Institute for Visualization and Interactive Systems

University of Stuttgart
Universitätsstraße 38
D–70569 Stuttgart

Bachelorarbeit Nr. not set

Establishing Public Display Infrastructure to promote user engagement using Avatars

Cagri C. Tasci

Course of Study:	Informatik
Examiner:	Prof. Dr. Albrecht Schmidt
Supervisor:	Romina Kettner M.Sc., Dr. Stefan Schneegaß
Commenced:	November 5, 2016
Completed:	April 27, 2017
CR-Classification:	H.5.2

Abstract

In recent years the deployment of interactive displays in public environments has rapidly increased and thus, public displays are now widely spread, offering a variety of applications. However various studies investigating the user engagement with public displays have found that these are mostly ignored by pedestrians. This corresponds to the phenomenon of ”display & interaction blindness”, identified in previous work. Our hypothesis is that by allowing the user to create and customize an avatar in order to interact with the display, the user establishes a more personal bond with the display application. Hence, in this bachelor thesis we investigated the impact of customisable, virtual avatars to promote user engagement. In a user study involving 20 participants, we explored visual preferences for designing such virtual avatars. With respect to future work, we implemented an infrastructure that, besides allowing the creation and customisation of avatars, visualises the avatars of nearby users and allows interaction with the display through an android application.

Contents

List of Figures

List of Tables

1 Introduction

Recently the deployment of large-screen, interactive, digital displays in public spaces has increased rapidly over the years. Therefore such public displays are becoming increasingly prevalent and ubiquitous in promoting 'everyday information encountering' [CND+04]. However, many studies have found, that user engagement regarding public display applications is generally rather low and public displays are widely ignored by passers-by [CND+04; MWE+09]. Previous work identifies two important phenomena that occur in the context of public display, which lead to the observed lack of user engagement. These phenomena are the so called "*display- & interaction-blindness*". Display blindness refers to users overlooking a public displays due to the high amount of digital content in public spaces, craving a persons attention. Interaction blindness on the other hand, refers to the struggle that users face, in order to tell passive and interactive display applications apart. Based on these identified issues, Kukka et al. [KOK+13] identifies the so called "first click problem", which refers to the phase before users begin the actual interaction with the display. It is when users discover the display, infer its interactivity and are motivated to begin interaction. He furthermore describes that the issue of low user engagement stems from the fact, that many researchers, when designing public display applications, tend to often implicitly assume that users have already overcome this phase, and therefore are not tending to the first click problem accordingly. Thus, one of the major challenge in designing public display applications is to ensure to overcome the first click problem and so, promote user engagement. Previous work identifies the key aspects of designing applications for increased user engagement by active focus on overcoming the described first click problem. Furthermore, it introduces the following key aspects, which have to be actively considered in the design of public display applications and any such application must contain some solution which attends to these aspects in order to overcome the identified problems. [KOK+13; MCA15; PKD14]:

Gaining passers-by attention and draw it towards the display application to avoid being overlooked.

Communicate the application's interactivity so that potential users can infer this interactivity, which generally struggle to tell passive and interactive display applications apart.

Persuade the users to become active by “communicating that the application may posses something of value to the user”. [KOK+13]

Goal of this thesis, was to further investigate the above given criteria in order to establish means of systematically designing public display applications, which overcome the first click problem and therefore promote user engagement. As such method of overcoming the first click problem, we have looked into the use of personalised avatars to motivate users to begin interaction. The work of Du et al. [DLKM16] has shown, that instead of trying to increase the perceived quality of visualised content, the approach of trying to establish an emotional bond between user and display application might be more promising. In fact, their study has shown that when such an emotional bond is evoked, users do tend to like the display application more and thence, this approach is more likely to suffice as means of overcoming specified problems.

Previous work in the field of public displays, such as that of Du et al. [DLKM16], Houben and Weichel [HW13], and Kukka et al. [KOK+13] or Christian and Avery [CA00] have investigated the use of "virtual avatars" to some degree in order to attract potential users. Though the effect of such avatars has not been investigated by the specified studies. Castronova [Cas03] as well as Ducheneaut et al. [DWYW09] suggest that virtual avatars can serve to establish a personal bond between users and virtual worlds. Young et al. [YMQQ07] proposes the use of a virtual avatars to establish such a relationship between user and display application. Furthermore, he notes that 'by allowing users the creation of personalized content users are encouraged to continual usage and emotional investment throughout regular intervals where such customization of ones own content is allowed' [YMQQ07]. However, the actual effect and the correlation between the use of such avatars and user engagement in the context of Public Displays has yet not been investigated thoroughly. Therefore, our goal is to implement a public display application which integrates virtual, customisable avatars and investigate the effect of such avatars on user behaviour and on user engagement.

Structure

This work is structured as follows. Firstly, Chapter 2 depicts previous work that has been made in the same or a related field of work. In Chapter 3, the idea and concept behind our implemented system will be outlined. This is followed by a more detailed depiction of the concrete implementation of both of our applications, as well as the established communication between both in Chapter 4. The procedure and the set-up of our study to explore visual preferences for designing avatars is described in Chapter 5. Later in that chapter, in Section 5.2, the results of our user study will be presented and visualised. The final section of that chapter, Section 5.3, focuses on discussing the results presented in the previous section. Finally, a summary of this work and discussion of the limitations, followed by a prospect of future work is given in Chapter 6.

2 RelatedWork

This chapter serves to introduce and discuss previous work that either investigates a similar or related research question, or of which the results and findings directly affected or influenced this work. Since this work focuses on the intersection of public display applications and virtual avatars, this chapter will include previous work to both fields of work. The first section focuses on illustrating the background and previous findings in regards to public displays and the analysis of user behaviour in such context. This is followed by Section 2.2, in which we will present previous work, that has specialised on overcoming the findings presented in the previous section. Section 2.4 presents previous work, which has investigated different input and interaction methods for public displays. In Section 2.4 we will discuss previous findings regarding virtual avatars and their effect on user engagement. It also focuses on previous studies that have used virtual avatars in the context of public display and have investigated how avatars can or might aid in promoting user engagement in given context.

2.1 User behaviour and public displays

A great variety of previous work in the hci-community has not only researched the interaction of users with public displays, but likewise investigated the factors associated with user engagement in such scenarios. Several different studies that have investigated user behaviour around public displays have had similar observations and thence have come to several conclusion which are quite close to each other [ATK16; HKB08; MCA15; MWE+09; PKD14]. We will take a look at a few of those to offer profound insights into user behaviour in the specified context.

The work of Memarovic et al. [MLA+12] has focused on investigating how public displays can serve to stimulate certain essential human needs in public spaces. More precisely, they have conducted a field study in order to explore and conceptualize processes of stimulating the following needs in public spaces through large screen displays:

(1) passive engagement, where users observe a certain content or behaviour

(2) active engagement through “intellectual challenges posed by the current space”.

(3) discovery, where novel discoveries within the current space result in excitement.

Two installations of displays in public spaces have been used to observe user behaviour in given context. Findings that have resulted from the observations of this study will be reoccurring throughout this section. As for passive engagement, the findings have shown that users generally (passively) interact only very briefly with the public display application. Most pedestrians do glance at the display or read a small portion of the content and then proceed to move on. During the study, they have observed a tiny amount of instances in which, pedestrians have engaged the display in order to explore any possible interactivity. Furthermore, they have noticed that users were actively trying to explore and learn more about the application once they were actively engaged with the display. From these findings, Memarovic et al. have constructed conceptual models to visualise and illuminate the observed stages in engagement and coordination for discovery.

Huang, Koster, and Borchers [HKB08] have conducted a field study examining different, common use practices of large screen displays in public settings. Their work has primarily investigated public display which were intended to provide “ambient or non-urgent content for a variety of purposes”. As such, they have chosen to look at informational, advertising, artistic or entertainment purposes for their field study and have further chose to test their installations in a variety of different public settings. While they have stated that for such public displays it is often assumed, that these are “inherently eye-catching and appealing” to passers-by. However, during their field study they have, amongst other things, observed, what they refer to as “brevity of glances”. Meaning that their findings suggest, that users mostly tend to only very briefly pay attention to such “public ambient information displays”, if at all. Furthermore, they have observed that only a tiny percentage of pedestrians tend to slow down while passing the display to observe it more closely. Equally few participants have actually stopped or changed their walking path in order to examine the displays content. From these results, they conclude that “glancing and attention at large displays is complex and dependent on many factors”, such as the positioning and type of the display, the type of presented content and the targeted audience.

Especially one of their conclusions has direct impact on our work. When designing such public ambient displays, one should design to give the users some degree of control over the content that is presented on the screen.

Mueller et al. [MWE+09] introduce the phenomenon called "*display blindness*", which results from their findings of a user-study investigating the effect of expectations on attention towards "digital signage". This refers to the effect that a users expectations on any digital content might have on his attention towards such digital content that is being displayed on a public display. Display blindness refers to the "disregard of low-priority inputs" stemming from public displays [MWE+09], meaning that passers-by tend to often overlook the content provided by displays in public environments. According to Mueller et al., the cause of this phenomenon is a *information overload* that users are presented with in public environments. This is due to the fact, that contents of public displays have to compete with many other applications and installations, craving passers-by attention. Therefore, when designing public display applications, one must consider certain methodologies in order to get users to notice the display and attract their attention. This phenomenon has been further investigated by Memarovic, Clinch, and Alt [MCA15], resulting in categorized approaches in addressing display blindness.

Based on the previous findings, Kukka et al. [KOK+13] and Parra, Klerkx, and Duval [PKD14] investigate and conceptualise methods of overcoming the identified problem of display blindness. Moreover, these works focus on interactive public displays. The work of Kukka et al. [KOK+13] explores visual signals in order to entice people to interact with public displays. In addition to the findings of Mueller et al. and based on their work, the work of Kukka et al. identifies an active and explicit avoidance of public display applications and content. These findings are the result of investigations on the behaviour of groups of users around public displays and how they socialize around such. Parra, Klerkx, and Duval [PKD14] investigate the same phenomenon, but focus on gaining profound knowledge about user behaviour around public displays and what can be done in order to attract awareness. Both of these works conclude in the finding of a phenomenon which they call "interaction blindness". This identified phenomenon refers to the fact, that passers-by may often struggle to tell passive and interactive display applications apart, leading to a generally low willingness to instantiate interaction with such displays. These previous works identify the key aspects of designing applications for increased user engagement and introduces the following phases in order to overcome the identified problems regarding public display applications [KOK+13][PKD14][MCA15].

(i) Gaining passers-by attention:
As previously stated, overcoming the issue of 'display blindness' [BR03] is critical for public display applications. Passers-by must be made aware of the display and it has to stand out in some shape or form to avoid being overlooked.

(ii) Visualising interactivity:
Potential users must be made aware of the interactivity, that the public display application offers. As described above, there is a great variety of different public display applications, which can either be passive or interactive. Therefore the application must ensure to visualise its interactivity properly.

(iii) Persuade to become active:
Once potential users attention has been gained, the application must persuade them to become active and make use of the applications interactivity. Kukka et al. [KOK+13] notes that an application should “communicate that it may posses something of value to the user”.

The work of Parra et al. introduces a fourth phase to this, which is the following:
(iv) Allow users to reach a certain goal in interaction. Meaning, that the application should aid the user in reaching whatever goal he or she is trying to achieve by interacting with the display.

Brignull and Rogers [BR03] investigate further based on these findings in order to explore methods to entice passers-by to interact with public display applications. Moreover, they also identifies a social affordance, which they call the “honey-pot effect”. This relates to an effect, where groups of people that are already interacting with the public display create a “sociable buzz” and by doing so, gain the attention of even more passers-by and attract additional potential users.

The findings of Brignull and Rogers [BR03], Kukka et al. [KOK+13], and Parra, Klerkx, and Duval [PKD14] show, that in order to promote user engagement in the context of public display, it is vital to consider methodology to overcome these identified problems, which have been coined as the phrase *first-click-problem*.

