrance, where it was common to see waiting people. During the observations, video was recorded using video cameras that looked similar to mobile phones (FLIP HD³³). Furthermore, field notes were kept. Interesting findings were presented and discussed in daily meetings with the entire research team. Eventually, the team agreed on a specific focus for following observations.

From the depth videos, we recorded roughly 1500 hours of videos. We selected 11 consecutive days for manual coding. We implemented an analysis software that automatically searched the log files for scenes in which a user was detected in the visual field of the camera for more than 4 seconds. In accordance to [216] and [281], interactions which followed each other within less than 20s were merged to single sessions. All sessions were then manually reviewed and annotated. The coded measures can be found in Table 8.1. We observed 363 interactions. Inter-rater reliability was substantial (Cohen's Kappa=.75) [199].

8.5.4 Findings

Mirror, Silhouette, and Call-to-action

The total number of interactions during the 11 coded days is shown in Table 8.2. We compared the number of interactions per day. An ANOVA reveals a significant effect for interactivity cue (call-to-action vs. inadvertent interaction) ($F_{1,11} = 12.6p < .001$). A post-hoc Tukey test shows that passersby interact more with the inadvertent interaction condition than with the call-for-action. The ANOVA also reveals a significant effect for user representation ($F_{2,22} = 13.1$).

³³ Flip website: <http://www.theflip.com>, last accessed February 13, 2013

Table 8.1: Measures Coded With Our MatLab Software

Measure	Explanation
InteractionStartFrame	First frame where somebody in this video interacted.
InteractionEndFrame	Last frame where somebody in this video interacted.
LandingEffect	The person who interacted walked past and came back to interact.
ButtBrushEffect	Somebody interacts, then other people pass by, and the interacting person leaves.
HoneypotEffect	Somebody interacts, then other people join who are from a different group.
PassingByInteraction	Somebody interacts without stopping.
MultipleRows	People seem to interact in multiple rows behind each other.
UserCollisions	Some people collide while interacting.
Interesting	Something interesting happens.

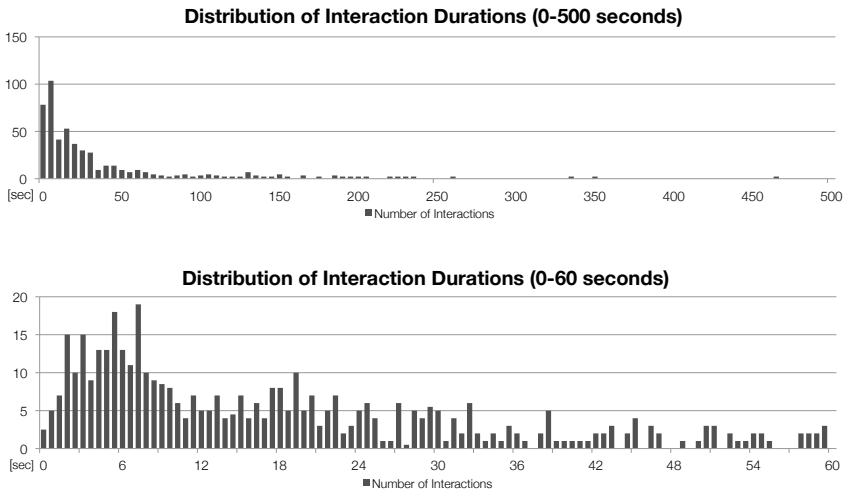


Figure 8.6: Interaction Durations: In order to investigate how well the different conditions communicate interactivity, we needed a large number of situations where nobody was currently interacting with the screen. We intentionally designed an interaction that demotivated extended play. The mean interaction duration was 31 s, but many interactions only lasted for a few seconds. Some users seemed to be motivated to play for many minutes.

Table 8.2: Total Number of Interactions (during 11 days of field study): Inadvertent interaction attracts significantly more interactions than call-to-action. Further, Mirror works significantly better than Silhouette and No-Representation.

	No-Representation	Silhouette	Mirror
Call-to-action	67	59	79
Inadvertent interaction	60	87	150




Figure 8.7: In the call-to-action condition people sometimes spent considerable time in front of the display (a) before stepping closer (b). In this case, the two women are surprised to see themselves and walk away (c). On the next window, they encounter inadvertent interaction in the silhouette condition and start playing (d).

A post-hoc Tukey test shows that Mirror is more efficient than Silhouette and No-Representation. Finally, the ANOVA also reveals a user representation * interactivity cue interaction ($F_{2,22} = 6.8, p < 0.005$). As expected, there are no significant differences between the user representations for call-to-action. User representations differ only in the inadvertent interaction condition. Many interactions with the display only last for seconds (Figure 8.6).

The interviews revealed different preferences for the user representations. The shop owner prefers the Silhouette as people are covered in company colors. There is no clear user preference, and many say that they like the representation they discovered first. Users who prefer the Mirror representation describe it as more ‘authentic’, more ‘fun’, and they like to see themselves and their friends. Users who prefer the Silhouette representation described it as more ‘anonymous’ and said that they like it when bystanders can not see their image. Some also say that they do not like to see themselves and prefer the Silhouette representation. In the Image representation, also some users mention that they do not like to be observed by a camera, which they do not say about the Silhouette representation.




Figure 8.8: Over time, knowledge about the display built up. In the morning we observed pupils expectantly waiting at the traffic light (a) before crossing the street (b) and playing with the display (c).

From our observations, we found that in the call-to-action conditions, people spend several seconds in front of the display before following the instructions (‘Step Close to Play’) (compare Figure 8.7). In this vignette, two women observe the display for some time, before one of them steps closer and activates the interaction in the Mirror condition. They are surprised to see themselves and walk away. A few meters further, they notice a second display running the inadvertent interaction Silhouette condition, where they start to play.

When interviewed on how they noticed interactivity, most people say that they saw themselves on the display. Some also say that they saw themselves and a friend / partner at the same time. Only very few stated to have seen the representation of another person walking in front of them.

When a crowd had gathered around the display, it was sometimes very difficult to distinguish who caused which effect. This was especially true for the Silhouette and obviously the No-Representation conditions. In these cases we observed people copying the movements of other users and seemingly interacting with the screen, even though they are not represented on the screen. Sometimes they are not even standing in the field of view of the camera. This can be an example of *overattribution* (compare Section 8.1.1), where people assume they are causing some effects although they are not.

Over time, knowledge about the presence built up and interactivity had built up among people who pass the location regularly. In the third week of deployment, a number of people who interacted said that they had seen somebody else interacting, e.g., ‘a few weeks ago’ or ‘earlier that day’, but had not tried to interact themselves. There were also a few regular players. For example, we noticed from the logs that between 7-8 am, there was considerable activity in front of the displays. Observations revealed that a number of children played regularly with the displays on their way to school. We observed them waiting expectantly



Figure 8.9: Landing Effect for a Group: A group of people passes the display (a). Only at the next shop windows person A stops (b), turns around, and walks back to the display (c). As he starts interacting (d) more and more people from the group join (e).

at the traffic light, then crossing the street directly to the display to play with it (Figure 8.8). Such interaction is obviously different from situations where people encounter the displays for the first time.

Design Recommendations

Inadvertent interaction outperforms the attract loop with call-to-action in attracting interactions. The Mirror representation also outperforms the Silhouette and interaction without user representation. In contrast to the lab study, the Mirror representation works significantly better than the Silhouette. From this we learn that Mirror representations is a powerful cue to communicate interactivity, although Silhouettes may have some benefits such as more artistic freedom in designing the content and provide more anonymity. As most people recognize themselves on the display rather than someone else, displays should be positioned so that people can see themselves well when passing by. Over time, as knowledge about the interactive device builds up, these interactivity cues become less important.

The Landing Effect

One striking observation regarding the moment when people start to interact was that often, people stop after they passed the display and walk back to it again (see Figure 8.9 for this effect in a group, and Figure 8.10 for this effect with a couple). In Figure 8.9, a group of young men is passing the display. One person (A) in the group looks at the display but keeps on walking with the group. After a few meters further, the person suddenly turns around and walks back, followed by a second person. They then start to interact, and are soon joined by other group members. In this paper we refer to these cases as the *landing effect*.




Figure 8.10: Landing Effect for a Couple: As the couple passes by, the woman notices the screen and stops. As her partner walks on, she drags him back to the screen. Both start interacting.



Figure 8.11: The Honeypot Effect: As people notice a person making uncommon gestures, they position themselves in a way allowing both the screen as well as the interacting person to be seen. They also often position themselves so that they are not represented on the screen.

Regarding the number of landing effects, ANOVA reveals a significant effect for interactivity cue ($F_{1,11}=23.1, p < 0.0001$). A post-hoc Tukey test shows that more landing effects are observed in inadvertent interaction (18.5% of all interactions) than in call-to-action (8%). There was no significant effect for visualization.

We observed this behavior only for people passing by the displays (not waiting), when no one is interacting with the displays, and who apparently do not know already that the displays were interactive (e.g., because they already interacted with them). The landing effect often leads to conflicts when one person in a group notices the interactivity. If the first persons in a group suddenly stops and turns around, the following people would sometimes bump into them. More often, the whole group stops rather than walking on. However, when a following person in a group notices interactivity, the first would usually walk on for some time before they notice that somebody stopped and stop themselves. This situation creates a

tension in groups. The people could either walk back and join the person who stopped or abandon the person interacting with the display and join the group. In some cases, the group simply walks on after waiting for some time, causing the interacting person to continue to play only for a short moment and then hurry (sometimes even run) to re-join the group. Interviews revealed more details about this behavior. One man who had walked back (Mirror condition) answered that he had seen from the corner of his eye two persons on the screen walking in the same direction. He was curious and walked back, accompanied by his wife. When he saw himself on the display, he understood that it was interactive and explained it to his wife. They both started to play with it. For another couple, the man stated that he saw something moving from the corner of the eye and walked back. His wife stopped, but did not follow him. He noticed that the display was interactive upon seeing himself, but only played very shortly before joining his wife. It is quite possible that users did not interact, because they only noticed interactivity after they had already passed the displays and did not want to walk back.

As we installed multiple displays along the same trajectory, passersby had the option to notice interactivity on one screen, but then interact with another one. When they noticed the second screen, they already expected that it was also interactive and stop earlier. One man said he noticed the balls jumping away on the first screen, but did not walk back. When he noticed the second screen, he decided to stop his friend. They saw their representations and played shortly. Often, after playing with one screen, people also searched the other windows for further screens and also played with those (Figure 8.7).

Design Recommendations

The landing effect is in line with our observation from the lab. People need approximately 1.2 s (Mirror) and 1.6 s (Silhouette) to recognize interactivity. They also need to notice the display first and be motivated to interact. With an average walking speed of 1.4 m/s, by the time passersby have decided to interact, they already passed the display. This effect is so strong that it should be designed for in any public display installation. Displays should be placed so that, when people decide to interact, they are still in front of the display and do not have to walk back. Optimally, friends walking in front of them should also still be able to interact with the display without walking back. This could be achieved by designing very wide displays, or more practically, a series of displays along the same trajectory. If possible, another solution would be to place displays in a way so that users walk directly towards them.




Figure 8.12: Multiple Rows: The girl from group A noticed interactivity first. Woman B positioned herself behind them to see what happens and also started interacting. Later, a couple C stopped behind them and started interacting in a third row.

Dynamics Between Groups

We observed many situations in which different groups started to interact. The first group (or person) usually causes what has been previously termed the ‘honey-pot effect’ [52]. We found that people passing by first observed somebody making unconventional movements while looking into a shop window (the manipulation [299]). They subsequently positioned themselves in a way that allowed them to see and understand the reason for these movements – usually in a location that allowed both the persons interacting as well as the display to be seen (Figure 8.11). In this figure, a man interacting with the display with expressive gestures attracts considerable attention. The crowd stops and stares at him and the display and ends up partially blocking the way for other passersby. Newcomers seem to be first attracted by the crowd, then follow their gaze, then see the man interacting, follow his gaze and eventually reposition themselves so they can see both the man and the display. They also seem to prefer to stand a bit to the side, so that they are not represented on the screen. The audience is mostly positioned behind the user. We observed this pattern regularly. When people in the audience decided to join the interaction, they accordingly did so *behind* the ones already interacting and not next to them (Figure 8.12). In this figure, the little girl in the front noticed the interactivity first, followed by her mother, who then stopped to explore the display together with the daughter (the father did not walk back and is standing behind the camera). The young woman behind them was attracted by their interaction and eventually also interacted behind them. This again attracted the couple behind them, and the girl finally also started interacting in a third row. In some cases, such multiple rows form again from people observing at the subway entrance. In

the few cases where other people started to interact in the same row as people already interacting, we were able to observe social interaction between the users, which we did not observe for different groups interacting behind each other.

People interacting with the screens were usually standing in the way of others. The resulting conflicts were solved in different ways. For the screen installed near the subway entrance, passersby usually tried to pass behind the ones already interacting, to not disturb them. When multiple rows of people interacted, this was not possible however, and they passed in front of them (Figure 8.12). When a large group passed by, we sometimes observed that the person interacting abandoned the display. This again sometimes led to that someone from the arriving group took the place. We also saw some occasions, where users deliberately moved between the display and the person interacting and interacted for a short moment.

Design Recommendations

The honeypot effect is a very powerful cue to attract attention and communicate interactivity. Displays which manage to attract many people interacting will be able to attract more and more people. The honeypot effect works even after multiple days, as people who have seen somebody interacting previously may also try the interaction in the future (see Section 8.5.4 Mirror, Silhouette, and Call-to-Action). To achieve this, displays should be designed to have someone visibly interacting with them as often as possible. This can be achieved by increasing the motivation and persuasion for people to play longer. Because the audience reposition themselves so that they can see both the user and the display, the environment needs to be designed to support this. In our case, both the subway entrance and the narrow sidewalk limited the possible size of the audience. In order to support more audience, displays should be visible from a wide angle, or considerable space should be available directly in front of the displays. This is also necessary as different groups start to interact behind each other. This interaction behind each other should also be supported, e.g., by increasing the maximum interaction distance beyond the distance from where single groups normally interact.

Dynamics Within Groups

We discovered that the vast majority of interactions are from people traveling in a group. The only cases of single people interacting we observed personally are the children before or after school hours, men waiting for a considerable amount of

time near the subway entrance, a man in rags, and a man filming himself while playing. One man for example waited for several minutes directly in front of one screen, while incidentally interacting with it through his movements. After some time, he was approached by an apparent stranger, who showed him the display and the fact that he was interacting. The man seemed surprised, and continued to play a little bit with the display. While a considerable number of single people pass by the store, they usually walk faster and look more straight ahead and downwards. When we interviewed some of them, only very few had noticed the screens at all, and nobody had noticed that the screens were interactive. Between one and five persons interacted simultaneously ($\mu = 1.5$). Often, the first person in a group noticed the display first, but this was not always the case.

We discovered that people strongly engage with the game and apparently identify more with their representation on the screen than the possible influence of their movements on people around them (see Section 8.1.1). This sometimes leads to situations where people are not aware anymore of their neighbors (people belonging to one group usually line up next to each other), even though they are able to see their representation on the screen. This focus on the virtual space leads in some situations to that people accidentally hit or bump into each other.

Another observation was that people usually start interaction with very subtle movements and continuously increase the expressiveness of their movements. This process sometimes takes just a few seconds and sometimes extends over many minutes. The subtle movements at the beginning are sometimes just slight movements of the head or the food. Later, people proceed with extensive gestures using both arms, jumping, and even acrobatic movements like high kicks with the legs.

Design Recommendations

The most important observation from this section is that very few persons who are alone, interact. This observation is supported by the results of the pre-study. Therefore it is important to understand how groups notice interactivity, and public displays should always be designed to support groups. Even if just one person is interacting, the display must provide some value for the other group members. When users strongly engage with their representation on the screen, they may forget about their real surroundings. According to our observations, slow moving objects make users move slower, which increases safety.

8.6 Summary

This chapter looked at how interactivity of public displays can be effectively communicated using a mirror representation of the user. Summarizing the findings from our lab and field study, the following can be learnt from this chapter:

1. Using the mirror image of users so that passersby inadvertently interact with public displays is an effective way of communicating interactivity. Mirror images are more effective than silhouettes and avatars, and more effective than a traditional attract loop with a call-to-action.
2. Noticing interactivity needs some time, which leads to the landing effect. When passersby decide to interact with public displays, they have often already passed them, so they have to walk back. This can be mediated e.g., by installing multiple displays in a row.
3. Users from a different group often start to interact behind the ones already interacting, forming multiple rows. Because the vast majority of interacting persons are also in groups, public displays should support multiple users, in particular when interacting behind each other.

We hope that mirror representations for inadvertent interaction will also be applied to other devices beyond public displays, e.g., tables or floors. Finally, we believe that public displays that effectively communicate their interactivity have the potential to make urban spaces all over the world more fun and engaging to be in.

Chapter 9

Enticing Interaction

As soon as a passerby notices that a display is interactive, the final step towards making public display advertising more attractive is to entice users to interact. This is a complex process where two requirements need to be fulfilled prior to the start of user interaction: on one hand, users need to be *motivated to interact*, on the other hand they need to *understand how to interact*. Unless both requirements are fulfilled, people will not start to explicitly interact. For example, a passerby might be motivated to interact but will not understand the employed interaction technique (e.g., they keep touching the display and do not understand that it supports only gesture-based interaction). On the other hand, a user might quickly understand how the interaction works, but not be motivated to interact (e.g., out of the fear that the application is not protecting their privacy well enough).

How to motivate people to use technology in general and to interact with displays in particular has been subject to research for many years. One of the most well-known models that explains why a person uses a technology or not is the so-called Technology Acceptance Model (TAM) developed by Davis at the MIT in the late 1980s [84; 85]. In its original form, the model postulates that a person's attitude (or motivation) towards using a technology depends on two variables: the perceived usefulness and the perceived ease of use. The model was refined in 2000 to also reflect social influences and cognitive processes [368]. With regard to public displays, Michelis investigated different factors that impact a user's motivation to use public displays, including challenge and control, curiosity and exploration, choice, fantasy and metaphor, and collaboration [236] (see Section 2.4.2). In this chapter we will show, that as novel technologies and

applications emerge, additional factors will gain importance. The level to which the content is of interest to the user determines how motivated the user will be to interact. Second, providing means that help to preserve the user's privacy will influence their motivation to interact. Third, future public display applications may require entering sensitive information such as user names and passwords. As smears on the display may reveal passwords, or shoulder surfers can spy upon passwords entered via an onscreen keyboard, ways to enable secure data input become more important.

As new sensing technologies are deployed, many novel interaction techniques and modalities make interaction with public displays possible. This creates an increasing need for the providers of public displays and content to ensure that passersby understand how to interact. As will be seen in the remainder of this chapter, this is challenging as people, on one hand, draw from their knowledge with familiar technologies such as surfing the Internet. On the other hand, novel devices enable interaction techniques previously unknown to the user, such as gesture-based interaction (e.g., using the Microsoft Kinect) or gaze-based interaction (e.g., using eye trackers).

In order to better understand the challenges of enticing users to interact with public displays, we draw upon findings presented in Chapter 4. In a descriptive study we elaborated on current practices that emerge as people interact with traditional public display space. The focus of the study was on content and how it can be managed, the interplay between space, stakeholders, and content, as well as the needs and motivations of content providers and display owners [11]. Following our design implications we implemented a prototype of a digital public notice area (PNA) called *Digifieds*. We evaluated the platform with regard to suitable interaction techniques, the users' privacy concerns, and preferred content. The evaluations were performed in the lab and during a real-world deployment. In addition, we looked more closely at security as an increasing challenge. We present an approach that allows users to authenticate with their gaze and report on a user study that aims at evaluating the security of the approach.

This chapter is based on the following publications:

- F. Alt, A. Sahami Shirazi, T. Kubitzka, and A. Schmidt. Interaction techniques for creating and exchanging content with public displays. In *Proceedings of the 2012 ACM Conference on Human Factors in Computing Systems* (Paris, France), CHI'13, New York, NY, USA, 2013. ACM

- F. Alt, T. Kubitzka, D. Bial, F. Zaidan, M. Ortel, B. Zurmaar, T. Lewen, A. S. Shirazi, and A. Schmidt. Digifieds: Insights into Deploying Digital Public Notice Areas in the Wild. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia* (Beijing, China), MUM'11, pages 165–174, New York, NY, USA, 2011. ACM
- F. Alt, N. Memarovic, I. Elhart, D. Bial, A. Schmidt, M. Langheinrich, G. Harboe, E. Huang, and M. P. Scipioni. Designing Shared Public Display Networks: Implications from Today's Paper-based Notice Areas. In *Proceedings of the 9th International Conference on Pervasive Computing* (San Francisco, CA, USA), Pervasive'11, pages 258–275, Berlin, Heidelberg, 2011. Springer-Verlag
- F. Alt, D. Bial, T. Kubitzka, A. S. Shirazi, M. Ortel, B. Zurmaar, F. Zaidan, T. Lewen, and A. Schmidt. Digifieds: Evaluating Suitable Interaction Techniques for Shared Public Notice Areas. In *Adjunct Proceedings of Pervasive 2011* (San Francisco, CA, USA), Pervasive'11, 2011
- A. Bulling, F. Alt, and A. Schmidt. Increasing The Security Of Gaze-Based Cued-Recall Graphical Passwords Using Saliency Masks. In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems* (Austin, TX, USA), CHI'12, pages 3011–3020, New York, NY, USA, 2012. ACM

9.1 Related Work

Various projects explore the technical requirements for networking and interacting with digital displays within and across offices [3], as well as in public spaces [256; 353]. MAGIC Broker allows people to interact using SMS, the WWW, and speech. It consists of separate gateways for each interaction method and allows several user interfaces to be used in parallel [98]. Paek et al. suggested I/O modules for using different techniques simultaneously, providing a similar solution [276]. In order to support parallel interaction, tunneling interaction via a server was explored, hence creating an indirect but abstract form of communication [145; 310].

Current research emphasizes the trend towards and potential for interconnecting displays for sharing information in an attempt to create a new communication medium. Its infrastructure aims at supporting multiple types of parallel interaction and attempts to reach as large an audience as possible. The impact is demonstrated by several studies. Churchill et al. assess the influence of introducing public displays into an office space [71; 72] and published insights into the augmentation of the user environment. McCarthy et al. designed CoCollage, a community supporting social network applications for public displays [223]. These studies demonstrate how user awareness is increased by introducing displays and analyzes what kind of content is used in communication. Moreover, public display networks also support the development of communities [295; 296] and can even act as a meeting point for users with common interests [360] – similar to traditional, (still) paper-based public notice areas.

When it comes to interaction with public displays, two classes of interaction techniques prevail. On one hand, a large proportion of the HCI community focuses on using mobile phones for interaction, hence allowing the user to browse and create data with their personal device. Usually, users are so familiar with their phones that little learning is required. Several mobile phone interaction techniques have been proposed. Sahami et al. use the mobile phone's flashlight [335], while Ballagas et al. use the camera to control a cursor on large displays [25]. PhoneTouch is a novel technique enabling phones to select targets by direct touch [321]. Hyakutake et al. embed transparent markers on the display allowing any camera-enabled device to interact with the display [152]. The touch projector enables interaction with remote displays using a live video image on the mobile device [45]. Several dedicated mobile applications have been developed, which make use of different connectivity options such as Bluetooth or WiFi for socket or web-based communication. Also SMS and MMS have been subject to research [83; 357]. More recently, the iPad also has been used as an interaction device for public displays [195].

The advent of cheap (multi-) touch technologies has shifted direct interaction with displays into the focus of recent research. Touch surfaces allow single or multiple users to interact with the display in parallel (e.g., collaboration or browsing at the same time). Nevertheless, keyboards, mice, and devices such as levers or buttons [144] turned out to be good alternatives, as they are well known to users and promise fast adaption. NFC technology can be used to simulate button-based interaction behavior in the digital world [133]. Finally, Nawaz et al. explore eye gaze and head gestures [257].

Furthermore, novel interaction methods have been studied with regard to user behavior [52; 238] and technologies [105; 125]. Of particular interest is Huang's finding that people spend less time learning about system capabilities when it does not support current practices [149]. Also, user motivation and interests in novel systems need to be taken into account [238].

Previous findings emphasize the need to embed existing routines into a novel system to support its usability. In our opinion, this entails not only catering to user needs, but also reflecting on other stakeholder motivations. PNA owners are typically interested in attracting people and reaching customers. Hence, a PNA can be seen as a central part of an environment. People stop by to read posts or leave messages while they are on the go. The specific interaction is not usually the main aim of a passerby. Typically, posting or reading is combined or triggered by the user's primary task, e.g., shopping or waiting. Furthermore, information on a PNA proposes locality. Information is related to people visiting the area and shops, increasing the attraction of certain activities. For example, a PNA in a music store is more attractable with respect to searching for a guitar, than it would be in the case of a computer store. Finally, interaction with a digital PNA must be simple and quick. Several technologies and modalities have been explored but have so far barely considered the user's current situation. Posting via email, SMS, or direct input as well as taking away content via Bluetooth and HTTP communication are familiar, but their suitability, advantages, and disadvantages have yet not been analyzed, which therefore constitutes the focus of the work presented in this chapter.

9.2 Lessons Learned From Traditional PNAs

Traditional public notice areas can be found in a wide variety of locations and are still very popular. People use them to post classifieds, information on events, community activities, and the like. At the same time, platforms such as Craigslist³⁴ or eBay³⁵ offer similar and successful services on a national or even global scale. Therefore, simply deploying these platforms for public displays seems obviously viable. However, the success of traditional, paper-based PNAs indicates that people often prefer them compared to online platforms.

³⁴ Craigslist website: <http://www.craigslist.org>, last accessed February 13, 2013

³⁵ eBay website: <http://www.ebay.com>, last accessed February 13, 2013

Previous work (see Chapter 4) has identified the following reasons: on one hand, traditional PNAs are *very easy to use* and they have low entry barriers, as paper and pen can easily be used, even when people are spontaneously passing by. The fact that no additional equipment, such as a PC or a mobile phone is required makes it possible to generate messages within an extremely short timeframe (often less than one minute) and without any need to be familiar with the technology (which is often an important issue, particularly for elderly persons). Furthermore, physically taking away information with one (e.g., tear-aways or flyers) contributes to high usability. On the other hand, PNAs have a strong *local character*, often addressing certain communities only. In contrast, web-based platforms allow a large number of people to be reached, even at a distance. Yet, there are various situations in which such platforms are inconvenient. When selling items, especially those which are difficult to handle due to size and weight and cannot easily be shipped (e.g., furniture, bikes), addressing the local community provides a better opportunity to find buyers who can easily pick up the items personally. The same is true for services such as babysitting or private lessons, which cannot be offered supra-regionally. Finally, the generally high level of trust in local communities, contributes to a good seller/buyer relationship.

These reasons indicate that many different aspects should be considered when designing digital, networked PNAs. Merely providing a digital version of current online solutions is not advisable and ignores important aspects of making traditional PNAs successful. In the following we summarize the core implications of the study that are of interest for the remainder of the chapter. Further information on the study can be found in Chapter 4.

Design for Different Types of Content

The study revealed different types of displays, that strongly relate to the specific content being published. Therefore, attention needs to be directed towards the PNA's location and intended purpose, as well as towards the type of content and how its design can be supported. Whereas for ad-hoc posters, it might be important to generate content as easily and quickly as possible, many people devote considerable effort to creating more eye-catching content. Hence it is important to offer means to create content with a default layout, while at the same time providing ways of styling and augmenting posts with different fonts, colors, images, videos, or (interactive) maps. For professional advertisers, formats such as presentations or videos should be supported.

Support Different Interaction Techniques

The main factor behind the success of a PNA is that content can be created, posted, and retrieved easily and flexibly. Content can be created at the display, but also prepared at home or on the go. Consequently, suitable interaction techniques need to be supported that keep the use of digital PNAs simple. Options include direct touch at the display, using mobile phones on the go, or the PC at home to create content. In order to exchange content with the display, means for easily transferring data to the display need to be provided. At the same time, an adequate way of retrieving content (similar to taking a copy of a flyer/tear-away) is required, e.g., allowing content to be sent to an email address, to be transferred to the phone, or to be printed.

Preserve Local Character of the Display

Many displays create a central location for community activities [11; 360]. Content observed on such displays in fact supports the assumption that exchanging goods locally is very popular. Hence, items can be found, which are either difficult to ship, or which are of special interest to the community. Although the potential for remote posting can add to the uptake of digital PNAs, users need to be allowed to specify where content should be posted.

9.3 Research Questions

New challenges arise as people start using public displays to create user-generated (advertising) content. Hence, we were mainly interested in content and its specific impact on the user's interest as well as in privacy as motivating factors. Furthermore, we were looking at suitable interaction techniques that could make people quickly understand how to interact with the display and that would potentially preserve the users' privacy.

9.3.1 Content

Traditional public notice areas are highly popular, even in the times of platforms such as eBay and Craigslist, which offer similar services. We found that content on traditional PNAs usually has a strong local character. For services, such as babysitting or cleaning, that cannot easily be offered supra-regionally, and for

offering items that are difficult to ship (e.g., a bed or a bike), public notice areas provide a good opportunity to find a local audience. As a result, buyers could easily pick up things by car. Furthermore, content informing about local events is especially interesting to the local community and tourists.

RQ 1 – Which content do people post and retrieve from digital PNAs?

To gain a better understanding of which types of content should be supported by a digital PNA, we interviewed people and asked them to fill out a questionnaire on content they are interested in on traditional PNAs and on content they would expect on digital PNAs. Furthermore, we analyzed the content posted during the evaluation period. Finally, we compared content on traditional PNAs from previous work with the content posted on Digifieds.

9.3.2 Interaction Techniques

The success of traditional PNAs lies in their high usability. Pen and paper allow content to be posted by everyone, also ad-hoc, and tear-aways as well as flyers available in multiple copies allow information to be retrieved quickly with ease. As a consequence, suitable interaction techniques that realize a similar functionality and are highly intuitive as well as easy to use need to be provided.

RQ 2 – What are suitable interaction techniques?

For the evaluation we implemented a display client and a mobile client. In a field trial we aimed at evaluating both clients with regard to usability and conducted semi-structured interviews in order to identify potential shortcomings and issues.

9.3.3 Privacy

Asking people to input private information (e.g., an email address) in public space as well as making such information available on the display might prevent many potential content providers from using the PNA. Though there is some evidence from traditional PNAs that in anticipation of the envisioned benefit (e.g., selling the advertised item) it is ok for people to provide this information, taking into account users' privacy concerns might add to a further and quicker uptake.




Figure 9.1: Digifieds – Conceptual Architecture: A central server stores the content. The display client periodically synchronizes with the server and both a mobile and a web client enable remote access.

RQ 3 – Which privacy issues arise while interacting with displays in public space?

In order to investigate privacy issues, we conducted the field trial in a public space, hence creating a realistic atmosphere. Afterwards, the users had to fill out a questionnaire in which we assessed privacy concerns with regard to traditional and digital PNAs and conducted semi-structured interviews.

9.4 Prototype

In order to tackle our research questions we developed a prototype of a digital PNA, called *Digifieds* – derived from digital classifieds. Note, that in the remainder of this chapter we use the term *Digifieds* to describe our platform, whereas *digified* describes a classified ad posted on the platform. The Digifieds platform consists of four components (Figure 9.2): (1) a central server back-end for the data management, (2) a web-based display client for visualizing information and direct interaction, (3) a mobile phone client as an alternative interface for interaction with the display, and (4) a public web client. Additionally, a web-based administration interface for content and configuration management is provided.

