MASTER'S THESIS 2019

Data Visualization of Product Relations
- An Interactive Virtual Reality Solution

Anton Kemvik Nils Goksör

jayway.
by devoteam

DEPARTMENT OF DESIGN SCIENCES
LTH | LUND UNIVERSITY

Data Visualization of Product Relations

(An Interactive Virtual Reality Solution)

Anton Kemvik
antonkemvik@gmail.com

Nils Goksör
nils.goksor@gmail.com

June 17, 2019

Master's thesis work carried out at
the Department of Design Sciences, Faculty of Engineering LTH, Lund University.

Supervisors: Joakim Eriksson, joakim.eriksson@design.lth.se
Pär Sikö, par.siko@jayway.com

Examiner: Günter Alce, gunter.alce@design.lth.se

Data Visualization of Product Relations

An Interactive Virtual Reality Solution

Published by

Department of Design Sciences
Faculty of Engineering LTH, Lund University
P.O. Box 118, SE-221 00 Lund, Sweden

Subject: Interaction Design (MAMM01)
Division: Ergonomics and Aerosol Technology
Supervisor: Joakim Eriksson
Examiner: Günter Alce

Abstract

Thanks to the rapid development in areas such as artificial intelligence companies now have the possibility to collect and analyze data to improve their business.

These algorithms quickly result in large volumes of data, which are to be processed and understood by humans. The stakeholder for this project, IKEA, cannot keep track of all data relations in their sales patterns by using regular tables or graphs.

Data visualization has always been a great tool to understand and communicate data. The purpose of data visualization is to convert slow reasoning tasks into into fast perception tasks.

The hypothesis of this project is that the third dimension available in Virtual Reality (VR) can simplify this for big data. We have created a VR application to visualize product relations within the product range at IKEA. The application combines known research in data visualization, such as visual encodings and gestalt principles, with the advantages of 3D visualization and VR.

Our conclusion is that the VR application provides a great overview of big data, and that tasks such as finding data relations in a big data set are greatly simplified. The testing sessions with the end users at IKEA confirms this. Furthermore, they believe that the application is a great tool for discovering patterns and relations within the data that was not possible before, and that it can be used to present research findings in an intuitive way.

Keywords: Virtual Reality, 3D, Data visualization, Data relations, Big data

Sammanfattning

Tack vare den snabba utvecklingen inom områden såsom artificiell intelligens har företag nu möjlighet att samla in och analysera data för att förbättra sin verksamhet.

Dessa algoritmer resulterar snabbt i stora datamängder, som ska bearbetas och tolkas av människor. IKEA, som är intressent i detta projekt, kan inte hålla reda på alla datarelationer i sina försäljningsmönster genom att använda vanliga tabeller eller grafer. Mängden data är för stor.

Datavisualisering är ett utmärkt verktyg för att förstå och kommunicera data. Syftet med datavisualisering är att förenkla analysering av data och göra snabb beslutsfattning möjlig.

Hypotesen för detta projekt är att den tredje dimensionen som finns tillgänglig i Virtual Reality (VR) kan förenkla just detta. Vi har skapat en VR-applikation för att visualisera produktrelationer inom IKEAs produktutbud. Applikationen använder sig av känd forskning inom datavisualisering, såsom visuell kodning och gestaltprinciper, och kombinerar det med fördelarna av att använda 3D-visualisering och VR.

Vår slutsats är att VR ger en bra översikt av samband inom stora mängder data, och att uppgifter såsom att hitta relationer i en stor dataset är mycket förenklade. Tester med slutanvändarna på IKEA bekräftar detta. De anser också att applikationen är ett bra verktyg för att upptäcka mönster och relationer som inte var möjliga tidigare, och att applikationen dessutom kan användas för att presentera forskningsresultat på ett intuitivt sätt.

Nyckelord: Virtual Reality, 3D, Datavisualisering, Datarelationer, Big data

Acknowledgements

This master thesis has been done in collaboration with Jayway, IKEA Range & Supply and the Department of Design Sciences at Lund Technical University.

First and foremost, we would like to thank Jayway for supervising the project and providing us with the proper tools. A special thanks to Pär Sikö and Hang Ruan for helping us in initiating the project, and to the whole of Studio 54 for supporting us throughout the project.

Secondly, we would like to thank Markus Metsola at IKEA Range & Supply for giving us this opportunity and helping us throughout the project. Furthermore, thanks to the whole of IKEA and to the end users in particular.

Lastly, we would like to thank Joakim Eriksson, our supervisor at Lund Technical University.

Contents

1 Introduction **9**
1.1 Background . . . 9
1.2 Purpose and vision . . . 10
1.3 Scope . . . 10

2 Theoretical Background **11**
2.1 The User Centered Design Process . . . 11
2.1.1 Concept Phase . . . 11
2.1.2 Design Phase . . . 12
2.1.3 Implementation Phase . . . 15
2.2 Design Principles . . . 16
2.3 Gestalt Principles . . . 17
2.4 Data Visualization . . . 18
2.4.1 Graphics To Understand And Explore Data . . . 18
2.4.2 The Basics Of Data Visualization . . . 19
2.4.3 Visual Encodings in Data Visualization . . . 20
2.4.4 3D Visualization . . . 21
2.5 Virtual Reality . . . 22
2.5.1 Walking . . . 22
2.5.2 Point and Teleport . . . 22
2.5.3 Object Interaction in VR . . . 22
2.5.4 Manipulating, Scaling and Rotating Objects . . . 23
2.5.5 Fading Transitions . . . 23
2.6 Related work . . . 23

3 Technical Background **25**
3.1 Software Tools . . . 25
3.1.1 Unity . . . 25
3.1.2 SteamVR . . . 25
3.1.3 VRTK . . . 26

3.1.4 Visual Studio Code . . . 26
3.1.5 Git and GitHub . . . 26
3.1.6 Adobe XD . . . 26
3.2 Hardware Tools . . . 26
3.2.1 Virtual Reality Headsets . . . 26

4 The Work Process **29**
4.1 Project Disposition . . . 29
4.2 Concept Phase . . . 30
4.3 Design Phase . . . 33
4.3.1 Iteration 1 . . . 34
4.3.2 Iteration 2 . . . 35
4.3.3 Iteration 3 . . . 36
4.3.4 Iteration 4 . . . 39
4.4 Implementation Phase . . . 41
4.4.1 Iteration 1 . . . 42
4.4.2 Iteration 2 . . . 44
4.4.3 Iteration 3 . . . 46
4.4.4 Iteration 4 . . . 51
4.4.5 Iteration 5 . . . 54

5 Discussion **61**
5.1 The results of the user centered design process . . . 61
5.2 Virtual Reality as a data visualization tool . . . 62
5.3 Future work . . . 63

6 Conclusion **65**

Bibliography **67**

Appendix A The System Usability Scale Survey **73**

Abbreviations

A list of common abbreviations used in the report:

VR - Virtual Reality.

2D - Two Dimensions.

3D - Three Dimensions.

UX - User Experience.

lo-fi Prototype - Low Fidelity Prototype, a very simple, low tech, prototype. Often sketches with limited functionality.

SUS - System Usability Scale.

BA - Business Area. IKEA's top hierarchy level.

HFB - Home Furnishing Business. IKEA's second hierarchy level.

PRA - Product Range Area. IKEA's third hierarchy level.

PA - Product Area. IKEA's last hierarchy level.

Chapter 1
Introduction

1.1 Background

Consumerism has changed drastically during the 21st century, and today, the market is bigger and competition is tougher than ever before. The growth of consumerism is leading to an increase in online- and in store purchases. Many companies have realised the potential in storing an increased volume of data from purchasing in order to gain a better understanding of customer patterns. What products do our customer purchase? When do they purchase it? Which products are purchased together? Through various algorithms, such as machine learning or other general AI, a company can calculate and predict supply, demand and more based on this data [1].

This project was done at Jayway, a design driven software studio in Malmö. It was done in collaboration with IKEA Range & Supply, a core business within Inter IKEA Group. Range & Supply is responsible for developing and supplying the global IKEA range. This means working throughout the whole value chain – from supplier to customer. In this thesis IKEA Range & Supply will be referred to as IKEA.

IKEA store data on how often two different products are sold together, or in other words the probability that you buy a certain product given you are buying another certain product. This is referred to as the affinity between two products. The data is used for better understanding and decision making in regards to production and sales. However, IKEA are currently having issues with presenting this data in an intuitive way that gives the user an easy overview of different product relations. Their current way of analyzing the data is also time consuming, especially when it comes to analyzing how one product relates to, and affect, multiple other products. Today, much of the data analyzing done by our end users at IKEA consists of scouring through a lot of data hoping to find relations and patterns that they suspect are there. In many ways the end user must have an idea of what he/she is looking for. Their current model leaves little room for the user to discover relations or patterns that he/she was unaware of beforehand.

1.2 Purpose and vision

The purpose of this project is to visualize product relations, i.e., the affinity between products, in an innovative and intuitive way.

The project investigates alternative ways of how data visualization can be used to visualize relations in big data. The vision is that the visualization will make it faster, easier and more comprehensive to understand data relations, which our stakeholder can use to make more precise and well-grounded decisions in regards to their product range. This can for example be:

- What happens if we stop producing a product?
- Which products should be advertised together?
- Is a compliment for a product required?
- Which products do customers usually combine when building a certain furniture?

1.3 Scope

This project is a master thesis project and the time spent on it is 20 weeks. Therefore, limitations and prioritizations were made to scope the task into something achievable given the time frame. The following limitations were made:

- Our application will mainly focus on visualizing product relations. Other desired data which might be wanted by our end users will not be prioritized.
- The application handles one specific scenario, whereas the finished product should be applicable for any relation data set.
- The relation data set is somewhat reduced and simplified compared to one of a real life scenario. Furthermore, some values in the application might be hard coded.

Chapter 2
Theoretical Background

2.1 The User Centered Design Process

We have followed a user centered design process in our project. It consists of three phases, which are described in this chapter.

2.1.1 Concept Phase

The purpose of the concept phase is for the design team to understand what issues a user faces with an everyday task or system and why these issues occur. Proper background information and an understanding of the user needs are vital when formulating a vision and goals to solve the issue. Additionally, in some projects the customer have already made up their mind on what is wrong and what kind of solution they want. This does not mean one can skip the concept phase. On the contrary, the design team must spend even more time getting to know the users and their needs, and experiment with different solutions, and possibly changing the customer's mind if they decide it will make for a better solution [2].

The concept phase can be initialized in two ways, and for best result, both these methods should be combined. The first method is solution-driven, where the design team speculates in possible solutions, and the second method is problem-driven, where the design team tries to originate in the user needs and formulate possible solutions which satisfy these. After the initial phase the design team will have a better understanding of the issue they are working on and who the end users of the solutions will be [2].

User studies are a big part of the concept phase and a perfect way into the mind of the end user to thoroughly understand what these users expect from the product and how they aim to use it. The participants of a user study should mainly be the end users. However, others who are connected to the issue should also partake in order to get a broader research, which hopefully will increase the chance that everyone who faces the issue today will be satisfied with the solution. The goal of user study is to gather even more information which will strengthen the design team's insight in the issue. There are different user study methods that can be used to get more insight into the work places, methods, products, people and more which might be relevant to the issue, and these are described below [2].

Interviewing

An open discussion when the participants are encouraged to discuss freely about how they are experiencing the current issue. The questions must be formulated to make the interviewee speak freely and try to make them recall a scenario or make them think out loud about their opinions on the issue. It is recommended that the design team take notes, or even do sketches, to remember what they are told and in what mood or setting it was said. Furthermore, recording the discussion might be useful, especially if the design team later wants to quote an interviewee [2].

Observations

Observing the end users in their current setting is a great compliment to interviewing in order to get an even better sense of the problem. The observing can be done silently, either from distance (anonymously) or by shadowing the user, or by asking them to explain what they are doing or thinking while performing the tasks [2].

Personas and Scenarios

The final part of the concept phase is to put all the gathered information together, preferably into one or multiple personas.

The advantages of creating personas are many. First and foremost, it greatly reduces the risk that the design team have different opinions and needs in mind when designing. Secondly, the communication within the designers as well as with other stakeholders will be easier if one can refer to personas. Lastly, personas makes it easier to establish what requirements the solution is fulfilling, since the persona's requirements are written down and easy to test in scenarios. Personas come to life through scenarios, which can be seen as a hypothetical event which the persona faces. This event focuses on an issue which the gathered information shows the persona commonly faces, and the event can be described by text, pictures, video or even storyboards [3].

A storyboard is a series of pictures with speech bubbles, almost like a cartoon. A great advantage with storyboards is how easy it is to show a user's feelings and the mood around a given issue which a persona faces in the sketch [2].

2.1.2 Design Phase

Following the conceptual phase is the design phase, an iterative phase that generates some sort of outcome or prototype that can be tested in every iteration.

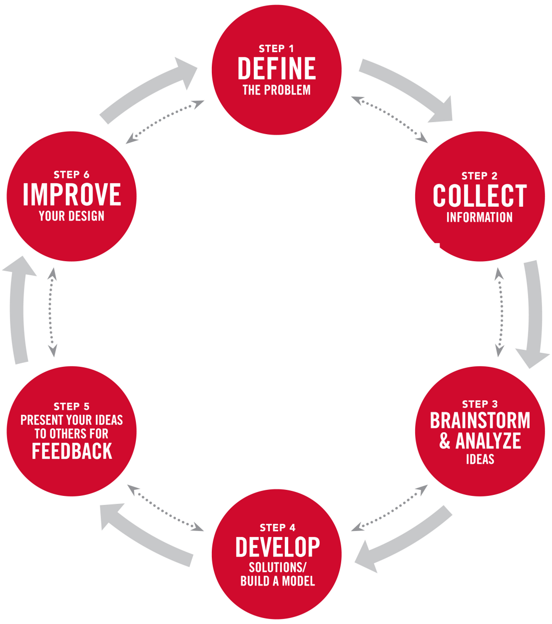

Figure 2.1: The six steps of the design process (adapted from Chicago Architecture Center [4]).

The Six Steps

An iteration in the design phase contains all, or some, of the six steps described in figure 2.1.

1. *Defining the problem*
   Make sure you have a clear idea what the problem is. The problem may change in every iteration as new ideas have been presented and feedback gathered.

2. *Collecting information*
   Make sure you have all the information you need in order to fully understand the current problem and start coming up with possible solutions. This can be done by analyzing the feedback and competitors, and by looking at the latest UX-trends. This step is quite similar to the conceptual phase but additional interviews and research can be conducted in order to stay relevant to the defined problem of the current iteration [5].

3. *Brainstorming*
   This step is the "idea factory". Brainstorming is used to generate as many and as innovative ideas as possible. This step is done on basic sketches that can easily be presented to stakeholders and quickly be re-drawn. The outcome from this step can be notes, sketches, wireframes and user flows [6].

4. *Developing*
   Here you take the best ideas and sketches from the previous step and continue to

work on these. The process here is to turn the initial sketches into prototypes that can be tested by the stakeholders or users. There are multiple online tools that can be used to quickly create interactive prototypes with the intended design. Many of these tools also contain features which allows tracking of user interaction.

5. *Testing*
This is the most important step of the design phase in order to keep iterating. This step is about evaluating the prototype developed in the previous step and gather feedback from the testers. There are multiple ways in which the tests can be conducted. By giving the users a challenge they are to solve using the prototype feedback can be gathered in both verbal form (from the tester in a later interview) and by observing the user interact with the prototype.

6. *Evaluation*
Reflect on the feedback that was gathered. Decide if, or to what extent, the proposed solution should be incorporated. Questions that need to be asked here is:

    - Is the system usable?
    - Does it provide a solution to the defined problem?
    - Is it flexible and easy to change?
    - Has it improved how they feel about, think about or perform their tasks?

Brainstorming

Brainstorming is to discuss all possible ideas, regardless of how crazy they might sound, that might fulfill the user needs. The more ideas that come up here the better, and no idea is too far fetched. Furthermore, it is helpful to experiment with other's ideas and try to build on them, mix ideas and so on [6].

Testing

When performing tests it is important that there are enough testers as the number of testers directly correlates to the percentage of usability problems found. It is described by Nielsen [8] that 15 participants normally find about 99% of all usability problems in an application during a usability test while 5 participants find about 85% of the problems (see figure 2.2). Still Nielsen recommends testing on only 5 participants in 3 iterations instead of testing on 15 participants just one time, why is that? He states that *You want to run multiple tests because the real goal of usability engineering is to improve the design and not just to document its weaknesses* [7]. After having tested the prototype on 5 persons about 85% of the problems have been found, the problems are corrected and then another iteration with 5 testers start. They test if the problems that were discovered have indeed been corrected and also discover most of the remaining 15% of the original usability problems. In the same manner, to make sure that most of the problems found in iteration 2 indeed were corrected a third iteration of testing is performed [8].