2.2 Overcoming the first-click-problem

This section will discuss previous work, that has either directly delved into finding methods and concepts to overcome the previously identified first-click-problem or of which, the findings have been considered as relevant in order to do so.

Mueller et al. [MWB+12] have conducted a lab and a field study investigating how passers-by notice interactivity of public displays. During that study, they have made several observations of the users behaviour. Their studies have shown an increased interest of passers-by in interacting with the display, once they have noticed the interactivity. Furthermore, the used visualisation methods of *avatar-like representations* or *mirrored user silhouettes* have proven to successfully communicate the displays interactivity.

However, their studies have shown that though both methods succeeded, *mirrored user silhouettes* are more effective than *avatar-like representations* in doing so. The studies conducted by Ackad, Tomitsch, and Kay [ATK16] and Tomitsch et al. [TAD+14], were quite similar to that of Mueller et al., but in addition to those findings, these studies have brought forth the observation, that the use of a avatar-like representation entices a more playful behaviour.

Du et al. [DLKM16] and Houben and Weichel [HW13] actively investigate methods of overcoming display and interaction blindness. Houben and Weichel explore attracting passers-by through physical objects, so called "curiosity objects". Du et al. [DLKM16] on the other hand, compare two different methods of overcoming interaction blindness.

The compared methods are through *animation and video* and have proven to be effective in visualising interactivity and therefore attracting users.

Christian and Avery [CA00] explore the use of an animated talking head in a public intelligent kiosk application in order to entice passers-by to interact. Their study has shown, that pedestrians are generally *attracted to such an animated face*.

Finke et al. [FTLB08] present the design and deployment of an interactive game on a public display, which provides interaction through mobile phones. A conducted field study has shown, that the designed system actively encouraged passers-by to interact with the system. Therefore, intersecting public display applications with elements of game design, can provide viable methods of overcoming the specified problems.

While most attempts of addressing the display or interaction blindness have focuses on increasing the perceived quality or value of the displayed content, Lee et al. [LCWD14] explore the approach of trying to establish an emotional bond between user and public display application. By doing so, they aim to give users a more meaningful and effective motivation to interact with a public display application. The approach they haven chosen, tries to use the participants emotional attachment to his mobile device in order to establish such an attachment to the public display as well. The conducted study has shown, that they were in fact able to evoke emotions on users towards the public display. These emotions that have been evoked in participants have lead to an increased liking of the display by participants. Furthermore, they state that establishing some sort of an emotional attachment between users and display applications presents a valid step towards a solution to the display and interactivity problem.

Young et al. [YMQQ07] propose the use of personalised, to some degree customisable, virtual avatars in a public display application to entice pedestrians to interact with that application. More precisely, Young et al. deploy a system, which allows users to generate their individual, virtual avatar in form of a virtual fish whose face can be customised through the selection of a photograph of oneself.

The face recognised in the photograph is then morphed onto the virtual fish. However, no study has been conducted in order to evaluate the effect of such avatars on user engagement.

2.3 Interaction methods

The next aspect we want to take a closer look at, is previous work, which has dealt with exploring varieties of interaction methods with public display application. There has been several endeavours to investigate the use of mobile devices in order to interact with public display applications. Examples for such are the works of Clinch et al. [CDKS12], Davies et al. [DFN+09], Holleis et al. [HROS07], Kaviani et al. [KFF+09], and Vajk et al. [VCBE08].

Kaviani et al. [KFF+09] explored methods of designing an interactive display applications for interaction via mobile devices. For their study, they have categorized the audience into three groups of users based on the current stage of interaction. These groups of users are listed below, along with the corresponding stage of interaction: *actors, spectators* and *bystanders*.

(i) *actors*: direct interaction

(ii) *spectators*: focal awareness

(iii) *bystanders* peripheral awareness

But more importantly, they have explored and analysed the use of "*different design strategies on the concepts of interactive dual displays*".

Dix and Sas [DS] explore the potential synergies between mobile devices and public displays and how content can be managed and visualised across both devices. Furthermore, they explore the possibilities of handling the interaction between those, meaning how to allocate both devices along the dimensions of input/display and public/personal. In their work, they provide a comprehensive categorisation of what type of content can be visualised on which device and what purpose each device could possibly fulfil to achieve a meaningful and satisfactory interaction. They identify several possible functionalities that a mobile device can provide, such as functioning as means of identifying users, identifying selected content and the input of custom content such as text.

On another note, there has also been previous work exploring the possibilities of allowing interaction between mobile and public display devices through the use of the bluetooth protocol. Davies et al. [DFN+09] and Le Grange et al. [LLLL10] focus on bluetooth device naming in order to allow the user the selection of content or desired action and so, interaction with the display.

Cheverst et al. [CDF+05] and Vajk et al. [VCBE08] on the other hand, use the bluetooth protocol to transfer data between devices or to communicate the users location to the public display.

2.4 Virtual avatars

In this section, we want to take a look at previous work that has explored the effects of virtual avatars on users. Following the definition of Castronova [Cas03], an avatar is "*the representation of the self in a given physical medium*". Kang and Yang [KY06] provides a comprehensive definition of virtual avatars. He defines a variety of different types, into which, virtual avatars can be categorised. Furthermore, the psychology behind these varying types and the use of avatars in general is illuminated. However, his work focuses on the set of avatars that can be generated in specified websites. Yong [Yon09] does provide a more general distinction between different types of avatars and does also delve into the psychology behind those. In addition to the work of Suler, Yong provides design principles to consider when allowing the use of virtual avatars. A few of those design principles are that the application should provide the ability to customise the avatar and offer the user to express creativity and individuality.

Boberg, Piippo, and Ollila [BPO08] have conducted user studies on avatars in order to propose design consideration for virtual avatars. Their evaluation has lead to the defining of several design considerations regarding the looks,functionality and context of such avatars. Ducheneaut et al. [DWYW09] has conducted a study in order to investigate the effects of avatar customisability and understand a users behaviour during a customisation process. Schwind et al. [SWHK15] developed an avatar customisation system in order to explore visual preferences in creating virtual human avatars. Their user study has led to the observation, that users show a "tendency towards designing feminine and attractive but still credible faces".

3 Concept

As mentioned in Chapter 1, our goal was to implement an infrastructure to investigate the effect of avatars on user engagement. Thence, our system needs to fulfil the following three tasks, stemming from the challenges for public display applications noted by Kukka et al. [KOK+13].

- Gain passers-by attention by visualising their personal avatar
- Clearly visualise the possible interaction to the user through avatars.
- Motivate the user to actually initiate the process of interaction with the display.

According to our hypothesis, the visualisation of a users personal avatar can serve to fullfil these tasks. However, in order to ensure their fulfilment, we need to erect certain requirements for our implemented system. These requirements and a description of how these will aid us in fulfilling the above mentioned tasks will be listed in the following section.

3.1 Requirements

Since our goal is to promote user engagement via virtual avatars, our System must actively gain passers-by attention and entice to interact, with the users being passive until initiation. Therefore, the arguably most important requirement for our application is that all of the implemented visualisation and interaction must be done without requiring any initiating action from the user. Otherwise our application would not suffice to investigate our hypothesis. This has a few important implications for our other requirements, first of which is how our System is structured in general.

3.1.1 System architecture & communication

We aspire to gain passers-by attention by having them recognize their own avatar and therefore drawing their attention towards the display. In order to do so, we need to visualise a users personal avatar on a nearby public display.

Thus, we need to constantly check whether a user is nearby, and if so, load and visualise the according avatar. Considering our first requirement, this means that these background checks need to be performable without any required initiating action by the user. To enable such proximity awareness, our public display application must be made aware of a users presence. This can only be made possible by introducing a counterpart to our public display application, which communicates a users current location to our application. Ergo, our system must consist of two applications, communicating with each other. One being the public display application, constantly checking for users in close proximity. The other being an application close to the user, indicating its location. Furthermore, the two applications must both establish a connection and communicate seamlessly to the user.

3.1.2 Application design

Another requirement stems directly from the second task : clearly visualising the application's interactivity. The visualisation of a users avatar is designed to serve for this purpose, as well as for providing a motivation to interact with our application. However, the visualisation of virtual avatars should be a secondary element of our application, which is designed to draw attention towards the application and its main functionality itself. Furthermore, the section of our application designed to visualise the Avatars must be clearly visible to passers-by in order to catch their attention.

The primary element of our application on the other hand, should serve to give users a main goal for interacting with our application. As such main interaction, we have chosen a user survey. The application therefore allows users to partake of a variety of different surveys and visualises the results on the display. Visualising the poll in combination with the avatar should aid in communicating the application's interactivity.

3.2 System infrastructure

This section serves to portray the high-level idea behind our implemented systems, which is designed to fulfil the requirements discussed in the previous section.

As it is directly inferred from the established requirements, our infrastructure will consist of two applications communicating with each other. Figure 3.1 serves to illustrate the desired process of interaction and the basic functionality of our infrastructure. One of these will provide the desired display interaction possibilities as well as our functionalities of detecting and visualising nearby users. The other application is purely on the client side and serves as a beacon, communicating the users whereabouts and

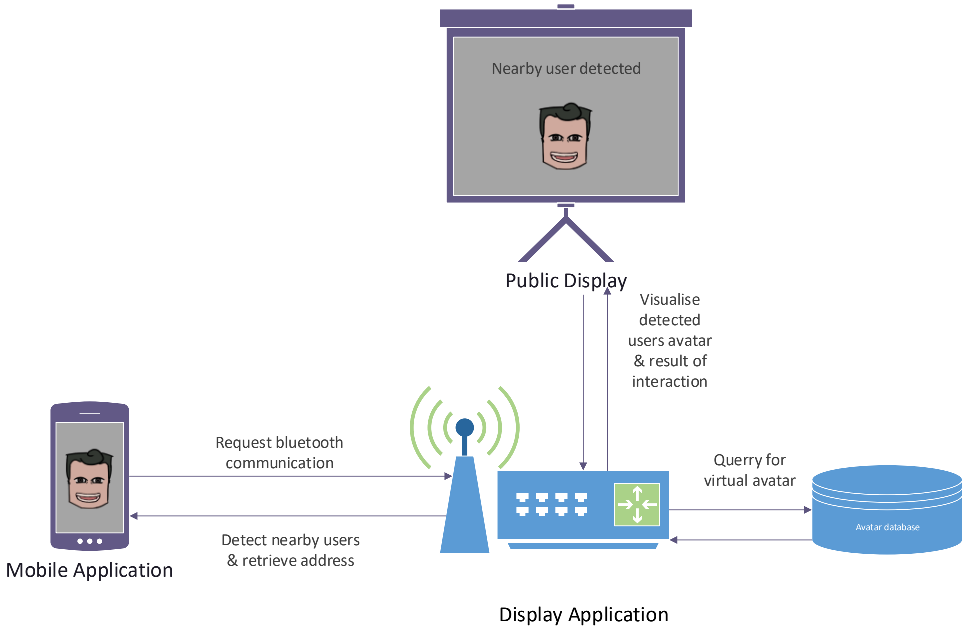

Figure 3.1: Set-up of the desired infrastructure.

proximity to specified public display. The mobile application is the main focus of this thesis, since it allows the user to generate and personalise his or her own virtual avatar. Thus, the mobile application serves to conduct our desired user study to explore visual preferences in designing such virtual avatars. Furthermore, the mobile application provides the basic functionalities required for any desired interaction with the display. The display application must at all times, visualise the avatars of all nearby users. Therefore, as soon as it detects a user who has newly entered the displays proximity, his avatar should be visualised. The visualisation of that avatar should stop, when that user has left a certain distance away from the display. Whenever the user wishes to engage the display and initiate interaction, he should be able to do so from his mobile phone. The client application will provide the required functionality to do so.