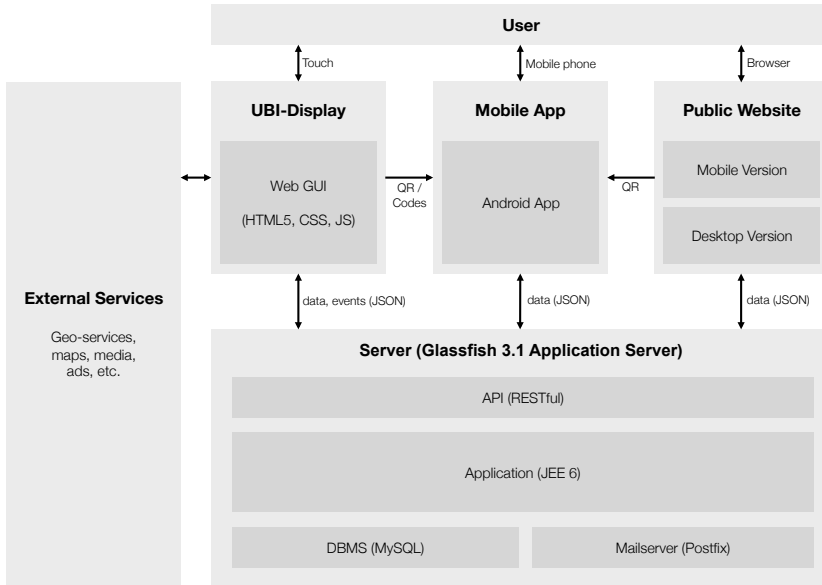


Figure 9.2: Digifieds – System Architecture: A Glassfish application server stores the data. The clients access the content via the RESTful API.

The prototype consists of a client-server infrastructure that allows arbitrary display clients to be connected. We also provide a mobile phone application that allows content to be created on-the-go and to be transferred to the display. Similarly, content can also be transferred from the display to the mobile phone.

9.4.1 Digifieds Server

The Digifieds server is the central component of the system. It is responsible for the data management and storage and provides access for arbitrary clients (display, mobile phone, web) through the RESTful API. In order to provide a robust server application we opted for the Java Enterprise Edition 6 Framework (JEE6). The Glassfish 3.1 application server ensures scalability (easy thread management and clustering) and trouble free updating of the running server application without compromising active sessions. A MySQL database stores data permanently and can be accessed through a Java Persistence API (JPA) layer. Caching optimizes database access and hence reduces CPU usage and overall access times.




Figure 9.3: Conceptual Layout of the Display Client: The display client provides different views that can be scrolled horizontally.



Figure 9.4: Digifieds Display Client: The display client shows an overview of the different views on the left, the active view, containing the actual digifieds, is located on the right.

Besides storing the content and layout information of each digified, the central database manages the information, configuration, as well as available categories for each connected display client. For evaluation purposes, all API interactions can be logged. The lightweight JSON data format transfers data between server and the connected display clients. An XML format, e.g., for use by external applications, can also be used by simply changing the corresponding HTTP request headers.

9.4.2 Digifieds Display Client

The main goal when designing the GUI of the display client is to preserve the advantages of traditional paper-based PNAs while at the same time enhancing it with digital features, such as multimedia content (pictures, videos), interactive content (maps), popularity-by-click count, sorting posts by various criteria (date, popularity), automated removal of outdated messages, search functionality, and novel retrieval techniques. For the display client's graphical layout, three challenges needed to be tackled. First, it has to be recognizable as a PNA, not just as a digital display; second, content has to be presented in a well-arranged manner (even if containing a lot of content); third, interaction has to be enabled in a very easy and intuitive way.

We adapted the layout of traditional PNAs, making content look like paper classifieds attached to a wall. In order to cope with scarce space, we decided not to display all content on one single screen but to split the PNA into several views. The concept is depicted in Figure 9.3. Each view holds posts related to a certain category, e.g., 'Housing', 'Sales', etc. Using buttons on the left and right side of the display enables switching between these views horizontally. In case a single view is overloaded with entries, it can also be scrolled vertically. The dimension of the active view adjusts automatically to the screen resolution. The background layout of a PNA can also be customized for each category. Using different views for scaffolding does not only help to solve the space issue but we also envision easing the use of the board and making browsing more convenient.

Finally, if displays are touch-enabled, the client provides an on-screen keyboard that allows users to create and send posts without using additional devices. Users can choose color and category of the digified from predefined values. Using the system does not require any registration or login process. Digifieds can also be retrieved in different ways, e.g., by sending them to an email address, printing them out, or using one of the mobile phone techniques described in the following section. A shopping cart function allows multiple digifieds to be retrieved easily at the same time.

The display client uses AJAX to create an interactive UI capable of attracting and enticing people through immediate feedback. HTML5 and CSS are used to layout the content. Using asynchronous HTTP requests, the display client periodically polls for data changes. If there is any new content, the corresponding GUI elements are updated. Currently, the default update rate is 30 seconds. However, since each display's configuration can be modified on the server, adjustment to arbitrary update rates as well as dynamic rates based on the data load can be easily



Figure 9.5: Digifieds Mobile Client: The mobile client allows for creating and retrieving classified ads. Furthermore, it stores all content that have been previously created or retrieved.



Figure 9.6: Retrieving Content: Users can send digifieds to their email address or transfer them to their mobile phone using either the QR code or the alphanumeric code.

realized. The internal browser cache minimizes the data traffic and is used for media documents (images, videos, HTML, CSS) and the browser’s local storage API saves the digified’s data in JSON format.

9.4.3 Digifieds Mobile Phone Client

In order to allow content to be created on-the-go, we developed an Android application. With this client the user can create new digifieds, containing a title and content (text, images and/or videos), and define additional information such as the expiration date, address, or contact data. The digifieds are permanently stored for future use on the phone and in the central database.

To enable exchanging content between phone and display in a transparent and understandable way, we implemented three techniques:

1. **Phone/Display Touch:** We implemented an interaction technique where the user can touch the display with the phone at an arbitrary position. In the posting mode, the digified created on the phone is transferred to the display and inserted at the touched position. In the retrieval mode, the digified located at the position where the user touches the screen with the phone is transferred to the mobile phone. The phone/display touch feature is implemented by synchronizing actions between the display and the phone. Once a touch gesture is detected, the phone and display are matched via timestamps. Subsequently, the selected post is transferred via the Digifieds server between the devices (Figure 9.10b).
2. **Alphanumeric Code:** Once the user creates a new post on the phone, it is stored on the server and assigned an ID. Then, on the server, a 5-character alphanumeric code (e.g., 4XB6A) is generated from this ID and displayed on the phone. This code can then be entered on the display (Figure 9.6).
3. **QR Code:** To transfer a digified from the display to the mobile phone, we provide a QR code next to each classified. The QR code can either be used to open the classified in the mobile browser, or, if it is scanned with the mobile phone client, be kept on the phone. Alike, entries on the screen can be transferred to the phone just by entering the alphanumeric code displayed next to each digified in the provided field on the display client (Figure 9.6).

The alphanumeric code and the QR code are used for two reasons. First, they identify the display or display group (see Section 9.6) the digified will be displayed on. Second, in order to preserve the locality of the display, we wanted people to personally come to the display. Note, that technically remote posting on a display could easily be implemented.



Figure 9.7: Digifieds Website: The website can display a digified based on the code retrieved from the display. Additionally, it provides a map of the display locations as well as a link to the mobile client in the Android market.

Finally, the mobile client provides an interactive map with the locations of all Digifieds-enabled public displays.

9.4.4 Digifieds Web Client

For people who do not own a smartphone or who prefer composing their digified on a PC at home or at work, we provide a public website³⁶. This website serves two purposes. First, it provides further information about the Digifieds platform, e.g., a tutorial about how to use it, information on where to find displays running Digifieds (interactive map), and a download link to the mobile app in the Android Market. Second, similar to the display and mobile client, the website can be used to create new digifieds or retrieve digifieds that have been found on one of the public displays. When creating a digified on the website, images and videos from the local PC can be embedded, and the PC keyboard as well as the computer monitor may be used to create sophisticated designs in a more flexible way. However, like with the mobile app, a created digified still has to be activated using the display client before becoming publicly visible. In order to retrieve

³⁶ Digifieds website: <http://www.digifieds.org/>, last accessed February 13, 2013

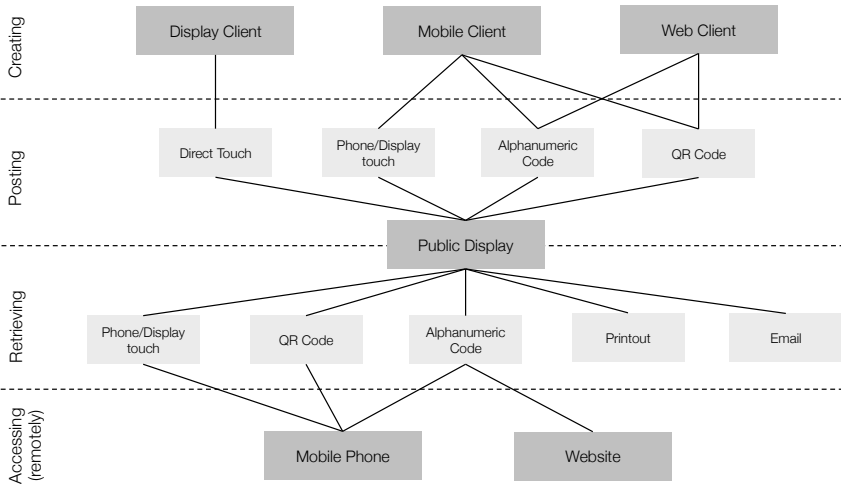


Figure 9.8: Overview of Interacting with Digifieds: Content can be created, posted, retrieved, and remotely accessed using different devices and techniques.

one or many digifieds that are on a public display, a user only needs to enter the alphanumeric code. Subsequently, the original digified, that includes all images, videos, maps, and a form to contact the owner is displayed on the website.

9.5 Lab Study

We evaluated the Digifieds client in a lab study in order to identify the most suitable interaction technique. We opted for an initial lab study, because for the anticipated measurement, a highly controllable environment was required for statistical data analysis to be performed post-hoc. This would have been difficult ‘in the wild’ due to external influences. Furthermore, to assess user behavior and to enable interaction, we used cameras, which would have been a major issue in public due to privacy concerns.




Figure 9.9: Creating Content for a PNA: (a) at the display, (b) on the phone, (c) on the PC.

9.5.1 Interaction Techniques

To identify the most suitable interaction technique we provide different interaction techniques to be used at the display, on the phone, and on a PC (Figure 9.8).

Creating Content

We support three situations for creating content. First, creating content using the *display client* is meant for ad-hoc posters, either coincidentally passing a display (e.g., on the way to the supermarket) or in a waiting situation (e.g., at a bus stop). The display client allows content to be created directly at the display by means of an on-screen keyboard (Figure 9.9a). Users could enter text, choose the background, and augment it with an image/video from a USB stick. Second, the *mobile phone client* can be used by people to create content on-the-go and prepare it for publishing it later at the display. The mobile client enables users to create a post by entering text (similar to writing an SMS) and transferring it later to the display (Figure 9.9b). The message could be augmented with images or videos taken with the mobile phone. Third, through the *web client*, users can create digifieds remotely at a PC, e.g., at home, at work, or on the go from a laptop (Figure 9.9c).

Posting Content

Content created at the display using *direct touch* input is stored directly and appears on the screen (Figure 9.10a). Apart from this, there are three different ways to transfer pre-generated content using additional interaction techniques. First, we provide the *phone/display touch feature*. After creating a post, users can touch the screen with the phone at the position where they want it to appear

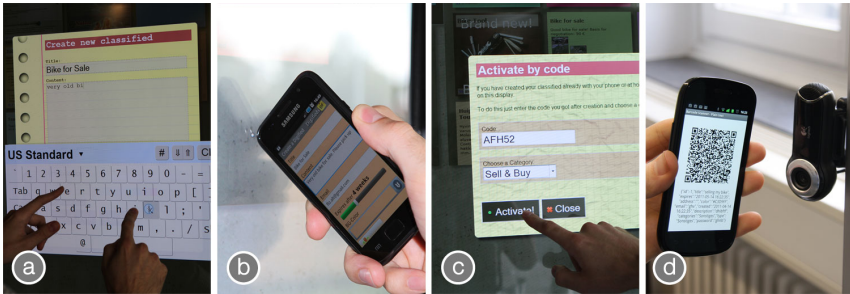


Figure 9.10: Posting Content on a PNA: (a) Direct touch at the display, (b) Phone/display touch, (c) Alphanumeric code, (d) QR code.

(Figure 9.10b). Second, *alphanumeric codes* can be used in a similar way to activate a previously generated digified at the display (Figure 9.10c). Third, we use the *QR code technique*. After users create a post on the phone, a QR code is generated which is then captured by the camera attached to the display (Figure 9.10d). The QR code is read and the digified placed on the screen. Similarly, QR codes could be generated from print-out posts created on the web site.

Retrieving Content

As users often want to take content with them from the display, we provide five options: *phone/display touch* – similar to posting on the display (Figure 9.11a), transferring it to the phone via *QR code* by scanning the code next to the post with the phone’s camera (Figure 9.11b), using the alphanumeric code (Figure 9.11c), sending via *email* by providing the address directly at the display (Figure 9.11d), and *printing* it out on the printer next to the display (Figure 9.11e).

Accessing Content

Using either the QR code, or the alphanumeric code, or the phone/display touch feature allows for remote access to content on the phone later (offline). The alphanumeric code can also be used to access information via the Digifieds website (online).




Figure 9.11: Retrieving Content from a PNA: (a) Phone/display touch, (b) QR code, (c) Alphanumeric code, (d) Email, (e) Paper printout.

9.5.2 Tasks

In order to evaluate the suggested interaction techniques with regard to usability and suitability, the users had to solve familiar tasks from traditional PNAs (e.g., selling items, etc.) in the lab study. We developed authentic tasks to simulate situations in which participants were free to behave around the display, as they would normally, such as in a supermarket. Hence, we created a controllable, yet realistic situation in which people behaved both naturally and were not aware of what was being measured, thus avoiding having any influence on their behavior.

For each of the tasks the users were asked to (1) generate a digified on a given topic (e.g., selling a bike, selling a mobile phone, renting their apartment, offering private lessons), (2) post it on the display, and (3) retrieve one of the digifieds. In order to include all interaction techniques, each task combined a set of three techniques (one each for creating, posting, and retrieving content³⁷). We used a within-subject design. The task order and the given topics were counter balanced.

Task 1: Display

For the first task, we asked the users to imagine a situation in which they wanted to spontaneously post a digified, e.g., when passing a display in the supermarket. They were requested to create a digified directly at the display using the on-screen keyboard. Once they were finished, we asked them to look for a particular older digified and send it to their email address (Figure 9.11a).

³⁷ Note, that the alphanumeric code techniques was developed as a substitute for the phone/display touch technique that would not work on capacitive displays. Hence, this technique was not included in the evaluation.



Figure 9.12: Study Arrangement: (a) Digifieds display client, (b) Digifieds web client on PC, (c) Scanning QR code from paper.

Task 2: Phone/Display Touch

Second, we asked users to create a digified on the phone with the knowledge that they would pass a display on the way. Once they were finished, they were to post it on the display using the phone/display touch feature. Before leaving the PNA, they had to pick up a specific digified by using their phone.

Task 3: QR codes

In the QR code task, users also created their digified on the phone. To post it on the display, they had to use the QR code technique. Therefore the mobile client generated a QR code, which users had to present to a camera attached to the display. The display client reads the QR code and publishes the digified on the screen. The users also had to look for a certain classified and take a picture of the associated QR code.

Task 4: Paper

For the final task, users generated their digified at a PC, simulating preparation of a post, at home or at work. The web client generated a QR code of the digified, which users had to print out and present to the display's camera. Users then had to search for a certain digified and print it out using the printer installed next to the display.

9.5.3 Setup and Data Collection

For the study, the Digifieds system was set up in our lab. The display client ran on a 42" touch-enabled public display (Figure 9.12a), containing an initial set of digifieds in different categories. The system was reset after each participant, so that everyone started with the same initial set of digifieds. Note that digifieds of different length would have influenced the subjects with regard to the post length.

In addition to the display, we setup a PC and a printer simulating a home/work environment. The PC initially displayed the Digifieds website in the Firefox browser (Figure 9.12b). A camera was attached to the display in order to allow scanning QR codes from paper or from the mobile phone (Figure 9.12c).

We collected the following quantitative data, derived from questionnaires, from a server logfile, and from video recording followed by a post-hoc data analysis.

- **Demographics:** We collected data on the age, gender, and profession of the participants.
- **Mobile Phone Usage:** We are interested in the subjects' habits when using their mobile phones. Therefore, we asked whether their phone supported multi-touch, if they surf the web on the phone, and whether they had a contract with an unlimited data package, and if they use third party apps.
- **System Usability:** After each task, the subject filled out a System Usability Scale (SUS) questionnaire [26], so that we can compare the perceived usability of all interaction techniques.
- **Task Completion Time:** For each technique, we measured how long it took the participants to create, post, and retrieve content. Measurement was conducted post-hoc, based on video footage.
- **Length of Content:** We analyzed the length (number of characters) of the digifieds the users created for the different interaction techniques.

Qualitative data was gathered via further questionnaires and semi-structured interviews, after the study. We asked about any problems, personal perceptions, likability, and areas for improvement. The interviews were videotaped for post-hoc transcription.

9.5.4 Procedure

Participants were recruited via mailing lists, Facebook, bulletins, and on an opportunity basis as people passed by the lab during the study. After a short briefing and signing a consent form, we asked them to fill out the demographic questionnaire. We then led them to the room where the Digifieds system was deployed. We gave each participant five minutes to explore the system and asked them to ‘think aloud’ and provide feedback on what they were exploring. Next, they were given a phone (Samsung Galaxy S) with the preinstalled mobile client and given another five minutes to make themselves familiar with the Digifieds application. The users then had to complete the four different tasks presented above in randomized order. We read out the task description to each user.

“Imagine that you are at the display in a shopping center and want to sell your bike. Please create a classified ad using the display client and post it in the ‘Sales’ category.” (Sample task)

There was no time limit for completing the task. The experimenter did not intervene and only answered questions if the subject got stuck or explicitly asked for help. After finishing each task, the user filled out the SUS and provided written feedback. We interviewed them afterwards.

9.5.5 Findings from the Lab

Prior to analyzing our data statistically, we performed a post-hoc video data analysis, and coded input time, time uploading content, and time downloading content. We transcribed the interviews and printed them out. Qualitative findings were pasted on a wall in a meeting room, so that each team member could familiarize themselves with the data. The data were then discussed and annotated, and patterns as well as high-level observations were collected on a separate board.

Quantitative Results

In total, 20 participants (10 male, avg. age 26.8 years) participated in the study. They were students, employees, and civil servants, most without a computer-science background. 9 participants had a touch-enabled phone, 10 used their phones for surfing the web (avg.=10.85 times/day, SD=13.1), 13 had unlimited Internet access on the phone, and 11 used third party apps.

The analysis of the *SUS* reveals the following scores: display (86.6), paper (82.1), QR code (73.5), phone/display touch (70.0). A series of paired t-tests indicate that the usability in the display task is rated significantly better than phone/display touch ($T=-4.25$, $p<.001$) and QR codes ($T=-3.37$, $p<.01$). Paper-based interaction is ranked significantly higher than phone/display touch ($T=-4.25$, $p<.05$).

Using paired t-tests (Table 9.1), we compare the *interaction times* of different techniques with regard to creating, posting, and retrieving content. We found that *creating content* on a PC/laptop is significantly faster than directly at the display ($T=-3.480$, $p<.01$) or on the phone ($T=2.897$, $p<.01$). More interestingly, there are no significant differences between phone and display. For *posting content*, scanning the printed QR code on paper is significantly faster than both QR codes on the phone ($T=3.374$, $p<.01$) and phone/display touch ($T=3.050$, $p<.01$). This could be due to the fact that on the phone, the functionality of posting content had to first be activated. With respect to *retrieving content*, printing is significantly faster than digitally through phone/display touch ($T=2.149$, $p<.05$) or QR code ($T=3.110$, $p<.01$); yet, sending via email is not significantly slower than printing.

With regard to the *length of post* (number of characters), we found that posts created on the PC/laptop are significantly longer than texts created on a phone ($T=-3.716$, $p<.001$) and at the display ($T=3.373$, $p<.01$). There is no significant difference in the length of digifieds created on the mobile phone and directly at the display ($T=-.707$, n.s).

Finally, we looked for *correlations* in our data using Pearson's "r" and tested for significance. The most important findings are:

1. People who use touch devices frequently write significantly less text at the display ($r=-.490$, $p<.05$) and on the phone ($r=-.054$, $p<.822$). Additionally, there is a significant correlation between the amount of text written on the phone and on the display ($r=.580$, $p<.01$).
2. The more often and the longer people use the phone, the better they perform with the phone-based techniques (positive correlation for amount of written text ($r=-.573$, $p<.01$), interaction time for retrieving content with QR codes ($r=-.448$, $p<.05$), rating of phone/display touch ($r=-.505$, $p<.05$)).
3. We found strong correlations between age and usage of the phone-based techniques. Younger participants perform significantly better when uploading content with phone/display touch ($r=.586$, $p<.01$) and when retrieving messages with the phone's QR scanner ($r=.512$, $p<.05$). Furthermore, the *SUS* rating for QR codes correlates negatively with age ($r=-.499$, $p<.05$).

Table 9.1: Differences in Task Completion Times (paired T-Tests; $df = 19$).

	Mean A[sec]	Mean B[sec]	Std. dev.	T	Sign.
Creating Content (A vs. B)					
Mobile vs. Display	182.10	166.80	85.10	.804	.431
PC vs. Mobile	184.54	114.75	107.7	2.897	.009
PC vs. Display	114.75	166.80	66.88	-3.5	.003
Posting Content (A vs. B)					
QR _{phone} vs. QR _{paper}	36.10	17.70	24.39	3.374	.003
P/D _{touch} vs. QR _{paper}	29.24	17.70	16.92	3.050	.007
Retrieving content (A vs. B)					
QR _{phone} vs. Paper	50.70	37.70	30.13	3.110	.006
P/D _{touch} vs. Paper	49.55	37.70	41.20	2.149	.045
Paper vs. Email	37.70	29.75	19.86	-1.8	.089

Next, we analyzed how performance impacts usability. We found that in cases where inputting and retrieving messages with the phone takes longer, participants rate the usability of this technique significantly lower (e.g., creating messages ($r=-.451$, $p<.05$), retrieving messages using QR codes ($r=-.829$, $p<.001$)).

To reveal differences in gender and for users with unlimited mobile Internet access, we performed an ANOVA. Participants who have unlimited mobile Internet access are significantly faster when creating content using the phone (on avg. 230% faster, $F=11.838$, $n=7$, $p<.01$) and touch display ($F=13.548$, $n=7$, $p<.01$) as well as retrieving content using the mobile techniques ($F=4.653$, $p<.05$).

Qualitative results

Overall, many participants feel that direct interaction at the display most closely matches their expectations from traditional PNAs (P2, P6, P15, P19).

“Touch input is most similar to writing a classified on paper”. (P2)

Yet, they feel that there are several privacy issues, e.g., when entering an email address publicly. Hence, numerous participants state that they would prefer using the mobile phone, since it is more private (P5, P7, P10, P11, P18, P19, P20).

“I don’t want the people standing behind me to know my email address”. (P18)

The participants identified several advantages of the digital PNA. Most important is the search functionality. Second, the filter and rank feature (e.g., by popularity/date) enables the users to search the classified ads easier and faster. Third, the digifieds can be enhanced with different designs, images, videos, and GoogleMaps. P13 also liked the fact that people could not simply remove or tear away classifieds like on traditional PNAs. Participants did however, mention the high ‘fun factor’ of using the phone-based techniques (P3, P6, P7, P9, P10, P14, P17, P18).

“Bumping and scanning were the biggest fun”. (P14)

Finally, we received feedback on how to enhance the system. This concerns the visualization (e.g., highlighting ‘new’ digifieds or providing a more casual layout (e.g., *“It [the layout] could be more old school.”* (P17)), but also ideas for new functionalities (speech input for creating digifieds, enable remote posting via web/email, linking Digifieds with Gmail).

9.5.6 Implications from the Lab

Based on the findings, the following implications were extracted for designing digital PNAs.

Design for different types of users and situations. We found two main aspects that influence user preferences for a certain interaction technique. First, users of different age groups or with diverse backgrounds and technical skills perform differently and prefer different techniques. Whereas young users like the mobile interaction, less mobile-savvy users favor the display or the PC. Second, qualitative feedback indicates that the preferred interaction technique often depends on the current situation of the user. Whereas participants incidentally passing by a PNA mainly want to use the display directly, the advantage of being able to prepare a digified at home allows more sophisticated designs to be created. Overall, there is no significant difference in length of text and duration when it comes to creating content. This indicates that several interaction techniques should be offered parallel.



Figure 9.13: Deployment of Digifieds in the Wild: Public Display at the market square (a) and in the public library (b).

Create robust, intuitive, and easy-to-use interaction techniques. The correlation analysis supports the assumption that, similarly to traditional PNAs, the acceptance and success of digital PNAs depends substantially on how easy and intuitive they are to use. While ad-hoc and occasional users are probably not willing to install a mobile app, people interested in more sophisticated designs are happy to do so. Yet, the study results in both cases indicate that if interaction techniques are difficult to understand or ‘flakey’ (e.g., requiring multiple attempts to use them successfully) users become frustrated and acceptance decreases significantly.

Preserve the user’s privacy. Our study revealed that some users are concerned about their privacy with regard to entering sensitive data, such as an email address, directly at the display. The reason for this is not just the fact that this data is shared with the PNA provider, but also that other people are able to see it. This indicates that methods for preserving privacy during interaction should be provided, e.g., phone-based techniques.

Overall, we observed that even though interacting directly with the display best resembles the functionality of traditional PNAs, there is considerable potential for using the mobile phone as an interaction technique. Content can be created on the go and in a manner that preserves privacy. Furthermore, our results indicate that there will be high uptake among younger people who perform very well in exchanging content with the display especially if they are familiar with touch-enabled devices and mobile Internet surfing.

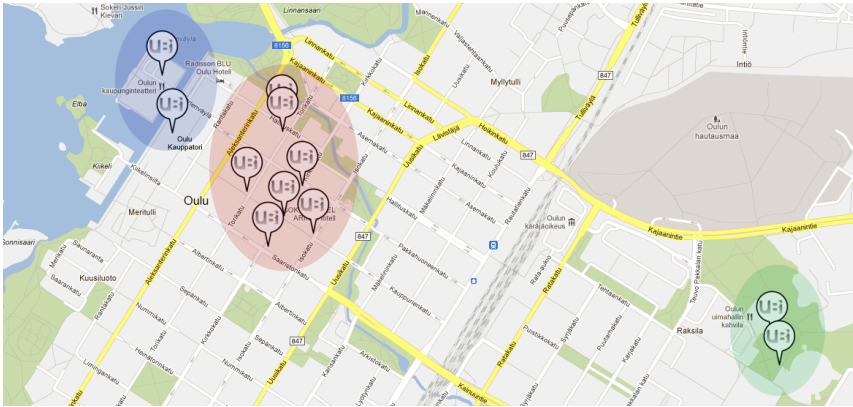


Figure 9.14: Display Groups: To preserve locality, we created a concept called *display groups*. This allows displays to be grouped based on arbitrary criteria. In this example, displays are grouped by location.

9.6 Deployment in the Wild

The lab study revealed the strength and weaknesses of different interaction techniques under controlled conditions. In order to observe natural usage of Digifieds and to ecologically confirm our lab findings, we deployed Digifieds in a real-world setting in Oulu, Finland. Oulu provides a city-wide network of six indoor and six outdoor displays [272], operated by the University of Oulu. The displays are deployed in public spaces, such as the library, the market square, the swimming hall, and the university. The display network was deployed in 2009, and provides an environment that allows public displays to be studied without the novelty effect that usually occurs when deploying new systems. Since the displays have been around for more than two years [272], many people are familiar with the displays and frequently use them. This is also reflected also by the fact that 73.3% of our participants stated to have used the displays before.

“I used this application [Digifieds] for the first time, but I have used the displays a few times before”. (Interviewee I13, technician)

Digifieds was deployed as a finalist of the UbiChallenge 2011 [270]. The Open Ubiquitous City Challenge (UbiChallenge) provides international researchers an opportunity to transfer their ideas from the lab into an urban environment. For the

deployment, a team of three researchers integrated the system with the existing infrastructure and thoroughly tested it before the public release on July 6, 2011³⁸. The application was intended to be available to the public until December 31, 2011 but is still active as of February 2013. To allow content to be posted on multiple displays but at the same time preserve the local character, we designed a concept called *display groups*. Each group consists of a number of displays with certain properties. Figure 9.14 depicts the concept. In this example, displays are grouped based on location – all displays at the market square (blue area), all displays in the pedestrian area (red area), and all displays at the sports center (green area). However, the concept is not limited to location and displays can be grouped based on arbitrary criteria (e.g., indoor/outdoor, size, orientation, etc.).

For the initial content, we collected content from traditional PNAs in the surroundings of the displays. We contacted the content providers and asked for their permission to feed their content into Digifieds. Furthermore, we asked the organizers of local events if they would be interested in advertising these events through our system. This way we had a considerable number of initial items in the week prior to the official release.

9.6.1 Evaluation

While the Digifieds service was available to the public (from July to December 2011) we performed a variety of evaluations, including observations, interviews, and a field trial. The observation and interviews were conducted on 11th and 14th of July 2011. The field trial ran over two weeks from 1st until 12th of August 2011. Additionally, user interaction was logged.

Observations and Interviews

We observed people for two reasons: first, we wanted to reveal usability issues; second we aimed at obtaining feedback on situations when people approached the display, on content they were interested in, and on how they used the system.

Observations were conducted around the displays in the public library and in the market place during two days. Overall, 60 people were observed. For the observations, a researcher would hide in a location close to the display and take notes on user behavior (Figure 9.15a). Additionally, gender and age was noted.

³⁸ Press Releases about the UbiChallenge 2011: <http://www.ubioulu.fi/en/UBI-challenge-documents>, last accessed February 13, 2013




Figure 9.15: (a) Observations at the public library: Two researchers observing people looking at or interacting with the display from a hidden location and took notes. (b) Interviews at the public library: As people turned away from the display after interaction, the interviewer approached them and requested an interview.

In parallel, we conducted semi-structured interviews. The interviewer would first observe any person approaching the display, regardless of whether they interacted or not (Figure 9.15b). As soon as the person turned away from the display, we asked whether the person would agree to an interview. For the interview, we used two different interview guidelines – one for people who interacted and one for people who did not. Besides following the guidelines, the interviewer would respond on interesting statements of the interviewees. In total, 29 of out 60 people agreed to an interview (21 knew the display before). The interviews were audio recorded for post-hoc analysis.

Field Trial

We conducted a field trial where people had to solve different tasks with the display and the mobile client. The trial was held in a public space (a university building), which helped us to create a realistic environment where people would be exposed to passersby watching as they performed the field trial. Hence, we were able to gather valuable feedback, especially with regard to privacy.

People were recruited from the street in front of a large department store and then sent to the university building. As they arrived, they got a brief introduction to the study and were asked to sign a consent form. Then, we provided them with the first part of a questionnaire where we asked them about their mobile phone usage (e.g., how often they use it, if it has a touch screen, if they use it to surf the web, and if they installed any third party apps) and whether they used the UbiDisplays



Figure 9.16: Field Trial: We asked people to look for and retrieve digifieds by sending them via email or to the phone (a) as well as to create digifieds (b), e.g., “You found an umbrella and want to place it on the display.”

before. Then we asked about their use of traditional PNAs (how many they knew, how often they use them, which type of information they usually look for as well as retrieve, and which information they post).

Afterwards, we asked them to test our application. For posting content, we gave them the task to place a digified about a phone they had (virtually) found into the ‘Lost & Found’ category. After that we wanted them to look for a bike and send the classified advertising the bike to their email address (Figure 9.16b). After finishing the tasks, they were asked to complete an SUS questionnaire [26].

Next, we asked them to test the mobile application we had pre-installed on a Samsung Galaxy phone. Therefore, they first had to create a digified to sell an umbrella we provided them (Figure 9.16a) and place it in the category ‘Sales’. After that we wanted them to search for and retrieve two digifieds, each time using one of the two mobile interaction techniques. First, they should retrieve the digified on the ‘Beach Tennis Cup 2011’ event using the QR code. Second, we asked them to transfer the ‘Sky Diving Oulu’ ad to the mobile phone by using the provided code. Then, they again had to fill out a standard SUS questionnaire.

After finishing the tasks, we asked them to fill out the last part of the questionnaire. We were first interested in which kind of information they would like to find or retrieve from Digifieds and which information they would post on Digifieds. Then we wanted to know what the users’ opinion was on entering (private) data in public space, on the display client, and on the mobile phone. Finally, they were asked to rate the different features and provide qualitative feedback.




