



NTNU – Trondheim
Norwegian University of
Science and Technology

Evaluation of an Interactive Campaign

Exploring the use of a motion-controlled
game in a public space

Mari Hansen Asplem
Mia Aasbakken

Master of Science in Computer Science
Submission date: June 2012
Supervisor: Alf Inge Wang, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Problem Description

In this project two students will develop an interactive campaign that is going to be placed in a crowded public space, for example a mall. The project will be done in cooperation with the Trondheim-based company Global Illumination (GI). The students will need to investigate what has been done previously and find relevant experiences with using interactive installations in a public space. The campaign should use relatively cheap motion sensing technology, like the Microsoft Kinect, and the students should use the open source framework OpenNI and develop a prototype for the campaign. Several empirical studies need to be conducted in order to gather data and evaluate the campaign, while focusing on how people in a public space relate to such campaigns and the possible marketing potential. The technology that is used should also be evaluated, as GI is interested in potentially using it in the future.

Assignment given: January 16th, 2012

Supervisor: Alf Inge Wang

Abstract

This thesis project explores and evaluates the use of a motion-controlled game as an interactive campaign in a public space. Through a collaboration with the Trondheim-based company Global Illumination, we were given the task of developing a prototype that would be tested in the field. The objective was to evaluate the relevant technology, how users relate to the prototype and what the marketing potential is.

Through a literature study and using the technology, we found that while it is still new, the technology is both mature enough to be used for several different platforms and languages, and cheap enough that the cost of getting started should not be a hindrance.

We developed a game prototype using the OpenNI and XNA frameworks, in which people who pass by the display would be reflected on the screen in the form of a silhouette and automatically be a part of the game. The prototype was tested at four different public locations in Trondheim, and was evaluated mainly through observation and questionnaires given to both participants and non-participants.

Our findings suggest that there is definite potential for using motion control in interactive campaigns in public settings. The game attracted a good amount of attention, and seemed to pique the curiosity of passers-by. We saw a trend emerge where participants were comfortable with playing the game in public and were easily engaged. Children and young people in groups were by far the most active participants. We also found that keeping a low threshold for interaction was essential, as adding an extra step in the form of a wave gesture to participate reduced the number of participants considerably.

Sammendrag

Dette masterprosjektet undersøker og evaluerer bruk av bevegelsesstyrte spill som en interaktiv kampanje i det offentlige rom. Gjennom et samarbeid med det nyetablerte firmaet Global Illumination i Trondheim, fikk vi i oppgave å utvikle en prototype som skulle testes ved hjelp av feltstudier. Ut i fra dette skulle vi evaluere den relevante teknologien, hvordan brukere forholdt seg til prototypen og hva markedsføringspotensialet er.

Gjennom et litteraturstudium og bruk av teknologien, fant vi at selv om man fremdeles er på et tidlig stadium, er teknologien moden nok til å brukes på flere ulike plattformer og med ulike programmeringsspråk, i tillegg til å være billig nok til at kostnad ikke er et stort hinder for å komme i gang.

Vi utviklet en spillprototype ved hjelp av rammeverkene OpenNI og XNA, hvor forbipasserende vil vises på skjermen i form av silhuetter og dermed automatisk er en del av spillet. Prototypen ble testet på fire ulike lokasjoner i Trondheim, og ble evaluert hovedsakelig gjennom observasjon og spørreskjemaer som ble gitt til både deltakere og ikke-deltakere.

Funnene våre indikerer at det definitivt eksisterer et potensiale for å bruke bevegelsesstyring i interaktive kampanjer i det offentlige rom. Spillet tiltrakk seg en god del oppmerksomhet, og virket å pirre nysgjerrigheten til forbipasserende. Vi så en trend der deltakere var komfortable med å spille spillet offentlig og de ble lett engasjerte. De mest aktive deltakerne var grupper av barn og unge. Vi så også at det er essensielt å ha en lav terskel for deltakelse, da det å legge til et ekstra steg i form av en vinkebevegelse førte til en betydelig nedgang i antall deltakere.

Preface

This report presents the work done for our master's thesis at the Department of Computer and Information Science at the Norwegian University of Science and Technology. The work was done during the spring of 2012 and concludes our master's degrees in Computer Science.

We want to thank Global Illumination, represented by Lars-Erik Mogenssen Ravn and Odd Joachim Aschim, for giving us an interesting assignment and the opportunity to explore the problem in a realistic business setting, and for providing us with the necessary equipment to perform the field studies. Thanks to Anders Knive Lassen for technical assistance and support throughout the semester.

We would also like to thank Øyvind Aralt on behalf of Location, the administration at Pirbadet, Bjørn Henning Hammer and Christian Sandstad from Studentsamskipnaden i Trondheim and Anette Vinje and Asbjørn Grønli at Mercursenteret for letting us test the prototype and providing us with essential data for our evaluation.

Last but not least, we would like to thank our supervisor Alf Inge Wang for invaluable guidance and feedback from day one.

Trondheim, June 7, 2012

Mari Hansen Asplem

Mia Aasbakken

Contents

1	Introduction	1
1.1	Context	1
1.2	Personal motivation	2
1.3	Goals	2
1.4	Structure of the thesis	3
2	Methodology	5
2.1	Goals and research questions	5
2.1.1	Method	5
2.1.2	Result	7
2.2	Evaluation with the DECIDE framework	9
2.2.1	Goals and questions	9
2.2.2	Evaluation approaches and methods	10
2.2.3	Practical issues	11
2.2.4	Ethical issues	12
2.2.5	Evaluating, interpreting and presenting data	12
2.3	Our process	13
2.3.1	Literature study	14
2.3.2	Prototype development	14
2.3.3	Evaluation approach	14
2.3.4	Observation	15
2.3.5	Questionnaires	16
2.3.6	Interviews	18
2.3.7	Ethical issues	19
2.3.8	Practical issues	19
2.3.9	Data gathering	19
3	Literature Study	23
3.1	Interactivity in applications	23

3.1.1	Short history of user interfaces	23
3.1.2	Using gestural interfaces	25
3.2	Interactive public displays	27
3.2.1	The Audience Funnel	27
3.2.2	Design considerations	29
3.2.3	Examples	33
3.3	Motion sensing devices	37
3.3.1	Hardware	38
3.3.2	Uses	39
3.4	Developing motion-controlled applications	43
3.4.1	OpenNI framework	43
3.4.2	Microsoft's Kinect SDK	46
4	Design and Development	49
4.1	Application concept	49
4.1.1	Brainstorming	50
4.1.2	Reusability	51
4.1.3	Low threshold for interaction	52
4.1.4	Final concepts	55
4.1.5	Other aspects	56
4.2	Chosen technologies	57
4.2.1	Why OpenNI, C# and XNA?	57
4.2.2	Using OpenNI	58
4.2.3	XNA Game Studio	61
4.2.4	Farseer Physics Engine	62
4.3	Game prototype	64
4.3.1	System overview	67
4.3.2	Modes	69
4.3.3	Versions	70
4.4	Interactive poster prototype	72
5	Empirical Study	75
5.1	Research context	75
5.1.1	Pirbadet waterpark	75
5.1.2	Stripa, NTNU campus	77
5.1.3	Location, Prinsen kino	78

5.1.4	Mercur shopping mall	80
5.2	Observation	81
5.2.1	Pirbadet	81
5.2.2	Stripa, NTNU	85
5.2.3	Location, Prinsen kino	90
5.2.4	Mercur	92
5.2.5	Interactive poster	98
5.3	Questionnaires	99
5.3.1	Demographics	99
5.3.2	Playing the game	102
5.3.3	Marketing aspect	111
5.4	Interviews	115
5.5	Input logs	116
5.5.1	Location data	117
5.5.2	Pirbadet waterpark	118
5.5.3	Stripa, NTNU campus	119
5.5.4	Location, cinema kiosk	119
5.5.5	Mercur shopping mall	120
6	Discussion	123
6.1	User experience	123
6.1.1	Comfort levels	124
6.1.2	User groups and engagement	125
6.1.3	Ease of use	127
6.1.4	Group dynamics	129
6.1.5	Other observations	130
6.1.6	Interactive poster	131
6.2	The intersection between user experience and marketing	131
6.2.1	Level of comfort with participating in a campaign	132
6.2.2	Placement of the display	133
6.2.3	Reward or no reward	134
6.3	Marketing potential	134
6.3.1	Effect on sales	135
6.3.2	Effect on user behavior	136
6.3.3	Interactive poster	137
6.4	Technology	138

6.4.1	Feature sets	138
6.4.2	Documentation and ease of use	140
6.4.3	Motion-controlled web applications	141
7	Conclusion	145
7.1	Summary	145
7.2	Feedback from Global Illumination	146
7.3	Our recommendations	147
7.4	Further work	148
	References	151
A	Questionnaires	159

List of Figures

2.1	The GQM approach	6
2.2	GQM Example	6
2.3	Goal 1 with its questions and metrics	8
2.4	Goal 2 with its questions and metrics	8
2.5	Goal 3 with its questions and metrics	9
3.1	The Audience Funnel (Müller et al., 2010a)	28
3.2	One of the displays from Magical Mirrors (Michelis, 2008)	34
3.3	CityWall installation in a store front window in Helsinki (Morrison et al., 2008)	35
3.4	Virtual fitting room using augmented reality and Kinect (AR Door, 2011)	36
3.5	Motion sensors: Microsoft Kinect and ASUS Xtion Pro Live	38
3.6	Interactive art installation by Seeper at a WIRED Event (seeper, 2011)	42
3.7	Abstract layered view of OpenNI (OpenNI, 2012)	44
4.1	Concept drawing 1: Attention grabber	50
4.2	Concept drawing 2: Change of styles	51
4.3	Concept drawing 3: Engaging crowds	52
4.4	Texture to polygon by using Farseer's PolygonTools	66
4.5	Overview of the system	68
4.6	Two-player mode with two active and two non-active players	70
4.7	Standard version of the game	71
4.8	Popcorn version of the game	71
4.9	Mercur version of the game	72
4.10	Poster prototype after two and six users, respectively	73
5.1	Environment at Pirbadet	76
5.2	Environment at Stripa	78
5.3	Environment at Location	79

5.4	Environment at Mercur	81
5.5	The installation in use at Pirbadet	82
5.6	Screenshot: Tricking with a ball	84
5.7	The installation in use at Stripa	86
5.8	Screenshot: A player using her backpack while playing	88
5.9	The installation in use at Location	91
5.10	The installation in use at Mercur	93
5.11	Screenshot: Two people collecting balls in their arms	94
5.12	Eager participant at Mercur	97
5.13	Gender distribution of our questionnaires	100
5.14	Age distribution of our questionnaires	100
5.15	Age distribution by participants and non-participants	101
5.16	Hours of gaming per week among respondents	102
5.17	Answer distribution I1: I felt comfortable with becoming a part of the game without giving explicit consent.	103
5.18	Answer distribution I2: I was comfortable with playing the game in a public space.	103
5.19	Answer distribution I3: I would have been more comfortable with playing with a controller (mobile phone, Xbox controller)	104
5.20	Answer distribution I4: I thought getting started with the game was easy.	105
5.21	Answer distribution I5: I thought it was easy to control the game.	105
5.22	Answer distribution I6: I think most people will be able to play the game.	106
5.23	Answer distribution I7: I became unaware of my surroundings while playing the game.	106
5.24	Answer distribution I8: Did you play with anybody else?	107
5.25	Answer distribution I9a: If yes: I think the game was made more fun by playing with somebody else.	107
5.26	Answer distribution I9b: If no: I think I would have played longer if I played with somebody else.	108
5.27	Answer distribution I10: I would consider playing with strangers.	109
5.28	Answer distribution I11: If your child played the game: I would be comfortable with letting my child participate in a game/campaign like this.	109
5.29	Answer distribution I12: Did you capture enough objects to win?	110

5.30	Answer distribution I13: I wanted to win while playing the game.	111
5.31	Answer distribution I14: The game/campaign caught my attention to a greater degree than a poster would have done.	112
5.32	Answer distribution I15: If you managed to collect enough objects: I wanted to buy the advertised product more than before I played the game.	113
5.33	Answer distribution I16: I would be comfortable with participating in a game like this if it was a part of an advertising campaign.	114
5.34	Answer distribution I17: My willingness to buy a product from the advertiser increased after playing the game.	114
5.35	Distribution of players, winning events and hours played at Pirbadet	118
5.36	Distribution of players, winning events and hours played at Stripa	119
5.37	Distribution of players, winning events and hours played at Location	120
5.38	Distribution of players, winning events and hours played at Mercur	121

Introduction

As the world around us has become filled with more and more advertisements in the form of posters or public displays, people have learned to ignore these types of campaigns, and grabbing the attention of passers-by has become even more challenging. LCD and plasma screens make it easy and cheap to create digital campaigns that can be switched out on the fly and they also make it possible to incorporate interaction with the user, for example in the form of a touch interface, instead of having the user be a passive spectator. The next big thing in user interfaces are most likely gesture-based interfaces, and with the introduction of Microsoft's Kinect, the technology to pull something like this off has become cheap and more available than before. Using this type of technology within interactive campaigns has not been the topic of much research, and it is an exciting area to explore.

1.1 Context

One of the authors wrote about interactive art installations in her specialization project and collaborated with a newly established Trondheim-based company called Global Illumination (GI) (Global Illumination, 2012). Global Illumination has their main focus on a communication platform which makes distribution of information and messages easier. They are also interested in being on the bleeding edge of technology and to use this in new and exciting ways for conveying messages to an audience. They were therefore very interested in exploring the possibilities of interactive campaigns, and wanted to continue to look into what was discovered in the specialization project further in this semester by involving themselves in a master thesis project. We had several meetings with GI in order to narrow down the problem domain, and we ended up with a focus on motion sensors and how these can be utilized in a public

2 INTRODUCTION

interactive campaign in order to better catch the audience's attention. The field of public interactive campaigns has not largely been a subject of exploration up to this date, and GI thought it would be a valuable insight to test something like a motion-controlled campaign in the real world to see if it can engage people in new ways, taking it to another level from the standard poster like displays.

1.2 Personal motivation

Both authors have always been genuinely interested in human-computer interaction, and both have the HCI-courses offered at NTNU highly ranked on their list of favorite subjects. Mia has always had a fascination for real-time graphical applications like games and demos¹, and this led her to choose game technology as her specialization. Mari, on the other hand, is more interested in graphical user interfaces and how people relate to these, as well as user-oriented designs. Her specialization is within Program and Information Systems, with a focus on HCI. She was convinced by Mia to join this project, as it combines these two fields of computer science, and also involves exploring, using and testing exciting and fairly new technology. It is also very intriguing to investigate how technology can be used in another context than it was originally made for. In addition to this we saw the value in being able to collaborate with a company and getting the opportunity to test a prototype on a large audience in a public space and getting hands-on experience with how things are done "in the real world".

1.3 Goals

Our overarching goals for the thesis can be separated into three different focus areas. First of all, we want to find out how people relate to a motion-controlled application, in our case a game, in a public space. This is interesting because

(1) A non-interactive software application that runs in real-time on your computer, presenting skills in programming, art and music, see [http://en.wikipedia.org/wiki/Demo_\(computer_programming\)](http://en.wikipedia.org/wiki/Demo_(computer_programming))

there are many particularities about public spaces that could affect how they choose to interact, and also because if it is to be used for marketing, public spaces allow the campaign to reach more people. Secondly, we want to find out how it can be used for marketing – if people are comfortable with being a part of a campaign in this way and if it has any effect on how people choose to spend their money. Separate from these two goals, but still a very interesting one, is the third goal, which is concerned with the state of the technology. Developing motion-controlled applications for use on a PC platform is a relatively new concept, and we want to find out which options exist and also whether or not the technology is mature enough to be used in a browser, which is what GI eventually want to do. Based on this, this study seeks to:

- Understand and evaluate how users relate to motion-controlled applications in a public space.
- Explore how motion-controlled applications can be used within interactive campaigns.
- Map out the possibilities for developing motion-controlled applications and evaluate the maturity of the technology.

Each of these goals have two to four research questions associated with them, and these can be found in Chapter 2.

1.4 Structure of the thesis

The structure of this thesis more or less correspond to the process the project has been through, with deciding upon the methodology and performing a literature study first, followed by the design and development process and then testing and evaluating the project.

1 Introduction Provides the reader with some information about the motivation and context for the project, and present our main goals.

2 Methodology This chapter describes our chosen methodologies for defining research questions and the framework we will use for our evaluation. We also present and describe how we will perform the various phases of the project and how we will perform the evaluation.

4 INTRODUCTION

3 Literature Study The literature study presents some relevant literature and findings about methods of interactivity in applications and using public interactive displays, and then describes the technology we will use and some of the options that exist.

4 Design and Development This chapter will describe the process through which we arrived at the concepts for our prototypes. We also present an overview of the implementation of each prototype.

5 Empirical Study This chapter presents the context and results from our evaluation. This includes input logs from the game itself, observations from each of the locations and results from the questionnaire and some short interviews.

6 Discussion In this chapter, we discuss the findings from our literature study and results. We will go through the research questions from Chapter 2 and give answers to these.

7 Conclusion This chapter presents our conclusions and recommendations and suggests some further work that we feel would be useful.

Methodology

This chapter gives insight into the methods that will be used in the different phases of the project. We start off by explaining the method we used to reach the goals and research questions necessary to analyze and evaluate the results of the study – the Goal Question Metric approach. The results we reached from this process were then used with the DECIDE framework for guiding evaluation of products. Following the description of the DECIDE framework is an overview of our process and the methods we used, such as a literature study and field studies of the prototype.

2.1 Goals and research questions

2.1.1 Method

To aid in the process of creating research questions that would be concrete enough to be answered using specific metrics while still giving insight into the overarching goals, the Goal Question Metric approach was chosen as a guide. The Goal Question Metric (GQM) approach formalizes how one can go from having high-level productivity or quality goals to formulating questions that say something about how this goal will be achieved and find out which metrics that are associated with answering each question (Basili et al., 1996). The overall structure of the approach can be seen in Figure 2.1.

An example of the GQM approach can be seen in Figure 2.2, which is a table from the article by Basili et al. The stated goal on the conceptual level can be defined in terms of what will be analyzed (object), which aspect is most interesting (issue), what the purpose is (understand, improve, evaluate), the relevant context and the point of view (van Solingen and Berghout, 1999, page 51). Templates are provided for this, but they are not necessary to use. On

6 METHODOLOGY

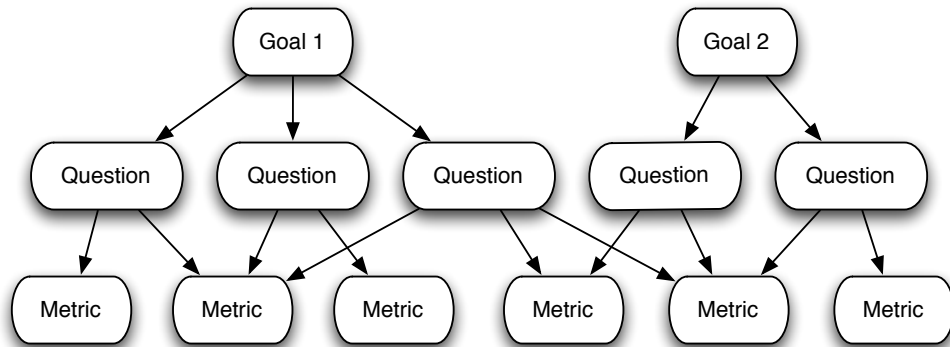


Figure 2.1: The GQM approach

the operational level, a set of questions for each goal are formulated. These questions characterize how a specific goal can be achieved or assessed. For each of these questions, a set of metrics are then defined. The metrics are a set of data that will provide quantitative answers to each question, and the data can be either subjective or objective.

Goal	Purpose Issue Object (process) Viewpoint	Improve the timeliness of change request processing from the project manager's viewpoint
Question		What is the current change request processing speed?
Metrics		Average cycle time Standard deviation % cases outside of the upper limit
Question		Is the performance of the process improving?
Metrics		$\frac{\text{Current average cycle time}}{\text{Baseline average cycle time}} * 100$ Subjective rating of manager's satisfaction

Figure 2.2: GQM Example

The approach may be used for a variety of different types of projects and is, for example, well suited to defining quality or productivity improvement goals within an organization, with very process-specific questions and related

metrics, as seen in Figure 2.2. However, GQM is also a suitable approach in the case of our study, where the project is more exploratory and the overarching goals relate to understanding and evaluating user behavior, for simply organizing our thoughts, goals and research questions. The metrics in this case will be the data that is necessary to evaluate the research questions, for example a literature study and observation of the application in the field as well as interviews with participants. These metrics will mean that the answers to the research questions may not always be based on quantitative data, like the GQM approach advises, but also qualitative data from the literature study and interviews. Because of the fact that our evaluation will be primarily focused on seeing if and how people discover and choose to interact with the application, measuring metrics like effectiveness or efficiency is difficult (Müller et al., 2010b), given that there is no specific, well-defined task to be performed.

This study seeks to provide insight into three different aspects of interactive campaigns with motion-controlled applications – the human-computer interaction aspect, the technology aspect and the marketing aspect. These three perspectives can be seen in the overarching goals of the study, and they separate the research questions into three different categories. To limit the scope of the study to fit within one semester, the research questions posed for each goal focus on a limited range of topics. Due to the exploratory nature of the project, they were selected based on which aspects we found to be most interesting. It was also important to us that they could be tested by making minor changes to the prototype, without having to develop several very different versions – which would be too time consuming. For example, the role collaboration plays is very interesting because the installation will be placed in a public space, and the effect it has can be tested through observation in the field and interviews with participants with only minor changes in the application (switching from single-player to multiplayer).

2.1.2 Result

The goals, research questions and metrics we arrived at can be seen in Figures 2.3, 2.4 and 2.5. Answering these questions is the main objective of the literature study, prototype development and field studies. From these

8 METHODOLOGY

results, we can then perform an evaluation and draw conclusions related to our goals.

G1: The study seeks to understand and evaluate how users relate to motion-controlled applications in a public space.	
RQ1.1	Are people comfortable with being a part of a motion-controlled game in a public space without explicitly giving their permission or actively making an effort?
Metric	Results from observation on site, questionnaires and interviews with participants
RQ1.2	Which user groups are most engaged in using the application?
Metric	Results from observation of participants and passers-by on site
RQ1.3	Do people find it easy to get started with and control the application?
Metric	Results from observation on site, questionnaire and interviews with participants
RQ1.4	Are people who play with a group more engaged in using the application than those who play alone?
Metric	Results from observation on site, questionnaires and interviews with participants.

Figure 2.3: Goal 1 with its questions and metrics

G2: The study seeks to explore how motion-controlled applications can be used within interactive campaigns.	
RQ2.1	Are people comfortable with being part of an interactive campaign like this?
Metric	Results from observation on site, questionnaires and interviews with participants.
RQ2.2	Does an interactive campaign of this type have any effect on people's purchasing behaviour?
Metric	Results from observation on site, questionnaires and interviews with participants.

Figure 2.4: Goal 2 with its questions and metrics

G3: The study seeks to map out the possibilities for developing motion-controlled applications and evaluate the maturity of the technology.	
RQ3.1	Which options with regards to APIs and programming environments exist if you want to develop motion-controlled applications?
Metric	Literature study of the relevant technologies
RQ3.2	Is the technology mature enough to be used within a browser for an easy to develop, cross-platform solution using modern web technologies?
Metric	Literature study of the relevant technologies

Figure 2.5: Goal 3 with its questions and metrics

2.2 Evaluation with the DECIDE framework

The DECIDE framework, as presented by Sharp et al. (2007), serves as a very useful guide for how to plan evaluations of products, such as software or hardware. The framework provides us with a checklist that contains important issues to consider during the planning and gives an outline for how to perform evaluation in general, regardless of which evaluation methods are used.

1. Determine the *goals*.
2. Explore the *questions*.
3. Choose the *evaluation approach* and *methods*.
4. Identify the *practical issues*.
5. Decide how to deal with the *ethical issues*.
6. Evaluate, analyze, interpret, and present the *data*.

2.2.1 Goals and questions

The first three parts of the DECIDE framework partially overlaps with our usage of the GQM framework, so the goals, questions and metrics can be transferred to the DECIDE framework without changes. Goal 1 (Table 2.3) and goal 2 (Table 2.4) and their research questions correspond to the first and second item in the above list, while goal 3 (Table 2.5) is not related to the users and will not be considered here. Because of the nature of our project, the metrics chosen in the GQM will be the results from our chosen evaluation methods, such as observation notes and questionnaire replies.

2.2.2 Evaluation approaches and methods

According to Sharp et al. (2007, chap. 12), there are three main *evaluation approaches*. These approaches are not necessarily used independently of each other – they may be combined, for example by using usability testing in a lab to check the usability of certain features or design aspects of a product by having potential users perform a set of tasks, followed by field studies to see how the application as a whole works in a natural setting where users are in control. Observation is often used in the beginning of a project to gain insight into what is needed (Shneiderman and Plaisant, 2005), and may be followed by any other approach.

Usability testing

A defining property of usability testing is that the environment and tasks that are performed are controlled by the evaluator, not the user. Usability testing often happens in a controlled environment where test users are asked to perform a set of tasks while being observed by evaluators. The focus is often on the types of errors or hiccups the user encounters while performing the tasks. Questionnaires and interviews may also be used to gauge user satisfaction. These types of tests usually remove the testers from their regular environment, and as a result, it is mostly suited to testing certain aspects of an application and how well they work, not how well the application works in the environment where it will usually be found.

Field studies

Field studies, as opposed to usability testing, are done in the natural environment where the product will normally be used. The goal is to understand how users naturally relate to it, and how the product mediate their activities. Field studies can be used for identifying opportunities for new technology and establishing design requirements early in the process, or they may be used for facilitating introduction of new technology and evaluating technology later on. Results from a field study may be notes from observation, interviews with users, scenarios and so on, and will mostly be qualitative in form.

Analytical evaluation

Unlike the two previously mentioned approaches, analytical evaluation does not involve end-users at all. Expert evaluators will use two categories of evaluation methods; inspections and theoretically based models. Inspections include heuristic evaluation – using knowledge of typical users and guidelines or standards to identify usability problems – and walkthroughs of the prototype using given scenarios. Theoretically based models are used mostly for comparing efficacy of different options for user interfaces in the same application, and choosing optimal arrangements and placements for features and components.

Methods

Each of these three approaches may use a variety of different methods, some of which are mentioned in the descriptions. The main methods used in evaluation are:

- Observing users
- Asking users for their opinions
- Asking experts for their opinions
- Testing users' performance
- Modeling users' task performance to predict efficacy

Using interviews, informal discussions and questionnaires are good ways of asking the users for their opinions, which can be done within both usability testing and field studies. Observation of users is mainly done within field studies, while testing of users' performance is done within usability studies, but simple observation can also be relevant for usability testing. Asking experts for their opinions and modeling users' performance are methods used within analytical evaluation. The chosen approach and methods for evaluating our goals will be expanded upon in Section 2.3.

2.2.3 Practical issues

Potential practical issues with the evaluation process are important to identify beforehand. They could be related to which users evaluators have access

to, availability of resources such as facilities and equipment, whether or not evaluators have sufficient knowledge about the process and whether or not they have realistic schedules and budgets. Testers or evaluators of a system should preferably be representative of the actual end-user target group, with respect to skill level, age, gender, cultural background and so on (Sharp et al., 2007), but practical restrictions could prevent this from being possible. Issues related to facilities and equipment may be whether or not to film, how comfortable users are with this filming, how to capture data, where to perform the tests and so on.

2.2.4 Ethical issues

Evaluating products often involve users and gathering information, so there are ethical issues to consider. People's privacy must be upheld, and any recorded personal data about users should be kept confidential. Users should be given sufficient information about the evaluation they are taking part in, which kind of information will be gathered and how it will be analyzed. They should also be free to withdraw from the evaluation process at any point. The type of information that is gathered influences how the information should be dealt with, but ideally, usability reports should not contain comments or descriptions that make it possible to identify users that participated in the test.

2.2.5 Evaluating, interpreting and presenting data

The final point of the DECIDE framework is concerned with the what and how of data gathering, analysis and presentation. The approach and method chosen will obviously affect which type of data is collected - observation of users or interviews will result in different data than theoretically based models for user interfaces. One method, such as interviewing a user, may also result in both quantitative data (age, years of experience, gender and other closed questions) or qualitative data (open questions and opinions). It is also important to consider the following properties of the methods chosen (Sharp et al., 2007, chap. 8):

Reliability A method is reliable if it could produce the same results under the same circumstances on separate occasions. An informal interview

with a user would be highly unlikely to be reliable (but could still provide interesting data), while a lab experiment would be more likely to, if performed correctly.