4 System

This chapter describes how our system was implemented and illustrates its functionality. Therefore, this chapter is a more in depth and technical description of the implemented infrastructure. Each component is individually described in the following sections. First Section serves to outline the system in general and depicts its setup. The established communication within our system is described in Section 4.2. Section 4.3 focuses on the functionality of the implemented android application. In Section 4.4, the application running on the public displays will be illuminated.

4.1 System outline

In this section, we want to give an outline of the actual implemented infrastructure and provide insights into the functionality of our system. The following description will be a more technical approach to illuminating the system as a whole. As our high-level idea suggests, our system consists of two applications. Both applications were implemented using native android and therefore run on any android device. Possible environments that are able to run these applications are regular android smartphones, android emulators, microcontrollers that can run any android based system image or smart TVs that are based on the android OS. In order to establish a reasonable infrastructure that could be tested during a feasibility test, we have chosen to implement our application based on a chosen use-case-scenario and provided the required functionality for the specified scenario. As such, we have chosen the implementation of a public survey application, that runs on a public display. Users in close proximity of the public display, can choose to participate in that survey and give their opinion on projected questions. In order to test our application fully, we have divided the functionality it offers into two conditions. The first condition provides a simple survey on which users can vote and all answers to currently active question are visualised in form of a bar chart.

The second condition provides the same basic functionality, however, it does not visualise the given answers in form of a bar chart but rather by displaying the avatars of every user, who has voted for the current question. The visualisation of these avatars is split into four segments so that each segment corresponds to one of the four possible answers and contains all the avatars whose users have voted for the corresponding answer. Users can only be assigned to either one of those conditions and once assigned, that assignment wont change unless the user reinstalls the application. The condition to which the user is assigned to, defines which of the two survey he can participate in. If a user decides to partake in one of the survey, he can simply initiate the connection from his mobile application, given that he is within the vicinity of the public display. The connection itself and how it functions is illuminated in the following section, Section 4.2. Functionality of the interaction procedure is fairly straight forward. If the connection request has been successful, the user can use his mobile application to enter his answer for the given question and transfer that answer to the public display application.

While active, the display application does not only visualise each of the survey-questions out of a given set for a specified amount of time, it also does the necessary background scans in order to detect users who have moved into the range of the display. Whence a user is detected, the application requests that users avatar from our online database and proceeds to visualise the retrieved avatar on screen. The visualisation of that avatar lasts as long as the user remains within the vicinity.

4.2 Communication between applications

In order to be able to visualise a users avatars once he is in close proximity, our two applications need to establish some sort of connection with each other. We chose to establish a bluetooth communication, for it is easy to implement and fitting for ranges up to 10 meters. The mobile application serves as a client. Once launched, the mobile application ensures that bluetooth is enabled and sets the device to be discoverable by bluetooth scans. The public display application on the other hand, repeatedly scans for nearby devices in the vicinity Androids bluetooth discovery process is highly resource-intensive process and therefore limits the bandwidth available for connections heavily. Therefore, we wanted to avoid having to manage multiple connections while the discovery process is still underway, which it constantly is. To solve this, the two display application does not connect to a newly discovered device. Instead, for each discovered device in the current scan process, the display application simply retrieves the devices MAC-address and saves these in a list. This list is then handed down to a background service, which iterates over the list. For each element, the service queries for the found MAC-address in the online database, which contains all of our users' avatars. If the queried MAC-address is found in the database, the corresponding avatar can be loaded and displayed on Screen. Devices which were discovered but do not occur in the online database are ignored for the current scan process. This way, a connection is not needed to display nearby users' avatars, saving valuable resources. A bluetooth connection is only established, if the user explicitly wishes to interact with the display. If so, the mobile application sends a connection-request to the display application. The display application opens a server-socket in order to listen for incoming connection requests. This server-socket listens to requests on two separate channels, each one associated with its own unique Id. These Ids correspond to the two conditions of our study to investigate the user engagement. If the socket receives an incoming connection request, both applications can exchange their unique Ids in order to check if this is the desired connection. If so, the connection is established and the applications can begin to communicate. Once communication is possible, user data is transferred from mobile to Display application to identify the users desired action. Based on the Id, which was used to establish the connection, the transferred data is handled differently. More detail on this is given in Section 4.4.2.

4.3 Android application

The android application, as mentioned above, primarily serves for interacting with the public display application. How this interaction is realised on the client side, as well as how this is visualised to the user is detailed in Section 4.3.2. Secondary use of the mobile application is the creation of ones personal avatar. The description of the creation process is split into three sections, one for each method of creation. In these sections, the setup of the creation process is illustrated and described. The first section on the other hand, is dedicated to describing the main activity, which firstly serves as a dashboard, from which the user initiates his desired action.

4.3.1 Main activity

The main activity serves a number of purposes, which will be described in this section. Firstly, it serves as a dashboard and starting point to the user. It visualises the current avatar and all interaction starts from this activity, such as editing the avatar or initiating interaction with a nearby display. The layout of this activity in each of the three forms is depicted in Figure 4.1. The basic layout of this activity, which was used in our study to explore visual preferences for designing avatars is depicted in Figure 4.1b. This layout does not yet include an option to interact with the display, but has later been optimized for our feasibility test. This leads us to the first important purpose that this activity serves to fulfil,which is allocating the user to one of the two conditions of our study investigating user engagement. This is done based on the current time of day, in which the user launches the application for the first time. Each hour, the condition to which new users are assigned to changes. By doing the assignment this way and not by

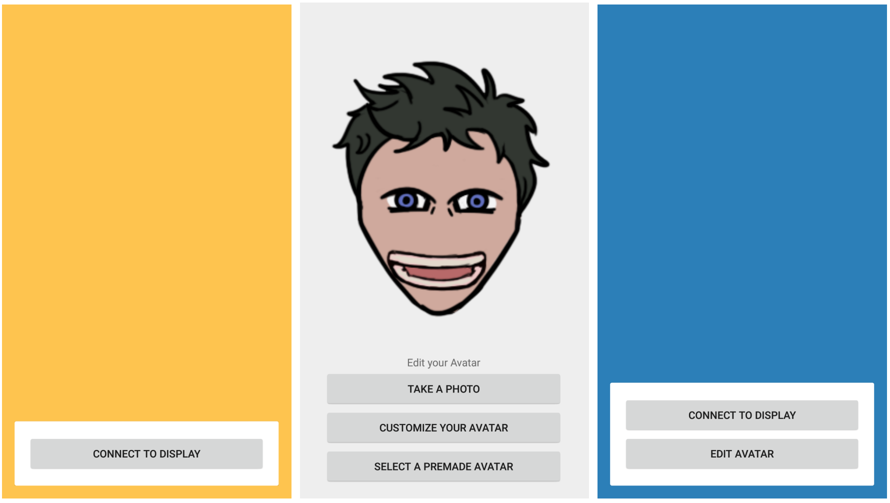

(a) Main menu of the condition without avatars

(b) Main menu during the study

(c) Main menu of the condition with avatars

Figure 4.1: Example images displaying the layout of the main menu

randomising the allocation, we ensure that groups of people that discover the public display application will be assigned to the same condition and therefore have the same user experience. The optimised version introduces to new layout, which correspond to these two conditions. The layout of the main activity changes, based on which condition the user is assigned to. Since our conditions define whether the user has access to the avatar or not, the functionality of the main activity changes as well. If the user is assigned to the condition in which no avatar creation is accessible, the layout changes to that depicted in Figure 4.1a. As can be seen in this figure, there is no option to create or edit an avatar. Merely the interaction button remains and the background colour is changed to clarify the condition. If the user is assigned to the condition in which he has access to the avatar creation, the layout changes to that of Figure 4.1c. Here, the user has access to both functionalities, interacting with the display and creating or editing his avatar. However, the option to edit the avatar has been optimised for ease of use. The three buttons corresponding to the three methods of creation from the previous layout have been dropped and instead, a single button which triggers the creation process is introduced. If the option to edit the current avatar is chosen, a dialogue menu shows up, allowing the user to chose between the 3 creation methods.

Another purpose of the main activity, is setting up some of the prerequisites of the desired bluetooth communication. When launched, a background service is started, which ensures that bluetooth is active on the current device. It then proceeds to set the device to be discoverable by other devices through bluetooth, so that our public display application can discover nearby users and visualise their avatar. To make sure that the device is discoverable, the android application has to ask the user for permission during runtime. More detail on this is presented in Section 4.4. A final addition to the optimised version created for the feasibility test is the introduction of a small demographic survey. This survey is shown once the application is launched for the first time. If the user fills out the form, the main activity saves the data in our online database and wont show that survey again. That is, unless the applications cache is deleted or the user reinstalls the application. Should the user cancel the survey, it will close for the current session, but show up at the next launch of the application.

4.3.2 Display interaction menu

This section describes how the interaction menu is set up and how interaction itself is managed. If the user presses the interaction button, the interaction activity launches and immediately sets up a background service to manage the display interaction.
Figure 4.3a shows the layout of this activity, which, as can be seen in the figure, solely contains 4 radio-buttons to make a selection and a button to confirm the choice. The background service sends a connection request to a specified MAC-address, which is that of the device running our public display application, paired with a unique user Id specific to the assigned condition. This way, we can ensure that the public display application only accepts those connections requested by users and not from any other devices. Furthermore, the Id serves to clearly identify, which condition the user is assigned to, without needing to look it up in the database. This speeds up the interaction process significantly. If no connection can be established, the interaction activity finishes itself and the application returns to the main activity. Otherwise, the connection between both applications is successfully established and so, the user has access to the interaction activity and can begin to input his choice. To do so, the user can select one of the four radio-buttons corresponding to one of the four answers, which will be displayed on the public display. The user can then proceed to confirm his choice with the confirmation-button. This creates a message containing an integer to identify the answer that the user has selected. Afterwards, this message is transmitted to the public display application and if successful, the connection between applications is closed. The current activity then proceeds to close itself and the application returns to the main activity.

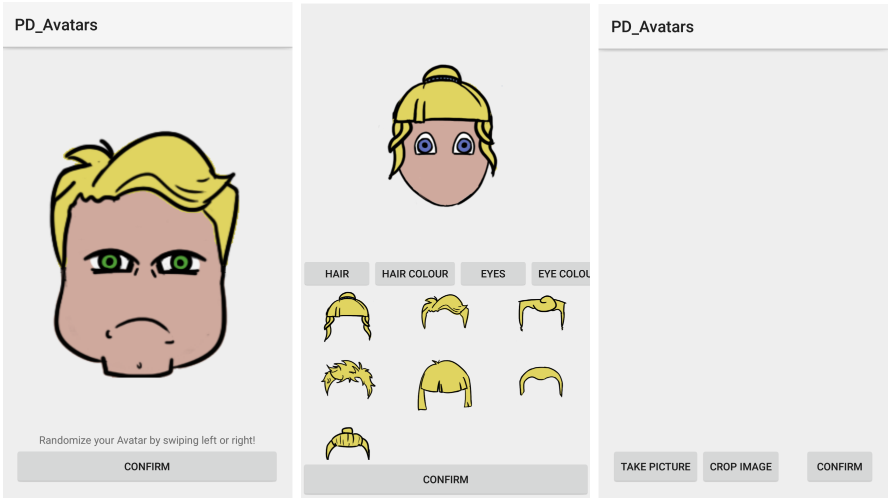

(a) Prefab creation screen **(b)** Full customisation screen **(c)** Photo creation screen

Figure 4.2: Example images displaying each of the three creation process screens.

4.3.3 full character customisation

We will now tend to the first method of creating or editing a virtual avatar, the full customisation method. The full customisation screen allows the user to change each visual aspect of their avatar. This is accomplished, by dividing the creation into a set of visual feature categories, which are the following:

1. Faceshapes
2. Hairstyle
3. Hair colour
4. Eyes
5. Eye colour
6. Mouth

Hereby, the faceshape category allows the user to specify the general head shape and jawline of the created avatar. Hair, eyes and mouth categories simply allow to choose a predefined feature of specified category. The user can then proceed to chose and confirm one of the presented features for the currently selected category. The categories; hair colour and eye colour, allow the user to alter the colour of corresponding feature. However, these categories are set to a default value, which is selected randomly.

