ART – Augmented Reality above the Tabletop

An immersive analytics tool for the visual analysis of mobile health data

Bachelor Thesis

by

Sebastian Hubenschmid
(01/794812)

at the

Universität Konstanz

Department of Computer and Information Science

First Reviewer:	Prof. Dr. Harald Reiterer
Second Reviewer:	Prof. Dr. Bastian Goldlücke
Adviser:	Simon Butscher

Constance, 2017

Zusammenfassung

Immersive Analytics untersucht inwiefern immersive Technologien, wie beispielsweise Augmented Reality Geräte und große berührungsempfindliche Displays, bei der Analyse von komplexen Daten behilflich sein können. Diese Arbeit untersucht das Potenzial von *Immersive Analytics* für die kollaborative Analyse von Mobile Health Daten, wobei sich das Vorgehen auf einen spezifischen Anwendungsfall konzentriert.

Zuerst werden Anforderungen und Einschränkungen von derzeitigen Analyseprozessen diskutiert, welche auf den Ergebnissen einer Fokusgruppe zwischen vier Interaktionsdesignern und sieben Fachleuten des SMART-ACT Projekts basieren. Die Fokusgruppe zeigte einige Ansatzpunkte für immersive Analysewerkzeuge die sich den derzeitigen Analyseansätzen, der Lesbarkeit der Visualisierung, Raum und Immersion sowie der Zusammenarbeit widmen.

Danach wird das kollaborative Analysewerkzeug ART (**A**ugmented **R**eality above the **T**abletop) für die visuelle Analyse von Mobile Health Daten präsentiert. ART befasst sich mit den wichtigsten identifizierten Anforderungen und Ansatzpunkten, indem es multidimensionale Daten in Augmented Reality mittels einer 3D Visualisierung darstellt. Die Visualisierung verbindet gleiche Datenpunkte zwischen verschiedenen 2D Scatterplots, womit ein 3D paralleler Koordinatenplot erzeugt wird. Um von etablierten Interaktionsmethoden zu profitieren ist die Visualisierung auf einem berührungsempfindlichem Tabletop verankert. Diese Arbeit demonstriert weiterhin die Umsetzung solch eines verteilten Systems, das aktuelle Virtual Reality Head-Mounted Displays mit stereoskopischen Kameras verwendet, um eine immersive Augmented Reality Umgebung zu erschaffen.

Zuletzt wurden für die Evaluation von ART zwei gruppenbasierte Experten Walkthroughs mit insgesamt zehn Teilnehmern durchgeführt. Die Walkthroughs offenbaren bereits mehrere Vorteile für die Benutzung des ART Systems gegenüber traditionellen, desktop-basierten Analyseansätzen, wie beispielsweise eine bessere Zusammenarbeit sowie eine einfachere Datenerkundung. Anhand dieser Ergebnisse wird eine Reihe von Designempfehlungen sowie weitere Forschungsrichtungen bereitgestellt für die Integration von immersiven Technologien in die kollaborative Analyse von multidimensionalen Daten.

Abstract

Immersive Analytics studies how immersive technologies, such as augmented reality devices and large touch-sensitive displays, can be instrumental in analysing complex data. This work investigates the potential of *Immersive Analytics* for the collaborative analysis of mobile health data, using a use-case-centred design approach.

Firstly, requirements and limitations of current analysis processes are discussed, based on the results of a focus group between four interaction designers and seven domain experts from the SMARTACT project. The focus group revealed several leverage points for *Immersive Analytics* tools addressing current analysis approaches, visualisation readability, space & immersion, and collaboration.

Secondly, the collaborative analysis tool ART (**A**ugmented **R**eality above the **T**abletop) for the visual analysis of mobile health data is presented. ART addresses the key identified requirements and leverage points by visualising multidimensional data in augmented reality using an interactive 3D visualisation. The visualisation links related data points between several 2D scatter plots to create a 3D parallel coordinates visualisation. To benefit from well-established interaction techniques, the visualisation is anchored to a touch-sensitive tabletop. This work further illustrates the implementation of such a distributed system, which uses contemporary virtual reality head-mounted displays with stereoscopic cameras to create an immersive augmented reality environment.

Thirdly, two group-based expert walkthroughs were conducted to evaluate ART, with ten participants in total. The walkthroughs already revealed several benefits of using the ART system over traditional, desktop-based analysis approaches, such as better collaboration and easier data exploration. Based on results from these walkthroughs, a set of guidelines and further research directions to integrate immersive technologies into the collaborative analysis of multidimensional data is provided.

Einreichung zur Veröffentlichung

Teile dieser Arbeit wurden zur Veröffentlichung auf der Konferenz

> *The ACM CHI Conference on Human Factors in Computing Systems 2018*

eingereicht unter dem Titel:

> *Clusters, Trends, and Outliers:*
> *How Immersive Technologies can Facilitate Collaborative Analysis of Multidimensional, Health-Related Data*

Danksagung

Hiermit möchte ich mich bei den Leuten bedanken, die zu dieser Arbeit beigetragen haben. Ein großes Dankeschön geht an Simon Butscher für die stetige Betreuung in allen Teilen dieser Arbeit. Weiterhin bedanke ich mich bei Dr. Johannes Fuchs, der zur Entstehung dieses Projektes beigetragen hat. Zudem möchte ich mich bei Prof. Dr. Harald Reiterer und allen Mitarbeitern des HCI-Lehrstuhls bedanken für das hilfreiche Feedback und für die Bereitstellung jeglicher notwendigen Hardware. Des Weiteren will ich mich auch bei den Psychologen des SMARTACT Projekts bedanken, mit denen diese Arbeit sehr anwendungsnah evaluiert werden konnte. Letzlich auch ein Danke an meine Freunde und Familie.

Contents

List of Figures

List of Tables

Abbreviations

AR . . . . augmented reality

ART . . . . **A**ugmented **R**eality above the **T**abletop

BMI . . . . body mass index

DOF . . . . degrees of freedom

FOV . . . . field of view

FPS . . . . frames per second

HMD . . . . head-mounted display

IMU . . . . inertial measurement unit

VR . . . . virtual reality

Chapter 1

Introduction

In the age of information and ubiquitous computers, the amount and complexity of data being generated is increasing rapidly. The prevalence of mobile devices, for example, allows users to track more and more information about their lives with applications such as fitness trackers or dietary journals. The tracking of such health-related data through mobile devices is commonly referred to as mobile health [1].

The SMARTACT project [41, 85], for example, uses mobile intervention studies to guide participants towards healthier lifestyles. For this, participants track their meals and physical activity through mobile devices, and are provided with visual presentations of their data [14]. This data can also be used by researchers to investigate trends and correlations between different participants. These researchers often rely solely on statistical methods for analysis, as traditional visualisations and desktop systems are too restricted for these data sets [63]. Yet, the visual analysis can help to gain an essential understanding of the data, which in turn can be fundamental for the discovery process [27].

Immersive technologies, such as virtual reality (VR), can be beneficial for such data visualisations. The research field of *Immersive Analytics* [17] explores these emerging technologies to facilitate the visual data analysis by combining intuitive interfaces with immersive technologies. The use of these technologies also offers advantages outside of data visualisation, which can further improve current analysis approaches. For example, the use of touch interaction can support natural collaboration between users, which in turn encourages discussions about current findings.

This work therefore studies the potentials of mixed reality environments for the visualisation of mobile health data. An investigation into the limitations of current analysis approaches used by researchers from the SMARTACT project guided the development of the immersive analytics tool **A**ugmented **R**eality above the **T**abletop (ART). ART displays abstract data in a 3D parallel coordinates visualisation above an interactive tabletop. Touch interaction allows for intuitive operation, while the augmented reality (AR) environment supports natural collaboration between domain experts.

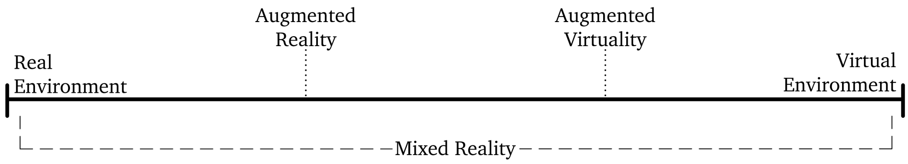

Figure 1.1: The Virtuality Continuum provides a taxonomy for classifying mixed reality applications [62].

1.1 Virtual and Mixed Reality

Although hardware for VR has existed for decades, previous research into VR has been largely concerned with – and limited by – technical details. With the recent advances in commercial VR headsets such as the Oculus Rift, this is no longer the case. However, commercial use of VR is often associated with the application's hardware, restricting its definition to certain technologies. By using Steuer's [79] definition of VR in terms of telepresence[1], VR can instead be characterised by the vividness and interactivity of the application:

> *'A "virtual reality" is defined as a real or simulated environment in which a perceiver experiences telepresence.' [79]*

Milgram and Kishino [62] further restrict VR to a completely virtual environment, and describe the fluidity between a purely real and a purely virtual world with a *Virtuality Continuum* (see Figure 1.1). Whereas VR contains purely digital elements, any mixture of real and virtual elements on a screen fall under the category of mixed reality. The mixed reality spectrum can be further refined into AR, where users are in a real world that is enriched by digital objects, and augmented virtuality, where users are in a virtual world with real elements.

Yet, this specific definition fails to distinguish between interactive and non-interactive mediums. For example, placing a button on top of a webcam-stream would satisfy the condition of mixing a real world with digital elements, and thus qualify as AR application. For this reason, Azuma [4] specified three key characteristics for AR applications:

1. Combines real and virtual
2. Interactive in real time
3. Registered in 3D

AR can also facilitate collaboration, as users can still perceive their surroundings and therefore communicate naturally with other users (i.e. using non-verbal communication cues such as gestures), whereas other mixed reality environments must rely on other means of communication (e.g. virtual avatars).

[1]Telepresence refers to the sense of being present through mediated means, such as immersing oneself in a video game [79].

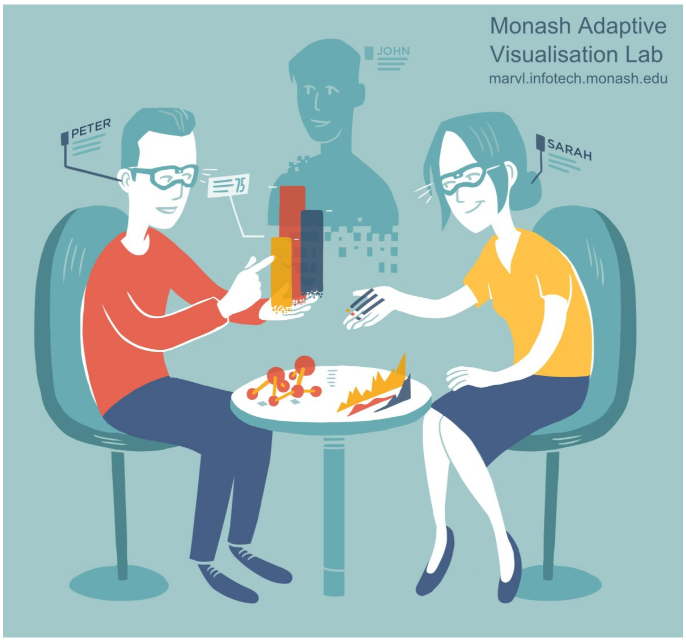

Figure 1.2: An exemplary immersive analytics tool: Two real users use wearable devices to augment their world with digital information, with which they can interact through hand gestures. Remote users can also join their session through a virtual presence [17].

1.2 Immersive Analytics

The research area of *Immersive Analytics* [17] aims to examine immersive technologies, such as mixed reality head-mounted displays (HMDs), and similar, emerging technologies in the context of data analysis (see Figure 1.2). In particular, *Immersive Analytics* brings up several research directions, such as:

- Investigating the collaboration in data analysis scenarios with abstract data sets, as this is largely unexplored. Mixed reality environments and novel interaction methods (e.g. touch interfaces) may be beneficial for collaboration.
- Reevaluating previous research on 3D visualisation, as better depth perception and egocentrical navigation may invalidate conventional wisdom.
- Identifying domain-specific requirements and workflows for immersive data analysis. The application of *Immersive Analytics* to certain use cases may offer insight that can be applied in other domains.

This work therefore investigates the specific use case of mobile health data for the application of a collaborative immersive analytics tool.

1.3 Use Case

Mobile intervention studies aim at improving people's health [43]. Study participants typically use mobile devices (e.g. smartphones) to track their health behaviour (e.g. eating behaviour and physical activity). Additionally, participants receive visual and/or textual information regarding their tracked behaviour on their smartphones. This feedback mechanism is referred to as an intervention. Current research in the area of behavioural and nutritional science is concerned with the analysis of user's behaviour signatures and how they relate to the effectiveness of intervention mechanisms. The investigation of this relation involves the analysis of multidimensional, health-related data and is a complex undertaking.

The SMARTACT project [41, 85] uses intervention studies to track eating behaviours through mobile devices. SMARTACT is a collaboration between the *University of Konstanz*, the *University of Mannheim* and the *Karlsruhe Institute of Technology*, funded by the German *Federal Ministry of Education and Research*. The project consists of an interdisciplinary team, including health and biological psychologists, sport scientists, computer scientists, and economists. Data is gathered with the use of mobile devices (e.g. smartphones and physical activity trackers) [14]. For every meal, participants may answer a questionnaire, as well as classify and take a picture of their food through a mobile device. Subsequent feedback about the meal's healthiness, as well as various interventions may help participants to improve and maintain a healthier lifestyle. To increase effectiveness, participants are presented with visualisations about their tracked data [16, 23].

The visualisation of this dataset is not only important as feedback mechanism for participants, but is also useful for analysts. However, these analysts are often experts in terms of statistical methods, not visualising the data. As such, they rely on traditional desktop systems for exploring and analysing data. These analysis methods have three major drawbacks [63]:

1. Relations between more than two dimensions are hard to visualise.
2. Display space limits the visualisation potentials.
3. The desktop environment is not designed to support collaborative analysis scenarios.

The research area of *Immersive Analytics* can therefore be instrumental in supporting these complex data analysis scenarios. Unlike traditional desktop systems, these technologies provide the means to visualise complex information in a physical space, allowing large amounts of data to be simultaneously investigated [27]. In addition, immersive technologies 'can facilitate the visual perception of users in a natural way, which, in turn, helps them quickly identify areas of interest, meaningful patterns, anomalies, and structures between artefacts that are hard to discover without spatial representations.' [66] Therefore, these technologies open up a new design space for the creation of interactive 3D visualisations.

1.4 Research Questions

Although there is a body of work investigating the general benefits of immersive technologies for data visualisation, their potentials for the domain-specific, collaborative analysis of multidimensional, abstract data has not yet been widely researched. This work explores a use-case-centred design, using a combination of traditional touch input devices in an immersive AR environment for an exploratory collaborative visual data analysis scenario. This offers especially two research questions:

[RQ.1] *How can immersive visualisations help overcome current limitations of traditional data analysis approaches?*

[RQ.2] *In what ways can an immersive, touch-based interaction facilitate the interaction with 3D visualisations?*

1.5 Contributions

For the creation of a use-case-centred immersive analytics tool, a concrete analysis of the use case at hand is necessary. Furthermore, the system has to be designed to improve the specific data analysis scenario, and implemented within the appropriate technical setting. Lastly, the finished system has to be evaluated by domain experts in order to determine if the system could facilitate and support the use case. As a result, this thesis contributes the following three aspects:

1. **Requirements & Leverage Points**

   A discussion of *requirements and leverage points* that help the collaborative analysis of multidimensional, health-related data. These findings are deduced from a workshop with domain experts and offer insight into problems and limitations with current analysis approaches.

2. **ART System**

   The *ART system* including its design, implementation, and technical setting. ART visualises multidimensional data in AR using multiple scatter plots with linked data points, creating a 3D parallel coordinates visualisation. The visualisation is anchored to a touch-sensitive tabletop, enabling a familiar operation.

3. **Design Guidelines & Further Research Directions**

   Design guidelines and further research directions for immersive technologies in the collaborative analysis of multidimensional data. Findings result from two group-based expert walkthroughs, within which experts from the domains of behavioural and nutritional science evaluated the usefulness of ART to collaboratively analyse clusters, trends, and outliers.

1.6 Methodical Approach

Following the structure of these contributions, a methodical approach consisting of three phases was used for this project: An *analysis phase* for examining the current use case; a *design phase* for creating the ART system; and lastly an *evaluation phase* to gather feedback.

Analysis Phase. A focus group with the domain experts from the SMARTACT team was organised in order to establish a detailed understanding of the use case at hand [63]. First, the data collection was analysed and current scenarios, in which data needs to be visualised and interpreted, were identified. In addition, associated limitations were discussed. Then, interaction designers provided a demonstration of several immersive technologies like HMD or AR tablets. Afterwards, possible use cases for immersive analytics interfaces as tools to enhance visual analysis and open issues were discussed.

Design Phase. A system prototype was designed and implemented for the target use case. The result is the collaborative analysis tool ART (see Figure 1.3), addressing identified key requirements and leverage points from the focus group. The tool visualises multidimensional data in AR using HMDs with front-mounted cameras. It uses a novel visualisation technique, combining 2D scatter plots with parallel coordinated by linking together related data points. To benefit from well-established interaction techniques, the visualisation is anchored to a touch-sensitive tabletop, where users can configure the scatter plots.

Evaluation Phase. To gather user feedback, the ART system was again presented to domain experts. For this, two expert walkthroughs were conducted, in which participants performed a collaborative analysis on real data using the ART system. Initial impressions show improvements in analysing high-dimensional data using an AR environment. Despite being non-experts in terms of visualisation, participants were quickly able to control the system and establish a visual analysis workflow.

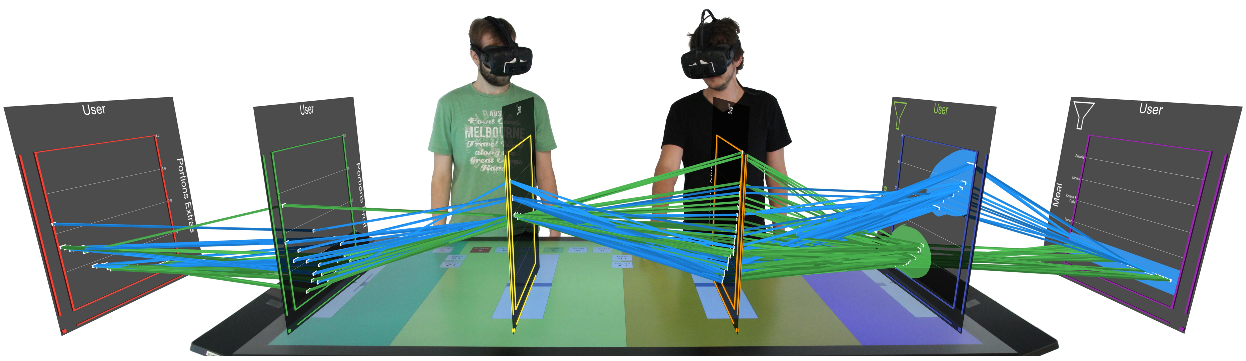

Figure 1.3: **A**ugmented **R**eality above the **T**abletop (ART) is designed to facilitate the collaborative analysis of multidimensional data. A 3D parallel coordinates visualisation in AR is anchored to a touch-sensitive tabletop, enabling familiar operation.

1.7 Outline

The structure of this document is based on the methodical approach, consisting of an analysis, design, and evaluation phase. First, Chapter 2 introduces the *analysis phase* by investigating the primary use case for the ART system. This chapter describes the initial focus group in detail and explains problems and limitations with current analysis approaches. In addition, these results were summarised as concrete requirements for the ART system.

Furthermore, Chapter 3 provides an overview over related work. Previous research into immersive interaction, collaboration, and visualisations illustrates several benefits for immersive data analysis tools. A comparison of contemporary immersive AR hardware provides a technical foundation for the ART system.

ART was created specifically for the purpose of analysing health-related data in an AR environment. Chapter 4 describes the *design phase* by providing an overview over ART's design and features, and outlines the implementation in terms of hardware and software.