Figure 9.17: Distribution of Digifieds Usage Over the Day: people mainly used Digifieds in the lunch break and in the after hours.

Log Files

Between 7th of July and 31st of August 2011, user interactions with the Digifieds platform were logged. We collected data about how often, when, and where Digifieds was launched. Furthermore, we logged which content users were interested in (based on when they opened the detailed view of it). Finally, we analyzed the posted items with regard to content and layout.

9.6.2 Findings from the Real World

During the 2-month evaluation period, we had a total of 1126 launches of the Digifieds application on the displays (avg. 125 per week). Figure 9.17 shows the distribution over the day. People used the application mainly during lunch break and in the afternoon, probably since they were shopping or on the way home from work. The analysis of the log file shows that 900 users looked at the content in more detail (based on the number of times the detail view of a digified was opened). In the following we provide in-depth findings with regard to content, privacy, and interaction techniques.

Content

Knowledge about the preferred content on PNAs provides useful hints with regard to which content should be presented where, when, and to whom (given that it is possible to identify the user).

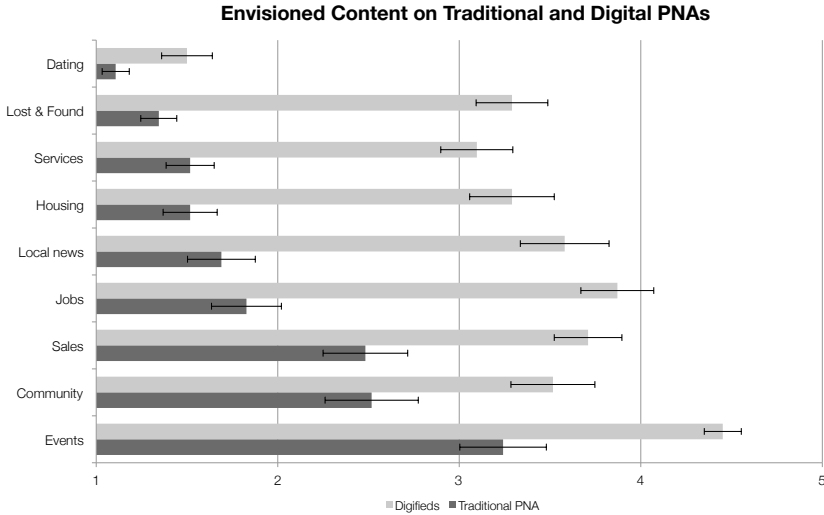


Figure 9.18: Preferred Content: Which information are you interested in on traditional / digital PNAs? (5-Point Likert scale, 1=not interested at all, 5=very interested)

First, we were interested in *comparing content on traditional and digital PNAs*. During the field trial, we asked users in a questionnaire which type of information they were usually looking for and retrieving from traditional PNAs. Figure 9.18 (dark bars) shows that most users are interested in ‘Events’, ‘Community-related Information’, and ‘Sales’, or in other words, content that is mostly locally relevant. When comparing these findings to digital PNAs, we found that people seem to also expect mainly event-related information but also other locally relevant content, such as ‘Community-based information’, ‘Sales’, and ‘Local News’. Also ‘Jobs’, ‘Services’, and the ‘Lost & Found’ categories are promising (Figure 9.18, light bars). Statements of the interviewees support these findings:

“I was expecting to find information on events and news from neighboring areas”. (I10, speech therapist)

“I wanted to know what was going on in the city at the moment.”
(I13, technician)




Figure 9.19: Comparison of content between traditional PNAs (based on data from previous work) and Digifieds: Sales and Events were the most popular content (in percent).

Several tourists were looking for both events as well as for directions to the city’s major sites (even though this functionality was not provided). Students often stated to be interested in housing, especially at the beginning of a new term.

Second, we were interested in the *content posted on Digifieds*. When we analyzed our log files, we found that between 7th of July and 31st of August, 49 classifieds were posted in eight categories. ‘Sales’ is the most popular category (23), followed by ‘Jobs’ (10) and ‘Events’ (8). Two posts with offensive content were removed. To compare this to traditional PNAs, we analyzed photo logs from a study that included 22 traditional PNAs and 300 pieces of content [11]. Figure 9.19 shows a comparison of the content. We found that in both cases, the PNAs contained mainly posts on sales as well as events and jobs. Housing was not as popular on Digifieds, which might be a result of the fact that the evaluation was conducted during the main holiday period. As Oulu is a student city, interest for housing will probably grow shortly before the new term starts.

Third, we analyzed the log files for assessing the *viewers’ interest*. To do so, we calculated for how many posts of the respective category, viewers opened the detailed view of the content, which we believe to be a good indicator for interest (Figure 9.20). We found that ‘Sales’ and ‘Events’ were most popular.




Figure 9.20: Viewer Interest: During our study, viewers were interested in detailed information on sales and in events.

Design Recommendations

Like on traditional displays, the envisioned content and the actual content posted on public displays depend on two factors: the location of the display and the viewer. Both often have a relationship (e.g., a tourist is more likely to be found in the vicinity of a display in the city center than in the swimming hall, whereas displays on the university campus are likely to attract mainly students). In many cases, content providers as well as viewers share the same place (e.g., one student offering his apartment and another student being interested in it). In cases where this is different (e.g., if an event organizer wants to advertise a rock festival on the market square to students), means have to be provided to allow content to be distributed to the intended places. One solution is our *area concept*, which provides a mechanism to target content towards a location. Future versions may allow distributing content not only based on location but also other types of context, such as time of day, communities, or demographics of the surrounding area (e.g., income or population density).

Privacy

In our interviews, several people expressed concerns that their privacy might be affected either if they leave personal information (e.g., email address) on a publicly available display that it could be found by everybody, or if people standing behind them watched as they entered this information.




Figure 9.21: Privacy Concerns with Regard to Content: Privacy is perceived as being strongly affected if inputting sensitive information on the display. Mobile phones can potentially overcome this issue.




Figure 9.22: Privacy Concerns with Regard to Technology: Concerns are similar for traditional / digital PNAs; phones and PCs preserve privacy better.

“I don’t want that people notice me interacting with erotic show classifieds on the public display.” (I14, IT worker)

“A phone can be used in private, without outside disturbance.” (I18, student)

Subsequently we aimed at further investigating this issue in the field study. In our questionnaire (31 participants) we tried to find out which information is subject to privacy concerns and how this relates to inputting information on traditional PNAs. Additionally, we were interested in whether the mobile client is able to cope with these concerns.

First, from analyzing the answers in the questionnaires, we found that inputting personal information (e.g., an email address) is perceived as being more privacy affecting than inputting more general information such as the title or the content of a digitified. This is true both for inputting information on the display (= publicly)

and on the mobile phone (= privately) (Figure 9.21). However, when we compared the perceived impact on privacy for emails, this was found to be significantly stronger on the public display than on the mobile phone ($t=-5.82$, $df=30$, $p<.001$).

Second, we were interested in whether there is a difference in the perceived impact on privacy between traditional and digital PNAs (Figure 9.22). We found that there is no significant difference ($t=1.2$, $df=30$, n.s.). The impact on privacy compared to digital PNAs is perceived to be significantly smaller compared to the mobile phone ($t=-6.69$, $df=30$, $p<.001$) and to the PC ($t=-6.83$, $df=30$, $p<.001$).

Design Recommendations

Our findings indicate that the users' privacy concerns – with regard to both other people lurking as well as private information being publicly available – have to be taken into account, as this is indeed being perceived as a major issue. Surprisingly, many people still use traditional PNAs where neither of these issues is tackled. This contradiction might be explained by the fact that reports in the media about the abuse and the loss of lots of personal information (e.g., credit card data [122]) lead to them changing their minds. Systems could address these issues as demonstrated by the Digifieds platform. First, the mobile client provides an alternative to enter sensitive data (e.g., contact data) in private. Second, in cases where contact with another user needs to be established, the user should be able to decide when to reveal their contact information, for example, by using a form.

Interaction Techniques

In the field trial, we had users compare both the display client and the mobile client with regard to usability. Some users reported that they consider the display client to sometimes be unresponsive (we believe that this is due to the capacitive display) whereas the phone client works smoothly.

Users see the option to take pictures as a strength of the mobile application (this is not possible with the display client as there is no camera integrated), and that it is possible to generate content on-the-go. Many of our trial participants were not familiar with QR codes. However, most were able to quickly figure out how this functionality worked and later stated that they like the concept (e.g., P5, P16, P20, P22). Furthermore, they like the opportunity to take multiple digifieds with them at the same time by means of the 'Digifieds basket' feature.

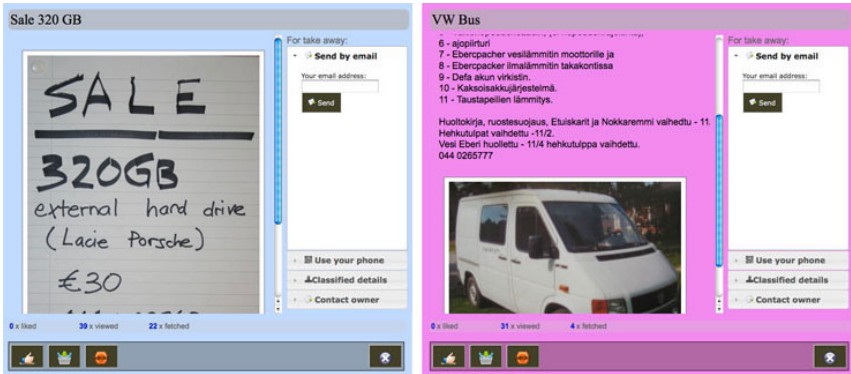


Figure 9.23: Content Posted with Digifieds: (1) Some people take pictures of traditional classified ads they have prepared with the Digifieds mobile client and post them on the display (left). (2) Another practices is to enter the text directly on the mobile client and then enhance it by adding a picture taken with the phone (right).

In order to quantitatively assess the usability, the users had to fill out an SUS both for the mobile client as well as for the display client during the field trial. A Pearson correlation analysis revealed, that people who rate the usability of the mobile client high, also rate the usability of the display client to be high ($r=.375$, $df=29$, $p<.05$), which is an indicator for that if people are familiar with either of the used technology they can easily adopt the other one. This is also supported by interviews where older people who do not use smart phones feel that the display is more suitable for younger people.

“I think this is for young people – I should bring my grandchildren.” (I25, pensioner)

With regard to practices, we found that people use different approaches to input information to the system. Whereas most people use the display client to type in text, others take a photo of their handwritten page using the mobile application (Figure 9.23). This mostly happens if people prepare a post with a sophisticated design (e.g., semi-professionally layout on the PC) or if they are in a hurry, and do not have enough time to type the text manually.

Design Recommendations

Similar to smart phones, public displays are perceived as a ‘new technology’. Interestingly, people who are familiar with smart phones have no problems at all using the public display. In contrast, though many older people show interest in the beginning, they state in the interviews that they see the benefit more for the younger generation. Consequently, one major challenge can be seen in finding ways to also entice the older generation to use the new technology. Furthermore, providing multiple ways of publishing content (typing, taking a picture, etc.) is crucial to support both ad-hoc posters as well as semi-professional or even professional content providers.

Limitations

The study ran during the summer holiday season. It can be assumed that a higher percentage of tourists might have interacted with the displays (and Digifieds) as in off-season times. Furthermore, all interviews and field trials were conducted during a weekday. Hence the number of tourists and older people might have been above average with regard to the users of the displays.

9.6.3 Lessons Learned from the Real World

From the observations, interviews, and field trial, we learned the following:

Communicating the value of the displays. As Müller et al. [255] report, displays in the wild are largely ignored, as people do not expect to find interesting content. Oulu provides an environment, in which displays have been available for several years and people have discovered that they are interactive. Only at this point it makes sense to investigate interaction techniques. During our observations and interviews, we found that people come back (44 out of 60 in our study) as they are aware of the display’s value. This strongly suggests that ‘display blindness’ can be overcome with services that really satisfy the people’s needs.

Managing user expectations. During our observations and interviews, we found that people interacted with Digifieds similarly to using a web browser. They expected, for example, that a login was required before being able to post and they were expecting a search functionality, as well as features

such as Google Maps integration. This indicates that one could draw from people's skills with other technologies when designing public displays. Hence, when using a UI design similar to that of a mobile phone or a web browser, people are likely to use interaction techniques that they know from these platforms and consequently, people's interactions could be steered.

Providing ubiquitous access. In contrast to the design ideas of locality and community [268], as well as location-based messaging [327], it was apparent that younger users have a clear expectation that access is ubiquitous and that displays are just one way of accessing content / information (others include WWW / mobile devices). At the same time, we see an opportunity to use the location as means of filtering information (e.g., displaying only locally relevant information) and easing access by providing multiple channels. Nevertheless, information is not perceived as if it is tied to a certain place and it is obvious for users that 'content is in the net'.

9.7 Case Study: Secure Authentication

In the remainder of this chapter, we present a case study that looks into secure authentication on public displays. As reported earlier, both findings from the lab as well as from the real-world deployment of Digifieds showed that strong concerns exist among the users of public display applications when it comes to inputting personal, potentially sensitive information. Such concerns are likely to grow in the future as more and more third party applications emerge and it may neither be obvious for the user who operates the display application nor who can access the data.

The mobile phone was shown to be able to partially cope with these concerns. However, installing a mobile client prior to being able to securely exchange data with a display may be a burden to many users and potentially exclude them if the software was not available for their phone. One alternative solution that allows users to authenticate in a shoulder surfing resistant and touch-free manner is gaze-based authentication. In one of the first works in this area, Kumar et al. present an authentication system that mitigates the issues of shoulder surfing by using gaze to enter a text-based password on an on-screen keyboard [194]. They found that gaze-based password entry requires marginal additional time and that the error rates are similar to those of using a keyboard. Overall, the users in their study preferred gaze to traditional keyboard-based password entry.

One problem with text-based passwords is that they are more vulnerable to guessing attacks, due to the predictability of passwords, particularly for weak user-chosen ones. Consequently, researchers are investigating graphical passwords as an alternative means of user authentication [37; 356]. Cued-recall graphical passwords (locimetric passwords) have considerable advantages over traditional approaches, such as text passwords, as they leverage the vast capacity and capabilities of the human visual memory system [86; 99]. In addition to improved memorability and thus usability [243], graphical passwords promise increased resistance to guessing attacks, due to the potentially larger theoretical password space. Despite these advantages, graphical passwords that rely on physical interaction with the authentication system are still susceptible to shoulder surfing attacks [110]. Particularly in public spaces, e.g., in front of an ATM, direct observation techniques such as video cameras or fake keypads can easily be used to eavesdrop and steal passwords or PINs.

One solution to increase the security of graphical passwords is to use authentication schemes that do not require any physical contact with the system. Amongst the methods investigated in the past, the human gaze is particularly promising for implementing such schemes [142]. By its very nature, gazing does not require any physical contact and therefore, potentially works over greater distances. In addition, the human eye moves rapidly, which makes eavesdropping gaze-based passwords more difficult than touch-based input.

A key challenge in user authentication generally, and in graphical schemes in particular, is to define secure passwords. Previous research has shown that such schemes lead to hotspots, i.e. areas of the image that are more likely to be selected by users as password points. A password point is defined as a single fixation that is detected by the system to be part of the graphical password. These hotspots render such schemes more susceptible to dictionary attacks [381]. The only viable solution so far has been to select single password points across a sequence of several images [69].

In this section, an alternative gaze-based authentication scheme is presented that supports users in selecting secure gaze-based graphical passwords. To tackle the problem of hotspots, our scheme uses a computational model of visual attention – also known as saliency maps – to mask out those areas of the image most likely to attract visual attention (Figure 9.24). We show that this approach significantly increases the security of gaze-based cued-recall graphical passwords.




Figure 9.24: Sample images used in the study without (a) and with (b) saliency mask.

9.7.1 Concept

Visual attention is constantly attracted by different parts of the visual scene. Bottom-up computational models of visual saliency aim at estimating the parts of a visual scene that are most likely to attract visual attention [157]. Given an input image or video, these models compute a so-called saliency map that topographically encodes for saliency at every location in the image. Visual saliency models are shown to predict human visual performance for a number of psychophysical tasks and have a large number of applications (see [114] for an extensive review). For example, in computer vision, saliency models are used for automated target detection in natural scenes, smart image compression, fast guidance of object recognition systems, or high-level scene analysis [158].

The key concept underlying this work is, that by encouraging users to select password points that do not fall inside salient regions of the image, the security of gaze-based graphical passwords can be increased significantly. This is similar to the characteristics commonly required for text-based passwords, such as a minimum number of different alphanumeric or special characters.

Figure 9.24a shows one of the normal login images used in this work. Because the dog's eyes and nose are most likely to attract the user's visual attention – as predicted by the visual saliency model – these parts are masked out in Figure 9.24b. In a real-world authentication system, such masked images would be shown to the user when selecting the initial password. During operation, such as for authentication at a public terminal, the same image but without a mask would be used instead.




Figure 9.25: Standard 10-digit keypad used in the studies for gaze-based PIN entry.

9.7.2 Apparatus

We designed and implemented a gaze-based authentication system consisting of a remote eye tracker and a custom login interface. The system allows users to login using their gaze with two types of graphical passwords. To record gaze data, we used a Tobii X120 remote eye tracker. The X120 has two integrated infrared cameras and was configured to track gaze at a sampling frequency of 60Hz. The login interface was shown to the participants on a 19" computer screen with a resolution of 1600x1200 pixels.

The login interface was implemented in Java and obtained gaze data from the eye tracker via TCP/IP. In addition, it recorded all events triggered by the user interacting with the system, such as the detection of password points or successful and failed login attempts. The interface included a standard 9-point calibration routine to adapt the tracking system to each user and a validation routine to assess the calibration quality. The testing routine involved the user looking at each of the calibration points in sequence until fixation was detected by the system. Calibration quality was calculated for each participant as the mean Euclidean distance between all calibration points and the detected fixation points.

In the user study, depending on the type of graphical password, the interface showed two different screens:

- PIN: the login screen showed a grid of 10 tiles resembling a standard 10-digit keypad (Figure 9.25). In order to perform a login attempt, the participant fixated at four of these tiles in sequence. Sequences of fixations to correct tiles in the correct order resulted in a successful login attempt; all other attempts were considered as failed.

- Image: the login screen showed a full-screen image. In contrast to PINs, the screen was not visually discretized by a grid.

We used a dwell-time-based method to select the password points of the graphical passwords, because this approach is more natural and less error-prone than, for example, double blinks [163]. To select a password point, the participant had to fixate on a certain point in the image. The software continuously analyzed the dispersion of the gaze points within a time window of 1.5 seconds. A fixation was detected by the system if at least 70% of the gaze points in this time window were inside a circular area with a radius of 1.7 degree of visual angle (about 75 pixels, the fixation detection threshold th_f).

9.7.3 Evaluation

The experiment consisted of a pre-study and a main study. In the pre-study, we asked one group of participants to define three different types of gaze-based graphical passwords (PIN, image with a saliency mask, image without a saliency mask). The goal of the main study was to assess the security of these passwords by asking a second group of participants to try to attack and guess these passwords by analyzing close-up videos of the eyes recorded from the pre-study participants. The main study was designed as a randomized, controlled user study (repeated measures design).

For the main study, the binary output of all login attempts per image (successful: 1, unsuccessful: 0) was the dependent variable. The dependent variable was measured under three different conditions of the within-participants factor ‘graphical password’: PIN-based, image-based without saliency mask, and image-based with saliency mask.

The results of our study demonstrate that image-based graphical passwords are significantly more secure than PIN-based passwords, both in an actual attack and in terms of participant perception. Using computational models of visual attention to mask the most salient areas of the images does significantly increase security, compared to the standard image-based approach. In combination with the much larger theoretical password space, these results make saliency masks a promising means of increasing the security of gaze-based graphical passwords. For a detailed description of the study design and the results we refer to [55].

9.7.4 Discussion

While image-based graphical passwords were perceived as significantly more secure than PIN-based passwords, the participants in our study rated the usability lower. Participants preferred image-based passwords for public terminals, while PIN-based passwords were preferred for mobile devices such as laptops and mobile phones. These responses may have been caused by the fact that participants could not imagine mobile devices equipped with robust eye trackers in the near future. While currently, application domains are indeed mostly limited to ATMs or similar stationary systems, the advent of mobile eye trackers will pave the way for gaze-based authentication on smaller and thus more mobile devices [56].

A proper analysis of password memorability requires a long-term study and was therefore beyond the scope of this work. However, when we asked two pre-study participants to log in with their image-based passwords two days later (whom we did not ask to remember their passwords in the first place) they correctly remembered 14 out of 40 passwords (five images with and nine without a saliency mask). 13 of these image-based passwords were remembered by the pre-study participant who had selected the password points in the direction of reading, i.e., from left to right. While using such strategies seems to improve password memorability, this may come at the cost of weaker security. We plan to investigate this trade-off between memorability and security in more detail in future work and particularly how password memorability can be improved without compromising security. In terms of security, it will also be interesting to see how saliency masks compare to other approaches, such as selecting password points on a sequence of images [69].

Finally, the study also reveals some of the issues researchers may face in the real-world implementation of gaze-based graphical passwords. Participants in our pre-study reported having used visual strategies for selecting their password points in the images. Two main-study participants noticed and exploited this behavior by specifically looking for characteristic eye movement sequences such as in a vertical or horizontal direction. This suggests that, in addition to the saliency masks presented here, measures need to be taken to prevent users from choosing closely related password points (similar to preventing PINs like ‘1111’ or ‘1234’). Additional user studies are required to investigate whether users should be allowed to choose their own graphical passwords (and potentially the images as well), or whether both should be provided by the authentication system during registration. In the latter case, it would be useful to identify what qualifies as ‘good’ passwords and images.

9.8 Summary

In the context of our central use case – digital classified ads – this chapter investigates suitable interaction techniques and different factors that impact a user’s motivation to interact with public displays. With regard to interaction techniques, we found that providing robust and easy-to-use interaction techniques is crucial for the uptake of display applications. The mobile phone and direct touch at the display turned out to be equally suitable for creating and exchanging content with the display. The strength of the mobile phone is that content can not only be created on the go but also in a privacy preserving way. Direct touch at the display does not require an additional device or installation of software but enables to walk up and post content ad hoc. Our findings suggest that public displays benefit from users’ prior experiences with technologies such as the Internet on mobile phones and that novel interaction techniques could draw from this knowledge. Yet, particular attention needs to be directed towards non-digital natives to not only make interaction techniques understandable but also to convince them to try and overcome any reservations.

With regard to motivation, we found creating a comprehensive model challenging. Previous models like the Technology Acceptance Model or Michelis’s building blocks provide a good basis for explaining the user’s motivation to interact with public displays. However, external influences, such as cultural aspects (e.g., level of privacy awareness in different countries), demographics (young users may be more open to use new technologies), or the type of application (informative vs. playful) may have a strong impact. The results of the Digifieds deployment and our case study yield, that in the future, the following aspects may be essential when the user decides whether or not to interact with a public display:

- **Interest:** Ever more displays try to catch the user’s every look. As more and more displays are networked and sensors are deployed, we see a major opportunity in tailoring the content towards the interest of the user. Müller et al. show that more than half of the users think that interesting content would make them look at public displays more [255]. This is supported by our findings that users are indeed interested and look mainly at locally relevant content.
- **Privacy:** Displays may affect the users’ privacy as they are interacting in public space. This applies to both cases where people implicitly or explicitly interact with the display. In cases where a display (implicitly) adapts its content to the user, it may indirectly reveal information about the user

(e.g., showing an advertisement for a Rolex watch may identify the person standing in front of the display to be rich). In a similar way, information about users explicitly interacting with the display may reveal to bystanders the content they are interested in or also more sensitive information such as an email address as they enter it on the display. Our findings suggest that devices such as mobile phones can be used to overcome such issues.

- **Security:** In the future, displays other than ATMs may enable access to sensitive, personal data, such as contact information, email, etc. In these cases, people may refrain from using a system if no secure authentication mechanisms are provided. We present an approach that can overcome issues such as shoulder surfing and significantly increase security.

Concluding, this part of the thesis has presented a user-centered design space for pervasive advertising on public displays. It identified three core challenges of promoting user interaction: attracting attention, communicating interactivity, and enticing interaction. Based on a number of prototypes that were evaluated both in the lab and in the field, we address each of these challenges and presented possible solutions. The design recommendations and lessons learned can help to inform the design of future advertising applications for public displays.

V

CONCLUSION

Chapter 10

Conclusion

10.1 Summary of Research Contributions

Public displays and pervasive advertising are still young research fields. Looking more closely at public display research, the literature survey presented in Section 2.8 shows that a lot of different application domains have been identified, ranging from artistic installations via collaborative settings to entertainment and advertising. Yet, prior work often explores particular aspects, mainly in descriptive studies. In this way, isolated spots in the design space are assessed. Such an approach does not allow findings to be compared and hence a comprehensive picture of the design space to be drawn. To set the scene, this thesis consequently aimed at describing a comprehensive design space for pervasive advertising, which is the prevailing application category for public displays.

We believe that advertising is likely to remain the major business model behind public displays in the future. At the same time it is clear that this model will succeed only if a benefit is provided for society. This can be achieved by presenting or providing access to relevant information at the right time, by involving the user, or by simply staying calm in the background as the user is not interested. The presented design space aims to support developers to shape this vision. In addition to this overall goal, smaller projects explored and tackled a set of specific challenges, mainly from an HCI perspective. The approaches we developed and evaluated led to a number of contributions described in the following. An overview based on the research questions is provided in Table 10.1.

Table 10.1: Overview of Contributions to Research Questions.

No.	Research Question
I. The Provider View	
R1	<p>How do systems need to be designed to cater to the providers' needs?</p> <p>In a first step we investigated traditional public displays. Stakeholders were identified, content was analyzed, and the space around the display was observed. By carefully looking into processes and practices, we were able to understand their interplay and we were able to extract concrete design implications that could help to design and develop public display applications in a way that addresses the needs of the providers (Chapter 4).</p>
R2	<p>How does the audience behave in front of public displays?</p> <p>Based on findings from a lab study with displays of different shape, a set of design guidelines was derived. The guidelines, first, explain how shape, content, placement of the display, and interactivity cue effect audience behavior. Second, we provided suggestions as to how to exploit or avoid these effects. For example, a cylindrical shape leads to that users are in constant motion and could hence be used in highly frequented areas to guide users towards non-crowded areas around the display (Chapter 5).</p>
R3	<p>How can the user be targeted in a meaningful, privacy-preserving way?</p> <p>As opposed to current targeting approaches that focus solely on information about user behavior or personal interests, we contributed an approach that blends both information to draw a more comprehensive picture of the user. Furthermore, we looked into how multiple users could be targeted at the same time. We provided the user with control over the content they are targeted with, making the process transparent and realized the profiling in a privacy-preserving way that does not require personal user data (Chapter 6).</p>
II. The Audience View	
R4	<p>How can the attention of the user be attracted?</p> <p>We contributed an approach that exploits the user's gaze information. We were able to show that adapting content based on gaze data is suitable for increasing attention and affecting attitude, independent of the user's interest (Chapter 7). This finding could be particularly valuable as eye tracking hits the mass market (e.g., Tobii Rex).</p>
R5	<p>How can interactivity be conveyed in an intuitive way?</p> <p>As people do currently not expect interactive public displays, communicating interactivity is a central challenge to make public displays more attractive. State-of-the-art interactivity cues include calls-to-action and attract sequences. We contributed a novel approach that conveys interactivity through the representation of the user. We investigated the effect of the degree of abstraction (lab) and the impact of this cue on audience behavior (field). Concrete design guidelines were provided (Chapter 8).</p>
R6	<p>How can users be enticed to interact?</p> <p>The final step to make the user interact with public displays is to provide easy-to-use interaction techniques while keeping the motivation of the user high. Besides introducing the design and development of the Digifieds platform, we contributed a comparison of different interaction techniques with display. As a result, guidelines were provided, which techniques should be provided in which situation and for which users. With regard to motivation we showed that creating value for the user is key (Chapter 9).</p>

10.1.1 Shared Use of Public Displays

For advertising on public displays to become more meaningful and provide a benefit for society, it is clear that content needs to be changed. Content has to be relevant and engaging for the viewer and hence a mixture of advertisements and non-advertising content is required. Having understood this, display providers already started to blend advertising with non-advertising content (e.g., information screens showing news, weather, cartoons, temporarily interleaved with advertisements). As displays are becoming interactive, many opportunities emerge for new types of content. Creating such content requires new expertise and graphics designers will soon be joined by filmmakers, interaction designers, game designers, composers, and programmers. To exploit these new types of content, new forms of distribution, such as application stores, may emerge that allow both display providers as well as passersby to access this novel content [74]. As a result, display providers may not be limited anymore to large outdoor advertisers but owners of individual displays (e.g., retailers, supermarkets, universities, municipalities) may enter the scene, ultimately creating a global scale public display network. This trend may be supported by the fact that prices for display hardware are dropping since years.

To understand the fundamental challenges of such open display networks we presented findings from an ethnographic study, which assessed factors that inhibit or promote the shared use of public display space in the analog world. Based on the findings and implications we developed and deployed an application that supports the shared use of public displays and evaluated it in the real world [10].

- **Design for specific uses of public displays:** Public displays are used to fulfill different purposes. As a result, different designs and practices emerged that support maintaining and exchanging content with the display. We identified different notice board styles, posting form factors, and content that could serve as a starting point for designers and developers. Furthermore, we suggest to design for multi-display interaction, which for example allows a story to be created along the path of the user.
- **Respect the neighborhood focus of public displays:** The vast majority of content on public displays relates to the local area or community using the space. As networked public displays are technically not limited with regard to the distribution, posting procedures need to be designed that support locality. For example, Digifieds supports content to be distributed to so-called display groups, that are based on arbitrary criteria, such as location, demographics, community, etc.

- **Profile of the display and agenda of the space owners:** Space owners use displays as decoration or a way to express or support the image of a place, to increase the importance of the venue to a community, and to disseminate information. To do so, they rely in many cases on third-party posters, but often have difficulties to articulate their agenda. Since the profile of a board often emerges from the interplay of different stakeholders' interests, mechanisms need to be provided that allow space owners to actively promote their priorities (e.g., an administration interface to remove undesired content).
- **Design for flexibility of input:** We found an impressive variety of content during our observations, ranging from hand-written notes to professionally designed posters. As a result we suggest to support three different types of posters: ad-hoc posters who coincidentally approach the board, sophisticated posters who prepare content in advance, and professional posters. Digifieds addresses this by supporting content to be created directly at the display, on the mobile phone, or on the PC at home.
- **Low overhead to post:** Analog public notice areas are successful due to their low entry barrier. Consequently we believe that for digital systems the success strongly depends on the ease of use of content creation mechanisms. In a lab study, we tested different interaction techniques, iteratively improved them and employed them for the use in a real-world setting in Oulu, Finland. In a nutshell, we found that direct touch at the display and using a mobile phone client to create content on-the-go were most promising.
- **Retrieving information:** For certain types of content, such as classifieds or promotions, the opportunity to retrieve the information is crucial for its success. We found the mobile phone to provide many suitable opportunities, e.g., taking a picture, scanning a QR code, or entering a short, alphanumerical code. Interestingly, more information but an expiration date or the name and address of the provider is associated with the retrieval of content. For example, missing tear-aways indicate high interest. This is reflected in Digifieds by providing feedback on how many people looked at the detailed information, how many liked the content, and how often the content was retrieved.

10.1.2 Making the User Interact

All interactive public displays face the problem of initially making the user interact. As knowledge builds up among users and enough value is provided for them to come back, this problem becomes less important. However, particularly for new installations or applications, such as interactive games, understanding how to make the user initially interact is crucial. Previous work often focused on sub-problems, such as how to attract attention or how to provide usable interaction techniques. The contribution of this thesis is to draw a comprehensive picture of the entire process and suggest solutions to each of the stages of this process. The process is anchored in the Audience Funnel.

Attracting attention is the fundamental prerequisite. Myriads of displays striving for the attention of the user and the effect of display blindness make it difficult for providers of display and content to compete. We tackle this challenge through an approach that analyses a user's gaze information and adapts the content on the screen accordingly. We were able to show a significant increase of attention towards the adapted content compared to randomly chosen content. Additionally, the increase in number of fixations has been shown in prior work to be a strong indicator for the peripheral route of the Elaboration Likelihood Model. This is interesting for advertisers because an effect on the attitude of the user towards the content he is attentive to can be assumed.