Confirmation of the created avatar is only possible, once a feature for each category, except for hair colour and eye colour, is selected. These two categories have been set to be optional, because in order to visualise the hair or the eyes, we do need to fall back to a default colour. Changing the corresponding colour is therefore not mandatory, in case that the user is satisfied with the default colour. How the layout of this screen is set-up, can be seen in Figure 4.2b.

4.3.4 Prefab selection

The prefab method screen serves to allow the user the selection of an avatar with already fully established features. When launched, this screen generates an avatar, by randomly selecting one feature out of every category and merging them. This screen layout contains only a frame to visualise the avatar, and a button to confirm the avatar. Figure 4.2a shows the set-up of this screen. The user can not alter the set of chosen features. He or she can, however, swipe left or right, to generate a new avatar with new, randomly chosen features. Upon confirming, the current features are locked in and the currently generated avatar is set as the users avatar.

4.3.5 Photo avatar

The final method screen allows the user to take a picture of themselves to use as avatar. Figure 4.2c depicts the layout of this activity. For this, this screen allows to launch the devices' camera activity to take an image. Once the photo button is pressed, the application calls the devices camera activity and expects the taken picture as result. If a picture has successfully been taken, the user is returned to the photo method screen, in which the newly taken image is displayed in the centre Additional, the user is given the option to crop and resize the image. By pressing the crop Button, again an external activity is launched, which displays the image and allows the user to select an image area to resize to. The resizeable area can be in the shape of a rectangle or an oval shape, which allows the user to crop the image to only contain the face and nothing else. If the user is satisfied with the image avatar, he or she can chose to confirm their creation.

4.3.6 Confirmation and survey screen

After confirming the current set-up of their avatar, the user is firstly presented with the "confirmation screen". This screens layout is depicted in Figure 4.3b. The purpose of it is to allow the user examine and rethink his or her creation once more, before uploading

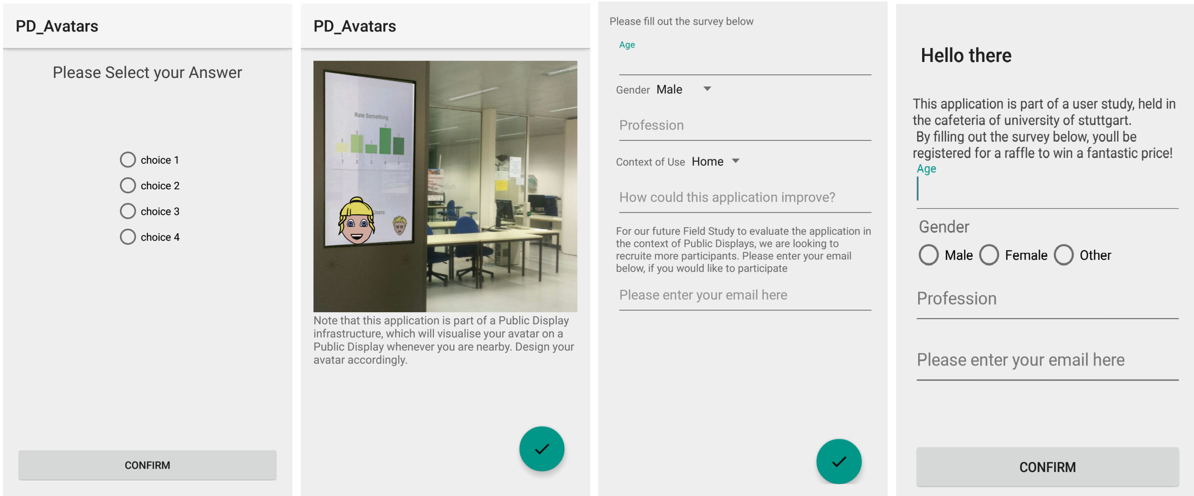

(a) Interaction screen **(b)** Confirmation screen **(c)** Survey screen **(d)** Optimised survey

Figure 4.3: Layout of the interaction-, confirmation- and the survey screen.

it to the database. Additional, it serves to communicate more clearly, that the users creation will be publicly visible. Thus, the Image View in the centre contains a picture of an exemplary public display, onto which the users avatar is overlayed. This serves to visualise, how the avatar might look like "in action", when its projected on a display in an public environment. The user is therefore presented with the opportunity to reiterate the creation, if he or she so wishes. Below the image, a text is presented, which serves as a final briefing. This text remarks that the created avatar will be displayed publicly and that this fact should be considered in the design of the avatar.

If the user chooses to confirm the current creation, one final screen is launched, the survey screen. As the name implies, this screen contains a user form and is depicted in Figure 4.3c. Purpose of this form, is to gather demographic data of our users which will be relevant in the evaluation, which is described in Chapter 5. It furthermore contains a field to select the current location of the user, which again will be relevant for the evaluation. Additional, we have included a field, in which the users can input suggested improvements for the creation process and the application in general. A few of these suggested improvements will be considered in Section 5.2 on page 43. However, the survey screen has been dropped in the optimised version that has been created for the feasibility test and the smaller demographic survey described in Section 4.3.1 is introduces instead. This improved demographic survey screen is depicted in Figure 4.3d.

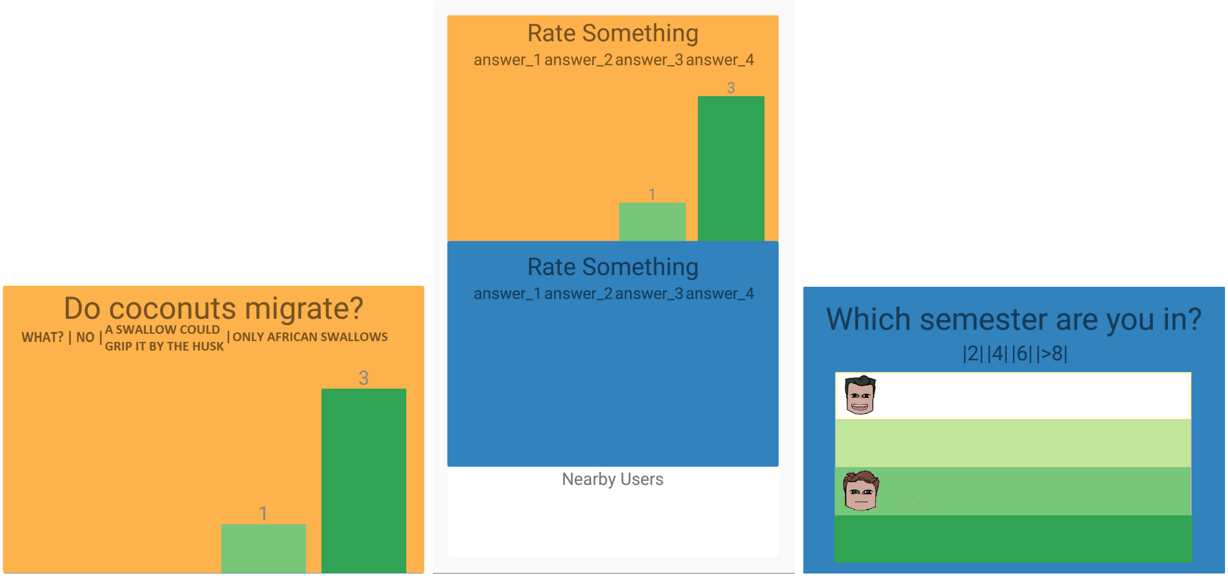

(a) Poll in the no-avatar-condition.

(b) Screenshot of the public display application

(c) Poll in the avatar-condition

Figure 4.4: Example images depicting the public display application.

4.4 Public display application

The public display application mainly serves two purposes. Firstly, it serves to visualise avatars of users in close proximity. The second purpose, which serves to perform the feasibility test, is to enable users to vote and visualise the gathered data.

4.4.1 Main set-up

This section serves to describe how the main activity is set up and which functionality it provides. Firstly, different from the mobile application, this application consists of one activity only, which is the main activity. It serves as a dashboard and provides all the required functionalities, mostly in form of background services. The most important purpose of this activity is to visualise our main interaction; the survey. Since we have two conditions, one with and one without avatars, our application also consists of two separate surveys. For our study, we have chosen to visualise both surveys at the same time on the screen. Therefore, the layout of this application contains two halves, each half corresponding to one of our conditions. The entire layout of our application however consists of two segments. One of these segments contains both survey, while the other segment is solely dedicated to displaying nearby users. The layout of the entire activity is depicted in Figure 4.4b. As can be seen in the figure, both surveys are visualised at the same time and have different background colours. These colours are the same colours used in the mobile application, to clarify which condition the user is assigned to and therefore in which survey the user can partake. How these surveys behave and how voting is managed in detail, will be described further in Section 4.4.2. The segment which deals with the visualisation of nearby users' avatars, will be described in Section 4.4.3.

One important functionality of the main activity is handling the bluetooth communication as a whole. For this, once the application is launched, this activity immediately starts two background services; one to discover nearby bluetooth devices and the other to listen to incoming connection-requests. Both of these were described previously in Section 4.2.

4.4.2 Survey & voting

Here, both surveys and how voting is handled in general will be illuminated. As described in the prior section, the application contains both surveys at the same time. For the survey, we have created a set of questions out of the following categories: politics, economy, sports, entertainment, local news and questions directly addressing the students. When the application is first launched, it randomises the order in which the questions will be presented. The application then cycles through these questions and visualises each one for exactly 3 minutes. If the last question has been reached and visualised for 3 minutes, the application randomises the order once again and starts the cycling from the beginning of the list. Note, that the current question is visualised for both of the surveys, meaning that they are identical at all times. This is to avoid effect on the users willingness to interact by questions in which the user has differing levels of interest.

As mentioned in Section 4.2, the handling of the transferred data differs based on the users assigned condition. We will firstly tend to the condition in which avatars can not be used. As can be seen in Figure 4.4a, this part of the application mainly contains a bar chart, to visualise current votes. Whenever a new question is displayed, the bar chart resets and displays random values for each possible answer. This is done so that potential users are more willing to interact once they get the sense that others have interacted with the application as well. If a user successfully establishes a connection and transmits his or her answer, this interaction is handled in a condition specific thread. Firstly, the user and his given answer is written into the log. This log contains the MAC-address of the users device, the time and duration of the interaction and given answer with the current questions identifier. Afterwards, the bar which corresponds with the users answer is incremented by one and the change in value is visualised.

The part of our application that handles the avatar-condition, behaves quite differently. As can be seen in Figure 4.4c, this part does not contain any chart to visualise past votes. Here, this is done by dividing the area in which votes should be visualised into four segments. Each segment corresponds to one of the four possible answers of each question. These segments are all coloured differently for clarity and serve to display a set of images in a grid. If a user that is assigned to this condition, successfully establishes a connection and transmits his answer, this is again written into the log. However, in addition to the values that the log of the other condition contains, this log also includes the users avatar. Furthermore, when the users answer is read, the condition specific thread queries for the users avatar from the database and stores it. This avatar is then added to that segment of the survey, which corresponds to the users given answer. Therefore, each of the four segments display every avatar, whose users have voted for the associated answer.

The generated log file is only uploaded to our online database, once a whole cycle through the current set of questions has been finished.

4.4.3 Avatar visualisation

The final important segment of our application, is that, which is dedicated to visualising avatars of users who are currently in the vicinity. This segment is located on the bottom of our application and its layout can be inferred from Figure 4.4b. It only contains a simple grid-view, which serves to visualise a set of images. The functionality is much more interesting. To ensure its functionality, the avatar visualisation is handled in a separate dedicated thread. This visualisation thread is started once the application is launched and runs permanently in the background. It serves to manage the detection of nearby users as is described in Section 4.2. After each bluetooth scan process, this thread retrieves the list of currently discovered devices. It then iterates over that list and queries in the online database for the MAC-Address of each device in that list. If the specified device is found in the database, the thread retrieves the users avatar and adds them to a new list of avatars. We will refer to this list as the "avatar-list".