Chapter 5 covers the final *evaluation phase*. For this, two expert walkthroughs were conducted during which domain experts acted out different analysis scenarios using the ART system. Based on these findings, design recommendations and directions for future research are provided. Likewise, ART is further assessed based on requirements from the initial use case analysis.

Lastly, Chapter 6 provides a summary for the initial analysis, the ART system design, and findings from the evaluation. Future work directions are provided based on findings and observations from the evaluation.

Chapter 2

Analysis

Before the ART system can be build, an analysis of the use case at hand is necessary. An examination of current workflows used by the target user group, as well as their data set allows ART to focus on specific problems. Section 2.1 presents the results of such an analysis. A focus group with domain experts was conducted, during which current problems and limitations, as well as possible mixed reality solutions were discussed. Section 2.2 condenses these issues into requirements which serve as guide for the development and evaluation of the ART prototype.

2.1 Focus Group

For a use-case-centred design of a visual analysis tool it is important to first get to know the data set, the target user group, and the analysis aims [64]. To address these topics and to identify leverage points, a focus group (2 h) was conducted with members of the SMARTACT project [63]. The group consisted of seven domain experts with a strong background in intervention studies (health psychologists and biological psychologists) and four interaction designers.

First, a typical data set resulting from an intervention study was analysed. Then, exemplary analysis scenarios and associated visualisation and interpretation steps were examined (approx. 45 min). In addition, the limitations of systems and approaches currently in use were discussed. Next, interaction designers provided a demonstration (approx. 30 min) of several state-of-the-art immersive technologies, including the *HTC Vive* HMD or the *Google Project Tango* AR tablet. Afterwards, possible use cases for immersive analysis interfaces and how they might enhance current visual data analysis approaches were discussed (approx. 45 min).

The results are structured according to the type of *data set*, the *target user group*, their *analysis aim*, and *leverage points* for immersive technologies.

2.1.1 Data Set

Applications for mobile intervention studies typically consist of a tracking part and an intervention part, both running on mobile devices. The data set considered in this work was collected during a mobile intervention study which investigated eating behaviour, psychological aspects, and context-related aspects. This data is combined with demographic information and clinical measures. Based on the tracked data, diverse intervention types can be provided. The data set therefore consists of a variety of data:

- **Demographic information** to account for differences outside of the intervention study (e.g. age and gender).
- **Subjective well-being** of individual participants (e.g. stress before and after a meal).
- **Context-related aspects** of individual meals (e.g. eating location and social context).
- **Psychological aspects** of why a meal was consumed (e.g. eating motives such as perceived healthiness of a meal).
- **Nutritional information** of the meal contents (e.g. portions of milk or vegetables).
- **Derived scores** for providing participants with feedback (e.g. calories).
- **Food pictures** including before and after eating the meal.

The data collected through the mobile intervention study is therefore of high complexity. An example data set contains information of 200 participants who tracked their food intake for five weeks. This results in 20 000 data records where each record represents one meal. Each data record has about 80 dimensions. Thus, the data set has the following properties:

1. It is large, because it typically contains information of several hundreds of participants who tracked data over the course of several weeks.
2. It is high-dimensional, because the intervention studies try to collect a holistic picture of the participants.
3. It contains mainly abstract information in terms of numerical or categorical values, but can also contain data such as food pictures.

2.1.2 Target User Group

Researchers who analyse the data set are domain experts in terms of nutritional and behavioural science (health psychologists and biological psychologists). The domain experts stated that they often validate hypotheses using statistical methods, but also apply explorative data analysis approaches to unveil unexpected effects and patterns. They typically use spreadsheets (e.g. *Microsoft Excel* [59]) or statistical applications (e.g. *IBM SPSS* [47]) which provide no or limited possibilities for an interactive visual analysis of the data. The analysts have only recently begun to use visual data analysis tools (e.g. *Tableau* [83]) to explore their data set. However, due to the high dimensionality of the data set, most visualisations are too limited for a visual analysis and exploration of the data set. This explorative part, especially, is sometimes conducted collaboratively. Therefore, the researchers are experts in terms of domain knowledge as well as statistical methods but are non-experts in terms of visual analysis.

2.1.3 Analysis Aims

Mobile health interventions are based on what people do (behavioural patterns), why people do what they do (psychosocial and contextual behavioural triggers), and when people do what they do (timing of behaviour and triggers) [31, 43]. These characteristics of a single participant are brought together to form an individual, high-dimensional behaviour signature. The domain experts emphasised three higher level analysis aims about the gathered, multidimensional data:

1. The identification of clusters of persons with similar behaviour signatures (high-dimensional clusters) and their correlation to related outcomes such as body mass index (BMI) and blood level.
2. The analysis of high-dimensional data on multiple aggregation levels (e.g. participant, day, meal).
3. The analysis of the effectiveness of an intervention in terms of chronological trends within the multidimensional data.

2.1.4 Leverage Points

Aside from the general analysis of the use case, the workshop revealed four limitations which can be considered as leverage points for immersive technologies.

1. **Approach to Analysis**

   Currently used tools provide only limited support for dynamic configuration and extension (e.g. no fluent switching between inter- and intra-individual analysis).

2. **Visualisation Readability**

   Currently used visualisations are not suitable for analysing dependencies between more than two dimensions and for identifying multidimensional clusters, trends, and outliers. Furthermore, they do not support displaying additional data such as food pictures associated with individual data entries.

3. **Space & Immersion**

   For desktop systems, display space is a limiting factor when visualising high-dimensional data. This hampers the users in getting an overview of the data in the collection.

4. **Collaboration**

   With traditional desktop systems, typically only one analyst interacts with the data while others take an observer role. Thus, such systems do not sufficiently support the collaborative, exploratory data analysis.

2.2 Requirements

The focus group unveiled numerous problems and limitations with currently used visual analysis tools, and showed several leverage points for immersive technologies. To better address these problems during the design phase of ART, they are assembled into specific requirements, grouped by the four leverage points of *Approach to Analysis*, *Visualisation Readability*, *Space & Immersion*, and *Collaboration*.

2.2.1 Approach to Analysis

Current analysis approaches of the target user group are often very static, and do not allow for further details or dynamic changes on demand. The aim of these requirements is thus to support and enhance the current analysis approaches by making the analysis more dynamic.

[A.1] *Support dynamic workflow*

'Visualizations can be static or dynamic. Interactive visualizations often lead to discovery and do a better job than static data tools.' [87] Currently used analysis approaches make it hard to dynamically investigate new variables and relations, and often require prior setup before a new analysis can take place. By supporting a more dynamic and interactive workflow, users can take a more exploratory approach when analysing the data set.

[A.2] *Switch between inter- and intra-individual analysis*

The target data set contains information about both individual meals of several participants, as well as information about the participants themselves. The system should support seamless switching between comparing values of a single participant (intra-individual analysis), and comparing these values to other participants (inter-individual analysis).

[A.3] *Reducing the data set*

The raw data set contains information about thousands of meals, which can make it difficult to view individual entries. The system should support methods (e.g. filters or aggregation methods) to reduce the amount of data displayed, so that users can remove unnecessary entries.

[A.4] *Provide easy, intuitive interface for non-experts*

The target users are non-experts in terms of visual analysis. The interface for creating and controlling the visualisation should therefore be obvious to users unfamiliar with data visualisation tools.

2.2.2 Visualisation Readability

Traditional 2D visualisations often do not support many dimensions, or offer little insight due to their high complexity. For a 3D immersive environment, the depth perception can provide a better, more intuitive understanding of a visualisation. Therefore, the aim is to find a visualisation that best supports the target data set.

[V.1] *Show arbitrary number of dimensions*

Due to the high-dimensional data set, the visualisation must be able to show an arbitrary number of dimensions simultaneously.

[V.2] *Display abstract data*

The data set contains mostly numerical and categorical data, such as BMI or calories consumed. Therefore, the visualisation should mainly focus on displaying abstract data.

[V.3] *Display photos for each entry*

Most meals in the data set have one or more photos of the meal associated with the entry. Displaying this extra information is usually not possible in traditional 2D visualisations, but may be practicable in an immersive environment.

[V.4] *Highlight clusters*

Interactivity plays an important role during the visual analysis of cluster patterns [19, 87]. Highlighting clusters allows users to analyse and categorise several clusters, and may encourage discussion with other users.

[V.5] *Identify clusters, trends, outliers*

One of the main analysis goals of the target user group is to find clusters, trends, and outliers within the data set. The visualisation should therefore support and facilitate the identification of such features, either in the visualisation itself or through the help of analysis tools.

2.2.3 Space & Immersion

Both AR and VR offer virtually unlimited space for visualisations, whereas currently used desktop systems are limited by monitor size. By immersing users in the data, the system can visualise more data at once and provide a higher telepresence, resulting in a more natural approach to analysis as digital objects may appear more present.

[SI.1] *Provide high immersion*

A higher immersion means that virtual elements in the scene appear more real, thus increasing telepresence. This narrows the hardware choice down to fully immersive hardware such as VR HMDs, rather than tablet-based AR environments.

[SI.2] *Provide large space for analysis*

Visualisations on traditional desktop systems are restricted by screen size and number of monitors. This limits the size and number of visualisations that can be displayed simultaneously. In contrast, immersive environments are not restricted by any screens, and can therefore offer a much larger space for more visualisations. By providing a large space for analysis, users can create as many visualisations as necessary, therefore facilitating the analysis process.

Although both AR and VR environments support collaboration, AR environments can better support natural collaboration, as users can still see each other without the use of virtual tools (e.g. avatars). Because collaboration is a key leverage point for immersive environments, the ART system uses AR instead of VR for this scenario. Azuma [4] defines three key characteristics of AR applications, which should be fulfilled by the ART system:

[SI.3] *Combines real and virtual*

A key characteristic of AR is to combine a real environment with virtual elements. For this, the system needs to capture a real scene and superimpose digital objects (e.g. the visualisation). This also creates a more familiar work environment for the target user group, and offers natural interaction with users outside of the ART system.

[SI.4] *Interactive in real time*

Interactivity is used to distinguish AR applications from other applications that combine the real world with virtual elements (e.g. movie effects). Interactivity may range from simply navigating the AR environment, to fully interacting with the digital elements by means of buttons or gestures.

[SI.5] *Registered in 3D*

To embed digital objects in a real scene, they have to be properly registered in 3D. This includes objects remaining in the correct physical position, regardless of where the camera is, as well as proper occlusion between real and virtual objects and emulating similar lighting conditions for virtual objects. Virtual objects with a correct 3D registration can appear real, thus increasing telepresence.

2.2.4 Collaboration

Current analysis approaches of the target user group are restricted to one active user, while others have to take a passive observer role. Immersive environments potentially allow for multiple active participants at the same time, allowing all users to be more engaged with the analysis process. The ART system therefore should aim to support and encourage this collaboration by allowing for multiple active users and provide relevant tools.

[C.1] *Allow multiple active users*

The exploratory data analysis can especially benefit from multiple users. Traditional desktop applications usually only allow for one active participant, whereas AR can support multiple users.

[C.2] *Support collaboration between multiple users*

Collaboration can be made more effective by offering tools to enhance the collaborative process. For example, collaboration tools may let users share annotations or highlight parts of the visualisation for discussion.

Chapter 3

Related Work

Although the field of *Immersive Analytics* is fairly new, there is a large body of prior knowledge for interaction, collaboration and visualisation in immersive environments. This chapter explores the previous work for these topics in Sections 3.1, 3.2, and 3.3, respectively. Furthermore, Section 3.4 explores several non-immersive visualisations that can handle high-dimensional data sets. Lastly, Section 3.5 evaluates and compares several hardware choices for immersive environments.

3.1 Interaction in Immersive Environments

'Exploration and analysis are most strongly supported when combining the best possible visual representations with the best possible interaction techniques.' [54] Despite recent advances in terms of HMD hardware, the interaction with visualisations in immersive environments is still an issue. Most commonly, spatially aware input controllers (e.g. *HTC Vive Controller*) or freehand gestures (e.g. with *Leap Motion*) are used. However, these input techniques suffer from the touching the void issue [12]. Other systems rely on gamepads to control and navigate the immersive environment, yet these gamepads often provide only limited interaction possibilities.

More recent research in the field of immersive visualisations considered the combination of immersive display technologies like HMDs or CAVE with touch as input style. Multi-touch devices provide the advantage that they give haptic feedback to the user which is of high importance for both real and virtual environments [71]. Whereas HMDs and stereoscopic displays facilitate depth perception and thus offer high visual immersion, touch-based interaction provides high immersion through interaction due to its directness [54]. Several research projects investigated the combination of AR devices with multi-touch tables to perform object positioning tasks [6, 7, 38, 81].

Toucheo. Hachet et al. [38] presented *Toucheo*, a system which used a stereoscopic mirror-based display to visualise 3D objects above a multi-touch table,

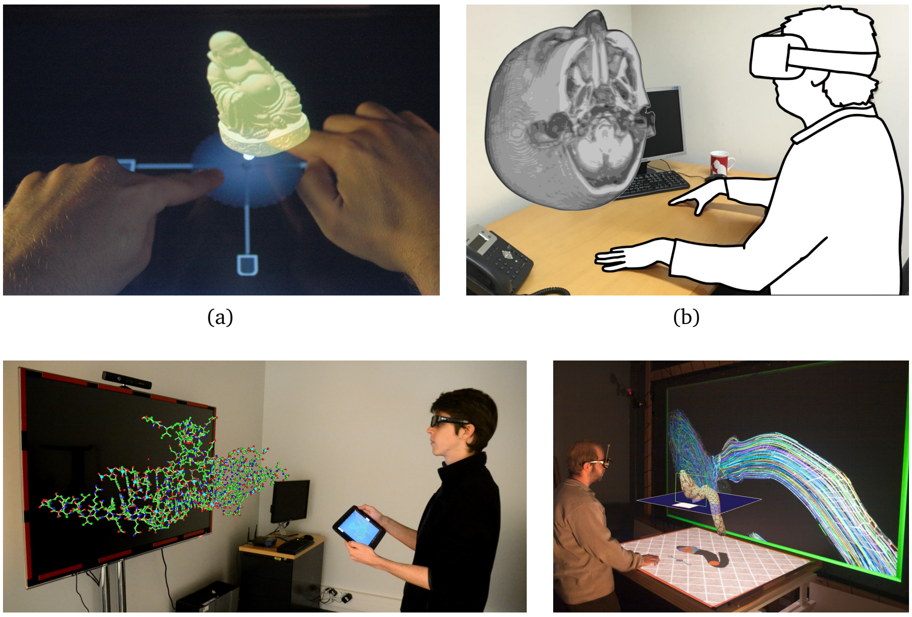

(a) (b) (c) (d)

Figure 3.1: Touch interaction in immersive environments. (a) A hologram is displayed above an interactive display, on which users can control the hologram [38]. (b) An interactive table allows users to control a visualisation inside a VR environment [77]. (c) Users can navigate the 3D visualisation with their tablet [54]. (d) An interactive tabletop allows users to navigate and control the visualisation, while a miniature version of the visualisation hovers above the tabletop [21].

yielding an occlusion-free visualisation for 3D objects (see Figure 3.1a). The system provides familiar touch gestures (e.g. pan and zoom) and widgets for interacting with 3D objects, allowing for translation, rotation, and scaling. In a user study, the interaction through well-known 2D metaphors on the multi-touch table received positive feedback in terms of user experience. However, due to its reliance on mirrors the system limits people's freedom to move around.

Virtual Reality for Radiologists in the Reading Room. Sousa et al. [77] used a similar approach for the analysis of medical images. Their VR setting combined a HMD with a touch-sensitive table, which is also represented in the VR environment (see Figure 3.1b). Since users were unable to see their exact hand position on the table, different gestures could be used in the left and right half of the display, with realtime visual feedback visible on the virtual table. The visualisation hovers in front of the table, allowing users to see the current status on the table and the visualisation simultaneously. Unlike midair controls, this setting and Toucheo provide the advantage that mapping controls to physical objects provides somesthesis[1] feedback, avoiding the touch the void issue.

[1] Somesthesis envelopes the cutaneous (skin) sensation and the capability to sense the movement and position of our limbs (proprioception). [71]

Interaction Modes for Tablet-based 3D Navigation. López et al. [54] combined 3D visualisations on a stereoscopic wall display with touch-based navigation on a tablet device (see Figure 3.1c). Their work addresses the issue of viewpoint desynchronisation between the wall display and the tablet, and reducing fatigue for navigating the visualisation with the tablet. For this, they identified a sets of different interaction modes and a workflow that helps users transition between these interaction modes. For example, the viewpoint can be manually transferred from the tablet to the wall display, and vice versa.

Interactive Slice WIM. Coffey et al. [20, 21] proposed *Slice WIM*, a combination of a vertical stereoscopic wall display to visualise medical volumetric data and a monoscopic horizontal multi-touch display which provides interaction widgets (see Figure 3.1d). This setup creates the impression that a miniature version of the 3D object is hovering above the table, while an interactive slice of the object is displayed on the table below. In a first evaluation users quickly learned the relation between the tabletop and the 3D content hovering above.

Summary. In summary, previous research has shown that combining multi-touch interaction in immersive technologies can provide an intuitive interface for manipulating 3D objects and visualisations. Different interaction modes provide intuitive 3D navigation on a 2D screen, while somesthesis feedback solves the touching the void issue. The combination of multi-touch input and immersive, mixed reality therefore seems suitable for the analysis of health-related data.

3.2 Collaborative Immersive Data Analysis

Decision-making based on data analysis is often the result of a collaborative effort [42, 48, 66]. Mixed reality environments seem to naturally support collaboration [24] as they provide the means to create a shared environment where the collaborators have the feeling of each other's presence. Some research already investigated the influence of immersive environments on collaboration [8, 11, 24, 27, 28, 32, 35, 65, 74, 82]. Billinghurst and Kato [11], for example, pointed out that collaboration can especially benefit from AR environments as these environments can decrease the cognitive and functional load on the user, while Fleury et al. [32] investigate remote collaboration and interaction between different immersive environments. Tanaya et al. [84] propose the *Cross Reality Collaboration Framework* to discuss and compare different immersive environments for mixed reality collaboration scenarios.

Studierstube. The project *Studierstube* [35, 74, 82] provides an AR environment in which users can collaboratively explore virtual objects situated in the space between them (see Figure 3.2a). The advantages of this AR setting are that users can interact with the real world and the virtual world simultaneously, that spatial cues are provided, and that natural collaboration is facilitated. The authors report that their interface is conducive to real world collaboration, because collaboration is mostly left to social protocols.

Visual Interaction Tool for Archeology. In a manner similar to this work Benko et al. [8] developed the collaborative mixed reality system VITA for an

(a) (b)

Figure 3.2: Collaboration during immersive data analysis. (a) Two users discussing a visualisation in an AR environment, using Studierstube [35]. (b) Two users collaborating on a 3D model. One user is presented with a miniature model, while the other user can inspect the larger model [8].

off-site visualisation of an archaeological dig (see Figure 3.2b). Domain experts appreciated the provided combination of a 3D visualisation of objects visible through HMDs with additional contextual information visualised in 2D, as well as the multi-modal interaction in terms of touch-input on a multi-touch table, speech, and 3D hand gestures.

Collaboration between HoloLens and Tablet Users. Chen et al. [18] proposed a collaborative interaction model whereby only one user wears a HMD, while others can observe the HMD user's view on a tablet. Although the observers cannot control the view directly, they can pause their view to add annotations to the scene, which are automatically added to the HMD user's scene.

Summary. In summary, previous works showed great potential for immersive environments for collaboration, whereby AR environments, especially, provide the means to facilitate the natural communication and coordination between users, even if only one user is inside the AR environment. This work therefore chooses an AR environment for the collaborative analysis of health-related data.

3.3 Immersive Visualisation

Although several kinds of visualisations have already been investigated for VR and AR, most projects focus on scientific visualisations[2] and demonstrate that immersive environments can improve the effectiveness [52] and memorability [58] of these visualisations and provide intuitive interfaces, allowing for new analysis approaches [13]. Examples can be found in domains such as brain tumour analysis [90], diagnostic radiology [77], archeology [8, 35, 51, 74, 76, 82], meteorology [91], and geographic information systems [9].

[2]Scientific visualisations can be defined as being 'primarily concerned with the visualization of 3-D+ phenomena [...], where the emphasis is on realistic renderings of volumes, surfaces, illumination sources, and so forth, perhaps with a dynamic (time) component' [33].