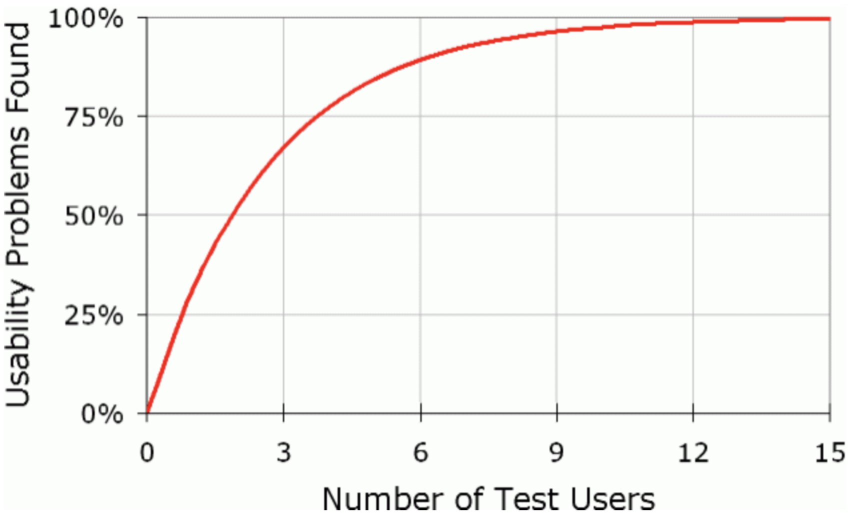

Figure 2.2: Probability graph for usability testing (taken from [7]).

SUS Survey

System Usability Scale, SUS, is a survey which measure the usability of an application. This is a good compliment to interviewing since those questions often are open-ended and more difficult to analyze, whereas the the survey uses a rating scale. SUS often uses a 5-point scale, from *Strongly Disagree* to *Strongly agree*. It consists of 10 questions, where every other question is worded in a negative way, and the other positive, to make the user really think instead of just answering with top score on each question. The amount of participants required for a meaningful SUS survey are very few. The collection of concrete usability data by surveys like SUS is best suited at the end of a study to get a rating of how usable the application is. It is preferred to perform the surveys anonymously over the web to avoid biases when collecting the data, also known as "The Social Desirability Bias". The data collected can be analyzed by converting it to numeric values, where *Strongly disagree* equals 1 and *Strongly agree* equals 5. There is an algorithm of how to calculate the collective SUS score based on these numeric values. Furthermore, a study of what a good SUS score is has been made. It states that 68 is the average score, and a value below 60 is relatively poor, whereas a value above 80 is pretty good [9].

2.1.3 Implementation Phase

This phase includes the transformation of a prototype into a real application. A good start to the implementation phase is to set up some measurable goals which can hopefully be achieved at the end of the phase. These goals could have to do with the success rate of end users performing a task in the application, or how they experience other aspects of the user experience.

The implementation of the application is then done iteratively. In each iteration, functionality can be added or removed, the interface and the flow of the application can be updated, until the application finally fulfills the goals that were set up at the start. This phase is of course user centered too, and follows many of the principles in the design phase.

2.2 Design Principles

These design principles, described by D. Norman [10], are used throughout the report.

Affordance

The term affordance refers to how easily a potential user of an object can define all ways in which this object can be used. It can be seen as a relationship between the object and the user of it, i.e., the affordance of an object depends on both these parties. In other words, affordance is that a physical object contain information about how a potential user can interact with it.

For example, a door knob afford turning, pushing and pulling. The signalling of an affordance, in this case the door knob, is called a *signifier*.

Signifier

Affordances and signifiers are closely connected, with the difference being that an affordance represent the interacting possibilities whereas a signifier can be seen as a signal of what actions can be done and how.

A signifier can be anything from a sign or label to a sound, as long as it indicates the behaviour of the object to the user.

Mapping

The term mapping is the concept of connecting two products that belong together, so the relationship between them becomes clear. It is often a control being mapped to a function of some sort.

The mapping will be much easier to understand if it uses spatial correspondence, which best can be explained by an example: the mapping between your steering wheel and the wheels of the car is spatial, since turning the former right will rotate the latter to the right too.

Feedback

It is important to communicate the result of an action, i.e., giving feedback. The goal of the feedback is to make the user understand when one of their actions have been registered. This can be done using any of the human senses.

Feedback can be done auditorial, visual, by using touch senses, and more. Different kind of feedback is suited for different scenarios, but they all have one thing in common: it should be instant. If the feedback takes too long the user might get frustrated and try other things, or even giving up.

Feedback should also contain sufficient information to explain what will happen next or what went wrong.

Lastly, too much feedback can be very frustrating and decrease the user experience, and too little feedback might be missed by the user.

Constraints

Interacting with a new object can sometimes be very tricky for a user. The chance of a user making an error with the application should be as slim as possible. Therefore, a designer should try to limit the interaction possibilities within the application. This is called adding constraints, which is done so that the user can figure out what do to without the risk of making errors.

2.3 Gestalt Principles

The psychologist M. Wertheimer formulated some principles about the human visual perception in the 1930's which are still very relevant today [11]. Many of these have been taken into account when designing our product, they are described below.

Proximity

The human brain perceives objects which are close together as being part of the same group, as shown in figure 2.3.

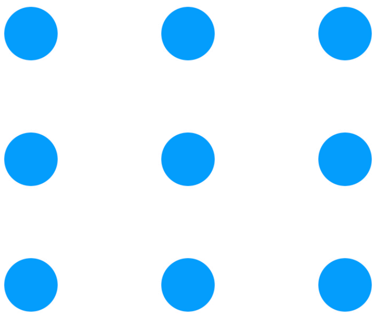

Figure 2.3: The law of proximity.

Similarity

The human brain perceives objects which have visual commonalities, such as color or shape, as being part of the same group, as shown in figure 2.4.

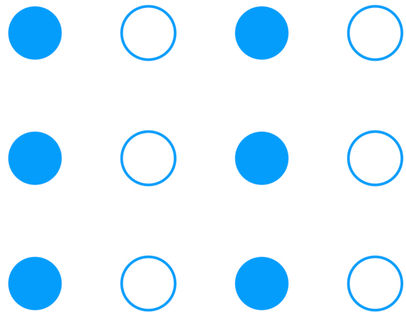

Figure 2.4: The law of similarity.

Common Fate

The human brain perceives objects which are moving at the same speed, and in the same direction, as being part of the same group, as shown in figure 2.5.

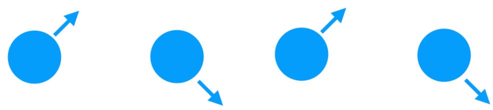

Figure 2.5: The law of common fate.

Direction and Continuation

The human brain perceive objects which appear to be part of a sequence or a path, as part of the same group, as shown in figure 2.6.

2.4 Data Visualization

2.4.1 Graphics To Understand And Explore Data

Graphics are excellent for exploring data as well as presenting the findings. The graphics used for the two different areas are usually different, and in this project focus was on exploring data.

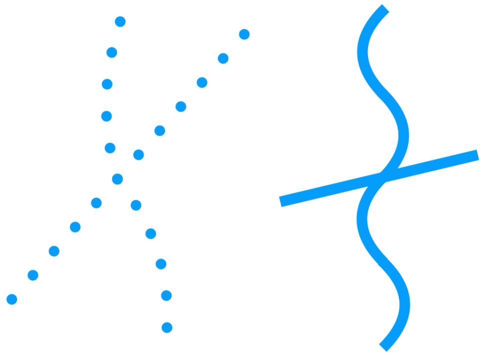

Figure 2.6: The law of direction and continuation.

Exploratory graphics are meant to be used by data analysts who look for results in huge volumes of data, in other words discovering and understanding data. Advancements in both computer software and hardware have opened up many opportunities in this field and a lot of methods have been developed, although it is far from fully exploited.

The amount of data collected by organizations is constantly increasing due to advantages in fields such as data collection and analysis, e.g., machine learning. An organization nowadays typically have a data warehouse containing various data sets from different sources. it is next to impossible for the human eye to, without help, analyse such huge volumes of data. Therefore, exploratory graphics are there to help gain insight in the data and to see patterns [12].

Regardless of what is to be shown in the data, there are (almost) exclusively these aspects available to show it: size, length, scale, color, position, shape and text.

2.4.2 The Basics Of Data Visualization

Analyzing data requires a lot of thinking as it is, and therefore, we want to shift some of the computations and reasoning which is done by the brain to the computer. For instance, when looking at a complex regular graph a user must first try and comprehend what the graph actually symbolizes, then understand the different parameters of the data, and finally analyze it. However, a more intuitive graph may bypass the first steps and let the user analyze the graph instantaneously [13].

The purpose of data visualization is to convert slow reasoning tasks into fast perception tasks. A good visualization can be used for analysis and decision making. Furthermore, it is of great advantage when communicating and presenting your research findings [13].

Data visualization happens automatically, we don't have to think to understand what we're seeing, whereas presenting data without proper visualization requires a lot of brain power to understand, analyze and make mental calculations on, as well as the need to store all this important information in the brain. All this is done for us by visualizing the data in a smart way [13].

For example, the process of identifying the largest number in a set of data is obviously much easier and faster done if the biggest number somehow is visualized. Consider figure 2.7 and try to identify the biggest number in the four data sets. One can easily spot the biggest number if it is plotted in a bar chart (figure 2.7 b), or if the number is written in bigger size (figure 2.7 c) or different color (figure 2.7 d). However, understanding the data takes significantly longer for a user which observes figure 2.7 a, since no form of visualization is used to help the observer [13].

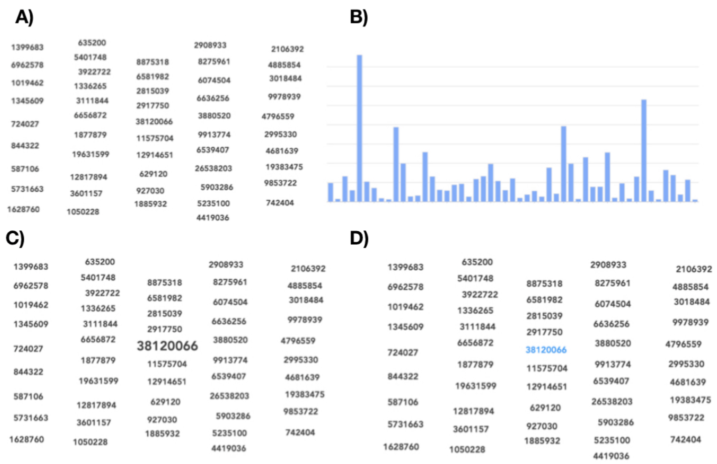

Figure 2.7: Why we need data visualization (adapted from [13].)

2.4.3 Visual Encodings in Data Visualization

There are some variables which make the foundation of data visualizations. These are position: size, shape, value, color, orientation and texture. Furthermore, there are some best practises for which scenarios these different variables are best suited [13].

Grouping by Category

When grouping data by category, one should use color, shape, icon, pattern, texture, enclosure and connection to visualize this [13].

Grouping by Relation

When grouping data by relation, one should use enclosure, connection, line pattern and line endings to visualize this. Furthermore, the same source states that the length of lines, as well as position and placement of objects, can be used to visualize quantities [13].

Using Size to show Quantities

The size of an object is a good way of visualizing quantities. Thus, size can be used to visualize how often a product is used, sold, and similar [13].

2.4.4 3D Visualization

Visualizing data in 3D is becoming increasingly popular. However, it is quite rare that the new dimension is used correctly. On the contrary, it often result in issues which easily could be avoided in 2D [14].

The main issue with 3D visualization is often associated with angles in some way. The human brain is quite bad at determining angles. Thus, it is often a disadvantage to look at data, for instance a pie chart, from an angle compared to from above. Consider the figure 2.8, d is prefered by far since the human brain have trouble interpreting the angles and determining the proportions between green and blue in the other figures [14].

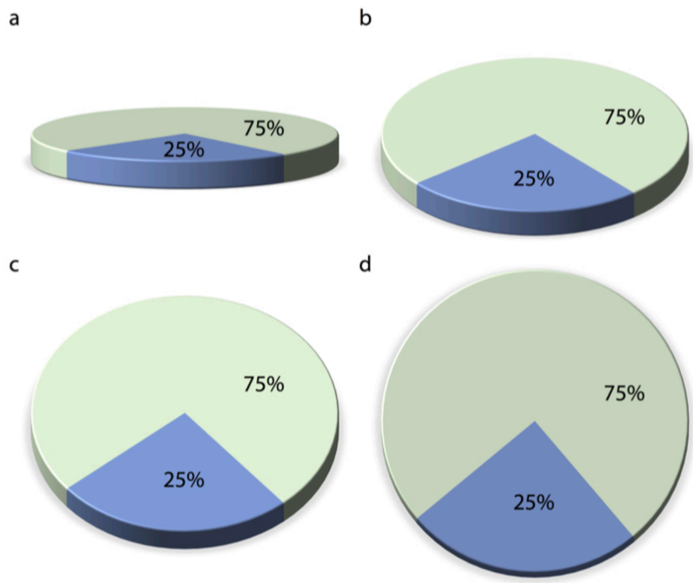

Figure 2.8: A bad use of 3D visualization (taken from [14]).

Additionally, visualizing data in 3D on a 2D displays is highly discouraged. Further misuse of 3D visualization will not be covered. Instead, the focus will be on how one can take advantage of 3D to better visualize data [14].

The pie charts discussed above does in fact not use the third dimension to something useful at all. On the contrary, it makes the chart worse. What the charts in figure 2.8 does wrong is that it should use all three dimensions to get the most out of a 3D visualization. Furthermore, for an optimal use of 3D visualization, one could present in Virtual- or Augmented Reality. The advantage here is that an observer can view the data from multiple angles, and the human brain is very good at reconstructing a 3D scene from a series of images taken from different angles. The more possibility for interaction with the data, i.e rotating, re-sizing and other tweaking, the better the understanding of the information will be [14].

2.5 Virtual Reality

Virtual reality (VR) is a computer generated digital environment that can be experienced and interacted with as if that environment was real. It is experienced through audio and visual effects provided by a computer, and one's interactions partially determines what happens in the virtual world [15].

There are plenty of different interaction mechanisms in VR. This section briefly describes the ones used in this project.

2.5.1 Walking

Walking is of course the most natural way of navigating around in a virtual environment. The user will walk in a real environment, and move in the same direction, and with the same distance, in the virtual world. It is a perfect *mapping* between the real world and the virtual one. The tricky part is keeping the user within the boundaries of the physical space. The usual VR set ups have an area of at least two square meters [16].

As soon as the virtual environment gets bigger than the area one has at disposal in real life, a need for some sort of locomotion is required. It is vital that this locomotion does not affect the user experience negatively. Thus, much research has been conducted and many locomotion techniques, such as teleporting, have been developed to move around in a virtual environment, each suitable for a specific scenario [17].

2.5.2 Point and Teleport

The point and teleport technique, which has proven to be user friendly, teleports the user to a location of his/her choosing, as the user aims at said spot in the environment. The technique does not visualize the locomotion, and thus, the user wont become nauseous. Furthermore, it is both fun and easy to use [17].

However, as it is easier to teleport than to walk, and more powerful in the way that you can move longer distances in shorter time, the user often becomes lazy and neglects walking entirely. This is not the purpose of teleportation, as this may remove the realistic and immersive feeling of being in the virtual environment.

2.5.3 Object Interaction in VR

Interaction with objects in a virtual environment can be done in multiple ways. Different mechanisms are suited for different purposes.

Two interaction functionalities relevant for this project are pressing virtual buttons and selecting objects.

There are of course many ways to interact with a button in VR. The most common, and perhaps most natural, uses physics. The best way to map a button click to how it is done in real life is to try to recreate the actual pressing of a button in the virtual world. This is done by creating a virtual object in form of a button, which can be pressed down with the VR controller. This becomes very natural to the user, aside from the fact that there is no physical resistance when pressing the virtual button, which obviously would be the case in a real environment.

Object selections can be done in various ways. Two selection mechanisms relevant for this project are grabbing and pointing. Grabbing, like all interaction mechanisms, requires two things. Firstly, indication of what product to be selected. Secondly, a signal or command to trigger the actual selection. The same goes for pointing, which is better suited for selecting items at a distance. The pointing is done with a laser beam, which the user can aim by using the VR controller [18].

It is vital to provide feedback of the action so the user knows when a product has been selected. The feedback can for example be to highlight the boundaries of the model or highlight the entire model [18].

2.5.4 Manipulating, Scaling and Rotating Objects

The interaction with objects in VR should be natural. A user should be able to grab and move an interactive object in the virtual environment just like he/she would do in a real world.

Scaling of an object can either be done with physical controls, like buttons, to increase or decrease the scale, or with hand motions. The latter is often preferred due to it being more intuitive and realistic. The method using hand motions is best suited for uniform scaling (all dimensions are scaled equally) and is built on the movement of the users hands (or controllers) in the virtual environment [18].

2.5.5 Fading Transitions

Scene transitions in VR should be done smoothly to not disturb the experience of being in the virtual world. It is important that the change of scenery does not influence spatial orientation, sickness, and the general sense of presence in an immersive experience.

A common way of doing this is the "Fade to Black" transition, which fades the screen to black, then changes scene and fades out again. This technique is described further and compared to other methods in [19].

2.6 Related work

Related work in the subject of data visualization in VR is rather limited. However, at Halmstad University in Sweden, M. Gustafsson and O. Odd have explored this in their thesis project [20]. Their thesis has many similarities to ours, even though their prototype is more focused on a general data visualization whereas ours focuses on visualizing relations between products, as requested by the stakeholders. Nevertheless, we have drawn the same conclusions on various vital design choices throughout the project, namely the importance of being able to walk freely in the virtual world, and the need for scalability and rotatability of the data.