When it comes to conveying the interactivity of a public display a set of strategies has been presented and is commercially used (e.g., attract sequences or calls-to-action). However, such interactivity cues require consciously processing complex information and are often culture and language dependent. As a solution we propose to use the representation of the user. Research in psychology suggests that using feedback to the passerby's incidental movements is promising, as humans are very efficient at recognizing human motion [80] as well as their own mirror image [241]. The results of our lab study show that (1) mirrored user silhouettes and images are more effective than avatar-like representations and (2) that it takes time to notice the interactivity (approximately 1.2 s). In the field study, three displays were installed during three weeks in shop windows, and data about 502 interaction sessions were collected. Our observations show that significantly more passersby interact when immediately showing the mirrored user image (+90%) or silhouette (+47%) compared to a traditional attract sequence with call-to-action. As most people recognize themselves on the display rather than someone else, displays should be positioned so that people can see themselves well when passing by.

As a final step we looked at how people can be enticed to interact. We identified two requirements from the user perspective in order to start interacting: understanding how the interaction works and motivation. In order to investigate suitable interaction techniques that would be easily understandable and usable, we tested different technologies and techniques with the Digifieds prototype in a lab study, including direct touch at the display and mobile phone-based interaction using QR codes and a display/phone touch feature. We found that according to an SUS questionnaire, usability of direct touch at the display is ranked significantly better than using mobile phone techniques. There are no significant differences for the time required to create posts and for the length of the posts. A correlation analysis using Pearson's 'r' revealed that using smartphones, having unlimited data access, and age have a strong effect on the amount of written text, interaction times, and perceived usability of the different interaction techniques. Some users state to prefer using the mobile phone in public settings as they better preserve privacy and allow content to be created on-the-go. Findings from the real world indicate that both techniques should be offered in parallel.

With regard to motivation, we first review existing models, such as the Technology Acceptance Model [84; 85] or Michelis's building blocks [236]. Our research reveals additional factors that impact on the user's motivation: relevance of content, privacy, and security. With regard to content, users expect locally relevant content both on analog as well as on traditional displays, according to a survey. Events, job offers, sales, and community-related information are most popular. This is also reflected by the actually posted content on Digifieds during the deployment. Consequently, public displays should support relevant content. For example, the ability of Digifieds to group displays based on location allows the content to be targeted towards a place. We aim to extend the concept in the future to support further types of context.

To further investigate privacy issues we conducted a survey in the context of a field study. With regard to content, people perceive inputting personal information in public as more privacy affecting than more general information. When comparing the input of personal information on public displays and on mobile phones, the impact on privacy is perceived significantly stronger on the display. We also compared traditional to digital public notice areas but did not find any significant difference. Hence, public displays should be designed to protect the user's privacy. For Digifieds, we do this in two ways. First, personal information can be input via the mobile phone in order not to reveal them to bystanders. Second, contact to the creator of a digified is established through a form that leaves the decision whether or not to reveal his email address to the user.

Finally, we investigated how secure authentication could be realized on public displays. We present a gaze-based approach where the user can define a password through a number of password points within an image. This hands-free technique can cope with current issues such as shoulder surfing in public spaces. As an extension to the approach, we showed that masking out areas in the image that are likely to be selected as passwords can further increase the security.

10.1.3 Audience Behavior Around Public Displays

This thesis makes a set of contributions towards understanding audience behavior around public displays. This is of particular importance for advertisers, as they want to deploy public displays that attract and eventually motivate passersby to enter their shops or that have a positive impact on brand perception. We investigated audience behavior in two contexts: first we compare audience behavior in front of planar and cylindrical displays. Second, we explore how inadvertent interaction impacts on audience behavior.

As displays of arbitrary shapes can be created, there is an inherent need to understand how people's behavior differs in front of planar and non-planar screens. This knowledge can then be used as a basis to develop applications that exploit the properties of the new format and investigate more elaborate topics such as multi-user interaction. As an example we chose cylindrical displays, which are, in their traditional form, still a popular advertising medium in the urban landscape. We compared a digital advertising column developed by Fraunhofer FIRST in Berlin against a planar display that would cover the same size of real estate. We created a museum-like situation in a lab which allowed precise measurements to be performed while still observing users in a realistic scenario.

We show that people walk significantly longer distances with the cylindrical display, they spend significantly more time walking (55.1%), and they cover significantly more locations in the display vicinity. For the cylindrical display, their shoulder position is parallel only 22% of the time as they walk (flat display: 46%) and 69% of the time as they are standing (flat display: 81%). With regard to time we found that users spend almost twice as much time with the flat display than with the column (2:53 m vs. 1:38 s), the difference being significant ($p < 0.002$). Furthermore, the results show that for the flat display there is a narrow area in front of the display where users get themselves in a frontal position (the sweet spot), about 1.5 m away from the display.

Based on these findings we derived the following recommendations.

- **Design for walking interaction:** Cylindrical displays provide a way to support real motion. Hence, content should be designed for walking by raising the attention of the viewer through implicit interaction and then encouraging them to move on. Also content that moves with the user around the column may be suitable. Flat displays are better suited for waiting situations, longer dwell times, and more complex content.
- **Place columns in the way of users:** Cylindrical displays are suitable to be used while moving and thus allow users to avoid bumping into people. Content should hence be designed to attract users to non-crowded areas.
- **Enable gesture-based interaction:** Previous research showed that for flat displays, touch, gestures/body position, and interaction through mobile phones are most suitable. The motion-fostering character of cylindrical displays makes them more suitable for gesture-like interaction. However, interaction with one arm only should be enabled as people spend most of the time with their shoulders not parallel to the display.
- **Use frameless or semi-framed content:** Due to the lack of a left and right frame it is challenging for both designers and audience to orient themselves. However, non-framed content should be used that integrates well with the column and exploits the frameless nature.

In the cylindrical displays study we were only able to observe single-user behavior. In contrast the Looking Glass deployment gave us the opportunity to focus on multiple users. We observed and in-depth investigated the following effects:

The Landing Effect. The passersby's walking speed of about 1.4 m/s leads to that by the time they recognized the interactivity of a displays, they already passed it and need to walk back. The landing effect only occurs if nobody is interacting yet. Furthermore, we observed more landing effects for interactivity cues. The landing effect is so strong that it should be designed for, e.g., by deploying very wide displays or multiple displays along the passersby's trajectories.

The Honeypot Effect. We found the honeypot to be a powerful cue to attract attention and communicate interactivity. It even works after multiple days, as people who have seen somebody interacting previously may also try interacting in the future. Hence, displays should be designed to have someone visibly interacting as often as possible. We observed that the audience repositions themselves to see both the display and the user. Hence, the environment needs to support this, for example either by positioning the display in such a way that it is visible from a wide angle, or that considerable space is available directly in front of it.

Immersion and Expressiveness. We found that single user interaction mainly occurred in waiting situations whereas in almost any other case people interacted in groups. Consequently, displays should provide some benefit also for other group members. As people interacted together, they got so immersed that they forgot what happened around them. This led to that they sometimes hit each other accidentally. Furthermore, people usually start with subtle movements before performing more expressive movements. Designs should therefore be reactive to a wide range of movements. To increase safety we found that slower moving objects could be used as they make users conduct slower movements.

10.2 Future Work

The thesis provides a common ground for future research on public displays and on pervasive advertising. At the same time, many open challenges have been identified during the described projects. The remainder of this chapter points out areas for future work, both in the short and in the long term.

10.2.1 Towards a Public Display Architecture

One reason for few current deployments of interactive public displays is the lack of a reference architecture. This general problem has been widely recognized by the community and an increasing number of researchers is currently concerned with developing frameworks, APIs, and conceptual architectures for interactive display networks. Initial results have been published at the *First International Symposium on Pervasive Displays (PerDis'12)*³⁹ and at the latest *Workshop on Infrastructure and Design Challenges of Coupled Display Visual Interfaces (PPD'12)*⁴⁰, for example [15; 59; 74; 119; 174]. Current approaches and high-level architectures serve as a good starting point, but so far fail to anticipate two major challenges in advertising environments.

First, there is an inherent tension between different stakeholders. Traditionally, the display owner and the content provider often used to be the same person (e.g.,

³⁹ Pervasive Display Symposium: <http://www.pervasivedisplays/>, last accessed February 13, 2013

⁴⁰ PPD Workshop website: <http://sachi.cs.st-andrews.ac.uk/activities/workshops/ppd12/>, last accessed February 13, 2013

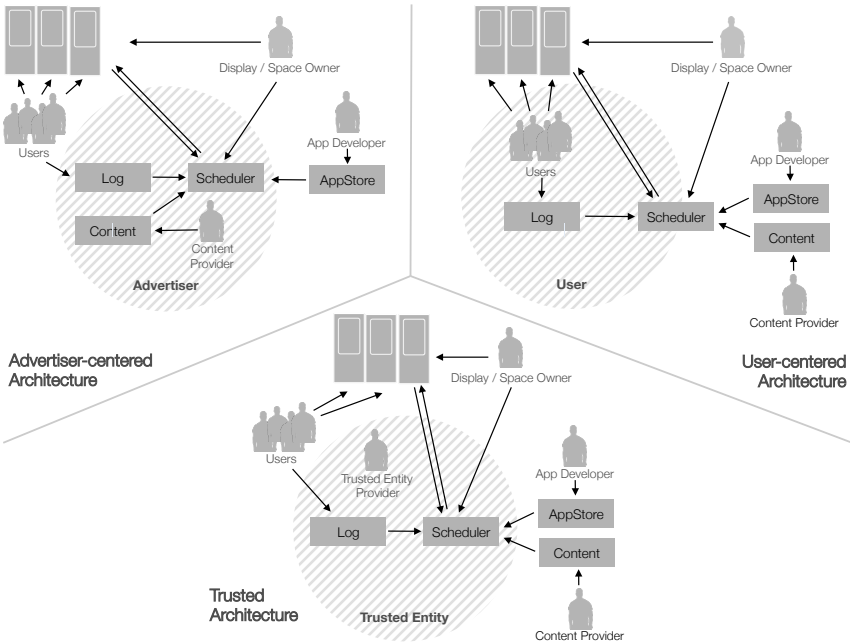


Figure 10.1: Conceptual Architecture Designs: advertiser-centered (top), user-centered (middle), trusted (bottom)

a retailer who deployed a display in his shop windows to advertise his products). As displays become interactive and networked, new stakeholders emerge and the boundaries between roles blur. For example, display owners may buy external content as they lack the expertise to create interactive content, developers may offer (customizable) apps, and viewers may at the same time become content providers as displays support user-generated content. These advances make it necessary to rethink the interplay between the different stakeholders [11; 74] and requires mechanisms to be installed that cater to their needs. For example, display owners may only be willing to share display space as they maintain ultimate control, app developers may want to know how often their apps have been downloaded, advertisers may want to know how successful their campaigns are, and users may want to know what happens with their data. Alt and Schneegaß propose different conceptual architectures [15] that try to reflect these needs (Figure 10.1). In *advertiser-centered architectures*, as we mainly find them today, advertisers try to gather as much data as possible by logging user data. This data is then used to measure success or to adapt content accordingly. On the downside,

users do not have any control or information about what happens to their data. A *user-centered architecture* provides the user control over his logged data and allows him to influence the presentation of content on the display. This approach well protects the user's privacy, but incentives are needed to make the user share his data with the advertiser. Finally, a *trusted architecture* establishes a trusted entity that stores the user data and only makes it available to the advertiser in an anonymized or abstracted form. The user could furthermore be provided access to his data and be enabled to delete them at any time.

Second, the system architecture needs to support designers as they consider two dimensions: how to present content – through time or space multiplexing, or in an integrated format – and whether the user or system should initiate interaction [12]. Table 10.2 provides examples of each possible scenario. *Time multiplexing* involves showing different types of content serially on a single display. For example, world and local news followed by weather forecast updates, cartoons, sports news, and advertisements. *Space multiplexing* involves presenting different kinds of content simultaneously, either in separate regions of the same display or in distinct displays – for example, restaurant advertisements next to a city map. Content and ads can also be *integrated*, meaning that they are seamlessly interweaved. Although in some cases legislation requires news to be clearly distinguished from advertising, ads can be embedded in other forms of content, such as an interactive game. In *user-initiated* public display systems, users start the interaction – for example, by touching the display. *System-initiated* displays present content without the need for user interaction – for example, a looping slideshow.

It is very clear that the public display architecture of the future needs to support advertising as one possible business model besides others. Still, we believe that ways need to be found to put the user more into the focus and unfold the potential of 'symmetric communication' [249]. Thus, public displays can be made more attractive while at the same time providing ways to fund the infrastructure. We believe that in the future large organizations may serve as trusted entities that provide the basic architecture consisting of an application store, a scheduler, and a central place keeping the user data.

Concretely, future work may include the development of a 'Display OS'. This development can be supported through the findings and recommendations presented in this thesis, and should entail addressing interaction challenges, supporting appropriate content and its integration with advertisements, catering to stakeholders' needs, and developing business models.

Table 10.2: Mixing Information and Ads: Sample scenarios (from [12]).

		Content Presentation Mode		
		Time multiplexing	Space multiplexing	Integrated
Interaction Type	User-initiated	Full-screen advertising display that switches to a store directory upon being touched	Browsable bus timetable with ads next to the schedule	Interactive ball game with a corporate logo attached to the balls
	System-initiated	Looping slideshow of various types of content including ads and information	Ads and information displayed side-by-side on the same screen	City map with embedded restaurant ads

10.2.2 Key Performance Indicators

In marketing, metrics similar to those known from traditional displays are currently used for their digital counterparts (Section 3.2.2). However, as public displays become interactive, these may not be appropriate anymore. Interactive applications have the potential to provide a benefit for the user (more interesting content, engaging experiences, quick access to relevant information, etc.). However, the benefit for the advertiser is less clear and whether this vision comes true may strongly depend on answering this question.

Currently, outdoor advertisers are rather reluctant to deploy interactive content on public displays. One possible reason is that the impact of interactivity has not been investigated yet. Can we assume, that because a user interacted with an advertising game, their attitude towards the advertised brand changed in a positive way or that the experience ultimately leads to a sale? To answer this question, researchers have to think about novel metrics, how to quantify them, and which technologies to use for gathering the data – preferably in an automated way.

Cognitive effects are of special interest. As we outlined in Chapter 3, attention is just a first step towards assessing more complex and profound constructs such as involvement, emotions, motives, or attitude. Eye tracking research provided some initial evidence, that changes in attitude and engagement could be measured through quantifying gaze behavior (e.g., number of fixations [287]), but further technologies may be available in the future to more comprehensively assess the states and processes, or their observable manifestations.

Future research could explore concrete metrics, both for observable actions as well as for cognitive processes. In a first step, this includes developing an understanding of the impact of interaction. Dimension to explore could entail the level of expressiveness, the used technology or technique, or how content and ads are interweaved. In a second step, a comprehensive model of key performance indicators could be developed, potentially being based on real-time data.

10.2.3 Payment Models

With novel sensing techniques, new payment models may be applied to public displays in the future. The CPM (Cost per Mille) is currently the prevailing payment model for public displays. As already explained, the size of the audience is being estimated based on previously gathered data and hence has a rather low quality as it can just be assumed – but not ensured – that the audience actually looked at the advertising message (similar to the Internet where a page load does not necessarily mean that an advertisement was seen).

With cameras and eye contact sensors, *Cost per View* models become possible where advertisers are being charged only in cases where users look at the display. Since displays may contain multiple types of content, the accuracy of gaze detection systems may play a major role and with eye trackers becoming commercially available, an analysis on content element level may become feasible. As public displays become interactive, *Cost per Interaction* or *Cost per Interaction Duration* models can be envisaged, which would resemble the Cost per Click model known from the Internet. Advertisers could be charged for users playing interactive games that include advertising content. Finally, a *Cost per Lead*, e.g., as users sign up at the display to use a service, or a *Cost per Sale* as users download an app or song from the display, could be realized.

As future work, suitable technologies to establish the suggested payment models could be explored. Furthermore, researchers could investigate, how such novel payment models could lead to a comprehensive business model.

10.2.4 New Technologies and User Experiences

With advances in technology, novel means for both input and output may change the way we interact with and perceive public displays in the future. On the input side, sophisticated camera technology will enable new interaction techniques as

well as the assessment of user performance and audience behavior. For example, eye trackers could be used to control content explicitly, potentially in combination with further input modalities. In contrast, implicit interaction techniques could make content follow the user (e.g., as he moves in front of very wide displays), or relevant content could be highlighted on large screens, similar to approaches presented in augmented reality (e.g., AR-Multipleye [275]). Cameras with high spatial and temporal resolution will enable measurements in different spectrums (e.g., thermographic cameras), hence enabling emotions or the level of arousal to be assessed. Further input modalities could also include brain computer interfaces that can be used to make selections on a display, for example, choosing a song from a list shown on the screen.

On the output side, novel display technologies may create highly immersive user experiences. This includes high resolution displays as well as 3-D displays. Current 3-D consumer devices use technologies that still require glasses (e.g., active shutter, polarization, or anaglyph 3-D), but autostereoscopic displays will in the future use cameras that can track the gaze of the viewer and adjust the visual angle of the rendered image accordingly. In this way 3-D effects can be created in public spaces with no further devices required. However, research on how to interact with 3-D content is still scarce. Interesting directions of future work could include exploring interaction techniques as well as investigating user acceptance, audience behavior, and user experience around 3-D displays.

10.2.5 Opportunities of Heterogenous Environments

Users are often exposed to multiple displays at the same time. For example, being at the airport, users look at information displays that show the departure timetable or gate information, perceive advertisements on nearby digital signage, and use their personal mobile phone. Experimental research could explore how to deal with such heterogenous display environments and how future architectures and infrastructures could design for this. Open access of these display spaces would allow content providers / advertisers to use all displays at the same time (including input and output), and hence create specific, persistent user experiences. As has been identified by Dey et al. [92], this raises questions as to how the user should deal with multiple input and output options as well as the fact that each display only receives little attention. Using this thesis as a starting point, future work could look at the increased design space of such environments.

10.2.6 Displays as Gateways

Rukzio et al. investigated physical posters as gateways related to mobile services [311]. Their basic idea is to allow users to interact with services indirectly presented through advertising posters via their mobile phone. Such services could provide further information about the advertised product or allow to buy it. Future work could advance this concept to digital displays, drawing from the findings presented in this thesis. Of particular interest may be identifying suitable interaction techniques and providing mechanisms that allow the needs of different users or user groups to be addressed.

10.3 Concluding Remarks

This thesis addresses some of the fundamental problems that designers and developers of future public displays advertising will face. At the same time it is very clear that the contributions we make can only be partial and temporary solutions in a young research field, which is in constant motion. As novel devices enter the market, arbitrary display shapes and sizes enable new advertising spaces. The displays' context will also change, for example when people learn that there is more than just advertising to public displays.

We believe that fundamental changes may happen when global players decide to provide a platform for the shared use of public displays, similar to what recently happened with Android on mobile phones. In the endeavor to prepare researchers and practitioners, interesting questions and challenges emerge from different disciplines, mainly from computer science and marketing. I hope that these questions will help to create a positive future of pervasive displays, full of new, enticing forms of interaction and amazing experiences.

VI

BIBLIOGRAPHY

Bibliography

- [1] R. Adams and C. Russell. Lessons from Ambient Intelligence Prototypes for Universal Access and the User Experiences. In C. Stephanidis and M. Pieper, editors, *Universal Access in Ambient Intelligence Environments*, volume 4397 of *Lecture Notes in Computer Science*, pages 229–243. Springer, 2007.
- [2] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating Contextual Information in Recommender Systems Using a Multidimensional Approach. *ACM Transactions on Information Systems*, 23(1):103–145, Jan. 2005.
- [3] S. Agamanolis. Designing Displays For Human Connectedness. In K. O’Hara, E. Perry, E. Churchill, and D. M. Russel, editors, *Public and Situated Displays*. Kluwer, Dordrecht, 2003.
- [4] A. Ajanki, D. R. Hardoon, S. Kaski, K. Puolamäki, and J. Shawe-Taylor. Can Eyes Reveal Interest? Implicit Queries from Gaze Patterns. *User Modeling and User-Adapted Interaction*, 19(4):307–339, Oct. 2009.
- [5] I. Ajzen. *Attitudes, Personality and Behavior*. Open University Pr, 2005.
- [6] R. Alexander. Olympic glory - a short history of olympic games timing. *watchworld.com*, 2008.
- [7] F. Alt, M. Balz, S. Kristes, A. S. Shirazi, J. Mennenöh, A. Schmidt, H. Schröder, and M. Gödicke. Adaptive User Profiles in Pervasive Advertising Environments. In *Proceedings of the European Conference on Ambient Intelligence* (Salzburg, Austria), AmI’09, pages 276–286, Berlin, Heidelberg, 2009. Springer-Verlag.

- [8] F. Alt, D. Bial, T. Kubitzka, A. S. Shirazi, M. Ortel, B. Zurmaar, F. Zaidan, T. Lewen, and A. Schmidt. Digifieds: Evaluating Suitable Interaction Techniques for Shared Public Notice Areas. In *Adjunct Proceedings of Pervasive 2011* (San Francisco, CA, USA), Pervasive'11, 2011.
- [9] F. Alt, C. Evers, and A. Schmidt. Users' View on Context-Sensitive Car Advertisements. In *Proceedings of the 7th International Conference on Pervasive Computing* (Nara, Japan), Pervasive'09, pages 9–16, Berlin, Heidelberg, 2009. Springer-Verlag.
- [10] F. Alt, T. Kubitzka, D. Bial, F. Zaidan, M. Ortel, B. Zurmaar, T. Lewen, A. S. Shirazi, and A. Schmidt. Digifieds: Insights into Deploying Digital Public Notice Areas in the Wild. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia* (Beijing, China), MUM'11, pages 165–174, New York, NY, USA, 2011. ACM.
- [11] F. Alt, N. Memarovic, I. Elhart, D. Bial, A. Schmidt, M. Langheinrich, G. Harboe, E. Huang, and M. P. Scipioni. Designing Shared Public Display Networks: Implications from Today's Paper-based Notice Areas. In *Proceedings of the 9th International Conference on Pervasive Computing* (San Francisco, CA, USA), Pervasive'11, pages 258–275, Berlin, Heidelberg, 2011. Springer-Verlag.
- [12] F. Alt, J. Müller, and A. Schmidt. Advertising on Public Display Networks. *IEEE Computer*, 45(5):50–56, 2012.
- [13] F. Alt, A. Sahami Shirazi, T. Kubitzka, and A. Schmidt. Interaction techniques for creating and exchanging content with public displays. In *Proceedings of the 2012 ACM Conference on Human Factors in Computing Systems* (Paris, France), CHI'13, New York, NY, USA, 2013. ACM.
- [14] F. Alt, A. Schmidt, and C. Evers. Mobile Contextual Displays. In *Proceedings of the 1st Workshop on Pervasive Advertising* (Nara, Japan), PerAd'09, 2009.
- [15] F. Alt and S. Schneegass. A Conceptual Architecture for Pervasive Advertising in Public Display Networks. In *Proceedings of the 3rd Workshop on Infrastructure and Design Challenges of Coupled Display Visual Interfaces* (Capri, Italy), PPD'12, 2012.
- [16] F. Alt, S. Schneegaß, A. Schmidt, J. Müller, and N. Memarovic. How to Evaluate Public Displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays* (Porto, Portugal), PerDis'12, pages 171–176, New York, NY, USA, 2012. ACM.

- [17] F. Alt, A. Shirazi, S. Legien, A. Schmidt, and J. Mennenöh. Creating Meaningful Melodies from Text Messages. In *Proceedings of the 2010 Conference on New Interfaces for Musical Expression* (Sydney, Australia), NIME'10, pages 63–68, 2010.
- [18] F. Alt, A. S. Shirazi, A. Schmidt, and J. Mennenöh. Increasing the User's Attention on the Web: Using Implicit Interaction based on Gaze Behavior to Tailor Content. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction* (Copenhagen, Denmark), NordiCHI'12, New York, NY, USA, 2012. ACM.
- [19] American Marketing Association (AMA). Definition of Marketing, 2012. <http://www.marketingpower.com/AboutAMA/Pages/DefinitionofMarketing.aspx>, last accessed February 13, 2013.
- [20] S. Antifakos and B. Schiele. Laughinglily: Using a flower as a real-world information displays. In *Adjunct Proceedings of the Fifth International Conference on Ubiquitous Computing* (Seattle, WA, USA), UbiComp'03, pages 161–162, 2003.
- [21] R. Arora. Validation Of An S-O-R Model For Situation, Enduring, And Response Components Of Involvement. *Journal of Marketing Research*, 19:505–516, 1982.
- [22] S. Asteriadis, P. Tzouveli, K. Karpouzis, and S. Kollias. Non-Verbal Feedback on User Interest Based on Gaze Direction and Head Pose. In *Proceedings of the Second International Workshop on Semantic Media Adaptation and Personalization* (London, United Kingdom), SMAP'07, pages 171–178, Washington, DC, USA, 2007. IEEE Computer Society.
- [23] R. Atterer, M. Wnuk, and A. Schmidt. Knowing the User's Every Move: User Activity Tracking for Website Usability Evaluation and Implicit Interaction. In *Proceedings of the 15th international Conference on the World Wide Web* (Edinburgh, Scotland), WWW'06, pages 203–212, New York, NY, USA, 2006. ACM.
- [24] R. Ballagas, J. Borchers, M. Rohs, and J. G. Sheridan. The Smart Phone: A Ubiquitous Input Device. *IEEE Pervasive Computing*, 5(1):70–77, Jan. 2006.
- [25] R. Ballagas, M. Rohs, and J. G. Sheridan. Sweep and Point and Shoot: Phocem-based Interactions for Large Public Displays. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA), CHI EA '05, pages 1200–1203, New York, NY, USA, 2005. ACM.

- [26] A. Bangor, P. T. Kortum, and J. T. Miller. An Empirical Evaluation of the System Usability Scale. *International Journal on Human Computer Interaction*, 24(6):574–594, 2008.
- [27] J. E. Bardram, T. R. Hansen, and M. Soegaard. AwareMedia: a Shared Interactive Display Supporting Social, Temporal, and Spatial Awareness in Surgery. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work (Banff, AB, Canada), CSCW '06*, pages 109–118, New York, NY, USA, 2006. ACM.
- [28] C. Bauer and S. Spiekermann. Conceptualizing Context for Pervasive Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 159–183. Springer Limited London, 2011.
- [29] R. B. Bechtel and A. Churchman. *Handbook of Environmental Psychology*, volume 24 (2). Wiley, 2002.
- [30] B. Begole. *Ubiquitous Computing for Business: Find New Markets, Create Better Businesses, and Reach Customers Around the World 24-7-365*. Ft Press, 2011.
- [31] M. Behr. First ever installation of electronic paper display on hamburg hochbahn train by vossloh it. *Funkwerk Information Technologies*, 2006.
- [32] H. Benko. Beyond Flat Surface Computing: Challenges of Depth-Aware and Curved Interfaces. In *Proceedings of the 17th ACM International Conference on Multimedia (Beijing, China), MM'09*, pages 935–944, New York, NY, USA, 2009. ACM.
- [33] H. Benko, A. D. Wilson, and R. Balakrishnan. Sphere: Multi-Touch Interactions on a Spherical Display. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology (Monterey, CA, USA), UIST'08*, pages 77–86, New York, NY, USA, 2008. ACM.
- [34] S. Berger, R. Kjeldsen, C. Narayanaswami, C. Pinhanez, M. Podlaseck, and M. Raghunath. Using Symbiotic Displays to View Sensitive Information in Public. In *Proceedings of the Third IEEE International Conference on Pervasive Computing and Communications (Kauai, HI, USA), PER-COM'05*, pages 139–148, Washington, DC, USA, 2005. IEEE Computer Society.
- [35] G. Beyer, F. Alt, J. Müller, A. Schmidt, K. Isakovic, S. Klose, M. Schiewe, and I. Haulsen. Audience Behavior Around Large Interactive Cylindrical

- Screens. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada), CHI'11, pages 1021–1030, New York, NY, USA, 2011. ACM.
- [36] D. Beymer and D. M. Russell. WebGazeAnalyzer: A System for Capturing and Analyzing Web Reading Behavior Using Eye Gaze. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA), CHI'05 EA, pages 1913–1916, NY, USA, 2005. ACM.
- [37] R. Biddle, S. Chiasson, and P. Van Oorschot. Graphical passwords: Learning from the first twelve years. *ACM Computing Surveys*, 44(4):19:1–19:41, Sept. 2012.
- [38] R. Biedert, G. Buscher, S. Schwarz, J. Hees, and A. Dengel. Text 2.0. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems* (Atlanta, GA, USA), CHI'10 EA, pages 4003–4008, New York, NY, USA, 2010. ACM.
- [39] W. Biehlig, C. Deter, S. Dube, B. Hill, S. Helling, K. Isakovic, S. Klose, and M. Schiewe. LaserCave - Some Building Blocks for immersive Screens. In *VR/AR Statustagung, Leipzig*, 2004.
- [40] G. Bilchev and D. Marston. Personalised Advertising – Exploiting the Distributed User Profiles. *BT Technology Journal*, 21(1):84–90, Jan. 2003.
- [41] A. Blackler, V. Popovic, and D. Mahar. The Nature of Intuitive Use of Products: An Experimental Approach. *Design Studies*, 24(6):491–506, 2003.
- [42] S. A. Bly, S. R. Harrison, and S. Irwin. Media Spaces: Bringing People Together in a Vvideo, Audio, and Computing Environment. *Communications of the ACM*, 36(1):28–46, Jan. 1993.
- [43] K. Boehret. Galaxy Nexus: An In-Your-Face Android Phone. *The Wall Street Journal*, 2012. http://online.wsj.com/article/SB100014-24052970204531404577054_233319145676.html, last accessed February 13, 2013.
- [44] M. Böhlen and M. Mateas. Office Plant# 1: Intimate Space and Contemplative Entertainment. *Leonardo*, 31(5):345–348, 1998.
- [45] S. Boring, D. Baur, A. Butz, S. Gustafson, and P. Baudisch. Touch Projector: Mobile Interaction Through Video. In *Proceedings of the 28th*

- international Conference on Human Factors in Computing Systems* (Atlanta, GA, USA), CHI'10, pages 2287–2296, New York, NY, USA, 2010. ACM.
- [46] R. Borovoy, F. Martin, M. Resnick, and B. Silverman. GroupWear: Nametags that Tell About Relationships. In *Proceedings of the 16th international Conference on Human Factors in Computing Systems* (Los Angeles, CA, USA), CHI'98, pages 329–330. ACM, 1998.
- [47] R. Borovoy, F. Martin, S. Vemuri, M. Resnick, B. Silverman, and C. Hancock. Meme Tags and Community Mirrors: Moving from Conferences to Collaboration. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (Seattle, WA, USA), CSCW'98, pages 159–168, New York, NY, USA, 1998. ACM.
- [48] R. Borovoy, M. McDonald, F. Martin, and M. Resnick. Things that Blink: Computationally Augmented Name Tags. *IBM Systems Journal*, 35(3-4):488–495, Sept. 1996.
- [49] M. Bosworth. Loyalty Cards: Reward or Threat?, 2005. http://www.consumer-affairs.com/news04/2005/loyalty_cards.html, last accessed February 13, 2013.
- [50] P. B. Brandtzæg, A. Følstad, and J. Heim. Enjoyment: Lessons from Karasek. In M. A. Blythe, K. Overbeeke, A. F. Monk, and P. C. Wright, editors, *Funology*, pages 55–65, Norwell, MA, USA, 2004. Kluwer Academic Publishers.
- [51] H. Brignull, S. Izadi, G. Fitzpatrick, Y. Rogers, and T. Rodden. The Introduction of a Shared Interactive Surface into a Communal Space. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, IL, USA), CSCW '04, pages 49–58, New York, NY, USA, 2004. ACM.
- [52] H. Brignull and Y. Rogers. Enticing People to Interact with Large Public Displays in Public Spaces. In M. Rauterberg, M. Menozzi, and J. Wesson, editors, *Proceedings of the IFIP TC13 International Conference on Human-Computer Interaction* (Zürich, Switzerland), Interact'03. IOS Press, 2003.
- [53] L. Brill. Led electronic message reader boards: Watching the world go by in streaming headlines. *signindustry.com*, 2008.
- [54] M. Bruhn. *Kommunikationspolitik*. Vahlen, München, 2009.