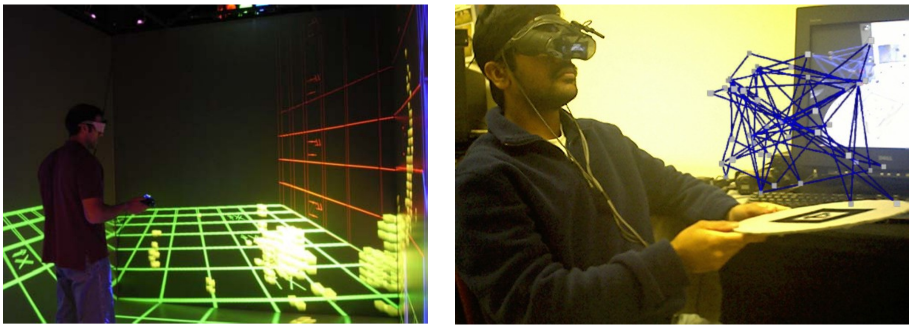

(a) (b)

Figure 3.3: Head-tracking in immersive environments. (a) 3D graph link analysis inside a CAVE environment with head tracking [70]. (b) An AR environment using a HMD with a tangible interface for interacting with the visualisation [5].

However, these visualisations are very domain-specific and their applicability to other domains is limited. Other researchers investigated less domain-specific visualisations like 3D scatter plots [27, 57, 70, 72], link graphs [5, 10, 24, 28] or parallel coordinates [72]. These works report on several benefits of immersive environments for information visualisations.

CaveDataView. Raja et al. [70] developed a 3D scatter plot visualisation for a CAVE[3] environment, assessing the impact on immersion and head tracking (see Figure 3.3a). They identified benefits, such as lower task completion time and higher usefulness ratings, with higher immersion and head tracking when analysing distances, trends, clusters, and outliers.

3D Graph Link Analysis. Ware and Franck [88] showed that depth and motion cues from a stereoscopic environment with head tracking increase spatial comprehension and accuracy when analysing three-dimensional graph links, resulting in three times as much perceived information. Belcher et al. [5] reproduced these results and revealed that a tangible AR interface is well suited for link analysis (see Figure 3.3b).

Network Graph Analysis in CAVE and HMD. Cordeil et al. [24] took a more technical point of view and compared a CAVE environment to HMDs in terms of collaboratively analysing a network graph. Whereas accuracy and affordances for communication did not differ between the two technologies, the HMDs lead to a significantly lower task completion time.

Summary. In summary, previous work showed the suitability of immersive environments for information visualisation and demonstrates that HMDs can already provide similar benefits as extensive CAVE setups. But in terms of visualising abstract data in immersive environment, research is limited to mainly

[3]In this work a CAVE refers to the generic type of Virtual Environment as described by Cruz-Neira et al. [25] and Febretti et al. [30].

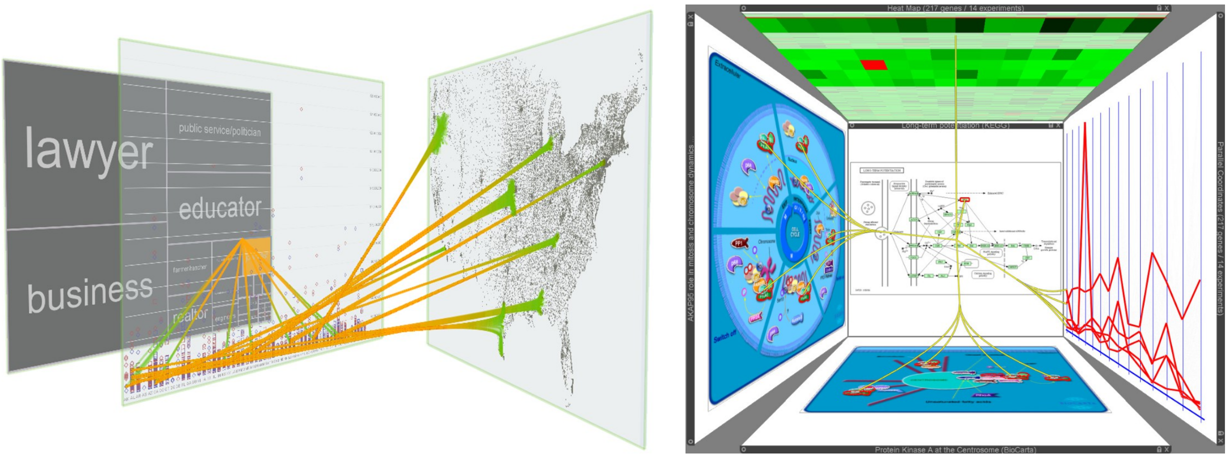

(a) (b)

Figure 3.4: 2D visualisations revealing relations in 3D space. (a) VisLink links related data entries between several 2D visualisation planes [22]. (b) Caleydo's Bucket view arranges several 2D visualisations in a square bucket and shows links between selected data entries [80].

three graph types: 3D scatter plots, link graphs, and parallel coordinates. Immersive environments, however, provide the means to experience alternative graph types or to combine known graph types which draw more benefits from these environments (e.g. 3D parallel coordinates consisting of 2D scatter plots). This work therefore aims at a HMD setup and a combination and extension of existing immersive visualisations.

3.4 3D Visualisations

Although research into immersive visualisations is currently limited, there is a large body of previous work on 3D visualisations for high-dimensional data aimed at desktop systems. This section explores several 3D visualisations that show promise for visualising high-dimensional data in an immersive environment.

VisLink. The *VisLink* [22] visualisation method allows for an interactive exploration of the relation between several 2D visualisations in a 3D environment (see Figure 3.4a). *VisLink* draws multiple different 2D visualisations onto different planes and generates lines between related entries. The resulting 3D visualisation allows for several viewing techniques, such as organising the individual planes like a book, behind each other in a line, or simply in 2D. Users can switch between these views with different hotkeys, explore the 3D navigation freely, or select individual 2D visualisations. The visualisation also allows for different interactions such as zooming, which can filter existing relations and reveal new patterns.

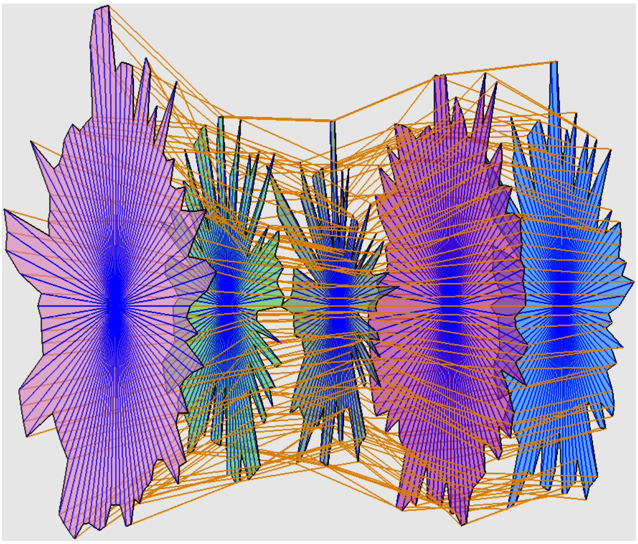

Figure 3.5: 2D parallel coordinates dimensions are unfolded into separate star glyphs, with polylines showing the relation of individual data entries [29].

Caleydo. Similar to *VisLink*, the *Caleydo Visualization Framework* [53, 80] can reveal relations between several 2D visualisations in a 2.5D environment called *Bucket* (see Figure 3.4b). *Bucket* provides a view of several 2D visualisations attached to the walls of a pseudo-3D room, in which relations between the different 2D visualisations can be depicted. The general workflow consists of filtering out irrelevant data, generating relations to other 2D visualisations in the 2.5D view for exploration, then zooming into a specific 2D visualisation for greater detail. Their results show that the *Bucket* visualisation performed significantly better than list-based methods.

Parallel Glyphs. Fanea et al. [29] developed *Parallel Glyphs*, a visualisation combining parallel coordinates with star glyps (see Figure 3.5). For this, a 2D parallel coordinates visualisation is unfolded in 3D into several star glyphs, with each variable of the parallel coordinates visualisation building one star glyph. These star glyphs therefore show the magnitude of the individual lines for their respective dimension. They adopted several interaction methods, such as reordering individual lines to compare them more easily, reordering individual star glyphs, and rotating the visualisation in 3D space.

Summary. In summary, previous research shows that combining 2D visualisations to create a 3D visualisation by linking together related data points can be beneficial for the analysis of multidimensional data. A similar approach can be practical for an immersive 3D visualisation, although interaction techniques have to be reconsidered as they might not be transferable to immersive environments.

3.5 Technical Foundation

Although recent advances in consumer-grade VR and AR technology have made the implementation of mixed reality systems easier, the specific choice of hardware is still important, as different systems offer different advantages. For this purpose, this work considers three contemporary mixed reality implementations and evaluates their usefulness for an immersive, collaborative data analysis setting: *video see-through* and *optical see-through* systems as classified by Milgram and Kishino [62], as well as *CAVE* systems as described by Cruz-Neira et al. [25] and Febretti et al. [30]. As this work aims to use highly immersive technologies, certain choices for mixed reality (e.g. tablet-based magic lenses) are not considered.

3.5.1 Video See-Through

Video see-through devices are classified by their combination of a video-based HMD with front-mounted cameras that project the real world back into the HMD, thus delegating the view of the outside world to the cameras. Due to its relative simplicity, such systems are cheap and easy to implement, especially now that VR HMDs are readily available.

For example, the *AR-Rift* project [78] and the *AR-Arm* project [39] use a VR HMD with front-mounted webcameras to create a video see-through device (see Figure 3.6a). Although such systems are usually bound by cable to reduce latency, this also means that this setup can use powerful desktop systems for generating visualisations. The use of two separate front-mounted cameras means that users still have stereoscopic vision inside the HMD, which allows for better depth perception. Yet, the field of view (FOV) of contemporary VR HMDs is still rather limited, and camera technology is still inferior to the human eye. This setup also introduces some latency between recording the real world and displaying it inside the HMD.

Alternatively, video see-through devices can be realised with mobile devices such as smartphones. Although such a system is usually more lightweight than a VR HMD and not bound to any cables, the limited processing power in these devices reduces the amount of data points that can be visualised. Moreover, as mobile devices only provide one back-facing camera, the view of the real world is usually monoscopic, thus hindering depth perception. The position of this camera further offsets the user's view and, combined with an increased latency for displaying camera images, can easily lead to motion sickness. Thus, projects [73] using mobile devices for immersive environments have focused on VR instead of AR.

3.5.2 Optical See-Through

In contrast to video see-through devices, optical see-through devices do not restrict the user's view of the real world, e.g. by using semi-transparent mirrors to superimpose digital objects. However, building such a device is rather complex, and commercially available hardware is severely limited. Many contemporary projects [2, 18, 40, 55, 75] use the *Microsoft HoloLens*, which has a severely limited FOV for digital content of about 30° horizontal and 17.5° vertical [49] (compared to about 110° vertical and 113° vertical [50] for current VR HMDs).

(a)

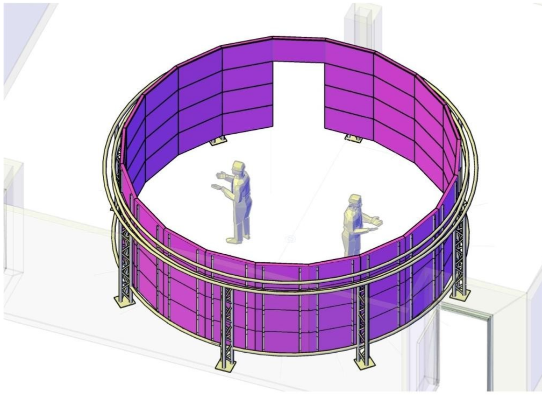

(b)

Figure 3.6: Hardware for immersive environments. (a) The *AR-Rift* project uses a VR HMD with two front-mounted cameras for stereoscopic vision [78]. (b) Concept drawing of a CAVE2 system, where users are fully surrounded by monitors with a thin bezel [30].

Alternatives, such as *Google Glass*, are even more restricted in their FOV for digital content, whereas other projects [68] use more promising devices such as the *Meta2*, which still has a very limited availability at the time of writing. Despite the *HoloLens*' limited FOV for virtual content, it offers stereoscopic vision with no perceivable latency for visual content. The *HoloLens* relies solely on a built-in computer and is therefore not bound by any cables, but offers only limited processing power.

3.5.3 CAVE

Unlike the previous options, a CAVE system does not rely on wearable displays, but immerses the user with the use of several displays or projectors [25, 30] (see Figure 3.6b). Although shutter stereo glasses are needed for proper head-tracking and stereoscopic vision, these glasses are extremely lightweight and do not restrict the user's view of the real world in any significant way. However, this method of stereoscopic vision restricts collaboration, as the system can only track one user at a time (though stereo vision is preserved due to a panoptic stereo mode). In addition, such systems are usually hard to set up and very expensive, both in terms of hardware cost and space requirements. This further restricts the amount of space for collaboration to the size of the CAVE.

3.5.4 Summary

In summary, contemporary mixed reality hardware can already provide a solid foundation for AR environments, despite several flaws. Comparing the different advantages and disadvantages of the available options (see Table 3.1), the most suitable option for a collaborative data analysis scenario is the video see-through display. Although a CAVE environment offers similar benefits as a video see-through HMD, the latter provides better support for collaboration, as the CAVE environment is limited to tracking one user. Furthermore, previous research [24] suggests several advantages of HMD-based systems over CAVE systems. Optical see-through systems conceptually provide a more natural environment, as users can perceive the real world without restrictions, but the available hardware is currently far too limited for the intended use case.

System	Advantages	Disadvantages
Video See-Through	• Room tracking • High computational capacity	• Cablebound • Limited FOV • Latency
Optical See-Through	• Unrestricted view of reality • Room tracking	• Extremely limited FOV for digital content • Limited computational capacity
CAVE	• High FOV • Lightweight wearables • High computational capacity • Unrestricted view of reality	• Space limited to CAVE size • Room tracking limited to one user • Expensive

Table 3.1: Comparison of current immersive mixed reality systems.

Chapter 4

Prototype

This chapter covers the development of the ART prototype, including its features and its implementation. First, Section 4.1 introduces the ART system and its visualisation, by describing available features and their intended use case. An AR system also requires specialised hardware to immerse users, which is presented in Section 4.2. Section 4.3 provides an overview over the software implementation, which was split into four core responsibilities: interaction, visualisation, video see-through, and tracking. Lastly, Section 4.4 describes the calibration process for aligning the real world with the virtual world.

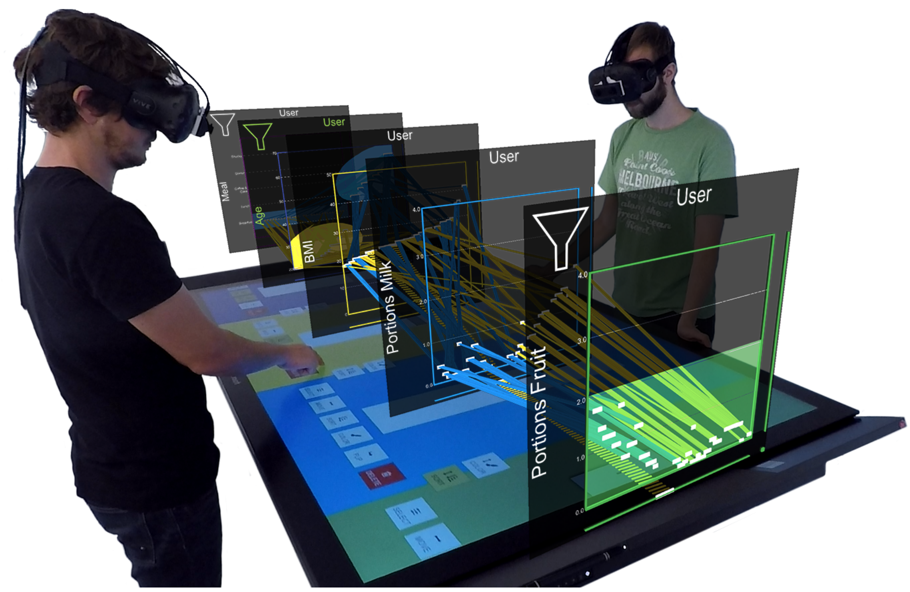

Figure 4.1: The ART prototype uses a VR HMD with front-mounted stereoscopic cameras to visualise 3D parallel coordinates above a touch-enable tabletop.

4.1 Concept

Conceptually, the ART system can be divided into two separate concerns: a 3D *visualisation* displaying data in AR and an *interaction* technique which allows users to customise and navigate the visualisation. The 3D *visualisation* combines a 3D parallel coordinates visualisation with scatterplots, whereby each dimension of a parallel coordinates visualisation is represented by a scatterplot. The visualisation hovers directly above a touch-enabled tabletop, so that users can *interact* with the visualisation.

4.1.1 Visualisation

Similar to previous approaches for non-immersive technologies [22, 29, 53, 80, 86], ART links individual plots to each other to create a multidimensional visualisation. ART consists of 2D scatter plots which are visualised in line at a fixed distance and linked to each other to create a 3D parallel coordinates visualisation (see Figure 4.2). Each data record is represented by a single line cutting through the 2D scatter plots at the corresponding positions. NULL values are represented by dashed lines which cut through a dedicated area below the respective axis in the scatter plot (see Figure 4.2 bottom right). This allows users to trace individual lines through the visualisation even if the entry contains missing values.

This visualisation can facilitate the analysis of clusters, trends, and outliers. Individual scatter plots allow for a familiar identification of clusters and outliers between one or two dimensions: outliers are recognisable by abnormal positions on the scatter plot, while clusters are noticeable by the amount of points in roughly the same position in the scatter plot. Multidimensional clusters and outliers can be identified by similar or abnormal line behaviour between scatter plots, respectively.

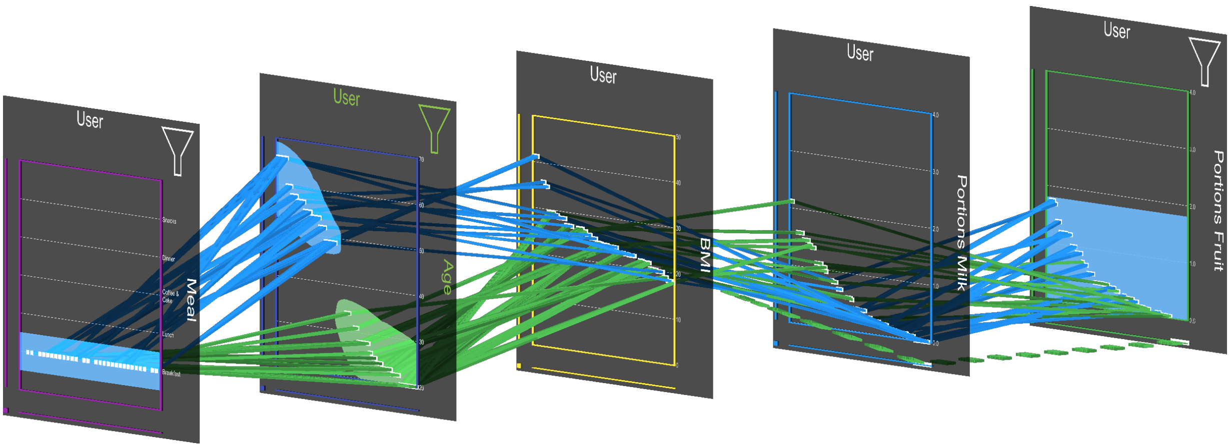

Figure 4.2: ART links individual 2D scatter plots together to create a 3D parallel coordinates visualisation. Dashed lines represent missing values, which are placed in a separate area below their respective axis.