G. Ryrlén's Master Thesis work at Lund Technical University used VR to visualize data for creating a financial instrument. The work process used is very similar to ours, and there are some resemblances between our prototypes even though the focus area is much different from ours [21].

Lastly, in 2014 at the University of California, a report was published which investigates how Virtual Reality can be used for immersive and collaborative data visualization. The report mentions some really important aspects which we have taken into account in our project and it gave us a lot of inspiration.

...it leads to a demonstrably better perception of a datascape geometry, more intuitive data understanding, and a better retention of the perceived relationships in the data.

This quote from said paper [22] is what made us consider solving the issue presented by the stakeholders with Virtual Reality in the first place.

Chapter 3

Technical Background

This chapter introduces the technical tools (hardware and software) which where used throughout the project.

3.1 Software Tools

3.1.1 Unity

Unity is a game engine which can be used to create both 2D and 3D games. It is a cross-platform engine with support to at least 27 platforms, including VR.

Unity game development includes scripting in C# and javascript, as well as drag and drop programming in the editor. C# was the prefered programming language in this project and Visual Studio Code (section 3.1.4) was used for script editing.

Unity has an asset store, which is a library with free and commercial assets created by either Unity or anyone in the community [23].

Before starting to develop, the following plugins were downloaded for free from the Unity Asset Store added to the project: SteamVR (section 3.1.2) and VRTK (section 3.1.3).

The latest Unity version is 2018.3.12, although version 2017.1.1f1 was used in this project to support VRTK.

3.1.2 SteamVR

SteamVR is available for free on the Unity Asset Store and contains scripts and models which eases the process of getting started with VR implementation, mainly with the integration of the VR hardware [24].

3.1.3 VRTK

VRTK is an open source VR toolkit containing scripts and concepts to simplify development of 3D games in Unity. To be more specific, it solves, or helps solving, the following aspects in a 3D game: locomotion, interaction with objects, interaction with UI, and body physics. Furthermore, it has a simulator to enable testing without connected VR hardware [25].

3.1.4 Visual Studio Code

Visual Studio Code is a code editor for Windows, Linux and MacOS which includes the following features: debugging, embedded git-control, syntax highlighting, intelligent code completion and code refactoring.

Visual Studio Code supports almost every major programming language, including C# which is used in this project.

3.1.5 Git and GitHub

Git is a version control system used in software development to track changes, enable collaboration among multiple programmers and more [26]. GitHub is an online hosting tool for version control using Git. It offers all functionality that comes with Git, along with some features of its own [27]. We used GitHub in our project to store our source code and to enable simultaneous development on the project.

3.1.6 Adobe XD

Adobe XD is a vector-based tool used for creating and prototyping user experience for web, mobile apps and more. It comes with features such as grid, animations and symbols [28].

3.2 Hardware Tools

3.2.1 Virtual Reality Headsets

An advantage with SteamVR is that it supports almost any VR headset, so we did not have to limit ourselves to a specific one.

We used three different headsets throughout the project, which all had its own advantages. These were the HTC Vive Pro, Pimax 5k Plus and the Samsung HMD Odyssey.

HTC Vive Pro is one of the most used VR headsets on the market. It consists of the headset, two hand controllers and two base stations, also known as lighthouses.

The resolution of the headset is 1440x1600 pixels per eye, it has a refresh rate of 90 Hz and a field of view of 110 degrees. Furthermore, it has integrated headphones and a microphone.

The controller have three interactive buttons, the trigger button, the grip button and the touch pad, together with two menu buttons, as displayed in figure 3.1.

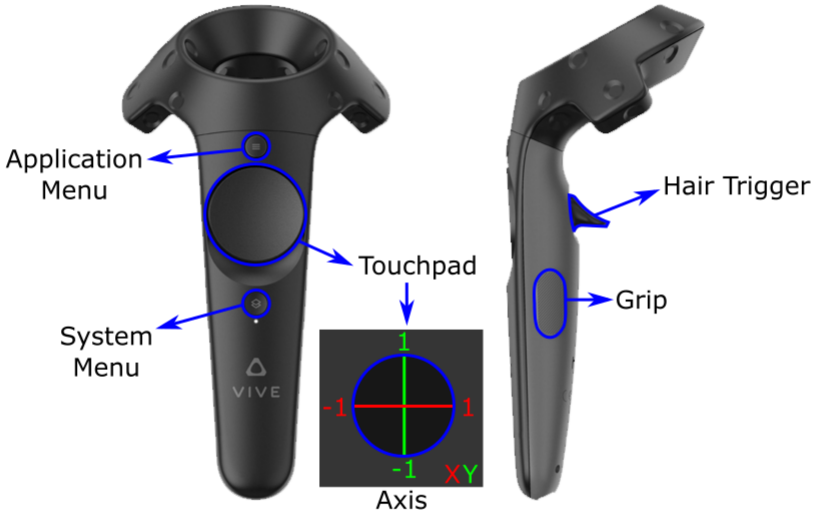

Figure 3.1: The HTC Vive controller which have been used together with the HTC Vive Pro- and Pimax 5k plus headsets throughout the project [29].

The lighthouses measures the play area and tracks the user in it. It provides a 360 degree play area, with an area of at least to 2x1.5 meters [29] [16].

We used the same controllers and lighthouses when using the Pimax headset. The Pimax headset is shown in figure 3.2 and it has a resolution of 2560x1440 pixels per eye and a refresh rate of 90 Hz. However, it does not have any built in audio, so external headphones were used together with this headset. The big difference here is the field of view, were the Pimax is significantly better with its 200 degrees compared to the 110 of the HTC Vive Pro [30].

Figure 3.2: The Pimax headset which have been the favored headset in this project, much because of its impressive 200 degree field of view which increases the sense of realism in the application [30].

The third headset, Samsung HMD Odyssey, is quite different to the other two. The screen resolution is 1440x1600 pixels per eye, the refresh rate is 90 HZ, the field of view 110 degrees, but most importantly, it does not require any lighthouses or other external sensors to track the user. Instead, two cameras on the headset track the user which can be seen in figure 3.3. This results in one big advantage compared to the other two: portability.

However, this way of tracking comes with some disadvantages as well. The main issue is that the cameras might lose track of the controllers if they are not visible [31].

Figure 3.3: The Samsung HMD Odyssey headset, used for testing outside of Jayway because of its great portability [31].

Chapter 4
The Work Process

4.1 Project Disposition

This chapter explains the entire work process, all the way from an idea to the finished application. Each section represents a phase during the project and explains what we achieved during it and how. The work process consists of three phases the are executed sequentially. These are: Concept Phase, Design Phase and Implementation Phase (see figure 4.1), where the latter two are done in iterations.

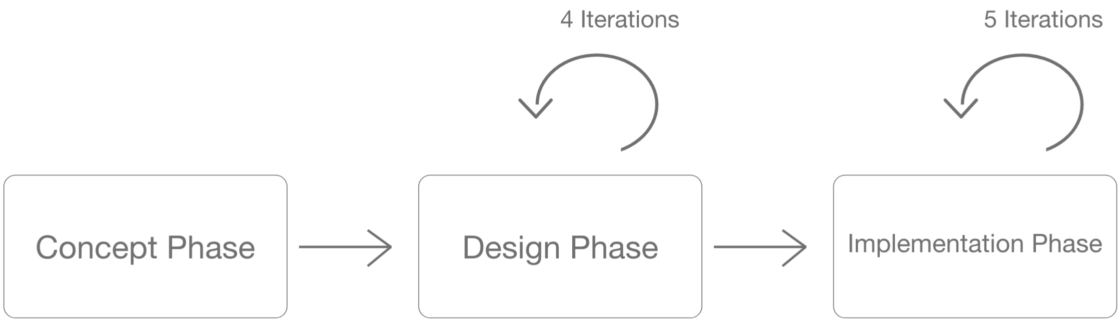

Figure 4.1: The three phases of the user centered work process used in this project.

The project originated from a vague idea that there probably are various improvements that can be made to the visualization of product relations in the sales data at IKEA. However, there had been no discussions between Jayway and IKEA regarding how this could be done. Thus, the beginning of the project included multiple meetings with IKEA, LTH and Jayway to get everyone on board with the idea and scoping of the project.

As we went into the project we knew very little about the problem at hand but after meetings with supervisors at Jayway and IKEA we were able to formulate an explicative project description.

During this phase we also set up goals, listed potential challenges and made a project plan, with the help of a MethodKit card deck, which made us discuss and prioritize almost every aspect of the project [32].

Figure 4.2: Some of the discussion points in the MethodKit project planning card deck [32].

MethodKit helped us define deadlines, limitations, goals, possible tools, methods and other topics, some of which can be seen in figure 4.2. We sat down with the card deck and discussed each card, or topic, during an afternoon. The goal was that each topic that might be relevant to our project should be discussed, mainly to make sure we were on the same page and had the same expectations on the whole process. The card deck, did among other things, make us discuss things like the work structure. It became clear the project should be divided into three phases: concept-, design- and implementation-phase, as described in chapter 2. The three following sections, 4.2, 4.3 and 4.4 describe how and what we did in each of these phases.

4.2 Concept Phase

The concept phase was initiated with a meeting with the project supervisors at IKEA. Some of the discussion points were scoping of the project and future project steps, such as interviewing, observing and testing.

We agreed upon the scope and that our end users would partake in interviews, observations and testing of the product during the project.

In this phase we were also introduced to our end users. The end users use data of product relations in their everyday work to make decisions in regards to what to produce and how and where to market different products.

- *Participant 1* has the title "Business Navigator", which means he uses the data to facilitate business development, plan for the future and explore future scenarios, in regards to product range.

- *Participant 2* and *3*, have the title "Sales Responsible". They focus on marketing and use the data to make decisions such as what to commercialise, provide price, quantity and in which countries new products should be sold. They also look at with which current products a new product should be sold in order to maximise sales and customer happiness.

In order to better understand IKEA's product range, and thus be able to make a more tailor suited solution to the end users, we got introduced to IKEA's product category hierarchy. A top down hierarchy that becomes more specific as you move down the hierarchy is used as can be seen in figure 4.3. Each BA has their own set of HFBs which, in their turn, each have their own set of of PRAs and so on. Under PA comes each actual product, described by a name and a unique product number.

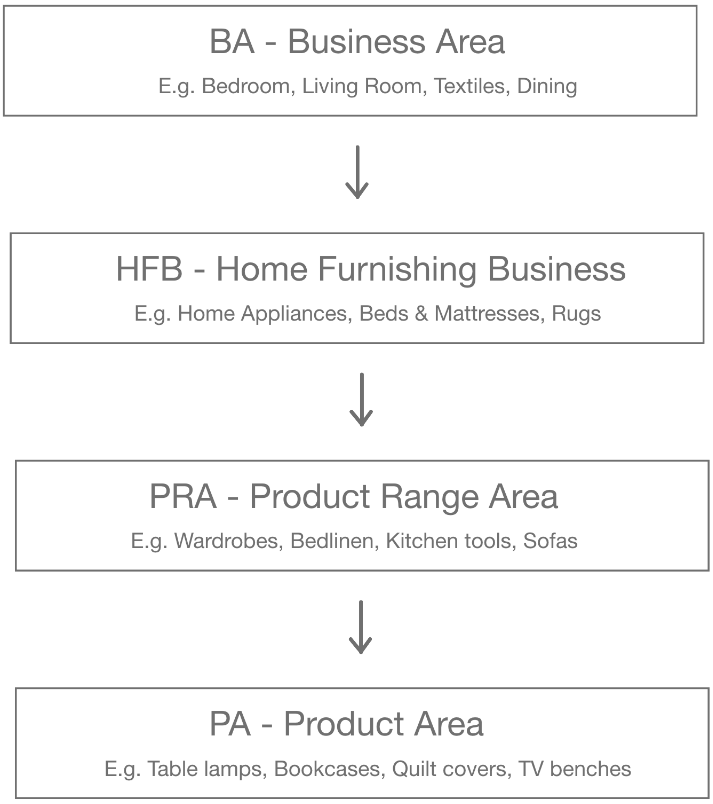

Figure 4.3: IKEA's product hierarchy. IKEA categories their products using these four categories.

The next step was to interview and observe how the end users currently analyze the data, and therefore, another trip to IKEA's facilities was made to conduct these interviews. Discussion points based on information from the previous meeting had been prepared beforehand. The questions were formulated according to the research presented in section

2.4.1 and 2.4.2. The goal of the interviews was to further increase our understanding of how the end users currently access and use the data and what they wish to be able to read from it.

Following the interview, we proceeded with observing the end users interact with the current system to see what issues they faced. They were encouraged to describe their thought process while performing their daily tasks, as described in section 2.1.1.

The most imperative information from the interviews and observations have been summarized below. These are their current frustrations:

- It is time consuming to look up multiple relations.
- It is difficult to keep track of data from multiple searches.
- There is no way to get an overview of how multiple data relations relate to each other. One must make multiple searches in the data.
- The huge amount of products makes it hard to extract valuable data when searching for relations.
- The end users sometimes have to make assumptions and/or use their gut feeling to make decisions in their work.

We also realised that they already had some quite precise expectations on the solution. They already knew what they wanted to achieve with the data, but the issue was their current solution could not fulfill these requirements. These are the goals they hope to achieve with our application:

- The end users want to use relation data to improve marketing and make it more specific towards the customers.
- The end users want to use relation data to make the product range more specific to the customers.
- By looking at relation data, the end users want to make decisions such as which products to start or stop producing, how to combine products when marketing in stores or online, and to understand what product range that should be available at their different stores.

One of the most important observations we made is that the end user often need to know what relation he/she is looking for beforehand. The current model for data analyzing makes it very hard for the end users to find relations or patterns that he/she was unaware of beforehand, which is not very optimal in terms of exploring new relations within the data 2.4.1.

From the limitations of the current system and with the input from our end users we created a short list with features we wanted our product to have:

- Provide an overview showing multiple product relations. Not just one-to-one relations.
- Innovative and intuitive navigation.

- Optimizing the probability of the user finding relations or patterns he/she did not know existed.

As a final step in this phase, we developed a persona from the information gathered, as seen in figure 4.4. The reason for this was to make sure all parties involved, for instance ourselvecs and the stakeholders, were on the same page regarding what the issues were and what expectations there were on a solution,.

In addition to the persona, we set up user scenarios to further empathise the frustrations the persona is facing, and to make sure all parties were on board with what the issues are.

We chose to describe the scenarios as user stories in story boards to better capture the feelings of the persona in different scenarios, as shown in figure 4.5.

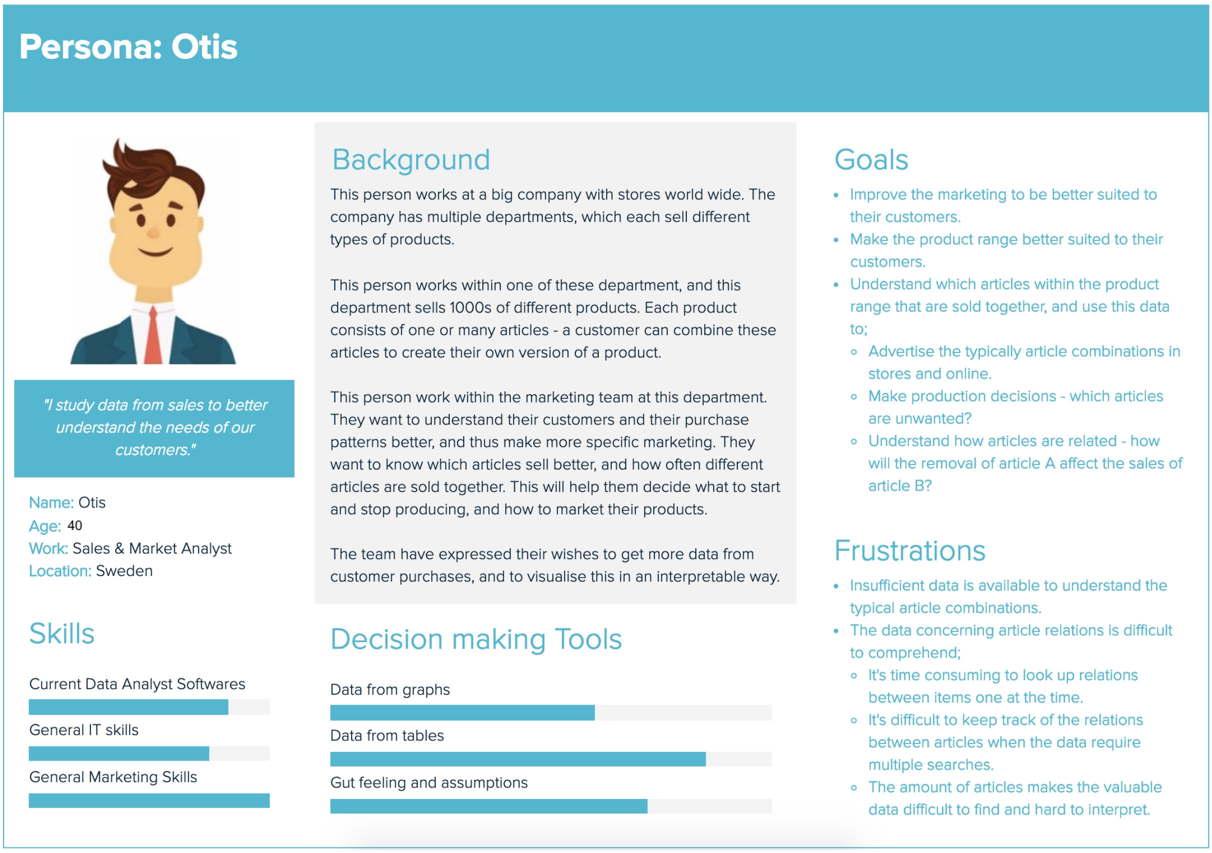

Figure 4.4: A persona describing a typical end user of our application.