Validity The validity of a method is concerned with how well a method measures what it is intended to measure. As an example, choosing a laboratory experiment to find out how certain users appropriate new technology in their workplace will most likely not result in relevant data.

Biases Bias happens when results are distorted. For example, asking leading questions may affect the results of interviews, and observers may fail to notice things they deem unimportant during observation.

Scope The scope of a study relates to how much the findings can be generalized. For example, a method that simply counts the number of keystrokes needed to achieve a goal without errors is not well-suited for describing how a novice may use the system.

Ecological validity The ecological validity of a study is concerned with how the environment in which the study is conducted affects the results. For example, a laboratory experiment happens in a very controlled environment compared to a normal workplace. User behavior may also be affected by being aware that they are being studied – this is known as the Hawthorne effect (Sharp et al., 2007, chap. 13).

2.3 Our process

This section will explain the choices we made concerning our evaluation process and why said choices were made. We start off by giving a short description of the literature study, which will be used to answer RQ3.1 and RQ3.2 about the maturity of the technology and to gain an insight into previous work with interactive campaigns. Following that is an explanation of which evaluation approach we chose, which methods we will use and how we will use them. As a result of these methods, we will end up with quite a bit of data, and this data must be handled well and we must be aware of biases and other previously mentioned properties.

2.3.1 Literature study

Our literature study will consist of a short, informal literature review on interactive campaigns and related concepts, such as interactive art installations, and following that we will look into the technology that will be used and which possibilities exist. Research questions RQ3.1 and RQ3.2 seek to find out how mature the technology is and which options there are for developing motion-controlled applications, and these questions will mainly be answered in the literature study. The implementation process will also shed light on the maturity of the solution we chose.

2.3.2 Prototype development

To be able to observe, evaluate and understand how people respond to and use interactive campaigns in a public space, we will need an application to put into the field and test. As a result, we will have to develop a prototype using the technology we chose based on the literature study. Global Illumination wanted us to use the OpenNI framework so that we would not be bound to one certain motion sensor device, like the Microsoft Kinect, so that is our starting point. However, OpenNI is multi-language and cross-platform, so there are several options for development environment.

Given that this type of technology is still at an early stage, it was obvious from the start that our development process would most likely result in a prototype, not a finished product. It is meant to test how well this type of development works and give us something to evaluate in the field. The focus in the development of the prototype will be to make it easily adaptable, so that multiple aspects can be tested without major software changes. For example, it would be beneficial to develop an application that works with both a single and multiple users being active at the same time. In addition to this, because of the marketing aspect, we want the application to cater to quick changes in what it advertises, for example by simply switching out the graphics.

2.3.3 Evaluation approach

Because of the nature of the project, some evaluation methods typically used in evaluating software may not be of much use to us. This type of interactive

application, whether in the form of a game or an art installation, has some different aspects (fun, immersion and so on) than a traditional information system, which is mainly focused on increasing effectiveness and efficiency in the completion of a task (Trifonova et al., 2008, p. 55). In addition to this, our project seeks to explore how users respond to using this type of interaction in advertising campaigns. The overall goal is to see how (or even if) users interact with the application, and controlled studies such as usability testing will not give us particularly useful data, because that approach involves telling users to do a specific thing in a controlled environment (Sharp et al., 2007, chap. 12). What we want to find out is how users relate to the application without instructions and in a public space. To evaluate this, we will have to build a robust prototype, put it out into the field and simply see what happens. The main evaluation methods will be observing users and asking for their opinions, through observation in the field, a questionnaire and interviews.

2.3.4 Observation

Observing users in the field allows us to understand the context of a user's activity and will give us insight that could be impossible to get from other methods (Sharp et al., 2007, chap. 7). The downsides are that it is very time consuming and can give us huge amounts of data to analyze, and often in a fairly unorganized form. That said, because of the fact that we seek to find out how our prototype is used in a public space, it is important to observe actual audience behavior and this cannot be replicated in a lab study (Müller et al., 2010b).

Placing an installation such as a game in a public space carries with it some special circumstances. For one, many of those who pass by it will never have seen it before and don't necessarily have any reason to use it (unlike an application the user has sought out or must use at work) or know how to do so. They might use the application in unexpected ways, and will also be influenced by their surroundings. For example, some might be embarrassed to use motion-controlled applications when other people can see them, and be afraid to look silly (Müller et al., 2010a). There will also be a huge variety of different people present, of all ages, cultural and social backgrounds, abilities and so on, and they can have different goals in mind and have different social roles. These things will affect how they use the application, which makes it

even more important to gather data from a natural setting instead of controlled experiments.

We will try to be more or less passive observers, but because we will be asking people to answer questionnaires and perhaps ask them some other questions, we will not be completely like “fly-on-the-wall” observers. Being a truly passive observer in the field is difficult, because being close enough to the installation to get useful information is difficult without interacting with the activities that are going on around us (Sharp et al., 2007, chap. 7), but we will not encourage people to try playing or actively encourage them to alter their behavior at all. We must, however, be aware that some might notice us and then alter their behavior regardless.

While observing, the focus areas will be the people currently interacting with the application and what they are doing. However, we should also keep an eye on their surroundings, see whether there is an audience, and try to understand their goals from their behavior, and pay attention if they express any particular feelings. The result of these observations will be observation notes, with information about time, place, which game version was tested and what happened, that we will need to go through and analyze afterwards.

2.3.5 Questionnaires

We chose to use a questionnaire for collecting some demographic data and for getting users’ opinions on certain questions related to their gaming habits, our game and their feelings about such a game. Using a questionnaire instead of interviews for the main method of asking users of their opinion allows us to ask them all the same questions and get answers in a quantitative form that allows us to easily analyze the responses.

Our questionnaire is inspired by the EGameFlow scale for measuring enjoyment of e-learning games (Fu et al., 2009) and the System Usability Scale (SUS) for measuring usability (Brooke, 1996). The SUS is a simple usability scale with ten statements where the respondents have to choose on a five-step scale from “Strongly disagree” to “Strongly agree”. The statements refer to a general “system” and are meant to work for a variety of different applications. The SUS could, in theory, be used as-is for our game, but we chose to also incorporate some statements from the EGameFlow scale and some of our own statements. This was done because we wanted answers

to some game-specific questions and because some aspects, such as the marketing and the fact that the game was placed in a public space, were not covered sufficiently by either scale.

As previously mentioned, the SUS is a very general scale, so while we did choose to use a couple of statements from it and were inspired by others, we had to rephrase (and later translate) a little bit. The following two statements were used directly, but instead of referring to the application as “the system” we switched to “the game” and we also switched out “use” with “control” or “play”.

- I thought the system was easy to use.
- I would imagine that most people would learn to use this system very quickly.

We also added two questions about comfort level while using the system, inspired by the SUS statement “*I felt very confident using the system*”, where they were asked to rank how comfortable they were with playing the game in a public space and how comfortable they were with being “thrown into” the game without giving their expressed permission.

The statements from the EGameFlow are very game-specific, and the ones we chose to use were:

- I become unaware of my surroundings while playing the game.
- The game grabs my attention.

The EGameFlow scale also has statements related to collaboration, and inspired by these, we added a few statements about that. We ask participants to answer whether or not they played with other people and how they thought that affected how long they played and the level of fun, and there is also a question about whether they would play the game with strangers (because the game will be set in a public space).

In addition to the previously mentioned questions, we added a couple related to the participant’s attitude towards the advertiser and the advertised product and asked those who participated without an advertising campaign whether they would be comfortable with playing the game if it was a part of an advertisement. We also wanted to know to what degree they wanted to win when playing the game, to perhaps gain an insight into their motivation for playing.

Although we drew inspiration from two already existing scales for measuring usability and enjoyment, where one goes from 1–7 where 1 is strongly disagree and 7 is strongly agree and the other uses the range 1–5, we chose to adapt our scale to be identical for all statements of this type. This was done because it would make the questionnaire much easier to fill out as the participant would only have to relate to one type of scale in addition to some yes/no and multiple choice questions (used for the gaming habits section). Because we only used a few statements from each scale and also adapted the range of answers, we will not be able to use the answers to calculate a SUS score, but using a consistent scale does allow us to analyze and compare numbers to a greater degree.

Because of the fact that the game will be tested both with and without an advertising campaign, we made two versions of the participant questionnaire, while the questionnaire used for non-participants is the same in all cases. Translated versions of the questionnaires can be found in Appendix A. These questionnaires will be presented to the users on a tablet computer, so that the data is quickly gathered using Google Docs.

2.3.6 Interviews

Our interviews will be in the form of unstructured or semi-structured informal interviews. This means that there will not be a set of defined questions we will ask participants, but the questions will be more open-ended and the interviewee will be encouraged to speak freely about the topic at hand (Cappellini, 2009). The structured questions we want answered will be in the questionnaire, and any interviews will act as a way to gain a more in-depth understanding of what some users experienced, and following up on any comments they may have about the application. They can also serve as a way to get some information from users who choose not to participate and are unwilling to fill out a questionnaire, but more willing to simply have a conversation. Because participants are encouraged to speak freely, it is important to have an agenda about which topics to cover, so that the interviewee stays more or less on topic (Sharp et al., 2007, Chap. 7).

2.3.7 Ethical issues

As mentioned in Section 2.2, several ethical issues may arise in our evaluation process. These mostly relate to how users and the collected data will be handled. In our interviews, questionnaires and observations, some data about how the users relate to the application will have to be gathered, as well as their gender and age, but names and other data that make users identifiable are not necessary. However, we may wish to be able to contact specific users later on, in which case some way of contacting them must be available and we will perhaps have to gather e-mail addresses from willing participants. It is important that we do not misuse this information in any way.

The motion sensing technology we use also means that we could, in theory, capture and store how the users moved and what they did so that we could have access to it later on, which would also pose some ethical issues related to privacy concerns. Our choice of using observation instead of recording users alleviates this issue, but can affect the data we gather.

2.3.8 Practical issues

Perhaps the biggest practical issue we may face in the evaluation process is finding a suitable place to test the application. Because of the advertising aspect, the initial goal is to find a public space, such as a mall, where we can conduct the field studies and also have the application be part of the advertising campaign. Some of the research goals can be tested wherever there are many people present, while the ones related to marketing may need to be tested in a setting like a mall or a store. To test various aspects of the application, we will do some testing in crowded places on campus in addition to the commercial location, because we have easy access and it is a more prototype-friendly environment.

2.3.9 Data gathering

From our field studies – in the form of observation, questionnaires and short interviews – we will be left with quite a bit of data in different forms. The questionnaire will result in quantitative data (because the questions are multiple choice), while the observation and interviews will result in qualitative data in

the form of observation notes, quotes and so on. When choosing evaluation methods, we have to keep in mind the properties of evaluation methods that were presented in Section 2.2.5 – reliability, validity, bias, scope and ecological validity – and how these affect the data we will gather.

Reliability

Because of the fact that our evaluation approach is a field study, reliability could be a problem (Sharp et al., 2007). A carefully controlled experiment would be easily replicable so one could check whether or not it produced the same results, while getting the same results on separate occasions in a field study might be difficult because there are so many factors that are uncontrollable. Observation also depends highly on those who perform the observation and what they deem important, which can vary based on their previous experience (Spyridakis and Wenger, 1989).

By keeping observer bias in mind, and performing multiple field studies in different locations, we hope to alleviate some of the reliability issues that could occur. Testing in multiple locations allows us to see whether certain behaviors occur in only one place or whether they are common overall, and getting a larger data set is good for both the observation and the questionnaires. Not only will we see whether or not the same observations occur multiple times and more easily be able to identify outliers, but we will also simply gather a larger data set that is less likely to be skewed by location-specific events.

Scope

We must keep in mind that making any generalizations about how users relate to a broad category like motion-controlled applications in a public space is difficult after studying only one type of application with a limited number of participants, independent of the choice of evaluation methods. This means that the scope is not very broad. Many factors come into play in field studies, particularly of an application people don't actually have to use, and they will affect to which degree generalizations can be made based on the results of our study.

The type of application, its features and even the graphical design of it can affect who it appeals to and which user groups choose to interact with

it. This will, in turn, affect the data that is gathered and thus the scope. It might, for example, be difficult to get the people who are most skeptical about the concept or who simply don't understand it to answer questionnaires or interviews, while those who actually used the application and enjoyed it could be more willing to answer. This will distort the data and give it a positive bias, and make it difficult to draw conclusions about potential user groups as a whole. We can, however, look for trends in the results and combine results from the questionnaires with data from our observation notes, to give a more accurate portrayal. It is also worth noting that as long as we are aware of the limitations in scope, plenty of interesting data can be gathered from the questionnaires alone.

Validity

The validity and ecological validity of our study will be affected by how we perform the evaluation. Because we are interested in motion-controlled applications in a public space, we should try to keep the testing environment similar. A public space serves functions that affect how people behave (are they taking a walk or on their way somewhere) and a large variety of different people are likely to pass by (Müller et al., 2010b). By placing the prototype in the field (mall, store, etc.) and observing what happens, the ecological validity should be fairly good, as the environment surrounding it will be more or less the same as the actual environment the application could potentially be used in.

The fact that one of our evaluation methods is observation means that we do have to be aware of the Hawthorne effect, where people's behavior is affected by being studied (Sharp et al., 2007, Chap. 13), and we might also see this in the interviews and questionnaires. However, observing how users actually relate to the game and asking them for their opinions is what is most interesting to us, so in order to measure what we actually want to measure – and maintain the study's validity – methods like observation and interviews are necessary. The problems relating to scope mentioned above also come into play here. If we want to measure what passers-by think of using a motion-controlled application in a public space, asking only those who choose to stop and play will not give us valid data. Using observation

of everyone's behavior around it can be a way to alleviate this problem, in addition to creating a questionnaire for those who chose not to play.

Bias

Bias occurs when those who perform the method, which in this case is us, cause the results to be distorted. In order to avoid distorting data, we should remember to avoid leading questions in questionnaires and interviews, as this is likely to affect how users answer the questions. Using statements in the questionnaire, such as we have done, can cause some bias in that the respondents might be careful about choosing the extremes on the scale (strongly disagree or agree) or agree with the statements as presented. Adding both positive and negative statements can alleviate the second problem, but will add more items to the questionnaire and make it longer, which can be a problem in a field study where the participants are often concerned with other things than answering a questionnaire.

Bias can also occur during observation, because this depends on the observers catching all the relevant details and they might have their own ideas about what is important to note and what is irrelevant, which can distort the data that is gathered (Sharp et al., 2007). Some interactions with a motion-controlled application in a public space can also be very short or subtle, which makes observation a bit more difficult (Müller et al., 2010b). Simply not noticing interactions or not seeing them as important enough can distort the data, and we must be aware of this.

An important and noteworthy bias that can occur in our project is the fact that people who have already tried or noticed our installation most likely will be biased when answering our questionnaire, as they have already been positively influenced by the installation and will give their answer thereafter. As mentioned, this will affect the scope of our study.

Literature Study

This chapter will present the relevant theory related to the project and the findings from our literature study. The literature study mainly covers two areas: interactivity in applications and installations and the technology that will be used. The goal of the literature study is to learn about various ways users can interact with applications of different types and what has been done previously with public interactive displays, and finding out which technology to use when developing our prototype. Two of the research questions, RQ3.1 and RQ3.2, are highly relevant in the literature study, as they are concerned with the maturity level and readiness of the technology.

3.1 Interactivity in applications

After over 20 years of having user interfaces be dominated by typical graphical user interfaces consisting of windows, icons, menus and pointing devices, there has been an increase in the interest in new interfaces based on multi-touch technology and gestures in recent years (Petersen and Stricker, 2009). This section will present a short overview of the history of user interfaces, and move on to talking about how gestural interfaces can be used and what the potential issues are.

3.1.1 Short history of user interfaces

According to van Dam (1997), we can identify four separate generations of user interfaces, where each lasted for several years and worked well on the hardware they had available at the time. Although there are probably other ways of structuring this history, this one shows the evolution that has taken place quite well. The first generation refers to using punch cards as the input

24 LITERATURE STUDY

method and line-printer output. Punch cards were used for data entry for several decades, but as far as more modern computers go, they were used in the 1950s and early 1960s. During the 60s, time-sharing on mainframes using terminals or minicomputers became possible, and the user interface changed from punch cards to typing in commands with parameters, using alphanumeric displays. This is still used in many Unix-based systems, for example, and using a command line shell is still preferred by many for certain uses.

While timesharing and command lines were still the dominating user interface, Xerox PARC were doing work that would eventually lead to the third generation of user interfaces, so-called WIMP graphical user interfaces (Windows, Icons, Menus, Pointing device) on raster graphics-based monitors. They were made popular by the release of the Macintosh in 1984, and although we have gone from monochrome to color displays and there have been improvements made to the graphics, these types of interfaces are still very popular. The fact that these types of user interfaces have been so immensely popular over a long period of time means that they have been tried and tested, and some fundamental principles and well-tested guidelines have emerged (Norman and Nielsen, 2010).

van Dam (1997) refers to the fourth generation as post-WIMP interfaces, which relies on input methods like gestures and speech. He does not mention touch interfaces, which in some cases are still WIMP interfaces with one or more fingers as the pointing device, for example where there is a button or menu on the screen that the user is supposed to click on. In other cases, they might not be WIMP-interfaces, for example in the case of moving photos around on a multi-touch table, much like one moves paper copies of photos on a real table (Petersen and Stricker, 2009). Some mobile devices also use things like shaking the device to switch which song is playing, rotating or tilting a device to control a game and other gestures as control mechanisms. Norman and Nielsen (2010) refer to these types of interfaces that rely on using the body, whether it's for pinching a screen with two fingers to zoom in or using your body as the controller in a game as gestural interfaces.

3.1.2 Using gestural interfaces

Gestures as an interaction method has, as Norman (2010) describes, been used for decades, but their popularity and the buzz surrounding them has increased in the last few years. Gestural interaction is often referred to as “natural user interfaces”, but Norman (2010) claim that they are not necessarily so. Gestures are not necessarily easy to learn or remember, and the same gesture can mean different things depending on where you are in the world.

As we will see, gestural interfaces also lack some properties that help make good user experiences and that make graphical user interfaces successful. These limitations can be overcome by introducing guidelines, standards and so on, but this means that they must be treated as other types of user interfaces, and we cannot rely on them simply being natural. Of course, some of these critiques depend on what the interface is used for – moving your body to control a golf game is more likely to feel natural than pushing a button on a screen by hovering your finger in the corresponding position in front of a sensor.

Norman (2010) describe some of the fundamental principles of interaction design that are independent of technology, and how gesture-based interaction might – at first glance, at least – have some problems with fulfilling these. For a complete list, the article is available online, but a couple of examples are visibility and discoverability, which rely somewhat on each other. Some gesture-based actions, such as swiping across an email in the inbox to get the option of deleting it, are completely invisible in many touch-based systems, so you can't really know when something is possible and when something is not because there is no visible system status to show you that something is an option. Swiping across an already open email does not give you the option of deleting it, for example. There is also a lack of discoverability. The main advantage of using menus in a system is that they make available commands obvious, while in a gesture-based system the user already has to know which alternatives he or she has in order to perform the commands.

Despite the potential issues with gestural interfaces, they do have some definite advantages in some areas. Those that use the entire body as a controller instead of relying on a touch screen interface, have been particularly embraced within games and art. As Nielsen (2010) explains, despite some usability issues, the Kinect as a controller in a social game works very well.

Users quickly become immersed in the game, and their main goal is entertainment. There is also often a group setting, so new users can be introduced to the concept by other users who are already playing – and it is evident what they are doing to control the game. There are also fewer consequences if a mistake is made in an application like this, compared to a business-oriented application, for example. Games are also very varied, so it makes sense that two very different games – for example a whitewater rafting game and a dance game – would have different user interfaces, as games have a wide variety of activities. An example of using gestures and movement in a game is the Nautilus group game (Strömberg et al., 2002), where a group of users play in an interactive virtual space. The game presents a group of users with a story and is given the mission of saving a dolphin, which they have to complete by collaborating in navigating a diving bell under water. The diving bell is controlled when the users move across the floor, which contains pressure sensors.

Gesture-based interaction has been embraced by artists behind interactive art installations as well. Trifonova et al. (2008) discuss how artists often utilize emerging technologies such as mobile devices or sensors to provoke a reaction from the audience, such as reflection. For these types of installations to work well, a mix of disciplines is necessary, and the final goal is often entertainment or reflection instead of a business-related goal. They want users to become immersed in using the installation, and this gives them freedom to explore various input methods that regular applications might not necessarily have the option of trying without risking too much. The “game_of_life” art installation (Satomi and Sommerer, 2007) is an example of an interactive art installation that tracks body movements, in this particular case a person’s eye gaze. Users can walk through a virtual city by using their eye gaze movements. By looking at a part of the city, a representation of the user on screen will walk to that position, which affects whether other representations of people were visible or not. Another example of an interactive art installation is the lamascope installation as described by Costello et al. (2005). The installation was in the form of a kaleidoscope that reacts to the movements of the participants in front of it and reflects back an abstracted view of the participant.

3.2 Interactive public displays

Public digital displays can be found all around us in public and semi-public spaces, and their main purpose is often advertising or showcasing information (José and Cardoso, 2011). Some experiments with art installations of this type have also been tried, which we will see in Section 3.2.3. In the past, these public displays have offered very few ways for interacting with or responding to people around them and have often more or less been a digital version of a traditional sign, but this is changing with (multi-) touch and motion sensing technology becoming cheaper and more common.

Many of the studies done on interactive public displays have revealed an important problem: it is hard to get people to interact with them, and a commonly cited reason is social embarrassment (Brignull and Rogers, 2003). Brignull and Rogers identify that the key bottleneck in their study of getting people to interact with the public display is transitioning from peripheral awareness (doing some activity away from the screen, but being peripherally aware of it) to focal awareness (focusing on the screen, watching it being used and talking about it).

The Honey Pot effect

One effect that helped users move from peripheral to focal awareness, occurred when some people were standing around the installation and showing an interest in it, which led to a progressive increase in the number of people in the immediate vicinity of the installation. This was referred to as the honey pot effect (Brignull and Rogers, 2003). People seemed to be affected by the fact that other people were showing an interest, and saw this as a tacit signal to engage with them. They would then exchange comments about the display and perhaps even move on to direct interaction activities.

3.2.1 The Audience Funnel

The Audience Funnel, as presented by Michelis and Müller (2011), provides a framework for investigating public displays by describing the different phases that make up the interaction process, and it is a more detailed version of the phases/activities described by Brignull and Rogers (2003). Between these

phases we find thresholds that the users must either cross or that cause them to abort the interaction. The Audience Funnel was a result of a study done with an installation called Magical Mirrors (see Section 3.2.3, where four sequential displays in a store-front window were used, so multiple interactions while walking was possible).

Michelis and Müller (2011) identified six phases of interaction (or lack thereof), and they are described below and can be seen in Figure 3.1. Although some of the phases are relevant for non-interactive displays and (multi-) touch displays, others work well and can be used to compare, for example, an installation that relies on motion and gesture-based interaction to a multi-touch display.




Figure 3.1: The Audience Funnel (Müller et al., 2010a)

- 1. Passing by** Everyone who happens to walk by the installation and are sufficiently close that they can realistically see the display. In this phase, they have not yet shown any reaction, such as looking at it.
- 2. Viewing & reacting** Once a passer-by shows any observable reaction, such as looking at, pointing or turning his or her head, they are no longer simply a passer-by, but a viewer. The change, which can be as simple as a glance, can be difficult to observe, but this depends on factors such as observer placement, observation method and even technology used (eye-tracking can be helpful).
- 3. Subtle interaction** Once a viewer performs some sort of movement that is meant to cause a change on the display, such as moving closer to

the screen, lifting a hand or briefly stopping, he or she is referred to as a subtle user. This often happens at a distance and can happen while others are also interacting with the display.

- 4. Direct interaction** Users who interact directly with the display often place themselves centrally in front of the screen and interact with it for some time, and in more depth, by for example focusing solely on the display and its immediate surroundings.
- 5. Multiple interactions** When more than one display is present, users can choose to interact with multiple displays. Multiple interaction users can also include users who stops the direct interaction, but then comes back later to reengage the display.
- 6. Follow-up actions** Follow-up actions means actions conducted by users after direct or multiple interactions, by for example taking a photo of themselves or their friends interacting with the display and uploading this online or similar activities.

After performing observations with the interaction phases in mind, they found that the conversion rate between passing by and subtle interaction was 33%, which means that a large number of those who passed by did not interact at all. However, once someone had subtly interacted, 95% of them would break the threshold to direct interaction. 70% of these would perform multiple interactions, but the installation in question had four displays in a row that users would walk by, which affects this number. They also observed that passers-by tended to not necessarily stop at the first display, but see that something was going on there and then stop at one of the next displays, which can be a progressive approach to capturing attention. Both Michelis and Müller (2011) and Brignull and Rogers (2003) find that the most difficult threshold to pass is that from peripheral to focal awareness, or passing by to subtle interaction, depending on which framework you use.

3.2.2 Design considerations

Developing new types of displays that include interactivity, whether in the form of informative displays, advertising or art, require some special design

considerations that we don't typically see in traditional PC settings. Motivation for using the technology and the social context are hugely important, so attracting, motivating and engaging users and dealing with the nature of public interaction become central design issues (Michelis and Müller, 2011; Müller et al., 2010a).

Compared to typical situations where computing technology is used, getting users to interact with a public display does not start with the interaction itself – the display first has to grab the user's attention while he or she is simply passing by, as explained in the previous section. Their curiosity must be piqued and they must then be motivated to choose to interact with the display. The public display must then deal with issues related to interaction in public, which can include embarrassment and people wanting to maintain their social role, by for example avoiding looking "silly", or simply not wanting to stop and be on other people's way. There are thresholds between the phases mentioned in the previous section, and grabbing attention, motivating users and engaging them can help them cross the thresholds.

Attention

Müller et al. (2010a) describe some general models for attracting attention, and these are *behavioral urgency*, *Bayesian surprise* and the *honey pot effect*. Behavioral urgency refers to the fact that certain things, such as the abrupt appearance of a new object, moving or looming stimuli and some luminary contrast changes have captured attention. Bayesian surprise refers to the difference between what someone expected and what they experienced in a situation (or prior and posterior beliefs about the world), and adding elements of surprise that will be unexpected to the user can grab their attention. The last factor, the honey pot effect, has been explained in Section 3.2, and refers to the fact that a grouping of people around an installation can grab the attention of passers-by.

Motivation

People who are moving around in a public space are generally not looking for a public display to interact with, so they need external motivating factors to do so if they come across one (Müller et al., 2010a). In some types of displays,

such as informative displays (e.g. a touch display presenting a map of an area), there is a set task they want to achieve, while in others things like fun and degree of immersion matter greatly (Trifonova et al., 2008, p. 55). Malone (1981) makes the distinction between toys and tools, where toys are systems used for their own sake with no external goal (games, puzzles and so on) while tools are used to achieve an external goal (text editors, for example), where achieving said goal is a motivation in itself. Toys need to create some sort of goal or challenge to be motivating to use.

Not a lot of work has been done to try and understand motivating factors behind a user's activity in a case like this, but following the Magical Mirrors study, Michelis identified a set of motivating factors (Müller et al., 2010a) building on the work done by Malone (1981).

Challenge and control Challenge and control being motivating factors is based on the belief that being able to master something while still being challenged can be motivating and make users want to continue interacting. Goals may emerge from the challenge present in the game, and this is again a motivating factor.

Curiosity and exploration Piquing a user's curiosity by presenting him or her with new stimuli, can cause them to want to figure something out, and the curiosity can thus be a precursor to exploratory behavior. Wanting to explore to "solve" something that might be uncertain or incomplete can motivate users to continue, but one must take care to continue to pique their curiosity while also not making anything incomprehensible.

Choice Adding choice to an application – even if the choice is only imagined by the user – has been shown to increase motivation. It allows users to feel like they have control and increases their autonomy. Choices also mean that an individual can adapt things more to his or her liking. Too many choices, however, can reduce the intrinsic motivation.