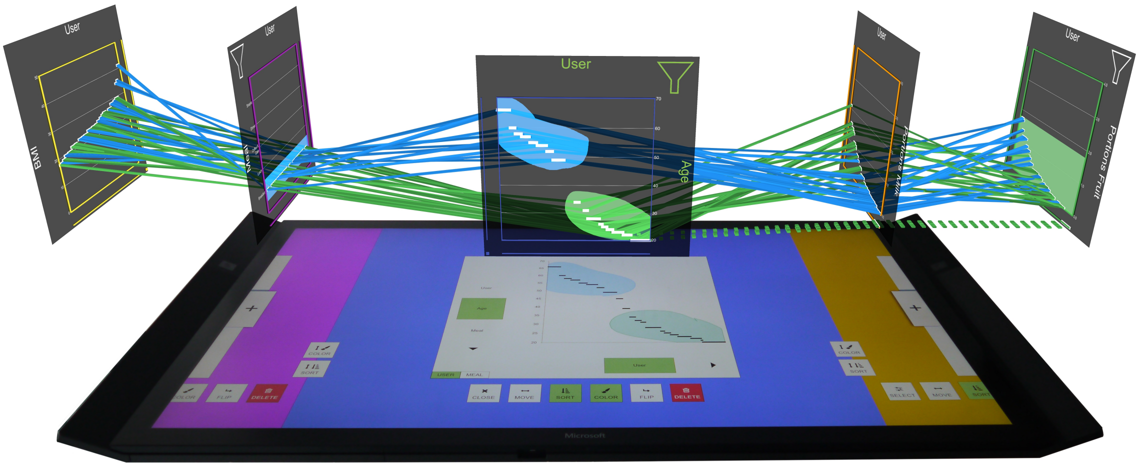

Figure 4.3: Selecting a scatter plot activates the detailed mode. In this mode, users can configure the selected scatter plot, for example by adding filters or changing dimensions. The selected scatter plot is rotated by 90° in AR to face the user.

Furthermore, visualising 3D parallel coordinates in AR provides three advantages over traditional desktop systems:

1. The AR environment provides a large space to visualise the information.
2. The AR environment offers better depth cues and thus simplifies the interpretation of distances.
3. The visualisation benefits from less occlusion (especially during navigation) and therefore facilitates line tracing across multiple scatter plots.

Each scatter plot has a representation in AR and a representation on the tabletop. Both representations are spatially linked to each other so that the AR representation hovers directly above the table representation (see Figure 4.3). Although the number of visible table representations is limited by the table's size, the representations in AR can exceed the size of the table and therefore provides a preview to all created scatter plots.

4.1.2 Interaction

Interactivity is an important part of viewing both 2D and 3D parallel coordinates, particularly for adding and rearranging dimensions. Thereby users can compare two dimensions, filter the data set to avoid clutter, and sort or highlight data records to reveal correlations. In addition to these general operations, ART supports both an egocentric and a non-egocentric navigation style. During egocentric navigation users can move in space to change their point of view. The non-egocentric style allows for navigating the visualisation by scrolling through the list of scatter plots on the table, or using the slider under a scatter plot to move the AR visualisation towards or away from the user.

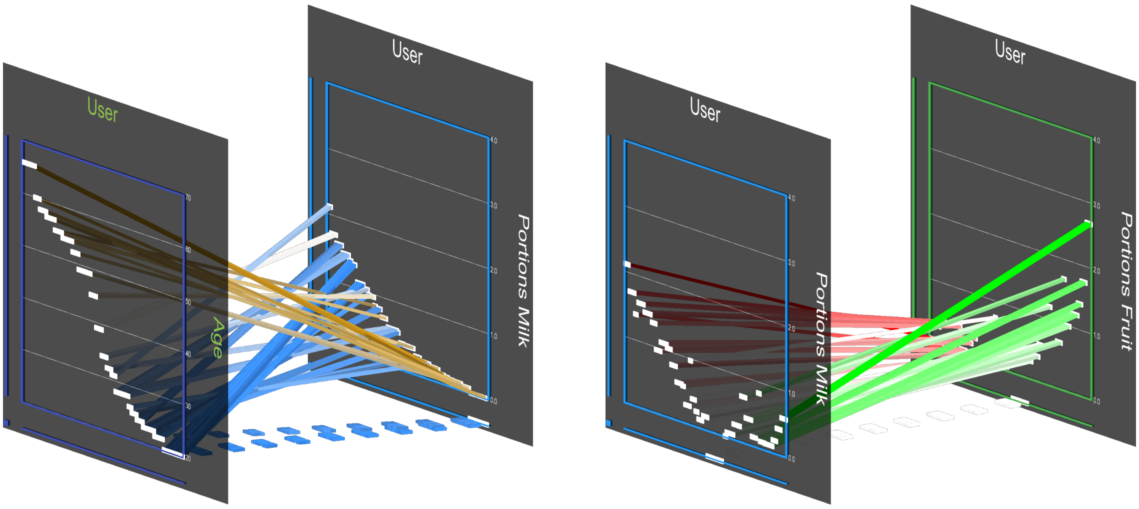

(a) (b)

Figure 4.4: Sorting values on the X-axis of two scatter plots. (a) Values are sorted based on their Y-axis value. (b) Values are sorted based on their horizontal inclination.

The table representations of the scatter plots provide two modes: an overview mode for exploring the visualisation (see Figures 4.2 and 4.5) and a detailed mode to configure the plot (see Figures 4.3 and 4.6). In the overview mode, users can add, remove, reorder, flip, sort, and colourise scatter plots through several buttons at the bottom of the interface.

- **Add**
  New scatter plots can be created by touching the 'plus' button located at either side of the tabletop, and can be directly dragged to the intended position.

- **Move**
  Scatter plots can be moved to another location to facilitate the analysis of relations between neighbouring plots.

- **Delete**
  Scatter plots can also be removed once they are no longer needed.

- **Sort based on absolute value**
  The X-axis of individual plots is sorted by the respective values on the Y-axis. This essentially disregards the X-dimension in favour of an easy to interpret and detailed visualisation of the distribution of the Y-values (see Figure 4.4a). In this mode, the visualisation also shows the rank of the value of a record within the data set (position on the X-axis). Neighbouring lines with the same Y-values appear as one line with a greater width on the X-axis.

- **Sort based on relative differences (inclination)**
  The X-axes of two neighbouring plots are sorted based on their horizontal inclination (differences between the dimensions on the Y-axis). This allows for an easy interpretation of correlations between two neighbouring plots, even when the distribution of Y-values in both plots is huge (see Figure 4.4b).

- **Colourise based on absolute values**
  One plot is selected to colourise the lines in the entire 3D visualisation. This makes tracing individual lines or the identification of correlations over longer distances and across multiple scatter plots easier. The colour of the lines is either set to a predefined gradient based on the Y-axis value, or set depending on the clusters defined in the detailed mode of the scatter plot configuration.

- **Colourise based on relative differences (inclination)**
  The colour of the lines in the visualisation is set by the relative difference of the Y-values between two neighbouring plots. Similar to sorting based on relative differences, this allows for an easy interpretation of correlations, or more specifically, clusters of records with similar correlations but different absolute values (see Figure 4.4b).

- **Flip**
  As certain perspectives, such as looking from above, are difficult to attain in this setup, users can flip individual plots – essentially swapping the X- and Y-axis. This is equivalent to rotating the visualisation by 90°, making a side view equivalent to a top-down view.

The ART system can also optionally replace the sliders in the middle of each scatter plot with a bar chart (see Figure 4.5). This bar chart provides a brief

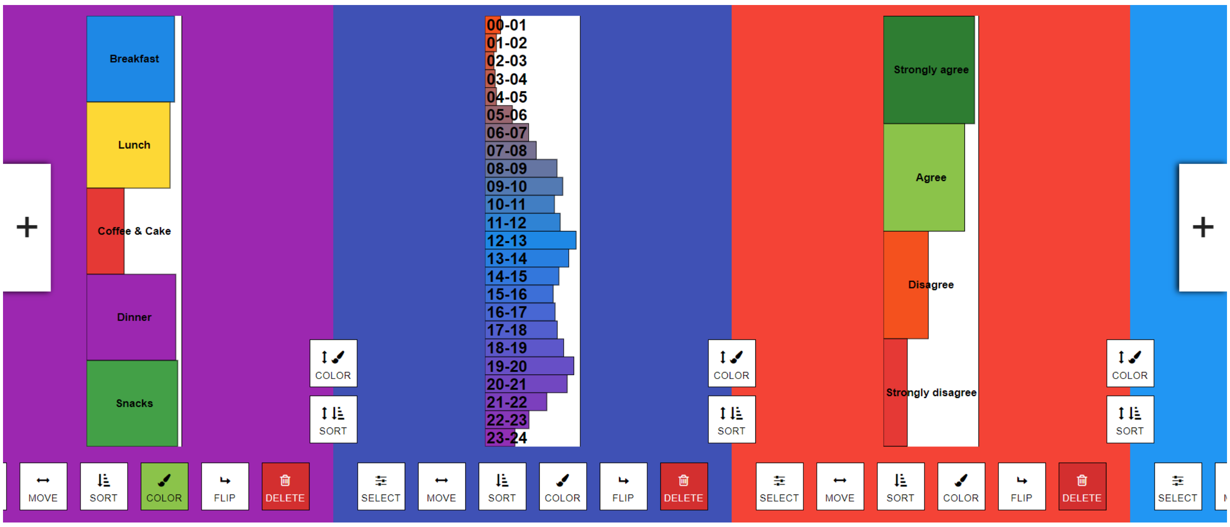

Figure 4.5: In overview mode, the ART system can also display bar charts about the distribution of values on the X-axis in place of the slider.

overview over the value distribution of the X-axis and allows users to quickly create data filters.

Each plot can be selected to open a detailed mode, containing an interactive representation of the scatter plot. Although most functionality from the overview mode is still available during the detailed mode, some functions (i.e. add, move, delete) will revert the system back to overview mode when used. In the detailed mode, the AR representation of the selected plot turns by 90° to provide an orthogonal view on the data and make it easier to mentally link the AR representation to the table representation (see Figure 4.3). The detailed mode provides two additional functionalities: to assign the dimensions of the plot, and to filter or specify clusters in the data set.

- **Assigning dimensions**
  Users can assign a dimension from the data set to the X- and Y-axis by selecting the dimension from a simple scroll list. The scroll list can be filtered depending on the aggregation level of the dimensions (e.g. calories per meal, day, week). The lines in the 3D graph visualisation split up or combine between scatter plots depending on the individual aggregation level (see Figure 4.7).

- **Filters and clusters**
  Filters can be created either by drawing directly into the scatter plot, or by dragging or tapping the relevant range on the axis (e.g. by tapping on the category label 'Breakfast', a filter containing all 'Breakfast' entries is created). Data records that do not belong to a filter are removed from the AR representation, but are still visible in the table representation.

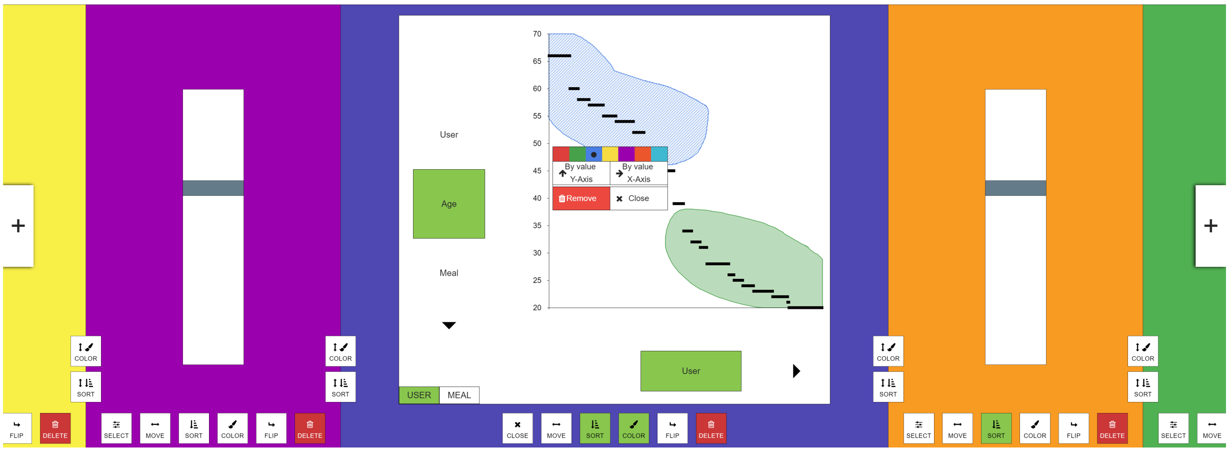

Figure 4.6: The ART tabletop interface. Each column represents one scatter plot in AR. During overview mode, the columns are collapsed (sides), allowing for limited interaction through buttons at the bottom. During detailed mode (centre), users can add filters and change dimensions. Tapping on a filter opens a pop-up menu.

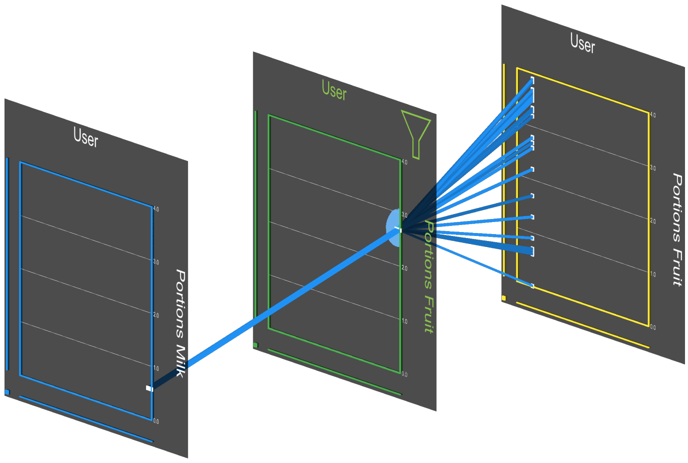

Figure 4.7: Different scatter plot aggregations cause lines to split up or combine. The two scatter plots on the left have a higher aggregation level (user average) than the scatter plot on the right (individual meals), therefore the lines split up from the average to the individual meal values.

Filters also act as cluster classification tool. Each filter is automatically assigned a default colour. If the scatter plot is colourised, lines will assume the colour of their containing filter from this scatter plot. This enables users to assign colours to different clusters, allowing users to track the cluster across the entire visualisation.

Tapping on the filter allows users to change the colour of the filter to one of seven predefined solid colours, or a colour gradient based on the X- or Y-values in the cluster (see Figure 4.6). This menu also offers an option to remove the filter.

4.2 Hardware

ART is based on video see-through AR devices, consisting of VR headsets and front-mounted stereo cameras, as well as a touch-enabled tabletop used for interaction. Every HMD is tethered to its own computer (Intel i7, 32 GB RAM, NVIDIA GTX 1080), which is mounted to the ceiling to maximise cable length. The following sections explain the rationale behind the individual hardware components used for ART, and discuss various technical details.

4.2.1 Displaying Virtuality

For the choice of HMD, the system supports both the *HTC Vive* (see Figure 4.8) and the *Oculus Rift*. However, due to the more complicated tracking setup of the *Oculus Rift* – involving an *Optitrack* system with infrared markers – this work focuses solely on the *HTC Vive*. More details about the *Oculus Rift* support can be found in the project report [46].

The *HTC Vive* [44] HMD provides a resolution of 1080×1200 pixels per eye (2160×1200 pixels total) with a diagonal FOV of 110° at a refresh rate of 90 Hz

Figure 4.8: *HTC Vive* HMD with front-mounted *OvrVision Pro* cameras.

on an OLED display. This HMD provides inaccurate, but low-latency built-in tracking with an inertial measurement unit (IMU) consisting of an accelerometer, gyroscope, and magnetometer. The *HTC Vive* also uses two external stations (*Lighthouses*) positioned in the corners of the tracked room for a more accurate, yet slower tracking (see Figure 4.9). The combination of these two tracking system yields accurate and low-latency 6 degrees of freedom (DOF) tracking, with a trackable area of roughly $15' \times 15'$.

4.2.2 Recording Reality

Due to the choice of a VR HMD, front-mounted cameras are needed to display the real world. The *OvrVision Pro* [89] cameras were specifically designed for both the *HTC Vive* and *Oculus Rift* and are therefore well-suited for this purpose (see Figure 4.8). Due to their small formfactor, the cameras do not add much offset between the user's eyes and the actual sensor viewpoint, resulting in very little difference between wearing the HMD and seeing the real world normally. In addition, the cameras are stereoscopic, and thus they allow for a better depth perception of the real world when wearing the HMD.

The cameras have an effective focal length of 2.5 and provide a maximum horizontal angle of 115° and a maximum vertical angle of 105°, depending on the chosen camera mode. The camera mode determines the resolution, framerate, and FOV, and ranges from high-resolution images at a low framerate (2560 × 1920 pixels per camera at 15 frames per second (FPS)) to low-resolution images at an extremely high framerate (320 × 240 pixels per camera at 120 FPS).

The ART system uses the cameras at a resolution of 960 × 950 pixels per camera (1920 × 950 pixels total) at 60 FPS to roughly match the HMD's refresh rate and resolution. Although this reduces the maximum horizontal and vertical angle to 100° and 98°, respectively, the cameras still provide a comparably high FOV.

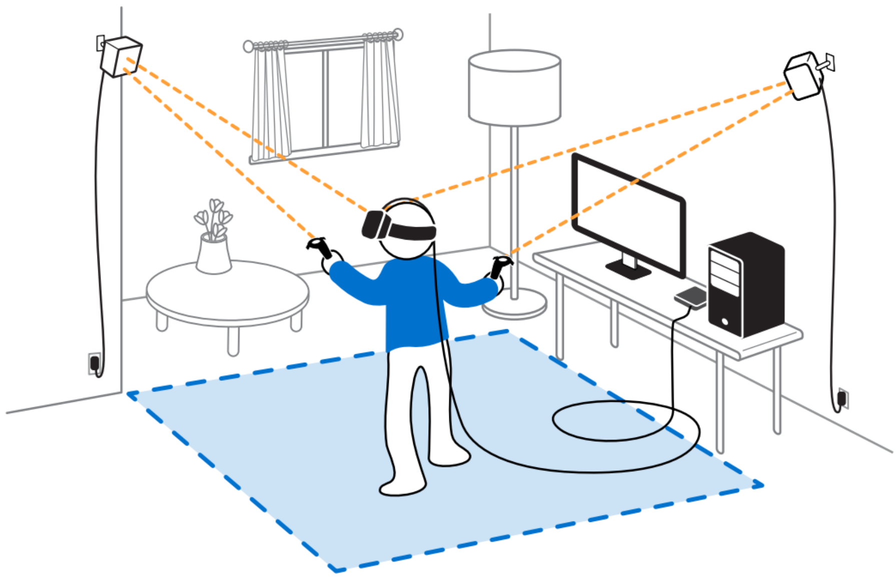

Figure 4.9: Two *Lighthouse* stations track the *HTC Vive* HMD and its controllers inside a predefined tracking area [45].

4.2.3 Interactive Tabletop

Although the ART system supports most multi-touch displays as interactive tabletops, the current setup uses the *Microsoft Surface Hub* [60] (see Figure 4.3). This monitor provides a large display surface with a size of 84″, which can easily accommodate multiple users. The large size along with its resolution of 3840×2140 pixels also helps to overcome the camera's and HMD's limited resolution. The *Surface Hub* comes with a built-in PC (Intel i7, 8 GB RAM, NVIDIA Quadro K2200) and a bespoke version of *Microsoft Windows*.

4.3 Implementation

The recent advances in VR technology, along with the rising popularity of AR for mobile devices make it easy to create mixed reality applications by hiding most technical details. Despite this, creating a multi-device AR system for data visualisation is still a complex task. To handle this level of complexity, the implementation is split into four separate, yet interconnected responsibilities: *interaction*, *visualisation*, *video see-through*, and *tracking* (see Figure 4.10). To make the implementation easier, some responsibilities are split over multiple frameworks and programming languages, or run on different hardware.

1. **Interaction**

   The interaction responsibility provides a touch-enabled interface on the tabletop for user interaction and processes any input from the user. This component is implemented as web application and runs on the interactive tabletop. A web server hosts the necessary web application files and establishes communication between the web application, *Unity3D*, and the SMARTACT database.