4.3 Design Phase

The design phase is iterative and follows the six steps described in section 2.1.2.

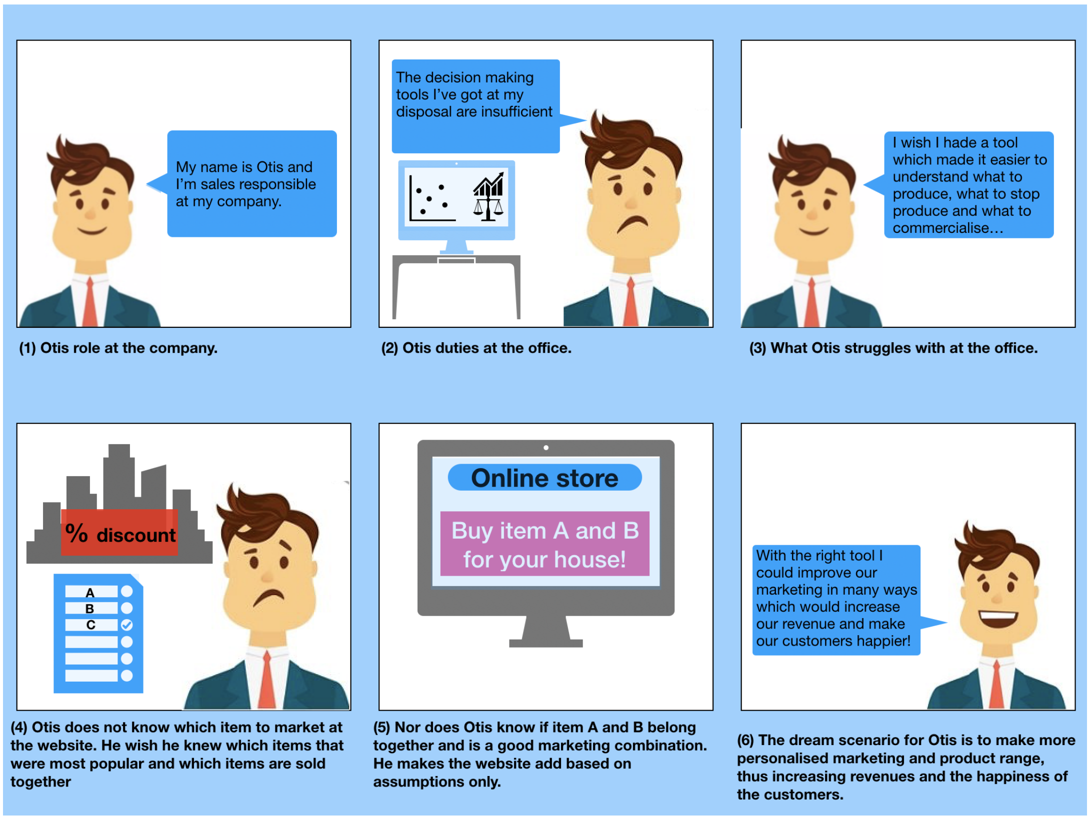

Figure 4.5: A storyboard describing the everyday work task of studying article relations which our persona Otis struggles with.

4.3.1 Iteration 1

Defining the problem

The concept phase gave us a good understanding of our end users. As we went into the design phase we knew what issues they faced with their current system and what expectations they had on a new one. The problem was defined as *How can we present data in an innovative way so the end users can better understand relations between items within big data sets?*

Collecting information

Since it was clear from the *Concept Phase* that the focus would be data visualization research in this are was conducted.

We got better understanding of the advantages of it, and what one have to keep in mind when visualizing data. This is covered in section 2.4.2.

Initially we tried not to focus too much on what had been done in the field already. On the contrary, what has not been done? We wanted to come up with something innovative and unique, as data visualization is a huge field with tons of research the past century already.

We started by looking at some of the latest UX-trends, with the goal of coming up with numerous unique ideas on how to solve the problem. We gathered information online and pursued to the next step in the design process: brainstorming.

Brainstorming

We were given a lot of freedom regarding how we wanted to solve the problems our end users currently face when analyzing product relations. A two hour brainstorming session was held with representatives from both Jayway and IKEA. Joining was our supervisor on Jayway, a senior developer from Jayway active in multiple IKEA projects and our supervisor on IKEA, who initially presented IKEA's problem to us.

One idea that stood out, and all parties supported was to visualize the data in 3D, by using VR. At this point there was no concrete idea how this would be done, and thus, no prototype was made. Instead, the idea was presented to employees at the Jayway office who are experienced in the field to get their opinion and feedback.

Evaluating

The idea we came up with in the brainstorming sessions was evaluated in a 30 min long discussion with four junior developers at Jayway. They were all aged 25-30 and none of them had been present in the brainstorming session. We presented our initial idea of visualizing data in VR. The idea was met with enthusiasm and many thought VR could solve the issues our end users faced. The most valuable feedback we received was that:

The advantage of doing it in 3D is the third dimension. You must find a way to use it to your advantage

4.3.2 Iteration 2

Defining the problem

The problem we wanted to solve in this iteration was how to use the third dimension to our advantage when visualizing product relations.

Collecting information

Further research about data visualization was made, and more specifically about visualizing relations. To visualize relations between two items, one should use enclosure, connection, line pattern and line endings, as described in section 2.4.3.

The next step was to take this information and apply it to a 3D environment. How can we combine this information with the possible advantages of 3D?

The most important aspect from our research on 3D and how it should be used to be an upgrade from 2D was that *The human brain is very good at reconstructing a 3D scene from a series of images taken from different angles, and therefore it's important to be able to view data from multiple angles* as described in section 2.4.4. Our hypothesis here is that by walking around in a 3D environment and viewing data from different angels, one will get a better understanding of it.

Lastly, research was conducted to find similar projects to see what had been done in the field already (see section 2.6). We took their conclusions into account when brainstorming on how to create our own application.

Brainstorming

A two hour brainstorming session with just the two of us was held at the Jayway office. The goal of the session was to come up with ideas for how a potential VR environment could look like. We used post-it notes and a whiteboard to speculate and combine different ideas we had.

Most of the ideas we discussed involved the user standing in a empty room, with some kind of data cloud around him/her. The person should be able to walk around in this area to view the data cloud from different angles. However, it was tricky to decide what the data cloud would actually look like.

An initial idea was a data cloud were the distance between two words, i.e, two products in the cloud would represent the relation between them, and that the words would be connected with a line. The length of the line would resemble their relation as described in *Grouping by Relation* in section 2.4.3. A longer line would mean a lower relation probability whereas a shorter line would mean higher relation probability, or higher affinity, i.e., the two products are often sold together.

With the persona described in section 4.2 in mind we figured a data cloud could potentially give the user an overview of multiple relations, which was both hard and time consuming in the system the end users currently use.

Evaluating

An evaluation session just like the one described in the previous iteration was held. A discussion, where we presented the ideas we had come up with in this iteration, was held with the same colleagues as last iteration. In this discussion we learned two very useful things about VR. Firstly, reading text in 3D can sometimes be tricky. Secondly, if we wanted our application to be in real-life-scale it would be problematic to walk around in it without hitting objects in real life.

4.3.3 Iteration 3

Defining the problem

The goal of this iteration was to solve the two issues mentioned above, as well as creating a lo-fi prototype we could bring to the scheduled meeting with our end users.

Collecting information

There were no real solution to the somewhat problematic issue of reading texts in VR. However, there had been much research made about locomotion in VR, see section 2.5, and we came across many different solutions which we discussed in the brainstorming session.

Brainstorming

There are multiple ways to move around in Virtual Reality already. The one we felt suited our application best was a combination of walking around and using a VR controller as a pointer to teleport around the room, which is described in section 2.5.

As for the issue of reading text in VR, we came up with the idea to have each product name be inside a sphere. Thus, it would make it a bit easier for the user to read as well as telling the objects apart from each other in the environment.

Developing

An interactive web based lo-fi prototype was made using Adobe XD (section 3.1.6). It was a bit tricky to create a 2D prototype that could simulate what we wanted our application to look like in 3D. We were a bit concerned that the end users would find it hard to comprehend what a VR implementation of this prototype might look like if they lacked VR experience. Therefore, we chose to focus on trying to make the interaction with the lo-fi prototype be just like it would in VR, for example by making objects that would be interactive in a VR implementation interactive in the lo-fi prototype as well, but with the computer mouse instead of a VR controller. For example, the user of the lo-fi prototype can use the mouse to click on the different products (or spheres, named A to G).

When clicking at a sphere, it becomes highlighted along with its relations to other spheres, i.e., other products. This type of interaction would be possible in a VR implementation as well: using the VR hand controllers as laser pointers to highlight spheres, which is illustrated in the prototype, see figure 4.6.

Each sphere's location in the room was determined by its affinity to the other spheres. Affinity here can be seen as the probability that two different products are sold together. Therefore, the greater the affinity between two spheres the closer to each other they appear in the room. In accordance with *direction and continuation* described in section 2.3 lines were used to symbolize which spheres had a connection. To better *signify* to what extent two products were related the lines were color coded. The color scale went from green to red, where green meant high affinity, and red low.

Evaluating

To evaluate our idea and prototype in this iteration we held a two hour meeting with our end users where we presented our prototype and our idea of visualizing data using VR. After our presentation an open discussion was held where we received feedback on our idea as well as our prototype. All three end users were positive to the idea of using a data cloud to symbolize data relations. They all agreed that it could solve some of the problems they presented to us in *Concept Phase*, such as:

- It is time consuming to look up multiple relations.
- There is no way to look at how multiple data relations relate to each other.

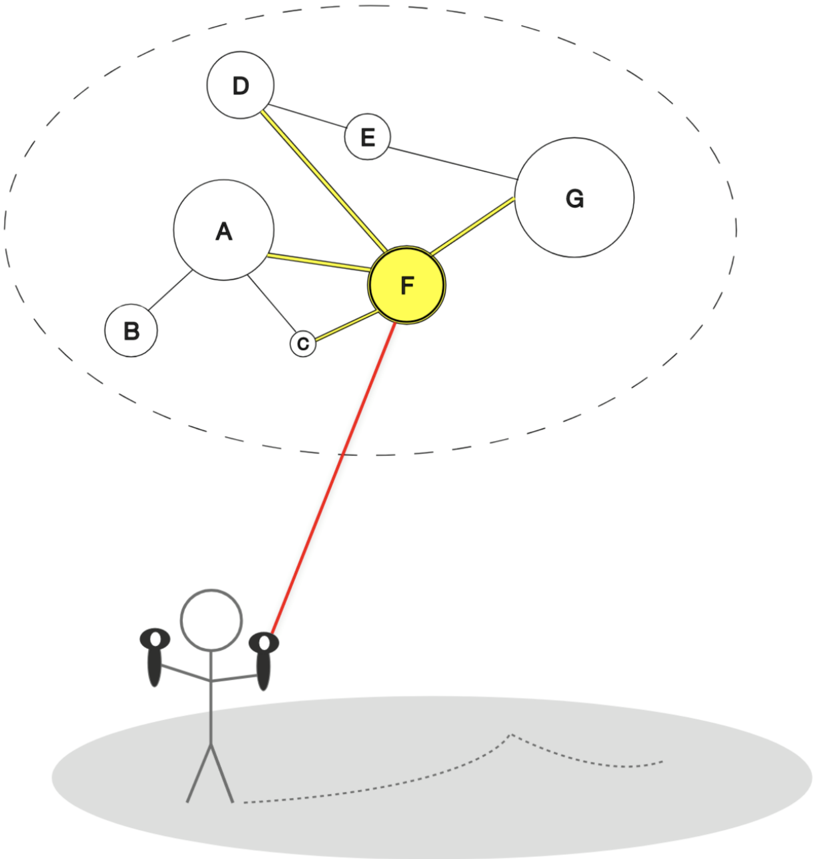

Figure 4.6: A lo-fi web based interactive 2D-prototype. The user can interact with the cloud using a mouse. The laser pointer will point to the sphere the user clicked with mouse at, showing how interaction using VR hand controllers would be done in a VR implementation

One participant said that *Using a data cloud showing multiple product relation at the same time would make it easy to see to which products a certain product have its strongest affinity with.*

Another participant explained that since every product relates to other product which in their turn relates to other products this quickly results in complex patterns that are impossible to understand by looking at multiple tables or 2D graphs (which the current system broadly can be described as). A graphical representation is needed if one is to understand these complex product relation patterns.

All three end users were in agreement that using the cloud to show multiple relations at the same time probably could save them time in their everyday work since less time will be spent simply navigating the system and selecting which products to visualize.

They also gave of ideas for further implementation. The two main things they wished for in the application was:

- Some sort of selection and filtration is needed. They have thousands of products in their product range and these are grouped by categories and subcategories. How should one select which product to visualize in our application? A big cloud of all articles would be impossible to get an overview of.

- To get the data as numbers too, not just graphically. For instance, they wish to see how often a product is sold and the affinity between two products.

4.3.4 Iteration 4

Defining the problem

There must be a way to select which products to visualize in our data cloud. Additionally, more clear numbers should be presented in the visualization.

Collecting information

From the previous iteration we knew that the end user wanted clear numbers in addition to lines and distances to represent product relation probability.

As for how often a product was sold, research was conducted to find an intuitive way to visualize this. We found, as stated in section 2.4.3, that the size of an object can be used as a measurement of how much/often something is sold.

Brainstorming

A new brainstorming session was arranged, and just like in iteration 2, it was between the two of us during an afternoon at the Jayway office. The goal was to find a way to select which products to visualize in the data cloud. It was clear we wanted to avoid regular 2D text based menus as this would be a bit exhaustive and slow to use in VR, not to mention how it ignores the third dimension which we have at our disposal.

With the possibilities of VR in mind we knew we had the right prerequisites to create a innovative and intuitive way of selecting what data to visualize. Using IKEA's products: furnitures, as a starting point one idea that came up was of the user walking around in a fully furnished house and simply picking which product model to visualize. The initial idea was a complete house with different products placed at their natural position in a home, and the user had to walk around and pick the desired products.

However, we quickly realised that this idea had one major flaw. In a house you usually have one table and one sofa, not 40 different articles of the same kind of furniture which is the case for IKEA. Our virtual house would be too overwhelming and complex with everything there at once.

We figured a better way would be to make a combination of a menu and a house. Each room in the house would represent one menu selection, i.e., choosing one product family. For example, the first room would have one model of a sofa representing the category *Textiles* and one model of a bed representing the category *Bedroom*. By selecting the bed the subcategories enabled in the next virtual room would relate to the *Bedroom* category. These subcategories would also be represented by models, for example pillows and covers. An overview of the selection procedure can be seen in figure 4.7. As described in section 4.2, BA, HFB, PRA and PA is how IKEA internally categories their products. It is a top down hierarchy in said order that becomes more specific as you move down in the hierarchy. To select a product family the according model is grabbed and placed on the pedestal. This unlocks the door to the room with subcategories (see figure 4.8). The same procedure is done in the next 3 rooms after which the specific product family the user chose is presented in the *Visualization Room* (see figure 4.7).

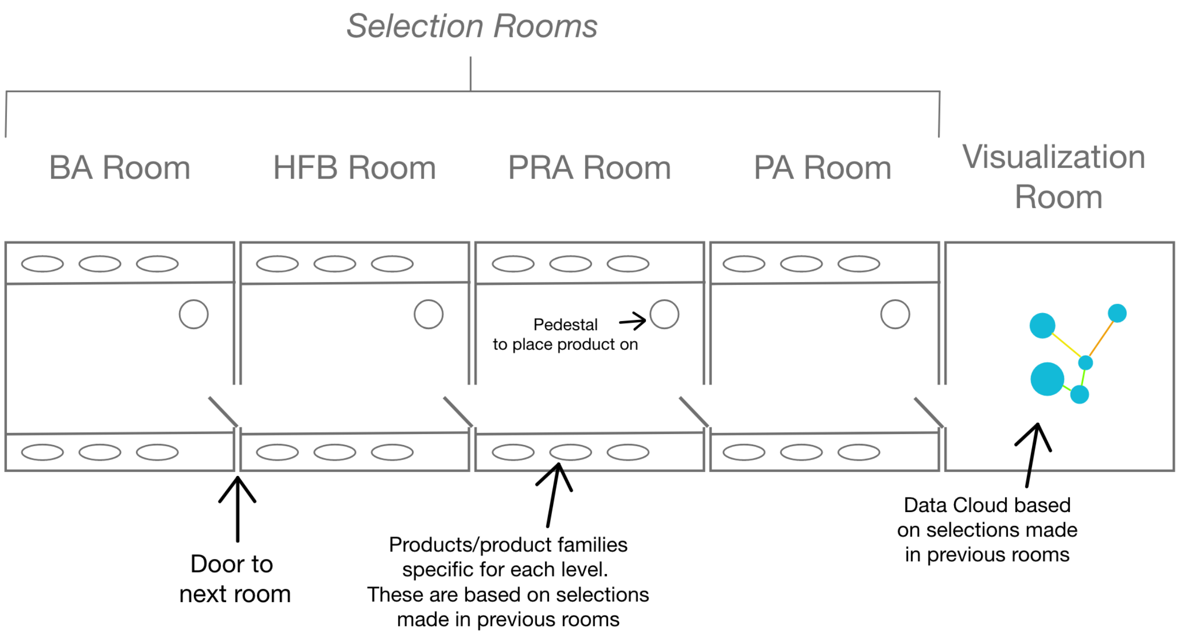

Figure 4.7: The selection procedure. Each room represent a category level where placing a product on the pedestal unlocks a new room with that product's subcategories

We believe that there are multiple advantages of selecting your product this way instead of from a list of product numbers for example. For starters, it's easier for inexperienced people within the company to get acquainted with all the products if you see them as 3D-models, compared to just a product name. Furthermore, this solution would increase the understanding of IKEA's product range for people outside the company, like partners, or even customers.