Fantasy and metaphor Adding metaphors to enrich an application with fantasy elements also appears to be a motivating effect. Metaphors can also make an application more familiar or easier to use, like with the Magical Mirrors study. Distortion mirrors are a known concept, and this was then used as a metaphor to remind passers-by of a known situation.

Collaboration Making it possible for users to collaborate, by having one user's interaction affect the other people's interaction, can be motivating if one sees that collaborating with others affects one's own experience or for example makes something easier or more entertaining. Making the collaborative behavior visible is important, so that individuals can recognize the efforts of each other.

Interaction in public

According to Müller et al. (2010a), the fact that the interaction happens in public has some consequences that designers must keep in mind. Not everyone will be comfortable with interacting with a public display, and this can be affected by everything from type of application to control mechanisms to the user's personality. Shy or introverted people, for example, might be less willing to interact than those who are outgoing. Some might also not want to stand out too much and draw attention to themselves, and risk being approached by others while interacting with the application. They could also be concerned about sharing private information if the application videotaped them or wanted them to make a choice based on personal interest, for example. If the display is placed in a crowded space, standing in front of it could also disturb other passers-by and some might see that as impolite and not want to do so.

Public displays can also be affected by things like people lingering in front of it without using it, which would block it from being possible to use by other people. These are factors that designers of the application cannot control, because it is part of the nature of a public setting. This also includes things like the sun reflecting off of the display and making it difficult to see, or the weather affecting whether or not people stop and interact, perhaps because it is raining or too cold.

Modalities of interaction

When users interact with a public display, the interaction can be found somewhere on the scale from implicit interaction to explicit interaction. Users can move from implicit interaction to explicit interaction, and displays can support a variety of different interaction modalities. Müller et al. (2010a) present ten such interaction modalities that includes both explicit and implicit interaction.

1. Presence
2. Body position
3. Body posture
4. Facial expression
5. Gaze
6. Speech
7. Gesture
8. Remote control
9. Keys (keyboard, mouse)
10. Touch

Public displays and advertising

According to José and Cardoso (2011), digital signage or public displays have some specificities as an advertising medium. In order to achieve a strong user engagement with the public display, the interaction has to be meaningful to the user, but it must also be relevant to the overall goal of placing the display in public – which can be very different depending on what that goal is. These types of displays will also be very place-sensitive and their overall success can depend on a variety of factors in their immediate physical and social environment. For example, geographic location and proximity to the thing that is being advertised can matter a lot. They will also potentially be viewed by multiple simultaneous users with varying interests, so advertisers have to decide who to target and whether they want to target a broad or narrow group.

As Schrammel et al. (2011) explain, quite a lot of research has been done on how people relate to advertisements in general (yellow pages, posters and so on), but much less has been done in relation to digital displays. Their study does, however, show that digital displays installed in a public transportation system where people are stationary tend to receive a fair bit of attention (61% of users looked at them once), while location and placement of logos in a shopping street was deemed very important in order to gain attention while people were walking.

3.2.3 Examples

To give a sense of what has been attempted with (semi-) public interactive displays, we will present some examples in this section. The Magical Mirrors study, which has been referred to frequently above because of how it has been used to create the Audience Funnel framework and identify some motivating factors and design considerations, is one such example. According to Michelis and Müller (2011), the installation consisted of four displays placed next to




Figure 3.2: One of the displays from *Magical Mirrors* (Michelis, 2008)

each other in a store front window in downtown Berlin. The display showed a mirror image of the scene in front of it, and by using motion detection from a simple video camera, it added optical effects to the image, for example a ribbon following a moving hand or flowers growing from your hand. Some pictures from the *Magical Mirrors* installation are shown in Figure 3.2. There were also three large projection displays that showed screenshots of people interacting with the displays. The multiple screens meant that people could walk by one, notice that something was happening and then choose to interact with the following displays instead of immediately stopping or having to walk back.

Another example that was also placed in an urban environment, was the *CityWall* installation in Helsinki (Peltonen et al., 2008), which can be seen in Figure 3.3. The installation consisted of a multi-touch screen in a store front window in central Helsinki, and provided users with a timeline full of photos of the city, downloaded from Flickr. Users could then zoom in on the timeline



Figure 3.3: CityWall installation in a store front window in Helsinki (Morrison et al., 2008)

and organize photos by moving, rotating and resizing them. This allowed passers-by to stop and play around with the installation and explore the city through the images in the multi-touch display.

A less public, but still collaborative and semi-public, interactive display is the BlueBoard device, intended for both personal and collaborative use in a work setting (Russell et al., 2004). The BlueBoard is a 1.3m plasma display with touch technology and an RFID reader so that users can be identified by swiping their card by the reader. It allowed users to pull up personal information quickly, and to collaborate on sketching ideas, sharing content and so on. The original intention was for it to be a facilitator for rapid, personal use, but it developed into a social collaborative platform that supported fast encounters and spontaneous meetings.

We have seen how public interactive displays can be used within art and for displaying information and working, but they have also be used within advertising. One example of this is when Volvo ran an interactive commercial – in the form of a game – prior to a screening of the movie *Ratatouille* at



Figure 3.4: Virtual fitting room using augmented reality and Kinect (AR Door, 2011)

12 cinemas across the UK (Tipping, 2009; Wang and Føllesdal, 2010). The audience drove a Volvo through an obstacle course by holding their arms up into the air and moving them left or right depending on which direction they wanted the car to drive, and a camera registered which direction the majority moved in and changed the direction of the car accordingly. There were objects they wanted to avoid and gather, and the 12 cinemas competed to get the highest score.

Another example of use within marketing is the virtual fitting room seen in Figure 3.4. This was a display developed by the Russian company AR Door, which was placed in the clothing store Topshop in Moscow, Russia (Sterling, 2011; AR Door, 2011). It used the Microsoft Kinect to track the body in the scene in front of it, and used augmented reality technology to superimpose a 3D model of a dress on to the person who was being filmed/tracked. The program was controlled by pushing virtual buttons by using gestures. The installation allowed customers to virtually try on clothes and they could see both the front and the back of the dress on their body.

3.3 Motion sensing devices

Although motion capture and analysis has been possible for a number of years using a variety of techniques, it has not been relevant for the consumer market for very long. There are two main approaches within visual human motion sensing: marker-based and marker-less. Generally, marker-based systems require that people and/or objects wear markers which can be detected and they are then filmed from multiple views. From this, the 3D position of each marker can be obtained and thus the position of the subject (Stone and Skubic, 2011). This method can be very precise and is widely used within professional settings such as for developing video games and movies, but it is not very practical in many circumstances. For example, the amount of expensive equipment and effort it requires to get started, it is not suitable for in-home use with games or other applications, or for settings where sessions are short-lived, such as in a public place where people come and go.

The marker-less approach is generally based on video capture and extracting the silhouettes of a subject from the image. If this is done from multiple camera angles, a skeletal model can be built. According to Stone and Skubic (2011), this method provides good results but it is both expensive and requires a very controlled environment or a large number of cameras to alleviate potential problems, such as clutter in the background, lighting and speed of movement (Gallo et al., 2011).

However, more recently a new and significantly cheaper way of capturing 3D information from a scene has emerged. This approach requires no markers or body suits, and uses a combination of an RGB camera and a depth sensor to provide information about a scene using a single small device. This technology is significantly cheaper than the previously mentioned approaches. The Microsoft Kinect was the first commercial device of this kind, released in November of 2010 (Knies, 2011), and ASUS has also released similar devices that utilize the same type of technology as the Kinect.

These devices have made motion sensing technology available to the masses. Not only is it easier than ever to get started with development, but the devices are widely available off-the-shelf, which increases the potential audience for applications. They are also very interesting within research areas such as user interfaces, interactive art, robotics, 3D modeling and within medical technology, as we will see in Section 3.3.2.

3.3.1 Hardware

Both the Kinect and the ASUS Xtion Pro Live use more or less the same 3D sensor for sensing their environment, developed by the company PrimeSense (Borenstein, 2012, chap. 1). The sensor consists of an RGB camera, a depth sensor and two or more microphones, and uses these components to retrieve information about the scene in front of the device (see Figure 3.5 for pictures of the sensors). The depth sensor uses an infrared light source that sends out a structured light pattern, and a monochrome CMOS image sensor detects the reflected light. The chip then analyzes how these patterns are deformed and calculates depth data to capture a 3D image of the scene (Villaroman et al., 2011). This is the depth image output, which provides the (x, y, z) position of any pixel measured in mm (which can then be converted to projected positions on a screen), and it is also used to find human forms in the scene. The RGB camera outputs an RGB image of the scene, and using this in combination with the depth image provides additional information about the scene.



Figure 3.5: Motion sensors: Microsoft Kinect and ASUS Xtion Pro Live

The main difference between the Microsoft Kinect and the ASUS Xtion Pro Live are that the Kinect has a motor to move the sensor and it also has an array of microphones, while the ASUS Xtion Pro Live only has two. There are also some differences in which resolutions and frame rates they provide with the two different image streams, as well as their field of views – but all of these differences are relatively minor. ASUS reports that the Xtion Pro Live has an optimal distance range of 0.8m – 3.5m, a horizontal field of view of 58° and a vertical field of view of 45°. The Kinect has an optimal distance range of 0.8 – 3.5m (although the Kinect SDK provides support for “Near mode” as close as 40cm (Microsoft, 2012a)) and a depth range of 57° horizontally and 43° vertically (Gallo et al., 2011). Both devices give developers access to

an RGB image and a depth image, but there are slight variations here as well depending on which version of the devices you have and which frame rate you're willing to use. For example, both devices provide a 640x480 pixel depth image at 30 frames per second, but while the Kinect's RGB image at 30 frames per second is 1280x960 pixels, the ASUS device has an RGB image with a resolution of 1280x1024 pixels for the same frame rate. This has changed recently, as the Kinect previously offered only 640x480 pixels at 30 frames per second for the RGB camera as well, and the resolutions and available frame rates are likely to change in the future.

3.3.2 Uses

Devices that enable cheap motion capture, gesture recognition and so on have been viewed as being potentially useful for a variety of purposes. Although we are still at an early stage, several uses have been proposed and tested, and this section highlights some of them. It is by no means a complete overview of the possibilities, but it gives the reader an idea of the versatility of these types of devices and thus why they are so interesting.

Natural user interfaces in general

There has been some debate about what constitutes a natural user interface and whether gestural interaction necessarily fall under this definition (Norman, 2010), but the overall point is that gestures or human motion can be used to control applications. Some of these motions, such as swiping for "next" may feel natural (although they lack feedback), while other gestures are less so. van Dam (1997) refer to these types of interfaces as post-WIMP interfaces ¹, where motion, speech and so on are used as control mechanisms.

He et al. (2011) present an example of real-time gesture recognition using a Kinect to control the game Angry Birds by Rovio (Rovio, 2012). This is game that is not originally developed for use with motion control, so they wrote software that recognized fingertip motion and fired mouse click events at the correct gestures. Boulos et al. (2011) demonstrate how the Kinect can be used to control Google Earth through the Kinoogle interface, which is implemented using OpenNI drivers. Through a set of gestures, for example

(1) WIMP interfaces are based on **W**indows, **I**cons, **M**enus, **P**ointers

moving your hands closer together or further apart, you control and move the earth in the application. Kinoogle also lets users “walk” and “turn” in Google Street View, as well as control the camera angle using shoulder motions.

The WIN&I Media Viewer (Callaham, 2012) is another example of using the Kinect to control a regular application, in this case a media viewer. Through arm motions and “clicking” certain points on the screen, users can control images, movies, PDFs and so on in the WIN&I application.

Games

The Kinect was originally launched as an add-on to the Xbox 360 gaming system, so it makes sense that games are a very popular use for this type of device. The selection of games range from sports and fitness games to action adventure to family-oriented games (Microsoft, 2012c). Most games for the Xbox 360 let players control an avatar in the game by using their body, so that the limbs of the avatar in the game correspond to the limbs of the people playing and will move accordingly. Bleiweiss et al. (2010) demonstrate how the skeleton data can also be used to partially animate an avatar on screen, while parts of the body perform predefined animation sequences. After the initial release of the Kinect, several frameworks were developed that made it possible to develop applications for multiple operating systems on a PCs, as we will see in Section 3.4. This means that developing games using the Kinect and similar devices is no longer restricted to the Xbox 360.

The medical field

Using motion sensing devices within the medical field has also proved to be of some interest, with very varied areas of usage. Rydén et al. (2011) demonstrate how the Kinect’s depth camera can be used to create real-time haptic virtual fixtures. This can be used for robotic or telerobotic surgery, to provide haptic feedback to the surgeon’s hand, and unlike using CT scans for the purpose, using a camera to provide a real-time representation does not rely on the section being operated on to be perfectly still.

The Kinect has also been suggested for use with in-home fall risk assessment, to measure a person’s gait passively in order to predict the fall risk, and compared to traditional web camera systems the Kinect showed

fairly good results (Stone and Skubic, 2011). Chang et al. (2011) also showed that the Kinect and similar devices could be used within physical rehabilitation. The system, called Kinerehab, let young adults with motor disabilities perform exercises in front of the Kinect, and measured how accurate these movements were. The system was very motivating for these young adults and they suggested improvements such as letting two or more people perform the exercises at the same time.

3D modeling and tracking

The Kinect's depth sensor does provide 3D data about the scene in front of it, and this has led to it being experimented with as an aid in 3D scanning as well. The Kinect has some advantages in that it can measure depth and color simultaneously, but it also has a relatively low resolution and accuracy compared to special-purpose 3D scanners. Nonetheless, Cui and Stricker (2011) has shown that with some extra work, the Kinect can be used for 3D shape scanning.

Oikonomidis et al. (2011) show how the depth and color images from the Kinect can be used to create a 3D model of a hand and to track this hand and its movements. They do this by isolating the hand and combining the depth and color information to create a model of the hand, which is then turned into a 3D model comprising of a set of geometric figures.

Public interactive displays

Section 3.2.3 gave some examples of public interactive displays, and although the majority of them did not utilize motion sensing devices, several projects have done so, most commonly in the form of a game or an art installation. The London-based art and technology collective Seeper developed interactive art installations for a WIRED Event (seeper, 2011), which can be seen in Figure 3.6 and for the launch of the XBox Kinect in Berlin (seeper, 2010). For the latter project, projections of users were shown on giant buildings in central Berlin, and participants took part in particle ball games and so on.

Another example of an art installation that uses motion sensing is the Looking Glas installation in Berlin, developed by Jürg Müller and his team (Michelis, 2011). Here a participant's silhouette is reflected on a smaller

screen, and they can use their silhouette to move footballs around, by kicking, hitting and so on. The display was placed in a storefront to encourage customers to play.



Figure 3.6: Interactive art installation by Seeper at a WIRED Event (seeper, 2011)

Robotics

Tölgyessy and Hubinský (2011) present the type of sensor used by the Kinect and ASUS Xtion Pro Live as a valuable sophisticated sensor within robotics in education. Traditional sensing techniques have been infrared, ultrasonic or laser sensors to sense obstacles, as well as RGB cameras in combination with blob detection, edge detection, object recognition and so on. The sensor present in the Kinect combines an RGB camera with the depth sensor, which can be very useful within robotics. It can be used for things like object avoidance and collision detection, object recognition, controlling robots using gestures, localization and navigation. Stowers et al. (2011) demonstrate that the Kinect sensor can be used to control the altitude a quadrotor helicopter in real time, with relatively good results.

3.4 Developing motion-controlled applications

Developing applications for 3D sensing devices like the Kinect and ASUS Xtion Pro Live is still in a relatively early stage. The Kinect was initially released for the Xbox 360, and as a result, multiple games for the Xbox 360 that used the Kinect were developed and released. However, there was significant interest in developing various motion-controlled applications for PCs, and several projects sought to make this possible right after the Xbox release – videos of the Kinect working with Windows 7 appeared a few weeks after it was released (Anthony, 2010). Although several people were working on hacking the Kinect, the first person to provide video proof that he had gotten access to video streams from the Kinect was Héctor Martín, and he open sourced his code, called libfreenect. This later developed into the OpenKinect project, which focused on creating free, open source drivers and libraries that enable the use of Kinect on several major operating systems such as several Linux distributions, Mac OS X and Windows and also has wrappers for Python, C++, C#, Java and several other languages (OpenKinect, 2012).

One of the first examples of an ambitious application using motion control on a PC is the DepthJS project from the MIT Media Lab (Bobleanta, 2010), which is a browser extension that allows users to control web pages by pointing, zooming, clicking and switching between tabs by using hand gestures (Zinman et al., 2010). They originally used libfreenect, OpenCV and worked only in Chrome, but has later expanded to supporting Chrome and Safari in OS X and Linux as well as Chrome in Windows, and now uses the OpenNI framework as a backend (Fritz et al., 2012). OpenNI is among the two most well-known ways of getting started with developing for these types of devices now, with Microsoft's Kinect SDK being the other one (Pheatt and McMullen, 2012).

3.4.1 OpenNI framework

OpenNI, where NI is an abbreviation of Natural Interaction, is both a framework and the organization behind said framework. The organization is industry-led and non-profit and was formed to actively encourage compatibility and interoperability of natural interaction devices, applications and middleware (OpenNI, 2011). The term natural interaction refers to interacting with a computer by

using body movement or sound, which can be seen as more natural than a graphical user interface. One of the founding members is PrimeSense, which provides 3D sensing technology for both the Kinect and the ASUS Xtion Pro Live (Villaroman et al., 2011). By providing the framework for developing applications for Natural Interaction devices, they hope to accelerate the introduction and adoption of these types of applications.

The OpenNI framework defines APIs developers can use for accessing Natural Interaction devices, and gives them access to vision and audio sensors, as well as perception software in the form of middleware that analyzes data from the vision and audio sensors. Some example uses are speech recognition, body motion tracking and capturing body/hand gestures (OpenNI, 2012). The framework supplies one set of APIs that sensor devices need to implement and one that middleware components need to implement. An abstract view of this can be seen in Figure 3.7. This architecture allows applications to be ported without additional effort to work on different middleware components and with different sensor devices.

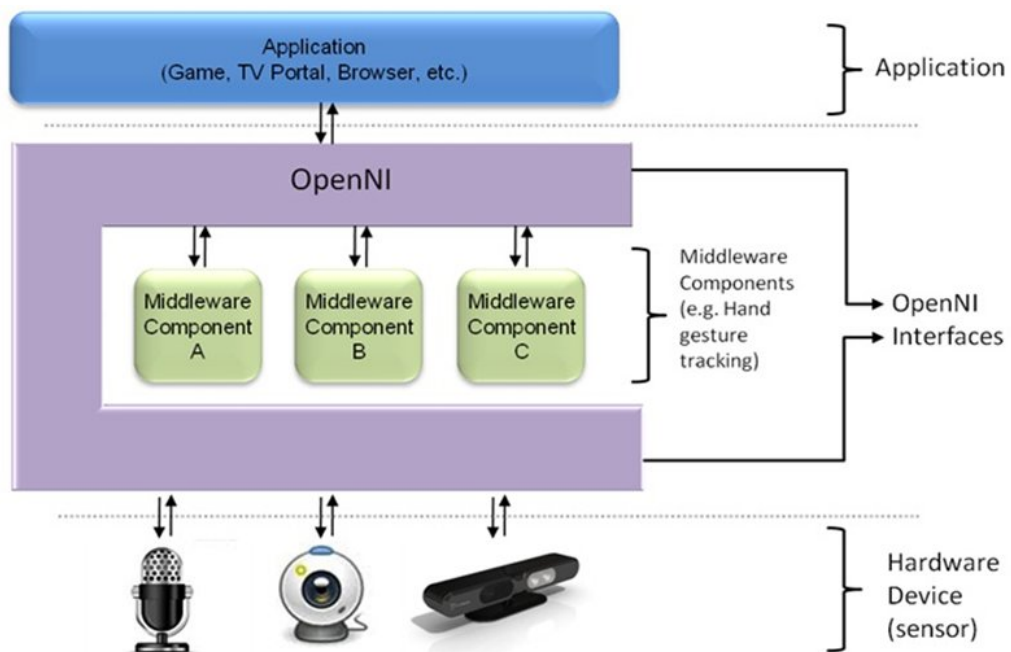


Figure 3.7: Abstract layered view of OpenNI (OpenNI, 2012)

The sensor components – or modules – currently supported by the

OpenNI framework are 3D sensors, RGB cameras, IR cameras and audio devices, as long as these devices implement the necessary interfaces that OpenNI requires. The supported middleware components are full body analysis, hand point analysis, gesture detection and scene analyzer middleware. The most important one is perhaps NITE, developed by PrimeSense (2011). NITE's middleware is used by OpenNI for providing support for skeleton tracking and hand gestures. Given OpenNI's modular nature, the NITE middleware could be replaced by, for example, Kinect SDK's skeleton tracking as long as a Kinect SDK to OpenNI bridge was used without needing to rewrite or even recompile any code, but using NITE is the most common approach. OpenNI without NITE provides developers with depth and RGB image data, but body/skeleton tracking and so on requires middleware such as NITE.

Although OpenNI is originally written for use with C or C++, wrappers for C#.NET and Java also come with the default installation (Pheatt and McMullen, 2012). The fact that the framework is multi-language makes it very versatile and available to a larger and more varied community. This also means that developers wishing to make applications aren't necessarily tied to a specific platform or device, unlike Microsoft's Kinect SDK (Microsoft, 2012a), where you're required to use Microsoft's systems and the Kinect. Independently written wrappers or bindings, such as the simple-openni wrapper for the Processing language and IDE (Rheiner, 2012) and the Zigfu Development Kit, also add to the versatility of OpenNI and makes it available for a multitude of purposes.

Zigfu Development Kit

One of the easiest ways to get started with developing using OpenNI is to utilize the Zigfu Development Kit, also known as the ZDK. Zigfu was founded by two previous PrimeSense employees, and focuses on making software to simplify the process of making motion-controlled applications for the web (Zigfu, 2012a). As a part of this process, they have made a very simple package that will install everything necessary to start developing using OpenNI – the OpenNI framework, NITE and sensors for the Kinect and the ASUS Xtion Pro Live. In addition to this, they provide a browser plugin with JavaScript bindings to make it possible to develop web applications using only HTML, CSS and JavaScript that utilize skeleton data and gesture recognition, as well

as bindings for using OpenNI/NITE or Kinect SDK with Unity3D. This means that it's possible to build games for PCs in Unity3D that utilize skeleton data and gesture recognition. Flash bindings are also planned. Unlike OpenNI, these bindings are not free to use, but trials are provided and using the browser plugin is free with an added watermark.

The ZigFu Development Kit is fairly new, and only a beta version without documentation and with limited functionality was available at the time when we chose which technologies to use for our implementation. However, more functionality and documentation has been released throughout the spring of 2012, with calibration-free skeleton recognition being an example of increased functionality.

3.4.2 Microsoft's Kinect SDK

In February of 2011, following the November 2010 release of Kinect for Xbox 360, Microsoft announced that a software development kit for Windows would be made available (Knies, 2011). A non-commercial beta version was released on June 16th 2011 and a commercial version (1.0) was later released on February 1st 2012 (Eisler, 2012).

The SDK contains drivers for Kinect that are compatible with Windows 7, and allows developers to develop applications using .NET compliant languages like C++, C#.NET or Visual Basic, using Visual Studio 2010 (Pheatt and McMullen, 2012). This means that the Kinect SDK can be used with the XNA Framework – which is Microsoft's set of tools to facilitate game programming for Windows, Xbox 360 and Windows Phone 7 (Miller and Johnson, 2010) – to develop games for PC that use skeleton data, audio and video from the Kinect device. Prior to the release of the SDK, this was only officially an option for the Xbox 360.

Like OpenNI, the Kinect SDK gives developers access to raw data streams, such as depth and RGB, and full-body skeleton tracking. It also comes with audio capabilities (Microsoft, 2012a). The audio capabilities are used to support speech recognition in multiple languages, including English, Italian, Spanish and Japanese. Unlike OpenNI and NITE, it does not provide gesture recognition capabilities for things like swiping, waving and pushing and specific hand tracking out of the box, but third-party developers can and have built their own gesture recognition libraries based on the skeletal data

that the Kinect SDK provides access to (Catuhe, 2011). An advantage of the Kinect SDK early on was that the skeleton tracking algorithm did not require a calibration pose in order to work and the Kinect SDK skeleton had a few additional joints compared to OpenNI (Borenstein, 2012, chap. 7), but OpenNI has since added support for calibration-free skeleton tracking.

The Kinect SDK is under continuous development, and Microsoft released version 1.5 on May 21st 2012 (Microsoft, 2012b). This included, among other things, support for seated skeletal tracking and skeletal tracking in Near Mode and an SDK for face tracking to provide a real-time representation of a person's facial features. With it they also released more code samples and *Kinect Studio*, which is a tool developers can use to record Kinect data and play it back later to aid in the development and testing process.

Design and Development

Throughout this chapter our prototype will be presented. In order to evaluate how users relate to an interactive campaign in a public space, it was crucial that we had some sort of application to use in our field studies. We start by discussing the prototype concept, the most important choices we made while coming up with an suitable concept and the reasoning behind these choices. Following that is a short section describing which technology we chose to use and why – based on information from the literature study, our desired prototype and our previous knowledge. We then present a short overview of the implementation of the prototypes.

4.1 Application concept

While developing the concept, there were two important overall requirements for the prototype – one of them being reusability and the other one being that getting started and using it should require very little effort by passers-by. These focus areas bring some limitations with them as to which type of application we could develop and which features were relevant, but they also have some considerable advantages, which we will expand upon later in this section. The idea of a simple game was suggested fairly early on in the process, as it would allow us to test some interesting research questions related to interaction between users and it could be made general enough to work for many different types of clients. GI also felt that a game would be a good way to capture people attention and give them a reason to stop in front of the screen.

4.1.1 Brainstorming

We developed a few concept drawings at the beginning of the semester to illustrate for GI what kind of usages our motion-controlled application could entail. These were mostly used by us when we brainstormed ideas for an interactive application, but they were also used further by GI as easy to understand explanations when they had meetings with malls that potentially would like to display such an installation (a shopping mall, for instance). Figure 4.1 shows our first concept drawing that tries to illustrate the concept of an attention grabbing poster. The idea behind this concept drawing was that people would become aware of being a part of something when walking by and seeing their own silhouette on the screen.



Figure 4.1: Concept drawing 1: Attention grabber

When we discovered the virtual fitting room, developed for the clothing store Topshop by the Russian company AR Door (Sterling, 2011), we played around with the idea that several stores in a mall could place overlays relevant to the store on a person's silhouette. For instance, adding clothes like they did at the Topshop store or changing the hair style to encourage a visit to the hairdresser's, as illustrated in Figure 4.2.

Another aspect that we wanted to consider was a multiplayer mode. If crowds of people are able to play at the same time, the hot pot effect will be




Figure 4.2: Concept drawing 2: Change of styles

prominent and more people will most likely notice the campaign. A concept drawing illustrating a simple ball game for multiple people was made, and can be seen in Figure 4.3.

After presenting the concept drawings to GI we discussed the features a bit more. It was decided that we should add a container for collecting balls, and there was some discussion around adding features like a high score list or social media integration, but as we will explain in Section 4.1.3 we did not choose to implement this.

4.1.2 Reusability

The reusability focus meant that the type of application we developed had to be fairly general. GI's initial idea was to place the application in a mall, and have multiple stores advertised within the same application, so things like the virtual change of styles would not fit our needs, so we soon moved towards a more general-purpose application and more specifically a simple game.

By choosing to develop a simple game with graphical elements that could be replaced easily without changing the gameplay, reuse of the applica-



Figure 4.3: Concept drawing 3: Engaging crowds

tion in multiple environments is straight-forward, and this would be helpful both in our evaluation phase, where we wanted to test our application in several different settings, and for GI when they use the application later on. This would also allow us to get answers to the two research questions related to marketing, as the graphics could be switched out to advertise for different advertisers or products (or no product at all for some tests) that we could then ask users about in questionnaires or interviews.

4.1.3 Low threshold for interaction

The other main focus was that the threshold for interacting with the application must be very low. Going from being simply a passer-by or viewing and reacting to subtle and direct interaction (Müller et al., 2010a) must require very little effort from the user. This meant that we should try to avoid needing a calibration pose or gesture to get started, which has previously been necessary with OpenNI when using the skeletal tracking. The control mechanisms also had to be very simple and should not require the user to have learn something before getting started with the motion-controlled game

or application. The overall concept also had to be simple, because users should be able to quickly understand what the objective of the game is and how to reach that objective. The fact that motion-controlled applications is a fairly new concept might be a hindrance in this case, because we cannot assume that people will immediately understand that their body controls the application.