- [55] A. Bulling, F. Alt, and A. Schmidt. Increasing The Security Of Gaze-Based Cued-Recall Graphical Passwords Using Saliency Masks. In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems* (Austin, TX, USA), CHI'12, pages 3011–3020, New York, NY, USA, 2012. ACM.
- [56] A. Bulling and H. Gellersen. Toward Mobile Eye-Based Human-Computer Interaction. *IEEE Pervasive Computing*, 9(4):8–12, Oct. 2010.
- [57] P. Burgoyne. Sao Paulo: The City That Said No To Advertising. Bloomberg Business Week, 2007. <http://www.businessweek.com/stories/2007-06-18/s-o-paulo-the-city-that-said-no-to-advertisingbusinessweek-business-news-stock-market-and-financial-advice>, last accessed February 13, 2013.
- [58] M. Burke, A. Hornof, E. Nilsen, and N. Gorman. High-Cost Banner Blindness: Ads Increase Perceived Workload, Hinder Visual Search, and are Forgotten. *ACM Transactions on Human Computer Interaction*, 12(4):423–445, Dec. 2005.
- [59] C. Burns, T. Seyed, T. D. Hellmann, J. Ferreira, and F. Maurer. Towards A Usable API for Constructing Interactive Multi-Surface Systems. In *Proceedings of the 3rd Workshop on Infrastructure and Design Challenges of Coupled Display Visual Interfaces* (Capri, Italy), PPD'12, 2012.
- [60] G. Buscher, A. Dengel, and L. Van Elst. Query Expansion Using Gaze-Based Feedback On The Subdocument Level. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Singapore, Singapore), SIGIR '08, pages 387–394, New York, NY, USA, 2008. ACM.
- [61] S. Carter, E. Churchill, L. Denoue, J. Helfman, and L. Nelson. Digital Graffiti: Public Annotation of Multimedia Content. In *CHI'04 Extended Abstracts on Human Factors in Computing Systems* (Vienna, Austria), CHI EA'04, pages 1207–1210, New York, NY, USA, 2004. ACM.
- [62] P. Chandon, J. W. Hutchinson, and S. H. Young. Measuring the Value of Point-of-Purchase Marketing with Commercial Eye-Tracking Data, 2007.
- [63] D. W. Chen. Imes sq. sign turns corner into silen. *The New York Times*, May 6, 1997, 2004.

- [64] M. C. Chen, J. R. Anderson, and M. H. Sohn. What Can a Mouse Cursor Tell us More? Correlation of Eye/Mouse Movements on Web Browsing. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, WA, USA), CHI EA'01, pages 281–282, New York, NY, USA, 2001. ACM.
- [65] K. Cheverst, A. Dix, D. Fitton, C. Kray, M. Rouncefield, C. Sas, G. Saslis-Lagoudakis, and J. G. Sheridan. Exploring Bluetooth-based Mobile Phone Interaction with the Hermes Photo Display. In *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services* (Salzburg, Austria), MobileHCI'05, pages 47–54, New York, NY, USA, 2005. ACM.
- [66] K. Cheverst, D. Fitton, and A. J. Dix. Exploring the Evolution of Office Door Displays. In K. O'Hara, E. Perry, E. Churchill, and D. M. Russel, editors, *Public and Situated Displays – Social and Interactional Aspects of Shared Display Technologies*, pages 141–169. Kluwer, Dordrecht, 2003.
- [67] K. Cheverst, N. Taylor, M. Rouncefield, A. Galani, and C. Kray. The Challenge of Evaluating Situated Display-based Technology Interventions Designed to Foster a Sense of Community. In *Proceedings of the Workshop on Ubiquitous Systems Evaluation* (Seoul, South Korea), USE'08, 2008.
- [68] A. Chew, V. Leclerc, S. Sadi, A. Tang, and H. Ishii. SPARKS. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA), CHI EA'05, pages 1276–1279, New York, NY, USA, 2005. ACM.
- [69] S. Chiasson, P. C. van Oorschot, and R. Biddle. Graphical Password Authentication Using Cued Click Points. In J. Biskup and J. Lopez, editors, *Proceedings of the 12th European Symposium On Research In Computer Security* (Dresden, Germany), ESORICS'07, pages 359–374. Springer, 2007.
- [70] A. D. Christian and B. L. Avery. Speak out and annoy someone: experience with intelligent kiosks. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (The Hague, The Netherlands), CHI '00, pages 313–320, New York, NY, USA, 2000. ACM.
- [71] E. Churchill, L. Les Nelson, P. Murphy, and J. Helfman. The Plasma Poster Network. In K. O'Hara, E. Perry, E. Churchill, and D. M. Russel, editors, *Public and Situated Displays – Social and Interactional Aspects of Shared Display Technologies*, pages 141–169. Kluwer, Dordrecht, 2003.

- [72] E. F. Churchill, L. Nelson, and L. Denoue. Multimedia Fliers: Information Sharing with Digital Community Bulletin Boards. In M. Huysman, E. Wenger, and V. Wulf, editors, *Communities and Technologies*, pages 97–117, Deventer, The Netherlands, 2003. Kluwer.
- [73] M. Claypool, P. Le, M. Wased, and D. Brown. Implicit Interest Indicators. In *Proceedings of the 6th International Conference on Intelligent User Interfaces* (Santa Fe, NM, USA), IUI'01, pages 33–40, New York, NY, USA, 2001. ACM.
- [74] S. Clinch, N. Davies, T. Kubitzka, and A. Schmidt. Designing Application Stores for Public Display Networks. In *Proceedings of the 2012 International Symposium on Pervasive Displays* (Porto, Portugal), PerDis'12, pages 10:1–10:6, New York, NY, USA, 2012. ACM.
- [75] P. Cohan. 5 Display Advertising Metrics you Should Know. iMedia Connection, March 2012. http://www.imediaconnection.com/article_full.aspx?id=31128, last accessed February 13, 2013.
- [76] M. Collomb, M. Hascoët, P. Baudisch, and B. Lee. Improving Drag-and-Drop on Wall-size Displays. In *Proceedings of Graphics Interface 2005* (Victoria, BC, Canada), GI'05, pages 25–32, Waterloo, ON, Canada, 2005. Canadian Human-Computer Communications Society.
- [77] B. Congleton. Prospero – A “Visual Commons” Framework for Public Displays, 2007. http://dlc.dlib.indiana.edu/dlc/bitstream/handle/10535/3816/prospero_wrap_up.pdf?sequence=1, last accessed February 13, 2013.
- [78] D. Cordova and M. Lepper. Intrinsic Motivation and the Process of Learning: Beneficial Effects of Contextualization, Personalization, and Choice. *Journal of Educational Psychology*, 88(4):715–730, 1996.
- [79] M. Csikszentmihalyi, C. Kolo, and T. Baur. Flow: The Psychology of Optimal Experience. *Australian Occupational Therapy Journal*, 51(1):3–12, 2004.
- [80] J. E. Cutting and L. T. Kozlowski. Recognizing Friends by their Walk: Gait Perception without Familiarity Cues. *Bulletin of the Psychonomic Society*, 9(5):353–356, 1977.
- [81] M. Czerwinski, S. Dumais, G. Robertson, S. Dziadosz, S. Tiernan, and M. van Dantzich. Visualizing Implicit Queries for Information Management and Retrieval. In *Proceedings of the SIGCHI Conference on Human*

- Factors in Computing Systems* (Pittsburgh, PA, USA), CHI'99, pages 560–567, New York, NY, USA, 1999. ACM.
- [82] D. Daley. Installation profile: Into the tunnel. *Sound and Video Contractor*, 2008.
- [83] N. Davies, A. Friday, P. Newman, S. Rutledge, and O. Storz. Using Bluetooth Device Names to Support Interaction in Smart Environments. In *Proceedings of the 7th international Conference on Mobile Systems, Applications, and Services* (Krakow, Poland), MobiSys'09, pages 151–164, New York, NY, USA, 2009. ACM.
- [84] F. D. Davis. *A Technology Acceptance Model for Empirically Testing new End-User Information Systems: Theory and Results*. PhD thesis, Massachusetts Institute of Technology, 1986.
- [85] F. D. Davis, R. P. Bagozzi, and P. R. Warshaw. User Acceptance of Computer Technology: A Comparison of Two Theoretical Models. *Management Science*, 35(8):982–1003, Aug. 1989.
- [86] A. De Angeli, L. Coventry, G. Johnson, and K. Renaud. Is a Picture Really Worth a Thousand Words? Exploring the Feasibility of Graphical Authentication Systems. *International Journal on Human-Computer Studies*, 63(1-2):128–152, July 2005.
- [87] J. E. de Castro and H. Shimakawa. Mobile Advertisement System Utilizing User's Contextual Information. In *Proceedings of the 7th International Conference on Mobile Data Management* (Nara, Japan), MDM'06, Washington, DC, USA, 2006. IEEE Computer Society.
- [88] E. Deci and R. Ryan. *Intrinsic Motivation and Self-Determination in Human Behavior*. Springer, 1985.
- [89] C. Dennis, A. Newman, R. Michon, J. Josko Brakus, and L. T. Wright. The Mediating Effects of Perception and Emotion: Digital Signage in Mall Atmospherics. *Journal of Retailing and Consumer Services*, 17(3):205–215, 2010.
- [90] L. Denoue, L. Nelson, and E. Churchill. AttrActive Windows: Dynamic Windows for Digital Bulletin Boards. In *CHI'03 Extended Abstracts on Human Factors in Computing Systems* (Ft. Lauderdale, FL, USA), CHI EA'03, pages 746–747, New York, NY, USA, 2003. ACM.

- [91] R. DePaula. A New Era in Human Computer Interaction: The Challenges of Technology as a Social Proxy. In *Proceedings of the Latin American Conference on Human-computer interaction* (Rio de Janeiro, Brazil), CLIHC'03, pages 219–222, New York, NY, USA, 2003. ACM.
- [92] A. K. Dey, P. Ljungstrand, and A. Schmidt. Distributed and Disappearing User Interfaces in Ubiquitous Computing. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, WA, USA), CHI EA'01, pages 487–488, New York, NY, USA, 2001. ACM.
- [93] A. Dix, J. E. Finlay, G. D. Abowd, and R. Beale. *Human-Computer Interaction*. Prentice-Hall, Inc., 3rd edition, 2003.
- [94] H. Drewes and A. Schmidt. The MAGIC Touch: Combining MAGIC-Pointing with a Touch-Sensitive Mouse. In *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II* (Uppsala, Sweden), INTERACT'09, pages 415–428. Springer, 2009.
- [95] A. Duchowski. A Breadth-First Survey of Eye-Tracking Applications. *Behavior Research Methods*, 34(4):455–470, 2002.
- [96] S. Dumais, E. Cutrell, R. Sarin, and E. Horvitz. Implicit Queries (IQ) for Contextualized Search. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Sheffield, United Kingdom), SIGIR'04, pages 594–594, New York, NY, USA, 2004. ACM.
- [97] J. Enns, E. Austen, V. Di Lollo, R. Rauschenberger, and S. Yantis. New Objects Dominate Luminance Transients in Setting Attentional Priority. *Journal of Experimental Psychology: Human Perception and Performance*, 27(6):1287, 2001.
- [98] A. Erbad, M. Blackstock, A. Friday, R. Lea, and J. Al-Muhtadi. MAGIC Broker: A Middleware Toolkit for Interactive Public Displays. In *Proceedings of the 2008 Sixth Annual IEEE International Conference on Pervasive Computing and Communications* (Hong Kong, China), PERCOM '08, pages 509–514, Washington, DC, USA, 2008. IEEE Computer Society.
- [99] K. M. Everitt, T. Bragin, J. Fogarty, and T. Kohno. A Comprehensive Study of Frequency, Interference, and Training of Multiple Graphical Passwords. In *Proceedings of the 27th international Conference on Human Factors in Computing Systems* (Boston, MA, USA), CHI'09, pages 889–898, New York, NY, USA, 2009. ACM.

- [100] J. Exeler, M. Buzeck, and J. Müller. eMir: Digital Signs that React to Audience Emotion. In *Proceedings of the 2nd Workshop on Pervasive Advertising* (Lübeck, Germany), PerAd'09, 2009.
- [101] J. Falk and S. Björk. The BubbleBadge: a Wearable Public Display. In *CHI'99 Extended Abstracts on Human Factors in Computing Systems* (Pittsburgh, PA, USA), CHI EA'99, pages 318–319, New York, NY, USA, 1999. ACM.
- [102] M. Farid, F. Murtagh, and J. Starck. Computer Display Control and Interaction Using Eye-Gaze. *Journal of the Society for Information Display*, 10(3):289, 2002.
- [103] A. Fass, J. Forlizzi, and R. Pausch. MessyDesk and MessyBoard: Two Designs Inspired by the Goal of Improving Human Memory. In *Proceedings of the 4th Conference on Designing Interactive Systems* (London, United Kingdom), DIS'02, pages 303–311, New York, NY, USA, 2002. ACM.
- [104] A. Fatah gen Schieck, V. Kostakos, and A. Penn. The Urban Screen as a Socialising Platform: Exploring the Role of Place Within the Urban Space. In F. Eckhardt, J. Geelhaar, L. Colini, K. S. Willis, K. Chorianopoulos, and R. Henning, editors, *MEDIACITY – Situations, Practices and Encounters*, pages 285–305. Frank & Timme GmbH, 2008.
- [105] A. Ferscha and S. Vogl. Pervasive Web Access via Public Communication Walls. In *Proceedings of the First International Conference on Pervasive Computing* (Zürich, Switzerland), Pervasive'02, pages 84–97, London, United Kingdom, 2002. Springer-Verlag.
- [106] J. Finney, S. Wade, N. Davies, and A. Friday. Flump: The Flexible Ubiquitous Monitor Project. In *Proceedings of the 3rd Cabernet Radicals Workshop*, 1996. <http://research.cs.ncl.ac.uk/cabernet/www.laas.research.ec.org/cabernet/workshops/radicals/1996/papers/flump-finney.html>, last accessed February 13, 2013.
- [107] R. S. Fish, R. E. Kraut, and B. L. Chalfonte. The VideoWindow System in Informal Communication. In *Proceedings of the 1990 ACM Conference on Computer Supported Cooperative Work* (Los Angeles, CA, USA), CSCW'90, pages 1–11, New York, NY, USA, 1990. ACM.
- [108] B. J. Fogg. *Persuasive Technology: Using Computers to Change what we Think and Do*. Morgan Kaufmann, 2003.

- [109] D. Fono and R. Vertegaal. EyeWindows: Evaluation of Eye-Controlled Zooming Windows for Focus Selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, OR, USA), CHI'05, pages 151–160, New York, NY, USA, 2005. ACM.
- [110] A. Forget, S. Chiasson, and R. Biddle. Shoulder-Surfing Resistance with Eye-Gaze Entry in Cued-Recall Graphical Passwords. In *Proceedings of the 28th international Conference on Human Factors in Computing Systems* (Atlanta, GA, USA), CHI'10, pages 1107–1110, New York, NY, USA, 2010. ACM.
- [111] S. Fox, K. Karnawat, M. Mydland, S. Dumais, and T. White. Evaluating Implicit Measures to Improve Web Search. *ACM Transactions on Information Systems*, 23(2):147–168, Apr. 2005.
- [112] S. Franconeri and D. Simons. Moving and Looming Stimuli Capture Attention. *Attention, Perception, & Psychophysics*, 65(7):999–1010, 2003.
- [113] H. Freter and J. Naskrent. *Markt- und Kundensegmentierung*. Kohlhammer Edition Marketing. Kohlhammer, 2nd edition, 2008.
- [114] S. Frintrop, E. Rome, and H. I. Christensen. Computational Visual Attention Systems and their Cognitive Foundations: A Survey. *ACM Transactions on Applied Perception*, 7(1):1–39, 2010.
- [115] K. Galloway and S. Rabinowitz. Hole-in-Space, 1980. <http://www.ecafe.com/getty/HIS/>, last accessed February 13, 2013.
- [116] G. G. Gallup JR. Chimpanzees: Self-Recognition. *Science*, 167(3914):86–87, 1970.
- [117] G. D. Garson. Significance. Statnotes: Topics in Multivariate Analysis, 2011. <http://faculty.chass.ncsu.edu/garson/PA765/signif.htm>, last accessed February 13, 2013.
- [118] G. D. Garson. Student's t-Test of Difference of Means. Statnotes: Topics in Multivariate Analysis, 2011. <http://faculty.chass.ncsu.edu/garson/PA765/ttest.htm>, last accessed February 13, 2013.
- [119] S. Gehring, F. Daiber, and C. Lander. Towards Universal, Direct Remote Interaction with Distant Public Displays. In *Proceedings of the 3rd Workshop on Infrastructure and Design Challenges of Coupled Display Visual Interfaces* (Capri, Italy), PPD'12, 2012.

- [120] H. Geiger. Building the Digital Out-Of-Home Privacy Infrastructure. In *Proceedings of the 3rd Workshop on Pervasive Advertising* (Helsinki, Finland), PerAd'09, 2009.
- [121] F. Gil-Castiñeira, A. Fernández-López, C. L. Bravo, N. Cid-Vieytes, D. Conde-Lagoa, E. Costa-Montenegro, and F. J. González-Castaño. Runwithus: a social sports application in the ubiquitous oulu environment. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia* (Beijing, China), MUM '11, pages 195–204, New York, NY, USA, 2011. ACM.
- [122] D. Goodin. TJX Suspect Indicted in Heartland, Hannaford Breaches. *The Register*, 2009. http://www.theregister.co.uk/2009/08/17/heartland_payment_suspect/, last accessed February 13, 2013.
- [123] G. O. Goodman and M. J. Abel. Collaboration Research in SCL. In *Proceedings of the 1986 ACM Conference on Computer Supported Cooperative Work* (Austin, TX, USA), CSCW'86, pages 246–251, New York, NY, USA, 1986. ACM.
- [124] A. Grasso, M. Muehlenbrock, F. Roulland, and D. Snowdon. Supporting Communities of Practice with Large Screen Displays. In K. O'Hara, E. Perry, E. Churchill, and D. M. Russel, editors, *Public and Situated Displays – Social and Interactional Aspects of Shared Display Technologies*, pages 261–282. Kluwer, Dordrecht, 2003.
- [125] S. Greenberg, M. Boyle, and J. Laberge. PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal. *Personal and Ubiquitous Computing*, 3(1):54–64, 1999.
- [126] S. Greenberg and M. Rounding. The Notification Collage: Posting Information to Public and Personal Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, WA, USA), CHI'01, pages 514–521, New York, NY, USA, 2001. ACM.
- [127] A. G. Greenwald. Audience Involvement in Advertising: Four Levels. *Journal of Consumer Research*, 11:581–592, 1984.
- [128] T. Grossman and R. Balakrishnan. Collaborative Interaction with Volumetric Displays. In *Proceedings of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy), CHI'08, pages 383–392, New York, NY, USA, 2008. ACM.

- [129] T. Grossman, D. Wigdor, and R. Balakrishnan. Multi-Finger Gestural Interaction with 3-D Volumetric Displays. In *Proceedings of the 17th annual ACM Symposium on User Interface Software and Technology* (Santa Fe, NM, USA), UIST'04, pages 61–70, New York, NY, USA, 2004. ACM.
- [130] F. Hamker. The Emergence of Attention by Population-based Inference and its Role in Distributed Processing and Cognitive Control of Vision. *Computer Vision and Image Understanding*, 100(1):64–106, 2005.
- [131] D. W. Hansen, D. J. C. MacKay, J. P. Hansen, and M. Nielsen. Eye Tracking off the Shelf. In *Proceedings of the 2004 Symposium on Eye Tracking Research & Applications* (San Antonio, TX, USA), ETRA'04, pages 58–58, New York, NY, USA, 2004. ACM.
- [132] D. R. Hardoon and J. S. Taylor. Information Retrieval by Inferring Implicit Queries from Eye Movements. In *Proceedings of the 11th International Conference on Intelligence and Statistics IS'07*, 2007.
- [133] R. Hardy, E. Rukzio, M. Wagner, and M. Paolucci. Exploring Expressive NFC-Based Mobile Phone Interaction with Large Dynamic Displays. In *Proceedings of the 2009 First International Workshop on Near Field Communication NFC'09*, pages 36–41, Washington, DC, USA, 2009. IEEE Computer Society.
- [134] S. Hart and L. Staveland. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Human Mental Workload*, 1:139–183, 1988.
- [135] M. Hassenzahl, M. Burmester, and F. Koller. AttrakDiff: A Questionnaire to Measure Perceived Hedonic and Pragmatic Quality. (Stuttgart, Germany), In *Proceedings of Mensch & Computer 2003 – Interaktion in Bewegung*, pages 187–196, 2003.
- [136] G. Hauland. Measuring Team Situation Awareness by Means of Eye Movement Data. In *Proceedings of HCI International 2003* (Mahwah, NJ, USA), HCII'03, pages 230–234. Lawrence Erlbaum Associates, 2003.
- [137] R. Heath, D. Brandt, and A. Nairn. Brand Relationships: Strengthened by Emotion, Weakened by Attention. *Journal of Advertising Research*, 46(4):410, 2006.
- [138] J. Heikenfeld, P. Drzaic, J.-S. Yeo, and T. Koch. A Critical Review of the Present and Future Prospects for Electronic Paper. *Journal of the Society for Information Display*, 19(2):129–156, 2011.

- [139] J. M. Heiner, S. E. Hudson, and K. Tanaka. The Information Percolator: Ambient Information Display in a Decorative Object. In *Proceedings of the 12th annual ACM Symposium on User Interface Software and Technology* (Asheville, NC, USA), UIST '99, pages 141–148, New York, NY, USA, 1999. ACM.
- [140] O. Hilliges, S. Izadi, A. D. Wilson, S. Hodges, A. Garcia-Mendoza, and A. Butz. Interactions in the Air: Adding Further Depth to Interactive Tabletops. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology* (Victoria, BC, Canada), UIST'09, pages 139–148, New York, NY, USA, 2009. ACM.
- [141] U. Hinrichs and S. Carpendale. Gestures in the Wild: Studying Multi-Touch Gesture Sequences on Interactive Tabletop Exhibits. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada), CHI'11, pages 3023–3032, New York, NY, USA, 2011. ACM.
- [142] B. Hoanca and K. Mock. Secure graphical password system for high traffic public areas. In *Proceedings of the 2006 Symposium on Eye tracking Research & Applications* (San Diego, California), ETRA'06, pages 35–35, New York, NY, USA, 2006. ACM.
- [143] D. Holman and R. Vertegaal. Organic User Interfaces: Designing Computers in Any Way, Shape, or Form. *Communications of the ACM*, 51(6):48–55, June 2008.
- [144] E. Hornecker. Interactions Around a Contextually Embedded System. In *Proceedings of the fourth international Conference on Tangible, Embedded, and Embodied Interaction* (Cambridge, MA, USA), TEI'10, pages 169–176, New York, NY, USA, 2010. ACM.
- [145] S. Hosio, H. Kukka, and J. Rieki. Leveraging Social Networking Services to Encourage Interaction in Public Spaces. In *Proceedings of the 7th International Conference on Mobile and Ubiquitous Multimedia* (Umea, Sweden), MUM'08, pages 2–7, New York, NY, USA, 2008. ACM.
- [146] N. Hristova and G. M. P. O'Hare. Ad-me: Wireless Advertising Adapted to the User Location, Device and Emotions. In *Proceedings of the Proceedings of the 37th Annual Hawaii International Conference on System Sciences* (Waikoloa Village, HI, USA), HICSS'04, Washington, DC, USA, 2004. IEEE Computer Society.

- [147] E. M. Huang, A. Koster, and J. Borchers. Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays? In *Proceedings of the 6th International Conference on Pervasive Computing* (Sydney, Australia), Pervasive'08, pages 228–243, Berlin, Heidelberg, 2008. Springer-Verlag.
- [148] E. M. Huang and E. D. Mynatt. Semi-Public Displays for Small, Co-located Groups. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, FL, USA), CHI'03, pages 49–56, New York, NY, USA, 2003. ACM.
- [149] E. M. Huang, E. D. Mynatt, and J. P. Trimble. When Design Just Isn't Enough: The Unanticipated Challenges of the Real World for Large Collaborative Displays. *Personal Ubiquitous Computing*, 11(7):537–547, 2007.
- [150] E. M. Huang, D. M. Russell, and A. E. Sue. IM Here: Public Instant Messaging on Large, Shared Displays for Workgroup Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vienna, Austria), CHI'04, pages 279–286, New York, NY, USA, 2004. ACM.
- [151] H. Hutchinson, W. Mackay, B. Westerlund, B. B. Bederson, A. Druin, C. Plaisant, M. Beaudouin-Lafon, S. Conversy, H. Evans, H. Hansen, N. Roussel, and B. Eiderbäck. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, FL, USA), CHI'03, pages 17–24, New York, NY, USA, 2003. ACM.
- [152] A. Hyakutake, K. Ozaki, K. M. Kitani, and H. Koike. 3-D Interaction with a Large Wall Display Using Transparent Markers. In *Proceedings of the International Conference on Advanced Visual Interfaces* (Roma, Italy), AVI'10, pages 97–100, New York, NY, USA, 2010. ACM.
- [153] IDC. Nearly 18 Million Media Tablets Shipped in 2010. Press Release, 2011. <http://www.idc.com/about/viewpressrelease.jsp?container-Id=prUS22737611>, last accessed February 13, 2013.
- [154] S. S. Intille. Change Blind Information Display for Ubiquitous Computing Environments. In *Proceedings of the 4th international Conference on Ubiquitous Computing* (Goteborg, Sweden), UbiComp'02, pages 91–106, London, UK, UK, 2002. Springer-Verlag.

- [155] H. Ishii and B. Ullmer. Tangible Bits: Towards Seamless Interfaces Between People, Bits, and Atoms. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, GA, USA), CHI'97, pages 234–241, New York, NY, USA, 1997. ACM.
- [156] L. Itti and P. Baldi. Bayesian Surprise Attracts Human Attention. *Vision Research*, 49(10):1295 – 1306, 2009.
- [157] L. Itti and C. Koch. Computational Modelling of Visual Attention. *Nature Reviews Neuroscience*, 2(3):194–203, 2001.
- [158] L. Itti, C. Koch, and E. Niebur. A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, Nov. 1998.
- [159] ITU. Measuring the Information Society. The ICT Development Index. Switzerland International Telecommunication Union, 2009. <http://www.itu.int/net/pressoffice/backgrounders/general/pdf/5.pdf>, last accessed February 13, 2013.
- [160] S. Iyengar and M. Lepper. Rethinking the Value of Choice: A Cultural Perspective on Intrinsic Motivation. *Journal of Personality and Social Psychology*, 76(3):349, 1999.
- [161] S. Izadi, H. Brignull, T. Rodden, Y. Rogers, and M. Underwood. Dynamo: A Public Interactive Surface Supporting the Cooperative Sharing and Exchange of Media. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology* (Vancouver, BC, Canada), UIST'03, pages 159–168, New York, NY, USA, 2003. ACM.
- [162] C. Izard. *Human Emotions*. Emotions, Personality, and Psychotherapy. Plenum Press, 1977.
- [163] R. J. K. Jacob. What You Look at is What You Get: Eye Movement-based Interaction Techniques. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, WA, USA), CHI'90, pages 11–18, New York, NY, USA, 1990. ACM.
- [164] G. Jacucci, A. Morrison, G. T. Richard, J. Kleimola, P. Peltonen, L. Parisi, and T. Laitinen. Worlds of Information: Designing for Engagement at a Public Multi-Touch Display. In *Proceedings of the 28th international Conference on Human Factors in Computing Systems* (Atlanta, GA, USA), CHI'10, pages 2267–2276, New York, NY, USA, 2010. ACM.

- [165] M. Jeannerod. The Mechanism of Self-Recognition in Humans. *Behavioural Brain Research*, 142(1-2):1–15, 2003.
- [166] J. Jonides and S. Yantis. Uniqueness of Abrupt Visual Onset in Capturing Attention. *Attention, Perception, & Psychophysics*, 43(4):346–354, 1988.
- [167] R. José and J. C. S. Cardoso. Opportunities and Challenges of Interactive Public Displays as an Advertising Medium. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 139–157. Springer Limited London, 2011.
- [168] L. Joyner and J. TerKeurst. Accounting for User Needs and Motivations in Game Design. *Interactive Convergence: Critical Issues in Multimedia*, 10:31–44, 2008.
- [169] W. Ju and D. Sirkin. Animate Objects: How Physical Motion Encourages Public Interaction. In *Proceedings of the 5th international Conference on Persuasive Technology* (Copenhagen, Denmark), PERSUASIVE'10, pages 40–51, Berlin, Heidelberg, 2010. Springer-Verlag.
- [170] S. Jigger. The Power of Packaging. *Admap Magazine*, 1999.
- [171] M. A. Just and P. A. Carpenter. Eye Fixations and Cognitive Processes. *Cognitive Psychology*, 8(4):441–480, 1976.
- [172] D. Kahneman. *Attention and Effort*. Prentice-Hall. Prentice-Hall, 1973.
- [173] M. Karam, T. Payne, and E. David. Evaluating Bluscreen: Usability for Intelligent Pervasive Displays. In *2nd International Conference on Pervasive Computing and Applications* (Birmingham, United Kingdom), ICPCA'07, pages 18–23. IEEE, 2007.
- [174] B. Kaufmann, M. Gratzner, and M. Hitz. 3MF – A Service-Oriented Mobile Multimodal Interaction Framework. In *Proceedings of the 3rd Workshop on Infrastructure and Design Challenges of Coupled Display Visual Interfaces PPD'12*, 2012.
- [175] P. Kazienko and M. Adamski. AdROSA-Adaptive Personalization of Web Advertising. *Information Sciences*, 177(11):2269–2295, June 2007.
- [176] M. Kellar, C. Watters, and J. Duffy. Motivational Factors in Game Play in Two User Groups. In *Proceedings of the Fourth International Conference of the Digital Games Research Association* (Tokyo, Japan), DIGRA'07, 2007.

- [177] J. Kendall and K. Kendall. Metaphors and Methodologies: Living Beyond the Systems Machine. *MIS Quarterly*, 17(2):149–171, 1993.
- [178] D. Kern, M. Harding, O. Storz, N. Davis, and A. Schmidt. Shaping How Advertisers See me: User Views on Implicit and Explicit Profile Capture. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy), CHI EA'08, pages 3363–3368, New York, NY, USA, 2008. ACM.
- [179] A. Khan, J. Matejka, G. Fitzmaurice, and G. Kurtenbach. Spotlight: Directing Users' Attention on Large Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, OR, USA), CHI'05, pages 791–798, New York, NY, USA, 2005. ACM.
- [180] A. Klami, C. Saunders, T. E. de Campos, and S. Kaski. Can Relevance of Images be Inferred from Eye Movements? In *Proceeding of the 1st ACM international Conference on Multimedia Information Retrieval* (Vancouver, British Columbia, Canada), MIR'08, pages 134–140, New York, NY, USA, 2008. ACM.
- [181] T. Kollmann. *Akzeptanz innovativer Nutzungsgüter und -systeme: Konsequenzen fuer die Einführung von Telekommunikations- und Multimediasystemen*. Dr. Th. Gabler Verlag, Wiesbaden, 1998.
- [182] P. Kotler. *Marketing Management*. Pearson one series. Pearson Prentice Hall, Upper Saddle River, NJ, 13th edition, 2009.
- [183] P. Kotler, G. Armstrong, J. Saunders, and V. Wong. *Grundlagen des Marketing*. wi wirtschaft. Pearson Studium, 2010.
- [184] L. T. Kozlowski and J. E. Cutting. Recognizing the Sex of a Walker from a Dynamic Point-Light Display. *Perception Psychophysics*, 21(6):575–580, 1977.
- [185] C. Kray, G. Kortuem, and A. Krüger. Adaptive Navigation Support With Public Displays. In *Proceedings of the 10th international Conference on Intelligent User Interfaces* (San Diego, CA, USA), IUI'05, pages 326–328, New York, NY, USA, 2005. ACM.
- [186] W. Kroeber-Riel, P. Weinberg, and A. Gröppel-Klein. *Konsumentenverhalten*. Vahlen, 9 edition, 2008.
- [187] M. W. Krueger. *Artificial reality*. Addison-Wesley, 1983.