Another function of the specified thread then retrieves that list and generates the required images of each avatar. This is done by loading each feature-image and then merging them all into one bitmap. Therefore, our list now consist of a set of bitmaps and can simply be handed over to a image-adapter, which handles the visualisation of images in the specified grid-view. The image-adapter retrieves the list, overwrites its current list of visualised images and visualises all images of the newly retrieved list in the grid-view.

Another important functionality, which is provided by this thread, is the handling of users who move out of range or who are simply not discovered in the current scan process. Note that a bluetooth discovery scan is not guaranteed to find all devices in the vicinity. Meaning that, devices which were discovered in the last scan process might not be discovered in the current, even though the device is still within close proximity. Since we do not want to loose the visualisation of a users avatar even though the user is still nearby, we have to handle the case of his device not being discovered in the current scan process. To do so, we have given each avatar an additional attribute, which serves to identify the number of times, in which a already discovered user, has not been discovered by consecutive scans. When a users avatar is newly discovered, this attribute is set to 0. Each time that a user is discovered , whose avatar is already stored in the avatar-list, this attribute is once again set to 0. However, each time that a already discovered user, meaning a user whose avatar is stored in the avatar-list, is not discovered by the current scan process, the identifier attribute is incremented by one. After each scan process, we iterate over the avatar-list and for each avatar whose identifier attribute is greater than 3, we remove that avatar from the list. Therefore, once we hand this list over to the image-adapter, the removed avatar will no longer be visualised on screen. This way, only avatars of users who have not been discovered in 3 consecutive scans and it therefore is safe to assume that they have left the vicinity, will be removed from the visualisation.

4.5 Feasibility test

In order to test our implemented infrastructure in a use-case-scenario, we have conducted a feasibility test. This section serves to describe the set-up, procedure and the outcome of the conducted test. As mentioned previously, we have chosen the implementation of a public survey application, that runs on a public display. Users in close proximity of the public display, can choose to participate in that survey and give their opinion on projected questions.

4.5.1 Deployment

The feasibility test was conducted in the semi-public environment, in a lab-room of the university. The application was projected on a large-screen display and was deployed for five days. In order to project the application onto the display, we installed our application on a raspberry pi 3 running the latest system image of “Android Things” [1]. Using the specified system image, provides the running android application with the required access to wifi and to the bluetooth functionality of the raspberry pi. Furthermore, it does

[1] https://developer.android.com/things/preview/download.html

provide a functioning bluetooth adapter, which, unlike the adapter provided by a regular android system image, does not return a null value and can therefore be accessed and used by android applications, just as it can on any android based smartphone device.

4.5.2 Results of the feasibility test

During the time that our application was deployed, we have registered a total amount of 10 occurrences, in which users have successfully interacted with the display application. However, these interactions are not all individual participants. In total, 6 individual participants have interacted with the display. Two of which have interacted twice during the entire interval of deployment. One of those 6 participants has even interacted three times. A remarkable fact is, that all of those participants who have chosen to interact more than once, are participants that have been assigned to the condition with access to avatars. All of the participants that were assigned to the avatar-less condition, did not interact with the display again. This could infer, that our hypothesis is right, and avatars do aid in promoting user engagement in the context of public displays. However, since this was only a feasibility test and not an in-depth study, these results might not be representative.

The infrastructure has proven its feasibility during conducted testing phase. No crashes have occurred during our test phase and interaction seems to have functioned without problems. However, when users that were assigned to the avatar-condition transmitted their vote, the application did require up to 7 seconds at times to visualise the vote with the corresponding avatar on screen. While this is an acceptable amount of delay, it does offer room for improvements. We have also noticed, that the bluetooth discovery process did rarely happen to not find the same user three times in a row and therefore stopping the visualisation of a users avatar even though that user was still within the vicinity. This has scarcely ever occurred but is unfortunately unavoidable without increasing the number of times, which a previously discovered but currently missing user is treated as a simple miss. Yet, this does obviously increase the time in which a avatar will still be visualised on screen despite the fact, that the user has already left the vicinity.

5 Evaluation: exploring visual preferences for designing avatars in a user-study

This chapter focuses on the entire evaluation process,that served to investigate visual requirements for virtual avatars and to explore visual preferences for designing such avatars. Firstly, the conducted study designed for this goal will be described. In the section after that, we will present and visualise results that have been collected during the conducted study. Afterwards, the results will be interpreted and the findings discussed in order to summarise desired requirements and preferences.

5.1 User study

Since our public display infrastructure is designed to use avatars to promote user engagement, we need to firstly investigate design aspects of avatars. Our user study served to investigate the visual criteria an avatar should fulfil, in order to appeal to users, and therefore, possibly promote user engagement. Furthermore, which visual aspects of avatars the users care the most and which aspects the least about. The results of this study will aid in providing certain design possibilities of virtual avatars, for which our hypothesis could hold. More specifically, we wanted to investigate the following criteria:

Method of Creation: Which method of creating an Avatar do users prefer? Do Users want to create an Avatar by choosing each facial feature individually or do they rather choose a pre-defined avatar that suits them the most.

General Representation: Closely tied to first bullet-point: do users prefer an abstract, cartoon-like avatar or having an image of their own face as representation of themselves.

Facial Features: Which facial features do users care the most about? Which can or should be neglected, or are all equally important.

Variety of facial features: How much variety should be given for each of the facial features? For which features is a great variety important and which features can or should have less.

Context of creation: How does the location or the context in which users create and interact with their avatar affect their general behaviour? Does the public visualisation of their avatar affect their behaviour and the level of personalisation of their avatar?

5.1.1 Study design and measures

As independent variable, we have chosen the method of creation: full customisation, prefab and photo avatar. These do serve as our conditions for the user study. During the creation process we record a variety of factors, which also serve as dependent variables, in order to analyse a users behaviour. First of all, we log the total time that a user has spent from start to finish for designing his or her avatar. Hereby, which features the user has chosen for his final avatar will also be transcribed. From this data, we are able to gather the distribution frequency of the selected avatar features over all participants as well. Furthermore, we have noted the time in each method of creation as well as which method the user chose, so that we establish a total frequency of distribution for the chosen methods over all participants. Once a method of creation was selected, the measures vary depending on the method. In the following listing, the measures will be listed in their associated method.

photo Avatar:
For the creation of the avatar via photo, we note the number of attempts until the user is satisfied with his image.

custom Avatar:
During the full customisation of the avatar, we log the total amount of clicks made. Each switch between sets of features, such as head, hair, mouth or eyes, is transcribed. For each feature-set, the time spend and the number of clicks performed choosing a feature from current set is logged. On every click, the selected item is also noted.

prefab Avatar:
In the selection of premade avatars, the total number of swipes the user has performed is logged.

As described in Chapter 3, once the user has created his or her avatar, the user enters a screen in which he is confronted with an exemplary image of his avatar being displayed on a public display. For this screen, we log the amount of time that the user spends in this screen until confirming. After confirmation, the user is presented a survey. This serves to collect demographic data of our users such as age, gender and profession. Furthermore, we ask for the users current location as in at home, at work or away, to identify any implications of the given context on the creation process.

5.1.2 Procedure and participants

Our user study was conducted by having participants create an individual virtual avatar using our implemented android application. The actual study is running in the background, transcribing the users behaviour during the procedure of creating an avatar. What kind of data is gathered, was listed in the previous section. Once the participant finishes the creation process, a small survey containing demographic questions is presented to the user. Content of this survey is as it was described and depicted in Section 4.3.6.

Participants for our study were either directly recruited or downloaded the application from the Google Play Store [1] by themselves, without any recruiting. The sample of recruited participants was drawn from students and employees from our university. We have recruited 20 participants in total to conduct the pre-study with. From the 20 Participants, 14 of such were Male and 6 participants were Female. The mean age of all our participants is 23 ($SD = 2,28$).

5.2 Results

This section tends the visualisation of data that has been gathered during our study to investigate the visual requirements for avatars. It serves, to structure the collected data based on the previously mentioned criteria to investigate. The collected data will not only be listed, but also visualised to allow clear references to the data during our discussion in the following chapter.

5.2.1 Completion time per method

Listed in this section, are creation and confirmation time. Hereby, creation time refers to the time in seconds it took the user to create their avatar, once they had chosen the corresponding button. Confirmation time on the other hand, refers to the time from

[1] https://play.google.com/store/apps?hl=de

when the confirmation screen is launched, until the user presses the confirm button. Therefore, the confirmation time refers to the total time a user has spent examining the mock-up image and possible reconsidering his avatar. These values are listed below and are separated into three tables, one for each creation method. Section 5.2 lists these gathered values, separated by the corresponding method. The gathered data is visualised as a box-plot in Section 5.2.1. This figure visualises the the time required to create, as well as the time required to confirm avatars. The boxplot serves to compare the average time per method, that users have spend to either create or confirm their avatar in respecting method.

We have exluded the data of the photo method from this figure, since both the creation and confirmation time for this method is overall too low to visualise it reasonably with the other plots. The number of samples gathered for this method are also too few to allow a meaningful comparison.

Prefab method		
Id	Create time	Confirm time
1	16.69	1.32
5	0.84	0.79
16	26.32	6.36
19	34.16	3.14

Table 5.1: Prefab method times

Photo method		
Id	Create time	Confirm time
4	1.96	0.925
12	1.83	23.692
17	2.14	3.862

Table 5.2: Photo method times

Customisation method		
Id	Create time	Confirm time
2	32.34	28.96
3	36.29	1.14
6	46.32	22.56
7	36.22	1.32
8	48.55	2.03
9	83.98	20.81
10	36.56	1.10
11	26.49	111.74
13	52.45	2.67
14	79.19	4.83
15	45.64	1.89
18	15	10.47
20	31.58	3.09

Table 5.3: Full customisation method times

Table 5.4: Total amount of time each participant required to create and to then confirm their avatar, separated by the corresponding method that was chosens.

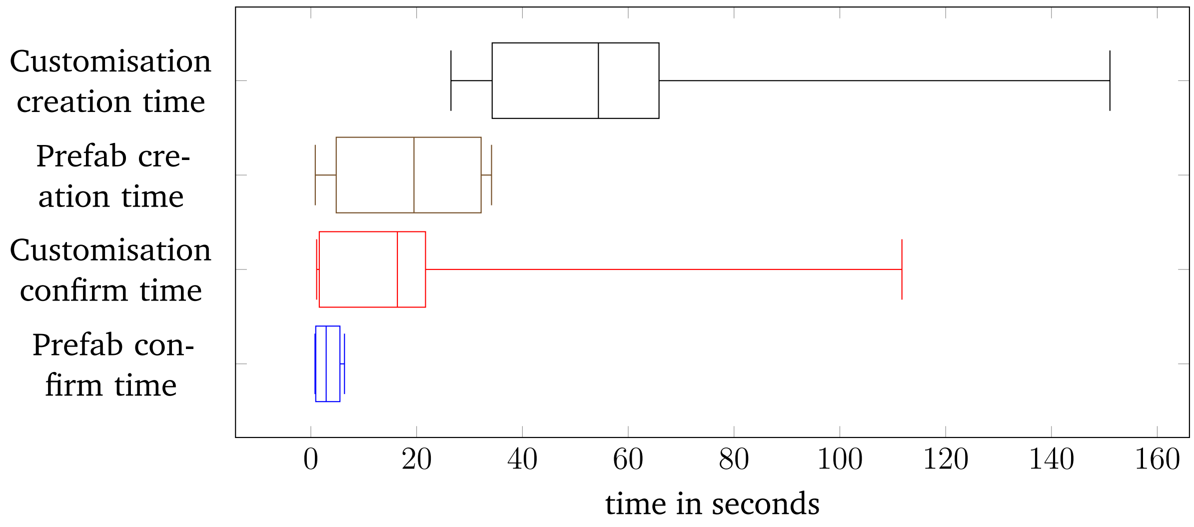

Figure 5.1: Boxplot of the creation and the confirmation times from prefab and customisation method.