We also speculated about how to display the clear numbers, and came up with the idea to display a menu next to a sphere once the user points at it. However, how this menu should look like and what else that could be displayed here was unclear at this point. We decided to prioritise other aspects of the design in this iteration and come back to this at a later stage.

Furthermore, we wanted to graphically visualize how often a product is sold. Inspired by the idea of *Using Size to show Quantities* described in section 2.4.3 we decided that the graphical visualization would be achieved with different sizes of the spheres. Thus, the bigger the sphere, the higher the sales numbers.

Developing

We added the following functionality to our prototype: spheres scaled in accordance with their sales numbers and a selection house with rooms for each category-level. The lo-fi

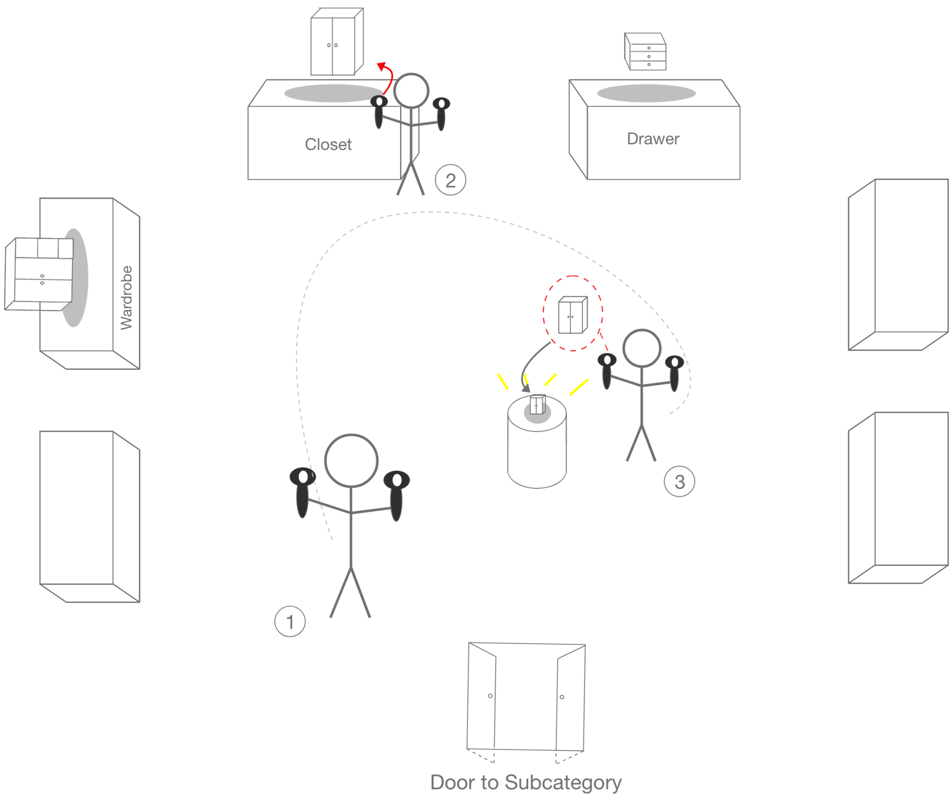

Figure 4.8: A lo-fi web based interactive prototype of how the user selects which products to visualize. The user can use the mouse to click on different products. The selected product is picked up and placed on the pedestal by the stick man to symbolise how interaction with products would take place in a VR environment.

prototype of the *Selection Rooms* can be seen in figure 4.7 and 4.8.

Evaluating

The updated lo-fi prototype was presented to our end users. They found the idea of selecting what to visualize by grabbing 3D models in rooms structured in accordance with their product hierarchy to be exciting and intuitive. The end users came to the same conclusions we argued for when designing the lo-fi prototype, namely that this selection procedure would give the the user great insight and overview of IKEA's sortiment.

With the lo-fi prototype tested and approved by the end users, and with the time limit of our project as a constantly present factor we found no need to test the lo-fi prototype on colleagues at Jayway. We therefore moved on to the implementation phase.

4.4 Implementation Phase

The following goals were set up before we started implementing:

- A user should be able to select which product to visualize in an intuitive way.
- A user should be able to tell which products are closely connected, or have high affinity to one another, without looking at clear numbers.
- A user should be able to tell if a product is sold more or less than another product without looking at clear numbers.
- A user should be able to move in the virtual space to explore the data and look at it from different angles.

4.4.1 Iteration 1

Defining the problem

The goal of the first iteration was to recreate the first part of our low-fi prototype in VR, the product selection mechanism.

Developing

Environment Setup - Before we began programming, we had to set up the environment and decide which tools to use. We did not choose one specific headset which we developed for, since Steam VR is compatible with multiple ones. Througout the project we used three different headsets, HTC Vive Pro, Pimax 5k Plus and Samsung HMD Odyssey, which are described in section 3.2.1.

The Pimax was actually prefered throughout the project since we believed its field of view would increase the feeling of being in a real house. However, the Samsung HMD Odyssey is significantly easier to bring with you and set up at a new location. Thus, it was used during the testing sessions at IKEA's facilities.

We chose to work in Unity 2017.1.1f1 (see section 3.1.1) since this would enable us to integrate it with VRTK (section 3.1.3), a toolkit to simplify the implementation of controller interaction, such as teleportation or pointing, within the project.

We found two packages within the Unity Asset Store which contained 3D models of furniture, which we used in our product. These were Big Furniture Pack [33] and Simple Home Stuff [34]. Furthermore, the packages Shooting Gallery - Environment Pack [35] and HQ Modern Desk [36] were used to decorate the house.

Implementing Selection Rooms - When all this was set up we began programming the *Selection Rooms*. We made simple white rooms with walls, lights, benches and a pedestal. The benches were located along the sides of the room. floating 3D models of products were placed on the benches, as can be seen in figure 4.9. We chose to have the models floating in the air to signify that they were grabbable, as opposed to have them standing on the benches which could confuse the user into thinking the models were stuck there.

Since we wanted to encourage interaction with objects and movement in the room we decided a good way would be for the user to walk (or teleport) to a model, then pick it up and place it on a pedestal. For more on how to move and interact with objects in VR see section 2.5. We used VRTKs dropzone script to keep the models in place on the benches,

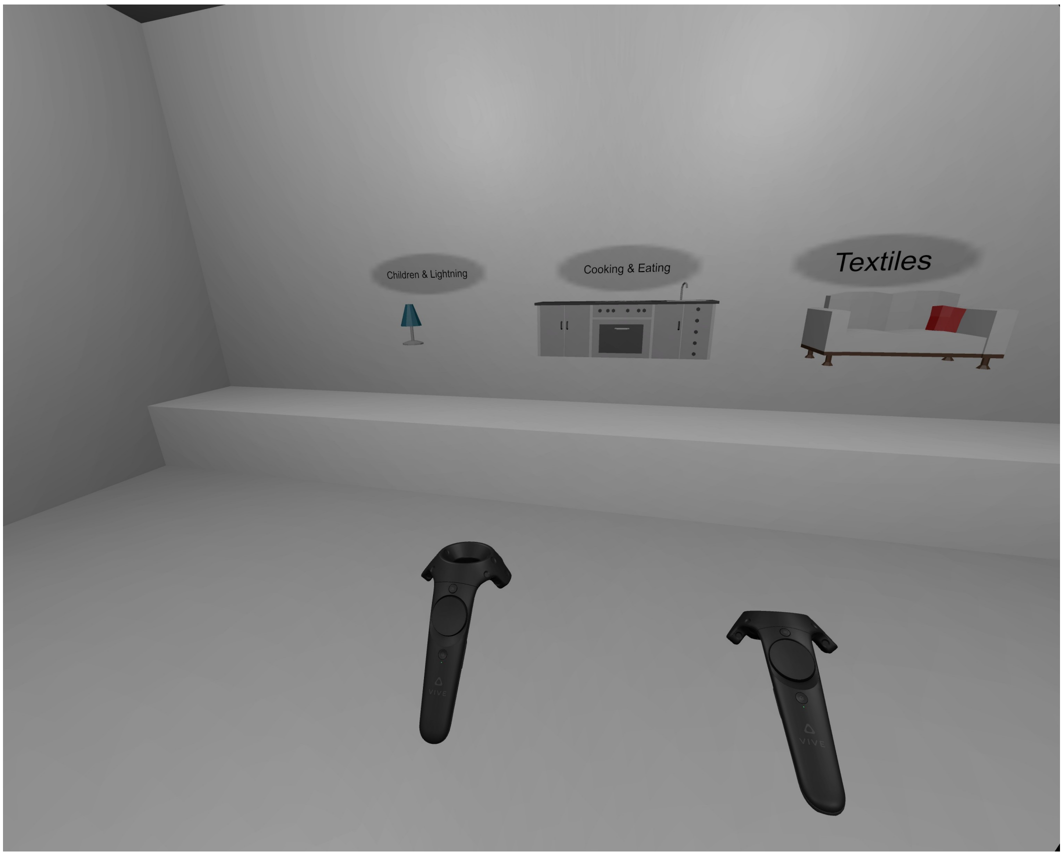

Figure 4.9: 3D models representing different BAs in the first *Selection Rooms*. The models are placed on the benches and have their name displayed above the model.

and another VRTK script to highlight the models which the user picked up. An outer highlighter was used when the hand controller was in contact with a model, this was used as a *signifier* (see section 2.2) to show the possibility of interacting with the model. A full, yellow highlighter of the initial start position was used when the user grabbed a model and moved it. The different highlighters are shown in figure 4.10. A final, empty, dropzone was placed on the pedestal, and when a model was dropped here, it would trigger what would happen next, which is that a door to the next room was enabled and shown and the user could now move on to the subcategory of the selected model.

Using dropzones reduced the risk of a user making mistakes, such as losing a product in the room or doing selections unwarily. The dropzones works as a *constraint*. For example it helps the user place the model correctly, in the center, on the pedestal which is required in order to lock in selection of a product. The dropzones makes it easy to place models on the pedestal, or back on its original position on the bench. Furthermore, it is not possible to place a model on another dropzone than the original. Additionally, the dropzones decreases the amount of precision needed to place a model somewhere, so the risk of dropping a model is also reduced.

Figure 4.10: The different highlighters used for product interaction. Left picture: no interaction with product. Middle picture: hand controller touching a product. Right picture: user grabbing and moving the product. The yellow highlighter showing the initial position for that product.

Evaluating

The application was tested continuously on employees at Jayway with experience in VR. Thus we got continuous feedback and tips on how to make interaction and navigation as intuitive and easy as possible. Some feedback and tips that resulted in changes in the application was how to highlight an object when touching, respectively moving it.

4.4.2 Iteration 2

Defining the problem

The goal of the next iteration was to make our VR-application ready for testing by the end users. Therefore, we decided to focus on implementing the second part of our lo-fi application, the *Visualization Room*.

Developing

A simple *Visualization Room* was created, containing just a plain floor with a data cloud in the air which the user could interact with. The data cloud is not a cloud per se, rather it is multiple spheres floating in the air.

The data cloud was created with a script. The script renders the data cloud by fetching data from a database, and then creating each sphere from it, using a reusable asset called a prefab.

The script scales the spheres depending on their sales value, as described in the lo-fi prototype. Inside each sphere was a text with the product name.

Lines were drawn between spheres to visualize that they have a high affinity to each other. Just like in the lo-fi prototype the length of a line tells the user how high the affinity is. The lines were color coded to further visualize to what extent the two spheres were related, as described in section 4.3.3.

Furthermore, we made a sphere highlight and rotate once it was aimed at by the laser pointer (see figure 4.11).

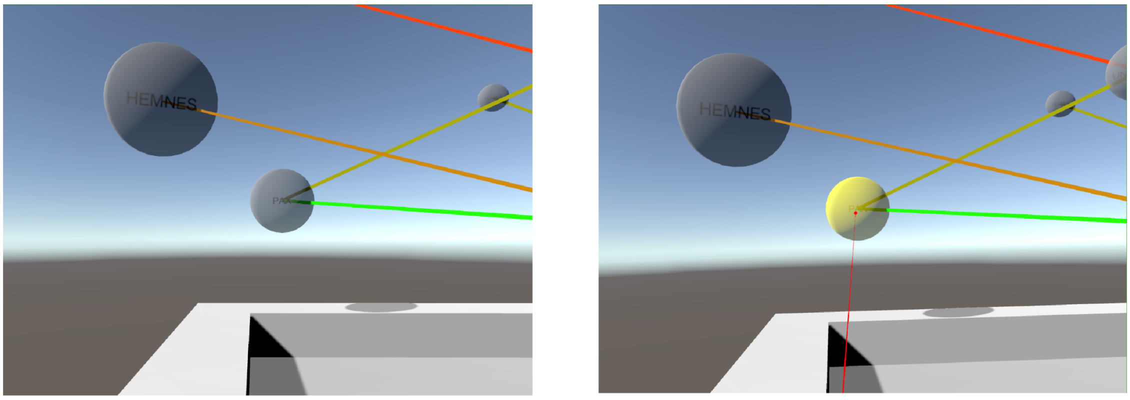

Figure 4.11: Highlighting of a sphere when aimed at with the VR hand controller.

Evaluating

To evaluate this current version of the application a one-hour combined testing- and brainstorming session was held with the three end users from IKEA, where they got to test the application. All end user had ten minutes each to try out the application. They all got the same two basic use cases they were to perform. These were:

1. Navigate the selection mechanism and select System Wardrobes from the PA to visualize these products.

2. Explore the data cloud to find which product has been sold the most and which product pair has the highest affinity.

From observing our end user we found that two of them, in the beginning of use case 1, had some slight problems grabbing objects in VR. The problem was mainly due to them having a hard time determining how far away an object they wanted to grab was. After a few trial and errors they both started to use the teleportation to get close to the object before trying to grab it. The end users struggle of determining the distance to an object is believed to have been due to inexperience in VR and not related the design of our application. This since the application was easily navigated by all end users once they got a few minutes to make themselves familiar with the headset, the hand controllers and a virtual environment.

All end users could easily read data from the data cloud. In observing our end user perform use case 2 no problems or difficulties where found

The feedback received in the discussion after the use cases was generally positive. For example, one participant said that:

It's very intuitive to understand connections compared to Excel. Now you don't have to read the numbers, analyze them, go back in the application, find a new relation, read the numbers, compare to previous number and so on. Instead you can see all connection at the same time. I think we'll be able to understand things which we could not before.

Another participant stated that:

It's much easier for people that are not used to reading and analyzing numbers, or not used to our product hierarchy, to understand product relations.

We also received some feedback of what could be improved and added to the application. These were the most important points:

- It should be possibly to visualize data relations from any category-level, not just the last one (level PA).
- It would be cool to grab the cloud to scale and/or rotate it in order to see things from different angels.
- There is a possibility the amount of spheres will be too many to be able to draw any conclusions in a real implementation.
- There should be some way to filter what's shown in the cloud (except for category), maybe just showing relations above a certain threshold?
- There were quite a bit of walking in the application, especially if you had gone down the hierarchy levels and then wanted to go back to the top and choose a new main category.
- To further increase the understanding for inexperienced users, there could be some sort of information in each room (e.g., what room you are in, what to do next and how to interact with objects).

4.4.3 Iteration 3

Defining the problem

In this iteration we wanted to add some functionality to simplify product selection, especially when going back from the *Visualization Room* to select a new product to visualize. Furthermore, it was requested that one should be able to visualize products from any hierarchy level. This means going to the *Visualization Room* from anywhere in the application.

Additionally, it was clear we needed a way to show which room one currently was in, since this was a bit confusing in the previous iteration.

Lastly, upon request from the end users, we wanted to add the feature of grabbing, scaling and rotating the data cloud.

Brainstorming

A major brainstorming session was held at the Jayway office. The session was divided into three parts, which focused on three different aspects. The two of us were involved during all parts. During the last part, which is regarding interaction in VR, we involved two colleagues at Jayway. Both of them work with VR and are very experienced.