- [188] M. W. Krueger. *Artificial reality II*. Addison-Wesley, 1991.
- [189] M. W. Krueger, T. Gionfriddo, and K. Hinrichsen. VIDEOPLACE – An Artificial Reality. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Francisco, CA, United States), CHI'85, pages 35–40, New York, NY, USA, 1985. ACM.
- [190] H. Krugman. The Measurement of Advertising Involvement. *Public Opinion Quarterly*, 30(4):583–596, 1966.
- [191] C. Küblbeck and A. Ernst. Face Detection and Tracking in Video Sequences Using the Modifiedcensus Transformation. *Image and Vision Computing*, 24(6):564–572, 2006.
- [192] C. Küblbeck, T. Ruf, and A. Ernst. A Modular Framework to Detect and Analyze Faces for Audience Measurement Systems. In *Proceedings of the 2nd International Workshop on Pervasive Advertising* (Lübeck, Germany), PerAd'10, 2010.
- [193] B. Kules, H. Kang, C. Plaisant, A. Rose, and B. Shneiderman. Immediate Usability: A Case Study of Public Access Design for a Community Photo Library. *Interacting with Computers*, 16(6):1171 – 1193, 2004.
- [194] M. Kumar, T. Garfinkel, D. Boneh, and T. Winograd. Reducing Shoulder-Surfing by Using Gaze-based Password Entry. In *Proceedings of the 3rd Symposium on Usable Privacy and Security* (Pittsburgh, PA, USA), SOUPS'07, pages 13–19, New York, NY, USA, 2007. ACM.
- [195] E. Kurdyukova, M. Redlin, and E. André. Studying User-defined iPad Gestures for Interaction in Multi-Display Environment. In *Proceedings of the 2012 ACM international Conference on Intelligent User Interfaces* (Lisbon, Portugal), IUI'12, pages 93–96, New York, NY, USA, 2012. ACM.
- [196] S. Kurkovsky and K. Harihar. Using Ubiquitous Computing in Interactive Mobile Marketing. *Personal Ubiquitous Computing*, 10(4):227–240, Mar. 2006.
- [197] A. Kuss and T. Tomczak. *Käuferverhalten: Eine marketingorientierte Einführung*. Grundwissen der Ökonomik. Lucius & Lucius, 2000.
- [198] Kyodo. Facial ID System Muddled to Speed up Airport Immigration. The Japan Times online, 2012. <http://www.japantimes.co.jp/text/nn20120318a2.html>, last accessed February 13, 2013.

- [199] J. R. Landis and G. G. Koch. The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33(1):159–174, 1977.
- [200] M. Langheinrich, N. Memarovic, I. Elhart, and F. Alt. Autopoiesic Content: A Conceptual Model for Enabling Situated Self-generative Content for Public Displays. In *Proceedings of the First Workshop on Pervasive Urban Applications* (San Francisco, CA, USA), PURBA'11, 2011.
- [201] R. Lawrence, G. Almasi, V. Kotlyar, M. Viveros, and S. Duri. Personalization of Supermarket Product Recommendations. *Data Mining and Knowledge Discovery*, 5(1):11–32, 2001.
- [202] J. Lazar, J. Feng, and H. Hochheiser. *Research Methods in Human-Computer Interaction*. John Wiley & Sons Inc, 2009.
- [203] J. C. Lee, S. E. Hudson, and E. Tse. Foldable Interactive Displays. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (Monterey, CA, USA), UIST'08, pages 287–290, New York, NY, USA, 2008. ACM.
- [204] J. Leikas, H. Stromberg, V. Ikonen, R. Suomela, and J. Heinila. 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* (Pisa, Italy), PERCOM'06, pages 66–70, Washington, DC, USA, 2006. IEEE Computer Society.
- [205] D. Li, J. Babcock, and D. J. Parkhurst. openEyes: A Low-cost Head-mounted Eye-Tracking Solution. In *Proceedings of the 2006 Symposium on Eye Tracking Research & Applications* (San Diego, CA, USA), ETRA'06, pages 95–100, New York, NY, USA, 2006. ACM.
- [206] G. L. Lilien, P. Kotler, and K. S. Moorthy. *Marketing Models*. Prentice-Hall, 1992.
- [207] J.-Y. Lin, Y.-Y. Chen, J.-C. Ko, H. Kao, W.-H. Chen, T.-H. Tsai, S.-C. Hsu, and Y.-P. Hung. i-m-Tube: An Interactive Multi-Resolution Tubular Display. In *Proceedings of the 17th ACM International Conference on Multimedia* (Beijing, China), MM'09, pages 253–260, New York, NY, USA, 2009. ACM.
- [208] G. Linden, B. Smith, and J. York. Amazon. com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing*, 7(1):76–80, 2003.

- [209] P. Ljungstrand, S. Björk, and J. Falk. The WearBoy: A Platform for Low-Cost Public Wearable Devices. In *Proceedings of the 3rd IEEE International Symposium on Wearable Computers ISWC'99*, pages 195–196, Washington, DC, USA, 1999. IEEE Computer Society.
- [210] H. Mahato, D. Kern, P. Holleis, and A. Schmidt. Implicit Personalization of Public Environments Using Bluetooth. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy), CHI EA'08, pages 3093–3098, New York, NY, USA, 2008. ACM.
- [211] S. Malik, A. Ranjan, and R. Balakrishnan. Interacting with Large Displays From a Distance with Vision-tracked Multi-Finger Gestural Input. In *Proceedings of the 18th annual ACM Symposium on User Interface Software and Technology* (Seattle, WA, USA), UIST'05, pages 43–52, New York, NY, USA, 2005. ACM.
- [212] T. Malone. Toward a Theory of Intrinsically Motivating Instruction. *Cognitive Science*, 5(4):333–369, 1981.
- [213] T. Malone and M. Lepper. Making Learning Fun: A Taxonomy of Intrinsic Motivations for Learning. *Aptitude Learning and Instruction*, 3(3):223–253, 1987.
- [214] J. Mankoff, A. K. Dey, G. Hsieh, J. Kientz, S. Lederer, and M. Ames. Heuristic Evaluation of Ambient Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, FL, USA), CHI'03, pages 169–176, New York, NY, USA, 2003. ACM.
- [215] L. Manovich. *The Language of New Media*. MIT Press, Cambridge, Massachusetts, 2002.
- [216] P. Marshall, R. Morris, Y. Rogers, S. Kreitmayer, and M. Davies. Rethinking 'Multi-User': An In-the-Wild Study of How Groups Approach a Walk-Up-and-Use Tabletop Interface. In *Proceedings of the 2011 annual Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada), CHI'11, pages 3033–3042, New York, NY, USA, 2011. ACM.
- [217] Y. Matsumoto and A. Zelinsky. An Algorithm for Real-Time Stereo Vision Implementation of Head Pose and Gaze Direction Measurement. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition* (Grenoble, France), FG'00, Washington, DC, USA, 2000. IEEE Computer Society.

- [218] T. L. Matthews, T. L. Rattenbury, and S. A. Carte. Defining, Designing, and Evaluating Peripheral Displays: An Analysis Using Activity Theory. *Human-Computer Interaction*, 22(1-2):221–261, 2007.
- [219] J. McCarthy. Using Public Displays to Create Conversation Opportunities. In *Proceedings of the Workshop on Public, Community, and Situated Displays*, New Orleans, LA, USA, 2002.
- [220] J. McCarthy, J. Müller, K. Cheverst, D. Fitton, N. Taylor, O. Paczkowski, and A. Krüger. Experiences of Supporting Local and Remote Mobile Phone Interaction in Situated Public Display Deployments. *International Journal of Mobile Human-Computer Interaction*, 1(2):1–21, 2009.
- [221] J. F. McCarthy, B. Congleton, and F. M. Harper. The Context, Content & Community Collage: Sharing Personal Digital Media in the Physical Workplace. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (San Diego, CA, USA)*, CSCW’08, pages 97–106, New York, NY, USA, 2008. ACM.
- [222] J. F. McCarthy, T. J. Costa, and E. S. Liongosari. UniCast, OutCast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays. In *Proceedings of the 3rd International Conference on Ubiquitous Computing (Atlanta, GA, USA)*, UbiComp’01, pages 332–345, London, United Kingdom, 2001. Springer-Verlag.
- [223] J. F. McCarthy, S. D. Farnham, Y. Patel, S. Ahuja, D. Norman, W. R. Hazlewood, and J. Lind. Supporting Community in Third Places With Situated Social Software. In *Proceedings of the Fourth International Conference on Communities and Technologies (University Park, PA, USA)*, C&T ’09, pages 225–234, New York, NY, USA, 2009. ACM.
- [224] J. F. McCarthy, D. W. McDonald, S. Soroczak, D. H. Nguyen, and A. M. Rashid. Augmenting the Social Space of an Academic Conference. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (Chicago, IL, USA)*, CSCW’04, pages 39–48, New York, NY, USA, 2004. ACM.
- [225] M. McLuhan. *Understanding Media: The Extensions of Man*. McGraw-Hill, New York, 1964.
- [226] H. Meffert, C. Burmann, and M. Kirchgeorg. *Marketing: Grundlagen Marktorientierter Unternehmensführung. Konzepte-Instrumente-Praxisbeispiele*. Gabler, 2011.

- [227] C. Mello-Thoms, C. F. Nodine, and H. L. Kundel. What Attracts the Eye to the Location of Missed and Reported Breast Cancers? In *Proceedings of the 2002 Symposium on Eye tracking Research & Applications* (New Orleans, LA, USA), ETRA'02, pages 111–117, New York, NY, USA, 2002. ACM.
- [228] N. Memarovic, I. Elhart, and M. Langheinrich. FunSquare: First Experiences with Autopoiesic Content. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia* (Beijing, China), MUM'11, pages 175–184, New York, NY, USA, 2011. ACM.
- [229] N. Memarovic, M. Langheinrich, and F. Alt. Connecting People through Content - Promoting Community Identity Cognition through People and Places. (Prato, Italy), In *Proceedings of Community Informatics*, 2011.
- [230] N. Memarovic, M. Langheinrich, and F. Alt. Interacting Places – A Framework for Promoting Community Interaction and Place Awareness Through Public Displays. In *Tenth Annual IEEE International Conference on Pervasive Computing and Communications – Workshop Proceedings* (Lugano, Switzerland), PerCom'12, 2012.
- [231] N. Memarovic, M. Langheinrich, and F. Alt. The Interacting Places Framework: Conceptualizing Public Display Applications that Promote Community Interaction and Place Awareness. In *Proceedings of the 2012 International Symposium on Pervasive Displays* (Porto, Portugal), PerDis'12, pages 71–76, New York, NY, USA, 2012. ACM.
- [232] N. Memarovic, M. Langheinrich, F. Alt, I. Elhart, S. Hosio, and E. Rubegni. Using public displays to stimulate passive engagement, active engagement, and discovery in public spaces. In *In Proceedings of Media Architecture Biennale 2012* (Aarhus, Denmark), MAB'12, New York, NY, US, Nov. 2012. ACM.
- [233] J. Mennenöh, S. Kristes, F. Alt, A. S. Shirazi, A. Schmidt, and H. Schröder. Customer Touchpoints im stationären Einzelhandel – Potenzial von Pervasive Computing. *Marketing Review St. Gallen*, 27(2):37–42, 2010.
- [234] A. Meschtscherjakov, W. Reitberger, M. Lankes, and M. Tscheligi. Enhanced Shopping: a Dynamic Map in a Retail Store. In *Proceedings of the 10th international Conference on Ubiquitous Computing* (Seoul, Korea), UbiComp'08, pages 336–339, New York, NY, USA, 2008. ACM.

- [235] A. Meschtscherjakov, W. Reitberger, T. Mirlacher, H. Huber, and M. Tschelligi. AmIQuin - An Ambient Mannequin for the Shopping Environment. In *Proceedings of the European Conference on Ambient Intelligence*, AmI'09, pages 206–214. Springer-Verlag, Berlin, Heidelberg, 2009.
- [236] D. Michelis. *Interaktive Großbildschirme im öffentlichen Raum: Nutzungsmotive und Gestaltungsregeln*. Gabler, 2009.
- [237] D. Michelis and M. Meckel. Why Do We Want to Interact With Electronic Billboards in Public Space? In *Proceedings of the 1st Workshop on Pervasive Advertising PerAd'09*, Nara, Japan, 2009.
- [238] D. Michelis and J. Müller. The Audience Funnel: Observations of Gesture Based Interaction with Multiple Large Displays in a City Center. *International Journal of Human-Computer Interaction*, 27(6):562–579, 2011.
- [239] D. Michelis and F. Resatsch. Alice Through the Interface – Electronic Mirrors as Human-Computer Interface. In *Proceedings of the 9th Conference on User Interfaces for all* (Königswinter, Germany), ERCIM'06, pages 88–98, Berlin, Heidelberg, 2007. Springer-Verlag.
- [240] M. Mitchell and K. Suggitt. iCapture: Facilitating Spontaneous User-Interaction with Pervasive Displays using Smart Devices. In *Proceedings of the Workshop on Mobile Devices as Pervasive User Interfaces and Interaction Devices PERMID'06*, 2006.
- [241] R. W. Mitchell. Mental Models of Mirror Self Recognition: Two Theories. *New Ideas in Psychology*, 11(3):295–325, 1993.
- [242] K. Miyaoku, S. Higashino, and Y. Tonomura. C-blink: A Hue Difference-based Light Signal Marker for Large Screen Interaction Via Any Mobile Terminal. In *Proceedings of the 17th annual ACM Symposium on User Interface Software and Technology* (Santa Fe, NM, USA), UIST'04, pages 147–156, New York, NY, USA, 2004. ACM.
- [243] W. Moncur and G. Leplâtre. Pictures at the ATM: Exploring the Usability of Multiple Graphical passwords. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, CA, USA), CHI'07, pages 887–894, New York, NY, USA, 2007. ACM.
- [244] O. Moore and A. Anderson. Some Principles for the Design of Clarifying Educational Environments. In D. A. Goslin, editor, *Handbook of Socialization Theory and Research*, pages 571 – 613, New York, NY, USA, 1969. Rand McNally (Chicago).

- [245] O. Mubin, T. Lashina, and E. Loenen. How Not to Become a Buffoon in Front of a Shop Window: A Solution Allowing Natural Head Movement for Interaction with a Public Display. In *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II* (Uppsala, Sweden), INTERACT'09, pages 250–263, Berlin, Heidelberg, 2009. Springer-Verlag.
- [246] F. Müller, S. Agamanolis, M. R. Gibbs, and F. Vetere. Remote Impact: Shadowboxing Over a Distance. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy), CHI EA'08, pages 2291–2296, New York, NY, USA, 2008. ACM.
- [247] F. Müller, S. Agamanolis, and R. Picard. Exertion Interfaces: Sports Over a Distance for Social Bonding and Fun. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, FL, USA), CHI'03, pages 561–568, New York, NY, USA, 2003. ACM.
- [248] J. Müller, F. Alt, and D. Michelis. Introduction to Pervasive Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*. Springer Limited London, 2011.
- [249] J. Müller, F. Alt, and D. Michelis. *Pervasive Advertising*. Springer London Limited, 2011.
- [250] J. Müller, F. Alt, D. Michelis, and A. Schmidt. Requirements and Design Space for Interactive Public Displays. In *Proceedings of the International Conference on Multimedia* (Firenze, Italy), MM'10, pages 1285–1294, New York, NY, USA, 2010. ACM.
- [251] J. Müller, J. Exeler, M. Buzbeck, and A. Krüger. ReflectiveSigns: Digital Signs That Adapt to Audience Attention. In *Proceedings of the 7th International Conference on Pervasive Computing* (Nara, Japan), Pervasive'09, pages 17–24, Berlin, Heidelberg, 2009. Springer-Verlag.
- [252] J. Müller and A. Krüger. Towards Situated Public Displays as Multicast Systems. In *Proceedings of the Workshop on Ubiquitous User Modeling* UbiqUM'06, 2006.
- [253] J. Müller and A. Krüger. MobiDiC: Context Adaptive Digital Signage with Coupons. In *Proceedings of the European Conference on Ambient Intelligence* (Salzburg, Austria), AmI'09, pages 24–33, Berlin, Heidelberg, 2009. Springer-Verlag.

- [254] J. Müller, R. Walter, G. Bailly, M. Nischt, and F. Alt. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In *Proceedings of the 2012 ACM Conference on Human Factors in Computing Systems* (Austin, TX, USA), CHI'12, pages 297–306, New York, NY, USA, 2012. ACM.
- [255] J. Müller, D. Wilmsmann, J. Exeler, M. Buzeck, A. Schmidt, T. Jay, and A. Krüger. Display Blindness: The Effect of Expectations on Attention towards Digital Signage. In *Proceedings of the 7th International Conference on Pervasive Computing* (Nara, Japan), Pervasive'09, pages 1–8, Berlin, Heidelberg, 2009. Springer-Verlag.
- [256] M. A. Nakamura. Creating a New Channel For Campus Communication. In *Proceedings of the 32nd annual ACM SIGUCCS Fall Conference* (Baltimore, MD, USA), SIGUCCS'04, pages 56–59, New York, NY, USA, 2004. ACM.
- [257] T. Nawaz, M. Mian, and H. Habib. Infotainment Devices Control by Eye Gaze and Gesture Recognition Fusion. *IEEE Transactions on Consumer Electronics*, 54(2):277–282, 2008.
- [258] R. Newman, Y. Matsumoto, S. Rougeaux, and A. Zelinsky. Real-Time Stereo Tracking for Head Pose and Gaze Estimation. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition* (Grenoble, France), FG'00, Washington, DC, USA, 2000. IEEE Computer Society.
- [259] D. Nguyen, J. Tullio, T. Drewes, and E. Mynatt. Dynamic Door Displays. <http://smartech.gatech.edu/jspui/bitstream/1853/3462/1/00-30.pdf>, last accessed February 13, 2013, 2000.
- [260] J. Nielsen. *Usability Engineering*. Morgan Kaufmann, 1994.
- [261] D. A. Norman. Affordance, Conventions, and Design. *Interactions*, 6:38–43, May 1999.
- [262] D. A. Norman. *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books, 1st edition, 2003.
- [263] D. A. Norman. The Way I See It: Signifiers, Not Affordances. *Interactions*, 15:18–19, November 2008.
- [264] T. P. Novak and D. L. Hoffman. The Fit of Thinking Style and Situation: New Measures of Situation Specific Experiential and Rational Cognition. *Journal of Consumer Research*, 36(1):56–72, 2009.

- [265] U. of Michigan. Stadium history. *Michigan Stadium Renovation*, 2013.
- [266] K. O'Hara, R. Harper, A. Unger, J. Wilkes, B. Sharpe, and M. Jansen. TxtBoard: From Text-to-Person to Text-to-Home. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA), CHI EA'05, pages 1705–1708, New York, NY, USA, 2005. ACM.
- [267] K. O'Hara, M. Lipson, M. Jansen, A. Unger, H. Jeffries, and P. Macer. Jukola: Democratic Music Choice in a Public Space. In *Proceedings of the 5th Conference on Designing interactive Systems* (Cambridge, MA, USA), DIS'04, pages 145–154, New York, NY, USA, 2004. ACM.
- [268] K. O'Hara, M. Perry, and S. Lewis. *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*. Kluwer, 2003.
- [269] T. Ojala, J. Korhonen, M. Aittola, M. Ollila, T. Koivumäki, J. Tähtinen, and H. Karjaluo. SmartRotuaari – Context-aware Mobile Multimedia Services. In *Proceedings of the 2nd International Conference on Mobile and Ubiquitous Multimedia* (Norrköping, Sweden), MUM'03, Norrköping, Sweden, 2003.
- [270] T. Ojala and V. Kostakos. UBI challenge: Research Competition on Real-World Urban Computing. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia* (Beijing, China), MUM'11, pages 205–208, New York, NY, USA, 2011. ACM.
- [271] T. Ojala, V. Kostakos, and H. Kukka. It's a Jungle Out There: Fantasy and Reality of Evaluating Public Displays in the Wild. *Proceedings of the First Workshop on Large Displays in Urban Life*, 4:1–4, 2011.
- [272] T. Ojala, H. Kukka, T. Lindén, T. Heikkinen, M. Jurmu, S. Hosio, and F. Kruger. UBI-Hotspot 1.0: Large-Scale Long-Term Deployment of Interactive Public Displays in a City Center. In *Proceedings of the 2010 Fifth International Conference on Internet and Web Applications and Services* (Barcelona, Spain), ICIW'10, pages 285–294, Washington, DC, USA, 2010. IEEE Computer Society.
- [273] R. Oppermann. User-Adaptive to Context-Adaptive Information Systems. *i-com*, 4(3):4–14, 2005.
- [274] Organized Shopping LLC. Definition of Conversion Rate, 2012. http://www.marketingterms.com/dictionary/conversion_rate/, last accessed February 13, 2013.

- [275] M. Ostkamp, G. Bauer, and C. Kray. Visual Highlighting on Public Displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays* (Porto, Portugal), PerDis'12, pages 21–26, New York, NY, USA, 2012. ACM.
- [276] T. Paek, M. Agrawala, S. Basu, S. Drucker, T. Kristjansson, R. Logan, K. Toyama, and A. Wilson. Toward Universal Mobile Interaction for Shared Displays. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, IL, USA), CSCW '04, pages 266–269, New York, NY, USA, 2004. ACM.
- [277] B. Paras and J. Bizzocchi. Game, Motivation, and Effective Learning: An Integrated Model for Educational Game Design. In *Proceedings of the Second International Conference of the Digital Games Research Association* (Vancouver, BC, Canada), DIGRA'05, 2005.
- [278] K. Partridge and B. Begole. Activity-Based Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 83–101. Springer Limited London, 2011.
- [279] A. O. PATRICK. Mass of messages lands at heathrow - new terminal to sport hundreds of ad screens to generate revenue. *The Wall Street Journal*, 2008.
- [280] T. Payne, E. David, N. R. Jennings, and M. Sharifi. Auction Mechanisms for Efficient Advertisement Selection on Public Displays. In *Proceedings of 17th European Conference on Artificial Intelligence ECAI'06*, pages 285–289, Amsterdam, The Netherlands, 2006. IOS Press.
- [281] P. Peltonen, E. Kurvinen, A. Salovaara, G. Jacucci, T. Ilmonen, J. Evans, A. Oulasvirta, and P. Saarikko. It's Mine, Don't Touch!: Interactions at a Large Multi-Touch Display in a City Centre. In *Proceeding of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy), CHI'08, pages 1285–1294, New York, NY, USA, 2008. ACM.
- [282] P. Peltonen, A. Salovaara, G. Jacucci, T. Ilmonen, C. Ardito, P. Saarikko, and V. Batra. Extending Large-Scale Event Participation with User-Created Mobile Media on a Public Display. In *Proceedings of the 6th international Conference on Mobile and Ubiquitous Multimedia* (Oulu, Finland), MUM'07, pages 131–138, New York, NY, USA, 2007. ACM.
- [283] R. Petty and J. Cacioppo. The Elaboration Likelihood Model of Persuasion. *Advances in Experimental Social Psychology*, 19(1):123–205, 1986.

- [284] R. Petty, J. Cacioppo, and D. Schumann. Central and Peripheral Routes to Advertising Effectiveness: The Moderating Role of Involvement. *Journal of Consumer Research*, 10(2):135–146, 1983.
- [285] C. Pinhanez. Using a Steerable Projector and a Camera to Transform Surfaces into Interactive Displays. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, WA, USA), CHI EA'01, pages 369–370, New York, NY, USA, 2001. ACM.
- [286] C. Pinhanez and M. Podlaseck. To Frame or Not to Frame: The Role and Design of Frameless Displays in Ubiquitous Applications. In *Proceedings of the 7th international Conference on Ubiquitous Computing* (Tokyo, Japan), UbiComp'05, pages 340–357, Berlin, Heidelberg, 2005. Springer-Verlag.
- [287] A. Poole and L. J. Ball. Eye Tracking in Human-Computer Interaction and Usability Research: Current Status and Future. In C. Ghaoui, editor, *Encyclopedia of Human-Computer Interaction*, Pennsylvania, PA, USA, 2005. Idea Group, Inc.
- [288] A. Poole, L. J. Ball, and P. Phillips. In Search of Saliency: A Response-Time and Eye-Movement Analysis of Bookmark Recognition. In *Proceedings of the 18th British HCI Group Annual Conference on HCI* (Leeds, United Kingdom), HCI'04, pages 363–378, Berlin, Heidelberg, 2004. Springer-Verlag.
- [289] Poynter Institute & Eyetools Inc. Eyetrack III: Online News Consumer Behavior in the Age of Multimedias, 2004. <http://www.poynterextra.org/eyetrack2004/index.htm>, last accessed February 13, 2013.
- [290] T. Prante, C. Röcker, N. Streitz, R. Stenzel, C. Magerkurth, D. van Alphen, and D. Plewe. Hello.Wall – Beyond Ambient Displays. In *Adjunct Proceedings of the 5th International Conference on Ubiquitous Computing* (Seattle, WA, USA), Ubicomp'03, pages 277–278, 2003.
- [291] N. M. Puccinelli, R. C. Goodstein, D. Grewal, R. Price, P. Raghuram, and D. Stewart. Customer Experience Management in Retailing: Understanding the Buying Process. *Journal of Retailing*, 85(1):15–30, 2009.
- [292] P. Qvarfordt and S. Zhai. Conversing With the User-based on Eye-Gaze Patterns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, OR, USA), CHI'05, pages 221–230, New York, NY, USA, 2005. ACM.

- [293] A. Ranganathan and R. H. Campbell. Advertising in a Pervasive Computing Environment. In *Proceedings of the 2nd International Workshop on Mobile Commerce* (Atlanta, GA, USA), WMC'02, pages 10–14, New York, NY, USA, 2002. ACM.
- [294] O. Ratsimor, T. Finin, A. Joshi, and Y. Yesha. eNcentive: A Framework for Intelligent Marketing in Mobile Peer-to-Peer Environments. In *Proceedings of the 5th International Conference on Electronic Commerce* (Pittsburgh, PA, USA), ICEC'03, pages 87–94, New York, NY, USA, 2003. ACM.
- [295] F. Redhead and M. Brereton. A Qualitative Analysis of Local Community Communications. In *Proceedings of the 18th Australia Conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments* (Sydney, Australia), OZCHI'06, pages 361–364, New York, NY, USA, 2006. ACM.
- [296] F. Redhead and M. Brereton. Designing Interaction for Local Communications: An Urban Screen Study. In *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II* (Uppsala, Sweden), INTERACT'09, pages 457–460, Berlin, Heidelberg, 2009. Springer-Verlag.
- [297] J. Redström, T. Skog, and L. Hallnäs. Informative Art: Using Amplified Artworks as Information Displays. In *Proceedings of Designing Augmented Reality Environments* (Elsinore, Denmark), DARE'00, pages 103–114, New York, NY, USA, 2000. ACM.
- [298] R. W. Reeder, P. Pirolli, and S. K. Card. WebEyeMapper and WebLogger: Tools for Analyzing Eye Tracking Data Collected in Web-use Studies. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, WA, USA), CHI EA'01, pages 19–20, New York, NY, USA, 2001. ACM.
- [299] S. Reeves, S. Benford, C. O'Malley, and M. Fraser. Designing the Spectator Experience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, OR, USA), CHI'05, pages 741–750, New York, NY, USA, 2005. ACM.
- [300] J. Reid, M. Lipson, J. Hyams, and K. Shaw. “Fancy a schmink?": A Novel Networked Game in a Cafe. In *Proceedings of the 2004 ACM SIGCHI International Conference on Advances in Computer Entertainment*

- Technology* (Singapore), ACE'04, pages 18–23, New York, NY, USA, 2004. ACM.
- [301] W. Reitberger, A. Meschtscherjakov, T. Mirlacher, T. Scherndl, H. Huber, and M. Tscheligi. A Persuasive Interactive Mannequin for Shop Windows. In *Proceedings of the 4th International Conference on Persuasive Technology* (Claremont, CA, USA), Persuasive'09, pages 41–48, New York, NY, USA, 2009. ACM.
- [302] W. Reitberger, C. Obermair, B. Ploderer, A. Meschtscherjakov, and M. Tscheligi. Enhancing the Shopping Experience with Ambient Displays: A Field Study in a Retail Store. In *Proceedings of the 2007 European Conference on Ambient intelligence* (Darmstadt, Germany), Aml'07, pages 314–331, Berlin, Heidelberg, 2007. Springer-Verlag.
- [303] D. Ringas, E. Christopoulou, and M. Stefanidakis. Clio: blending the collective memory with the urban landscape. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia* (Beijing, China), MUM '11, pages 185–194, New York, NY, USA, 2011. ACM.
- [304] R. Rodenstein. Employing the Periphery: The Window as Interface. In *CHI'99 Extended Abstracts on Human Factors in Computing Systems* (Pittsburgh, PA, USA), CHI EA'99, pages 204–205, New York, NY, USA, 1999. ACM.
- [305] Y. Rogers. Moving on from Weiser's Vision of Calm Computing: Engaging Ubicomp Experiences. In *Proceedings of the 8th international Conference on Ubiquitous Computing* (Orange County, CA, USA), UbiComp'06, pages 404–421, Berlin, Heidelberg, 2006. Springer-Verlag.
- [306] Y. Rogers, W. R. Hazlewood, P. Marshall, N. Dalton, and S. Hertrich. Ambient Influence: Can Twinkly Lights Lure and Abstract Representations Trigger Behavioral Change? In *Proceedings of the 12th ACM international Conference on Ubiquitous Computing* (Copenhagen, Denmark), UbiComp '10, pages 261–270, New York, NY, USA, 2010. ACM.
- [307] M. Rohs and B. Gfeller. Using Camera-Equipped Mobile Phones for Interacting with Real-World Objects. In *Advances in Pervasive Computing* volume 176, pages 265–271, 2004.
- [308] G. Roussos. *Ubiquitous and Pervasive Commerce: New Frontiers for Electronic Business*. Springer SMB, 2005.

- [309] E. Rukzio, M. Müller, and R. Hardy. Design, Implementation and Evaluation of a Novel Public Display for Pedestrian Navigation: The Rotating Compass. In *Proceedings of the 27th international Conference on Human Factors in Computing Systems* (Boston, MA, USA), CHI'09, pages 113–122, New York, NY, USA, 2009. ACM.
- [310] E. Rukzio, A. Schmidt, and H. Hussmann. An Analysis of the Usage of Mobile Phones for Personalized Interactions with Ubiquitous Public Displays. In *Proceedings of the 1st Workshop on Ubiquitous Display Environments* (Nottingham, UK), UbiDisplay'04, 2004.
- [311] E. Rukzio, A. Schmidt, and H. Hussmann. Physical Posters as Gateways to Context-Aware Services for Mobile Devices. In *Proceedings of the 6th IEEE Workshop on Mobile Computing Systems and Applications WMCSA'04*, pages 10–19, Washington, DC, USA, 2004. IEEE Computer Society.
- [312] D. M. Russell, J. P. Trimble, and A. Dieberger. The Use Patterns of Large, Interactive Display Surfaces: Case Studies of Media Design and Use for BlueBoard and MERBoard. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences HICSS'04*, Washington, DC, USA, 2004. IEEE Computer Society.
- [313] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, and M. Cohen. Gaze-based Interaction for Semi-Automatic Photo Cropping. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montreal, Quebec, Canada), CHI'06, pages 771–780, New York, NY, USA, 2006. ACM.
- [314] N. Sawhney, S. Wheeler, and C. Schmandt. Aware Community Portals: Shared Information Appliances for Transitional Spaces. *Personal and Ubiquitous Computing*, 5(1):66–70, Jan. 2001.
- [315] J. Scheible and T. Ojala. MobiLenin – Combining a Multi-Track Music Video, Personal Mobile Phones and a Public Display into Multi-User Interactive Entertainment. In *Proceedings of the 13th annual ACM International Conference on Multimedia* (Singapore), MM'05, pages 199–208, New York, NY, USA, 2005. ACM.
- [316] B. Schmid and B. Lyczek. *Unternehmenskommunikation*. Gabler, Wiesbaden, 1st edition, 2007.
- [317] A. Schmidt. Implicit Human Computer Interaction Through Context. *Personal and Ubiquitous Computing*, 4(2/3):191–199, 2000.