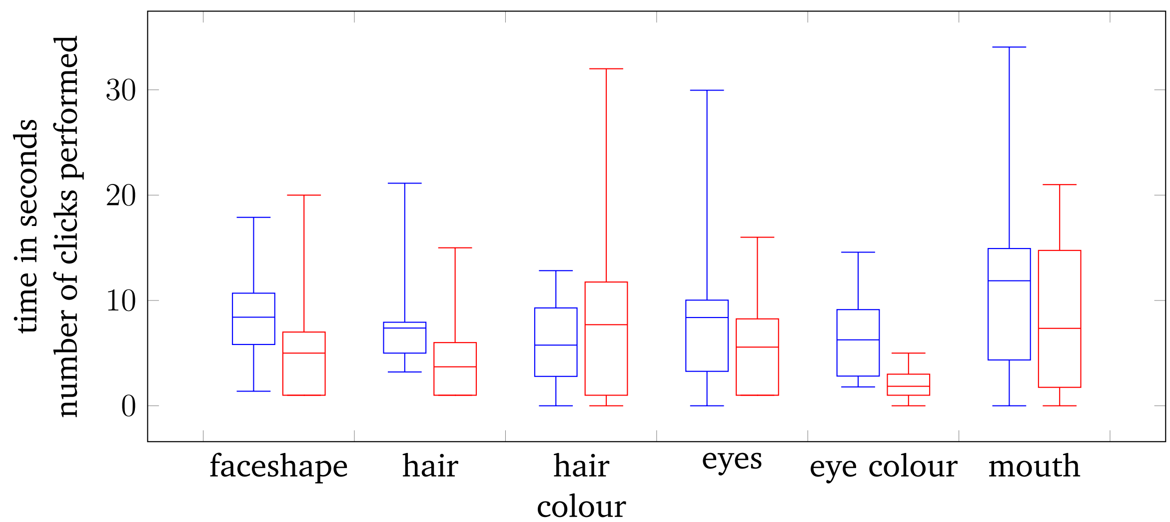

Figure 5.2: Boxplot of the time spend and clicks performed in each category.Blue plots visualise the time, while red plots the clicks performed.

5.2.2 Feature categories

The next aspect we wanted to investigate, is which categories of facial features users tend to care more and which less about. For this, we have gathered both the time required for each category, as well as the total number of clicks, which were performed, while the specified category was active. Table A.2 lists the time in seconds for each category. However, the given values for each participant, are the total amount of time, the specified participant has spent in given category throughout the entire creation process.

Meaning that each time a participant has entered a certain category, the time until he enters another, contributes to the time listed in the table. Even if no click has been performed during that time. Furthermore, values result in a mean time of 8.41 seconds in faceshapes, 7.38 in hair, 5.76 in hair colour, 8.38 in eyes, 6.26 in eye colour and 11.88 the mouths category. These values are visualised in form of a boxplot in Figure 5.2. Table A.3 on the other hand, lists the number of clicks performed in each category. Again values listed in specified table, are the total amount of clicks, that the corresponding participant has performed in each category. Therefore these values are the sum of all clicks performed in the specified category throughout the entire creation process. For the number of clicks in each category we get the following mean values: faceshapes : 5 clicks, hair : 3.1 clicks, eyes : 5.57 clicks and mouths : 7.35 clicks. These values are again visualised in the form of a boxplot in Figure 5.2.

5.2.3 Chosen features

This section serves to list and compare the chosen facial features. The data listed in this section does include the avatars generated by the prefab method, even though the users had limited control over the features of the chosen avatar. This separation between the two methods will however be made in Section 5.3. Table 5.5 contains each chosen feature of the avatar, sorted by participant‘s Id‘s (from top to bottom ; 1 to 20). The corresponding method in which the avatar was generated is listed as well. This data results in the following frequency of distribution, which is visualised in Figure 5.3, showing the number of participants who have chosen each feature. This allows for a more simple and fast comparison of the frequency distribution of selected features for each category.

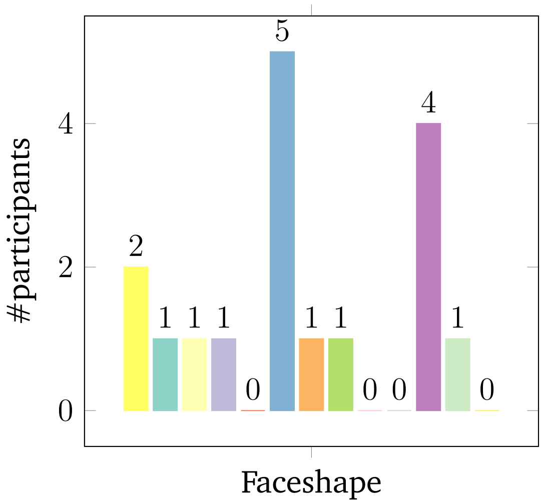

(a) Faceshapes selected by participants

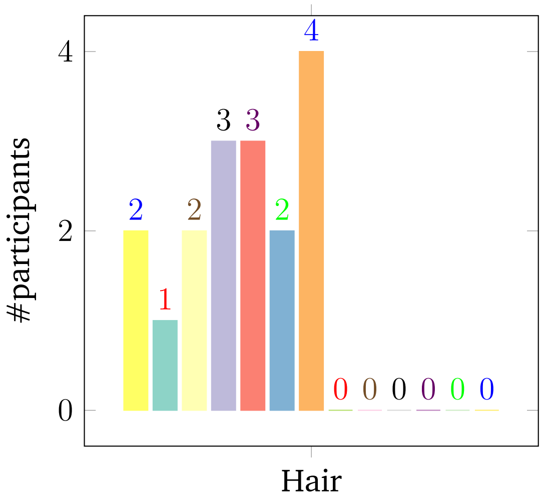

(b) hair selected by participants

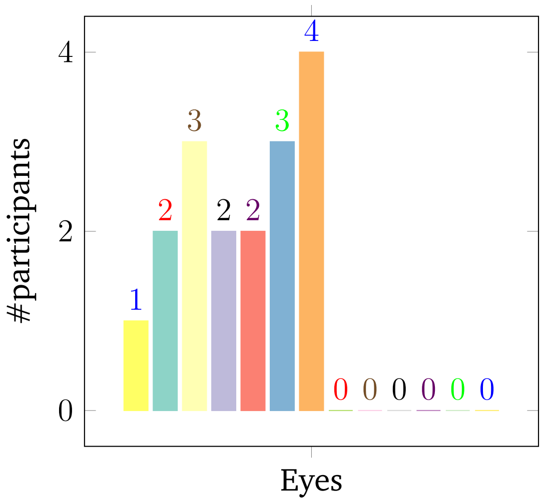

(c) Set of eyes selected by participants

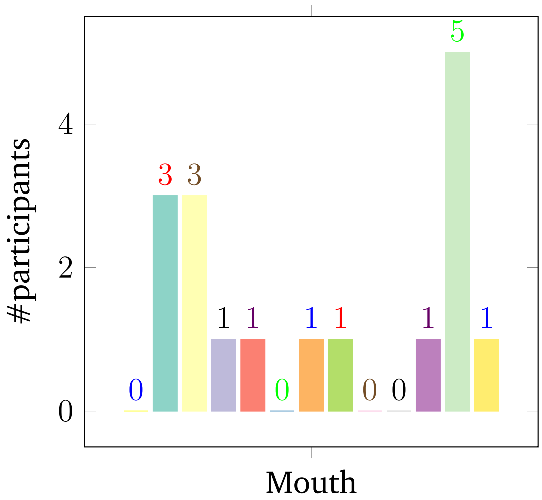

(d) Mouth selected by participants

Figure 5.3: Bar chart visualising for each category, the number of participants choosing each feature. X axis corresponds to the feature Id ranging from 1 to 13 (from left to right)

	Chosen features					
Method	Faceshape	Hair	Hair colour	Eyes	Eye colour	Mouth
Prefab	12	4	black	6	brown	13
Customisation	4	7	black	3	brown	12
Customisation	11	7	blond	7	blue	12
Prefab	1	4	orange	6	brown	3
Customisation	1	3	black	7	green	12
Customisation	6	5	brown	5	brown	3
Customisation	7	5	orange	3	blue	4
Customisation	3	1	brown	1	blue	2
Customisation	11	6	blond	4	blue	7
Customisation	2	1	brown	2	green	11
Customisation	11	7	blond	6	green	5
Customisation	6	5	orange	4	green	3
Customisation	6	2	brown	7	green	12
Prefab	11	3	orange	3	brown	2
Customisation	6	4	brown	7	green	8
Prefab	6	6	black	2	green	2
Customisation	8	7	brown	5	green	12

Table 5.5: Table listing the chosen features for each participant who chose the customisation or the prefab method. The chosen features are listed with their id, represented by a number ranging from 1 to 13, or the name of the colour as String

5.2.4 Context of creation

This section orders the data by the context in which the data was generated. For this, we look into the correlation between the context of creation and the amount of time the participant required to generate and to confirm his or her avatar. Table A.1 lists these values separated by the context of creation. The time is again given in seconds and for both, creation and confirmation. The boxplots in Figure 5.4 serve to visualise the required time and allows comparison between the different required times based on the context of creation.

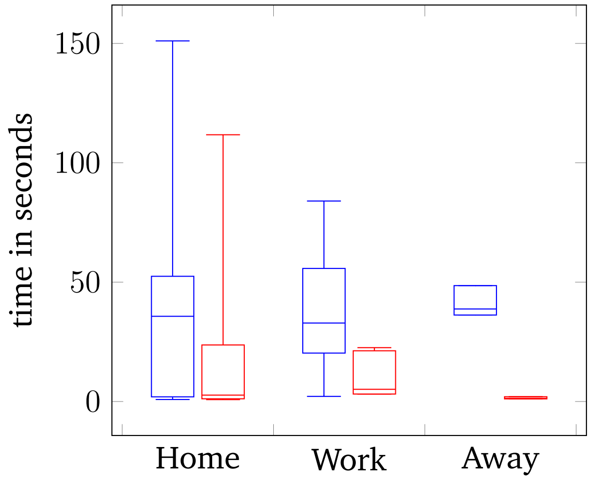

Figure 5.4: Boxplot of the time spend to create and to confirm the avatar, separated by the context of creation. Blue plots visualise the creation times, while red plots visualise the confirmation time

5.2.5 Exemplary created avatars

This section serves to depict a few of the user generated avatars in each of the three methods. We have chosen a small number of avatars out of the entire set of avatars, which were generated by our participants during the study, to be depicted in Figure 5.5. The depicted avatars from the customisation and from the prefab method were selected randomly.

(a) Avatars generated by the customisation method

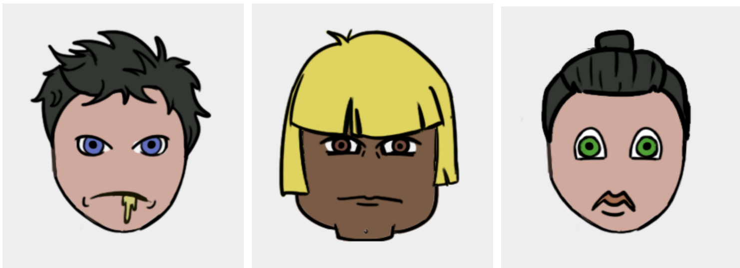

(b) Avatars generated by the prefab method

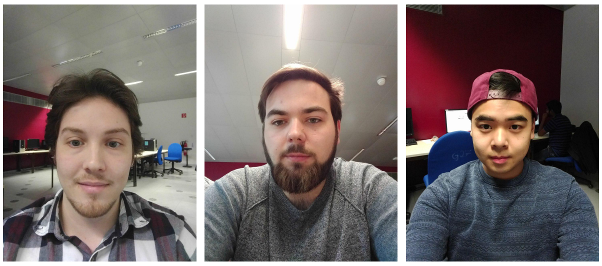

(c) Avatars generated by the photo method

Figure 5.5: Exemplary avatars for each of the three methods

5.3 Discussion

The results concluded in a number of findings, which will be discussed in this section. Structure of this section is loosely based on the order in which the investigated criteria were listed in the previous section.