2. **Visualisation**

   The visualisation renders the actual 3D data visualisation in AR. The *Unity3D* game engine is used, as it provides tools for creating 3D objects and renders to the HMD natively. This component listens for any updates from the interactive tabletop through the web server and modifies the visualisation accordingly.

3. **Video See-Through**

   The video see-through component is responsible for fetching images from the camera hardware and displaying these images on the HMD, while introducing as little latency as possible. Although images are fetched and processed using a separate library, the images are displayed on the HMD using *Unity3D*. These images are also used to determine the position and rotation of the real cameras, which is passed on to the tracking responsibility.

4. **Tracking**

   To combine the real and virtual world, the system needs to track the HMD's position in relation to the interactive tabletop. The tracking component therefore receives tracking data from both the HMD hardware through *Unity3D*, as well as position estimations using computer vision from the video see-through responsibility.

The following sections will explore the technical implementation of these responsibilities in greater detail. However, given the complexity of the system, this work provides only a brief high-level overview over the system's implementation. A more detailed description is available in the project report [46].

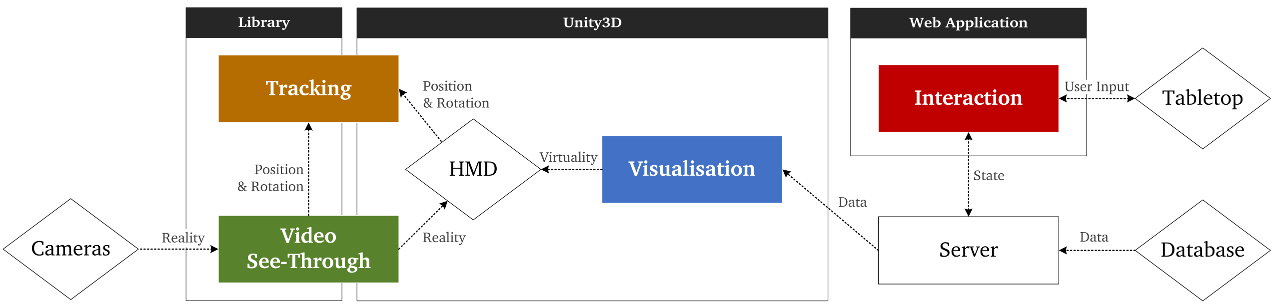

Figure 4.10: A high-level overview of how the different responsibilities (rectangles) are broken up into separate software applications (named boxes), as well as what data is transferred between internal applications and hardware (diamonds).

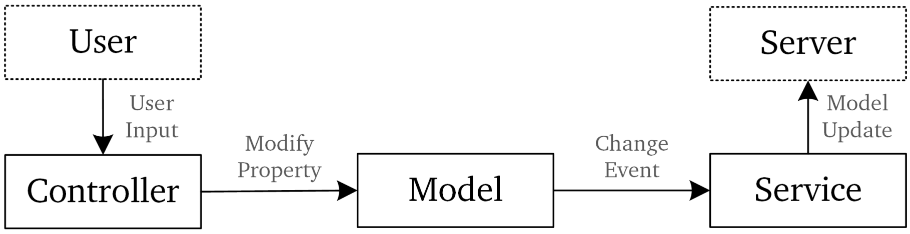

Figure 4.11: Standard event sequence for handling user input on the tabletop. User input is usually processed by a controller, which modifies a property on a model. This triggers a property change event, to which a service reacts by sending a model update message to the web server.

4.3.1 Interaction

Users interact with the system through a touch-capable tabletop, which necessitates some software running on this display. The *Surface Hub* runs a bespoke *Windows* version, making the deployment of custom software difficult. Therefore, a web-based approach was chosen for the tabletop software, which also offers greater compatibility with other systems (e.g. using a tablet instead of a tabletop).

The implementation uses *TypeScript* [61], with HTML and CSS for structuring and designing the visual elements. The *Angular* [36] framework was used, as it provides many features for creating interactive web applications. *Angular* follows a standard model-view-controller pattern, whereby the view is implemented in HTML and CSS with a matching controller and models written in *TypeScript*. In addition, each model in this system implements a property change listener, which triggers an event if any of the model's public properties are changed. This allows for different services to dynamically react to individual model changes. For example, a service listening for model changes sends incremental updates to a web server to synchronise the model with the main application (see Figure 4.11).

The web application uses different models to manage the visualisation:

- **Surface**

  A surface holds metadata about the current tabletop, such as pixel density and resolution of the display, which may be used for calculating the display's physical size.

- **Dimension**

  Dimensions contain data provided by the database, as well as metadata to display the data correctly in a scatter plot (e.g. category names, aggregation level).

- **Graph**

  A graph represents a single scatter plot, and is thus a centre piece of this system. Graphs reference one dimension model per axis and hold metadata such as position on the table and column colour.

- **Filter**

  Filters can be added to any graph model, fading out data not contained within the filter. They are mainly defined by a path (consisting of several 2D points in the scatter plot) and a colour.

Where possible, the web application calculates computationally intensive operations locally to increase performance for the main application, though at the cost of increased network traffic. Communication with the web server is handled with HTTP requests and *WebSockets*.

1. HTTP requests are used for fetching static data, such as the initial application state or data from the database.
2. *WebSockets* establish a real-time channel for incremental updates, such as propagating model changes.

The web server itself is written in *TypeScript* and uses the *Node.js* [67] runtime. The server fulfils three main purposes: *Hosting the web application*, *establishing a connection between web application and Unity3D*, and *distributing data from the database*.

1. **Host web application**

   A web application must be served from a web server, so that the client browser may fetch all relevant data.

2. **Establish communication between web application and Unity3D**

   To connect the web application with *Unity3D*, the web server offers a *WebSocket* connection for the web application, and a TCP connection for *Unity3D*. Data sent through the *WebSocket* connection is forwarded to the TCP connection and vice versa. This architecture therefore allows for multiple *Unity3D* clients (HMD users) and multiple web applications (interactive tabletops).

3. **Distribute data from database**

   Both the *Unity3D* and the tabletop client need data for their visualisation. The web server serves data directly from the SMARTACT MSSQL database, including custom metadata such as colour gradient, data domain, and category names. This metadata has to be created for each dimension in the database in a special mappings file.

4.3.2 Visualisation

The visualisation responsibility is entirely handled within the *Unity3D* game engine. The focus herein is to maintain a stable framerate of 90 FPS, matching the HMD's refresh rate and thus keeping the application responsive. Individual scatter plots consist of a mixture of 2D UI elements (e.g. labels, ticks), primitive shapes (e.g. background, borders), and custom, dynamically generated meshes (see Figure 4.12). A Delaunay triangulation allows the system to convert a concave filter path into an accurate 3D mesh.

For performance reasons, both data points and all lines between two scatter plots use a single mesh each. Furthermore, the actual mesh is generated through

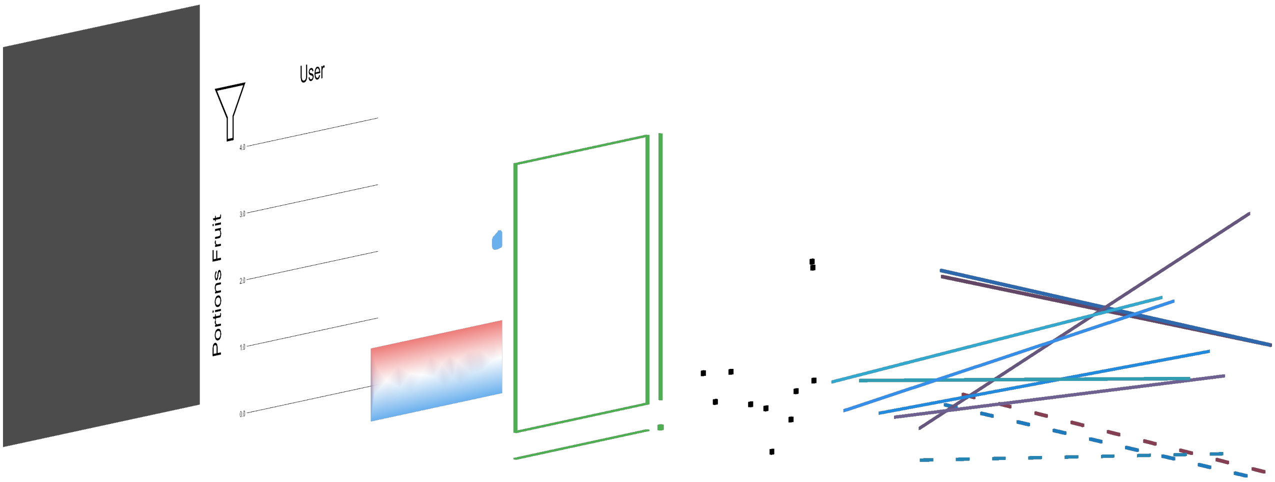

Figure 4.12: Scatter plot composition in ascending render order: Semi-transparent background, UI text elements, filter meshes, border indicators, data values (one mesh), parallel coordinates lines (one mesh).

geometry shaders, trading off CPU load for GPU performance. Lines between two scatter plots use a skinned mesh with both scatter plots acting as bones, which automatically applies appropriate geometry transformations to the mesh on the GPU.

Objects that are too far away from the user are automatically disabled, under the presumption that users cannot see the object anymore. With these methods, the system can support up to 32767 data points per scatter plot, though performance and visualisation readability may suffer when visualising more than several thousand lines.

4.3.3 Video See-Through

The video see-through module is responsible for fetching and processing images from the camera and displaying it on the HMD. Camera images are fetched and processed in a separate library written in C++, then passed to *Unity3D*.

The library uses a pipeline structure for processing camera images. Raw, undistorted images from a camera are passed onto all pipelines, which run in separate threads to maximise performance. Each pipeline can contain multiple processors and outputs, allowing for a modular structure. The system currently uses two pipelines for image processing (see Figure 4.13): One pipeline for delivering the camera image as quickly as possible to the HMD, and one pipeline for detecting AR barcode markers using the *ARToolKit 5* [26], which will be used for tracking purposes.

Different outputs can pass information to *Unity3D*. For example, textures are uploaded to the GPU using a multithreaded *DirectX 11* environment, then synchronised in *Unity3D*'s render thread, reducing rendering time and therefore decreasing latency.

Unity3D uses two sets of cameras for combining reality with virtuality:

1. **Reality**
   One set of cameras renders the real world first, with one camera for each eye. Both cameras have a simple rectangle (quad) directly in front of them, on which a real-world texture is displayed.

2. **Virtuality**
   After the cameras have rendered the real world, a third camera renders the virtual scene, writing the virtual objects on top of the real-world textures. *Unity3D* automatically splits the virtual camera's view into two separate cameras internally, with one camera for each eye.

4.3.4 Tracking

For an immersive AR environment, both the physical and virtual world must be aligned. For this, the system needs to perform three actions:

1. **Calibrate tabletop position**
   In order to correctly place the visualisation above the tabletop, the system has to know the tabletop's position. This is done by using a *HTC Vive Controller* to calibrate each corner of the tabletop once.

2. **Track virtual camera location**
   To correctly provide the users with a view of the virtual world, the system needs to track the HMD's position. For this purpose, the *HTC Vive* provides 6 DOF tracking, which is available in *Unity3D*.

3. **Track real camera location**
   The real camera location must be aligned correctly with the virtual camera's location. Although the cameras are attached to the HMD for tracking purposes, the cameras are positioned at a slight offset in front of the HMD. Thus, an offset has to be applied when tracking the real camera's location.

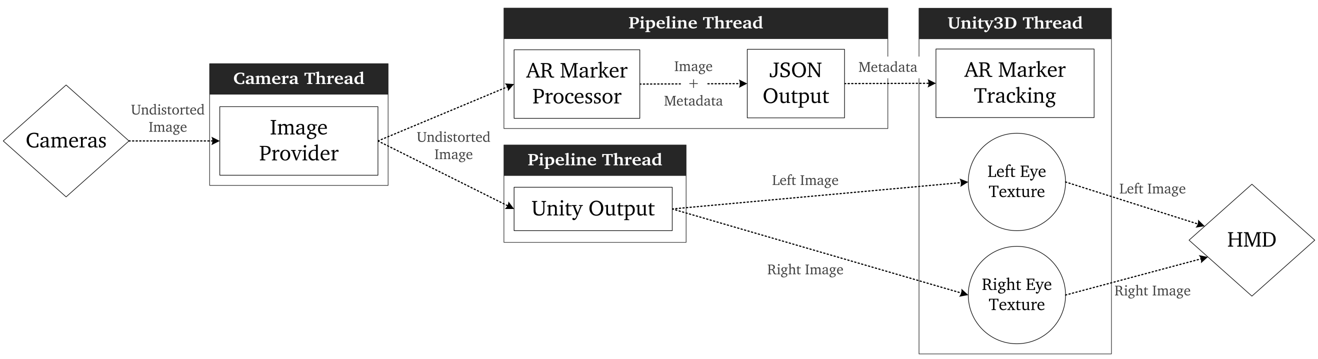

Figure 4.13: Overview over the video see-through structure. An image provider distributes images from the camera hardware (diamonds) to pipeline modules (rectangles), which process the image in different threads (named boxes). Images are written into textures (circles), which are then displayed in the HMD.

4.4 Calibration

For a correct alignment of the real and virtual world, the position and rotation of both the real cameras and the virtual cameras need to be exactly aligned. However, the virtual camera is bound to the HMD's centre, while the lenses of the real cameras are attached with a slight offset at the front of the HMD. In addition, the real camera's rotation may not be completely aligned with the HMD's tracked rotation (e.g. the real cameras may point slightly downwards). Therefore, an offset for both position and rotation has to be applied to compensate for the differences. This can be hard to do manually, and has to be repeated if the camera configuration is changed. Thus, the system uses AR markers to estimate the real camera's position, which allows the system to calculate the offset automatically. For this, the camera needs to be calibrated[1] beforehand, and the marker's physical size and position needs to be known.

Given that the system knows the tabletop's location (due to earlier calibration for placing the visualisation), the marker's exact position and size when displayed on the tabletop can be calculated using the tabletop's resolution and pixel density. The system uses the *ARToolKit 5* framework to detect and extract the camera pose from several AR barcode markers (see Figure 4.14). Although calibration is possible with only one marker, the system generally uses multiple markers to increase accuracy. Once both the real camera's location ($\text{Position}_{\text{Camera}}, \text{Rotation}_{\text{Camera}}$) and the HMD's location ($\text{Position}_{\text{HMD}}, \text{Rotation}_{\text{HMD}}$) are known, the offsets for position Δ_{Position} and rotation Δ_{Rotation} can be calculated given two formulas[2]:

$$\begin{aligned}\Delta_{\text{Position}} &= \text{Rotation}_{\text{HMD}}^{-1} * (\text{Position}_{\text{Camera}} - \text{Position}_{\text{HMD}})\\ \Delta_{\text{Rotation}} &= \text{Rotation}_{\text{HMD}}^{-1} * \text{Rotation}_{\text{Camera}}\end{aligned}$$

To avoid inaccuracies, the HMD should not be moved during calibration, as varying latencies between tracking system and camera images can yield incorrect results. For this reason, the system will only perform this calibration if no HMD movement is detected. Likewise, the offsets are measured over several seconds, then calculated using an average value of all measured offsets. This eliminates most outliers, making the offset more accurate.

Once both Δ_{Position} and Δ_{Rotation} are known, the tracking system can simply apply the static offsets to the current HMD tracking data to correctly align the virtual world with the real world:

$$\begin{aligned}\text{Position}_{\text{Camera}} &= \text{Position}_{\text{HMD}} + \text{Rotation}_{\text{HMD}} * \Delta_{\text{Position}}\\ \text{Rotation}_{\text{Camera}} &= \text{Rotation}_{\text{HMD}} * \Delta_{\text{Rotation}}\end{aligned}$$

[1] This calibration determines, for example, the intrinsic camera parameters, and has to be performed only once per camera.

[2] In this context, positions are stored as three-dimensional vectors, rotations as quaternions.

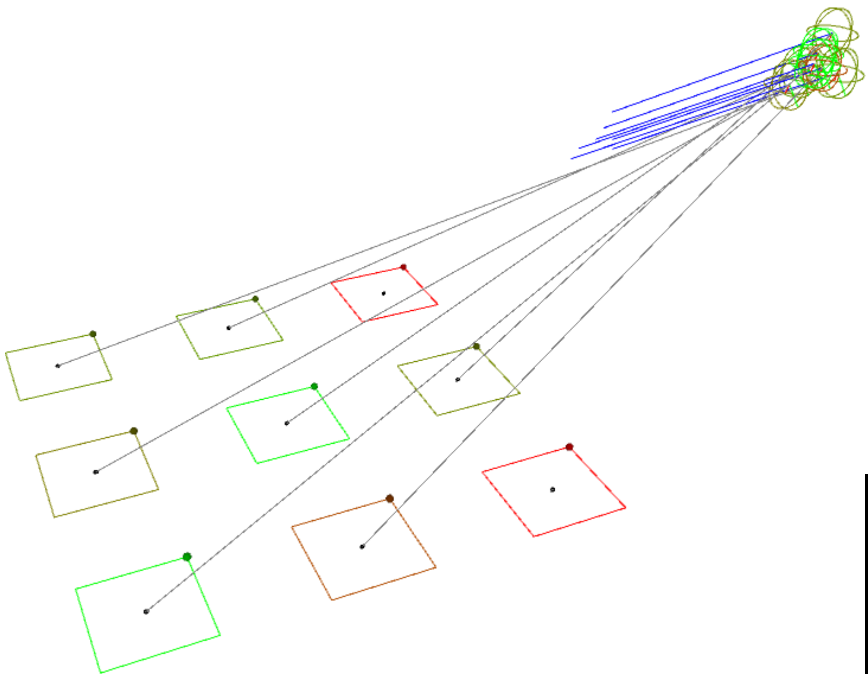

(a)

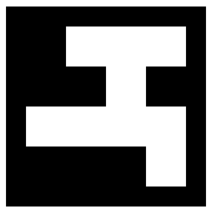

(b)

Figure 4.14: Camera pose estimation using AR markers. (a) Visual representation of camera pose estimation (spheres) using AR markers (squares) in *Unity3D*. The marker colour varies from red to green according to a confidence value. Blue lines indicate the approximate camera rotation. (b) A 4×4 *BCH_13_9_3* barcode marker for use with the *ARToolKit 5* [26].

Chapter 5

Evaluation

To evaluate ART system's usefulness in an analysis workflow for the target use case, an expert walkthrough was conducted with domain experts from the SMARTACT team. Participants could assess the system and simulate a real analysis scenario. Section 5.1 provides an overview over the expert walkthrough setting. Feedback from observations and discussions is presented in Section 5.2, together with design recommendations and further research directions for immersive data analysis tools. Furthermore, there still are some limitations regarding both the evaluation and the technical setting, which are discussed in Section 5.3. Lastly, Section 5.4 reviews the initial design requirements based on the first focus group, and discusses to what extent ART can answer the initial research questions.

5.1 Expert Walkthrough

To evaluate the ART system prototype, two group-based expert walkthroughs with domain experts (health psychologists and biological psychologists) were conducted. Ten domain experts participated in the walkthrough session (five per session) and each session took approximately two hours. Seven domain experts had already participated in the initial focus group to identify requirements and leverage points. The sessions started with a recap of the previous requirements focus group (15 min) followed by an introduction to the ART system (15 min).

In the subsequent discussion (60 min), the three analysis aims that were identified through the first focus group (identification of clusters with similar behaviour patterns, analysis of high-dimensional data on multiple aggregation levels, and analysis within chronological trends) were performed using ART. In a first use case, domain experts explored the relation between food consumption (e.g. portions of milk, meat, grains) and clinical measures (e.g. BMI). They further investigated different aggregation levels (e.g. meal, day, user level) and tried to identify clusters of people with similar behaviour signatures.

In a second use case, domain experts analysed the effectiveness of the applied intervention over the study's duration. Domain experts created a timeline visualisation with which they were able to track different measures over several weeks. Each scatter plot was assigned a single week, so that the lines showed the changes over time.

During the session, two experts analysed the data, while others observed the analysis process on two large screens situated in the same room, showing the current actor's augmented view. Domain experts rotated from observers to actors approximately every 10 min. This actual analysis walkthrough was concluded with a group discussion in which the leverage points – which were identified in the first focus group – were discussed (30 min).