Simplify movement between rooms - We needed to come up with an idea of how to simplify movement between different rooms since the user in some scenarios want to move up or down multiple levels in IKEA's product hierarchy (see section 4.2). The tricky part was to keep the "house metaphor" while still enabling movement across all levels.

We eventually came up with the solution to replace the regular house with a skyscraper, where each floor is a *Selection Room* for that category level, and the roof top is the *Visualization Room* (see figure 4.12). This would enable us to use an elevator to move between different levels in quick and intuitive way.

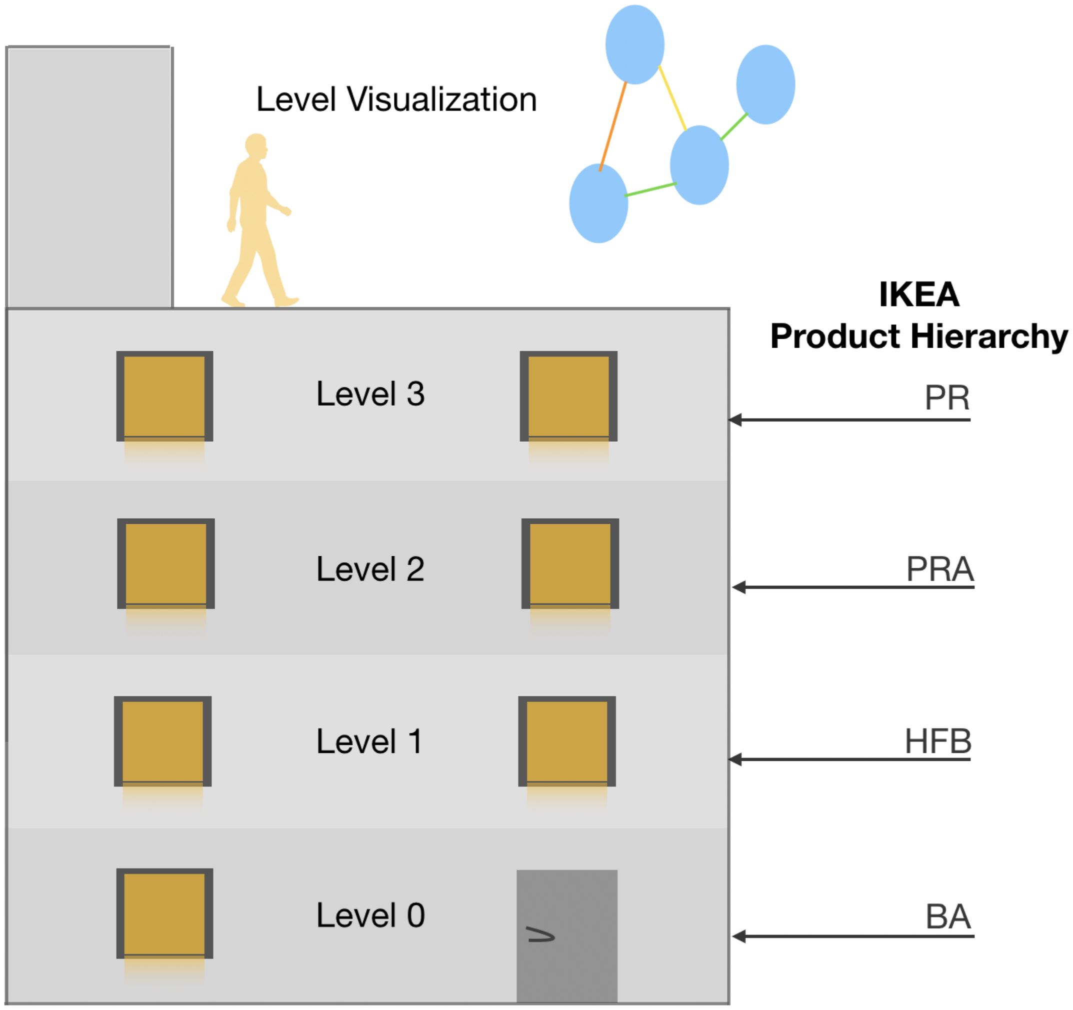

Figure 4.12: Each floor in the skyscraper is a product category in IKEA's hierarchy of categories. The rooftop is where the user visualizes the data.

However, unlike a normal skyscraper, one shouldn't be able to move between all levels at once. It is first when the user has selected a product that he/she should be able to go to the subcategory of that product, i.e., the next floor. The same goes for the visualization button, since the user also should have the possibility to access the *Visualization Room* and study relations to a selected product. Thus, the elevator buttons should appear as the user make choices on the current floor.

Figure 4.13 shows a scenario where all buttons are visible. The Visualization Button takes you to the top floor, or the skyscraper rooftop to fit the metaphor, where the data cloud is visualized.

Figure 4.13: The buttons in the elevator, which are enabled one by one as the user navigates up the hierarchy.

In addition to the elevator a desk was added in the middle of the *Selection Rooms* to make the room feel more like a control room. As the user selects a product at the current level, by placing a product on the pedestal, a button appears on the desk which takes the user directly to the *Visualization Room*. When a model of a product family from one of the lower levels is placed on the pedestal and the *Visualization Room* button is pushed all subcategories to this product family is shown in the data cloud. For example, if *Kitchen Appliances* is chosen on level HFB and placed on the pedestal, and the button on the desk is pushed, all product families on PRA level that was under- *Kitchen Appliances* will be shown in the data cloud. In this case that would mean there would be spheres like *Hobs* and *Ovens and cookers* in the data cloud.

Since the button on the desk simply takes the user to the *Visualization Room* it had the same purpose as the *Visualization Room* button in the elevator. The assumption was that more buttons would be required at a later stage, if the application should support the selection of multiple products, like a basket and therefore the desk was used as a type of control desk. Furthermore, the desk reduces the walking required by the user since he/she did not have to enter the elevator. A sketch showing the overview of the *Selection Rooms* can be seen in figure 4.14. As can be seen in the figure the door that previously led to the next room still exists. It is now stairs that leads the user to the next floor. Thus the skyscraper has both an elevator and stairs, where the stairs are suited for moving up just one floor and the elevator is suited for moving multiple floors.

How to visualize what room the user is in - The previous iteration suggested that it was not always clear for the user which level he/she was on. Since the elevator would let the user go to any floor instantaneously this would probably make it even harder for the user to keep track of what floor they were actually on. We therefore decided to have signs in the elevator, but also in the *Selection Rooms*, to help the user navigate around the building. The *Selection Rooms*-sign for the PA room can be seen in figure 4.15.

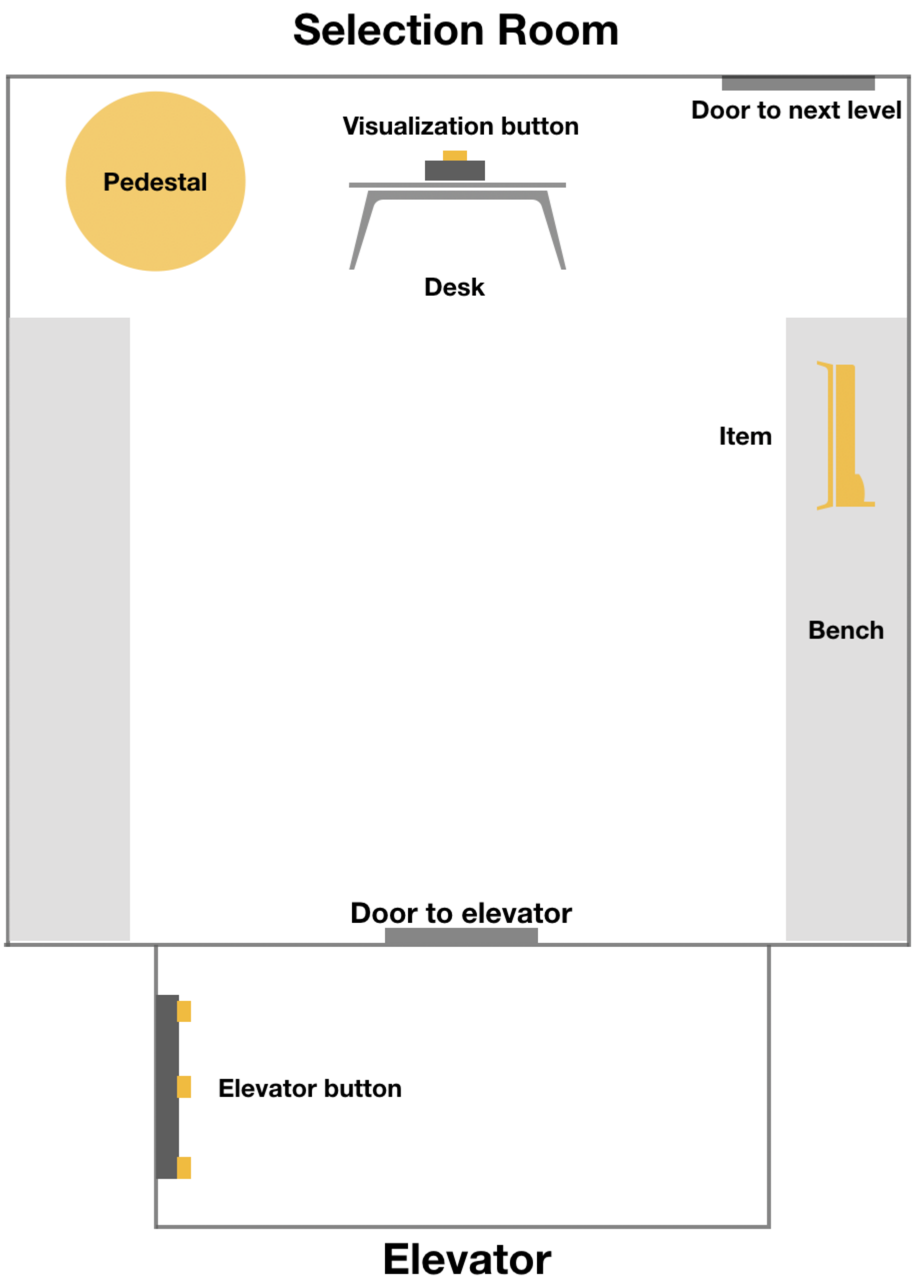

Figure 4.14: The *Selection Rooms*. The user selects a product from the current category by grabbing it at the benches and placing it at the pedestal. Selecting a product enables the user to visualize that product, or to go down the hierarchy to the product's subcategory.

Scaling and rotation of the data cloud - As for the scaling and rotation of the data cloud, research was conducted to find out how others have solved this in Virtual Reality, see section 2.5. We solved the scaling and rotation request in the following way: an object is selected and picked up using both hands. By moving your hands away from each other the object is enlarged. By moving your hands closer together the object is shrunk. In other words it works like a larger version of a pinch zoom. Furthermore, by grabbing an object with one hand and twisting and moving the hand-controller the object moves accordingly in the virtual world. Thus, the user can manipulate objects in an intuitive and easy way. This functionality exists in VRTK scrips, which we added to the data cloud gameobject, thus enabling the user to scale and rotate the whole clouds once he/she touches it. This method was strongly recommended by one of the colleagues from the brainstorming session.

Developing

An *Elevator Room* was created, which was directly connected with the regular room. Yellow elevator buttons were created and set to inactive and invisible at the start. As the user placed an object on the pedestal, i.e., selected a product on that floor, the button to take the user to the subcategory of that product was enabled in the elevator. Also the button to the *Visualization Room* on the desk was enabled, as seen in figure 4.15.

Figure 4.15: The *Selection Rooms* once a product has been selected. This action prompts the following to be displayed: the visualization button on the desk, the door to the next level, and the elevator button to the next level. The "Level Hierarchy"-sign is there to help the user navigate around the different floors.

The buttons were, just like the one on the desk, physically press-able with the hand controller, which is described more in *Object Interaction in VR*, section 2.5. As one began pushing down the button it was highlighted to give the user some feedback of the interaction. The buttons are displayed in figure 4.13.

Lastly, the appropriate VRTK scripts for rotating and scaling an object was placed on the entire data cloud.

Evaluating

To evaluate the work done in this iteration and the current version of the application a testing session was held at Jayway. Three master thesis worker at Jayway, all with experience in VR, tried out our application. The test was not structured with use cases as in the previous iteration. Instead each tester got to explore the application freely for ten minutes with just a short introduction to the idea and the application itself. This in hope of finding design flaws and implementation flaws that may go unnoticed when a tester is fully focused on completing a certain use case. After they all had tested the application we held a joint discussion where they could give us feedback on the application. The most important findings and feedback from the testing session is listed below:

- Two of the participants did not understand that anything happened once they pressed an elevator button, as the elevator looked the same and no feedback was given. More feedback was required here to make sure that it was clear the elevator moved to another level

- The participants found it far from obvious that the data cloud was interactive. Furthermore, it was confusing for the user to scale and rotate an object which was so big that the user sometimes stood inside it. A new way of interacting with the data cloud was needed.

4.4.4 Iteration 4

Defining the problem

Firstly, a new way of interacting with the data cloud was needed since the cloud object was so big, and without borders. It was hard to tell when one was close enough to interact with it, and sometimes, the user stood in the cloud without knowing, which made the interaction even harder. Additionally, the *affordance* that the data cloud is interactive must be improved.

Secondly, the elevator transportation must give some feedback to the user to make sure the level change is obvious.

Brainstorming

Interacting with the data cloud - First priority was to solve the interaction problem with the data cloud. It was problematic since the data cloud is so big and it does not have any borders that surrounds all spheres, thus making it difficult to make it look interactive. We came to the conclusion that having a miniature interactive data cloud *mapping* the whole cloud could solve the problem. By rotating the miniature cloud the big cloud would rotate in the same way, enlarging the miniature enlarges the big data cloud etc. This would save the user the problem of having to interact with the ungainly big data cloud.

The miniature data cloud is directly related to the data cloud, which is emphasised by showing a miniature of the cloud. Furthermore, the *common fate*-principle (see section 2.3) is applied too, since all the spheres in the big cloud is moved in the same direction and at the same speed as the interactive miniature data cloud.

In order to not confuse our user with "yet another" data cloud the miniature data cloud was placed inside a pink sphere. The bright pink color also worked as a *signifier* and made the sphere look interactive. The sphere was placed on a pedestal in the center of the room. Our hope was that this would increase the *affordance* of the interaction. In addition to the pedestal, a button to reset the cloud (rotation wise and scale wise) was placed next to it, to enable the user to start over when required.

A sketch showing the overview of the *Visualization Room* with the interactive sphere can be seen in figure 4.16.

Elevator Feedback - As for the elevator feedback, we tried to recall how an elevator ride feels like in real life. Since the physics of the transportation is impossible to recreate in VR, we had to use other senses. The first that came to mind was the sound, and thus we decided to play an elevator sound when the user used the elevator. The idea was to use a machine sound, followed by a 'ding!' when the transportation is complete.

There was not much we could do in terms of vision to imitate a real life elevator transportation. However, we came up with the idea to use a 'computer game'-metaphor, which is to fade the screen for a second or two when changing level, which is common in many video games when changing scenery. More information about *Fading Transitions* can be found in section 2.5.

Developing

Interacting with the data cloud - A pedestal was created and a rotation symbol from the unity store was painted on top of it, to increase *affordance* further. A miniature of the data cloud was made and placed inside a pink sphere. The same scripts that were described in last iteration were added to the pink sphere, making it grabbable, scalable and rotatable. When scaling or rotating the pink sphere the big data cloud would change accordingly.

The pink sphere was placed on top of the pedestal, and next to it a reset button was added to enable resetting the cloud to its original position, as seen in figure 4.17.

Elevator Feedback - The headset fade and sound were implemented in a script that simultaneously changed the floor. A sound as described in the brainstorming section was downloaded and used in the faded transition [37].