5.3.1 Method of creation

This section focused on the method of creation, that users have chosen for their avatar and discusses why that decision might have been made. Furthermore, we want to find out, which method is the most reliable in order to investigate our hypothesis. The large majority of participants preferred the full customisation of their avatar, as can be seen in Section 5.2. Only 5 participants selected a randomised avatar, and only 3 participants used a photo of themselves as avatar. When asked about why they chose the selected method, participants replied in the case of the full customisation, that they found this method to be the most exiting, since it allows the user the most freedom over their avatar. Some of these participants also noted that they found the other methods simply to be "boring" (P.18) and "not satisfactory" (P.2). From these results, we can conclude that allowing users the creation and full individualisation of their own avatar yields the highest satisfaction for users. Therefore a fully customisable avatar is best suited in order to promote user engagement. We further conclude, that any system aiming to achieve a similar goal using personalised content, should incorporate full customisability of given content. Regarding the photo avatar, participants noted that they had imagined “something more playful” (P.9) like “an automated avatar generation based on the photo” (P.7). Another Participant has stated, that would have enjoyed the photo method more, “if it allowed to add stickers on top of the image, somewhat like BitMoji does” (P.17). We therefore excluded this method of creation from our final study, since it felt dis-satisfactory to most participants and therefore will most likely not suffice as means of investigating our hypothesis. Same goes for the randomised avatars. 9 of 20 Participant have selected this method at least once during the creation of their Avatar. However, from these 9, only 4 Participants have chosen to confirm their avatar, generated by this method. Those participants who have aborted this method of creation and chosen another, mentioned that, while they do not think the prefab method is "bad" in any way, wanted to have more freedom over personal content, and thus, have chosen the full customisation instead.

We have therefore chosen, to remove the prefab method from the final study as well, since it also would most likely not suffice. However, these methods could be reconsidered and reinvestigated more thoroughly in future work, since our sample is too small to make general assumptions over how these methods would suffice in a fully fleshed out infrastructure.

5.3.2 Feature categories

Now, we want to take a closer look at which categories were of importance to users and which were less so. In doing so, we hope to achieve better understanding of which visual aspects of virtual avatars should be focused on in the design of such an application and which categories should offer a greater variety and more detail in the implemented features. For this, we will take a look at the time and the amount of clicks that the participants have spend or performed in each category. Figure 5.2 visualises in which category the users on average spend the most and in which the least amount of time. Looking at the mean values of these boxplots, we can infer that the categories in which

Figure 5.6: Exemplary avatar, which depicts the most popular features

the users spend the least time in, are hair colour and eye colour. The categories with the most time spend are mouths, followed by eyes and faceshape. A similar trend can be observed in Figure 5.2. Here, the most clicks are performed in the mouth-category and the least in eye colour. From this, we can conclude that special focus should be given in designing these categories, since these values imply that users care more about these categories than others. Furthermore, we conclude that the greater the number and variety of features, the higher the amount of time spend and clicks performed in given category. Participants have given the same reasoning as for why they have spend so less or so much in some categories. Some have suggested, that they tend to spend less time in the colour-categories since they have only a very limited set of selections, which poses no difficult choice to make. Furthermore, it has been implied that choosing hair- or eye-colour is generally a easier choice to make, since most participants tend to "just pick the colour, that their real eyes or hair has" (P.2). Generally, we suggest that the larger the impact of the feature on the avatars general aesthetics, the more carefully users will go on about choosing such feature. For example, the mouth of a face has quite a large impact on the general aesthetics, the mood and emotion that the face conveys. Therefore, users have spend more time and clicks on deciding which feature most closely resembles or conveys what the user has in mind.

5.3.3 Facial features

The next visual aspects we want to focus on, are the individual features, which we provided. More precisely, we want to take a closer look at which of these features were relevant to the users and which were not. For this, we took a look at the amount of participants, who have chosen each individual feature as part of the final set of features of their personal avatar. Figure 5.3 visualises the data, which we want to investigate more closely. What is clearly visible in this figure, is which features have been selected the most, and which the least. In each of the 4 categories, we can see a set of features, which have not been selected by a single participant. We chose to categorise these features as "non-relevant", and have therefore chosen to exclude them from the field study. Reasons as to why they might have been so unpopular, are that most participants have found the specified features to either "not be distinguishable enough from other features" (P.8) or "not really conveying much emotion" (P.16). As for the individual features, who were quite popular amongst our participants; these were kept for the field study. When questioned why the participants chose specified features which were amongst the most popular ones, participants have answered, that they found these to "contain the most detail" or to be "the most elaborate and well designed"(P.5) features. To empathise this, we have put together an exemplary avatar out of the features, which were the mostly picked in their corresponding category. This avatar is depicted in Section 5.3.3.

5.3.4 Context of creation

Another aspect we want to investigate, is how the context in which the user creates his avatar, affects the creation process. More precisely, we want to look at the possible effect of the context on the creation time and the confirmation time. Table A.1 serves to compare these values based on the location in which the avatar was generated. From this Table, as well as from Figure 5.4, we can infer that users tend to spend the most time on the creation process when at home and therefore uninterrupted. We conclude that users tend to care about the creation, given that they have the time to do so. Therefore, when assuming that most users will tend to the creation process while being on the go, for example directly after noticing the application in a public environment, it is advised to provide a generally smaller set of features, so that users have a more easy time choosing between features. Otherwise, when assuming that users will be creating their avatar in an environment in which they can spend more time on the creation process, a greater variety can be given to satisfy the demands of users for individuality of their content.

5.3.5 Possible improvements

The final set of data that we want to look into, is the set of suggested improvements for our application, which users provided through the user-form of our study. Many users have noted, that they would prefer a larger set of features in the full customisation of their personal avatar. Additionally, some users have suggested to add more feature categories to the creation process, such as noses or the addition of a body for the avatars. A remarkable category, which many participants have mentioned to be desirable, was the category of facial hair.

6 Conclusion and Future Work

Our final chapter will now draw a conclusion based on our user study and from the conducted feasibility-test. Moreover, we will take a look at the factors that have lead to limitations and restrictions on this work. Finally, we delve into the field of possible future work that might be based on what has been established within this thesis.

6.1 Conclusion

In order to find a methodology of overcoming the problem of display and interaction blindness identified by previous work, we have implemented an infrastructure, designed to explore preferences for designing virtual avatars with the goal of promoting user engagement by incorporating them into public display applications. To do so, we have implemented an android application, allowing users the creation and personalisation of such virtual avatars, which served to conduct a user study in order to explore the visual preferences in designing avatars. Our study has lead to the observation, that users value high levels customisability, prefer avatars that let them express their creativity and individuality over predefined avatars or a simple photograph, and that when creating an avatar, certain visual aspects of an avatar are of higher value to the user than other. The system we have implemented consists of two android applications that communicate with each other in order to allow users to interact with the public display and to detect users in the close vicinity of the public display. This is done, so that our public display application can visualise the avatars of users nearby in hopes to attract their attention and motivate them to become active. The communication between our applications is established by using the bluetooth protocol. The application we have implemented to do our user study with, has been expanded to function as the client application in the implemented infrastructure. Additionally, we have conducted a feasibility test for our implemented infrastructure. The system has proven to be stable and to offer the desired functionality and effect, at least to some degree. Since this was not a full-blown user study, the findings are not representative and so, the implemented system should be evaluated in a user study in future work.

6.2 Limitations

In this section, we would like to take a closer look at the limitations and shortcomings of this work. One of the major limitations of this work is the size of our sample in the user study to explore visual preferences in designing avatars. The number of participants in our study is rather low and therefore, our evaluation might not be representative. This is partially due to the lack of time we had at our hands.

There are also a few factors, which did not only limit this work, but also the extend of our implemented infrastructure. Another factor which limits the infrastructure quite heavily, is the use of the android bluetooth protocol. This protocol has a number of shortcomings for the desired task. As mentioned previously, discovering nearby devices and handling connections simultaneously is quite resource-heavy and requires a lot of orchestrating of different tasks to manage it. Also, handling multiple connections simultaneously is generally a difficult task, since each established connection needs to be handled separately. We do believe, that by looking into other protocols to manage the connection between applications, one might be able to overcome these issues. Basing the infrastructure on Android has also brought many restrictions about. The initially desired process of interaction could not be fully achieved due to restrictions of the android framework on developers. As for android version 6.0 and above; any action that requires access to bluetooth must be explicitly granted by the user. Furthermore, devices running such an android-version can not be set to be discoverable for an indefinite amount of time through applications. Therefore, the desired communication is quite heavily limited and not as seamless as we had hoped it to be.

Furthermore, the extend of our mobile application, especially the extend of the image library, was heavily limited by the amount of time that could be used for this task. With more time and also higher artistic skills, the feature set would have been larger and so, the effect of the avatar generator on user behaviour could have been more thoroughly investigated. There are a number of libraries specifically made for use in avatar generators. However, these tend to be quite pricey, which is why we chose to create our own set of features.

Unfortunately, the time did not proof to be enough in order to conduct an additional in-depth user study after our feasibility test. Otherwise, a full blown user study might have been conducted in order to investigate the effects of avatars on user behaviour and user engagement in the context of public display applications.

6.3 Future work

This section serves to discuss possible future work, based on the findings of this thesis or investigating a similar hypothesis. As mentioned previously, we did not conduct a final study in order to investigate the effect that our system could have on user engagement. Therefore, in future work, a full blown user study, in order to investigate the effect of virtual avatars on user behaviour and engagement, should be conducted. Such a user study could also investigate the incorporation of virtual avatars into the different forms, a public display application could occur in, and how avatars can be incorporated to achieve the utmost efficiency. Further exploration and deployment of avatars in the context of public display applications, such as the use of animated avatars compared to static ones or additional methods of avatar generation could also be part of future work. The use of alternative frameworks to build such an infrastructure should also be explored in following research.

Furthermore, as has been implied in the previous section, the established communication could be more efficient by using different near field communication protocols. Therefore, future work investigating the use of other NFC-protocols, such as RFID, wifi-direct or even bluetooth-low-energy, could potentially provide better results.

A Appendix

A.1 Tables

Location: away	
Create Time	Confirm time
46,32	22,56
83,98	20,81
26,32	6,36
2,14	3,86
34,16	3,14
31,58	3,09
Location: work	
Create Time	Confirm time
1,96	0,925
1,83	23,692
2,14	3,862

Location: at home	
Create Time	Confirm time
16,69	1,32
35,69	28,96
37,52	1,14
1,96	0,93
0,84	0,79
26,49	111,74
1,83	23,69
52,45	2,67
79,19	4,83
45,64	1,89
151,06	10,47

Table A.1: The total amount of time in seconds, each participant required to create and to then confirm their avatar, separated by the context in which the avatar was generated.

Time (in seconds) spend per category					
faceshape	hair	hair colour	eyes	eye colour	mouth
15,26	6,75	2,40	5,36	1,79	4,13
8,38	11,39	3,96	6,51	2,85	4,43
11,35	5,83	5,35	0,00	4,96	0,00
7,03	5,20	2,92	6,88	14,58	9,68
6,44	4,771	5,04	2,89	2,30	14,76
6,95	9,363	10,07	7,26	4,95	9,33
10,47	21,129	12,83	20,43	7,84	10,62
1,38	6,615	2,25	3,4	13	12
9,77	3,678	4,67	1,07	3,05	2,84
3,98	3,2078	9,03	3,52	6,48	9,11
17,89	6,978	6,22	10,32	4,73	32,99
7,28	5,094	5,07	9,95	2,73	15,45
7,80	5,832	10,85	29,96	14,10	34,06
3,71	7,457	0	9,79	4,24	6,84

Table A.2: Table listing the amount of time in seconds, each user spend in each category while creating their avatar using the full customisation method.