5.2 Results

The results of the group-based expert walkthrough are structured based on the four leverage points identified in the initial focus group: *Approach to Analysis*, *Visualisation Readability*, *Space & Immersion*, and *Collaboration*. In addition, each point contains design recommendations and further research directions for immersive data visualisations, which are based on observations of and feedback from the domain experts.

5.2.1 Approach to Analysis

The domain experts mentioned that it was easy to familiarise oneself with ART as well as to quickly identify the operations required to follow their analysis approach.

> *'I found it stunningly easy to get into the workflow [...] somehow, everything was totally plausible.'* [G1/P4]

To limit complexity, most often all scatter plots were configured with the *UserId* dimension on the *X*-axis. This allowed for the creation of a visualisation in which each depth value represented a single user. Domain experts constantly added scatter plots either to investigate relations between dimensions or solely to filter the data set (e.g. plot with gender on the *X*-axis and age on the *Y*-axis). To organise the visualisation, plots that were intended to filter the data set were placed in the leftmost position of the AR visualisation (see Figure 5.2b). Plots which did not reveal any findings were instantly reconfigured or removed.

Domain experts further changed aggregation levels during the analysis. If interesting effects were identified on a higher aggregation level (e.g. participant level), they added plots with a lower aggregation level (e.g. day level) to conduct a more detailed analysis. For example, these detailed analyses were performed to decide whether a data record is an outlier. These plots were often removed afterwards to continue with the higher aggregation level. During the analysis, domain experts dynamically created clusters, and colourised and sorted the data records. They judged these functionalities as essential for an efficient analysis of the data set.

Domain experts agreed that ART supports an explorative analysis workflow in which findings can be fluently investigated in more detail without discarding the previous analysis.

> *'The tool allows performing quick and easy actions. Thus, the dynamic somehow remains in the workflow.'* [G1/P3]

> *'If you see something interesting you can directly investigate it in more detail. If you have a look at relations between multiple variables with other tools it instantly gets very complicated.'* [G2/P6]

However, ART currently only supports a linear analysis workflow. Domain experts wanted to be able to save snapshots of the analysis state at hand. This would allow to open up new analysis branches without losing previous ones. The AR environment seems to be well-suited for this, because snapshots could be laid out in physical space, relations between the snapshots could be visualised, and they could be accessed easily to continue the analysis at a previous state.

Design recommendations

Support fluent workflows
Extensible visualisations in combination with an easy to learn and fluent way to configure the visualisations can provide the means to dynamically analyse data.

Further research directions

Support of non-linear analysis workflow
(e.g. snapshots to allow for new analysis branches)

5.2.2 Visualisation Readability

Domain experts gave positive feedback about the visualisation's readability. They mentioned that the 3D parallel coordinates are well suited to identify multidimensional trends and relations, and to visualise timelines.

> *'A pretty cool thing, because I can look at the data in a different way. You get a better understanding of trends, and you can see more than one relation at the same time.'* [G2/P8]

Domain experts also mentioned that ART allowed them to identify multidimensional clusters of persons with similar behaviour signatures. Another benefit identified by the domain experts is that multidimensional outliers are easy to detect. Outliers are not only visible if their value is quite different to the rest of the distribution, but also if the relation to other dimensions is different.

> *'If the data point is out of line, you can see it easily. For many data points, it is interesting to see if a data point has a different inclination than all others.'* [G1/P1]

Domain experts made frequent use of the sort, colourise, and filter functionalities (see Figure 5.2). The experts sorted most of the scatter plots to reduce the complexity, as the lines would otherwise create a lot of clutter, making the interpretation more difficult. The colourise functionality was especially used to compare relations over multiple – but not necessarily neighbouring – scatter

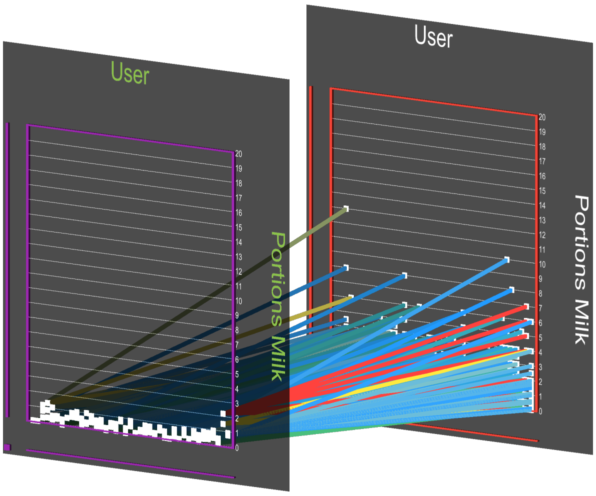

(a) Absolute scale

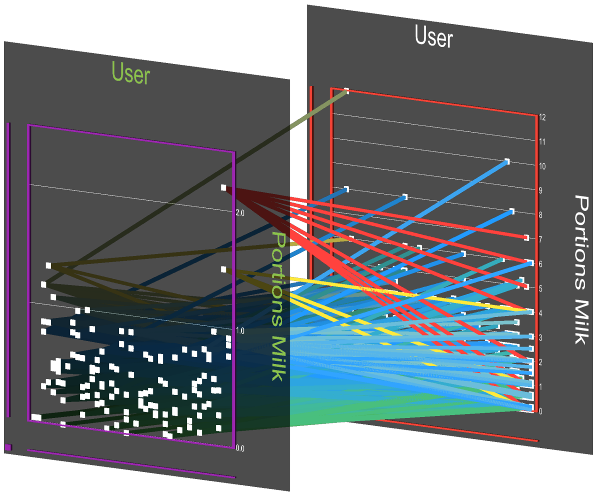

(b) Relative scale

Figure 5.1: Absolute and relative scale used for the Y-axis. The left scatter plot has a higher aggregation level (phase average), the right scatter plot has no aggregation (individual meals). (a) Using an absolute scale (0–20) allows for comparisons between plots, but makes values on the scatter plot itself indistinguishable. (b) Using a relative scale (0–2.5 left, 0–12 right) for each scatter plot, the values become distinguishable, but makes line comparisons misleading.

plots, and to highlight created clusters. Also, the visualisation of relative differences between data points on two scatter plots in terms of sorting or colouring by difference was assumed to be an important feature for the analysis.

One problem that the experts mentioned was related to the fixed scales of the scatter plots axes (see Figure 5.1). A fixed scale facilitates the comparison between plots, but the distribution of the data records span only a small part of the scale, making it difficult to see smaller differences between the data points. An additional function to either automatically apply a suitable scaling based on the currently filtered data over multiple scatter plots, or to manually adapt the scale (either globally for all scatter plots or locally for individual scatter plots) is required.

Another difficulty occurred during the comparison of clusters with different numbers of data records. The domain experts therefore recommended the integration of visual representations for descriptive statistics.

> *'You should include an additional line showing the average. This line could be thicker or smaller depending on the number of lines within the cluster.'* [G1/P2]

In addition to visualising more abstract information, the domain experts also recommended integrating non-abstract information like images of the single meals. This would allow the analysts to get an even better understanding of the data.

Design recommendations

Provide functionality to reduce data complexity
Participants made frequent use of the sorting functionality, indicating that the lines were too cluttered and therefore too complex.

Provide colourise functionality
Colourising the data set is essential to compare values over larger distances, or follow individual data records in the visualisation.

Highlight relative differences
For the analysis of multidimensional relations, users need to investigate both absolute values as well as relative differences between dimensions.

Further research directions

Integration of descriptive statistics for clusters
(e.g. clusters' average line)

Integration of additional non-abstract information
(e.g. pictures of meals)

Integration of statistical significance tests
(e.g. significant correlations or differences of means)

5.2.3 Space & Immersion

The domain experts perceived the immersive technology as valuable for data analysis. They reported getting a better feeling for the data compared to traditional desktop tools:

> *'I think this gives you a different feel for the data. [...] One has a feeling faster of what is in there or how they behave.'* [G1/P2]

The experts further appreciated the large space that was available to visualise information. This allowed them to visualise a large amount of data simultaneously.

> *'I found it really great that you can really see all the data at once for each person, because otherwise we are not able to.'* [G2/P6]

Participants did not actually see all the information at once, but perceived it as laying in physical space and therefore being available all the time. In terms of the used AR devices the domain experts stated that although the see-through functionality lowers the immersion compared to a VR environment, they would prefer an AR environment because of three reasons:

1. The orientation in the room is facilitated.
2. Co-located collaboration is supported, as users can still see each other naturally.

3. The analysis could be better integrated in their holistic workflow and daily working routine.

The domain experts further reported that not only the visualisation, but also the familiar and fluent interaction, which allows for a dynamic adaption of the visualised information, increased their feeling of being immersed in the data.

> *'You can change things so easily, so you can really just move [the visualisation] back and forth, or somehow choose another type of aggregation, and thereby you can immediately solve [the question you have]; otherwise it takes an eternity and one is out of the actual process already, here it's somehow done quickly and then you can continue with what you actually aim for.'* [G1/P2]

The mapping between the table and the AR visualisation was perceived as easily understandable. Some difficulties occurred during the creation of filters and clusters. To create a cluster, users opened the detailed view on the table. During that time, the AR scatter plot rotated to match the orientation of the detailed view. As users focused on the table they did not observe this rotation and therefore had difficulties identifying the areas of interest in the rotated scatter plot. In general, however, the domain experts judged the interaction on the tabletop as being familiar:

> *'The touch interaction feels very natural and easy and provides instant feedback.'* [G2/P2]

In terms of navigation, the domain experts stated that egocentric navigation facilitates the interpretation of the multidimensional information:

> *'You can change the perspective to see the correlation between the dimensions.'* [G1/P5]

However, due to the high immersion, users had the urge to additionally navigate the 3D visualisation through gestures:

> *'While moving and turning you have the urge to directly grasp the visualisation to interact. But things like selection [dimensions] and the filtering is fine on the table.'* [G1/P2]

In addition, navigating the visualisation via scrolling on the tabletop was frequently used. Although the experts liked the combination of the two navigation styles and had no difficulties to combine them, they suggested to offer additional gestural input to navigate the chart (e.g. to vertically flip or horizontally turn a scatter plot, or to scroll through the scatter plots with a gesture) or simple voice commands for selecting dimensions.

Design recommendations

Combine visualisations in AR and interaction on touch-enabled devices
AR environments provide depth cues and large spaces to visualise information. Touch displays allow for a familiar and fluent interaction, which is required for dynamic operation. The combination can create the feeling of being immersed in the data, while still allowing for natural communication with other people.

Be aware of changes outside the user's view
Special attention has to be put to guiding users' attention between the visualisation in AR and the configuration work on the touch devices. Interacting with the table can temporarily limit the user's focus, resulting in unnoticed changes to the AR environment.

Further research directions

Allow for navigating the AR space through gestures
(see e.g. [15] for a combination of input styles in AR environments)

Comparison of immersive interaction methods
(e.g. voice commands, gestures, touch interaction, and combinations thereof)

Experimental comparison of AR and VR environments for collaborative data analysis scenarios
Although AR offers a more natural collaboration, VR can offer further functionalities that could increase collaboration.

5.2.4 Collaboration

The domain experts perceived ART as being suitable for collaboration. They mentioned that with classic desktop applications, one collaborator performs the actions whereas the other collaborator only observes. The possibility for both collaborators to access the tabletop as well as the possibility to select the individual points of view gave the domain experts the feeling of being an active part of the analysis process. This collaboration provided three advantages:

1. It fostered discussion.
2. It is helpful to explain findings to others, hence it helps to get a shared understanding of the data.
3. Mistakes during the analysis process or in the data can be identified.

During the analysis, participants switched between phases of tightly-coupled collaboration and phases of loosely-coupled collaboration (see Figure 5.2). During tightly-coupled collaboration the domain experts configured the visualisation and discussed findings or next steps. They often tried to get the same

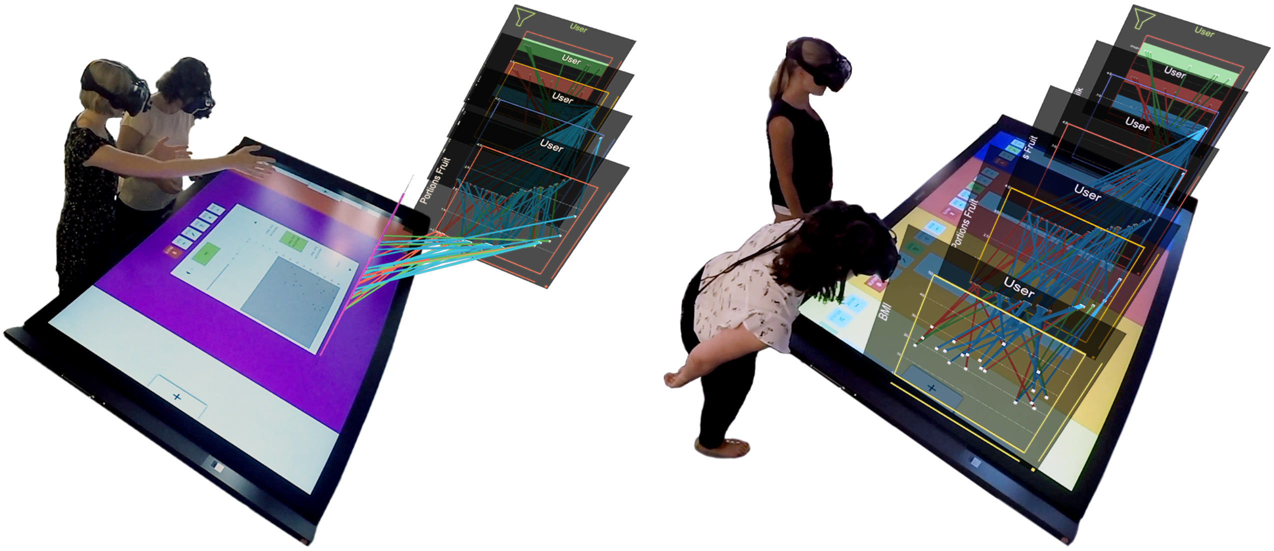

(a) Tightly-coupled collaboration (b) Loosely-coupled collaboration

Figure 5.2: Participants switched between phases of tightly- and loosely-coupled collaboration. (a) During tightly-coupled collaboration, participants worked closely together to configure the visualisation and discuss findings by pointing. (b) During loosely-coupled collaboration, participants explored the visualisation individually.

viewpoint, and made heavy use of gestures (e.g. pointing). In phases of loosely-coupled collaboration participants mainly tried to get an initial understanding of the currently visualised information. For this, they individually navigated the visualisation egocentrically, often walking around the table to get a different view. However, the support of loosely-coupled collaboration was limited, as reconfiguring the visualisation would also change the collaborator's view. Further support for loosely-coupled collaboration is required to overcome this limitation. AR environments allow to visualise distinct information to the collaborators, if required. Still, the operation on the table also has to be adapted to the collaboration style.

In terms of navigation styles, non-egocentric navigation was applied more often during tightly-coupled collaboration. It was preferred, because it guides users' attention to the same part of the visualisation (see Figure 5.2a). For a loosely-coupled collaboration, the egocentric navigation was preferred, because the points of view can be chosen individually (see Figure 5.2b), allowing for individual exploration. During the evaluation, the experts often applied these two methods in the corresponding situations, because they were aware of the social protocols, and tried not to negatively influence the collaborator's environment.

During tightly-coupled collaboration, participants applied two strategies: they either tried to take similar points of view on the visualisation to discuss the information, or they took different points of view (e.g. one collaborator on each side of the table) to combine their individual findings. In both situations, they frequently applied deictic gestures to make spatial references, but faced difficulties because the virtual content occluded their hands. Therefore, precise pointing was hindered and it was difficult to see where exactly the collaborator was pointing at. A 3D registration of the hand, or a virtual pointing and highlighting functionality was requested by the experts and could help overcome this limitation.

Although participants not wearing a HMD could not actively explore the visualisation by themselves, they were still an active part during discussions, as they were able to see the HMD user's perspective through a monitor. They often requested a specific viewpoint from the HMD-wearing domain experts, or wanted to view the visualisation themselves:

> *'Could you hold still for a moment?'* [G1/P3]

Allowing for a separate navigation for these observers (without wearing a HMD) can further foster discussions, for example by showing the visualisation on a tablet.

Design recommendations

Combine navigation styles
Egocentric navigation has benefits during loosely-coupled collaboration as it allows collaborators to select individual points of view. Non-egocentric navigation has benefits during tightly-coupled navigation as it helps to set the same focus to the collaborators.

Allow for precise pointing / selection of individual data entries
This can quickly guide the discussion towards an outlier or a cluster without having to match points of view to get an understanding what entry is being discussed.

Allow non-HMD observers to take a more active role
Using HMD and egocentric navigation can be inconvient for some users, but excludes them from the analysis approach. By offering a view from alternative devices such as tablets (e.g. similar to Chen et al. [18]), these users can explore, annotate and discuss the visualisation with the fully immersed users.

Further research directions

Extended support for loosely-coupled collaboration
(e.g. individual AR visualisations that allow for reconfiguration and navigation)

Extended support for tightly-coupled collaboration
(3D registration [34] and pointing)

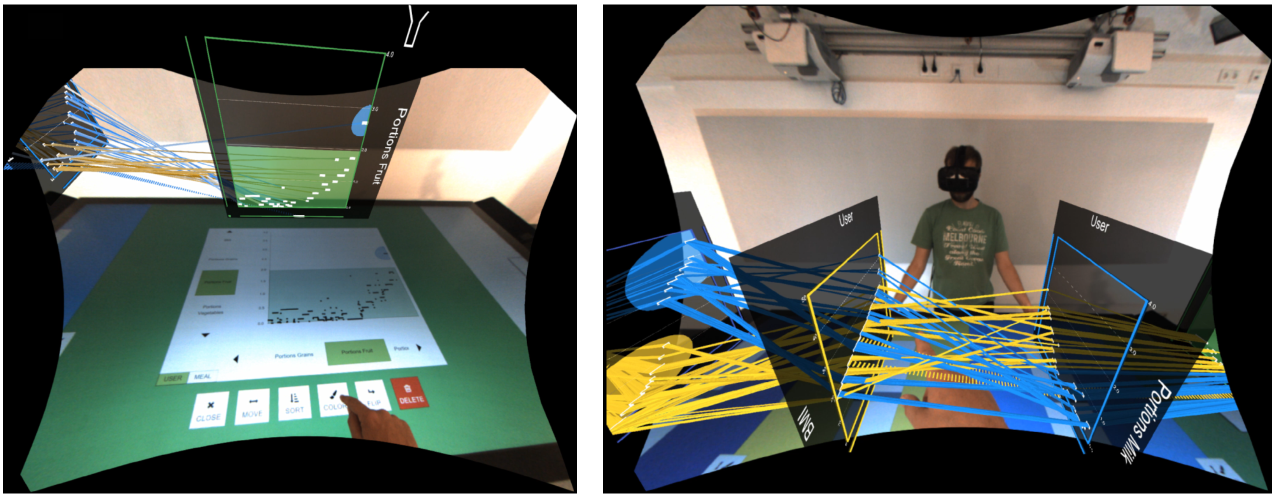

(a) (b)

Figure 5.3: User view from inside a HMD. Distortions at the edges of the images are due to correcting the radial lens distortions. (a) The cameras can distort the tabletop's contrast, resulting in text that may be hard to read from certain angles. (b) Collaboration with other HMD users is possible, though limited to voice and gestures, as the HMD blocks most facial expressions.

5.3 Limitations

The ART system comes with several limitations concerning evaluation, hardware, and software.

5.3.1 Evaluation Limitations

While the expert evaluation provided valuable insights for the design of immersive data analysis tools, the evaluation did not reveal any measurable statistics. For example, no statement can be made about the speed or quality of evaluation when comparing the ART system to similar desktop systems. Further user studies are necessary to compare such measures and to uncover other possible advantages (or shortcomings) of using the ART system.