Silhouette versus skeleton-based model

One of the choices we made to make it more obvious to people that their body was the controller was the use a silhouette instead of a 2D or 3D model of their “character” on the screen. We hoped this would provide a greater sense of recognition, in that the passers-by would be able to see a real-time representation of themselves on the screen and see that it moved as they moved. Using silhouettes instead of a skeleton-based model also required no calibration, which is very useful for keeping the threshold for interaction low. At the beginning of 2012, calibration-free skeleton recognition had not been implemented in the Zigfu Development Kit (described in Section 3.4.1), but this has been implemented later. Leaving out a calibration pose or gesture and using silhouettes that were a direct representation of the people in the scene would also hopefully attract the attention of passers-by who glanced at the screen, because they would see themselves immediately without having to perform a pose or a gesture.

Attention

As Müller et al. (2010a) points out, within human-computer interaction one often assumes that the user is aware of the computer or application initially. In our case where the application is placed in a public space, this is not really the case, so grabbing a user’s attention is important. In a public space like a mall or a store, a display containing a motion-controlled application must compete with other stimuli, such as other people, signs or the user’s initial objective (going somewhere, meeting someone).

By immediately displaying the user’s silhouette on the large screen when he or she passes the display, we hope to catch their attention visually. Not only is this a visual change and appearance of a new object, it can also be

surprising because it might be unexpected. Having objects fall down towards the silhouettes might also instigate some behavioral urgency and make people want to react to the object. Depending on how far away from the display they are, this might happen in their peripheral vision or in their actual field of view. The presence of a large colorful display/projection in an unexpected place is also a way we can hope to grab the attention of passers-by. One of the ways in which we could have attracted attention to the game, but chose not to, is by using sound. In-game sounds can make passers-by who do not see the screen aware of its existence, but it can also be an annoyance to both the general public and the people who work nearby the installation all day. There are also a lot of other sounds in a public space that could drown out the game sounds.

Motivation

People in a public space often have other goals in mind than playing a game they come across, so motivating factors matter greatly (Müller et al., 2010a). By developing a game, there is an inherent challenge in trying to beat the game that might be motivating to players. Using motion sensing technology in a public space is not a common thing, so hopefully it will pique the curiosity of some of the passers-by and cause them to want to explore what is going on. The game will also let people collaborate freely (or compete if they so choose) because in the multiplayer mode users can come and go as they please, and they can work together on accomplishing the set goal. Users can also choose how much they want to interact with the game, so that can give them a sense of control and motivate them to continue playing in the way they want to.

Sacrifices

Although keeping the threshold for interaction as low as possible was a major goal of ours, it also means that we give up some potentially interesting features. A public high score list could be a huge motivating factor in making people use the game for an extended period of time, or even just as an incentive to use it in the first place. However, using any form of high score list would require some form of registration, either in the form of writing in their name or perhaps taking a photo using the ASUS sensor, which we would need to

get permission from the player to do. This would be too time-consuming for a simple passer-by and perhaps require them to give up information they did not want to share.

Implementing things like social media could also increase the attention and make people motivated to play given to the application. Some early suggestions included adding a Facebook component, where people could post to their wall that they had played and how well they did. The same problem as with the high score list appears here, though, where it requires a much greater effort from the user and they cannot easily come and go as they wish, join others who are playing and so on.

4.1.4 Final concepts

Game

The concept we ended up using was a game in which one or more users have to gather similar objects that fall down from above into a container of some sort. Once N similar objects had been gathered, they would be shown some form of advertising campaign, such as an image with an offer or a coupon code that they could use in a store that was a part of the campaign. This is a fairly simple game, but it is very easily adaptable and allows us to test multiple modes.

The objects that fall down are controlled by moving your body so that the silhouette collides with them. Because your silhouette appears on the screen as soon as you pass by the sensor, you're immediately a part of the game and this fulfills the goal of having a low threshold for interaction. The objects that fall down, the background, the container and the winning image can be switched out easily, thus changing the appearance of the game and fulfilling the goal of reusability.

Interactive poster

In order to test a form of interactive campaign that did not require as much effort from the user as a game does, we also implemented a form of interactive poster. This poster started out as a black screen, and every time a new user was registered as passing by, a puzzle piece would fall down. This would

reveal a poster underneath, slowly or quickly depending on how many users were present or passing by the screen at the time. Unlike the previously mentioned game, this could easily be implemented as a web application using ZigJS, which is described later in this chapter.

This additional prototype would allow us to test whether or not something simply happening on the screen as a person was passing by would grab their attention and make them curious about what was going on.

4.1.5 Other aspects

In addition to focusing on reusability and requiring very little effort from users to get started, we also wanted the game to be easily adaptable so that we could test a few different modes and user behaviors. For example, we wanted to test multiplayer mode, where everyone who passed the sensor was automatically a part of the game, and compare it to a one or two player mode where people in the background appeared on the screen but were not a part of the game. To accomplish this, we had to add some form of mechanism to select which users were active players.

One alternative here was to use the typical calibration pose often used in Kinect, but we chose to make users who wanted to play wave to the screen to be chosen, as a wave gesture seemed more natural than the Kinect pose (which consists of standing straight and holding your arms out to the side, bent upwards in a 90° angle at the elbow). This increases the effort required to join and goes against one of our main goals of simplicity, but testing how well it works is interesting and implementing it is fairly straight-forward. While the multiplayer mode is the simplest, it also means that people who pass by in the background can “ruin” things for the ones who are actively playing because their silhouettes will also be a part of the game, which the one or two-player mode avoids. Seeing how these two modes affect people’s behavior and opinion of the game is interesting.

Another interesting aspect is whether people who play the game choose to collaborate with others or compete, which can be tested in both the multiplayer mode and in the two-player mode, although the two-player mode enhances the competitive aspect by providing two containers that objects can be caught in. In multiplayer mode, only one container is provided and the players choose freely how they want to play. We think it is important that the

gameplay allows users to decide this for themselves by not providing any set limitations. This means that by observing how players use the application without direction, we can see how they behave and use this as a basis for answering our research questions.

4.2 Chosen technologies

After researching the various technologies available to us and discussing the options with Global Illumination, we made a choice of using OpenNI as the framework for communicating with the ASUS device and C# and the XNA framework for the game development.

4.2.1 Why OpenNI, C# and XNA?

GI initially wanted the application to run in the browser, but at the beginning of 2012, the technology for this was quite limited, with Zigfu's ZigJS (Zigfu, 2012a) or DepthJS from the MIT Media Lab (Zinman et al., 2010) being the most relevant options. Both of these utilize OpenNI on a lower level, but neither would be sufficient for implementing our primary game prototype.

ZigJS was still in such an early development phase that it was a closed beta without documentation, which meant that we were skeptical about basing our thesis project on it because of its maturity level. It would also not give us direct access to the RGB or depth image streams, but focused on providing gesture recognition and skeletal data, which makes sense given that applications would be written in JavaScript. In May of 2012, ZigJS has been further developed and building a JavaScript game using the skeleton tracking and/or gesture recognition could be feasible, depending on your choice of game. Our wish of using the players' silhouettes in the game instead of creating a model of their skeleton would not be achievable. DepthJS's primary focus seems to be on navigating the web, and a natural consequence of this is that the features in their browser extension and developer API were related to recording gestures, which did not fit our choice of game concept.

Because of these limitations and the fact that our application would be a prototype to be used in an evaluation process, it was decided that making

it into a standalone application would be acceptable. OpenNI was selected because it is open source, and supports multiple programming languages and multiple platforms (as explained in Section 3.4.1), and also because GI wanted us to use it. The Kinect SDK was only available in a non-commercial beta version in the beginning of 2012 – the commercial SDK was released as late as February 1st 2012 (Eisler, 2012).

The choice of using C# and the XNA framework was based on the fact that the OpenNI wrapper was available for C#.NET and using XNA simplifies many aspects of game development (Miller and Johnson, 2010). One of us also had previous experience with developing games using XNA. This means that we are locked to Windows as a platform for development and execution, which is a downside, but one both GI and we were willing to accept given that the overall goal was to develop a prototype mainly for evaluation purposes.

Unlike the game, the proposed interactive poster did not require much more interaction with the motion sensing device other than being made aware of when a new user appeared, which ZigJS was capable of doing quite well through its JavaScript API. This meant that the interactive poster would be a good way to test both an application that required less interaction from a user and to examine how mature the technology for developing motion-controlled applications using OpenNI in the browser were.

4.2.2 Using OpenNI

To use OpenNI, one must first download the necessary drivers, which often includes both OpenNI and NITE as well as the relevant sensor drivers. The quickest way to do this is, as mentioned, through the package provided by Zigfu, but they are also available from OpenNI's website (www.openni.org). To get started with development, getting the relevant wrapper (depending on choice of language) and referencing this in the code is sufficient. OpenNI comes with a configuration file that must also be used, and in this things like frame rate, resolution, mirroring of image streams and which nodes you want are configured. The configuration file is then used when creating the `Context` object.

The `Context` object in OpenNI is the main object that keeps track of the state of the production chains that are used by applications. One application can have multiple contexts, but they cannot share information,

which means that a middleware node such as a user generator cannot use the depth generator from a different context. The context is used when creating nodes, and each of these nodes will also have a *metadata* object. As an example, the metadata object for a depth node will have information about the resolution of the depth map.

The following is a list of the data generators provided by OpenNI, with a short description of which data and functions they provide access to, with some of them requiring middleware (scene analyzer, gesture and user recognition) while others being sensor-related (depth and image maps, audio). They are also referred to as production nodes, where some production nodes (such as a user generator) depend on data produced by another node (in this case the depth generator) which again depends on raw sensory data. This could be called a production chain, where one generator or node is dependent on another.

Map Generator This is a basic interface for the generators that produce a map, and has functionality such as cropping and alternative viewpoint.

Depth Generator This object generates a depth map, and includes functionality to return said depth map, get the position of a user and device max depth.

Image Generator This object generates the current RGB image map and returns this.

IR Generator Generates and returns the current IR image map.

User Generator Generates data related to figures in the scene, and provides support for finding the number of users, current users, the center of mass location for each user and a map of pixels that represent the scene where pixels that correspond to a user are labeled with the user ID.

Scene Analyzer The scene analyzer gets raw sensory data and provides a map with labels on each pixel to make the sensory data into a meaningful scene, such as labeling figures, the floor and the background.

Audio Generator Generates audio data, and has support for configuring the output, such as sample rate and bits-per-sample.

Gesture Generator This object enables gesture tracking, so once a gesture is added, the generator will look for this gesture being performed.

Hand Point Generator Enables hand detection and tracking, with methods for starting or stopping the tracking of a specific hand.

To give the reader an idea of how to use the framework, we will present a few minor code examples. They are by no means complete and does not include separation into methods, error handling and the like, but will hopefully provide some insight into how it is used. The code examples are in C#, as that is what we used. Listing 4.1 shows the basic structure used when drawing the depth map to the screen. A complete example can be found in OpenNI's programmer guide (OpenNI, 2012).

```

private Context context;           1
private ScriptNode scriptNode;    2
private DepthGenerator depth;     3
                                   4
// Upon initialization, create context and the depth node
context = Context.CreateFromXmlFile("path_to_folder/SampleConfig.xml", 5
    out scriptNode);
depth = context.FindExistingNode(NodeType.Depth) as DepthGenerator; 7
DepthMetaData depthMD = new DepthMetaData(); 8
                                   9
// Get resolution and frame rate of depth map for later use 10
                                   11
// Start drawing, either in a loop or thread 12
while(shouldDraw) { 13
    // Update depth map 14
    try { 15
        this.context.WaitOneUpdateAll(this.depth); 16
    } catch (Exception) {} 17
    // Get fresh metadata 18
    depth.GetMetaData(depthMD); 19
                                   20
    // Get depth map to use while drawing the image to the screen 21
    ushort* pDepth = (ushort*)depthMD.DepthMapPtr.ToPointer(); 22
    // Use information from the pDepth pointer to draw image to the 23
        screen using, for example, a histogram
    } 24

```

Listing 4.1: Getting the depth map of a scene in OpenNI

While drawing the depth map is interesting, being able to find users in the scene is even more so. To do this, we must use the `UserGenerator`, and from this we can do things when a new user is lost or found, and we can also find how many users the sensor sees in the scene, get their positions, skeletons, pixels and so on. The example in Listing 4.2 shows how to start the

user generator and perform some action when a new user is discovered and when one is lost. Our game relies heavily on `UserGenerator`, and Section 4.3 will explain how we used it.

```

private Context context;           1
private ScriptNode scriptNode;    2
private UserGenerator userGenerator; 3
                                     4
// Upon initialization, create context and user generator           5
context = Context.CreateFromXmlFile("path_to_folder/SampleConfig.xml",6
    out scriptNode);
userGenerator = new UserGenerator(context);                         7
                                     8
// Bind events to methods                                           9
userGenerator.NewUser += userGenerator_NewUser;                   10
userGenerator.LostUser += userGenerator_LostUser;                  11
                                     12
// Start generating user information                                13
userGenerator.StartGenerating();                                   14
                                     15
void userGenerator_NewUser(object sender, NewUserEventArgs e) {    16
    // Do something when a new user is found, their ID can be found 17
    through e.ID
}                                                                     18
                                     19
void userGenerator_LostUser(object sender, UserLostEventArgs e) {  20
    // Do something when a user is lost, their ID can be found through e21
    .ID
}                                                                     22

```

Listing 4.2: Using the UserGenerator in OpenNI

This is only a little bit of what the OpenNI framework can do, but they are the two generators that are most relevant for us. The `UserGenerator` can be used to work with a user's skeleton as well, by registering a calibration pose and setting the skeleton capability. More information about this and working with hand points and audio generators can also be found in the OpenNI Programmer Guide.

4.2.3 XNA Game Studio

There are a ton of approaches one can take when making a game. You can code everything from scratch and develop your own framework, or you can use existing tools and frameworks for easier and faster development. Since

we were going to make a prototype we decided that it was a wise choice not to spend precious time on inventing the wheel for the Nth time, but rather focus on implementing the game prototype. We decided to use the XNA Game Studio 4.0, which is a programming environment that enables you to use Visual Studio to develop games for Windows (and also for Windows Phone and Xbox 360) (Microsoft, 2012d). With XNA Game Studio you get access to a set of managed libraries from the XNA Framework that is designed for game development and follows the design patterns and idioms from the Microsoft .NET Framework. XNA strives to have everything you need to develop a game in one single system. The programming language that was used developing our game was C#, but any .NET-compliant language of your choice can be used when developing a game with XNA.

Being an intuitive and good framework that simplifies game making is not the only motivation for choosing this framework. One of the authors has previous experience in working with the XNA Game Studio, both from personal projects and projects and assignments given by various courses at our university. The learning curve for using XNA would thus not be very steep, and we could get going with making our prototype quickly.

When creating a new game, XNA will automatically generate a number of classes for you, enabling you to run the program and immediately show a window displaying a blue background. It will give you a `Game` class that inherits from `Microsoft.Xna.Framework.Game` which already overrides some important methods that you will need, like the `LoadContent`, `Update` and `Draw` methods. You now have a good starting point for creating your game, enabling you to display stuff on the screen from the very beginning.

4.2.4 Farseer Physics Engine

With our chosen game concept and style, we needed simple physics to make the game playable and more fun. As we had chosen to work with XNA, there were a few physics libraries that were recommended to use with it. One of them was the Farseer Physics Engine, an open source 2D physics engine made for the Microsoft .NET platform (Physics, 2010). It is based on the Box2D physics engine (Box2D, 2012), an open source C++ engine that has been widely used in 2D games, for instance Angry Birds by Rovio (Rovio,

2012). Farseer uses four central concepts (introduced by Box2D) in order to simulate a moving body that can collide and interact with another body:

World is a collection of bodies, fixtures and constraints that can interact. Manages all the objects, updates them and makes sure the simulation is consistent and stable. A call to the world's `Step` will iterate over all the objects in the world and update them according to the simulation.

Body keeps track of a position in the world and is the backbone for fixtures. The body is affected by forces like impulses and torque from the simulation. A body can be:

Static - does not move under the simulation, behaves like it has infinite mass and does not collide with other static or kinematic bodies.

Kinematic - moves according to it's velocity, does not respond to forces, have infinite mass and does not collide with other static or kinematic bodies.

Dynamic - is fully simulated and normally moves according to forces, has a finite nonzero mass and can collide with all other bodies.

Shape is a 2D geometrical object in space, like a circle or a polygon. Things like area, inertia and the centroid are calculated for the shape.

Fixture acts like glue and attaches the shape to the body and adds material properties like density, friction and restitution.

After you have created a world, you can create a body, give it a 2D form by attaching a shape to it by using a fixture. The centroid of the shape will be fixed to the body's position. The body will then carry the shape around in the world and once the shape collides with another shape in the world, the forces are calculated and applied to the body. Listing 4.3 illustrates how this was done.

```
//Create a World object with zero gravity           1
World world = new World(Vector2.Zero);             2
                                                    3
//We create a body object and make it dynamic (movable) 4
Body myBody = world.CreateBody();                 5
myBody.BodyType = BodyType.Dynamic;              6
                                                    7
//We create a circle shape with a radius of 0.5 meters 8
```

64 DESIGN AND DEVELOPMENT

```
CircleShape circleShape = new CircleShape(0.5f);          9
                                                            10
//We fix the body and shape together using a Fixture object 11
Fixture fixture = myBody.CreateFixture(circleShape);      12
```

Listing 4.3: Creating a world with a single body

In Farseer 3.3.1 (which we used) enables you to make use of a `BodyFactory` in order to make this process easier. It takes care of creating a shape and fixing it to the body for you, as shown in Listing 4.4.

```
//Create a World object with zero gravity                1
World world = new World(Vector2.Zero);                   2
                                                         3
//We create a circle shaped body object and make it dynamic (movable) 4
Body myBody = BodyFactory.CreateCircle(world, radius, density); 5
myBody.BodyType = BodyType.Dynamic;                     6
```

Listing 4.4: Creating a world with a single body in Farseer 3.3.1

In addition to this you can create joints and constraints that can affect the body in various ways. A ton of cool features that will make your simulation interesting are also included, like cutting shapes and converting a texture into a polygon.

4.3 Game prototype

The game prototype was made from scratch by using XNA and utilizing the OpenNI and Farseer libraries. We began with setting up a normal Windows Game Project with XNA Game Studio, linking and exploring the OpenNI library by showing our silhouettes on the screen. This was fairly straight forward to do, as the `UserGenerator`'s `GetUserPixels` method gives a pointer to a pixel map where each pixel is assigned a label that is either 0 (no user) or the ID of a user. These labels were then assigned a color and written to a texture that was displayed on the screen by using XNA's `SpriteBatch`.

After getting the silhouettes shown on the screen, it was time to look at the physics engine. As we had decided that the game was going to be a "collect N items into a container by using your silhouette"-game, we needed some falling and collectible items. We made a `Ball` class that used

`BodyFactory.CreateCircle` for its body and spawned them above the screen. We now had falling objects that could interact with each other, but how could they interact with a silhouette that is in the form of a texture? In order to get Farseer to include something in its simulation, it needs a body from its `BodyFactory`. Farseer has an interesting functionality that enables a texture to be converted into a polygon and then into a body (which we will explain more later in this section), but it is too time consuming to run every frame. Our texture with the user silhouettes needs to be updated each frame in order to smoothly show the users' movements on the screen. We decided that we had to write our own collision detection for objects against the silhouette, and Listing 4.5 shows a short pseudo code of the approach.

```

for each Ball                                     1
  pushVector = (0,0)                               2
  for each nonzero Pixel in LabelBuffer           3
    if the distance to the Pixel from Ball.Center is less than Ball. 4
      Radius
      calculate the Ball's motion vector towards the Pixel           5
      apply this velocity to the Ball in the opposite direction with a6
        dampening factor
      pushVector += Vector pointing from the Pixel to the Ball       7
  add the pushVector to the Ball's position        8

```

Listing 4.5: Pseudo code for our collision detection

This approach led to the balls bouncing off the silhouette, and in addition we added a push vector that made the ball flee from the silhouette if the object ended up trapped inside of it. There was a lot of tweaking of factors, normalizations and clamping involved to get a decent looking approximation to physics, but in the end we were pleased with the behavior and went with it. As an illustration of how their behavior was experienced by the user, the balls behaved like heavy balloons.

The container that the objects were to be collected in should also be very easy to replace, as the graphical style may vary greatly. We discovered a nice set of algorithms in Farseer where you could convert a texture into a polygon, and then to a `Body`, that worked together with the other elements in the physics engine. Figure 4.4 shows a shape in texture form to the left and its polygon shape to the right.

Listing 4.6 shows the code for converting a texture to a `Body`:

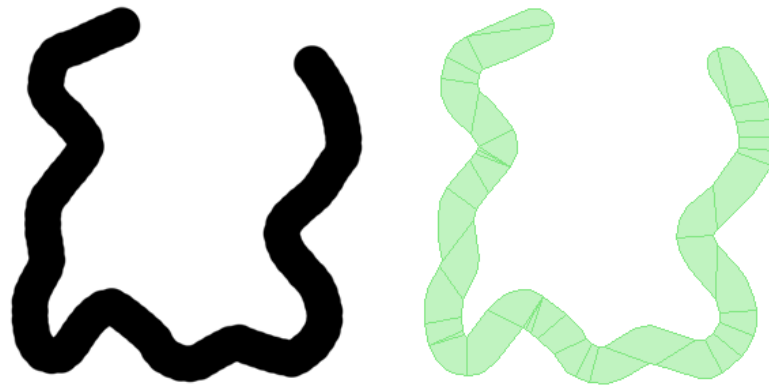


Figure 4.4: Texture to polygon by using Farseer's PolygonTools

```

uint[] data = new uint[polyTex.Width * polyTex.Height];           1
polyTex.GetData(data);                                           2
vertices = PolygonTools.CreatePolygon(data, polyTex.Width);      3
decomposedVertices = BayazitDecomposer.ConvexPartition(vertices); 4
body = BodyFactory.CreateCompoundPolygon(world, decomposedVertices, 5
    density);

```

Listing 4.6: Converting a texture to a body using Farseer

The `BayazitDecomposer` is one of five algorithms used for decomposing concave polygons into convex ones, as the Farseer engine only supports convex bodies. This was an easy way to add obstacles and containers to our liking, and you could attach another texture on top of the polygon (the polygon texture would normally not be visible, as the texture input to the polygon converting process should be as clean as possible in order to get best results) to get a nice and fancy looking obstacle.

With this we had a very basic game where you could use your body to collect balls into a container. The next step was to identify what was in the container, in order to get some sort of gameplay. Farseer supports bodies that act like sensors by setting `Body.IsSensor = true`. Each body has two event handlers: `OnSeparation` and `OnCollision`. We used both to keep track of balls inside the container. The `OnCollision` fires once when another body collides with this body. The same happens for `OnSeparation` except it only fires when a body as a whole leaves the event handling body. We created an `ObjectCounter` class that was a sensor type and registered

each object that triggered the `OnCollision` event. We had some problems figuring out the different types of objects that triggered the event, as you only got access to a `Fixture`, which is something that remains from the Box2D engine and now is merged with `Body` in Farseer 3.3.1, as discussed in Chapter 4.2.4. We solved this by going through every ball that is currently in the scene and checking against the `Fixture` ID that is connected to the ball's body. The gameplay can be summarized in Listing 4.7.

```

if there are less than two objects in the container           1
  spawn new balls at a random position above the screen     2
if there are two or more objects in the container           3
  if the objects are of the same type                       4
    if the number needed to "win" is reached                5
      show the winning overlay (campaign/congratulations)  6
      empty and reset the container                         7
      continue                                              8
    else                                                    9
      display the "try again" overlay                       10
      empty and reset the container                         11

```

Listing 4.7: Pseudo code summary of gameplay

4.3.1 System overview

The game architecture is very simple. A brief overview of the system can also be seen below in Figure 4.5. It consists of a simple monolithic core where the different functions are split into annex classes just for the sake of basic organization. The core and the annex classes are organized as follows:

Game The core class of the system. This class is responsible for basic I/O and event handling, and instantiating and organizing the other classes. The `Game` class contains a reference to exactly one instance of each of the following annex classes:

TextureCollection Maintains a collection of images used to theme the game to fit different customer cases. Other parts of the game access the texture resources through this class. This class is typically changed to adapt the game to a new campaign or new customer.

Logger Responsible for logging selected user activity to a file on disk, so statistics can be generated later.

Sensor Responsible for managing interoperation with OpenNI. Sensor also keeps track of users joining and parting from the game by appearing or disappearing from the sensor's range, and registering users by recognizing their wave movements to the sensor, if enabled in the Game class.

Informator Responsible for watching the state of the game and informing the user whenever the game is "won", typically by drawing a layer of information on top of the game graphics. This class can be changed based on the theme.

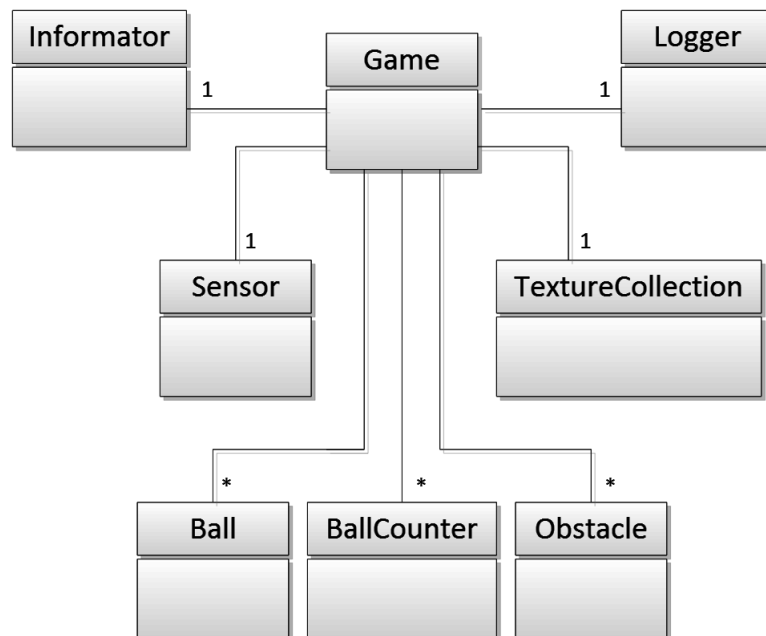


Figure 4.5: Overview of the system

In addition to the core and the annex classes, the Game object manages collections of the following classes of objects:

Ball A ball represents a free moving object in the game that a player can interact with. A ball is linked to a circle shape in the physics engine.

BallCounter Represents the "goal" in which the player is supposed to make the ball end up. BallCounter raises events whenever a ball intersects it's body.

Obstacle Represents a physical obstacle that balls collide with. These can be placed throughout a level according to the theme. This was typically used to create a "container" which the balls are supposed to fall into. A `BallCounter` was then placed inside a concave obstacle. Obstacles were also used to block off the walls of the level, for instance.

4.3.2 Modes

We wanted to make modes to see if people behaved differently if they had to play alone or with others. But in order to choose players that were willing to play from a crowd of people, we needed some kind of registering mechanism. We chose the wave gesture in our prototype, where a `GestureGenerator` scans the field of view to detect wave (or other specified) gestures and generates gesture data that can be used further. Whenever a wave gesture is picked up by the `GestureGenerator`, an event is triggered and uses the position of the hand to look up the user ID in the `LabelBuffer`. This user with this ID is then set as active, and as a result of a small test in the collision detection, the user may now interact with the falling objects. A short description of each mode follows:

Multiplayer

The multiplayer mode allows every user that has a silhouette to interact with the falling objects. Only one container is provided, as the screen tends to be crowded and people will end up standing in the way of the container. No registration is needed, all you have to do is move in front of the sensor so that it detects you, which will display your silhouette on the screen and make you a part of the game. This also leads to a lot of disturbance from passers-by in the background, as they will be picked up by the sensor and affect the game, whether they want to or not. This fact also led to the two-player mode, as we saw that a lot of people in the background easily could disturb and "ruin" the gameplay, in addition to testing how people behaved when they had to play alone or as a pair.