- [318] A. Schmidt, F. Alt, P. Holleis, J. Mueller, and A. Krueger. Creating Log Files and Click Streams for Advertisements in Physical Space. In *Adjunct Proceedings of the 10th International Conference on Ubiquitous Computing Ubicomp'08*, pages 28–29, Seoul, South Korea, 2008.
- [319] A. Schmidt, M. Langheinrich, and K. Kersting. Perception Beyond the Here and Now. *IEEE Computer*, 44(2):86–88, Feb. 2011.
- [320] A. Schmidt and K. Van Laerhoven. How to Build Smart Appliances? *Personal Communications, IEEE*, 8(4):66–71, 2001.
- [321] D. Schmidt, F. Chehimi, E. Rukzio, and H. Gellersen. PhoneTouch: A Technique for Direct Phone Interaction on Surfaces. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA), UIST'10, pages 13–16, New York, NY, USA, 2010. ACM.
- [322] B. J. Scholl and P. D. Tremoulet. Perceptual Causality and Animacy. *Trends in Cognitive Sciences*, 4(8):299 – 309, 2000.
- [323] J. Schönböck, F. König, G. Kotsis, D. Gruber, E. Zaim, and A. Schmidt. MirrorBoard - An Interactive Billboard. (Lübeck, Germany), In M. Herzog and M. C. Kindsmüller, editors, *Mensch & Computer*, pages 217–226. Oldenbourg Verlag, 2008.
- [324] J. Schrammel, E. Mattheiss, S. Döbelt, L. Paletta, A. Almer, and M. Tschelligi. Attentional Behavior of Users on the Move Towards Pervasive Advertising Media. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*. Springer Limited London, 2011.
- [325] S. D. Scott, M. Sheelagh, T. Carpendale, and K. M. Inkpen. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, IL, USA), CSCW'04, pages 294–303, New York, NY, USA, 2004. ACM.
- [326] T. Selker. Visual Attentive Interfaces. *BT Technology Journal*, 22(4):146–150, Oct. 2004.
- [327] A. Sellen, R. Harper, R. Eardley, S. Izadi, T. Regan, A. S. Taylor, and K. R. Wood. HomeNote: Supporting Situated Messaging in the Home. In *Proceedings of the 20th Conference on Computer Supported Cooperative Work* (Banff, AB, Canada), CSCW'06, pages 383–392, New York, NY, USA, 2006. ACM.

- [328] R. Shannon, M. Stabeler, A. Quigley, and P. Nixon. Profiling and Targeting Opportunities in Pervasive Advertising. In *Proceedings of the 1st Workshop on Pervasive Advertising* (Nara, Japan), PerAd'09, 2009.
- [329] M. Sharifi, T. Payne, and E. David. Public Display Advertising Based on Bluetooth Device Presence. In *Proceedings of the 1st Workshop on Mobile Interaction with the Real World* (Espoo, Finland), MIRW'06, 2006.
- [330] C. Sharma. Mobile Advertising Metrics: What Matters Most. Gigaom, 2009. <http://gigaom.com/2009/09/07/mobile-advertising-metrics-what-matters-most/>, last accessed February 13, 2013.
- [331] R. Sharp, A. Madhavapeddy, R. Want, and T. Pering. Enhancing Web Browsing Security on Public Terminals Using Mobile Composition. In *Proceedings of the 6th international Conference on Mobile Systems, Applications, and Services* (Breckenridge, CO, USA), MobiSys'08, pages 94–105, New York, NY, USA, 2008. ACM.
- [332] R. Sharp, J. Scott, and A. R. Beresford. Secure Mobile Computing via Public Terminals. In *Proceedings of the 4th international Conference on Pervasive Computing* (Dublin, Ireland), PERVASIVE'06, pages 238–253, Berlin, Heidelberg, 2006. Springer-Verlag.
- [333] M. Sherif and H. Cantril. *The Psychology of Ego-Involvements, Social Attitudes and Identifications*. John Wiley & Sons Inc, London, United Kingdom, 1947.
- [334] R. Shiffrin and W. Schneider. Controlled and Automatic Human Information Processing: II. Perceptual Learning, Automatic Attending and a General Theory. *Psychological Review*, 84(2):127, 1977.
- [335] A. S. Shirazi, C. Winkler, and A. Schmidt. Flashlight Interaction: A Study on Mobile Phone Interaction Techniques with Large Displays. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Bonn, Germany), MobileHCI'09, pages 931–932, New York, NY, USA, 2009. ACM.
- [336] B. Shneiderman. Designing for Fun: How Can We Design User Interfaces to be More Fun? *Interactions*, 11(5):48–50, Sept. 2004.
- [337] G. Shoemaker and K. Inkpen. Single Display Privacyware: Augmenting Public Displays with Private Information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, WA, USA), CHI'01, pages 522–529, New York, NY, USA, 2001. ACM.

- [338] G. Shoemaker, A. Tang, and K. S. Booth. Shadow Reaching: A New Perspective on Interaction for Large Displays. In *Proceedings of the 20th annual ACM Symposium on User Interface Software and Technology* (Newport, RI, USA), UIST'07, pages 53–56, New York, NY, USA, 2007. ACM.
- [339] G. Shoemaker, T. Tsukitani, Y. Kitamura, and K. S. Booth. Body-centric Interaction Techniques for Very Large Wall Displays. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction* (Reykjavik, Iceland), NordiCHI'10, pages 463–472, New York, NY, USA, 2010. ACM.
- [340] L. E. Sibert and R. J. K. Jacob. Evaluation of Eye Gaze Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands), CHI'00, pages 281–288, New York, NY, USA, 2000. ACM.
- [341] G. A. Silberer. Digital Signage im stationären Handel – Das Anwendungs- und Wirkungspotenzial eines neuen POS-Mediums. *Der Markt*, 49(1):3–16, 2010.
- [342] A. Sippl, C. Holzmann, D. Zachhuber, and A. Ferscha. Real-time gaze tracking for public displays. In *Proceedings of the First international joint Conference on Ambient intelligence* (Malaga, Spain), Aml'10, pages 167–176, Berlin, Heidelberg, 2010. Springer-Verlag.
- [343] T. Skog and L. E. Holmquist. WebAware: Continuous Visualization of Website Activity in a Public Space. In *CHI'00 Extended Abstracts on Human Factors in Computing Systems* (The Hague, The Netherlands), CHI EA'00, pages 351–352, New York, NY, USA, 2000. ACM.
- [344] T. Sluckin. About the Nature of Crystallised Liquids and Liquid Crystals. *Bunsen Magazine*, 7, 2005.
- [345] W. R. Smith. Product Differentiation and Market Segmentation as Alternative Marketing Strategies. *Journal of Marketing*, 21(1), 1956.
- [346] D. Snowdon and A. Grasso. Diffusing Information in Organizational Settings: Learning from Experience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Changing our world, changing ourselves* (Minneapolis, MI, USA), CHI'02, pages 331–338, New York, NY, USA, 2002. ACM.

- [347] B. Spethmann. Shoppers Like Grocery Loyalty Cards Despite Privacy Worries. *Promo Magazine*, 2005. http://chiefmarketer.com/incentives/shopper_loyalty_cards_020205, last accessed, February 13, 2013.
- [348] E. D. Spiegler, C. Hildebrand, and F. Michahelles. Social Networks in Pervasive Advertising and Shopping. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 207–225. Springer Limited London, 2011.
- [349] C. Stahl, J. Baus, B. Brandherm, M. Schmitz, and T. Schwartz. Navigational and Shopping Assistance on the Basis of User Interactions in Intelligent Environments. (Colchester, UK), In *IEEE Seminar on Intelligent Building Environments*, pages 182–191, 2005.
- [350] U. Stalder. Digital Out-of-Home Media: Means and Effects of Digital Media in Public Space. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 31–56. Springer Limited London, 2011.
- [351] L. Story. Anywhere the Eye Can See, It’s Likely to See an Ad. *NY Times*, 2007. <http://www.nytimes.com/2007/01/15/business/media/15everywhere.html>, last accessed February 13, 2013.
- [352] O. Storz, A. Friday, and N. Davies. Supporting Content Scheduling on Situated Public Displays. *Computers & Graphics*, 30(5):681–691, 2006.
- [353] O. Storz, A. Friday, N. Davies, J. Finney, C. Sas, and J. Sheridan. Public Ubiquitous Computing Systems: Lessons from the e-Campus Display Deployments. *IEEE Pervasive Computing*, 05(3):40–47, 2006.
- [354] N. Streitz, C. Röcker, T. Prante, R. Stenzel, and D. van Alphen. Situated Interaction With Ambient Information: Facilitating Awareness and Communication in Ubiquitous Work Environments. In *Proceedings of the 10th International Conference on Human-Computer Interaction (Crete, Greece)*, HCI International’03, 2003.
- [355] M. Strohbach, M. Bauer, M. Martin, and B. Hebgén. Managing Advertising Context. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 185–205. Springer Limited London, 2011.
- [356] X. Suo, Y. Zhu, and G. S. Owen. Graphical Passwords: A Survey. In *Proceedings of the 21st Annual Computer Security Applications Conference (Tucson, AZ, USA)*, ACSAC’05, pages 463–472, Washington, DC, USA, 2005. IEEE Computer Society.

- [357] A. Tang, M. Finke, M. Blackstock, R. Leung, M. Deutscher, and R. Lea. Designing for Bystanders: Reflections on Building a Public Digital Forum. In *Proceedings of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy), CHI'08, pages 879–882, New York, NY, USA, 2008. ACM.
- [358] N. Taylor and K. Cheverst. Exploring the Use of Non-Digital Situated Displays in a Rural Community. (Cairns, Australia), In *Proceedings of the 2008 Workshop on Public and Situated Displays to Support Communities*, 2008.
- [359] N. Taylor and K. Cheverst. Rural Communities and Pervasive Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 269–286. Springer Limited London, 2011.
- [360] N. Taylor, K. Cheverst, D. Fitton, N. J. P. Race, M. Rouncefield, and C. Graham. Probing Communities: Study of a Village Photo Display. In *Proceedings of the 19th Australasian Conference on Computer-Human Interaction: Entertaining User Interfaces* (Adelaide, Australia), OZCHI'07, pages 17–24, New York, NY, USA, 2007. ACM.
- [361] S. Telschow and N. Loose. *Digital Signage: die globale Studie; Chancen und Risiken*. Gesellschaft für Innovative Marktforschung, Heidelberg, 2008.
- [362] A. Thawani, S. Gopalan, and V. Sridhar. Context-aware Personalized Ad Insertion in an Interactive TV Environment. In *Proceedings of the 4th Workshop on Personalization in Future TV* TV'04, Eindhoven, The Netherlands, 2004.
- [363] S. Toto. From NTT's Showroom In Tokyo: Next-Generation Digital Signage System, 2010. <http://techcrunch.com/2010/08/23/from-ntts-showroom-in-tokyo-next-generation-digital-signage-system-video>, last accessed February 13, 2013.
- [364] L. Trevino and J. Webster. Flow in Computer-Mediated Communication Electronic Mail and Voice Mail Evaluation and Impacts. *Communication Research*, 19(5):539–573, 1992.
- [365] V. Trommsdorff. *Konsumentenverhalten*. Kohlhammer, Stuttgart, 7th edition, 2009.

- [366] P. van Waart, I. Mulder, and C. de Bont. Meaningful Advertising. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 57–81. Springer Limited London, 2011.
- [367] U. Varshney. Location Management for Mobile Commerce Applications in Wireless Internet Environment. *ACM Transactions on Internet Technology*, 3(3):236–255, Aug. 2003.
- [368] V. Venkatesh and F. D. Davis. A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies. *Management Science*, 46(2):186–204, Feb. 2000.
- [369] R. Vertegaal. Introduction. *Communications of the ACM*, 46(3), 2003.
- [370] R. Vertegaal, J. Shell, D. Chen, and A. Mamuji. Designing for Augmented Attention: Towards a Framework for Attentive User Interfaces. *Computers in Human Behavior*, 22(4):771–789, 2006.
- [371] N. Villar, G. Kortuem, K. Van Laerhoven, and A. Schmidt. The Pendle: A Personal Mediator for Mixed Initiative Environments. (Colchester, United Kingdom), In *IEE International Workshop on Intelligent Environments*, 2005.
- [372] D. Vogel and R. Balakrishnan. 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), UIST’04, pages 137–146, New York, NY, USA, 2004. ACM.
- [373] D. Vogel and R. Balakrishnan. Distant Freehand Pointing and Clicking on Very Large, High Resolution Displays. In *Proceedings of the 18th annual ACM Symposium on User Interface Software and Technology* (Seattle, WA, USA), UIST’05, pages 33–42, New York, NY, USA, 2005. ACM.
- [374] P. Vorderer, T. Hartmann, and C. Klimmt. Explaining the enjoyment of playing video games: The role of competition. In *Proceedings of the second international Conference on Entertainment Computing* (Pittsburgh, PA, USA), ICEC ’03, pages 1–9, Pittsburgh, PA, USA, 2003. Carnegie Mellon University.
- [375] I. Wakeman, A. Light, J. Robinson, D. Chalmers, and A. Basu. Deploying Pervasive Advertising in a Farmers’ Market. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, pages 247–267. Springer Limited London, 2011.

- [376] J. Wang, S. Zhai, and H. Su. Chinese Input with Keyboard and Eye-Tracking: An Anatomical Study. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, WA, USA), CHI'01, pages 349–356, New York, NY, USA, 2001. ACM.
- [377] R. Want, A. Hopper, V. Falcão, and J. Gibbons. The Active Badge Location System. *ACM Transactions on Information Systems*, 10(1):91–102, Jan. 1992.
- [378] M. Weiser. The computer for the 21st century. *Scientific American*, 265(3):94–104, 1991.
- [379] M. Weiser and J. S. Brown. The Coming Age of Calm Technology. In P. J. Denning and R. M. Metcalfe, editors, *Beyond Calculation*, pages 75–85, New York, NY, USA, 1997. Copernicus.
- [380] P. Wellner. Interacting with Paper on the DigitalDesk. *Communications of the ACM*, 36(7):87–96, July 1993.
- [381] S. Wiedenbeck, J. Waters, J.-C. Birget, A. Brodskiy, and N. Memon. PassPoints: Design and Longitudinal Evaluation of a Graphical Password Systems. *International Journal of Human-Computer Studies*, 63(1-2):102–127, 2005.
- [382] A. D. Wilson and H. Benko. Combining Multiple Depth Cameras and Projectors for Interactions On, Above and Between Surfaces. In *Proceedings of the 23rd annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA), UIST'10, pages 273–282, New York, NY, USA, 2010. ACM.
- [383] M. Wilson. Perceiving Imitatable Stimuli: Consequences of Isomorphism Between Input and Output. *Psychological Bulletin*, 127(4):543–553, 2001.
- [384] Y. Yanagisawa, T. Maekawa, Y. Kishino, K. Kamei, Y. Sa-kurai, and T. Okadom. A Relative Positioning Technique with Low-frequency Waves. In *Adjunct Proceedings of the 7th International Conference on Pervasive Computing* (Nara, Japan), Pervasive'09, 2009.
- [385] N. Y. Yankees. Yankee stadium history. yankees.com, 2013.
- [386] J. Zaichkowsky. Measuring the Involvement Construct. *Journal of Consumer Research*, 12(3):341–352, Dec. 1985.

- [387] S. Zhai, C. Morimoto, and S. Ihde. Manual and Gaze Input Cascaded (MAGIC) Pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Pittsburgh, PA, USA), CHI'99, pages 246–253, New York, NY, USA, 1999. ACM.

VIII

APPENDIX

Appendix I: Evaluation of Public Displays

Paper	Stu.	Obj.	Meth.	Part.	Dur.
F. Alt, N. Memarovic, I. Elhart, D. Bial, A. Schmidt, M. Langheinrich, G. Harboe, E. Huang, and M. P. Scipioni. Designing Shared Public Display Networks: Implications from Today's Paper-based Notice Areas. In <i>Proceedings of the 9th International Conference on Pervasive Computing</i> (San Francisco, CA, USA), Pervasive'11, pages 258–275, Berlin, Heidelberg, 2011. Springer-Verlag	ETH	AB	Obs.	50	28
	AU	UA	Int.	10	28
F. Alt, T. Kubitzka, D. Bial, F. Zaidan, M. Ortel, B. Zurmaar, T. Lewen, A. S. Shirazi, and A. Schmidt. Digi-fieds: Insights into Deploying Digital Public Notice Areas in the Wild. In <i>Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia</i> (Beijing, China), MUM'11, pages 165–174, New York, NY, USA, 2011. ACM	FS	AB	Obs.	60	2
	FS	UA	Int.	29	2
	FS	UX	Int.	29	2
	FS	Eff.	Log.	–	90
	FS	UP	Qu.	32	7
	FS	Pri.	Qu.	32	7
R. Ballagas, M. Rohs, and J. G. Sheridan. Sweep and Point and Shoot: Phocem-based Interactions for Large Public Displays. In <i>CHI'05 Extended Abstracts on Human Factors in Computing Systems</i> (Portland, OR, USA), CHI EA '05, pages 1200–1203, New York, NY, USA, 2005. ACM	LS	UP	Log.	10	1
	LS	UA	Qu.	10	1
G. Beyer, F. Alt, J. Müller, A. Schmidt, K. Isakovic, S. Klose, M. Schiewe, and I. Haulsen. Audience Behavior Around Large Interactive Cylindrical Screens. In <i>Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems</i> (Vancouver, BC, Canada), CHI'11, pages 1021–1030, New York, NY, USA, 2011. ACM	LS	AB	Log.	15	2
	LS	AB	Int.	2	2
	LS	AB	Obs.	15	2
	LS	UX	Qu.	15	2
S. A. Bly, S. R. Harrison, and S. Irwin. Media Spaces: Bringing People Together in a Vvideo, Audio, and Computing Environment. <i>Communications of the ACM</i> , 36(1):28–46, Jan. 1993	DBR	AB	Obs.	20	–
	DBR	AB	Log.	20	–
S. Boring, D. Baur, A. Butz, S. Gustafson, and P. Baudisch. Touch Projector: Mobile Interaction Through Video. In <i>Proceedings of the 28th International Conference on Human Factors in Computing Systems</i> (Atlanta, GA, USA), CHI'10, pages 2287–2296, New York, NY, USA, 2010. ACM.	LS	UP	Log.	12	1

Paper	Stu.	Obj.	Meth.	Part.	Dur.
H. Brignull, S. Izadi, G. Fitzpatrick, Y. Rogers, and T. Rodden. The Introduction of a Shared Interactive Surface into a Communal Space. In <i>Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work</i> (Chicago, IL, USA), CSCW '04, pages 49–58, New York, NY, USA, 2004. ACM	ETH	AB	Obs.	120	14
	FS	UA	Log.	120	14
	FS	UA	Qu.	120	14
	FS	UA	Obs.	120	14
	FS	SI	–	120	14
H. Brignull and Y. Rogers. Enticing People to Interact with Large Public Displays in Public Spaces. In M. Rauterberg, M. Menozzi, and J. Wesson, editors, <i>Proceedings of the IFIP TC13 International Conference on Human-Computer Interaction</i> (Zürich, Switzerland), Interact'03. IOS Press, 2003	FS	SI	Obs.	40	1
	FS	SI	Int.	23	1
	FS	UA	Obs.	23	1
	FS	AB	Log.	40	1
S. Carter, E. Churchill, L. Denoue, J. Helfman, and L. Nelson. Digital Graffiti: Public Annotation of Multimedia Content. In <i>CHI'04 Extended Abstracts on Human Factors in Computing Systems</i> (Vienna, Austria), CHI EA'04, pages 1207–1210, New York, NY, USA, 2004. ACM	AU	UA	Int.	4	1
	FS	AB	Log.	2	5
K. Cheverst, A. Dix, D. Fitton, C. Kray, M. Rouncefield, C. Sas, G. Salsis-Lagoudakis, and J. G. Sheridan. Exploring Bluetooth-based Mobile Phone Interaction with the Hermes Photo Display. In <i>Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services</i> (Salzburg, Austria), MobileHCI'05, pages 47–54, New York, NY, USA, 2005. ACM	LS	UA	Qu.	17	1
	LS	UP	Qu.	17	1
	LS	SI	Qu.	17	1
K. Cheverst, N. Taylor, M. Rouncefield, A. Galani, and C. Kray. The Challenge of Evaluating Situated Display-based Technology Interventions Designed to Foster a Sense of Community. In <i>Proceedings of the Workshop on Ubiquitous Systems Evaluation</i> (Seoul, South Korea), USE'08, 2008	ETH	–	FG	60	700
	AU	–	Qu.	60	700
	DBR	SI	Obs.	60	700
	DBR	AB	Log.	60	700
	DBR	UA	Obs.	60	700
	DBR	SI	Log.	60	700
M. Collomb, M. Hascoët, P. Baudisch, and B. Lee. Improving Drag-and-Drop on Wall-size Displays. In <i>Proceedings of Graphics Interface 2005</i> (Victoria, BC, Canada), GI'05, pages 25–32, Waterloo, ON, Canada, 2005. Canadian Human-Computer Communications Society	LS	UP	Log.	12	–
	LS	UA	Qu.	12	–
E. F. Churchill, L. Nelson, and L. Denoue. Multimedia Fliers: Information Sharing with Digital Community Bulletin Boards. In M. Huysman, E. Wenger, and V. Wulf, editors, <i>Communities and Technologies</i> , pages 97–117, Deventer, The Netherlands, 2003. Kluwer	DBR	UX	Log.	28	180
	DBR	UA	Int.	28	180
	DBR	SI	Qu.	28	180
	ETH	AB	Inr.	17	–
	ETH	AB	Obs.	17	–

Paper	Stu.	Obj.	Meth.	Part.	Dur.
L. Denoue, L. Nelson, and E. Churchill. AttrActive Windows: Dynamic Windows for Digital Bulletin Boards. In <i>CHI'03 Extended Abstracts on Human Factors in Computing Systems</i> (Ft. Lauderdale, FL, USA), CHI EA'03, pages 746–747, New York, NY, USA, 2003. ACM	LS	UP	Obs.	6	1
	LS	UP	Log.	6	1
J. Exeler, M. Buzeck, and J. Müller. eMir: Digital Signs that React to Audience Emotion. In <i>Proceedings of the 2nd Workshop on Pervasive Advertising</i> (Lübeck, Germany), PerAd'09, 2009	FS	AB	aObs.	569	30
	FS	UA	Int.	8	30
J. Falk and S. Björk. The BubbleBadge: a Wearable Public Display. In <i>CHI'99 Extended Abstracts on Human Factors in Computing Systems</i> (Pittsburgh, PA, USA), CHI EA'99, pages 318–319, New York, NY, USA, 1999. ACM	FS	UA	Int.	some	some
R. S. Fish, R. E. Kraut, and B. L. Chalfonte. The VideoWindow System in Informal Communication. In <i>Proceedings of the 1990 ACM Conference on Computer Supported Cooperative Work</i> (Los Angeles, CA, USA), CSCW'90, pages 1–11, New York, NY, USA, 1990. ACM	FS	AB	Obs.	20	90
	FS	AB	Log.	20	90
A. Grasso, M. Muehlenbrock, F. Roulland, and D. Snowdon. Supporting Communities of Practice with Large Screen Displays. In K. O'Hara, E. Perry, E. Churchill, and D. M. Russel, editors, <i>Public and Situated Displays – Social and Interactional Aspects of Shared Display Technologies</i> , pages 261–282. Kluwer, Dordrecht, 2003	DBR	UA	Log.	20	365
	DBR	UA	Qu.	–	–
S. Greenberg, M. Boyle, and J. Laberge. PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal. <i>Personal and Ubiquitous Computing</i> , 3(1):54–64, 1999	FS	AB	Obs.	–	–
	FS	AB	Log.	–	–
U. Hinrichs and S. Carpendale. Gestures in the Wild: Studying Multi-Touch Gesture Sequences on Interactive Tabletop Exhibits. In <i>Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems</i> (Vancouver, BC, Canada), CHI'11, pages 3023–3032, New York, NY, USA, 2011. ACM	FS	AB	Obs.	40	8
	FS	AB	Log.	40	8
E. M. Huang and E. D. Mynatt. Semi-Public Displays for Small, Co-Located Groups. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems</i> (Ft. Lauderdale, FL, USA), CHI'03, pages 49–56, New York, NY, USA, 2003. ACM	FS	AB	Qu.	8	14
	FS	AB	Qu.	8	14

Paper	Stu.	Obj.	Meth.	Part.	Dur.
E. M. Huang, D. M. Russell, and A. E. Sue. IM Here: Public Instant Messaging on Large, Shared Displays for Workgroup Interactions. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems</i> (Vienna, Austria), CHI'04, pages 279–286, New York, NY, USA, 2004. ACM	FS	AB	Log.	11	43
	FS	UA	Obs.	11	43
	FS	SI	Int.	11	43
E. M. Huang, A. Koster, and J. Borchers. Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays? In <i>Proceedings of the 6th International Conference on Pervasive Computing</i> (Sydney, Australia), Pervasive'08, pages 228–243, Berlin, Heidelberg, 2008. Springer-Verlag	ETH	AB	Obs.	17 /	24
	ETH	AB	Log.	105 3 / 50	24
S. Izadi, H. Brignull, T. Rodden, Y. Rogers, and M. Underwood. Dynamo: A Public Interactive Surface Supporting the Cooperative Sharing and Exchange of Media. In <i>Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology</i> (Vancouver, BC, Canada), UIST'03, pages 159–168, New York, NY, USA, 2003. ACM	AU	Int.	–	65	1
	FS	Obs.	UP	–	4
	FS	Int.	UP	–	4
	LS	Log.	UP	12	1
	LS	Log.	AB	12	1
	LS	Log.	SI	12	1
G. Jacucci, A. Morrison, G. T. Richard, J. Kleimola, P. Peltonen, L. Parisi, and T. Laitinen. Worlds of Information: Designing for Engagement at a Public Multi-Touch Display. In <i>Proceedings of the 28th International Conference on Human Factors in Computing Systems</i> (Atlanta, GA, USA), CHI'10, pages 2267–2276, New York, NY, USA, 2010. ACM	FS	AB	Log.	101	3
	FS	UX	Qu.	101	3
	FS	SI	Int.	101	3
M. Karam, T. Payne, and E. David. Evaluating Bluscreen: Usability for Intelligent Pervasive Displays. In <i>2nd International Conference on Pervasive Computing and Applications</i> (Birmingham, United Kingdom), ICPCA'07, pages 18–23. IEEE, 2007	FS	Eff.	Log.	–	365
	FS	UA	Int.	8	7
	FS	UA	Qu.	8	7
A. Khan, J. Matejka, G. Fitzmaurice, and G. Kurtenbach. Spotlight: Directing Users' Attention on Large Displays. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems</i> (Portland, OR, USA), CHI'05, pages 791–798, New York, NY, USA, 2005. ACM	LS	UP	Log.	12	–
P. Marshall, R. Morris, Y. Rogers, S. Kreitmayer, and M. Davies. Rethinking 'Multi-User': An In-the-Wild Study of How Groups Approach a Walk-Up-and-Use Tabletop Interface. In <i>Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems</i> (Vancouver, BC, Canada), CHI'11, pages 3033–3042, New York, NY, USA, 2011. ACM	DBR	AB	Obs.	297	32
	DBR	AB	Int.	297	32
	DBR	AB	aObs.	297	32

Paper	Stu.	Obj.	Meth.	Part.	Dur.
J. F. McCarthy, B. Congleton, and F. M. Harper. The Context, Content & Community Collage: Sharing Personal Digital Media in the Physical Workplace. In <i>Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (San Diego, CA, USA), CSCW'08</i> , pages 97–106, New York, NY, USA, 2008. ACM	FS	UA	Log.	45	28
	FS	AB	Qu.	31	7
	FS	SI	Int.	5	1
J. F. McCarthy, T. J. Costa, and E. S. Liongosari. UniCast, OutCast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays. In <i>Proceedings of the 3rd International Conference on Ubiquitous Computing (Atlanta, GA, USA), UbiComp'01</i> , pages 332–345, London, United Kingdom, 2001. Springer-Verlag	FS	AB	Log.	–	–
	FS	AB	Int.	–	–
A. Meschtscherjakov, W. Reitberger, T. Mirlacher, H. Huber, and M. Tscheligi. AmlQuin - An Ambient Mannequin for the Shopping Environment. In <i>Proceedings of the European Conference on Ambient Intelligence, Aml'09</i> , pages 206–214. Springer-Verlag, Berlin, Heidelberg, 2009	AU	Eff.	Int.	–	3
	AU	UA	Int.	–	3
O. Mubin, T. Lashina, and E. Loenen. How Not to Become a Buffoon in Front of a Shop Window: A Solution Allowing Natural Head Movement for Interaction with a Public Display. In <i>Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II (Uppsala, Sweden), INTERACT'09</i> , pages 250–263, Berlin, Heidelberg, 2009. Springer-Verlag	LS	UP	Log.	12	–
	LS	AB	Log.	12	–
J. Müller, J. Exeler, M. Buzeck, and A. Krüger. ReflectiveSigns: Digital Signs That Adapt to Audience Attention. In <i>Proceedings of the 7th International Conference on Pervasive Computing (Nara, Japan), Pervasive'09</i> , pages 17–24, Berlin, Heidelberg, 2009. Springer-Verlag	FS	AB	Log.	38612	30
	FS	AB	aObs.	38612	30
	FS	UA	Int.	15	1
J. Müller, D. Wilmsmann, J. Exeler, M. Buzeck, A. Schmidt, T. Jay, and A. Krüger. Display Blindness: The Effect of Expectations on Attention towards Digital Signage. In <i>Proceedings of the 7th International Conference on Pervasive Computing (Nara, Japan), Pervasive'09</i> , pages 1–8, Berlin, Heidelberg, 2009. Springer-Verlag	FS	AB	Int.	91	–
	FS	UA	Int.	91	–

Paper	Stu.	Obj.	Meth.	Part.	Dur.
F. Müller, S. Agamanolis, and R. Picard. Exertion Interfaces: Sports Over a Distance for Social Bonding and Fun. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems</i> (Ft. Lauderdale, FL, USA), CHI'03, pages 561–568, New York, NY, USA, 2003. ACM	AU	SI	Qu	32	56
	LS	SI	Qu.	32	56
	LS	UA	Int..	32	56
J. Müller, R. Walter, G. Bailly, M. Nischt, and F. Alt. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In <i>Proceedings of the 2012 ACM Conference on Human Factors in Computing Systems</i> (Austin, TX, USA), CHI'12, pages 297–306, New York, NY, USA, 2012. ACM.	LS	UP	Log.	16	2
	LS	UP	Int.	16	2
	FS	AB	Log.	363	11
	FS	Eff.	Int.	363	11
J. Müller and A. Krüger. MobiDiC: Context Adaptive Digital Signage with Coupons. In <i>Proceedings of the European Conference on Ambient Intelligence</i> (Salzburg, Austria), Aml'09, pages 24–33, Berlin, Heidelberg, 2009. Springer-Verlag	AU	UA	Qu.	39	–
	FS	UP	Int.	24	–
	FS	UP	Qu.	24	–
	DBR	UA	Int.	26	–
K. O'Hara, R. Harper, A. Unger, J. Wilkes, B. Sharpe, and M. Jansen. TxtBoard: From Text-to-Person to Text-to-Home. In <i>CHI'05 Extended Abstracts on Human Factors in Computing Systems</i> (Portland, OR, USA), CHI EA'05, pages 1705–1708, New York, NY, USA, 2005. ACM	ETH	AB	Log.	4	60
	ETH	UA	Log.	4	60
K. O'Hara, M. Lipson, M. Jansen, A. Unger, H. Jeffries, and P. Macer. Jukola: Democratic Music Choice in a Public Space. In <i>Proceedings of the 5th Conference on Designing interactive Systems</i> (Cambridge, MA, USA), DIS'04, pages 145–154, New York, NY, USA, 2004. ACM	ETH	AB	Int.	–	1
	ETH	UA	Obs.	–	1
	FS	UA	Obs.	–	1
	FS	SI	Log.	–	1
	FS	AB	Int.	–	1
T. Ojala, V. Kostakos, and H. Kukka. It's a Jungle Out There: Fantasy and Reality of Evaluating Public Displays in the Wild. <i>Proceedings of the First Workshop on Large Displays in Urban Life</i> , 4:1–4, 2011	DBR	AB	Log.	–	294
	DBR	AB	Obs.	–	294
T. Ojala, H. Kukka, T. Lindén, T. Heikkinen, M. Jurmu, S. Hosio, and F. Kruger. UBI-Hotspot 1.0: Large-Scale Long-Term Deployment of Interactive Public Displays in a City Center. In <i>Proceedings of the 2010 Fifth International Conference on Internet and Web Applications and Services</i> (Barcelona, Spain), ICIW'10, pages 285–294, Washington, DC, USA, 2010. IEEE Computer Society	ETH	UA	Int.	74	2
	FS	UP	Log.	80	90
	FS	UA	Qu.	711	90
	FS	SI	Qu.	711	90
	FS	UP	Qu.	711	90
	FS	Eff.	Int.	81	90
	FS	UX	Int.	81	90