5.3.2 Hardware Limitations

The choice of using a video see-through over an optical see-through device comes with several benefits, such as a larger FOV for virtual objects, but also has several limitations. The biggest limitation is the reduced image quality of the video see-through functionality like the relatively low resolution, the low contrast and the deviation from true, real-world colours (see Figure 5.3). Despite the use of large, high-contrast text and icons, some participants were unable to read the UI elements on the tabletop, and had to take off the HMD for interacting with the table (see Figure 5.4).

(a) (b)

Figure 5.4: Due to the limited camera quality, some users had to take off the HMD when interacting with the table. (a) User on the left takes off HMD to interact with table. (b) User on the left wears HMD again to look at the AR visualisation.

Initially, the offset between the user's eyes and the camera position caused some uncertainty in terms of navigating the physical space. Especially during collaboration, users reported missing their innate feeling of where the other person is when wearing the HMD. This effect subsided quickly, but did not vanish completely. The experts also stated that if the camera vision would be better, they would probably forget about the fact that it is a video see-through device.

Furthermore, the video see-through setting causes some latency between recording and displaying the real world. Although keeping this latency as low as possible was of high priority in the ART system, there is still a chance that it can cause motion sickness for sensitive people. During the expert walkthrough, however, none of the domain experts were affected by motion sickness – rather, some experts reported slight motion sickness when viewing the augmented view through the displays, which subsided once they put on the HMD themselves. One reason could be that the table served as a kind of anchor point.

Another limitation is that the HMDs are tethered. During the evaluation, a long cable allowed users to walk around the table. Still, the wired connection negatively influenced users' perception of being able to move freely in space. The ceiling-mounted computers alleviated this problem to some degree, though users still had to take care of the cable during extensive egocentrical navigation.

The technical setting also had an influence on collaboration, as the HMDs hindered the domain experts from seeing each other's faces. This caused them to focus on the visualisation and not on the collaborator during discussions.

5.3.3 Software Limitations

Given the scope of this project, there are also many software limitations that could not be implemented. Although one of the requirements for an AR application is the correct 3D registration of objects, the current ART system fails to fully implement such a registration, as the system always renders the visualisation above the real world. However, support for partial 3D registration is possible, since the system can use the stereoscopic cameras to extract a rough depth map of the real world. This missing 3D occlusion affected participants negatively, as it hindered them when pointing at specific data values.

The inflexibility in terms of the available data set is another limitation. For every dimension in the data set, a separate mapping must be created in the web server, containing the dimension's name, domain, and other metadata. This restricted participants to predefined dimensions, even though different dimensions were requested.

5.4 Discussion

Despite several limitations, the ART system received very positive feedback, thus demonstrating the usefulness of immersive data analysis tools for the target use case. The evaluation also provided insight into the research questions, and fulfilled most of the initial requirements from the focus group in terms of *Approach to Analysis*, *Visualisation Readability*, *Space & Immersion*, and *Collaboration* (see Table 5.1).

Research questions.

[RQ.1] *How can immersive visualisations help overcome current limitations of traditional data analysis approaches?*

Using an immersive visualisation provided various benefits over traditional desktop-based visualisations: For example, the visualisation readability was improved, as better depth perception and egocentric navigation facilitated the interpretation of three-dimensional data. This also allowed for integrating statistics into the visualisation (e.g. sorting showed a value distribution), or allowed for quicker investigation into details (e.g. by switching aggregation levels between scatter plots). However, these immersive visualisations should also offer functions to reduce complexity, as too many lines or too much clutter can render the visualisation unreadable.

[RQ.2] *In what ways can an immersive, touch-based interaction facilitate the interaction with 3D visualisations?*

The touch-enabled tabletop allowed for a familiar interaction with the visualisation. Although participants wanted to perform certain actions (e.g. rotating a scatter plot) through gestures, the touch interaction was appreciated during more precise configuration tasks (e.g. creating a filter). The touch interaction avoids the touch the void issue, which is still a problem with current spatial controllers. A combination between several immersive interaction techniques (e.g. voice commands, spatial controllers, touch input) might offer additional benefits in scenarios where one of the methods is inadequate (e.g. selecting individual lines can be difficult with touch and voice commands, but intuitive with a spatial controller).

Approach to Analysis. Requirements relating to the approach to analysis have been completely fulfilled. The ART system supports a dynamic workflow ([A.1]) as users can quickly add new scatter plots, therefore investigate correlations to other dimensions. However, the system could further support this, as the current workflow is still fairly linear (e.g. by adding the ability to save current snapshots of the visualisation and return to them later on). The combination of scatter plots with parallel coordinates, and the ability to choose between different aggregation level allows users to switch between inter- and intra-individual analysis ([A.2]). Reducing the data set ([A.3]) is possible through the filter function. Lastly, the touch-enabled tabletop provides an easy, intuitive interface for non-experts ([A.4]).

Visualisation Readability. Concerning the visualisation readability, the ART system can fulfil all but one requirement. The scatter plots can display abstract data ([V.2]) in a familiar way, while the 3D parallel coordinates allow for an arbitrary number of dimensions ([V.1]) by adding as many scatter plots as needed. Although the visualisation could conceptually support non-standard data such as displaying photos for each entry ([V.3]), this is currently not implemented in the ART system. The visualisation does, however, support the identification of clusters, trends, and outliers ([V.5]), while the filter and colourisation tool allows users to highlight clusters ([V.4]).

Space & Immersion. Despite some technical limitations, the ART system fulfils almost every requirement for space and immersion. The HMD hardware provides high immersion ([SI.1]), and the AR environment allows for a rather large space that can be used for analysis ([SI.2]). This also fulfils two of the three key characteristics for AR environments, as it combines real and virtual ([SI.3]) and is interactive in real time ([SI.4]). Although the visualisation is anchored to the tabletop and is therefore partially registered in 3D ([SI.5]), the system does not provide any occlusion yet, meaning that the visualisation is always rendered above the real world.

Collaboration. In terms of collaboration, the ART system is a clear improvement over the target user's previous analysis approaches. However, while the system allows for multiple active HMD users ([C.1]), observers without HMD are still in a passive position. Moreover, the current system does not sufficiently support active collaboration, nor does it offer collaborative tools ([C.2]).

Status	Identifier	Title
		Approach to Analysis
✓	[A.1]	Support dynamic workflow
✓	[A.2]	Switch between inter- and intra-individual analysis
✓	[A.3]	Reducing the data set
✓	[A.4]	Provide easy, intuitive interface for non-experts
		Visualisation Readability
✓	[V.1]	Show arbitrary number of dimensions
✓	[V.2]	Display abstract data
✗	[V.3]	Display photos for each entry
✓	[V.4]	Highlight clusters
✓	[V.5]	Identify clusters, trends, outliers
		Space & Immersion
✓	[SI.1]	Provide high immersion
✓	[SI.2]	Provide large space for analysis
✓	[SI.3]	Combines real and virtual
✓	[SI.4]	Interactive in real time
✗	[SI.5]	Registered in 3D
		Collaboration
✓	[C.1]	Allow multiple active users
✗	[C.2]	Support collaboration between multiple users

Table 5.1: Overview over fulfilled requirements. ✓ denotes a fulfilled requirement, ✗ an unfulfilled requirement.

Chapter 6

Conclusion

This work addresses the challenge of collaboratively analysing mobile health data in immersive AR environments. For this, a methodical approach consisting of three steps was chosen:

1. Analysis of limitations in current analysis approaches and extraction of leverage points for *Immersive Analytics*.
2. Design and implementation of the ART system.
3. Evaluation of the ART system with domain experts.

For the initial use case analysis, a focus group between four interaction designers and seven domain experts from the SMARTACT project was conducted. The focus group revealed several limitations in current analysis approaches, and four leverage points for immersive data analysis tools concerning the approach to analysis, visualisation readability, space & immersion, and collaboration were identified. Moreover, specific requirements belonging to these leverage points were gathered. Based on these requirements the immersive analytics tool ART was proposed.

ART combines the familiar interaction on a touch-sensitive tabletop with a 3D visualisation in AR above the tabletop. A novel visualisation technique using scatter plots and 3D parallel coordinates is used to visualise multidimensional data. Users can configure the visualisation through familiar touch controls on the tabletop. This work also presented the technical setting and implementation of ART. The system uses the *HTC Vive* VR HMD with front-mounted cameras to create an immersive AR environment. As interaction is handled on a touch-enabled tabletop, the system is divided among several components, necessitating a distributed system, using a combination of an interactive web application with a visualisation in *Unity3D*.

The evaluation results of group-based expert walkthroughs show evidence that a setting similar to that used in ART can facilitate immersion in the data and collaboration, when compared to traditional desktop systems. The system further allows for the analysis of multidimensional clusters, trends, and outliers, and supports a fluent extensible analysis workflow. Based on these findings, design guidelines for the design of immersive analytics tools were provided. In addition, several research directions were identified to further facilitate the collaborative analysis of multidimensional data in AR environments. Despite several limitations regarding the evaluation, hardware, and implementation, the ART system fulfils almost all requirements and demonstrates the usefulness of immersive data analysis tools for this use case.

There are, however, still many possibilities for improvements and additional features that can further facilitate the analysis process. The following future work directions are collected from observations and suggestions during the evaluation, as well as outstanding requirements.

Gesture interaction. One of the most requested feature during the expert evaluation was gesture interaction. Devices like *Leap Motion* are small enough to be fitted in front of the HMD, offering gesture recognition which can increase immersion and make the interaction more intuitive and natural. Although the table interaction was very useful for some interactions that required precision (e.g. filter creation), other interactions can be more intuitive through gestures (e.g. rotating individual scatter plots). This would also allow for new interaction methods (e.g. selecting individual lines), which may be hard to perform on a 2D interface. However, proper gesture interaction requires proper 3D occlusion, so that the visualisation is not rendered above the user's hand.

Dynamic workflow improvements. Even though ART supports a rather dynamic workflow, many individual components are still very static. Features such as different aggregation levels have to be precalculated in the database, as the ART system does not perform any calculations for this on its own.

Moreover, the available variables are currently limited to the system's predefined variables. Allowing for the dynamic creation of new variables (e.g. by specifying a range from an existing dimension, or including a new data source) could facilitate the workflow.

The visualisation can also be adapted more dynamically to the data at hand. For example, all scatter plots currently use a static global scale to make lines comparable between similar variables (e.g. portions of milk, vegetables, meat). This makes it easier to compare differences between two scatter plots, but squeezes data points together. Dynamically scaling or expanding the scatter plots based on their surrounding scatter plots could therefore increase visualisation readability.

Lastly, ART is currently limited to a fairly linear workflow. Features for nonlinearity, such as saving a snapshot of the current configuration, could make the analysis more dynamic and useful.

Collaboration support. The system currently allows for collaboration during exploration, but is still fairly restricted when it comes to configuration on the tabletop. Further research is required on how to support diverging workflows for individual data analysis (i.e. when users want to individually analyse different scenarios), and how these workflows can later be recombined to resume collaboration.

Furthermore, users that do not wear the HMD are currently forced into passive observer roles, and can only glimpse at the visualisation through monitors showing the HMD's view. For a more active role in the analysis, these observers could view the visualisation on tablets (e.g. like Chen et al. [18]), allowing for a simplified interaction, which can encourage discussions and boost collaboration possibilities. With recent advances in AR frameworks such as *ARCore* [37] and *ARKit* [3], these tablets could potentially also use an AR environment for a more active collaboration between HMD and tablet users.

More visualisations. Currently, individual plots in ART are restricted to 2D scatter plots, allowing the visualisation to link related data points between plots. Projects like *VisLink* [22] and *Caleydo* [53, 80] demonstrate the practicality of applying similar linking techniques to different visualisations. Since ART is not restricted to a 2D environment, the system could incorporate different data visualisations, such as 3D bar charts or a map visualisation for displaying the meal's geospatial location. This would also allow the inclusion of non-abstract information, such as pictures of the meal.

Visualisation improvements. There is a wealth of research into the optimal organisation and usage of visualisations. Due to its novel combination of 2D scatter plots with 3D parallel coordinates, such methods have not been considered in the current ART system. For example, research shows various beneficial rendering techniques for parallel coordinates [56], or clutter reduction through dimension sorting [69]. Further research is necessary on how these methods can be applied to ART.

Inclusion of statistical measures. The visual analysis can further be facilitated by including statistics such trendlines or averages in the visualisation. ART's sort functionality already shows some information concerning the value distribution, which added contextual information. Yet, the AR environment offers much more space and possibilities to include similar statistical information.

Interoperability with other applications. Currently, ART is a closed system that does not integrate with other tools outside of fetching data from a database. For a data analysis workflow, however, the interoperability between different tools is crucial, as it allows data analysts to move seamlessly between the appropriate tools. Therefore, ART should offer functionalities such as exporting, saving, and loading current visualisation configurations as well as taking screenshots.

Bibliography

[1] Sasan Adibi. *Mobile Health: A Technology Road Map*. Ed. by Sasan Adibi. Springer International Publishing, 18th Feb. 2015. 1172 pp. isbn: 978-3-319-12817-7. doi: 10.1007/978-3-319-12817-7.

[2] Judith Amores and Jaron Lanier. 'HoloARt. Painting with Holograms in Mixed Reality'. In: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '17*. ACM Press, 2017. doi: 10.1145/3027063.3050435.

[3] Apple. *ARKit - Apple Developer*. 2017. url: https://developer.apple.com/arkit/ (visited on 2017-11-04).

[4] Ronald T. Azuma. 'A Survey of Augmented Reality'. In: *Presence: Teleoperators and Virtual Environments* 6.4 (Aug. 1997), pp. 355–385. doi: 10.1162/pres.1997.6.4.355.

[5] Daniel Belcher, Mark Billinghurst, S.E. Hayes and Randy Stiles. 'Using augmented reality for visualizing complex graphs in three dimensions'. In: *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.* San Diego, CA, USA: IEEE Comput. Soc, 2003, pp. 1–9. isbn: 0-7695-2006-5. doi: 10.1109/ismar.2003.1240691.

[6] Hrvoje Benko and Steven Feiner. 'Balloon Selection: A Multi-Finger Technique for Accurate Low-Fatigue 3D Selection'. In: *2007 IEEE Symposium on 3D User Interfaces*. San Diego, CA, USA: IEEE, 2007, pp. 79–86. isbn: 1-4244-0907-1. doi: 10.1109/3dui.2007.340778.

[7] Hrvoje Benko, Edward W. Ishak and Steven Feiner. 'Cross-dimensional gestural interaction techniques for hybrid immersive environments'. In: *IEEE Proceedings. VR 2005. Virtual Reality, 2005.* San Diego, CA, USA: IEEE, 2005, pp. 209–216. isbn: 0-7803-8929-8. doi: 10.1109/vr.2005.1492776.

[8] Hrvoje Benko, E.W. Ishak and Steven Feiner. 'Collaborative Mixed Reality Visualization of an Archaeological Excavation'. In: *Third IEEE and ACM International Symposium on Mixed and Augmented Reality*. San Diego, CA, USA: IEEE, 2004, pp. 132–140. isbn: 0-7695-2191-6. doi: 10.1109/ismar.2004.23.

[9] Rebecca Bennett, David J. Zielinski and Regis Kopper. 'Comparison of interactive environments for the archaeological exploration of 3D landscape data'. In: *2014 IEEE VIS International Workshop on 3DVis (3DVis)*. San Diego, CA, USA: IEEE, Nov. 2014, pp. 67–71. isbn: 978-1-4799-6826-8. doi: 10.1109/3dvis.2014.7160103.

[10] Alberto Betella, Rodrigo Carvalho, Jesus Sanchez-Palencia, Ulysses Bernardet and Paul F. M. J. Verschure. 'Embodied interaction with complex neuronal data in mixed-reality'. In: *Proceedings of the Virtual Reality International Conference (VRIC'12)*. New York, NY, USA: ACM Press, 2012, p. 1. isbn: 9781450312431. doi: 10.1145/2331714.2331718.

[11] Mark Billinghurst and Hirokazu Kato. 'Collaborative Mixed Reality'. In: *International Symposium on Mixed Reality (ISMR '99). Mixed Reality - Merging Real and Virtual Worlds*. Berlin, DE: Springer, 1999, pp. 261–284. isbn: 9783642875144. doi: 10.1007/978-3-642-87512-0_15.

[12] Gerd Bruder, Frank Steinicke and Wolfgang Stuerzlinger. 'Touching the Void Revisited: Analyses of Touch Behavior on and above Tabletop Surfaces'. In: *Human-Computer Interaction – INTERACT 2013*. Berlin, DE: Springer, 2013, pp. 278–296. doi: 10.1007/978-3-642-40483-2_19.

[13] *VTK4AR: An Object Oriented Framework for Scientific Visualization of CAE Data in Augmented Reality.* Jan. 2006, pp. 76–81.

[14] Simon Butscher, Yunlong Wang, Katrin Ziesemer, Karoline Villinger, Deborah Wahl, Laura König, Gudrun Sproesser, Britta Renner, Harald T. Schupp and Harald Reiterer. 'Lightweight Visual Data Analysis on Mobile Devices: Providing Self-Monitoring Feedback'. In: *VVH 2016-1st International Workshop on" Valuable visualization of healthcare information"*. 2016, pp. 28–34.

[15] Andreas Butz, Thomas Höllerer, Steven Feiner, Blair MacIntyre and Clifford Beshers. 'Enveloping users and computers in a collaborative 3D augmented reality'. In: *Proceedings 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. San Diego, CA, USA: IEEE Comput. Soc, 1999, pp. 35–44. isbn: 0-7695-0359-4. doi: 10.1109/iwar.1999.803804.

[16] Maurizio Caon, Stefano Carrino, Federica Prinelli, Valentina Ciociola, Fulvio Adorni, Claudio Lafortuna, Sarah Tabozzi, José Serrano, Laura Condon, Omar Abou Khaled and Elena Mugellini. 'Towards an Engaging Mobile Food Record for Teenagers'. In: *New Trends in Image Analysis and Processing – ICIAP 2015 Workshops*. Springer International Publishing, 2015, pp. 417–424. doi: 10.1007/978-3-319-23222-5_51.

[17] Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Cagatay Goncu, Matthias Klapperstueck, Karsten Klein, Kim Marriott, Falk Schreiber and Elliot Wilson. 'Immersive Analytics'. In: *2015 Big Data Visual Analytics (BDVA)*. IEEE, Sept. 2015, pp. 1–8. isbn: 9781467373432. doi: 10.1109/bdva.2015.7314296.

[18] Henry Chen, Austin S. Lee, Mark Swift and John C. Tang. '3D Collaboration Method over HoloLens™ and Skype™ End Points'. In: *Proceedings of the 3rd International Workshop on Immersive Media Experiences - ImmersiveME '15*. ACM. ACM Press, 2015, pp. 27–30. doi: 10.1145/2814347.2814350.

[19] Keke Chen. 'Optimizing star-coordinate visualization models for effective interactive cluster exploration on big data'. In: *Intelligent Data Analysis* 18.2 (2014), pp. 117–136. doi: 10.3233/IDA-140633.

[20] Dane Coffey, Nicholas Malbraaten, Trung Bao Le, Iman Borazjani, Fotis Sotiropoulos, Arthur G. Erdman and Daniel F. Keefe. 'Interactive Slice WIM: Navigating and Interrogating Volume Data Sets Using a Multisurface, Multitouch VR Interface'. In: *IEEE Transactions on Visualization and Computer Graphics* 18.10 (Oct. 2012), pp. 1614–1626. issn: 1077-2626. doi: 10.1109/tvcg.2011.283.

[21] Dane Coffey, Nicholas Malbraaten, Trung Le, Iman Borazjani, Fotis Sotiropoulos and Daniel F. Keefe. 'Slice WIM: a multi-surface, multi-touch interface for overview+detail exploration of volume datasets in virtual reality'. In: *Symposium on Interactive 3D Graphics and Games on - I3D '11*. New York, NY, USA: ACM Press, 2011, pp. 191–198. isbn: 9781450305655. doi: 10.1145/1944745.1944777.

[22] Christopher Collins and Sheelagh Carpendale. 'VisLink: Revealing Relationships Amongst Visualizations'. In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (Nov. 2007), pp. 1192–1199. issn: 1077-2626. doi: 10.1109/tvcg.2007.70521.