Figure 4.6: Two-player mode with two active and two non-active players

Two-player

Only two users can interact with the falling objects. The users need to do a wave gesture in order to play. When this gesture is successful their silhouette will change color, a small burst of particles will emerge from the users waving hand as an indication of a successful registration and the wave instruction will disappear. All other users are "grayed out" and are unable to interact with the game, as can be seen in Figure 4.6. The game will provide two containers for collecting the falling objects.

4.3.3 Versions

Standard

The game was initially made with a crayon like look to match the rather rough contours of the user's silhouette. The first version of the game had a gameplay where you were supposed to collect three balls of the same color into a container. A screenshot of the standard version can be seen in Figure 4.7. The standard version was used at the locations Pirbadet and NTNU, which we will describe in Chapter 5.1 (along with the other two locations).



Figure 4.7: Standard version of the game

Popcorn

The popcorn version of the game was used at the Location kiosk at Prinsen Kino. Location wanted us to use their graphical style on the popcorn basket (and the rest of the game), which can be seen in Figure 4.8. This design proved to be troublesome when it came to displaying what was inside the container (as we did on the other two locations), so we added a small counter on the container in order to track how many popcorns you already had collected. We discovered that it was necessary to make "winning" easier, and changed the concept from "collect three similar colored objects" to "collect five objects".




Figure 4.8: Popcorn version of the game

Mercur shopping mall

Mercur's version of the game continued to use the concept of collecting five objects that we introduced at Location. They also wanted to give away gift certificates during the final two days of the time period, so we needed to add the functionality for this as well. The screenshot in Figure 4.9 shows the game after it was altered to give out gift certificates, as the text in the screenshot reads: "Collect five diamonds! Get a chance to win a gift certificate!". We distributed a random four digit code each time a user managed to collect five diamonds, and a predefined winning code was distributed after a given probability was met. In the beginning there was about a 1/50 chance of getting a winner code, but this was lowered to a 1/25 chance after observing that the winner code came too seldom based on how many different people that were playing. In order to check if their code was a winning one, the player had to bring his/her code to a nearby store.



Figure 4.9: Mercur version of the game

4.4 Interactive poster prototype

The interactive poster prototype was implemented as a website that uses Zigfu's JavaScript library, and would thus work on both Windows and Mac OS X, with the prerequisite that OpenNI and the ZigJS browser plugin are


```
    } 18  
  }); 19
```

Listing 4.8: Using zig.js in the poster prototype

Listing 4.8 shows how we initialized the Zig object and added an event listener to perform the animation of the dropped tiles. We also check that the position of the user that was found is closer to the sensor than 3000mm, or 3m. The `dropTile` function is in charge of animating the tiles, and it uses jQuery's `animate` function and the jQuery Easing plugin for simplicity's sake. As the code sample shows, we use a namespace called `Poster` and the `dropTile` and `reset` belong to this namespace, but this is not strictly necessary and is simply done to keep the code cleaner.

Empirical Study

This chapter will present the reader with the results from our evaluation phase. The first section will explain the research context for our four field studies, which includes which location we tested the game at, how the game was situated within that location and setting, why we chose to test at this location and which version of the game was presented to the users. We then summarize the results from our observations of the field studies, and then move on to present the results of the questionnaires that were presented to participants and non-participants. We also investigate the log files that were produced by the game.

5.1 Research context

Our installation was displayed at four different public locations. Each of them has varieties when it comes to audience, their motives, atmosphere, placement and setup of the installation, and look and feel of the campaign. In this section we will present these different locations.

5.1.1 Pirbadet waterpark

Pirbadet is a waterpark located at the pier in Trondheim. It offers a variety of attractions like a 50 metre sports pool, jacuzzis, therapeutic pools, a wave pool, slides and saunas. There are also a fitness center and a physiotherapy institute in the same building. All these facilities share a common entrance, and our installation was placed near the main entrance. The pictures in Figure 5.1 shows the environment where our installation was displayed. The upper two shows the main entrance, the lower right shows the entrance to the waterpark and the lower left shows the entrance to the fitness center.

Pirbadet has all kind of visitors ranging from small children to youth, families and grown ups. Their motive is often very clear, on their way in or out from either the fitness center or the waterpark. People also spend some time in the main entrance area waiting for other family members or friends to leave the changing rooms. A 60 inch display was placed on a small cupboard which contained a PC running the game. The display was angled towards the main entrance and was placed between the entrances to the waterpark and the fitness center, and we tested it over three days during our eight days at Pirbadet, from the 17th to the 24th of April.

The game presented at Pirbadet was our standard version of the game. We used Pirbadet mostly as a test location to see how the game behaved in a proper setting and to eliminate and fix bugs that we had not seen before while testing the game in private. The game did not have any advertising elements at this location, and the goal was to collect three similar colored balls into a basket.



Figure 5.1: Environment at Pirbadet

5.1.2 Stripa, NTNU campus

Stripa is a long hallway located below the two central buildings at the NTNU Gløshaugen campus. This hallway connects several large auditoriums, a kiosk, administrative offices, a coffee shop and a huge canteen. A lot of students travel through this hallway every day, especially when lectures are held, as they walk between facilities. There are also a lot of students that enter this hallway in order to get lunch, either at the canteen or from the only kiosk at the Gløshaugen campus. A 60 inch display on a table was placed to the left of the southern entrance, fairly close to the coffee shop and canteen. The kiosk laid further down the hallway to the left of the installation. The placement of the installation was not optimal, as students that were coming from the southern entrance and were headed for the coffee shop or canteen did not easily spot the installation. The reason for this was that the “sweet spot” in front of the entrance was taken by another stand. There was also a lot of pillars in the area, so it was necessary to angle the screen and sensor towards the entrance in order to remove blocking pillars from the sensor’s field of view.

The pictures in Figure 5.2 shows (from the top left) the entrance seen from the "sweet spot", the "sweet spot" and the stairway to the bookstore seen from the entrance, troublemaking pillars and the hallway leading to the kiosk and auditoriums and the hallway leading to the canteen and coffee shop. The students traveling through Stripa are used to stands and people trying to give you fliers or recruit you. Huge screens, however, are not that common, and our screen caught some attention by itself. We stood at Stripa for three days in the time period from 30th of April to the 3rd of May.

We contacted SiT (Studentsamskipnaden i Trondheim) about advertising their kiosk’s food products at Stripa, and they wanted us to do a campaign for their “monthly lunch deal” (baguette sandwich and juice). The look of the game was not changed from our standard version, except for when the player managed to collect three similar colored balls into a basket - then the SiT campaign would appear on the screen.



Figure 5.2: Environment at Stripa

5.1.3 Location, Prinsen kino

Location is a kiosk located in the foyer of the cinema Prinsen in Trondheim. This kiosk is the first thing you see when you enter the cinema from both main entrances, one of them is shown in Figure 5.3 at the top left. There are chairs, tables and sofas surrounding the kiosk, as well as a lot of TVs displaying both information, movie trailers and advertisements for various goods from the kiosk, as shown in the two bottom pictures of Figure 5.3. People come and go in chunks, as they often enter the kiosk to buy candy only minutes before the movie begins. We got access to use a permanent 40 inch display that is normally used for advertisements. This display is placed in the only entrance to the kiosk, which is narrow and downhill, making the entrance appear like a funnel. This "forces" every visitor to face the display when walking into the kiosk. The kiosk also has 16 other advertisement displays mounted on strategic places inside, in the top right picture in Figure 5.3 you may see half of the kiosk. We tested the game at Location for four days, from the 7th to the

10th of May.

Location wanted the game to have a look that was based on their standard graphical profile for advertising. We modified the game's graphical look to match their graphical profile by adding their preferred background and logo, and making the collectible objects (previously balls) to popcorn that should be collected into a popcorn container. Location wanted to have a campaign for "gourmet" lollipops. The game was also modified from collecting three similar colored objects to collecting five objects, after feedback from users during previous testing and from representatives for Location. This was meant to make it easier to win, and thus make the advertisement be displayed more frequently.




Figure 5.3: Environment at Location

5.1.4 Mercur shopping mall

Mercur is a shopping mall located in the southern end of the Nordre street in the city center of Trondheim. This mall strives for a youthful image, with mostly clothing stores inside. Mercur bought and installed a projector for our game, projecting on a white wall near the main entrance. The projected image got washed out during the day, as the roof of the mall is made from glass. There was no place to hide the PC and place the sensor, so the PC was placed on the floor and the sensor was mounted on the wall. The projected image was not easily seen when entering the mall, as it was in a corner and people needed to turn their heads when entering to be able to see it. However, it was easier to see when people were leaving the mall through the main entrance. The pictures in Figure 5.4 shows the environment near our installation at Mercur. The top left picture is taken from the main entrance right before you would encounter our installation (the projection was behind the wall on the left side), the top right shows the main entrance seen from the escalator, the bottom left shows the projector and the play area when coming from the second entrance, and the last picture shows the benches that we sat on while observing (and a restaurant/cafe behind them). We tested at Mercur for five days, from the 8th to the 12th of May.

Mercur wanted the game to match their mall's graphical profile, but did not have a campaign that they wished to promote. The game was modified accordingly with Mercur's logo and mascot, and the collectible items were in the form of their diamond trademark. We continued using "collect five objects" rather than "collect three objects of similar color" at this location, as we did at Location. At the end of the test period, Mercur decided that they wanted to give an actual prize to some of the winners. This was done by giving a four digit code to each person that managed to collect five diamonds, where only a few codes would yield a prize. The participants went to a nearby store to check their code and received a gift certificate worth 100 NOK if their code was a winner code.



Figure 5.4: Environment at Mercur

5.2 Observation

5.2.1 Pirbadet

Pirbadet was our first location for testing and observation where we had access to more than a couple of users at the time. For this reason, the first couple of days were spent observing whether or not things worked as expected, and if there were any major problems. If so, these should be fixed before we progressed with the testing. Although there is a relatively constant stream of people here, it is definitely less crowded than a typical mall. People also tend to spend a short time passing through the reception, so it's often relatively empty. This might have led to people being more willing to play enthusiastically, because they were not being observed by a lot of people. The installation in use can be seen in Figure 5.5.



Figure 5.5: The installation in use at Pirbadet

Technical issues

We had initially placed the container in the center of the screen in the multi-player mode, so that two or more users could be close to it and collaborate on collecting enough balls. However, we quickly saw that many people chose to stand in the center of the screen thus covering the container with their silhouette, so this was less than ideal. This caused us to decide to move it to one side after the first day of testing, so that people could stand next to each other in the center of the screen without blocking the container or standing in each other's way.

Another thing we noticed was that people seemed to want there to be walls around the screen so that balls would not disappear out of frame, because they kept trying to bounce them off the edges of the screen. We also received comments on the fact that the balls moved too quickly and that it was too difficult to gather the balls into the container, so before the last day of testing at Pirbadet, we widened the container and slowed down the balls. On

the last day, one of the employees who had played a lot previously exclaimed “Oh, it’s easier now!” and seemed satisfied.

User behavior

Quite a lot of people did not notice the installation at all and simply walked right past it on their way into the waterpark or the fitness center. This could be caused by the fact that they were on their way somewhere and focused on the reception area of the waterpark/fitness center, not the surroundings, and that the screen – while facing the entrance – was in an otherwise empty area where not a lot usually happened. Some did notice it, seemed curious, but did not want to try. This was especially true for adults, some of them commented that “[they] won’t be able to do this”. Others tried moving the balls for a few seconds, seemed to just want to see what it was all about, and then moved on quickly without really trying to gather anything in the container.

The most eager groups were definitely children of elementary and middle school ages, as well as the younger employees (life guards, receptionists and so on, around 20-25 years old) at Pirbadet. The employees were in a familiar environment and were surrounded by people they knew, and they seemed fairly confident about playing it enthusiastically in the middle of the hallway. For the first couple of days, some of them would play a little bit every time they passed the screen, and they often played alone. Other than these two groups of people, some adults would play for a little while, but they seemed to be more self-conscious about doing so in the public space.

Elementary and middle school-aged children who noticed the screen seemed to be drawn towards it, and many would run over to play. Sometimes parents would drag them away from the screen and towards the waterpark, but other times the parents would wait and let them play for a little while. In a few cases, children as young as 3-4 were encouraged by their parents to try the game, and this often happened on their way out of the waterpark/fitness center. While they did not understand the objective of the game, they seemed to have fun and sometimes the parents would explain that they should try collecting three balls. One young boy commented that he wished they had a game like this at home. Most of the children who played were very active, and they would run around, jump and yell, without seeming to be affected by the fact that the game was placed in a public space.

When groups of children played, there would often be some initial light arguing before they tried to organize themselves and collaborate. Some immediately understood the goal of the game, while others simply played with the silhouettes and the balls. Many of the children seemed satisfied by collecting three balls, regardless of color, or cheered after simply collecting one. At one point while a group of younger children were playing, two men in their early twenties walked past, commented to each other that they kind of wanted to play, but did not want to disturb the kids who were already playing.

Although some did comment that controlling the balls was difficult, we also saw that those who played quickly learned some tricks for catching balls more easily, by for example placing their arm in an angle next to the container so that the balls that hit the arm would just slide down into the container, or by gathering balls in their arms and then dropping these into the container. Some also made up their own mini-games and tried catching as many balls as possible between their arms over their head, or doing typical football tricks with the balls, as shown in Figure 5.6. Two of the waterpark employees tried kicking the ball to each other as if it was a football.



Figure 5.6: Screenshot: Tricking with a ball

Aside from some comments about the level of difficulty, the most problematic thing seemed to be that people walking in the background would sometimes ruin things for the people playing, since they would also automatically be a part of the game. For example, sometimes someone would pass by in the background and disturb a ball that was on its way into the container.

When two or more people were playing, they would sometimes also disturb each other, but they often quickly figured out how to collaborate and avoid this.

Two-player mode

We tested the two-player mode for a couple hours one of the days, and quickly saw that people struggled with getting started. Some younger children stood too close to the sensor and tried waving without being able to register. One teenage boy was able to be recognized, but quickly stopped playing. Among a group of kids who wanted to play, a couple were able to do the wave gesture and cheered when they were recognized and treated it like an achievement. They would also call their friends over and try to have them join, but since this was two-player mode, only two players could be recognized and this seemed to confuse them. Using gesture recognition for participation in the multiplayer mode is entirely possible – just not something we had implemented.

5.2.2 Stripa, NTNU

As mentioned in Section 5.1, the game was tested at Stripa with an advertisement for a baguette and juice deal. The testing lasted for three days, with about 5 hours of testing per day. Most of the testing was done using multiplayer mode, but the two-player mode was also tested for a while in the most busy hours around lunch time. As one can expect on a university campus, most of the people who passed by the installation were in their early or mid-twenties, but there were also some university employees and older students. Figure 5.7 shows four students that enjoyed playing at Stripa on the 2nd of May.

User behavior

Perhaps due to the placement of the screen, most people who passed it did notice it and went from simply being passers-by in the audience funnel to viewing and reacting (Müller et al., 2010a). Unless people were heading in the opposite direction after entering, not noticing the screen would be fairly difficult because of its size and placement. However, a fairly large portion



Figure 5.7: The installation in use at Stripa

of those who looked at the screen did not try to participate in any way, as expected. Some would walk by while doing a superfluous movement, such as waving an arm, without stopping, while others would exclaim “Look, we’re on TV!” or “Is that us? No? Yes, it’s us!”. A couple of people also stood in front of the screen while talking on the phone, and did some slight movements and reached for the balls without really playing. They seemed to treat it as a way of passing time during the phone call.

Although the goal of the game is presented in the center of the screen, several people chose to simply have fun in front of the screen. A small group of people danced and watched their silhouettes, while two other people tried playing volleyball with the balls in the game. Some tried grabbing a single ball with their hands, while others played around with capturing as many as possible between their arms, like at Pirbadet – either as one person or while holding on to other people. A couple of times, people would win while playing around like this, and they simply ignored the advertisement and kept

on playing.

The large majority of those who stopped and started playing were in groups of two or more people, and most of them were boys. A couple of groups of girls walked by and commented that it looked cool or fun, but did not stop to play. Mixed groups of both girls and boys stopped several times, though, more than once on the initiative of a girl in the group. The ones who played alone – only about six or seven people – were also boys. We also noticed that while a group were playing, other people who passed by would either stop and watch them or slow down while checking out what was going on.

When groups of people stopped, most of them seemed to be quick to decide that they wanted to collaborate. One person exclaimed “We’re supposed to collaborate, this thing isn’t a contest!” when others in the group disturbed him in the process of catching a ball. Among those who played for a little while, they often found their own techniques for quickly catching enough balls by organizing the “team”, and sometimes the collaboration turned into a friendly mini-contest. Although most of the players did not move around as much as the younger children at Pirbadet, they were fairly active and seemed to be having fun – some also used objects they were holding to control the game, like in Figure 5.8 where a girl is using her backpack to hurl the balls around. An older man who tried playing said that he really did not have time, but that it looked fun so he had to try playing.

One of the problems people encountered was that they stood too close to the sensor, at which point they will either not appear on the screen at all or cover almost all of it with their silhouette. Some understood this and moved back, while others gave up. Giving up quickly was quite common, which we expected because people are usually on their way somewhere when they walk through Stripa. Some also seemed to think that controlling the balls and hitting the container was quite difficult (especially with three similar balls, not just three in general), and expressed a bit of frustration at this. Another cause of frustration was people walking by in the background, as they would disturb the players and quite often “ruin” the trajectory of a ball heading towards the container. Several players expressed annoyance at this, but when asked about whether they thought some form of registration that would alleviate this problem (waving, pose) would be better, they seemed skeptical.

Surprisingly, quite a few of those who tried playing did not see that there was a container that was meant for gathering the balls. A few groups of people tried catching three balls in their arms, and one person even asked if he was not supposed to get something for winning, thinking he had already won after holding three balls above his head. One group played for a couple of minutes before noticing the container. In cases where one person in the group noticed it and others did not, they were quick to share information, telling each other about the container and trying to find out how to hit it. We also noticed that people who had understood/read the rules before others would share their knowledge, for example by explaining that they needed to collect three of the similar color or by asking people to move if they were standing in front of the container without realizing it.



Figure 5.8: Screenshot: A player using her backpack while playing

Two-player mode

After we switched to the two-player mode where a wave gesture was required to play, we quickly noticed that people had problems with getting started. Several groups of people tried waving without getting recognized and when nothing happened, they were quick to leave. Some did try playing even though they could see that the balls were not bouncing off them, just falling through their silhouettes, but these also gave up quickly and seemed to not understand what they were doing wrong.

One group tried playing without waving, because it took them about 30-40 seconds to read that they were supposed to wave to get started. They eventually understand that they had to wave to make it work, but one of the players had played the game previously and knew that it should work, so he encouraged the others to continue trying. Others, who were enthusiastic and exclaimed that the game seemed cool but couldn't figure out how to wave correctly, ended up leaving after trying for a little bit.

Although getting started was difficult because the wave gesture generator did not easily understand the variety of ways in which people waved, once they actually managed to get registered, the game seemed to be easier to control. As mentioned, one of the main sources of annoyance for those who played in multiplayer mode was the people in the background, and this is no longer a problem when people have to wave to become a part of the game. Because people in the background could not disturb the balls, it seemed like the players found it easier to gather enough balls to be shown the advertisement.

Responses to the advertisement

Although winning did produce joy among the participants, many of them seemed disappointed when they saw that the "prize" was an advertisement for an already existing offer – especially those who had tried for a longer period of time to gather three similar balls. Some commented that they wanted a special offer, while others wanted to win something. One person thought he had won a baguette after quickly looking at the advertisement without really understanding what it said.

Quite a few of those who collected enough balls did not seem to care about the advertisement at all. They gave it a brief look, and seemed more satisfied with actually having collected enough balls than what the consequence was. Another person did not understand that he had won, and yelled "No, they're disappearing from my basket!" when he was shown the advertisement and the balls bounced out of the container.

One person commented that this was the first time he had been happy to be stopped by someone at Stripa (it is a common place for businesses and organizations to have stands where they try to recruit or inform people), so although he did not find the advertisement very exciting, he was positive about it. Many other people also seemed positive until they were presented

with the advertisement, at which point they expressed some disappointment at the offer they were given.

5.2.3 Location, Prinsen kino

Our field study at Prinsen kino lasted for four days, for four hours every afternoon (which is when the cinema is busiest). The amount of observations here is smaller than at the other locations, for the simple reason that fewer people noticed that something was happening on the screen, which resulted in fewer players and events to observe.

As mentioned in Section 5.1, there were 16 other screens in the area surrounding the screen with the game that all had the same graphic profile, so the game did not stand out very much, even though everyone who was entering the store had to pass by it. We could quite obviously see that the large majority of people did not even notice that there was a game to be played, even if they stopped in the middle of the entrance and stood right next to the screen. Some did look at the screen and continued walking by. The screen we used had previously been a regular digital signage screen, so people could have learnt to ignore it, and a group of boys commented that they did not notice that the screen was different from the other 16. The installation in use can be seen in Figure 5.9, where three teenage girls are playing.

At one point, a young man stood in front of the screen for a while, observing what is going on. Coincidentally, five popcorns were gathered into the container while he was watching, and he saw that the advertisement was displayed. This caused him to start moving his head to touch the popcorn and eventually his arms, and when two of his friends came over, they joined him in playing for a little while. We also observed a couple of other times that when people saw that others were playing, they seemed to become more curious and try it for themselves, as if they were being made aware that it was actually a game by watching others play. Two of the employees at Location, who had seen other people play, seemed to have fun playing during a quiet period in the kiosk. They used their heads actively, as well as jumping around and waving. Other people who played were fairly stationary, as the screen was placed in the main entrance and people who played could be in the way of those walking in and out.



Figure 5.9: The installation in use at Location

A surprising phenomenon, which we had not seen at any other locations, was that three people, at separate occasions, thought it was a touch screen and tried collecting the popcorn by touching it on the screen. Both two young boys and an adult woman did this. One of the boys eventually noticed that if he moved back, the popcorn bounced off his silhouette, so he tried playing for a little while. The woman eventually tried walking back and forth in front of the screen, but did not understand that the silhouette was her until she asked and we answered. She then tried to play for a little bit, managed to collect five popcorns, and was disappointed that she only “won” an offer that everyone else also gets. She said that she felt a bit tricked, but would like it better if it was a special offer just for her.

A group of four adults in their fifties found the game quite intriguing, and the two women encouraged the two men to play, but did not want to play themselves. One of the men commented that it was difficult at first, but that he quickly got better at playing. This group hung out near the game for

quite a while and seemed interested in the concept. The rest of the area was fairly empty while they were playing, though, and we noticed that others who stopped to play were quite quick to move on once passers-by walked by them in the entrance.

Interactive poster

The interactive poster was tested for about two hours, but no one gave it any attention. These two hours were relatively quiet, though, and we decided to test it for a bit longer at a different location later.

5.2.4 Mercur

During the first three days at Mercur, the game was tested without any particular prize or advertisement, but the fourth and fifth day a gift certificate was added as motivation, which greatly increased the amount of time people spent playing and the amount of people who wanted to play. Figure 5.10 shows the installation in use at Mercur on Friday the 11th of May, after we had introduced the gift certificate.

User behavior without gift certificate

During the first three days at Mercur, the amount of people who played was fairly low. This might have been caused by the placement or the fact that the game was projected onto a wall in bright daylight, so it did not draw much attention. It could also simply be caused by the fact that many people in a mall have a goal in mind, and they simply did not want to play. There are also a limited amount of people in the mall early in the day, so most of the people who played did so later in the day. We also noticed a clear trend where most of the people who stopped were groups of young people around 8–14, and unlike the other locations, many of them were girls. However, given that Mercur is a mall, groups of teenagers who are simply hanging out is quite common, so this makes sense. Some young children also played, but it varied whether parents let them try (one dad even taught his children what to do) or simply saw it as a disturbance and made them leave.




Figure 5.10: The installation in use at Mercur

Not everyone who chose to stop and play actually played in the way we intended. Two-three groups of boys of around 13-14 stopped for a little while and used various body parts to move the balls around, but they did not seem to have a goal in mind while doing this and rarely tried depositing the diamonds in the container. A couple of other groups played around and collected diamonds while standing with their arms in a Y shape, like we have seen on the other locations, illustrated in Figure 5.11, and some figured they had completed the game after collecting five diamonds this way. Two young girls stopped in front of the screen, exclaimed "Ooh, this looks fun!", pushed some diamonds around with their heads, but quickly gave up. We also saw that some men in their early twenties kicked the diamonds to each other and had fun in this way.

During these first three days, the amount of time most people actually played once they had stopped was pretty short, often well under a minute. There were some exceptions, but most of them were either groups of young



Figure 5.11: Screenshot: Two people collecting balls in their arms

teenagers or people who were sitting on the benches surrounding the game, having their lunch break. Two of these people played for a long time and seemed to enjoy it a lot – they were also very interested in both how the game worked and winning. They also challenged each other and explained the rules to those who did not understand (one woman had not noticed the container).

As with the other locations, the largest source of annoyance to people seemed to be passers-by in the background who would throw diamonds off their trajectory. Because the mall was quite crowded, this was a bigger problem here than other places. Another issue was people stopping in the background to talk, and some of these would block the basket. Quite a few of the teenagers seemed like they did not want to ask people to move, but they also became annoyed. We also noticed the usual light quarreling between people who were playing together if one person ended up ruining for others, but this was mostly friendly.

We also noticed that some seemed hesitant to play in the setting, perhaps because it was quite public. One girl who had actually played was embarrassed about being watched when she was asked about filling out a questionnaire afterwards. An adult man with a stroller stopped in front of the screen for a while, watched what was going on and posed slightly, but did not actually play. Another young girl looked like she wanted to play, but then looked over at us and did not want to. We moved into the nearby coffee shop to see whether this led to more people playing, but this did not seem to have

any effect (we were sitting on a bench nearby originally).

Compared to at Pirkadet and Stripa, where people who won were shown an advertisement and some expressed disappointment, no one seemed to express significant disappointment at not winning anything. Some friends of one of the mall employees commented that a prize would be cool, but out of the users who actually played and tried winning, actually winning seemed to give enough satisfaction. Two young boys who played for quite a while and were very active, were thrilled when they finally collected five diamonds and high-fived each other before moving on. However, the game remained largely unused for long periods of time, and we found it interesting to test the impact of adding a motivating factor, such as a potential prize.

Two-player mode

We tested two-player mode for about two hours one of the days, to see whether the tendency from the previous locations continued. One girl tried playing, but did not read that she was supposed to wave to the sensor, so collecting the diamonds was impossible and she quickly gave up. Another person, this time a boy, understood that he was supposed to wave and tried it, but it did not work. A friend of him tried to help him, but he ended up asking one of us for help. After showing him how to get registered, he played for quite a long time and expressed that it was fun. While he played, an adult woman watched in the background and seemed to find it interesting. One person, who had been sitting in a nearby coffee shop and watched other people play, came running out and tried to wave to the game. However, he waved with both hands and moved too much, so he never got recognized and left before trying to play.

User behavior with gift certificate

After adding the gift certificate, we immediately saw an increase in the amount of people who chose to stop and play, but it still required that people actually read what was written in the background of the game to know that there was something to win. However, the increased number of people who played also meant that more people noticed the game and stopped to watch those who

were already playing, which made them more likely to notice the possibility of a prize and also to become aware of how to actually play.

We did see that people seemed skeptical about playing while others were already playing. Some young adults appeared to find it interesting and expressed that they wanted to play, but did not want to disturb the young girls who were already playing. At one point, a group of young girls who had played earlier in the day were sitting on a bench watching a group of boys play, and the moment the boys left, the girls ran over to play. Some seemed to be aware of keeping others from playing, as one girl, who had played with her friends for a while, said to the others “Come on, we have to let others play as well” and dragged her friends away. While some groups would send one person to check winning codes, others would go together, and this opened up the game for other people who had been watching.

In addition to having more people stop and play, perhaps the thing that changed the most when we added the potential prize was that people played much longer. Instead of being “finished” after collecting enough diamonds once, they would do this repeatedly for a long time, and they would often organize themselves so that they took turns running over to the nearby store to check whether the latest code was a winning code or not, while the rest continued playing to get a new code. Several times, one of the girls in the group would say “Okay, this is the last time!” before playing, but after checking the code and not winning, they would return again. One person who played alone for a long time wrote down codes he collected on his hand, and then went over to check afterwards if he had won, as seen in Figure 5.12. After playing for a long time and checking more than 25 codes, he eventually won and was thrilled – and commented that he wanted to play even more.