Paper	Stu.	Obj.	Meth.	Part.	Dur.
P. Peltonen, A. Salovaara, G. Jacucci, T. Ilmonen, C. Ardito, P. Saarikko, and V. Batra. Extending Large-Scale Event Participation with User-Created Mobile Media on a Public Display. In <i>Proceedings of the 6th international Conference on Mobile and Ubiquitous Multimedia</i> (Oulu, Finland), MUM'07, pages 131–138, New York, NY, USA, 2007. ACM	FS	SI	Qu.	12	3
	FS	UP	Obs.	12	3
	FS	AB	Int.	12	3
	FS	AB	Log.	12	3
P. Peltonen, E. Kurvinen, A. Salovaara, G. Jacucci, T. Ilmonen, J. Evans, A. Oulasvirta, and P. Saarikko. It's Mine, Don't Touch!: Interactions at a Large Multi-Touch Display in a City Centre. In <i>Proceeding of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems</i> (Florence, Italy), CHI'08, pages 1285–1294, New York, NY, USA, 2008. ACM	FS	SI	Obs.	1199	30
	FS	AB	Log.	1199	30
	FS	AB	Int.	1199	30
F. Redhead and M. Brereton. Designing Interaction for Local Communications: An Urban Screen Study. In <i>Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II</i> (Uppsala, Sweden), INTERACT'09, pages 457–460, Berlin, Heidelberg, 2009. Springer-Verlag	DBR	AB	Log.	–	240
	DBR	SI	Obs.	–	240
W. Reitberger, A. Meschtscherjakov, T. Mirlacher, T. Scherndl, H. Huber, and M. Tscheligi. A Persuasive Interactive Mannequin for Shop Windows. In <i>Proceedings of the 4th International Conference on Persuasive Technology</i> (Claremont, CA, USA), Persuasive'09, pages 41–48, New York, NY, USA, 2009. ACM	FS	UA	Int.	102	3
D. M. Russell, J. P. Trimble, and A. Dieberger. The Use Patterns of Large, Interactive Display Surfaces: Case Studies of Media Design and Use for Blue-Board and MERBoard. In <i>Proceedings of the 37th Annual Hawaii International Conference on System Sciences</i> HICSS'04, Washington, DC, USA, 2004. IEEE Computer Society	FS	UA	Log.	163	1
	FS	AB	Int.	163	1
	FS	UP	Obs.	163	1
	AU	UP	Int.	–	1
	AU	AB	Int.	–	1
A. S. Shirazi, C. Winkler, and A. Schmidt. Flashlight Interaction: A Study on Mobile Phone Interaction Techniques with Large Displays. In <i>Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services</i> (Bonn, Germany), MobileHCI'09, pages 931–932, New York, NY, USA, 2009. ACM	LS	UP	Log.	11	1
	LS	UX	Qu.	11	1

Paper	Stu.	Obj.	Meth.	Part.	Dur.
J. Scheible and T. Ojala. MobiLenin – Combining a Multi-Track Music Video, Personal Mobile Phones and a Public Display into Multi-User Interactive Entertainment. In <i>Proceedings of the 13th annual ACM International Conference on Multimedia</i> (Singapore), MM'05, pages 199–208, New York, NY, USA, 2005. ACM	FS	UX	Obs.	14	1
	FS	UX	Int.	14	1
	FS	UX	Log.	14	1
D. Schmidt, F. Chehimi, E. Rukzio, and H. Gellersen. PhoneTouch: A Technique for Direct Phone Interaction on Surfaces. In <i>Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology</i> (New York, NY, USA), UIST'10, pages 13–16, New York, NY, USA, 2010. ACM	LS	UP	Log.	12	1
J. Schönböck, F. König, G. Kotsis, D. Gruber, E. Zaim, and A. Schmidt. MirrorBoard - An Interactive Billboard. (Lübeck, Germany), In M. Herczeg and M. C. Kindsmüller, editors, <i>Mensch & Computer</i> , pages 217–226. Oldenbourg Verlag, 2008	AU	UA	FG	5	1
	FS	AB	Obs.	12	1
	FS	UP	Int.	12	1
M. Sharifi, T. Payne, and E. David. Public Display Advertising Based on Bluetooth Device Presence. In <i>Proceedings of the 1st Workshop on Mobile Interaction with the Real World</i> (Espoo, Finland), MIRW'06, 2006	FS	AB	Log.	458	90
	FS	Eff.	Log.	458	90
G. Shoemaker, T. Tsukitani, Y. Kitamura, and K. S. Booth. Body-centric Interaction Techniques for Very Large Wall Displays. In <i>Proceedings of the 6th Nordic Conference on Human-Computer Interaction</i> (Reykjavik, Iceland), NordiCHI'10, pages 463–472, New York, NY, USA, 2010. ACM	LS	UP	Obs.	6	–
D. Snowdon and A. Grasso. Diffusing Information in Organizational Settings: Learning from Experience. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Changing our world, changing ourselves</i> (Minneapolis, MI, USA), CHI'02, pages 331–338, New York, NY, USA, 2002. ACM	DBR	AB	Log.	20	365
	DBR	UA	Qu.	20	365
D. Vogel and R. Balakrishnan. Interactive public ambient displays: Transitioning from implicit to explicit, public to personal, interaction with multiple users. In <i>Proceedings of the 17th annual ACM Symposium on User Interface Software and Technology</i> (Santa Fe, NM, USA), UIST'04, pages 137–146, New York, NY, USA, 2004. ACM	LS	Eff.	Obs.	4	1
	FS	UP	Obs.	4	1

Paper	Stu.	Obj.	Meth.	Part.	Dur.
D. Vogel and R. Balakrishnan. Distant Freehand Pointing and Clicking on Very Large, High Resolution Displays. In <i>Proceedings of the 18th annual ACM Symposium on User Interface Software and Technology</i> (Seattle, WA, USA), UIST'05, pages 33–42, New York, NY, USA, 2005. ACM	LS	UP	Log.	12	1
	LS	UX	Qu.	12	1
A. D. Wilson and H. Benko. Combining Multiple Depth Cameras and Projectors for Interactions On, Above and Between Surfaces. In <i>Proceedings of the 23rd annual ACM Symposium on User Interface Software and Technology</i> (New York, NY, USA), UIST'10, pages 273–282, New York, NY, USA, 2010. ACM	AU	Eff.	Int.	(800)	3

Study Types:

ETH: Ethnographic Study, AU: Asking Users, LS: Lab Study, FS: Field Study, DBR: Deployment-based Research

Objectives:

AB: Audience Behavior, Eff: Effectiveness, P: Privacy, SI: Social Impact, UA: User Acceptance, UP: User Performance, UX: User Experience,

Method:

FG: Focus Group, Int: Interview, Log: Logging, Obs: Observations, Qu: Questionnaire

Appendix II: PD-Net Ethics Process

The PD-Net project requires following a strict ethics process since the studies conducted involve human subjects. The goal of this process is to “minimize potential threats to human subjects stemming from project-related user studies” (PD-Net Ethics Handbook). To demonstrate the process, this appendix provides the sample documents used for the Digifieds field study, conducted in Oulu (see also Chapter 9). The ethical worksheet and the study design process template were created by Marc Langheinrich and Nigel Davies as a part of PD-Net.

PD-Net Ethical Worksheet

The ethical worksheet covers the study we conducted in the context of the UbiChallenge 2011 between July to December 2011. The study involved a number of researchers from the Universities of Duisburg-Essen and Oulu.

The document provides a description of the study, including research questions, intended methods, and information on the subjects, the recruitment process, and the data to be collected. Potential risks are identified and planned precautions are described. Finally, information data storage and processing is provided.

Study Design Process Document – Public Trials

The Study Design Process Document provides guidelines for a particular type of study. It outlines potential risks and informs experimenters how to appropriately select and brief participants. Furthermore, requirements for collecting, storing, and processing (personally identifiable) data are described.

Consent Form, Study Instructions, and Questionnaires

A set of documents was provided to the participants in the context of the study. First, and most importantly, the information sheet introduces the purpose of the study, how their data is treated, how long the study will take, and informs them that they can withdraw at any time. Second, each participant signed a consent form allowing documentation of the briefing by the experimenter. Third, a questionnaire handed out to each participant during the field trial.

V2.4

PD-Net Ethical Worksheet

This worksheet documents a particular experiment within PD-Net that is covered by one or more *Study Process Templates (SPTs)*. It is used for documentation purposes and forms part of the corresponding deliverable in which this experiment took place. It helps researchers within PD-Net to ensure that all personal information collected and processed in PD-Net is treated in accordance with the project's ethical guidelines and the feedback from its ethical advisory board (EAB).

1 Study Information

1.1 Study Title

Give a concise title to your study that describes the particular problem you are investigating in your experiment and/or field study.

Digifieds deployment at the UbiChallenge 2011

1.2 Brief Description

Describe your experiment in a few sentences (no more than 3-4 sentences). The description should include the problem you are trying to address and the methods you are planning to use.

Digifieds is a digital public notice area, which allows users to create, post, and retrieve classifieds from public displays. Digifieds provides both a display and a (mobile) web interface, hence supporting both ad-hoc posting but also preparing content on-the-go.

Examples for content

- Classifieds
- Event information
- 3rd party advertisements

Digifieds emerged out of our ethnographic study on traditional shared public notice areas. With the Oulu deployment we aim at understanding challenges and elaborating on suitable practices with regard to use and management of shared display space in the wild.

1.3 Planned Duration

How long (start, finish, duration) do you plan to run these experiments and/or field studies for?

The Digifieds application is scheduled to run from July 6 until December 31 2011. Depending on the success of the deployment and mutual agreement between the Digifieds team, City of Oulu, and University of Oulu the application could be running even after December 31 2011.

1.4 Work Package

What PD-Net work package(s) does this work fall in?

WP1

2 Project Staff

2.1 Principle Researcher, Institution

Who are the principle researchers responsible for this experiment/study (incl. institution)?

Prof. Albrecht Schmidt, University of Duisburg-Essen (UDE)

2.2 Other Staff (Project Members)

Name all personnel involved in this set of experiments/field studies. This may also include project members from other institutions,

Florian Alt, UDE

Dominik Bial, UDE

Thomas Kubitzka, UDE

2.3 External Staff

Name all personnel involved in this set of experiments/field studies. This may also include project members from other institutions, as well as external researchers.

Markus Ortel, UDE

Firas Zaidan, UDE

Björn Zurmaar, UDE

Alireza Sahami Shirazi, USTUTT

Prof. Timo Ojala, University of Oulu

Jarkko Iisakka, University of Oulu

Simo Hosio, University of Oulu

3 Aims and Methods

3.1 Goals and Research Questions

What are the goals of this research effort? What research questions do you hope to be able to answer with this set of experiments? Please be as specific as possible.

The goal of the Oulu deployment is to investigate digital shared public notice areas in the wild. We are especially interested in (1) the motivation of the people to use public displays as a communication medium (e.g., local community, ease of use, etc.), (2) suitable interaction techniques (e.g., how to create, post, retrieve information), and (3) suitable management practices with regard to different stakeholders (e.g., how to avoid / deal with abuse, stale content, inappropriate content). We envision a qualitative description of the deployment based on observations and interviews, as well as quantitative data on usage and uptake.

3.2 Envisioned Methods (Planned Experiments and Studies)

Describe the set of experiments and/or field studies that you are planning to conduct within the scope of this research effort. If there is an order to the experiments, enumerate them in order. Otherwise list them as bullet points in no particular order. Be as specific as possible (e.g. instead of "interviews" write "online interviews through SurveyMonkey" or "Student focus groups on campus"...)

We plan to conduct the following:

- Offline surveys (SUS, NasaTLX)
- Walk-up interviews
- Observations through notes and pictures
- Application usage analysis (through data logs)

We are also considering

- Online surveys

3.3 Study Type and ID

Select the type of study and if appropriate enter a reference to the guidelines used. If the experiment does not fit into one of the existing PD-NET categories select "Other" and contact the coordinator in order to verify if additional EAB input is needed.

Document ID: [Click here to enter text.](#) (see PD-Net Wiki)

- Observational Studies and Interviews (follow local best practices)
- Volunteer Studies (follow local best practices and see **guidelines**)
- Public Trials (follow local best practices and see **guidelines**)
- Others (describe and contact coordinator in case additional ethics advice req.):
[Click here to enter text.](#)

3.4 Adequacy of Methods

Briefly explain why the experimental methods indicated above are adequate for the research questions

Digifieds is going to be deployed on 11 public displays in a real world setting in the city of Oulu, Finland. We plan to take observational notes and photos and conduct offline interviews. In this way we will gather valuable feedback from the users.

Additionally we will analyze application usage by examining data logs in order to discover emerging practice patterns and to get insights into content that is considered to be most appropriate by the stakeholders

Finally, we envision providing online questionnaires (using tools such as LimeSurvey) to be filled in by the stakeholders (content providers, readers).

3.5 Subjects and Recruitment Process (if any)

Who are the subjects of your research, and how are you planning to recruit them? Be as specific as possible. Note that you might not require recruitment if you are simply planning observations in public.

Digifieds is intended to be a public walk-up and use service, that at least once established, does not require any recruiting. However, we believe that in order to make the service known to the public we need to advertise, e.g., in local newspapers, via campus notice boards but also through mailing lists, via Facebook.

If we decide to recruit participants (e.g., for a controlled study) all participants will receive an information sheet with study background information, as well as contact information of the PI in charge and onsite staff. Non-anonymous participants moreover will sign a consent form.

3.6 Data to be Collected

What information are you collecting in your experiments? Try to be as inclusive and specific as possible, listing all potential data types that you might be interested in collecting. If you are not collecting identifiable information from your subjects, state this here, but still list the (anonymous) data that you are planning to collect.

On-staff researcher will take insitu observations with written notes and photos and optionally video recordings. Walk-up interviews will be recorded with written notes and optionally also a voice recorder (with the consent of the interviewed subject).

User interactions, i.e., UI clicks, will be logged for the later analysis.

4 Risks and Precautions

4.1 Potential Risks

Try to envision the risks that your data collection might pose to your data subjects. What if the data you collected would wind up on the Internet, together with your subjects' real name and contact info? Could they suffer problems at work if their employers would find this information? What if hackers would be able to break into your system and steal your data collection? Would they be able to commit criminal acts with this information?

Participants could be identified in the observational pictures. The participant's opinion could be traced back to them as interviews may be recording voice.

4.2 Planned Precautions

Given the risks you identified in question 4.2, what precautions will be taken by you and your team in order to prevent data leakage? Try to be as specific as possible. You can also refer to your description of data storage in the next section.

No names are recorded electronically – we use only random identifiers that require access to the paper consent forms (which are kept in a locked office) to associate them with the participant's name. The demographics collected do not have enough discriminatory power to identify individuals (age, gender, profession), given the city of Oulu has a population of about 150000. Pictures and voice recordings will be stored in encrypted files on a University server with limited access.

5 Data Storage and Processing

5.1 Storage Locations

Describe where each of the data enumerated under 3.6 will be stored. Also consider potential backup processes.

All raw data is kept in a locked office (paper) or in encrypted files (voice, pictures). Anonymous results and transcripts are stored on a server with access control

5.2 Access Control

Explain how access to the data is regulated. Try to be as specific as possible. Use the [PD-Net Guide to Secure Storage](#) Document for guidance, but make sure to describe your actual implementation of these guidelines here.

Access to electronic data is password protected and available only to researchers involved in the project (incl. backups). Physical access to the office is controlled by keys available to project staff and a selected number of department staff only.

5.3 Data Processing

Describe how the collected data will be used. What kind of statistics will be assembled, what kind of qualitative information extracted, what kind of information combined?

The results will be used anonymously for a qualitative and quantitative analysis with regard to the specified research questions. These results will be published in academic papers. Anonymous quotations will also be published.

5.4 Data Anonymization

If some or all of your data will be anonymized or pseudonymized, explain how you do this. What algorithms will be used for the anonymization, and what guarantees do they offer? How are pseudonymous identifiers generated and where is lookup information (if any) for those pseudonymizers kept?

Subject names are recorded only on the consent forms. All experimental data is keyed with a random identifier matched to the consent form to support deletion requests by individual participants. Access to the electronic data would not provide any mechanism for identifying subjects without access to the physical consent forms. The demographic data collected is not enough to identify an individual, given the relatively large population size.

5.5 Data Retention

How long will you keep personally identifiable information (PII)?

- No PII will be kept
- 3 months after study finishes
- 3 months after results have been first published*
- 3 months after PD-Net finishes
- Others (describe): [Click here to enter text.](#)

**Or within 3 months after PD-Net finishes, if results are not published until end of the project*

6 Ethical Checklist

Use the checklist below for a quick overview of the ethical issues in your planned experiment.

Aspect	Yes	No
1. Informed Consent Form Needed? Due to the short evaluation time (2 days) and as we expect to have very short timeslots for the interviews (in the order of minutes) it is not feasible to have consent forms filled in prior to the questions.	<input type="checkbox"/>	<input checked="" type="checkbox"/>
2. Deception Used? <i>If yes, justify here by overwriting this text</i>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
3. Private Information collected? <i>(c.f. to 3.6. for list of data items)</i>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
4. Subjects Remunerated? <i>If decided the amount will be roughly between 2 and 10 euros</i>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
5. Involvement of Children, Patients, People with Cognitive Disorders? <i>If you answer yes to this question, you must explicitly consult with the EAB!</i>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
6. Internal IRB Review Needed? <i>If yes, attach feedback after obtained</i>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

7 Worksheet Versioning

Provide date, authorship, and changes made to *this specific worksheet* in the history table below.

Date	Author	Comment
28.06.2011	Florian Alt	Initial Version
06.07.2011	Florian Alt	Update on Consent Forms
...
...
...
...

V2.3

Procedures for Public Trials

A PD-Net Study-Design Process Document

The PD-Net project will conduct a number of studies to explore how members of the public interact with networked public display systems in the wild, i.e., in a public or semi-public setting. In this document we describe our operating guidelines for such studies. All such studies will follow local guidelines and additionally conform to the rules and procedures set out in this document.

Definition of Public Trials

Public trials allow researchers to investigate actual use of a public display system in people's everyday lives. Public trials do not involve an explicit user recruitment step but instead observe the interaction of a potentially large set of passers-by that happen to encounter a specifically fielded system by chance. However, public trials might trigger (or run concurrently to) a volunteer study, where people can personally register as volunteer participants. Public trials can be conducted in a range of locations including public places, offices, or universities. In all cases explicit permission of the corresponding authorities (e.g., university or city council) will be obtained.

Nature of Studies

Public trials in PD-Net typically involve observing non-recruited participants interacting with publicly (or semi-publicly) accessible public displays. For example, passers-by in a shopping street might come upon a public display installation running a PD-Net application, such as a digital notice board where they can post classifieds via a touch screen. Note that such applications might also entail an online component, where non-recruited participants can, e.g., post messages on a Web form. Observations of such interactions can be direct, i.e., using a researcher on-site taking written notes, making audio-recordings, or capturing video footage (in accordance with local rules and regulations), or indirect, i.e., based on interaction logs, click-stream information, and network traffic data. Information about both direct and indirect observations would be available at or through the installation, e.g., in the form of an information brochure or an online notice. PD-Net public trials never involve activities that place members of the public at risk of bodily harm. The experimental procedure for the studies varies according to the nature of the research but will always be overseen by a local PI.

Participants

Careful thought will always be given to the location of our public trials in order to ensure that applications and content are appropriate for those who become part of the study. No vulnerable participants (e.g., children or people with cognitive disorders) will be targeted.

A brief description of the PD-Net project, as well as a comprehensive description of the study and contact details on appropriate headed paper will be available to members of the public on request.

V2.3

Data

Data collection, storage, and use of personally identifiable information (PII) in PD-Net in general will follow the EU legal framework, as well as individual national legislation on data protection. The project has produced information sheets to help ensure researchers are familiar with these requirements. Where local legislation requires it information will be posted notifying members of the public about the ongoing study. This information might, e.g., describe the data being collected, the purpose of this data collection (i.e., the particular research aspect under investigation), the recipients of this data (i.e., only researchers involved in PD-Net), the use of the data (i.e., that data will only be published in anonymous form), the name and full contact information of the PI responsible for the data collection, access information (i.e., how to get a copy of the data collected about oneself), and how long data will be retained.

All collected data will use *pseudonymous identifiers* for all subjects whenever possible, or use appropriate secure storage procedures to safeguard PII where such anonymisation at collection time is not possible (such as photographs). PII collected through indirect observations (e.g., clickstream data, form filling data) that pertain to non-recruited participants will be deleted or anonymized prior to archival. We have produced an information sheet to provide researchers with guidance on secure data storage. Collected data will be deleted at most 3 months after PD-Net ends, though earlier times are possible. All publications will only use fully anonymized data when reporting qualitative and quantitative data.

Documentation

Details of each public trial will be recorded on a PD-Net ethics worksheet and these worksheets will be uploaded to and stored on the project's secure wiki.

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



Information Sheet

Thank you for your participation in this study.

The purpose of the study is to explore users' motivations, suitable interaction techniques, and applicable management practices for the use of digital shared public notice areas. In the following questionnaire you will be asked questions about the Digifieds service and your motivations, feelings, and concerns of using it in a public space.

The data gathered during this study will be used only for research and teaching purposes, and only in anonymous form.

It is expected that the study will take approximately 45 minutes. Participation is entirely voluntary and you may withdraw from it at any time during or after the study. If you have any further questions about the study, feel free to contact

- **Responsible Investigator:** Prof. Albrecht Schmidt (albrecht.schmidt@acm.org)
- **Onsite Investigators:** Thomas Kubitzka (thomas.kubitzka@uni-due.de), Florian Alt (florian.alt@uni-due.de)
- **Local Contact Point:** Prof. Timo Ojala (timo.ojala@ee.oulu.fi)
- **Ethical Advisory Board Member:** Prof. Bertil Cottier (bertil.cottier@usi.ch)

If you are using Digifieds in a group each member is required to fill in a separate consent form and questionnaire!

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)

UNIVERSITÄT
DUISBURG
ESSEN

FALUNO
The Ruhr Institute for Software Technology

Examples of (Traditional) Public Notice Boards

Public Notice Boards in a Supermarket



Public Notice Board on a University Campus



UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



“Digifieds” Consent Form

- I have read and understood the information sheet.
- I understand the nature of this study and agree to participate.
- I understand that I can withdraw from this study at any time.

Participant ID _____ (to be filled out by researcher)

If you agree to be contacted for further questions, please write down your contact info (email/phone):

Participant(s) _____ Date _____

Researcher _____ Date _____

This study is funded in part by European Union Seventh Framework Programme (FP7/2007-2013) under the project “pd-net - Towards Future Pervasive Display Networks”

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



1. Background

Gender: male female

Age: _____ Profession: _____

How many hours per day do you (on average) use your computer? _____ (h/day)

How many hours per day do you (on average) use your mobile phone(s)? _____ (h/day)

Does your mobile phone have a touch screen? yes no I don't know

How often do you (on average) access the Internet with your mobile phone?

- Every Day or More
- 2-6 Times/Week
- About Once/Week
- About Once/Month
- Never

Do you have an unlimited data package on your mobile phone? yes no I don't know

Do you use apps on your phone (e.g., Android, iPhone, ...)? yes no don't know

In case you have other devices with touch screen (e.g., an iPad), how many hours do you (on average) use them daily? How often do you use the Ubi Displays?

- | | |
|--|--|
| <ul style="list-style-type: none"> <input type="checkbox"/> Every Day or More <input type="checkbox"/> 2-6 Times/Week <input type="checkbox"/> About Once/Week <input type="checkbox"/> About Once/Month <input type="checkbox"/> Never | <ul style="list-style-type: none"> <input type="checkbox"/> Every Day or More <input type="checkbox"/> 2-6 Times/Week <input type="checkbox"/> About Once/Week <input type="checkbox"/> About Once/Month <input type="checkbox"/> Never |
|--|--|

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



2. Use of Notice Boards

How many places providing public notice boards
(similar to those on page 2) do you know? _____ places

How often did you look for / take away any content
(classifieds, etc.) from these boards within the last 6
months? _____ times

Which types of information did you look for / take away from these boards?

Events	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
For Sale	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Lost & Found	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Housing	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Jobs	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Dating	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Community	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Services (e.g., Babysitting)	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Local News	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Other: _____	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often

How often did you post any content (classifieds,
etc.) to these boards within the last 6 months? _____ times

Which types of information did you post?

Events	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
For Sale	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Lost & Found	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Housing	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Jobs	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Dating	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Community	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Services (e.g., Babysitting)	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Local News	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often
Other: _____	Never	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Often

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



3. Testing the Display Application

3.1 Tasks

Task 1a:

You found a mobile phone and want to bring it back to its owner - please, go to the display and launch the Digifieds application. Then, create a Digified including a title, description, and email address. Place it in the category Lost & Found!

Task 1b:

You are looking for a bike! Open again the Digifieds service (if you have not done so yet). Browse through the offers and find an appropriate bike. Send the Digified to your email address.

Hint:

You can use digifieds@ubichallenge.fi in case you do not want to provide your own address.

When you finished both tasks, please proceed with filling in the questionnaire on the following page.

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



3.2 Evaluation of the Display Client

Please answer the following questions with regard to the display application you just used.

1. I think that I would like to use this system frequently
 strongly disagree strongly agree
2. I found the system unnecessarily complex
 strongly disagree strongly agree
3. I thought the system was easy to use
 strongly disagree strongly agree
4. I think that I would need the support of a technical person to be able to use this system
 strongly disagree strongly agree
5. I found the various functions in this system were well integrated
 strongly disagree strongly agree
6. I thought there was too much inconsistency in this system
 strongly disagree strongly agree
7. I would imagine that most people would learn to use this system very quickly
 strongly disagree strongly agree
8. I found the system very cumbersome to use
 strongly disagree strongly agree
9. I felt very confident using the system
 strongly disagree strongly agree
10. I needed to learn a lot of things before I could get going with this system
 strongly disagree strongly agree

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



4. Testing the Phone Application

4.1 Tasks

Task 2a:

You want to sell the pen you are using for filling in this questionnaire - please take the mobile phone and launch the Digifieds application. Create a new Digified, fill in title, description and an email address. Then, take a picture of the pen and select the category "Sales and Buys". Press "Save".

You will now receive a code - enter it on the display to activate your Digified and make it appear on the screen.

Task 2b:

You are interested in the *Beach Tennis Cup 2011* event. Open again the Digifieds service (if you have not done so yet). Browse through the offers and find the according event. You want to take it away with your mobile phone using QR codes. Therefore, start your mobile application and select "Take away". Then select "via QR code". Scan the QR code at the display.

Task 2c:

You are interested in the *Sky Diving Oulu* ad. Open again the Digifieds service (if you have not done so yet). Browse through the offers and find the according event. You want to take it away with your mobile phone using the provided code (which looks something like this: "xZ63"). Start again your mobile application and select "Take away". Then enter the code to transfer the Digified to your phone.

Hint:

You can use digifieds@ubichallenge.fi in case you do not want to provide your own address

When you finished both tasks, please proceed with filling in the questionnaire on the following page.

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



4.2 Evaluation of the Mobile Phone Application

Please answer the following questions with regard to the phone application you just used.

1. I think that I would like to use this system frequently
 strongly disagree strongly agree
2. I found the system unnecessarily complex
 strongly disagree strongly agree
3. I thought the system was easy to use
 strongly disagree strongly agree
4. I think that I would need the support of a technical person to be able to use this system
 strongly disagree strongly agree
5. I found the various functions in this system were well integrated
 strongly disagree strongly agree
6. I thought there was too much inconsistency in this system
 strongly disagree strongly agree
7. I would imagine that most people would learn to use this system very quickly
 strongly disagree strongly agree
8. I found the system very cumbersome to use
 strongly disagree strongly agree
9. I felt very confident using the system
 strongly disagree strongly agree
10. I needed to learn a lot of things before I could get going with this system
 strongly disagree strongly agree

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



5. Content on Digifieds

Which kind of information would you like to find on / take away from Digifieds?

Events	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
For Sale	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Lost & Found	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Housing	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Jobs	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Dating	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Community	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Services (e.g., Local News)	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Other: _____	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Other: _____	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Other: _____	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all

Which other type of information would you like to find on Digifieds?

Which kind of information would you offer via Digifieds?

Events	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
For Sale	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Lost & Found	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Housing	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Jobs	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Dating	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Community	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Services (e.g., Local News)	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Other: _____	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Other: _____	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all
Other: _____	Definitely	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not at all

Which other type of information would you like to offer via Digifieds?

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



6. Privacy

6.1 Scenario 1

Imagine you are standing in a public space in front of the UbiDisplay. Please rate whether you agree with the following quotes or not.

Inputting the **title** of a Digified affects my privacy. Totally agree Totally disagree

Inputting the **content / description** of a Digified affects my privacy. Totally agree Totally disagree

Inputting my **email address** affects my privacy. Totally agree Totally disagree

I am feeling uncomfortable interacting with Digifieds. Totally agree Totally disagree

6.2 Scenario 2

Imagine you are using the Digified mobile application in a public space to create your Digified. Please rate whether you agree with the following quotes or not.

Inputting the **title** of a Digified affects my privacy. Totally agree Totally disagree

Inputting the **content / description** of a Digified affects my privacy. Totally agree Totally disagree

Inputting my **email address** affects my privacy. Totally agree Totally disagree

I am feeling uncomfortable interacting with Digifieds. Totally agree Totally disagree

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



My privacy is being affected when inputting private data (e.g., phone number / email address) ...

... on a traditional paper-based notice board. Totally agree Totally disagree

... on a digital notice board (e.g., Digifieds). Totally agree Totally disagree

... on the mobile phone. Totally agree Totally disagree

... at home on the PC. Totally agree Totally disagree

UbiChallenge 2011: Digifieds

- Field Trial in Oulu, Finland -

(July/August 2011)



7. Digifieds Features

Please rate how useful you consider the following features:

Searching the Digifieds	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all
Images associated with a Digified	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all
Videos associated with a Digified	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all
Map associated with a Digified	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all
Editing Digifieds	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all
Deleting Digifieds	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all
Creating Digifieds at home	Very useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Not useful at all

Did you miss any features?

How could Digifieds support you when posting / looking for / taking away information?

Now that you know the application, do you intend to use it again?

Definitely Not at all

Do you have further comments?

Florian Alt

A Design Space for Pervasive Advertising on Public Displays

Today, people living in cities see up to 5000 ads per day and many of them are presented on public displays. More and more of these public displays are networked and equipped with various types of sensors, making them part of a global infrastructure that is currently emerging. Such networked and interactive public displays provide the opportunity to create a benefit for society in the form of immersive experiences and relevant content. In this way, they can overcome the display blindness that evolved among passersby over the years. We see two main reasons that prevent this vision from coming true: first, public displays are stuck with traditional advertising as the driving business model, making it difficult for novel, interactive applications to enter the scene. Second, no common ground exists for researchers or advertisers that outline important challenges. The provider view and audience view need to be addressed to make open, interactive display networks, successful.

The main contribution made by this thesis is presenting a design space for advertising on public displays that identifies important challenges - mainly from a human-computer interaction perspective. Solutions to these core challenges are presented and evaluated, using empirical methods commonly applied in HCI. The results of this thesis are envisioned to provide a basis for future research and for practitioners to shape future advertisements on public displays in a positive way.