Evaluating

In this iteration we performed testing on four colleagues at Jayway. Preferably we would have liked more participants in this test session but due to our time limit and other colleagues unavailability we had to settle with four testers. The testing group consisted of two female master thesis students (aged 25), one male developer from the VR and Web department (aged 30), and the chief of the department (aged 45). The testing was performed at the Jayway facilities. All participants had individual testing sessions which lasted about 20 minutes. In the beginning of the test session they were introduced to IKEA's problem and our approach on how to solve it. They got a short introduction to the application and how IKEA categorize their products. They were then asked to navigate the application freely and visualize and read data they found interesting. During this they were asked to

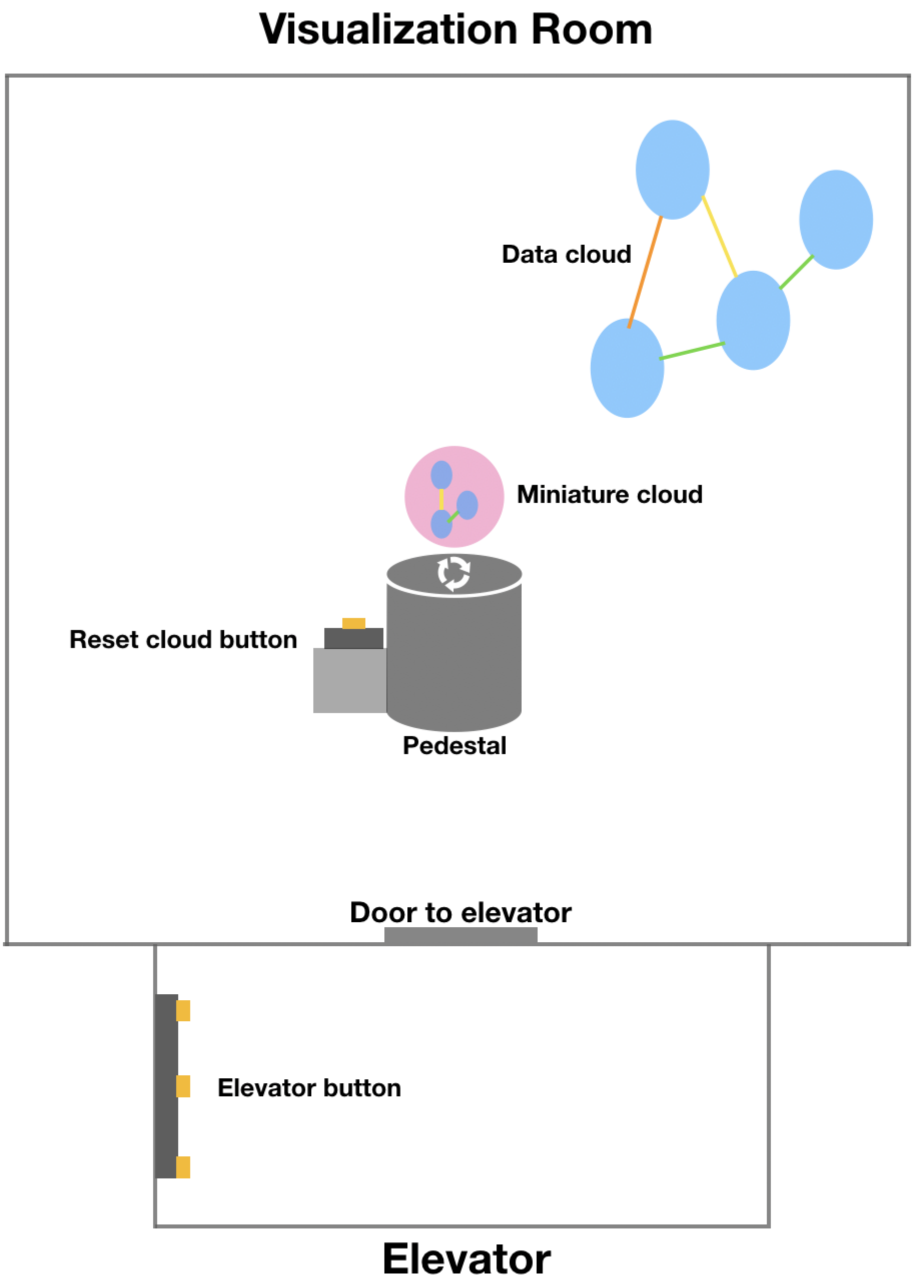

Figure 4.16: The *Visualization Room*. Each product is represented by a sphere, and the lines between the spheres show the affinity. The data cloud can be scaled and rotated by grabbing the miniature cloud at the pedestal.

describe their thought process and give spontaneous feedback.

It was the first time each of the testing persons tried our application. They all immediately understood that the data cloud was interactive. Furthermore, there where less confusion about which floor they where on and how the elevator worked than in the previous iterations with other participants.

However, we received a lot of valuable feedback on other aspects of our application.

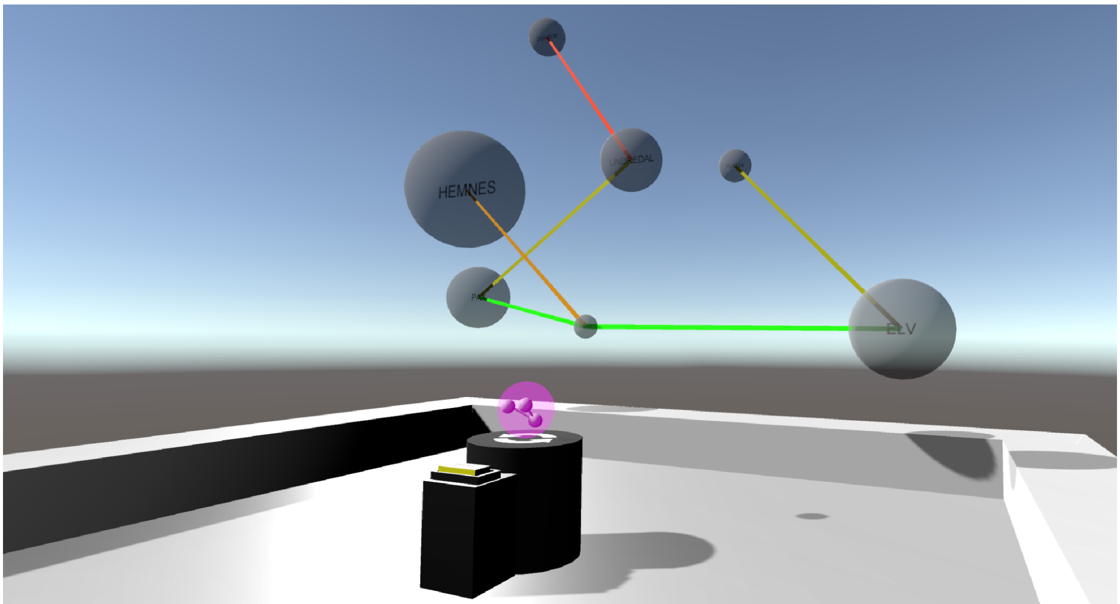

Figure 4.17: The interactive miniature cloud is placed on the pedestal in the middle of the room. The purple sphere can be grabbed, rotated, and scaled - and each action is mapped to the actual cloud instantly.

For starters, some believed the wrong button on the controller was used for grabbing objects. Furthermore, some thought a better pedestal object could have been chosen to increase the *signifier* that one should drop a product there.

As for the *Visualization Room*, it was suggested that texts inside the spheres always should be rotated towards the user, since it currently was a bit tricky to read from some angles.

Some suggestions were raised about how the product selection process could be improved. The doors were not used in a natural way, since one had to teleport in front of it and then walk right into it. Instead, it was suggested that the elevator door should open automatically when someone walked nearby. As for the regular door to the next floor, one participant suggested that an open door with visual stairs would increase the understanding of changing to another level, and that this would remove the rather unnatural way of walking straight into a door.

4.4.5 Iteration 5

Defining the problem

Using feedback received in the previous iteration a list of fixes to prioritize was made. We chose to ignore the unnatural doors for now, and instead focus on the *Visualization Room*.

The goal was to fix the following:

- Change grab button
- Implement a menu about data for a sphere which is shown once a sphere is selected

- Auto-rotating sphere text and menus towards the user

Furthermore, we wanted to add some new functionality to the *Visualization Room*. An idea which was brought up in the lo-fi prototype already (see section 4.3.4) was to have a sphere menu with information about each product. We decided that now was the time to implement it.

Lastly, we wanted to add more spheres to our cloud, ideally from different product families, to make it look more like it would in a real scenario so that the end users could see the true value of analyzing relations with our application. A natural way of doing this would have been selecting multiple products in the *Selection Rooms* and visualizing all spheres from different products in a big joint cloud. However, the feature of selecting multiple products had not been implemented, nor was there a concrete idea of how this actually would be done.

Collecting information

Research was made on how to group items that belong together in the cloud, and thus be able to visualize products from different categories. *Grouping by Category* in section 2.4.3 together with *Similarity* and *Proximity* described in section 2.3 were used to implement this.

Furthermore, research about how to evaluate our application was made. The SUS survey, described in *SUS Survey*, section 2.1.2, was chosen since it suited our purpose perfectly. It was a good compliment to our previous interview and furthermore, it was a good way to remove any "Social Desirability Bias". Additionally, SUS is best suited at the end of study and with few participants.

Brainstorming

Whether or not to change the grab button - The original thought was to use the controller side button for this (the grip button), since pressing it better resembles a hand squeeze which is what one would do to grab an object in real life. However, many participants had wished for the trigger button instead, since they believed the side button was difficult to find on the controller. Thus, we decided to change to the trigger to get a better user experience.

Information menu for each sphere - The most obvious feature, which had also been expressed by the end users, was the sales number. This since the end users sometimes want the exact numbers and not just the approximation given by the scale of the sphere. We figured more information, and even buttons, could be added here later on. However, since this was outside of our scope, and we did not know exactly what the end users wanted, we decided to show the sales numbers only.

How to make the Visualization Room as realistic as possible - We discussed on how to to add spheres from different products to the data cloud without completely changing our current product selection mechanism. Since this was going to be our last iteration in the implementation phase and we wanted our end users to get the full experience of what the *Visualization Room* could look like we decided to add spheres to the cloud via a script to

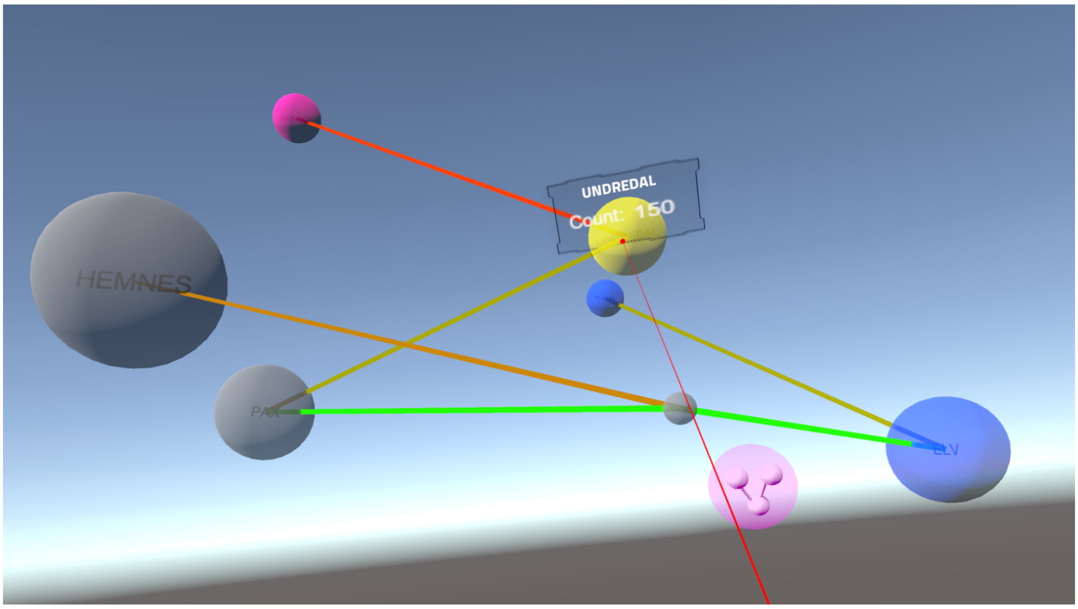

Figure 4.18: The data cloud is now color coded - the grey spheres belong to one product family, the blue ones to another and the purple to a third. The yellow sphere is highlighted by the user, which displays the menu for that sphere.

make it look like multiple products had been selected. We also discussed what would be the best way of visualizing which spheres belonged to which product family.

Developing

Grab button - With the motivation described in the brainstorming session the button for grabbing objects was changed to the hand controller's trigger button.

Sphere Menu - When implementing where and when the sphere menu was to be displayed, we chose to build on the idea from the lo-fi prototype, since the end users were positive towards it at the time. The idea was that once a user selects a sphere with the laser pointer, a menu should appear next to it with relevant data. As soon as the user deselects it, the menu will disappear. Since the end users recently expressed how they wanted the spheres rotated towards the user at all times, we figured the same should apply for the menus too. The menu can be seen in figure 4.18.

The auto rotation of the menus and texts to face the user was implemented by making them turn towards an empty object, which in turn was told to follow the user.

As explained in the Brainstorming section only the sales number is displayed in the menu as we did not know what type of data our end user wanted to be shown. However, ideally more information about the product and perhaps even its relations to other products would have been shown.

Adding more spheres to the cloud and visually categorize them - Spheres from different product families were added to the data cloud using a scrips since there were no way of doing this manually in the current version of the application. Next step was to visually categorise which spheres belonged to which product family.

Grouping by Category in section 2.4.3 clearly states that color is an excellent way of showing that items belong together. Thus, we decided that all spheres belonging together should be of the same color, since it emphasises *similarity*. Furthermore, grouping them closer together is also a way of showing that they are similar, i.e., *proximity*. However, this could only be done to a certain extent, as the affinity between all spheres in the cloud needs to be taken into account when placing spheres in the cloud.

The color coding of the spheres was done in the same script that creates the cloud. All product models selected in the *Selection Rooms* are passed to the script, which renders the cloud. In this case some additional products were chosen automatically as the feature of selecting multiple product had not been implemented. Spheres which are of the first product family will be blue, those of the second product family will be purple, and so on, as shown in figure 4.18.

Evaluating

As this was our final iteration, we decided to perform a more thorough evaluation. There were three major things we wanted to evaluate:

- Are the goals which were set up at the beginning of the implementation phase fulfilled?
- Does our application solve the end users initial frustrations which are listed in the concept phase?
- Is our application user friendly?

Firstly, we wanted to evaluate whether the goals set up at the beginning of the implementation phase were fulfilled. These goals were:

- A user should be able to select which products to visualize in an intuitive way.
- A user should be able to tell which products are closely connected, or have high affinity to one another, without looking at clear numbers.
- A user should be able to tell if a product is sold more or less than another product without looking at clear numbers.
- A user should be able to move in the virtual space to explore the data and look at it from different angles.

To test this, we conducted a testing session at IKEA's facilities with *participant 1*, *2* and *3*. The testing session contained four major use cases, each with some minor tasks connected to them. The participants were given the same use cases, which they performed one at the time, as the others were asked to leave the room in the meantime. The use cases are listed below:

1. Case 1 - Sales Amount

1.1. Find out which product within Bedroom Furniture is sold the most? How many times has it been sold?
1.2. Find out which product within Bedroom Furniture is sold the least? How many times has it been sold?
1.3. Which product of Pax and Platsa have been sold the least?

2. Case 2 - The Affinity Between Products

2.1. Find which product within System Wardrobes has the highest affinity to Pax?
2.2. Which two products within System Wardrobes have the least affinity between each other?

3. Case 3 - Multiple Product Families

3.1. How many different product families are visualized?
3.2. Which product family has the least affinity within itself?
3.3. Which product family has the most product relations to other families?

4. Case 4 - The Manipulation Of The Cloud

4.1. Make the data cloud bigger.
4.2. Rotate the data cloud.
4.3. Reset the data cloud.
4.4. Make the data cloud smaller.

Here follows the result of the use cases: Use case 1 was pretty straight forward, and all participants understood the three questions and answered them correctly without hesitation.

As for use case 2, which handles questions regarding affinity, it was clear for all participants what the distance- and colors- of the lines between the spheres meant. They were encouraged to describe their thought process during the testing, and three interesting things were noted here. Firstly, one participant speculated if the line thickness has a meaning. Secondly, another participant mentioned the importance of combining color coding with distance, since the distance sometimes was difficult to estimate in 3D. Furthermore, one participant raised the concern that it was not clear what affinity value was required between two spheres in order to draw a line between them. This came up when performing the task of identifying which pair of products had the lowest affinity, which technically was all those product pairs without an affinity line (or a value close to 0). This will be discussed further in the *Future Work*, Chapter 5. This use case involved using the elevator, and many participants were pleased with the addition of elevator sounds.

The third use case focused on visualizing multiple product families in the same cloud. All three sub tasks were completed by all participants, although some feedback were given along the way. For starters, all found the color coding very clear, and understood which products belonged to the same family. However, two of the participants expressed the wish that it somewhere should be stated which color represented which product family. Another feature that was required was the possibility to add and remove products, or whole product families, from the cloud without having to go back to the *Selection Rooms*. This will be discussed in more detail in chapter 5.

The final use case involved manipulating the cloud, i.e, scaling or rotating it. This had been tested before, but the *affordance* of the miniature cloud was changed, and the cloud was bigger, so we wanted to test it again. Moving and rotating the cloud was easy for all participants, but some found the scaling a bit tricky. However, after trying it out for a few minutes, everyone got a hold of that too. Lastly, the "reset cloud" button was used effortlessly.

Other new functionality which had been added, but not directly included in a specific use case, seemed to be intuitive. For example, no one had any issues reading the sphere menus or sphere names, as they were turned toward the user at all times.

Figure 4.19: A discussion among the end users after the test session.

After the testing session we talked about the application while it was running, so the end users could demonstrate what they liked as well as what they thought could be improved, as seen in figure 4.19.

The results of the use cases proved that the goals set up in the start of the implementation phase were met. The participants managed to complete the use cases described above which they were given, which were formulated to test the all goals which had been set up.

The question of whether the application was a solution to the frustrations listed in the concept phase (section 4.2) still stood. To answer this question an interview was held after the testing session. The questions in the interview were formulated from the frustrations the end users described in the concept phase. Below are the questions and the summarized answers listed.

1. Would you say it's more or less time consuming to look up multiple relations to a specific product using this application?

This application could actually help a lot. It's much easier to see many relationships, and especially to discover new. However, for looking up one specific relation between A and B - I think our current system is faster.

2. Would you say it's more or less difficult to keep track of relations between products from multiple product families using this application?