Even after adding the gift certificate as a prize, groups of teenagers were still by far the most active and engaged user group. One of the days, three groups of girls returned several times to play, and played for a long time every time. Two of them eventually won gift certificates, and were very pleased. In one of the groups, where they shared the gift certificate after winning, stopped playing afterwards, while another – where the girl who ran over to check the winning code kept the gift certificate for herself – continued playing because the others wanted to win as well. Not everyone who managed to collect enough diamonds to display a code bothered to actually check whether it was



Figure 5.12: Eager participant at Mercur

a winning code, though. Some simply ignored it and continued playing, while others left. Some of the players were too young to understand, but among those old enough to understand, most of those who did not bother with the code were male.

A consequence of playing for a longer time was that they eventually got better at controlling the diamonds and figured out their own techniques for winning quicker and easier. We also noticed less focus on having fun and making up their own mini-games, however, because the focus was on collecting five diamonds to produce a new code, and nothing else. While one teenage boy was playing alone and simply catching diamonds in his arms, one of the young girls who had played a lot – and who did not know the boy previously – went over to him and explained to him what he had to do in order to win. We also saw that these girls would explain to others who had played where the nearby store was and so on, and took on a “teacher role”.

The last interesting thing we noticed was that although quite a few

people, especially girls, expressed some embarrassment about playing, that did not necessarily stop them from doing so. One girl who had watched friends playing without participating eventually wanted to play, too, but asked a friend of her to play with her so she would not be standing there alone. One boy, who was playing with some friends, commented that "We look like morons", but continued playing anyway. A couple of times, two or three girls discussed among themselves whether or not they should try, hesitating because it was embarrassing, but eventually ended up trying, and one of them even won a gift certificate. An older woman, about 70 years old, spent quite a bit of time watching others play, and once one of the mall employees took it upon himself to show her how to play, she carefully played for a while as well.

5.2.5 Interactive poster

We only tested the interactive poster two places, quite simply because those were the only two places where had something to advertise. As previously mentioned, the tests at Location were not very successful. The amount of people at the cinema varies greatly with when movies start, and the poster was tested while it was fairly quiet, because we wanted to test the game more thoroughly while it was crowded. Very few people looked at the screen while passing it, and thus did not notice that anything was happening.

To get some more information about how people related to the interactive poster, we tested it again at school later, and compared it to simply displaying the poster that was underneath the puzzle pieces. The 60 inch screen was placed in an area where quite a lot of people walk by (the "sweet spot" mentioned in Section 5.1), and we noticed that a lot of people looked at it and became curious when something fell down as they were passing. Some tried moving and waving in front of the sensor to see whether something would happen, and they would often try to interact more than what was intended. Others watched for a while and eventually seemed to understand that a piece fell down when other people walked past it.

After testing the interactive poster for a while, we compared it to simply displaying the poster on the screen without puzzle pieces, to check whether people noticed the interactive poster simply because it was on a large screen. Although some did still look over at the poster on the screen, not very many

did this compared to how many who looked and reacted to the interactive poster.

5.3 Questionnaires

This section will present the relevant results from our questionnaire. We start off by giving an overview of the demographics of our respondents, as this is very relevant for the scope of the study and for how we interpret the data. We then give an overview of the questionnaire items related to the gameplay and using the application in a public space and show the answer distribution. Finally, the statements related to the marketing aspect are presented with their answer distribution. It is important to note that the questionnaire answers are only related to the game prototype, not for the interactive poster.

We often refer to *participants* and *non-participants* in this section. The participants are those who chose to participate in the game and were given the questionnaire with statements related to playing the game, while the non-participants are the respondents who became aware of the game, but chose not to play, and were given the shorter questionnaire. This was done to get some information from those who did not want to play, but because we still needed them to be aware of the game and the concept, this questionnaire was not given to those who simply passed by without looking at the screen at all.

5.3.1 Demographics

Overall, we had a total of 105 respondents for all three questionnaires. This can be further separated into 28 non-participants, 29 participants in the game with an advertising campaign and 48 participants in the game without a specific advertising campaign. The first questions in our questionnaire had to do with age, gender and gaming habits in general. This gives an insight into the age and gender distribution of our respondents, and allows us to draw lines between their gaming habits and their answers to the other statements.

Figure 5.13 shows that the majority of our respondents were male, which is a direct consequence of the fact that the majority of those who actually




Figure 5.13: Gender distribution of our questionnaires

played were male. These numbers are similar within the non-participant and participant groups (about one third of the non-participants were female), so the numbers presented here is the overall distribution.




Figure 5.14: Age distribution of our questionnaires

We can also see in Figure 5.14 that the largest age group are those between 18 and 24. This makes sense given that we tested the game at the university, where there are many people in their early twenties. Only 14% of the respondents were of the ages 31 or above, which is also a consequence of the fact that mostly younger people chose to play or watch other people play. However, if we look at Figure 5.15 we can see that among

non-participants, there are no children between 7 and 12 and a greater percentage of respondents over the age of 25, particularly among the 31–50 and the 50+ age ranges.




Figure 5.15: Age distribution by participants and non-participants

In order to find out whether or not there is a correlation between a respondent's gaming habits in general and how they related to our game in a public space, we also asked respondents how many hours per week they spent playing games, including games on their mobile phones, tablets, consoles and PCs.

As we can see in Figure 5.16, the majority answered that they played 1–4 hours per week. Although it is not shown in this chart, we could also see that there was a tendency where those above 25 played less than those below. We did see some minor differences in the hours spent playing games between participants and non-participants, but the sample size is too small for us to draw any conclusions from this. Overall, the extremes (0 hours per week and more than 30 hours per week) were represented to a greater degree among those who did not choose to play. For example, among our 28 non-participants two people said that they play for more than 30 hours per week, while that number was one person among the 77 participants. The percentage of respondents who played 0 hours per week was 21% among the non-participants and 9% among the participants, but this could be caused by the fact that, as Figure 5.15 shows, a larger percentage of the non-participants were above the age of 25.

As we have shown here, our demographic is skewed towards a male majority and people between the ages of 18 and 24, but other age groups are also represented. Our youngest respondents were 7 (their parents filled out




Figure 5.16: Hours of gaming per week among respondents

the questionnaire for them), while the oldest was 70. 75% of the respondents say that they play games between 1 and 9 hours per week.

5.3.2 Playing the game

The following statements and questions were all concerned with how respondents felt about playing the game. Some are related to the gameplay it self while others are related to the fact that the game is a public installation. Although we had 105 respondents overall, not everyone chose to answer every question, so the numbers here might reflect that. There is also a different number of participants and non-participants, but they are shown on a bar chart based on percentages, to show the overall differences between the two groups. Each section of the bar chart also shows how many actual respondents it corresponds to.

I1: I felt comfortable with becoming a part of the game without giving explicit consent.

The first question, which appeared on both the questionnaires for the participants and the non-participants, was concerned with how comfortable the respondents were with being “thrown into” a game while passing by it, without

explicitly giving their consent to their silhouette being projected on the screen and suddenly being a part of the game. As we can see from Figure 5.17, 47 of the 75 participants (63%) who answered said that they strongly agree with feeling comfortable. For non-participants, on the other hand, only 20% (5 out of 25) strongly agreed. We also see that 16% of the non-participants strongly disagreed, while only 4% of the participants did so.




Figure 5.17: Answer distribution I1: I felt comfortable with becoming a part of the game without giving explicit consent.

I2: I was comfortable with playing the game in a public space.

Again, this item was present on both the questionnaires for participants and for non-participants and the answers can be seen in Figure 5.18. The statement




Figure 5.18: Answer distribution I2: I was comfortable with playing the game in a public space.

was phrased slightly differently on the non-participant questionnaire (“I am comfortable with playing such a game in a public space.”) because of the fact that the non-participants hadn’t actually played, but what we want to know is

the same thing – how comfortable are they or would they be with playing a motion-controlled game in a public space? Among the participants, 61 out of the 77 respondents – 79% – who filled out the questionnaire agreed or strongly agreed that they felt comfortable with playing the game in a public space. This share was quite a bit lower for the non-participants with 46%, while 31% of the non-participants felt neutral about it compared to 16% of the participants.

I3: I would have been more comfortable with playing with a controller (mobile phone, Xbox controller).

This question was only given to the non-participants, in order to keep the participant questionnaire from being too long. We wanted to find out whether they would be more comfortable if they did not have to use their body as a controller since this was in a public space and some might be worried about “looking silly”, but the answers shown in Figure 5.19 give no clear indications that this would be preferred. Nine respondents say that they disagree or strongly disagree, while nine others agree or strongly agree. Seven, as we can see, are neutral.




Figure 5.19: Answer distribution I3: I would have been more comfortable with playing with a controller (mobile phone, Xbox controller)

I4: I thought getting started with the game was easy.

With this question we wanted to find out whether the participants found it easy to understand how to get started with playing the game. Using one’s own body as a control mechanism might not be familiar to everyone, so it is interesting to see whether they felt that they understood this. Only those who actually

participated in the game were asked this question, and all of them played the multiplayer mode where no calibration to start was required. As shown in Figure 5.20, 87% of those asked agreed or strongly agreed that getting started was easy (strongly agreed being the majority answer with 68% of the total votes), while only one person out of 76 disagreed.




Figure 5.20: Answer distribution 14: I thought getting started with the game was easy.

15: I thought it was easy to control the game.




Figure 5.21: Answer distribution 15: I thought it was easy to control the game.

Unlike the previous question, this was more concerned with controlling the game once the player had gotten started. Again we see in Figure 5.21 that the majority agreed or strongly agreed that it was easy. However, 8% of the players disagreed that controlling the game was easy, which is a larger share than those who disagreed that it was easy to get started. 41% of the players strongly agreed that getting started was easy, while 34% simply agreed.

16: I think most people will be able to play the game.

Figure 5.22 shows that 61% of those asked strongly agreed that most people would be able to play the game, while 25% agree, which means that a 86% of

the respondents think most people will be able to play the game. 5% of those asked disagreed or strongly disagreed (three and one respondent respectively), but two of them commented jokingly while filling out the questionnaire that they did not think babies or paralyzed people would be able to play and answered accordingly.




Figure 5.22: Answer distribution 16: I think most people will be able to play the game.

17: I became unaware of my surroundings while playing the game.

We wanted to know how unaware players became of their surroundings in order to find out a bit about how immersed they became in the game and Figure 5.23 shows the answer distribution for this statement. Because the game installation was placed in a public space, there is a lot happening around players all the time. 65% of those who played – 50 out of 77 respondents – strongly agreed or agreed that they became unaware of their surroundings. 13% strongly disagreed or disagreed and were not unaware of their surroundings while playing.




Figure 5.23: Answer distribution 17: I became unaware of my surroundings while playing the game.

I8: Did you play with anybody else?

We were interested in seeing how many chose to play alone compared to those who played with other people, so we first asked whether or not they played with anybody else and then had a follow-up question in each case, which is presented below. As we can see in Figure 5.24, 56 out of 77 respondents who answered the participant questionnaire – 73% – played with somebody else, while only 21 respondents did not.




Figure 5.24: Answer distribution I8: Did you play with anybody else?

I9a: If yes: I think the game was made more fun by playing with somebody else.

This data set shown in Figure 5.25 contains only those who answered yes to I8. No one disagreed or strongly disagreed that they thought the game was more fun because they played with somebody else, and only 12.5% were neutral. The rest of the participants all agreed or strongly agreed that the game was made more fun by playing with someone else, with 35 out of 56 strongly agreeing.




Figure 5.25: Answer distribution I9a: If yes: I think the game was made more fun by playing with somebody else.




Figure 5.26: Answer distribution 19b: *If no: I think I would have played longer if I played with somebody else.*

19b: If no: I think I would have played longer if I played with somebody else.

Because this question was given just to those who played alone and one of them chose not to answer, the data set represented in Figure 5.26 only consists of 20 respondents. Ten of those 20 people agreed or strongly agreed that they would have played longer if they played with somebody else, while three people did not think so. The remaining seven did specifically agree or disagree.

110: I would consider playing with strangers.

In addition to finding out whether people thought it was more fun to play with other people, we also wanted to know whether they would consider playing with strangers or whether they simply wanted to play with friends while passing by the installation together. As we can see in Figure 5.27, the amount of respondents who strongly agreed – 12% – was quite low, while 39% agreed that they would consider it. The amount of respondents with no opinion one way or the other was fairly high here at 24 people or 32%. 17% – 13 people – did not think they would consider playing with strangers.

111: If your child played the game: I would be comfortable with letting my child participate in a game/campaign like this.

Because of the fact that sometimes the game was tested with a campaign and sometimes without, this item exists in two different forms. One asks just about the game, while the other asks about the campaign as a whole, which




Figure 5.27: Answer distribution I10: I would consider playing with strangers.

includes the game but also the advertising aspect. For this reason, we have split up the answers into two groups, one with the campaign and one without. Only three people answered the question without a campaign and all of them strongly agreed that they were comfortable, so this is not displayed in the chart. We also asked those without the campaign about how they would feel if the game was a part of a campaign, the answers to this item can be found in Section 5.3.3.




Figure 5.28: Answer distribution I11: If your child played the game: I would be comfortable with letting my child participate in a game/campaign like this.

This item was originally meant for those who had children who actually played the game, but a surprising number of respondents without children answered this question as well. In order to have usable numbers, we decided to only count the answers given by those above the age of 25, as we had observed that no one younger than this had had children who played the game. We also felt that those at 25 years or over would be better suited to suggesting how they would feel even if they did not have children.

As we can see in Figure 5.28, only three people above the age of 25 (two of which had kids who played the game) answered the questionnaire

for the version of the game without a campaign. All of these strongly agreed that they were comfortable with letting their child participate in such a game, but the sample size is very small. Among those 13 who answered for the game with a campaign, the answers were a bit more spread out, with nine respondents agreeing or strongly agreeing that they were comfortable with their child participating in such a campaign, three who were neutral and one who disagreed. There is a difference here between those who answered with the campaign and without it, but it is difficult to say whether or not it is relevant because of the small sample size.

I12: Did you capture enough objects to win?

Figure 5.29 shows the answer distribution to this question, which is meant both as a means of judging the game's difficulty level and as a prelude to a question related to the marketing aspect (I15). In the chart, there is a "before" section and an "after" section, because we made winning easier about half way through our observation because of feedback that it was too difficult. The numbers might also be affected by the fact that most of the people who filled out the questionnaire had played for a little while and thus had a bigger chance of winning, because quite a few people played for about 20 seconds, gave up and then left without filling in a questionnaire.




Figure 5.29: Answer distribution I12: Did you capture enough objects to win?

We can see that after we made winning easier (changed from collecting three of the same kind to simply collecting five in total). Before changing the difficulty level, 30 out of 44 – 68 % – managed to collect enough. Afterwards, the number was 27 out of 33, or 82%.

I13: I wanted to win while playing the game.

Figure 5.30 shows that 59% of those who answered the questionnaire strongly agreed that they wanted to win while playing. This is interesting to see whether or not respondents had a goal in mind while playing, or if they simply played because it was fun/interesting in and of itself. Only 5% of the respondents answered that they disagree or strongly disagree that they wanted to win, which points towards winning being a big motivation.




Figure 5.30: Answer distribution I13: I wanted to win while playing the game.

This question is particularly interesting because it varied whether or not people won anything or if they got a “Congratulations! You won!” screen when they collected enough objects. At Location and at NTNU, winning would show an already existing advertisement, but this was unknown to players while they were playing. At Mercur, a “Congratulations! You won!” screen was shown for the first three days, while players could win a gift card for the last two days. This fact was written in the background of the game, which could increase their wish to win. However, perhaps due to a limited sample size, this did not affect the questionnaire data (the percentage of respondents who wanted to win at Mercur before and after we added the prize was the same), but we observed different behavior from users. These observations are described in Section 5.2.

5.3.3 Marketing aspect

I14: The game/campaign caught my attention to a greater degree than a poster would have done.

Figure 5.31 compares the responses from participants and non-participants about to which degree the game caught their attention better than a poster

would have done. As could be expected, a much bigger portion of those who played compared to those who did not play say that they agree or strongly agree. 91% of those who played (70 out of 77 respondents) say that they agree or strongly agree, while 56% (14 out of 24) among the non-participants say the same. 16% of the non-participants disagree or strongly disagree, while this number is only 2,6% among the participants.




Figure 5.31: Answer distribution I14: The game/campaign caught my attention to a greater degree than a poster would have done.

These numbers are to be expected, given that the people who chose to play obviously had to have noticed the game and paid enough attention to realize it was a game (and that they were a part of it) in order to do so. What we should keep in mind is that the non-participant part of these numbers does not include those people who simply passed by without looking (or barely looking) at the game, but chances are that these would pass by a poster without looking at it as well.

I15: If you managed to collect enough objects: I wanted to buy the advertised product more than before I played the game.

The first chart in Figure 5.32 shows the participants who played the game with an advertising campaign and their answers to whether or not they collected enough objects to win. As we can see, 36 respondents collected enough objects to be shown the advertising campaign, either at Location or at NTNU. Among these 36, only a single person strongly agreed that he or she wanted to buy the product more than before playing. However, 11 people – 31% – agreed, so one third (12 out of 36) either agree or strongly agree. We




Figure 5.32: Answer distribution I15: *If you managed to collect enough objects: I wanted to buy the advertised product more than before I played the game.*

also see that a fairly large percentage of people – 44% – either disagree or strongly disagree that playing the game increased their willingness to buy the advertised product.

I16: I would be comfortable with participating in a game like this if it was a part of an advertising campaign.

This question was asked to those who played the game while it was not part of a campaign advertising a specific product, which means the participants at Pirbadet and Mercur. We wanted to know whether they would feel comfortable playing the game if it was part of an advertising campaign, and as we can see in Figure 5.33, 14 out of 29 respondents (48%) strongly agreed that they would feel comfortable. Only one person disagreed, while six were neutral and eight agreed.

I17: My willingness to buy a product from the advertiser increased after playing the game.

Unlike I15, this is more concerned with buying any product from the advertiser instead of just the advertised product. The data set reflected in Figure 5.34 is




Figure 5.33: Answer distribution I16: I would be comfortable with participating in a game like this if it was a part of an advertising campaign.

from NTNU and Location, where the game included an advertising campaign. Although the game was not themed for the campaign at NTNU, winning would show you the advertised product and the advertiser. The game was themed completely to Location's graphic profile. As we can see, only two people strongly agree and only four agree (12.5% in total) that their willingness to buy a product from the advertiser increased. The largest percentage is neutral, at 46%, while 14.5% disagree strongly and 27% disagree that it increased their willingness.




Figure 5.34: Answer distribution I17: My willingness to buy a product from the advertiser increased after playing the game.

I18a: If your child played the game: I would be comfortable with letting my child participate in a game/campaign like this..

The answer distribution here can be found in Figure 5.28, with the relevant bar being the one with respondents who played the game with a campaign. 9 out of the 13 respondents agreed or strongly agreed that they were comfortable with letting their child participate in a campaign of this kind.

118b: If your child played the game and it was not part of an advertising campaign: I would be comfortable with letting my child be a part of an advertising campaign like this.

This question was asked to those who had a child who participated in the game while it was not advertising anything, that is to those at Pirbadet or at Mercur. This includes only three respondents, and all three respondents strongly agreed that they would be comfortable with letting their child participate in an advertising campaign in the form of this game.

5.4 Interviews

The interviews that were performed were mainly short conversations with participants or non-participants, as a follow-up to a questionnaire or as an alternative if they did not want to fill one out. They were not the main focus, but we did not want to miss potentially interesting information that a questionnaire cannot capture.

While we were testing at the campus, one person commented that after playing and winning the game, people are in a positive state of mind, so perhaps those positive feelings could be tied to the advertisement or offer that appeared when winning. He did say that a special code or something similar would be better than a general advertisement, because many tend to see advertisements as something negative and after winning, you expect something good to happen – which could lead to being let down and removing the positive feelings the game had drawn forth. The expectation of something good, like a prize, was presented by several people at the testing locations.

During a short conversation with a 13 year old girl at the shopping mall before we added the gift certificates, she said that she thought it sounded really fun if the game would be part of an advertising campaign. She also commented on the question related to playing with strangers, and said that it sounded fun, but she was not sure how to actually start and she was also worried that they might not be “nice” to her.

Both while testing at the campus and at cinema kiosk (Location), a few people commented that they were not entirely comfortable with playing in such

a public space. One of the guys who played on campus said that he would most likely not have stopped if the area around the installation had been as crowded as it tends to be during lecture breaks, because he would not be comfortable with all those people around. A non-participant at the cinema kiosk also said that the game was a bit too public when asked, which had led to him not choosing to play. He also did not want to fill out the questionnaire.

Through a conversation with a man in his fifties at Location, who had never played a Kinect-based game before, we learned that while he initially felt that the game was quite difficult to control, he felt that he improved quickly. Controlling the objects was unfamiliar at first, but once he got used to how they moved, he said that he learned how to control them quickly and he commented that the learning curve was quite steep.

5.5 Input logs

In this section we will present data that were logged while the game was running. Statistics from each location will be presented in its own subsection. These data mostly show how many players each location had and how well they played, together with an overview of the game's uptime.

The logged data proved to be both redundant and inaccurate, as the sensor often tended to loose and re-register participants while they were playing. The user information we logged was:

ID Each user get an ID from OpenNI when detected by the sensor. This ID is a number between 1 and 10 and represents a unique user that is currently present in the scene.

Time stamp The time stamp when the user was first detected.

Start time The time elapsed since the game started when a new user is detected

Spent time The time the user spent in the scene. We recorded the time elapsed when the user was lost and subtracted the start time from this to get the time spent.

Number of times won The number of times the winning conditions (collect three balls of the same color/five objects etc.) was fulfilled while this user was present.

Number of times lost The number of times the losing conditions was fulfilled while this user was present.

Number of users present The maximum amount of users that have been present together with this user.

Because of the redundancy in the data set (caused by the sensor and our choice on how we handled users), there was a need to filter them to get realistic numbers. We began with merging all users that had identical IDs and were active in approximately the same time period (four records in a row with user ID 3, for instance). Then we filtered on play time, as we were most interested in users that actually spent some time in front of the installation, not just quickly passing by in the background. We chose the minimum play time to be 40 seconds. Even with a filtered data set, we still perceive the numbers to be a bit too high, as our observations does not match with how many people the logs says played the game. One reason for this is that the sensor will register people that are in the background, even though they are not playing.

We logged for each game mode (two-player and multiplayer), but since we rarely tested the two-player mode, the numbers from the two modes have been merged together. The numbers from the two-player mode are small enough to be negligible. The logs from Pirbadet, for example, showed 1502 users of the multiplayer mode versus 37 users of the two-player mode.

5.5.1 Location data

Each of the following sub sections includes one table and one graph. The table gives numbers from the whole period and the graph visualizes how many users played each day (legend "Players"), how many wins (three similar colored objects or five objects collected, legend "Wins") and how many hours (legend "Hours") that were recorded that day. The vertical axis shows the percentage of the total recordings that were recorded that particular day. If we use Pirbadet's graph as an example (Figure 5.35), we see that 20,9% of all the winning events and 18,1% of all the users was recorded on the 17th of

April. The game was up and running for 7 hours, or 11,4% of the total uptime for this location. The purpose of the graph is to see which days were the most popular, if more users gave more winning activity and to compare these with the uptime.

5.5.2 Pirbadet waterpark

Number of days on display:	8
Number of users after merging IDs:	4880
Number of users after filtering:	1539
Users that got at least one win/loose event:	476
Average time spent in front of the installation:	98 seconds

Table 5.1: Statistics from Pirbadet

Pirbadet has the highest amount of input logs, as the installation was displayed at this location for over one week. As mentioned in Section 5.1, this location served as a test site, and we changed the game a couple of times during the time period that it was displayed. Numbers from Pirbadet can be found in Table 5.1 and an overview of each day can be found in Figure 5.35. The first day was mostly dominated by eager employees that had a lot of fun challenging each other. The two next days illustrate what we saw in our observations, namely that the basket was too small. People struggled to win, and the basket size was increased on the 20th of April. After this day, more people were able to score.




Figure 5.35: Distribution of players, winning events and hours played at Pirbadet

5.5.3 Stripa, NTNU campus

Number of days on display:	3
Number of users after merging IDs:	3578
Number of users after filtering:	449
Users that got at least one win/loose event:	207
Average time spent in front of the installation:	104 seconds

Table 5.2: Statistics from Stripa

The installation at Stripa was displayed for the least amount of time of all the locations. As it was in the beginning of the exam period, people was not as stressed as they normally would be when traveling to and from lectures. This meant that people had the time to stop and play on their way to their lunch break. Statistics from Stripa can be found in Table 5.2, and an overview of each day can be found in Figure 5.36.




Figure 5.36: Distribution of players, winning events and hours played at Stripa

5.5.4 Location, cinema kiosk

The beginning of the week is not the most popular time of the week to visit the cinema, as can be seen from the graph in Figure 5.37. Thursday is closer to

Number of days on display:	4
Number of users after merging IDs:	3702
Number of users after filtering:	343
Users that got at least one win/lose event:	94
Average time spent in front of the installation:	76 seconds

Table 5.3: Statistics from Location

the weekend, and thus has more visitors. But the numbers may be artificially high, as the sensor often picks up people that are waiting and talking in front of the sensor and not playing the game. The reason for the low uptime on the 9th of April was because we displayed the interactive poster the first two and half hours of the period in stead of the game. Statistics from Location can be found in Table 5.3. An overview of each day can be found in Figure 5.37.




Figure 5.37: Distribution of players, winning events and hours played at Location

5.5.5 Mercur shopping mall

The graph from Mercur is by far the most interesting one out of the four locations. The three first days were run as the normal game, themed for Mercur and with the "collect five objects mode. At the fourth day Mercur decided to add a prize that you could win. The visitor numbers are of course higher on a Friday and a Saturday than the rest of the week, but we also saw

Number of days on display:	5
Number of users after merging IDs:	23033
Number of users after filtering:	1750
Users that got at least one win/loose event:	836
Average time spent in front of the installation:	77 seconds

Table 5.4: Statistics from Mercur

an increase in activity by observing, which clearly also can be seen in Figure 5.38. The high percent of wins was caused by small groups of people that desperately wanted to win the prize who played the game repeatedly for a long time. Statistics from Mercur can be found in Table 5.4.




Figure 5.38: Distribution of players, winning events and hours played at Mercur

Discussion

The overall aim of our thesis could be separated into three major areas – user experience, marketing and technology. This section will consider each of these and discuss our findings in relation to the goals and the research questions that belong to each goal. We will start off by discussing the various aspects of the user experience and behavior, and move on to how this relates to the marketing value of such an installation. Following this is a section on the marketing potential of the installation, where we present relevant findings from the previous chapter and some sales data. Finally, we discuss our technology related findings from the literature study and our experience with working with OpenNI and try to answer the research questions related to this.

6.1 User experience

One of our three major goals was understanding and evaluating how users would relate to an interactive campaign in the form of a motion-controlled game, and the four research questions that relate to this goal are listed below. We will discuss each of them in this section by looking at results from Chapter 5, and any other relevant observations will be discussed at the end.

- RQ1.1** Are people comfortable with being a part of a motion-controlled game in a public space without explicitly giving their permission or actively making an effort?
- RQ1.2** Which user groups are most engaged in using the application?
- RQ1.3** Do people find it easy to get started with and control the application?
- RQ1.4** Are people who play with a group more engaged in using the application than those who play alone?

6.1.1 Comfort levels

Participants were asked how comfortable they were with being a part of the game without giving explicit permission and how comfortable they were with playing a motion-controlled game in a public space. Over 80% of the participants reported that they felt comfortable with being a part of the game without giving permission, but it is also unlikely that they would have stopped to play if they had not been. Several people, both among participant and those who simply passed by, expressed excitement at seeing themselves represented on the screen.

Among the non-participants, the answers were a bit mixed. There was a greater portion of people who did not feel comfortable, but also a relatively large portion (over 40%) who said that they were comfortable, but still chose not to play. This decision could be affected by several factors, such as whether they thought the game seemed fun or whether they were busy. On a couple of occasions, we observed a reluctance from some adults to join in while young kids or teenagers were playing, which could also be a contributing factor.