Clicks performed per category					
faceshape	hair	hair colour	eyes	eye colour	mouth
20	6	0	1	1	2
1	1	1	1	3	1
13	4	2	7	1	0
1	1	4	1	3	1
1	2	1	1	1	2
1	4	17	8	0	5
2	6	1	14	0	10
3	15	22	4	3	14
1	1	32	2	2	17
5	1	6	1	2	3
1	6	10	7	5	17
2	1	8	6	2	7
18	1	3	16	1	21
1	3	2	9	2	3

Table A.3: Table listing the number of clicks each user spend in each category while creating their avatar using the full customisation method.

Bibliography

[ATK16] C. Ackad, M. Tomitsch, J. Kay. “Skeletons and Silhouettes: Comparing User Representations at a Gesture-based Large Display.” In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), pp. 2343–2347. DOI: 10.1145/2858036.2858427 (cit. on pp. 15, 19).

[BPO08] M. Boberg, P. Piippo, E. Ollila. “Designing Avatars.” In: *Proceedings of the 3rd International Conference on Digital Interactive Media in Entertainment and Arts*. DIMEA ’08. Athens, Greece: ACM, 2008, pp. 232–239. ISBN: 978-1-60558-248-1. DOI: 10.1145/1413634.1413679 (cit. on p. 21).

[BR03] H. Brignull, Y. Rogers. “Enticing people to interact with large public displays in public spaces.” In: *Proceedings of INTERACT* 3.c (2003), pp. 17–24. DOI: 10.1.1.129.603 (cit. on pp. 17, 18).

[CA00] A. D. Christian, B. L. Avery. “Speak out and Annoy Someone: Experience with Intelligent Kiosks.” In: CHI ’00 (2000), pp. 313–320. DOI: 10.1145/332040.332449 (cit. on pp. 12, 19).

[Cas03] E. Castronova. “Theory of the Avatar.” In: (2003). DOI: https://ssrn.com/abstract=385103 (cit. on pp. 12, 21).

[CDF+05] K. Cheverst, A. Dix, D. Fitton, C. Kray, M. Rouncefield, C. Sas, G. Saslis-Lagoudakis, J. G. Sheridan. “Exploring Bluetooth Based Mobile Phone Interaction with the Hermes Photo Display.” In: *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices &Amp; Services*. MobileHCI ’05. Salzburg, Austria: ACM, 2005, pp. 47–54. ISBN: 1-59593-089-2. DOI: 10.1145/1085777.1085786 (cit. on p. 21).

[CDKS12] S. Clinch, N. Davies, T. Kubitza, A. Schmidt. “Designing Application Stores for Public Display Networks.” In: *Proceedings of the 2012 International Symposium on Pervasive Displays*. PerDis ’12. Porto, Portugal: ACM, 2012, 10:1–10:6. ISBN: 978-1-4503-1414-5. DOI: 10.1145/2307798.2307808 (cit. on p. 20).

[CND+04] E. F. Churchill, L. Nelson, L. Denoue, J. Helfman, P. Murphy. "Sharing Multimedia Content with Interactive Public Displays: A Case Study." In: *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*. DIS '04. Cambridge, MA, USA: ACM, 2004, pp. 7–16. ISBN: 1-58113-787-7. DOI: 10.1145/1013115.1013119 (cit. on p. 11).

[DFN+09] N. Davies, A. Friday, P. Newman, S. Rutlidge, O. Storz. "Using Bluetooth Device Names to Support Interaction in Smart Environments." In: *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services*. MobiSys '09. Kraków, Poland: ACM, 2009, pp. 151–164. ISBN: 978-1-60558-566-6. DOI: 10.1145/1555816.1555832 (cit. on p. 20).

[DLKM16] G. Du, L. Lohoff, J. Krukar, S. Mukhametov. "Comparing Two Methods to Overcome Interaction Blindness on Public Displays." In: *Proceedings of the 5th ACM International Symposium on Pervasive Displays*. PerDis '16. Oulu, Finland: ACM, 2016, pp. 243–244. ISBN: 978-1-4503-4366-4. DOI: 10.1145/2914920.2940339 (cit. on pp. 12, 19).

[DS] A. Dix, C. Sas. "Mobile Personal Devices meet Situated Public Displays: Synergies and Opportunities." In: (). URL: http://www.hcibook.com/alan/papers/MPD-SPD-2009/ (cit. on p. 20).

[DWYW09] N. Ducheneaut, M.-H. Wen, N. Yee, G. Wadley. "Body and mind." In: *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09* (2009), p. 1151. ISSN: 00278424. DOI: 10.1145/1518701.1518877 (cit. on pp. 12, 21).

[FTLB08] M. Finke, A. Tang, R. Leung, M. Blackstock. "Lessons Learned: Game Design for Large Public Displays." In: DIMEA '08 (2008), pp. 26–33. DOI: 10.1145/1413634.1413644 (cit. on p. 19).

[HKB08] E. M. Huang, A. Koster, J. Borchers. "Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays?" In: *Pervasive Computing: 6th International Conference, Pervasive 2008 Sydney, Australia, May 19-22, 2008 Proceedings*. Ed. by J. Indulska, D. J. Patterson, T. Rodden, M. Ott. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 228–243. ISBN: 978-3-540-79576-6. DOI: 10.1007/978-3-540-79576-6_14 (cit. on pp. 15, 16).

[HROS07] A. P. Holleis, E. Rukzio, F. Otto, A. Schmidt. "Privacy and Curiosity in Mobile Interactions with Public Displays." In: *CHI* (2007). URL: http://uie.bit.uni-bonn.de (cit. on p. 20).

[HW13] S. Houben, C. Weichel. “Overcoming Interaction Blindness Through Curiosity Objects.” In: *CHI ’13 Extended Abstracts on Human Factors in Computing Systems*. CHI EA ’13. Paris, France: ACM, 2013, pp. 1539–1544. ISBN: 978-1-4503-1952-2. DOI: 10.1145/2468356.2468631 (cit. on pp. 12, 19).

[KFF+09] N. Kaviani, M. Finke, S. Fels, R. Lea, H. Wang, K. Nima, F. Matthias, F. Sidney, L. Rodger, W. Hua. “What goes where? Designing interactive large public display applications for mobile device interaction.” In: *Proceedings of the First International Conference on Internet Multimedia Computing and Service* (2009), pp. 129–138. DOI: 10.1145/1734605.1734637 (cit. on p. 20).

[KOK+13] H. Kukka, H. Oja, V. Kostakos, J. Gonçalves, T. Ojala. “What Makes You Click: Exploring Visual Signals to Entice Interaction on Public Displays.” In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’13. Paris, France: ACM, 2013, pp. 1699–1708. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2466225 (cit. on pp. 11, 12, 17, 18, 23).

[KY06] H.-S. Kang, H.-D. Yang. “The Visual Characteristics of Avatars in Computer-mediated Communication: Comparison of Internet Relay Chat and Instant Messenger As of 2003.” In: *Int. J. Hum.-Comput. Stud.* 64.12 (Dec. 2006), pp. 1173–1183. ISSN: 1071-5819. DOI: 10.1016/j.ijhcs.2006.07.003 (cit. on p. 21).

[LCWD14] K. Lee, S. Clinch, C. Winstanley, N. Davies. “I Love My Display: Combatting Display Blindness with Emotional Attachment.” In: *Proceedings of The International Symposium on Pervasive Displays*. PerDis ’14. Copenhagen, Denmark: ACM, 2014, 154:154–154:159. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611038 (cit. on p. 19).

[LLLL10] T. Le Grange, D. Leow, L. Laxman, A. Lau. “Real-time content translation framework for interactive public display systems.” In: *Proceedings - 2010 International Conference on User Science and Engineering, i-USEr 2010*. 2010, pp. 307–310. ISBN: 9781424490493. DOI: 10.1109/IUSER.2010.5716771 (cit. on p. 20).

[MCA15] N. Memarovic, S. Clinch, F. Alt. “Understanding Display Blindness in Future Display Deployments.” In: *International Symposium on Pervasive Displays*. 2015. ISBN: 9781450336086. DOI: 10.1145/2757710.2757719 (cit. on pp. 11, 15, 17).

[MLA+12] N. Memarovic, M. Langheinrich, F. Alt, I. Elhart, S. Hosio, E. Rubegni. “Using Public Displays to Stimulate Passive Engagement, Active Engagement, and Discovery in Public Spaces.” In: *Proceedings of the 4th Media Archi-*

tecture Biennale Conference on Participation (MAB '12) (2012), pp. 55–64. DOI: 10.1145/2421076.2421086 (cit. on p. 16).

[MWB+12] J. Mueller, R. Walter, G. Bailly, M. Nischt, F. Alt. "Looking glass: a field study on noticing interactivity of a shop window." In: *Proc. Conf. on Human Factors in Computing Systems (CHI)* (2012), pp. 297–306. DOI: 10.1145/2207676.2207718 (cit. on p. 18).

[MWE+09] J. Mueller, D. Wilmsmann, J. Exeler, M. Buzeck, A. Schmidt, T. Jay, A. Krüger. "Display Blindness: The Effect of Expectations on Attention Towards Digital Signage." In: *Proceedings of the 7th International Conference on Pervasive Computing*. Pervasive '09. Nara, Japan: Springer-Verlag, 2009, pp. 1–8. ISBN: 978-3-642-01515-1. DOI: 10.1007/978-3-642-01516-8_1 (cit. on pp. 11, 15, 17).

[PKD14] G. Parra, J. Klerkx, E. Duval. "Understanding Engagement with Interactive Public Displays." In: *Proceedings of the International Symposium on Pervasive Displays (PerDis'14)* (2014), pp. 180–185. DOI: 10.1145/2611009.2611020 (cit. on pp. 11, 15, 17, 18).

[SWHK15] V. Schwind, K. Wolf, N. Henze, O. Korn. "Determining the Characteristics of Preferred Virtual Faces Using an Avatar Generator." In: *Proceedings of the ACM SIGCHI Annual Symposium on Computer-Human Interaction in Play (CHI PLAY)*. Vol. 2. Oct. 8, 2015. ISBN: 978-1-4503-3466-2. DOI: 10.1145/2793107.2793116. published (cit. on p. 21).

[TAD+14] M. Tomitsch, C. J. Ackad, O. Dawson, L. Hespanhol, J. Kay. "Who Cares About the Content? An Analysis of Playful Behaviour at a Public Display." In: *International Symposium on Pervasive Displays*. 2014, pp. 160–165. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611016 (cit. on p. 19).

[VCBE08] T. Vajk, P. Coulton, W. Bamford, R. Edwards. "Using a mobile phone as a 'Wii-like' controller for playing games on a large public display." In: *International Journal of Computer Games Technology* 2008 (2008), p. 6. ISSN: 1687-7047. DOI: 10.1155/2008/539078 (cit. on pp. 20, 21).

[YMQQ07] G. T. Young, N. Y. Matthes, K. G. Qld, K. G. Qld. "Virtual Fish : Visual Evidence of Connectivity in a Master-Planned Urban Community." In: November (2007), pp. 28–30 (cit. on pp. 12, 19).

[Yon09] J. Yong. "The Understanding of Avatar Design : Various Types and Styles of Avatars and the Design Considerations." In: 18 (2009). ISSN: 1738-1541 (cit. on p. 21).

All links were last followed on April 10, 2017.

Declaration

I hereby declare that the work presented in this thesis is entirely my own and that I did not use any other sources and references than the listed ones. I have marked all direct or indirect statements from other sources contained therein as quotations. Neither this work nor significant parts of it were part of another examination procedure. I have not published this work in whole or in part before. The electronic copy is consistent with all submitted copies.

place, date, signature