[23] Sunny Consolvo, Predrag Klasnja, David W. McDonald and James A. Landay. 'Designing for Healthy Lifestyles: Design Considerations for Mobile Technologies to Encourage Consumer Health and Wellness'. In: *Foundations and Trends® in Human–Computer Interaction* 6.3–4 (2012), pp. 167–315. issn: 1551-3955. doi: 10.1561/1100000040.

[24] Maxime Cordeil, Tim Dwyer, Karsten Klein, Bireswar Laha, Kim Marriott and Bruce H. Thomas. 'Immersive Collaborative Analysis of Network Connectivity: CAVE-style or Head-Mounted Display?' In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 441–450. issn: 1077-2626. doi: 10.1109/tvcg.2016.2599107.

[25] Carolina Cruz-Neira, Daniel J. Sandin and Thomas A. DeFanti. 'Surround-screen projection-based virtual reality: the design and implementation of the CAVE'. In: *Proceedings of the 20th annual conference on Computer graphics and interactive techniques - SIGGRAPH '93*. New York, NY, USA: ACM Press, 1993, pp. 135–142. isbn: 0897916018. doi: 10.1145/166117.166134.

[26] DAQRI. *Open Source Augmented Reality SDK | ARToolKit.org*. 2017. url: https://artoolkit.org/ (visited on 2017-11-04).

[27] Ciro Donalek, S. G. Djorgovski, Alex Cioc, Anwell Wang, Jerry Zhang, Elizabeth Lawler, Stacy Yeh, Ashish Mahabal, Matthew Graham, Andrew Drake, Scott Davidoff, Jeffrey S. Norris and Giuseppe Longo. 'Immersive and collaborative data visualization using virtual reality platforms'. In: *2014 IEEE International Conference on Big Data (Big Data)*. San Diego, CA, USA: IEEE, Oct. 2014, pp. 609–614. isbn: 978-1-4799-5666-1. doi: 10.1109/bigdata.2014.7004282. arXiv: 1410.7670.

[28] Daniel Drochtert and Christian Geiger. 'Collaborative magic lens graph exploration'. In: *SIGGRAPH ASIA 2015 Mobile Graphics and Interactive Applications on - SA '15*. New York, NY, USA: ACM Press, 2015, pp. 1–3. isbn: 9781450339285. doi: 10.1145/2818427.2818465.

[29] Elena Fanea, Sheelagh Carpendale and Tobias Isenberg. 'An interactive 3D integration of parallel coordinates and star glyphs'. In: *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.* San Diego, CA, USA: IEEE, 2005, pp. 149–156. isbn: 0-7803-9464-X. doi: 10.1109/infvis.2005.1532141.

[30] Alessandro Febretti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, J. D. Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, Dan Sandin, Luc Renambot, Andrew Johnson and Jason Leigh. 'CAVE2: a hybrid reality environment for immersive simulation and information analysis'. In: *The Engineering Reality of Virtual Reality 2013*. Ed. by Margaret Dolinsky and Ian E. McDowall. International Society for Optics and Photonics. SPIE, Mar. 2013, pp. 864903–864903. doi: 10.1117/12.2005484.

[31] Edwin B. Fisher, Marian L. Fitzgibbon, Russell E. Glasgow, Debra Haire-Joshu, Laura L. Hayman, Robert M. Kaplan, Marilyn S. Nanney and Judith K. Ockene. 'Behavior Matters'. In: *American Journal of Preventive Medicine* 40.5 (May 2011), e15–e30. issn: 0749-3797. doi: 10.1016/j.amepre.2010.12.031.

[32] Cedric Fleury, Nicolas Ferey, Jean-Marc Vezien and Patrick Bourdot. 'Remote collaboration across heterogeneous large interactive spaces'. In: *2015 IEEE Second VR International Workshop on Collaborative Virtual Environments (3DCVE)*. IEEE, Mar. 2015. doi: 10.1109/3dcve.2015.7153591.

[33] Michael Friendly. *Milestones in the history of thematic cartography, statistical graphics, and data visualization*. 2009. url: http://www.datavis.ca/milestones/ (visited on 2017-08-28).

[34] Anton Fuhrmann, Gerd Hesina, François Faure and Michael Gervautz. 'Occlusion in collaborative augmented environments'. In: *Computers & Graphics* 23.6 (Dec. 1999), pp. 809–819. issn: 00978493. doi: 10.1016/s0097-8493(99)00107-7.

[35] Anton Fuhrmann, Helwig Löffelmann, Dieter Schmalstieg and Michael Gervautz. 'Collaborative visualization in augmented reality'. In: *IEEE Computer Graphics and Applications* 18.4 (1998), pp. 54–59. issn: 02721716. doi: 10.1109/38.689665.

[36] Google. *Angular*. 2017. url: https://angular.io/ (visited on 2017-11-04).

[37] Google. *ARCore - Google Developer*. 2017. url: https://developers.google.com/ar/ (visited on 2017-11-04).

[38] Martin Hachet, Benoit Bossavit, Aurélie Cohé and Jean-Baptiste de la Rivière. 'Toucheo: multitouch and stereo combined in a seamless workspace'. In: *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '11)*. New York, NY, USA: ACM Press, 2011, pp. 587–592. isbn: 9781450307161. doi: 10.1145/2047196.2047273.

[39] Ping-Hsuan Han, Kuan-Wen Chen, Chen-Hsin Hsieh, Yu-Jie Huang and Yi-Ping Hung. 'AR-Arm. Augmented Visualization for Guiding Arm Movement in the First-Person Perspective'. In: *Proceedings of the 7th Augmented Human International Conference 2016 on - AH '16*. ACM Press, 2016. doi: 10.1145/2875194.2875237.

[40] Ping-Hsuan Han, Yang-Sheng Chen, Yilun Zhong, Han-Lei Wang and Yi-Ping Hung. 'My Tai-Chi coaches. an augmented-learning tool for practicing Tai-Chi Chuan'. In: *Proceedings of the 8th Augmented Human International Conference on - AH '17*. ACM Press, 2017. doi: 10.1145/3041164.3041194.

[41] HCI group University of Konstanz. *SMARTACT / SMARTMOBILITY*. 2015. url: http://hci.uni-konstanz.de/smartact (visited on 2017-09-13).

[42] Jeffrey Heer and Maneesh Agrawala. 'Design Considerations for Collaborative Visual Analytics'. In: *Information Visualization* 7.1 (Feb. 2008), pp. 49–62. issn: 1473-8716. doi: 10.1057/palgrave.ivs.9500167.

[43] Kristin E. Heron and Joshua M. Smyth. 'Ecological momentary interventions: Incorporating mobile technology into psychosocial and health behaviour treatments'. In: *British Journal of Health Psychology* 15.1 (Feb. 2010), pp. 1–39. issn: 2044-8287. doi: 10.1348/135910709x466063.

[44] HTC Corporation. *Vive*. 2017. url: https://www.vive.com/ (visited on 2017-11-10).

[45] HTC Corporation. *Vive PRE User Guide*. https://www.htc.com/managed-assets/shared/desktop/vive/Vive_PRE_User_Guide.pdf (visited on 2017-11-04). 2016.

[46] Sebastian Hubenschmid. *Project Report: Explorative Data Analysis in Augmented Reality above the Tabletop*. HCI University of Konstanz. 31st July 2017.

[47] IBM. *SPSS Statistics*. 2017. url: https://www.ibm.com/products/spss-statistics (visited on 2017-11-10).

[48] Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma and Hans Hagen. 'Collaborative visualization: Definition, challenges, and research agenda'. In: *Information Visualization* 10.4 (July 2011), pp. 310–326. issn: 1473-8716. doi: 10.1177/1473871611412817.

[49] Oliver Kreylos. *On the road for VR: Microsoft HoloLens at Build 2015, San Francisco*. 1st May 2015. url: http://doc-ok.org/?p=1223 (visited on 2017-07-13).

[50] Oliver Kreylos. *Optical Properties of Current VR HMDs*. 1st Apr. 2016. url: http://doc-ok.org/?p=1414 (visited on 2017-09-13).

[51] Gregorij Kurillo and Maurizio Forte. 'Telearch - Integrated Visual Simulation Environment for Collaborative Virtual Archaeology'. In: *Mediterranean Archaeology and Archaeometry* 12.1 (2012), pp. 11–20.

[52] Bireswar Laha, Doug A. Bowman and James D. Schiffbauer. 'Validation of the MR Simulation Approach for Evaluating the Effects of Immersion on Visual Analysis of Volume Data'. In: *IEEE Transactions on Visualization and Computer Graphics* 19.4 (Apr. 2013), pp. 529–538. issn: 1077-2626. doi: 10.1109/tvcg.2013.43.

[53] Alexander Lex, Marc Streit, Ernst Kruijff and Dieter Schmalstieg. 'Caleydo: Design and evaluation of a visual analysis framework for gene expression data in its biological context'. In: *2010 IEEE Pacific Visualization Symposium (PacificVis)*. San Diego, CA, USA: IEEE, Mar. 2010, pp. 57–64. isbn: 978-1-4244-6685-6. doi: 10.1109/pacificvis.2010.5429609.

[54] David López, Lora Oehlberg, Candemir Doger and Tobias Isenberg. 'Towards An Understanding of Mobile Touch Navigation in a Stereoscopic Viewing Environment for 3D Data Exploration'. In: *IEEE Transactions on Visualization and Computer Graphics* 22.5 (May 2016), 1616–1629. issn: 1077-2626. doi: 10.1109/tvcg.2015.2440233.

[55] Eli Mahfoud and Aidong Lu. 'Gaze-directed Immersive Visualization of Scientific Ensembles'. In: *Proceedings of the 2016 ACM Companion on Interactive Surfaces and Spaces - ISS Companion '16*. ACM Press, 2016. doi: 10.1145/3009939.3009952.

[56] Kevin T. McDonnell and Klaus Mueller. 'Illustrative Parallel Coordinates'. In: *Computer Graphics Forum* 27.3 (May 2008), pp. 1031–1038. issn: 0167-7055. doi: 10.1111/j.1467-8659.2008.01239.x.

[57] Bianchi Serique Meiguins, Ricardo Melo Casseb do Carmo, Leonardo Almeida, Aruanda Simões Gonçalves, Sérgio Clayton V. Pinheiro, Marcelo de Brito Garcia and Paulo Igor Alves Godinho. 'Multidimensional information visualization using augmented reality'. In: *Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications - VRCIA '06*. New York, NY, USA: ACM Press, 2006, pp. 14–17. isbn: 1595933247. doi: 10.1145/1128923.1128996.

[58] Leonel Merino, Johannes Fuchs, Michael Blumenschein, Craig Anslow, Mohammad Ghafari, Oscar Nierstrasz, Michael Behrisch and Daniel A. Keim. 'On the Impact of the Medium in the Effectiveness of 3D Software Visualizations'. In: *2017 IEEE Working Conference on Software Visualization (VISSOFT)*. IEEE, Sept. 2017. doi: 10.1109/vissoft.2017.17.

[59] Microsoft. *Excel*. 2017. url: https://products.office.com/en-gb/excel (visited on 2017-11-10).

[60] Microsoft. *Microsoft Surface Hub | Product Specs*. 2017. url: https://www.microsoft.com/microsoft-surface-hub/en-gb/product-specs (visited on 2017-11-10).

[61] Microsoft. *TypeScript*. 2017. url: https://www.typescriptlang.org/ (visited on 2017-11-04).

[62] Paul Milgram and Fumio Kishino. 'A taxonomy of mixed reality visual displays'. In: *IEICE TRANSACTIONS on Information and Systems* 77.12 (1994), pp. 1321–1329. issn: 0916-8532.

[63] Jens Müller, Simon Butscher and Harald Reiterer. 'Immersive Analysis of Health-Related Data with Mixed Reality Interfaces: Potentials and Open Questions'. In: *Workshop Immersive Analytics 2016 (in conjunction with ISS '16)*. University of Konstanz, Nov. 2016. doi: 10.1145/3009939.3009951.

[64] Tamara Munzner. *Visualization Analysis and Design*. CRC Press, 15th Sept. 2015, pp. 117–130. 428 pp. isbn: 9781498759717.

[65] Michael Narayan, Leo Waugh, Xiaoyu Zhang, Pradyut Bafna and Doug Bowman. 'Quantifying the benefits of immersion for collaboration in virtual environments'. In: *Proceedings of the ACM symposium on Virtual reality software and technology - VRST '05*. New York, NY, USA: ACM Press, 2005, pp. 78–81. isbn: 1595930981. doi: 10.1145/1101616.1101632.

[66] Huyen Nguyen, Peter Marendy and Ulrich Engelke. 'Collaborative Framework Design for Immersive Analytics'. In: *2016 Big Data Visual Analytics (BDVA)*. San Diego, CA, USA: IEEE, Nov. 2016, pp. 23–30. isbn: 978-1-5090-5272-1. doi: 10.1109/bdva.2016.7787044.

[67] Node.js Foundation. *Node.js*. 2017. url: https://nodejs.org (visited on 2017-11-04).

[68] Isabel Pedersen, Nathan Gale and Pejman Mirza-Babaei. 'TombSeer. Illuminating the Dead'. In: *Proceedings of the 7th Augmented Human International Conference 2016 on - AH '16*. ACM Press, 2016. doi: 10.1145/2875194.2875201.

[69] Wei Peng, Matthew O. Ward and Elke A. Rundensteiner. 'Clutter Reduction in Multi-Dimensional Data Visualization Using Dimension Reordering'. In: *IEEE Symposium on Information Visualization*. IEEE, 2004, pp. 88–96. doi: 10.1109/infvis.2004.15.

[70] Dheva Raja, Doug A. Bowman, John Lucas and Chris North. 'Exploring the Benefits of Immersion in Abstract Information Visualization'. In: *Proceedings of Immersive Projection Technology Workshop*. 2004.

[71] Gabriel Robles-De-La-Torre. 'The Importance of the Sense of Touch in Virtual and Real Environments'. In: *IEEE Multimedia* 13.3 (July 2006), pp. 24–30. issn: 1070-986X. doi: 10.1109/mmul.2006.69.

[72] René Rosenbaum, Jeremy Bottleson, Zhuiguang Liu and Bernd Hamann. 'Involve Me and I Will Understand!–Abstract Data Visualization in Immersive Environments'. In: *International Symposium on Visual Computing: Advances in Visual Computing (ISVC'11)*. Vol. 6938 LNCS. PART 1. Springer Berlin Heidelberg, 2011, pp. 530–540. isbn: 9783642240270. doi: 10.1007/978-3-642-24028-7_49.

[73] Sam Royston, Connor DeFanti and Ken Perlin. 'A Collaborative Untethered Virtual Reality Environment for Interactive Social Network Visualization'. In: *CoRR* abs/1604.08239 (2016). arXiv: 1604.08239.

[74] Dieter Schmalstieg, Anton Fuhrmann, Gerd Hesina, Zsolt Szalavári, L. Miguel Encarnação, Michael Gervautz and Werner Purgathofer. 'The Studierstube Augmented Reality Project'. In: *Presence: Teleoperators and Virtual Environments* 11.1 (Feb. 2002), pp. 33–54. issn: 1054-7460. doi: 10.1162/105474602317343640.

[75] Beat Signer and Timothy J. Curtin. 'Tangible Holograms: Towards Mobile Physical Augmentation of Virtual Objects'. In: *CoRR* abs/1703.08288 (2017). arXiv: 1703.08288.

[76] Neil G. Smith, Kyle Knabb, Connor DeFanti, Philip Weber, Jurgen Schulze, Andrew Prudhomme, Falko Kuester, Thomas E. Levy and Thomas A. DeFanti. 'ArtifactVis2: Managing real-time archaeological data in immersive 3D environments'. In: *Proceedings of Digital Heritage International Congress (DigitalHeritage'13)*. IEEE, Oct. 2013, pp. 363–370. isbn: 978-1-4799-3170-5. doi: 10.1109/DigitalHeritage.2013.6743761.

[77] Maurıcio Sousa, Daniel Mendes, Soraia Paulo, Nuno Matela, Joaquim Jorge and Daniel Simões Lopes. 'VRRRRoom: Virtual Reality for Radiologists in the Reading Room'. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. New York, NY, USA: ACM Press, 2017, pp. 4057–4062. isbn: 9781450346559. doi: 10.1145/3025453.3025566.

[78] William Steptoe. *AR-Rift*. 20th Nov. 2013. url: http://willsteptoe.com/post/66968953089/ar-rift-part-1 (visited on 2017-09-13).

[79] Jonathan Steuer. 'Defining Virtual Reality: Dimensions Determining Telepresence'. In: *Journal of Communication* 42.4 (Dec. 1992), pp. 73–93. doi: 10.1111/j.1460-2466.1992.tb00812.x.

[80] Marc Streit, Alexander Lex, Heimo Müller and Dieter Schmalstieg. 'Gaze-based focus adaption in an information visualization system'. In: *IADIS International Conference Computer Graphics, Visualization, Computer Vision and Image Processing*. 2009, pp. 303–307.

[81] Sven Strothoff, Dimitar Valkov and Klaus Hinrichs. 'Triangle cursor: interactions with objects above the tabletop'. In: *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces - ITS '11*. New York, NY, USA: ACM Press, 2011, pp. 111–119. isbn: 9781450308717. doi: 10.1145/2076354.2076377.

[82] Zsolt Szalavári, Dieter Schmalstieg, Anton Fuhrmann and Michael Gervautz. '"Studierstube": An environment for collaboration in augmented reality'. In: *Virtual Reality* 3.1 (Mar. 1998), pp. 37–48. issn: 1359-4338. doi: 10.1007/bf01409796.

[83] Tableau Software. *Tableau*. 2017. url: https://www.tableau.com/ (visited on 2017-11-10).

[84] Michael Tanaya, KouKeng Yang, Taran Christensen, Shen Li, Michael O'Keefe, James Fridley and Kelvin Sung. 'A Framework for analyzing AR/VR Collaborations: An initial result'. In: *2017 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*. IEEE, June 2017. doi: 10.1109/civemsa.2017.7995311.

[85] University of Konstanz. *SMARTACT*. 2015. url: https://www.uni-konstanz.de/smartact/ (visited on 2017-09-13).

[86] Christophe Viau, Michael J. McGuffin, Yves Chiricota and Igor Jurisica. 'The FlowVizMenu and Parallel Scatterplot Matrix: Hybrid Multidimensional Visualizations for Network Exploration'. In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (Nov. 2010), pp. 1100–1108. issn: 1077-2626. doi: 10.1109/tvcg.2010.205.

[87] Lidong Wang, Guanghui Wang and Cheryl Ann Alexander. 'Big data and visualization: methods, challenges and technology progress'. In: *Digital Technologies* 1.1 (2015), pp. 33–38.

[88] Colin Ware and Glenn Franck. 'Viewing a graph in a virtual reality display is three times as good as a 2D diagram'. In: *Proceedings of 1994 IEEE Symposium on Visual Languages*. San Diego, CA, USA: IEEE Comput. Soc. Press, 1994, pp. 182–183. isbn: 0-8186-6660-9. doi: 10.1109/vl.1994.363621.

[89] Wizapply Team. *Ovrvision Pro*. 2017. url: http://ovrvision.com/entop/ (visited on 2017-11-10).

[90] Song Zhang, Çağatay Demiralp, Daniel F. Keefe, Marco DaSilva, David H. Laidlaw, Benjamin D. Greenberg, Peter J. Basser, Carlo Pierpaoli, Ennio Antonio Chiocca and Thomas S. Deisboeck. 'An immersive virtual environment for DT-MRI volume visualization applications: a case study'. In: *Proceedings Visualization, 2001. VIS '01.* San Diego, CA, USA: IEEE, 2001, pp. 437–440. doi: 10.1109/visual.2001.964545.

[91] Sean Ziegeler, Robert J. Moorhead, Paul J. Croft and Duanjun Lu. 'The MetVR case study: meteorological visualization in an immersive virtual environment'. In: *Proceedings of the Conference on Visualization (VIS'01)*. San Diego, CA, USA: IEEE, 2001, pp. 489–596. isbn: 0-7803-7200-X. doi: 10.1109/visual.2001.964559.

Appendix A

Contents of DVD

The attached DVD includes all files related to this work:

- Digital version of this document
- Project report
- Seminar
- Complete code for the ART system
- Video showcasing the ART system
- Expert walkthrough videos