The answers to the question about being comfortable playing in a public space were distributed very similarly to the previous question, and this was to be expected as we suspect that people might have a general sense of whether or not they are comfortable, without attributing it to factors such as playing in public or not being asked for explicit permission. While observing in the field, we did notice that some commented on feeling silly and being a bit embarrassed at playing in public, this was often an initial reaction and seemed to pass. The majority of the participants reported that they became unaware of their surroundings while playing (see Figure 5.23), and this might be a contributing factor in making it less embarrassing. We also observed several teenage girls who felt it was embarrassing, but still wanted to play in order to see if they would win the gift certificate. They would encourage friends to join them, so that they would not have to play alone.

The non-participants were asked whether or not they would be more comfortable with using a controller, and as we can see in Figure 5.19, no clear trend emerged here. We did see a slight tendency towards women answering that they would have preferred a controller while it did not seem to matter for men, but the sample size was probably too small for this to be significant.

Overall, there were very few negative responses to our questionnaire,

but this could easily be a consequence of the fact that those who were uninterested would just walk past the game or not want to fill out the questionnaire if asked. Because of some of the properties of a public space, which we will expand upon in the next section, many people who passed the installation simply did not move out of the passing-by phase, as defined by Michelis and Müller (2011). However, once someone actually saw the game and reacted to it, they would often perform some subtle interaction, for example in the form of a wave or slowing down to look for their silhouette.

6.1.2 User groups and engagement

Our second research question had to do with which user groups were most active and engaged in using the application. This can be affected by several factors, such as the type of application, the design of it, potential rewards, location/context of the installation, comfort level with playing in public and so on. Knowing which user groups are most engaged is also beneficial from a marketing perspective.

As we expected, the overall distribution of replies to our questionnaires was skewed towards males, and they were the ones who were most willing to stop and play at three of our test locations and thus were given the questionnaire most often. At the shopping mall Mercur, however, the majority of those who played seemed to be girls, but the same groups of girls came back several times to play in order to try to win the gift certificate. This could be attributed to the fact that malls are popular places for teenagers to hang out, and you will often see groups of girls spending quite a bit of time there. Unlike at Mercur, we did not see any groups of only girls stop to play at the university campus, but girls who were in groups with boys would stop several times and they were very active.

The age distribution of the respondents shows a majority of respondents in their twenties, but this is primarily caused by the fact that one of our testing locations was a university where participants were primarily students. At locations like the waterpark and the shopping mall, we noticed that those who played the most and seemed to become most engaged were definitely children and young teenagers. While teenagers often were a bit skeptical at first, children seemed to have no inhibitions about playing and would become very immersed and seemed to become completely unaware of their

surroundings. Adults were more restrained, but some seemed curious and those who did play were surprisingly willing to move around a lot and seemed to be quite focused on the game and trying to win. This was also confirmed by our questionnaire, where over 60% of respondents said they became unaware of their surroundings while playing. Over 80% replied that they wanted to win while playing the game, which seems to correspond to the observations we made about older players, but younger children seemed to care more about having fun.

We asked participants about their gaming habits in order to see whether or not this had an effect on how they used the installation, but a large majority of our participants (67 out of 77 respondents) reported that they play between 0 and 9 hours on average per week. This made categorizing based on how much people played difficult, as the group of those who played more was too small for us to see any correlations. The same is true for which platforms people use – most of the answers included mobile phones/smartphones, for example, and there were no clear correlations between platform use and answers to other questions.

For most of the tests, people generally did not play for very long. In Section 5.5 we see that after filtering out those who were in front of the display for more than 40 seconds, the average playing time varied between 76 and 104 seconds, and these numbers include people who stood in front of the display walking without being aware of it and so on. We suspect the short play times have to do with the fact that people in a public space are generally heading somewhere and have another goal of being there besides playing with an interactive campaign. There is also the chance that they feel like they are obstructing others' path. These findings correspond to findings by Michelis and Müller (2011) for the *Magical Mirrors* installation. However, after we added the possibility of winning a gift certificate, we noticed a definite increase in how long people played and how engaged they were.

Michelis and Müller (2011) describe that 70% of those who had some direct interaction with the installation did multiple interactions, but they had four displays in a row where a user could interact with all of them. We did not see numbers anywhere close to this, as turning around to interact more after passing the display the first time seems to be quite a large threshold to pass. We also think that once people had explored the game for a little while and

satisfied their curiosity, there was no good reason to continue playing unless there was some external motivation, for example in the form of a possible prize. Once we added the possibility of a prize, several groups of people played many times in a row and the number of multiple interactions increased greatly. Some users also commented that they wanted a high score option to make it more exciting to play and to motivate them further, but this would also increase the threshold for interaction and perhaps cause greater annoyance when other people affected their playing.

6.1.3 Ease of use

Our questionnaire contained three questions about the ease of use of the game – one concerned with getting started, one with controlling the game and one about whether or not the respondent thought most people would be able to play the game. The answers to these were generally positive, and a large majority thought the game was easy to get started with and that most people would be able to play it, while a slightly smaller majority thought it was easy to control. We did see that quite a few people initially had problems with gauging how their movements would affect the balls on the screen – as an example, many would hit the balls too hard. However, once they had played for a little while, their control increased quickly and we also had some comments saying that the learning curve was steep, and that while it was difficult to begin with, learning how to move went quickly. Balancing the feeling of control with challenging aspects is one of the main motivating factors described in Section 3.2.

Although some initial problems with controlling the balls were present, most people seemed to realize that they could control the game with their body once they saw the silhouette of themselves. A few people did not seem to notice the silhouette initially and then walked towards the screen, which moved them too close to the sensor and caused it to no longer draw their silhouette, and tried touching the objects on the screen. This might have something to do with the fact that while (multi-) touch technology has become very common in the past few years, using motion-controlled application is still a relatively new concept to many people. The Kinect has been out for about a year and a half, but it has more or less just been used for games and not everyone has tried playing it.

The game initially required players to gather three objects of the same color into the container, but this proved to be quite difficult and after switching to capture simply five objects, we saw that the number of wins went up. Another problem with gathering three objects of the same color was that many people simply did not read the text in the middle of the screen that explained what the objective of the game was. Several people would play for a little while before reading the text, and others did not even notice the container, which was placed on the bottom left of the screen. This might have been because we saw a tendency to stand in the middle of the screen and focus on that area. Making the game easy enough to immediately understand seems to be important, but there must also be a balance between making it easy to control and easy enough to understand and at the same time difficult enough that it provides a challenge. Once the game became easy to control and they had managed to win, people were quick to leave – especially those who played alone.

While testing the two-player mode, we saw that it was much easier to capture three similar objects or five objects in general, because people in the background would not disturb the object's trajectory and be in the way of the container. This was the main source of annoyance among users, and they often complained that it made things difficult. While having to perform some sort of gesture in order to play made the control easier, it made it much more difficult to get started playing. Many people seemed to not read that they had to perform a wave gesture in order to play, and among those that did perform a wave gesture, the sensor did not recognize them very easily, so they quickly gave up and left.

Finding a good way to select only a few players could be important (depending on the type of application), and there are several ways to go about this, where some are easier for the user than others but might have other disadvantages. One could use a standard calibration pose, but chances are good that players would have to have this demonstrated to them, which increases the difficulty. Another option is basing the selection process on how close to the sensor someone is or whether someone stops in front of it. We tried the wave gesture, as waving is a pretty natural thing, but it turned out that the sensor's gesture recognition was not good enough to recognize all the variations of waving people did. Selecting people who are close to the

sensor would be easier for users, but people who walked between those who were playing and the sensor might ruin things. Basing the selection process on whether or not someone stopped in front of the sensor could work for other purposes, but in a game like ours, people tended to move around and not stop for very long. Overall, it is important to keep in mind that adding a step like this might improve the control mechanism, but also increase the threshold for playing significantly, and it might not be worth it for some types of applications.

6.1.4 Group dynamics

As expected, passers-by who were in groups were much more likely to stop and play, and also played for a longer time when they stopped. People who were walking alone were much less likely to pass over from the viewing and reacting and subtle interaction phases of the Audience Funnel to the direct interaction phase, while those who were in groups and discovered themselves on the screen very often interacted directly, although the length of these interactions varied greatly. Some groups at the shopping mall, after we added the prize, would play for as long as an hour spread out over multiple interactions, while others interacted for less than a minute.

Over 70% of the respondents to our questionnaire played with others, and an overwhelming majority of them said that they thought the game was more fun because they played with someone else. Half of those who played alone thought they would have played for a longer time if they had been playing with someone else. The answers to whether or not respondents would consider playing with strangers were fairly mixed, but about 50% agreed that they would consider it. We also observed that those who played in groups seemed to be much less self-conscious while playing, which makes sense because when you play in a group, each single person stands out less than the group as a whole. We sometimes observed that if one person in a group wanted to play, he or she would convince the others to join in because that would be less embarrassing.

When groups played together, they would very frequently choose to collaborate instead of competing with each other. They would interact with each other in front of the sensor and see how this looked in the game, and would often come up with their own strategies for winning the game more efficiently. They were also more likely than those who played alone to simply

play around with the objects, and coming up with their own mini-games and simply having fun instead of necessarily trying to reach the objective of the game. We also noticed that if one person in the group understood what the goal of the game was before the others, they would share the knowledge with the others. A couple of times at the shopping mall we also noticed that people who were watching strangers play, and who had played themselves previously, would explain how the rules if the people playing did not immediately understand what to do. Collaboration is described as an important motivating factor in Section 3.2.2, and our findings definitely support this.

6.1.5 Other observations

We also found some interesting results that do not necessarily relate to one of the research questions, but which we still felt were useful. As previously mentioned, players would often play around with the game without really trying to win (especially after having already won once and being satisfied with that) and create their own mini-games, such as gathering as many balls as possible in their arms and so on. After adding the prize, however, we noticed that the focus shifted drastically towards winning as many times as possible and almost no one simply played around for fun. This must be taken into consideration when designing an installation of this type, because how users behave will vary greatly depending on, for example, whether the goal is to simply engage the crowd and create an attraction or selling a product.

As mentioned in Section 3.2.2, grabbing the attention of users and passing the threshold from peripheral awareness to focal awareness can be difficult. Except in the setting at the cinema kiosk Location where our game did not stand out from all the other screens, those who did view it seemed to react to seeing their silhouette and thus giving it their attention. Quite a few people of those who viewed it and reacted ended up subtly interacting with it, which is a way of crossing that threshold. The fact that an installation is unusual and stands out also helps attract attention, as well as the honey pot effect mentioned in Section 3.2, and immediately seeing yourself on the screen seemed to motivate users.

Getting adults to stop and play seemed difficult, which we attribute to their priorities while in a mall or a similar setting. They often have a specific

goal in mind and attempt to ignore distractions. For this reason, we think it might be interesting to see how they relate to an interactive installation like this in a place where many people sit around and wait, for example in a train station or an airport. These are still relatively public places, but people tend to spend a longer time there and often don't have a specific goal in mind, except for waiting for something. One potential problem could be that many people are alone in circumstances like these and the most active players were definitely groups, but it would still be an interesting experiment.

6.1.6 Interactive poster

As we mentioned in Section 5.2.5, the interactive poster was not tested for very long, but we did see that in an environment like at the cinema kiosk, the poster did not stand out and no one seemed to notice that it was different from all the other screens. When we tested at the school, where the screen itself stood out, people noticed it to a greater degree than a static poster, and seemed to want to interact with it more than what was intended. This might be caused by the fact that they realized that it was different than a regular poster and wanted to explore and see what it could do. Some could also have seen us test the game earlier. We initially wanted to see whether people would try to interact with the poster in a more passive way than a game requires, but our observations pointed towards people wanting more active interaction.

6.2 The intersection between user experience and marketing

One of our goals was to explore the marketing aspect of a public interactive display, like the one we developed, and we wanted to find out how the user's experience and behavior affected the marketing value and how the marketing aspect affected the user's behavior and experience.

G2 The study seeks to explore how motion-controlled applications can be used within interactive campaigns in a public space.

RQ2.1 Are people comfortable with being part of an interactive campaign like this?

One of the research questions connected to this goal had to do with how comfortable the participants were with being involved in an interactive campaign like this, while the other one was more focused on which effect the campaign had on users and their willingness to buy something from the advertiser. The last question will be discussed in Section 6.3, while the most relevant question for this section was RQ2.1. In this section we will also discuss some other observations we made that relate to the intersection between the marketing aspect and the user's behavior and experience.

6.2.1 Level of comfort with participating in a campaign

Given that we tested the game with and without an advertising campaign, those that were a part of the game with a campaign were asked about their comfort-level with the campaign as a whole, while those who played the game without a campaign were asked about their comfort-level with being a part of the game and then about how comfortable they would be if it had been a part of a campaign. We observed no discernible difference in how people approached the game with or without an advertiser's graphic profile – the differences we saw between locations (which is what the game's look depended on) seemed to have more to do with how busy users were, whether or not they saw the game at all and which user group they were a part of.

46 out of 48 respondents who played the game with a campaign answered that they felt comfortable with being a part of the campaign without explicit permission, while the two remaining respondents were neutral on the topic. Out of those who played without an advertising campaign and were asked if they would be comfortable with playing the game if it was a part of an advertising campaign, we can see in Figure 5.33 that a large majority of the respondents agreed or strongly agreed that they would be comfortable. These numbers could, of course, be affected by the fact that those who were asked had already played the game.

We also asked participants about how they would feel or how they felt about their children being a part of a campaign of this type. The answers

we got were surprisingly positive, as we can see in Figure 5.28, 9 out of 13 participants who played the game while it was a campaign agree or strongly agree that they would be comfortable with letting their child participate in this kind of campaign. Those who played the game without a campaign also agreed that they would be comfortable with letting their children participate in the game if it was a part of a campaign, but the sample size here was very small. The number of respondents to these questions were quite low, so the scope is not large enough for us to generalize and draw any conclusions based on them. In order to keep the questionnaire short, we did not ask non-participants this question, so the fact those who answered had already tried or watched their kids try the game could definitely have affected their answers. We did observe that several parents dragged their children away from the game, both when it was an advertising campaign and not, which we mainly attributed to the parents seeing the game as a disturbance, as they did not pay enough attention to it to likely have realized whether it was simply a game or an advertising campaign.

6.2.2 Placement of the display

During the observation at the cinema kiosk, we noticed that very few people even noticed that something unusual was happening on the screen, even though it was placed in the middle of their field of view when entering the kiosk. As explained in Section 5.1.3 and in the observation, there were 16 other screens in the kiosk, and they all had the same graphical profile, which we suspect led to people thinking that the game was just another regular advertisement. We asked some of the people sitting outside the kiosk if they had noticed that there was a game they could play, and most of them said that they had simply thought it was another regular advertisement. Because of the fact that very few people noticed the screen, the marketing value for such a game might be limited when it does not stand out from its surroundings.

At locations like the campus and the shopping mall, the game or the screen stood out because it was not surrounded by other screens showing very similar graphics, and a much larger portion of people looked at it and reacted to it. It seemed obvious that people had learned to more or less ignore the multitude of screens displaying advertisements, unless they were looking for an interesting offer, while something that looked out of place or was

surprising (i.e. the game at other locations) caught their attention. Surprise in the form of something being different than expected is one of the general models for attracting attention, per Section 3.2.2, and we definitely saw how important it was to attract attention when we compared how the passers-by acted at the cinema kiosk compared to the other places where we tested the application.

6.2.3 Reward or no reward

Although the effect of adding an actual prize will be discussed in Section 6.3, we made another interesting observation related to user experience and rewards. At two of the locations, the campus and the cinema kiosk, players were “rewarded” with a standard advertisement when winning, which several people expressed disappointment with, and one woman said she felt tricked. However, at the waterpark and the shopping mall where there was a simple “Congratulations! You won!” instead of a reward, most people seemed satisfied. Many expressed joy at winning, and some continued playing while others gave the impression that playing had been fun and then left. A couple of people commented that it would be cool if there was a prize or reward, but no one seemed as disappointed as when they were shown advertisements for already existing offers.

6.3 Marketing potential

The research question that is directly addressing the marketing aspect is research question 2.2:

RQ2.2 Does an interactive campaign of this type have any effect on people’s purchasing behavior?

In order to address this research question, we will use data gathered from the questionnaire and the observation, as well as some sales figures we received from the cinema kiosk (Location). The advertised product was gourmet lollipops and the game had a popcorn theme, so we received sales data for these products from two weeks prior to the game being tested, for

the time period we tested the game and for two weeks after. None of these time periods included public holidays, which is why they were chosen over the weeks immediately prior to and after the testing period. We do not have sales data from the campaign on campus, where a baguette deal was advertised.

6.3.1 Effect on sales

Tables 6.1 and 6.2 present the numbers we got from Location. As we can see in Table 6.2, the changes in gourmet lollipop sales were not very noticeable and they kept increasing after we left. Popcorn sales, on the other hand, increased with 18.2% despite a slight decrease in the number of paying customers. We also see that despite a 31.5% increase in the number of paying customers two weeks later, the popcorn sales dropped again compared to the week we were testing our game.

Week:	Time period:	Paying customers:
17	23.04 - 26.04	988
19	07.05 - 10.05	969
21	21.05 - 24.05	1274

Table 6.1: Paying customers from three time periods at Location

Compared weeks:	Customer difference:	Increase in gourmet lollipop sales:	Increase in popcorn sales:
17 and 19	-2.0%	4.7%	18.2%
19 and 21	31.5%	1.2%	-8.2%

Table 6.2: Comparison of sales data from Location

The fact that the popcorn sales increased significantly while the gourmet lollipop sale difference was almost negligible makes sense given the fact that both players and passers-by spent more time watching the popcorn theme of the game than the lollipop advertisement that would appear when winning. We also observed at other locations that several people would play around with the game without necessarily trying to win, which meant that the potential advertisement would rarely be seen. Because of this, it might be beneficial to provide some form of advertisement within the game itself, like with the

popcorn at Location, because this is visible to both the players and the passers-by most of the time. Unlike with a poster, passers-by cannot glance quickly at the display and see the advertised product if it is only displayed when someone wins.

6.3.2 Effect on user behavior

The most relevant items from the questionnaire are I15 and I17, which refer to wanting to buy the advertised product after playing the game and whether or not their willingness to buy any product from the advertiser increased after playing the game. As we can see in the answers to I15 in Figure 5.32 and the answers to I17 in Figure 5.34, people's responses were quite mixed with a lot of people disagreeing or strongly disagreeing. These numbers do not reflect well on the campaign's effect on marketing, but during our observation we noticed that several people were disappointed with being shown an already existing offer when winning, so the data could be affected by this disappointment.

The people who expressed disappointment when managing to collect the right amount of objects, seemed to think that they would be rewarded with something. As an example, a couple of players on campus thought they had won a free baguette, and several others expressed a wish for some kind of prize. We believe that the prize does not need to be extraordinary, but big enough that the participants are motivated and feel rewarded while playing. As we saw in Figure 5.30, a desire to win while playing was very common among the participants. A small discount or codes for winning stuff can be enough to make them feel like the time was worth spending. During our time at the shopping mall (Mercur), we saw a clear jump in both active players and number of times that the players won on the last two days where we distributed codes for winning a gift certificate. This increase can clearly be seen in Figure 5.38 in Section 5.5.5. Even though they were not guaranteed a win, the chance seemed to be enough motivation.

After adding the prize, we noticed that some groups of people played for a long time and would come back to play multiple times. Although this in some cases seemed to prevent others from trying, it also created a honey pot effect (explained in Section 3.2) which leads to a lot of people noticing the campaign because others are giving it their attention. In Figure 5.31 we

presented the answer distribution for the statement: “The game/campaign caught my attention to a greater degree than a poster would have done”. Not surprisingly, most of the respondents that played the game agreed that the game caught their attention better than a poster. This is most likely caused by a combination of the honey pot effect and the game itself grabbing their attention because they could see that something they did not expect was happening - i.e. that a game was present and that their silhouettes were appearing in the game.

Some of the results from both our observation, the questionnaire and the sales figures provided by Location point towards the usefulness of such a game as an attraction instead of as an alternative to a poster. The honey pot effect became very obvious, especially at Mercur, where people would stop and watch others playing, and the responses to whether or not the campaign caught their attention better than a poster were largely positive. This combined with the sales figures from Location, where we saw a significant increase in the amount of popcorn sales, lead us to think that it could be useful as a way of attracting a crowd and highlighting a product. At the current time, this type of game is a fairly new thing, and the novelty factor might play a big role. It is difficult to predict how users will relate to an installation when the novelty factor wears off, but we did see some repeat interactions at Mercur when there was sufficient motivation in place.

6.3.3 Interactive poster

While testing the interactive poster, we observed that it drew more attention than a static poster, except in a location where it was surrounded by similar-looking screens. This could point towards motion sensing technology being useful in advertisement without having to be used in such an active and engaging way as it would be in a game. Passers-by seemed to notice that something moved on the screen and looked over at it, and some stopped to find out how it worked. This leads to both the participant and passers-by looking at the screen with the advertisement more carefully than they would with a regular poster, and could be beneficial for marketing purposes.

6.4 Technology

One of the goals of this thesis was to find which possibilities exist for developing motion-controlled applications and to find out how mature the technology is. We also wanted to see if it was possible to develop motion-controlled web applications, and these objectives were formulated into research question 3.1 and 3.2:

RQ3.1 Which options with regards to APIs and programming environments exist if you want to develop motion-controlled applications?

RQ3.2 Is the technology mature enough to be used within a browser for an easy to develop, cross-platform solution using modern web technologies?

Given that GI wanted us to use OpenNI, this is the technology we have become most familiar with and will be most capable of evaluating. We have, however, also looked at Microsoft's Kinect SDK as an alternative and also briefly on the open source alternative libfreenect by the OpenKinect community. These three are, as far as we can tell, the most relevant alternatives for developing motion-controlled applications. Which you choose depends on your platform, preferences and your desired functionality, but all in all they make it possible to develop motion-controlled applications for both Windows, Mac OS X and Linux, using a variety of languages.

6.4.1 Feature sets

Microsoft's Kinect SDK and the OpenNI framework are fairly similar feature-wise, while libfreenect is more bare bone and lacks many of the high-level features provided by the other two frameworks. The high-level API provided by libfreenect gives access to the depth and RGB image streams, but not much else (OpenKinect, 2012). Both OpenNI and the Kinect SDK also provide this information, in addition to things like information about users in the scene, their skeleton data and audio capabilities. libfreenect and the Kinect SDK only support the Microsoft Kinect device, while OpenNI supports both the Kinect and sensors by ASUS. OpenNI in general seems to be more modular and allows for adding components and modules from different vendors, as explained in Section 3.4.1. Given that OpenNI is an initiative to promote

compatibility and interoperability between natural interaction devices, applications and middleware (OpenNI, 2011), it makes sense that they would choose support multiple types of devices, while both the Kinect SDK and the Kinect device is released by Microsoft.

Although developing motion-controlled applications is entirely possible at the moment, using it for tasks that require a lot of precision might be difficult. Things like the frame rate and resolution of the depth map stream offered by both the Kinect and the ASUS Xtion Pro Live are, as we mentioned in Section 3.3, only 640x480 pixels at 30 frames per second, which is not very high. The silhouettes of people that can be produced using the depth map are also quite jagged, and while this can be improved by developers who are using the SDK or framework, this requires quite a bit of extra work. It is likely that as the sensors develop and improve, the resolution will increase and it can be better suited for things like 3D modeling, robotics and so on.

Both the OpenNI framework and the Kinect SDK are in a relatively early stage, so both are being developed continuously and the offered functionality and their APIs might change frequently. For example, version 1.5 of the public Kinect runtime and SDK was released by Microsoft on May 21st (Kinect for Windows Team, 2012) and added several new features, while OpenNI is starting to design and develop a new version of the API and is seeking feedback from developers at the time of writing this thesis (Hendel, 2012). Because of this, the functionality offered by each of them will vary and some of the comparisons made here might not be accurate in a month or two from now.

Among the functionality that they both offer, for example skeleton tracking, there are also some differences in which algorithms they use and so on. This might lead to differences in the accuracy and, for example, how many joints are available for tracking. With the 1.5 release of the Kinect SDK, Microsoft also added a Face Tracking SDK, which is functionality that OpenNI does not (yet) have. As mentioned in Section 3.4.2, OpenNI with the NITE middleware comes with some default gestures, such as swipe, push and wave, which Kinect does not. Several gesture libraries are available for the Kinect SDK, though, which means that developers do not have to rely on doing this work for themselves. OpenNI also comes with hand tracking by default, which the Kinect SDK does not appear to do, while the Kinect SDK supports seated

skeleton tracking.

One drawback of the Kinect SDK that would have been problematic for us is that the Kinect SDK can only track six users at the same time and only provide skeleton data for two (Microsoft, 2012b). OpenNI can track (display the silhouette of) as many people as it can fit in its field of view, and provides skeleton data for an unlimited number of users. This means that in a space where there are a lot of potential users, for example a public setting such as ours, not everyone who wanted to would be able to participate, and if we had chosen to use skeleton data instead of silhouettes, we would be limited to two players at the time. This might be okay for a setting like a living room, but is far from ideal if you want to engage crowds.

One major advantage of OpenNI over the Kinect SDK, is that it is open source. This has led to there being several projects, like Zigfu's development kit, DepthJS and simple-openni, that build upon it and make their own APIs for use in different programming languages than OpenNI provides wrappers for. For example, simple-openni has simplified the OpenNI API a lot and made it available for use with the language and IDE Processing, and there is no reason why this could not be done for other languages as well.

6.4.2 Documentation and ease of use

The first major difference between the OpenNI framework and the Kinect SDK that we noticed was the perceived difference in ease of use and available documentation. Although we have not used the Kinect SDK for developing anything, we have read up on it, looked into how to get started with it and taken a look at the available resources and documentation, to get some insight into what the major alternative to OpenNI is and how to use it.

Unlike Microsoft, OpenNI is a not-for-profit organization, and while its founding members are companies like PrimeSense and ASUS, with definite economic interests in advancing the use of this type of interaction, they are not one single company that provides the whole stack (device, device drivers, development platform, software development kit). Microsoft benefits greatly from motivating developers to use their products and their SDK instead of the competitor, and they provide developers with development tools, step-by-step instructions, multiple code samples and other resources in the form of a toolkit to get developers started. Because they are focused on the Windows

platform and only .NET-compliant languages, this is perhaps easier than it is for OpenNI, which supports multiple languages and platforms and also encourages third-party middleware.

We found that understanding how to use the OpenNI framework was a bit cumbersome, because of insufficient documentation and outdated information online. No specific guide states what you need to start developing, aside from the OpenNI binary. The NITE middleware is very useful when using OpenNI, but there is no clear explanation of what NITE is in relation to OpenNI and why one might want to use it. This is why Zigfu's install package is very useful, as it provides you with everything you need, including sensor drivers for the ASUS and the Kinect, OpenNI and NITE.

OpenNI's website provides instructions on how to get started with an empty project that uses OpenNI, but only for C++ or C development using Visual Studio. Some blogs and other websites have provided instructions on how to get started with the development in, for example, C# or Java and how to use the API in your code, but because of changes in the API after the initial release and excitement (which led to several guides and tutorials), these will often be outdated.

The OpenNI website does provide a programmer guide, which documents the framework and some sample programs. This documentation makes it evident that the focus has been on C and C++, but especially the sample programs are well documented and explain what the lines of code do and why they do so. Eight out of the 14 sample programs are written in C++ with the others being four in Java and two in C#, but the explanations of the code makes them useful even if the sample you are looking for is not present in your choice of programming language. The source code for these sample programs are a part of the OpenNI and NITE downloads, which means that they can be explored and used by developers to get started. The API reference is also written for C or C++, and while it is definitely understandable even without good knowledge of these languages, it would benefit from being presented in Java and C# as well.

6.4.3 Motion-controlled web applications

Although we had originally hoped to be able to develop a web-based application, we quickly realized that this would be risky. At the beginning of

the semester, the option that now exists – Zifgu’s JavaScript API – was not released past a closed beta version, which made us skeptical about basing the entire project on it. There was also a Unity wrapper available, which we could use to develop a game in Unity and present it in a browser using Unity Web Player, but this was not relevant for us. We wanted to see if it could be done using free, standard web technologies like HTML5, CSS and JavaScript.

During the spring of 2012, Zifgu has released their JavaScript API and browser plugin to the public, and it would now be possible to use it for developing a game or a motion-controlled web applications – with some limitations. For example, they do not provide access to image streams, so we cannot get the silhouette of a user. They do, on the other hand, provide information about a user’s joints and their positions, which could be used to create a model of a user in, for example, a `canvas` element using HTML5 and JavaScript. The browser plugin also fires events when a user performs certain gestures, which should make it possible to make web applications that could be controlled with gestures. The previously mentioned project DepthJS is also an option for making web applications that use gestures as a control mechanism, for example by swiping, pushing or pulling. Both of these projects provide their own APIs for developers to use, and they use OpenNI as their backend.

Zifgu makes it very easy to get started with the development process. Not only do they provide the necessary sensor drivers, the OpenNI framework and NITE as a simple to install package, but their API is fairly small and the documentation is very focused on allowing developers to get started quickly. For example, they provide a tutorial for writing a “user radar” application from start to finish, which will display a top-down view of users in the scene and how they move. In addition to this and the documentation of the API, they also provide recipes for coding common tasks like finding and losing users, reacting to push or swipe events and waiting for a user with full body skeleton data. There is also recipes for implementing a single user UI session, which can be practical if only one person in front of the screen should be able to control a web application. DepthJS has a demo file in their GitHub repository that shows how to use their API in a simple web page, with events like `onRegister`, `onPush`, `onMove`, `onSwipe` and so on. The developer can then add functions for what they want to happen when an event occurs.

Although DepthJS's API is very simple, they do also require installing a browser plugin and extension to get access to the Kinect data. Based on the documentation, this seems more complicated than it is for Zigfu's browser plugin, and the developers behind DepthJS admit that making the install process easier for end users is necessary. They also support fewer browsers (no Internet Explorer, for example, and they are working on Firefox support), while Zigfu supports the major browsers like Firefox, Chrome, Internet Explorer and Safari. An advantage of DepthJS over Zigfu's JavaScript API and browser plugin is that DepthJS supports Linux. Both would make it possible to develop motion-controlled web applications that relied on registering and reacting to gestures. However, if you want information like skeleton data for use with games or anything else than hand gestures, Zigfu is the obvious choice.

Conclusion

7.1 Summary

In this thesis project, we have implemented and evaluated an interactive campaign. We started by performing a literature study and evaluating the relevant technology, as well as deciding on our methodology for the evaluation. After this we implemented two prototypes that we wanted to evaluate, and put these into the field to study how they were used. The field studies consisted of mainly observation and questionnaires, supplemented by some short interviews and logs from the game itself. After studying how the prototypes were used in four different public locations, we discussed our findings in relation to our research questions.

Our field study indicated that the most active users of such an installation would be groups of people, and in particular children, teenagers and young adults. This does depend on the choice of location and type and look of the application, but it was a tendency at all our testing locations. Those who chose to play expressed that they were comfortable with both the public aspect of the game and the fact that they were automatically a part of the game. Both of our prototypes seemed to grab the attention of passers-by but the threshold for getting people from viewing, reacting and subtly interacting with the game to directly interacting with it was quite high. Most of the people who noticed the game seemed to immediately understand that their body was the controller, but adding a selection process in the form of a wave gesture caused significant problems and confused participants.

Although most interactions with the game were relatively short, we saw a considerable increase in the time spent playing after adding a potential prize for winning, which indicates that some form of external motivation is useful past the initial curiosity has been satisfied. The testing at two of our

locations rewarded users with an already existing offer when they won, while they were simply congratulated on the two other locations. Participants that were presented with the existing offer seemed disappointed when winning and expressed a desire for some form of special offer or prize, while those who were simply congratulated seemed satisfied. Global Illumination thinks that this observation is particularly valuable, as it means that unless the advertiser has something special to give away to winners, using the game itself as the advertisement is probably a good choice.

We have also spent some time evaluating the technology and looking at the options that exist, with the focus being on OpenNI and the Kinect SDK. Depending on what you are trying to accomplish, both have their advantages and disadvantages. For example, Kinect SDK is locked to Windows and Microsoft's Kinect device, but has the major advantage that the available documentation is much better, as Microsoft controls the whole stack. OpenNI, on the other hand, works on several platforms and supports devices from several vendors, and has also been used as the backend for some browser implementations, like Zigfu's JavaScript API and DepthJS. These projects have made it possible to develop motion-controlled web applications using a browser extension and standard, modern web technologies, but have some feature limitations compared to stand-alone applications using OpenNI or the Kinect SDK.

7.2 Feedback from Global Illumination

The feedback that we have received from GI highlights the following points as the most valuable results from our project.

- They have been given an indication of the complexity and cost involved with developing motion-controlled applications, where the conclusion is that they think there is potential for using this type of technology in interactive campaigns.
- We have uncovered which types of locations and venues work well for interactive campaigns of this kind, and given them some insight into what would be the best choice of equipment and technology.

- Our observations revealed which user groups are easiest to reach through this type of campaign, which is useful for marketing purposes.
- We have made some interesting observations about how people respond to the marketing aspect of interactive campaigns. For example, rewarding participants with an already existing offer led to more negative reactions than giving them no reward other than congratulating them.

7.3 Our recommendations

Based on our findings throughout the project, we have arrived at some recommendations that we feel would be useful for others to consider if they are doing similar work. They are grounded in our personal thoughts and experiences, and will be affected by choices we made concerning game play, appearance of the game and so on.

- Children, teenagers and young adults were the most active user groups, so campaigns that target these groups might have a better chance of succeeding than those that do not.
- Groups of people were more active and more likely to stop and play than people not in a group, and this can be taken advantage of when designing campaigns.
- The threshold for interaction should be very low. If the application includes some form of selection process to separate between active and inactive users, this process should be transparent for the user.
- Placement is important for grabbing the attention of passers-by, and the installation should stand out from its surroundings. If it is surrounded by other advertisements in the form of posters or non-interactive screens with the same graphical profile, we recommend that some changes are made to the interactive campaign to make it stand out more.
- If the amount of time spent in the game itself is high compared to the time spent presenting the reward when a player wins, we recommend that the advertisement message should be present in the game itself to maximize the exposure.
- The campaign must be designed with the environment in mind. If it is placed in a location where passers-by are in a hurry, the interaction

mechanism must be simpler than in a place where people tend to spend more time.

- Adding some form of prize upon completion of the game works as a motivating factor and increases the time people spend playing. However, it also shifts their focus from playing around in front of the installation and exploring it to simply trying to win as quickly as possible, which one must be aware of when developing this type of campaign.
- If participants are presented with something upon winning, we think it is important that this prize feels like a genuine reward for spending time and energy on playing the game. Treating an existing offer as a reward seemed to produce negative reactions, which should be avoided.

7.4 Further work

Given that combining motion sensing technology with public interactive campaigns is a fairly unexplored area, there is a definite need for further studies to be done. We think that this type of technology has the potential to engage users in a new and exciting way, and that we might see it used for a variety of purposes in the near future. The list below contains some areas that we feel should be researched further to gain a deeper understanding of how and when to use the technology.

- Explore and evaluate alternative methods for selecting one or just a few users from a crowded scene in front of the sensor, without increasing the threshold for interaction too much.
- Compare campaigns using RGB images of users (i.e. like a mirror), silhouettes and skeleton-based models, and how users react to each of them.
- Investigate the effect of adding social features like a high score list and integrating social media into a interactive campaigns.
- Perform more studies that involve actual advertisers and investigate the effect the campaign has on sales figures and the public's opinion on the advertiser.

- Explore methods for involving user groups that are not immediately drawn to the installation, for example by finding out how to make it more appealing to adults.

References

- Anthony, S. (2010). Microsoft Kinect controlling Windows 7, exciting proof of concept. Website. <http://downloadsquad.switched.com/2010/11/22/microsoft-kinect-controlling-windows-7-exciting-proof-of-concept/> [Accessed on April 16th 2012].
- AR Door (2011). Virtual Fitting Room for Topshop. Website. <http://ar-door.com/2011/05/virtualnaya-primerochnaya-dlya-topshop/?lang=en> [Accessed on May 29th].
- Basili, V. R., Caldiera, V. G., and Rombach, H. D. (1996). The Goal Question Metric Approach. *Encyclopedia of Software Engineering*, 2:528–532.
- Bleiweiss, A., Eshar, D., Kutliroff, G., Lerner, A., Oshrat, Y., and Yanai, Y. (2010). Enhanced Interactive Gaming by Blending Full-Body Tracking and Gesture Animation. In *ACM SIGGRAPH ASIA 2010 Sketches*, SA '10, pages 34:1—34:2, Seoul, Republic of Korea. ACM Press.
- Bobleanta, V. (2010). DepthJS for Chrome lets you use a Microsoft Kinect to browse the web. Website. <http://downloadsquad.switched.com/2010/11/26/depthjs-for-chrome-lets-you-use-a-microsoft-kinect-to-browse-the/> [Accessed on April 16th 2012].
- Borenstein, G. (2012). *Making Things See: 3D Vision with Kinect, Processing, Arduino and MakerBot*. O'Reilly Media, Inc., first edition.
- Boulos, M. N. K., Blanchard, B. J., Walker, C., Montero, J., Tripathy, A., and Gutierrez-Osuna, R. (2011). Web GIS in practice X: a Microsoft Kinect natural user interface for Google Earth navigation. *International journal of health geographics*, 10(1):45.
- Box2D (2012). Box2D - A 2D Physics Engine for Games. Website. <http://box2d.org/> [Accessed on May 16th 2012].
- Brignull, H. and Rogers, Y. (2003). Enticing People to Interact with Large Public Displays in Public Spaces. In Rauterberg, M., Menozzi, M., and Wesson, J., editors, *Human-Computer Interaction, INTERACT '03*, pages 17–24.

152 REFERENCES

- Brooke, J. (1996). SUS - A quick and dirty usability scale. In *Usability Evaluation in Industry*. CRC Press.
- Callaham, J. (2012). Kinect for Windows third party media viewer released. Website. <http://www.neowin.net/news/kinect-for-windows-third-party-media-viewer-released> [Accessed on May 15th 2012].
- Cappellini, L. (2009). Interactive installation art in Italy: an analytical survey.
- Catuhe, D. (2011). Gestures and Tools for Kinect. Website. <http://blogs.msdn.com/b/eternalcoding/archive/2011/07/04/gestures-and-tools-for-kinect.aspx> [Accessed on April 16th 2012].
- Chang, Y.-j., Chen, S.-f., and Huang, J.-d. (2011). A Kinect-based system for physical rehabilitation: A pilot study for young adults with motor disabilities. *Research in Developmental Disabilities*, 32(6):2566–2570.
- Costello, B., Muller, L., and Amitani, S. (2005). Understanding the Experience of Interactive Art : Iamascope in Beta_space. *Symbolic Interaction*, pages 49–56.
- Cui, Y. and Stricker, D. (2011). 3D Shape Scanning with a Kinect. In *ACM SIGGRAPH 2011 Posters*, SIGGRAPH '11, pages 57:1—57:1, New York, New York, USA. ACM.
- Eisler, C. (2012). Kinect for Windows is now Available! Website. <http://blogs.msdn.com/b/kinectforwindows/archive/2012/01/31/kinect-for-windows-is-now-available.aspx> [Accessed on April 13th 2012].
- Fritz, D., Zinman, A., Elliott, G., and Shilkrot, R. (2012). DepthJS @ GitHub. Website. <https://github.com/doug/depthjs> [Accessed on April 16th 2012].
- Fu, F.-L., Su, R.-C., and Yu, S.-C. (2009). EGameFlow: A scale to measure learners' enjoyment of e-learning games. *Computers & Education*, 52(1):101–112.
- Gallo, L., Placitelli, A. P., and Ciampi, M. (2011). Controller-free exploration of medical image data: Experiencing the Kinect. In *Computer-Based Medical Systems (CBMS), 2011 24th International Symposium on*, pages 1–6.
- Global Illumination (2012). Velkommen til Global Illumination. Website. <http://www.globalillumination.no/> [Accessed on June 2nd].

- He, G.-f., Kang, S.-k., Song, W.-c., and Jung, S.-t. (2011). Real-time Gesture Recognition using 3D Depth Camera. In *2011 IEEE 2nd International Conference on Software Engineering and Service Science (ICSESS)*, ASSETS '11, pages 187–190.
- Hendel, Z. (2012). Next generation API for OpenNI - we want your feedback! Website. https://groups.google.com/forum/?fromgroups#!topic/openni-dev/paLBzL_v7zg [Accessed on May 30th].
- José, R. and Cardoso, J. C. S. (2011). Opportunities and Challenges of Interactive Public Displays as an Advertising Medium. In Müller, J., Alt, F., and Michelis, D., editors, *Pervasive Advertising*, Human-Computer Interaction Series, pages 139–157. Springer London, London.
- Kinect for Windows Team (2012). Kinect for Windows: SDK and Runtime version 1.5 Released. Website. <http://blogs.msdn.com/b/kinectforwindows/archive/2012/05/21/kinect-for-windows-runtime-and-sdk-version-1-5-released.aspx> [Accessed on May 20th].
- Knies, R. (2011). Academics, Enthusiasts to Get Kinect SDK. Website. <http://research.microsoft.com/en-us/news/features/kinectforwindowssdk-022111.aspx> [Accessed on April 13th 2012].
- Malone, T. W. (1981). Toward a theory of intrinsically motivating instruction. *Cognitive Science*, 5(4):333 – 369.
- Michelis, D. (2008). *Interaktive Großbildschirme im öffentlichen Raum*. PhD thesis, University of St.Gallen.
- Michelis, D. (2011). Magical Mirrors 2.0. Website. <http://magicalmirrors2006.wordpress.com/2011/10/11/magical-mirrors-2-0/> [Accessed on May 16th 2012].
- Michelis, D. and Müller, J. (2011). 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.
- Microsoft (2012a). Kinect New Features. Website. <http://www.microsoft.com/en-us/kinectforwindows/develop/new.aspx> [Accessed on April 16th 2012].
- Microsoft (2012b). Natural User Interface for Kinect for Windows. Website. <http://msdn.microsoft.com/en-us/library/hh855352> [Accessed on May 30th].

- Microsoft (2012c). Xbox 360 + Kinect - Games. Website. <http://www.xbox.com/en-US/kinect/games> [Accessed on April 10th 2012].
- Microsoft (2012d). XNA Game Studio 4.0. Website. [http://msdn.microsoft.com/en-us/library/bb200104\(v=xnagamestudio.40\).aspx](http://msdn.microsoft.com/en-us/library/bb200104(v=xnagamestudio.40).aspx) [Accessed on May 31st].
- Miller, T. and Johnson, D. (2010). *XNA Game Studio 4.0 Programming*. Addison-Wesley, Pearson Education, first edition.
- Morrison, A., Jacucci, G., and Peltonen, P. (2008). CityWall: Limitations of a Multi-Touch Environment. In *Proceedings of Workshop on designing multi-touch interaction techniques for coupled public and private displays (PPD'08)*, volume 45 of *AVI 2008 (The International Working Conference on Advanced Visual Interfaces)*, pages 31–35, Naples, Italy.
- Müller, J., Alt, F., Schmidt, A., and Michelis, D. (2010a). Requirements and Design Space for Interactive Public Displays. In *Proceedings of the international conference on Multimedia*, number Figure 1 in MM '10, pages 1285–1294, New York, New York, USA. ACM.
- Müller, J., Michelis, D., and Kray, C. (2010b). Evaluating Natural User Interfaces for Public Displays in the Wild (Position Paper). In *Workshop Natural User Interfaces, CHI 2010*.
- Nielsen, J. (2010). Kinect Gestural UI: First Impressions. Website. <http://www.useit.com/alertbox/kinect-gesture-ux.html> [Accessed on May 31st].
- Norman, D. and Nielsen, J. (2010). Gestural Interfaces: A Step Backwards In Usability. *interactions*, 17(5):46–49.
- Norman, D. A. (2010). Natural user interfaces are not natural. *interactions*, 17(3):6–10.
- Oikonomidis, I., Kyriazis, N., and Argyros, A. (2011). Efficient model-based 3D tracking of hand articulations using Kinect. In Hoey, J., McKenna, S., and Trucco, E., editors, *Proceedings of the British Machine Vision Conference 2011*, pages 101.1–101.11. British Machine Vision Association.
- OpenKinect (2012). OpenKinect Main Page. Website. http://openkinect.org/wiki/Main_Page [Accessed on May 29th].
- OpenNI (2011). About OpenNI. Website. <http://www.openni.org/About.aspx> [Accessed on March 27th 2012].

- OpenNI (2012). OpenNI Programmer Guide. Website. <http://openni.org/Documentation/ProgrammerGuide.html> [Accessed on March 27th 2012].
- Peltonen, P., Kurvinen, E., Salovaara, A., Jacucci, G., Ilmonen, T., Evans, J., Oulasvirta, A., and Saarikko, P. (2008). It's Mine, Don't Touch!: interactions at a large multi-touch display in a city centre. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, CHI '08*, pages 1285—1294, New York, New York, USA. ACM.
- Petersen, N. and Stricker, D. (2009). Continuous natural user interface: Reducing the gap between real and digital world. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*, pages 23–26. IEEE.
- Pheatt, C. and McMullen, J. (2012). Programming for the xbox kinect™ sensor. *Journal of Computing Sciences in Colleges*, 27(5):140–141.
- Physics, F. (2010). Farseer Physics Engine. Website. <http://farseerphysics.codeplex.com/> [Accessed on May 16th 2012].
- PrimeSense (2011). NITE: Making the magic happen. Website. <http://primesense.com/technology/nite3> [Accessed on April 13th 2012].
- Rheiner, M. (2012). Introduction to Simple-OpenNI. Website. <http://code.google.com/p/simple-openni/> [Accessed on April 13th 2012].
- Rovio (2012). Angry Birds. Website. <http://www.rovio.com/en/our-work/games/view/1/angry-birds/> [Accessed on May 16th 2012].
- Russell, D. M., Trimble, J. P., and Dieberger, A. (2004). The Use Patterns of Large, Interactive Display Surfaces: Case Studies of Media Design and Use for BlueBoard and MERBoard. In *Proc HICSS 2004. IEEE*, volume 00, pages 1–10.
- Rydén, F., Chizeck, H. J., Kosari, S. N., King, H., and Hannaford, B. (2011). Using Kinect and a Haptic Interface for Implementation of Real-Time Virtual Fixtures. In *2nd Workshop on RGB-D: Advanced Reasoning with depth cameras (in conjunction with RSS 2011)*.
- Satomi, M. and Sommerer, C. (2007). “game _ of _ life ”: An Interactive art installation using an eye-tracking interface. *Screen*, pages 246–247.
- Schrammel, J., Mattheiss, E., Döbelt, S., Paletta, L., Almer, A., and Tscheligi, M. (2011). Attentional Behavior of Users on the Move Towards Pervasive Advertising Media. In Müller, J., Alt, F., and Michelis, D., editors, *Pervasive*

- Advertising*, Human-Computer Interaction Series, pages 287–307, London. Springer London.
- seeper (2010). XBox Kinect. Website. <http://seeper.com/xbox-kinect-2/> [Accessed on May 16th 2012].
- seeper (2011). WIRED Event. Website. <http://seeper.com/wired-event/> [Accessed on May 16th 2012].
- Sharp, H., Rogers, Y., and Preece, J. (2007). *Interaction Design: beyond human-computer interaction*. John Wiley & Sons, 2nd edition.
- Shneiderman, B. and Plaisant, C. (2005). *Designing the User Interface*. Pearson Education, 4th edition.
- Spyridakis, J. H. and Wenger, M. J. (1989). The Relevance of Reliability and Validity to Usability Testing. *IEEE Transactions on Professional Communication*, 32(4):265–271.
- Sterling, B. (2011). Augmented Reality: Kinect fitting-room for TopShop, Moscow – Wired.com. Website. http://www.wired.com/beyond_the_beyond/2011/05/augmented-reality-kinect-fitting-room-for-topshop-moscow/ [Accessed on May 3rd 2012].
- Stone, E. E. and Skubic, M. (2011). Evaluation of an Inexpensive Depth Camera for Passive In-Home Fall Risk Assessment. In *2011 5th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, pages 71–77.
- Stowers, J., Hayes, M., and Bainbridge-Smith, A. (2011). Altitude control of a quadrotor helicopter using depth map from Microsoft Kinect sensor. In *2011 IEEE International Conference on Mechatronics (ICM)*, pages 358–362.
- Strömberg, H., Väättänen, A., and Rätty, V.-P. (2002). A group game played in interactive virtual space. In *Proceedings of the 4th conference on Designing interactive systems: processes, practices, methods, and techniques, DIS '02*, pages 56–63, New York, New York, USA. ACM Press.
- Tipping, L. (2009). Volvo makes waves at the cinema. Website. <http://www.prweek.com/uk/research/895241/Volvo-makes-waves-cinema/?DCMP=ILC-SEARCH> [Accessed on May 29th].

- Tölgyessy, M. and Hubinský, P. (2011). The Kinect Sensor in Robotics Education. In Stelzer, R. and Jafarmadar, K., editors, *Proceedings of 2nd International Conference on Robotics in Education (RiE 2011)*, pages 143–146. INNOC - Austrian Society for Innovative Computer Sciences.
- Trifonova, A., Jaccheri, L., and Bergaust, K. (2008). Software engineering issues in interactive installation art. *Int. J. Arts and Technology*, 1(1):43–65.
- van Dam, A. (1997). Post-WIMP user interfaces. *Communications of the ACM*, 40(2):63–67.
- van Solingen, R. and Berghout, E. (1999). *The Goal/Question/Metric Method: a practical guide for quality improvement of software development*. McGraw-Hill, Cambridge.
- Villaroman, N., Rowe, D., and Swan, B. (2011). Teaching Natural User Interaction Using OpenNI and the Microsoft Kinect Sensor. In *Proceedings of the 2011 conference on Information technology education*, pages 227–232. ACM.
- Wang, A. I. and Føllesdal, E. A. (2010). Evaluation of a Social Multiplayer Game Featuring Multimodal Interaction. In *lasted International Conference on Software Engineering and Applications*, SEA 2010, Marina Del Rey, CA, USA.
- Zigfu (2012a). Zigfu – Kinect Development in HTML, Unity3D and Flash. Website. <http://zigfu.com/en/> [Accessed on April 13th 2012].
- Zigfu (2012b). Zigfu Browser Plugin. Website. <http://zigfu.com/en/downloads/browserplugin/> [Accessed on May 29th].
- Zinman, A., Elliott, G., Fritz, D., and Shilkrot, R. (2010). DepthJS @ Fluid Interfaces Group – MIT Media Lab. Website. <http://fluid.media.mit.edu/people/roy/exploration/depthjs.html> [Accessed on April 16th 2012].

— A —

Questionnaires

Questionnaire for non-participants

* Required

Personal information

Gender *

- Male
- Female

Age *

How many hours per week on average do you play games? *

- 0 hours
- 1-4 hours
- 5-9 hours
- 10-19 hours
- 20-30 hours
- More than 30 hours

Which platforms do you play on?

- Mobile phone/smartphone
- Tablet
- Web
- Console
- PC
- Mobile console

What is your favorite game genre?

- Adventure games
- FPS
- Strategy games
- RPG
- MMORPG
- Racing games
- Sports games
- Puzzle games (Sudoku, Wordfeud and so on)
- Fighting games
- Simulation
- None

None

Other:

What is your favorite game?

About the campaign

I felt comfortable with becoming a part of the game without giving explicit consent.

1 2 3 4 5

Strongly disagree Strongly agree

I was comfortable with playing the game in a public space.

1 2 3 4 5

Strongly disagree Strongly agree

I would have been more comfortable with playing with a controller (mobile phone, Xbox controller).

1 2 3 4 5

Strongly disagree Strongly agree

The game/campaign caught my attention to a greater degree than a poster would have done.

1 2 3 4 5

Strongly disagree Strongly agree

Other comments

Submit

Powered by [Google Docs](#)

[Report Abuse](#) - [Terms of Service](#) - [Additional Terms](#)

Questionnaire for participants without a campaign

* Required

Personal information

Gender *

- Male
- Female

Age *

How many hours per week on average do you play games? *

- 0 hours
- 1-4 hours
- 5-9 hours
- 10-19 hours
- 20-30 hours
- More than 30 hours

Which platforms do you play on?

- Mobile phone/smartphone
- Tablet
- Web
- Console
- PC
- Mobile console

What is your favorite game genre?

- Adventure games
- FPS
- Strategy games
- RPG
- MMORPG
- Racing games
- Sports games
- Puzzle games (Sudoku, Wordfeud and so on)
- Fighting games
- Simulation
- None

None

Other:

What is your favorite game?

About the game

I felt comfortable with becoming a part of the game without giving explicit consent.

1 2 3 4 5

Strongly disagree Strongly agree

I was comfortable with playing the game in a public space.

1 2 3 4 5

Strongly disagree Strongly agree

I would be comfortable with participating in a game like this if it was a part of an advertising campaign.

1 2 3 4 5

Strongly disagree Strongly agree

The game caught my attention to a greater degree than a poster would have done.

1 2 3 4 5

Strongly disagree Strongly agree

I thought getting started with the game was easy.

1 2 3 4 5

Strongly disagree Strongly agree

I thought it was easy to control the game.

1 2 3 4 5

Strongly disagree Strongly agree

I think most people will be able to play the game.

1 2 3 4 5

Strongly disagree Strongly agree

I became unaware of my surroundings while playing the game.

1 2 3 4 5

Strongly disagree Strongly agree

Did you play with anybody else?

Yes

No

If yes: I think the game was made more fun by playing with somebody else.

1 2 3 4 5

Strongly disagree Strongly agree

If no: I think I would have played longer if I played with somebody else.

1 2 3 4 5

Strongly disagree Strongly agree

I would consider playing with strangers.

1 2 3 4 5

Strongly disagree Strongly agree

If your child played the game: I would be comfortable with letting my child participate in a game like this.

1 2 3 4 5

Strongly disagree Strongly agree

If your child played the game and it was not part of an advertising campaign: I would be comfortable with letting my child be a part of an advertising campaign like this.

1 2 3 4 5

Strongly disagree Strongly agree

I wanted to win while playing the game.

1 2 3 4 5

Strongly disagree Strongly agree

Did you capture enough objects to win?

Yes

No

Submit

Powered by [Google Docs](#)

[Report Abuse](#) - [Terms of Service](#) - [Additional Terms](#)

Questionnaire for participants with a campaign

* Required

Personal information

Gender *

- Male
- Female

Age *

How many hours per week on average do you play games? *

- 0 hours
- 1-4 hours
- 5-9 hours
- 10-19 hours
- 20-30 hours
- More than 30 hours

Which platforms do you play on?

- Mobile phone/smartphone
- Tablet
- Web
- Console
- PC
- Mobile console

What is your favorite game genre?

- Adventure games
- FPS
- Strategy games
- RPG
- MMORPG
- Racing games
- Sports games
- Puzzle games (Sudoku, Wordfeud and so on)
- Fighting games
- Simulation
- None

None

Other:

What is your favorite game?

About the campaign

I felt comfortable with becoming a part of the campaign without giving explicit consent.

1 2 3 4 5

Strongly disagree Strongly agree

I was comfortable with playing the game in a public space.

1 2 3 4 5

Strongly disagree Strongly agree

The campaign caught my attention to a greater degree than a poster would have done.

1 2 3 4 5

Strongly disagree Strongly agree

I thought getting started with the game was easy.

1 2 3 4 5

Strongly disagree Strongly agree

I thought it was easy to control the game.

1 2 3 4 5

Strongly disagree Strongly agree

I think most people will be able to play the game.

1 2 3 4 5

Strongly disagree Strongly agree

I became unaware of my surroundings while playing the game.

1 2 3 4 5

Strongly disagree Strongly agree

Did you play with anybody else?

Yes

No

If yes: I think the game was made more fun by playing with somebody else.

1 2 3 4 5

Strongly disagree Strongly agree

If no: I think I would have played longer if I played with somebody else.

1 2 3 4 5

Strongly disagree Strongly agree

I would consider playing with strangers.

1 2 3 4 5

Strongly disagree Strongly agree

If your child played the game: I would be comfortable with letting my child participate in a campaign like this.

1 2 3 4 5

Strongly disagree Strongly agree

I wanted to win while playing the game.

1 2 3 4 5

Strongly disagree Strongly agree

Did you capture enough objects to win?

Yes

No

If you managed to collect enough objects: I wanted to buy the advertised product more than before I played the game.

1 2 3 4 5

Strongly disagree Strongly agree

My willingness to buy a product from the advertiser increased after playing the game.

1 2 3 4 5

Strongly disagree Strongly agree

Submit

Powered by [Google Docs](#)

[Report Abuse](#) - [Terms of Service](#) - [Additional Terms](#)