   Our current system definitely requires more effort, whereas the VR application does most of the work for you.

3. Would you say this application makes it easier to get an overview of all product relations to a specific product?

   Absolutely. This is the biggest advantage of the VR application.

4. Do you think this application could be useful in your daily tasks at IKEA?

   Yes, it would be a good compliment for us in performing our tasks, especially in discovering customer behavior and browsing for new product relations which are unaware of. I think the biggest use for the application is to show data for coworkers who are not used to analyzing numbers. For example to make people at other departments, like marketing, understand product range decisions, since it's very accessible anyone regardless of knowledge of numbers or the product range. It would be a really useful tool in presenting data to others.

Thus, many of the frustrations listed by the end user in the concept phase (see section 4.2) had been solved with our application. However, there is still functionality that has not been implemented in this project which would be desired by the end users in a real application. These will be discussed in chapter 5.

As a final step, we performed an evaluation to see how our application performed usability wise. We sent out an online SUS survey to our three participants who participated in the user tests. The survey was sent immediately after the testing session, and they were encouraged to answer right away, when the experience still was fresh in their minds. The survey was done online, and therefore also anonymously, to avoid biased answering. We collected the results the same afternoon, and calculated the SUS score according to their algorithm.

The survey can be seen in Appendix A, and these were the results: The average SUS score was 71.25, which is neither particularly good nor bad, according to *SUS Survey*, section 2.1.2. However, one should note that there were two questions in particular that brought our score down. These were *I felt very confident using this application* and *I found the website very cumbersome/awkward to use*, and it is quite likely that the end users inexperience in Virtual Reality made them insecure and a bit awkward while trying out the application. It is possible that the two questions would have been answered differently once they feel more comfortable using Virtual Reality.

This concludes the implementation phase, the final phase in our work process. In the next chapter, the discussion, the results will be analyzed.

Chapter 5
Discussion

In this chapter we will analyze the test results. Additionally, the concept of using VR to visualize data will be discussed. Lastly, future work that could be added to the application will be brought up.

5.1 The results of the user centered design process

This section will review the user centered design process and whether we achieved the goals set up at the start of the project thanks to this process.

We believe that the overall result of the project is positive, and that the user centered design process helped us in achieving this. Our iterative way of working have continuously improved the application, and according to the last testing session (section 4.4.5), the end users are pleased with the application.

However, we believe an even better result could have been achieved if a higher number of participants could have been part of the testing sessions, but unfortunately, IKEA could only provide us with three. We tried to compensate for this by inviting others to other testing sessions. Especially during the implementation phase when the interaction in the application was the focus, which could be tested just as well by someone else. However, many of these tests were informal and mainly focused on interaction, and not the application as a whole. We do wish we involved these people in the more structured testing session with use cases, as well as in the SUS survey described below. This would probably have given us a more trustworthy result.

The SUS survey resulted in an average score above 70, which means the user experience is quite good. Furthermore, our interviews and discussion in the same iteration shows that they believe most of their requirements at the start of the project are fulfilled. To summarize: they believe it is easier and less time consuming to study multiple relations, and the overview of all products and their relations is much clearer. They are pleased with the possibilities that comes with VR, and they believe this application could be used for many purposes at IKEA, such as discovering new relations within data and presenting data for others with less knowledge in the area. Unfortunately, we only involved our three test persons in this survey. As stated in the paragraph above, we wish we had involved more people here, regardless if they are our end users or not, and if they are inexperienced in VR. The reason for this is that three participants are too few to give a trustworthy general evaluation, as one a deviation among those three will affect the result enormously. To summarize, the result does of course give us a hint about the usability of the application, but more participants would give us a more precise score.

5.2 Virtual Reality as a data visualization tool

This section will discuss whether VR is a good tool for visualizing data, and more specifically data relations in big data sets.

The goal of this project was to take huge amounts of data, which is the outcome of machine learning algorithms applied on purchase patterns to get the product relations, and visualize this in such a way that a human can grasp it and get insight in these patterns.

By visualizing this, we hope to achieve the general goal of converting slow reasoning tasks into fast perception tasks. Furthermore, a we hope our visualization fulfills the two requirements:

- It helps with analyzing data and making decisions.
- It helps with communicating and presenting your research findings.

We believe that our VR application achieves this. Take these quotes from iteration 2 in the implementation phase (section 4.4.2) for example:

It's very intuitive to understand connections compared to Excel. Now you don't have to read the numbers, analyze them, go back in the application, find a new relation, read the numbers, compare to previous number and so on. Instead you can see all connection at the same time. I think we'll be able to understand things which we could not before.

It's much easier for people that are not used to reading numbers, or not used to our product hierarchy, to understand the relations.

It is clear the application provides good data visualization since it allows the user to intuitively interpret the data. The user don not have to do the same amount of active analyzing of data as when presented with data in a regular 2D application. Instead the data is, in many ways, already analyzed by the application and presented in a way that promotes decision making, which is the purpose of data visualization.

We believe we have achieved a good result by combining known research of general data visualization, such as gestalt principles (section 2.3) and visual encodings (section 2.4.3), with the advantages of 3D and more specifically VR (section 2.6 and section 2.5).

A downside of using VR for data visualization is that VR generally is not very including to multiple users at the same time. Often, like in our application, only one user can be in the virtual world at the same time. It is possible for people without the headset to see what the user sees in the virtual world on a separate screen and in that way multiple users could collaboratively analyze data. However, this is tricky since depth and immersiveness cannot be perceived in the same way without the VR headset, and many advantages of using VR is thereby lost. Also communication with someone in the virtual world from the outside is often problematic which makes this kind of collaborative analyzing even harder. However, we knew from the concept phase that there was no need for us to develop a multiplayer application that allowed collaborative analyzing as next to all analyzing by our end users were done individually. But if this type of application were to be developed for a company where this is not the case one would have to take the struggle of inclusiveness for multiple users in a VR application into account.

5.3 Future work

This section will list future work which could be added to a real version of the application. This includes improvements to the current application, as well as new features.

First and foremost, real data sets of the product relation data should obviously be used. This would preferably include real 3D models of each product in the product range, something which would improve the application, especially as a presentation tool.

Secondly, the unnatural doors which are discussed in section 4.4.5 should be reworked. One idea is to have the elevator open automatically once the user approaches it, which would increase *mapping* to a real elevator. The door to the next room should also be reworked, to increase the *affordance* of going down the hierarchy, by showing some stairs down the house, instead of just a door which one cannot see through. However, one could argue that the selection mechanism should be reworked entirely. The pros of keeping it this way, with the addition of the proposed changes above, is that the user can see 3D models and get a better knowledge of the whole product range, how the products are categorized and IKEA's hierarchy structure in general. Furthermore, from our tests and feedback we can conclude this way of selecting what to visualize is quite fun and intuitive but not very efficient. On the contrary, there are quite many steps one has to go through to select a product in some cases, especially if the feature to select multiple products at the same time were to be added, as requested by the end users. During the last brainstorming with the end users, it was suggested there could be a selection mechanism inside the cloud in some way. Here it would be possible to add new products/spheres with the click of a button. This would save the user a lot of time, since he/she never would have to leave the data cloud to go up or down the hierarchy searching for an item.

Data visualization can be categorized into two fields: presenting and exploring, which are described in section 2.4.2. We believe that the current selection mechanism is perfect for the former, whereas the latter category would need a reworked mechanism. There are some advantages with keeping everything in one scene. If the user could stay in the room while alternating the cloud, this would mean that he/she would see the cloud changing. Animating additions or removal of products, and their effect to other products, would probably increase the understanding of the relations even further. Additionally, adding the possibility to view the history of the relations, and see how these have changed over time, or even their predicted change in the future, would be extremely powerful for the end users.

The affinity line, drawn between two products, to show their relation to one another, is well understood by the end users. However, there are two minor details that should be considered for the future work. Firstly, it should be clear what level of affinity is required in order for a line to be drawn, since technically all products are probably connected (i.e., have been sold together at least once), but far from all connections are visualized. Secondly, the thickness of a line, which is the same for all relations, does currently not have any meaning. However, lines close to the user appear thicker, which might risk confusing the user.

The feature of selecting multiple products to visualize is definitely something that should be prioritized if one would continue developing the application. In fact, it has already been initiated. The color coding of the cloud has been done to show how it would look if one could select multiple products. The use cases shows that this is intuitive, but that there should be some explanation as to which color relates to which product family. According to the participants a simple sign in the *Visualization Room* would be sufficient. This, along with a reworked selection mechanism would be great steps towards multiple selections.

Another feature that should be added, especially if multiple selections is implemented, is the possibility to filter what data to be shown. The end users were concerned about the lack of filtration in section 4.4.2, where they said that too many spheres in the cloud might be difficult to comprehend. One filtration method that we believe would be great is the possibility to filter what affinity value that is required to draw a relation line between a product pair. This could for example be done with a lever next to the pedestal in the *Visualization Room*. This feature would solve the concern that was raised in section 4.4.5 that it is unclear what value of affinity is required in order for a relation line to be drawn.

Lastly, the sphere menu should contain more information. The end users have requested that data such as total revenue and gross margin should be added, and even functionality such as buttons to remove the product from the cloud, view data history or even emailing an excel document with data of the highlighted product to the user. These features are not part of our scope for this project, but should definitely be part of a finished application.

Chapter 6
Conclusion

Data visualization has always been a great tool to understand and present data. The purpose of data visualization is to convert slow reasoning tasks into fast perception tasks. A good visualization can be used for analysis and decision making. Furthermore, it is of great advantage when communicating and presenting your research findings.

However, the incredible amount of data volumes that IKEA possess in form of product relations required something new and innovative in data visualization. The volumes are simply too big for a human to comprehend all of it. Looking in multiple tables or 2D graphs results in too much information the human must keep track of, often because it require multiple searches to get all relevant data.

The purpose of this project was to investigate alternative ways of how data visualization can be used to visualize relations in big data. We chose to study whether a 3D presentation of the data in Virtual Reality could help in exploring the data and finding the hidden gems within the big data that simply would take to long to find and understand in regular tables or graphs.

A VR application was created in which the user can select from which products data is to be visualized. The data is presented in a data cloud where the user in an intuitive way can see sales number of products and the affinity between multiple products

Testing sessions with the end users at IKEA show that this application provides a great overview of big data, and tasks such as finding data relations in a big data set are greatly simplified. The vision of making it easier, faster and more comprehensive to understand data using our visualization application is fulfilled. They believe this application will be used to discover new things within the data that was not possible before, and that it can be used to present research findings in an intuitive way to inexperienced colleagues at other departments.

Bibliography

[1] A. Uzialko. How businesses are collecting data (and what they're doing with it), *Visited June 2019*. https://www.businessnewsdaily.com/10625-businesses-collecting-data.html.

[2] M. Arvola. *Interaktionsdesign och UX*. Studentlitteratur, 2014. ISBN 9789144097640.

[3] A. Cooper. The origin of personas, *Visited February 2019*. https://www.cooper.com/journal/2008/05/the_origin_of_personas.

[4] Chicago Architecture Center. Discoverdesign handbook, *Visited February 2019*. https://www.discoverdesign.org/handbook.

[5] S. Minhas. User experience design process, overview of stakeholders and activities involved in each stage, *Visited January 2019*. https://uxplanet.org/user-experience-design-process-d91df1a45916.

[6] S. Gibbons. Design thinking 101, *Visited January 2019*. https://www.nngroup.com/articles/design-thinking/.

[7] J. Nielsen. Why you only need to test with 5 users, *Visited February 2019*. https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/.

[8] T. Landauer J. Nielsen. A mathematical model of the finding of usability problems. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*, pages 206–213. ACM, 1993.

[9] B. Albery T. Tullis. *Measuring The User Experience*. Morgan Kaufmann Publishers, 2008. ISBN 978-0-12-373558-4.

[10] D. Norman. *The Design of Everyday Things*. Basic Books, 2013. ISBN 9780465050659.

[11] M. Wertheimer. Laws of organization in perceptual forms. *A source book of Gestalt psychology*, pages 71–88, 1938.

[12] A. Unwin C. Chen, W. Härdle. *Handbook of Data Visualization*. Springer, 2008. ISBN 978-3-540-33036-3.

[13] T. Wayson F. Vale. Charts and custom visualizations - beyond the map, *Visited February 2019*. http://proceedings.esri.com/library/userconf/devsummit18/papers/dev-int-101.pdf.

[14] O. Wilke. Fundamentals of data visualization, *Visited February 2019*. https://serialmentor.com/dataviz/index.html.

[15] J. Jerald. The vr book - human-centered design for virtual reality, *Visited June 2019*. http://www.morganclaypoolpublishers.com/catalog_Orig/samples/9781970001136_sample.pdf.

[16] Vive. What is the recommended space for the play area?, *Visited June 2019*. https://www.vive.com/eu/support/vive/category_howto/what-is-the-recommended-space-for-play-area.html.

[17] S. Katkoori R. Dubey E. Bozgeyikli, A. Raij. Point and teleport locomotion technique for virtual reality. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play*, 2016. URL http://doi.acm.org/10.1145/2967934.2968105.

[18] M. Mine. Virtual environment interaction techniques, *Visited April 2019*. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.294.7700&rep=rep1&type=pdf.

[19] E. Ragan K. Moghadam. Towards understanding scene transition techniques in immersive 360 movies and cinematic experiences, *Visited May 2019*. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.294.7700&rep=rep1&type=pdf.

[20] O. Odd M. Gustafsson. Virtual reality data visualization, *Visited May 2019*. http://hh.diva-portal.org/smash/get/diva2:1222037/FULLTEXT02.pdf.

[21] G. Ryrlén. Advantages of using virtual reality as a financial instrument, *Visited May* 2018. http://lup.lub.lu.se/luur/download?func=downloadFile&recordOId=8939254&fileOId=8939267.

[22] S. Davidoff A.Cioc A. Wang G. Longo J. Norris J.Zhang E.Lawler S. Yeh A. Mahabal M. Graham A. Drake C. Donalek, S. Djorgovski. Immersive and collaborative data visualization using virtual reality platforms. *2014 IEEE International Conference on Big Data (Big Data)*, pages 609–614, 2014.

[23] Unity. Unity game engine, *Visited April 2019*. https://unity.com/.

[24] SteamVR. Steamvr plugin, *Visited April 2019*. https://www.assetstore.unity3d.com/en/#!/content/32647.

[25] VRTK. Vrtk plugin, *Visited April 2019*. https://www.assetstore.unity3d.com/en/#!/content/64131.

[26] Git. Git, *Visited May 2019*. https://git-scm.com/.

[27] GitHub. Github, *Visited May 2019*. https://guides.github.com/activities/hello-world/#what.

[28] Adobe XD. Adobe xd, *Visited May 2019*. https://www.adobe.com/se/products/xd.html.

[29] HTC Vive. Vive vr system, *Visited April 2019*. https://www.vive.com/au/product/vive-pro/.

[30] Pimax. Pimax 5k plus, *Visited April 2019*. https://pimaxvr.com/products/5kplus?variant=19912548122683.

[31] HMD Odyssey. Hmd odyssey - windows mixed reality headset, *Visited April 2019*. https://www.samsung.com/us/computing/hmd/windows-mixed-reality/xe800zaa-hc1us-xe800zaa-hc1us/.

[32] MethodKit for Projects. Methodkit for projects, *Visited February 2019*. https://methodkit.com/shop/methodkit-for-projects/.

[33] Vertex Studio. Big furniture pack, *Visited May 2019*. https://assetstore.unity.com/packages/3d/props/furniture/big-furniture-pack-7717.

[34] Mohelm97. Simple home stuff, *Visited May 2019*. https://assetstore.unity.com/packages/3d/simple-home-stuff-69129.

[35] J. Miguel. Shooting gallery - enviroment pack, *Visited May 2019*. https://assetstore.unity.com/packages/3d/environments/shooting-gallery-enviroment-pack-105306.

[36] M. Ocklze. Hq modern desk, *Visited May 2019*. https://assetstore.unity.com/packages/3d/props/furniture/hq-modern-desk-66134.

[37] Music Technology Group. Freesound, *Visited May 2019*. https://freesound.org/.

Appendices

Appendix A

The System Usability Scale Survey

Participant ID: _____ Site: ____________________________ Date: ___/___/___

System Usability Scale

Instructions: For each of the following statements, mark <u>one</u> box that best describes your reactions to the website *today*.

		Strongly Disagree				Strongly Agree
1.	I think that I would like to use this website frequently.	☐	☐	☐	☐	☐
2.	I found this website unnecessarily complex.	☐	☐	☐	☐	☐
3.	I thought this website was easy to use.	☐	☐	☐	☐	☐
4.	I think that I would need assistance to be able to use this website.	☐	☐	☐	☐	☐
5.	I found the various functions in this website were well integrated.	☐	☐	☐	☐	☐
6.	I thought there was too much inconsistency in this website.	☐	☐	☐	☐	☐
7.	I would imagine that most people would learn to use this website very quickly.	☐	☐	☐	☐	☐
8.	I found this website very cumbersome/awkward to use.	☐	☐	☐	☐	☐
9.	I felt very confident using this website.	☐	☐	☐	☐	☐
10.	I needed to learn a lot of things before I could get going with this website.	☐	☐	☐	☐	☐

Please provide any comments about this website: