UTRECHT UNIVERSITY
DEPARTMENT OF INFORMATION AND COMPUTING SCIENCES

MASTER THESIS IN GAME AND MEDIA TECHNOLOGY

Visualising Lifelogging Data in Spatio-Temporal Virtual Reality Environments

Author
Joris VAN ABEELEN

Supervisors
Wolfgang HÜRST
Lynda HARDMAN

Student Number
ICA-6035566

Abstract

The goal of this research was to extend an existing spatial visualisation of lifelogging data in virtual reality with the visualisation of temporal information. For this, we created two novel spatio-temporal visualisations of lifelogging data in virtual reality. A qualitative (n = 12) study was performed to find the effect of the temporal aspects of the visualisations on the enjoyment, usability, effectiveness, and nauseousness of the user. The continuous visualisation we created outperformed the other ones on all evaluation points and was therefore considered to be the most promising of the three. Furthermore, we demonstrated that the temporal aspects of the visualisations helped de-cluster images that were spatially too close to each other on the map, making the data more approachable for users.

August 25, 2019

Contents

I. Introduction

Ever-increasing amounts of images and other data ask for new ways of browsing data sets. This is particularly true in the domain of lifelogging. Lifelogging is capturing data about one's life and can take the form of people wearing cameras to visually capture moments of their life. In this research, we will discuss the creation of three novel visualisations that improve upon a preexisting visualisation of lifelogging data in virtual reality [28]. This preexisting visualisation is a map-based interface of the data where images are placed on a map. Since this visualisation is in virtual reality, it benefits from the large and immersive field-of-view of the virtual reality headset. This is helpful to display the large amounts of lifelogging data in one accessible visualisation.

A spatial visualisation has been proven to be beneficial in certain contexts, however lifelogging data also has a temporal aspect to it which is underused in such spatial visualisations. Presenting a user with information about time can be a useful addition to a system that allows for browsing lifelogging data. This research adds a temporal aspect to the preexisting spatial visualisation to help users browse the data. This research has two main goals:

- Use temporal data to improve a map-based visualisation of lifelogging images to help users explore the data.
- Use temporal data to spread out clusters of lifelogging images in a map-based visualisation to make the data more accessible and help users find specific data.

The first goal aims on improving the preexisting visualisation by introducing time-related features. This can be of benefit to the user as seeing temporal information of the data allows for a new way of looking for specific data. The second goal is related to a problem that arises when a lot of photos have been taken at approximately the same location. The resulting visual clutter leads to problems with selecting the right data and we propose solutions to this problem by including temporal information to the preexisting visualisation of lifelogging images. There are infinite ways to add temporal information to the spatial representation of lifelogging images. We propose three different visualisations based on intuitive and educated reasoning. Those three visualisations will be evaluated in a comparative study which focuses on four parameters. These parameters are enjoyment, usability, effectiveness and nauseousness which we argue are essential in our context. Therefore, we address the following research question:

- What is the effect of different temporal visualisations based on a spatial data representation of lifelogging visualisation in virtual reality on the enjoyment, usability, effectiveness, and nauseousness of the visualisation?

The contributions of this research are:

- The creation of three novel spatio-temporal visualisations of lifelogging data in virtual reality. These visualisations extend upon previous research [28] by adding a temporal aspect to a spatial visualisation. This benefits the users of the system by enabling the them to browse the data based on the temporal features of the data which allows for new ways of browsing this data. We introduce a baseline interface which adds some temporal features to the visualisation of Ouwehand [28], a continuous visualisation which allows for browsing the data in a continuous and chronological stream, and a discrete visualisation which allows users to browse specific time frames of the data.
- The analysis of these novel visualisations by performing a qualitative study which focuses on the enjoyment, usability, effectiveness and nauseousness of the visualisations. Using lifelogging data from [25] we performed a study with twelve participants which compared the three visualisations. Our results show that the continuous interface outperforms the other two on all evaluation points.

II. Theoretical Background

In this section, we will provide a theoretical background of image browsing for lifelogging data in a virtual reality environment. We will start with an introduction into lifelogging by giving a definition of lifelogging and by looking at motivations for people to lifelog. Subsequently, we will look at image browsing and the visualisations that help users browse images. We take a close look at map based image browsing and image browsing that uses the temporal aspect of the metadata of the images. Finally, we examine browsing lifelogging images. First, in a normal setting and finally in a virtual reality environment. We end with some conclusions about the current state-of-the-art of lifelogging image browsing in virtual reality (VR).

A. *Lifelogging Data*

Lifelogging is the recording of data about one's life. There can be different types of data that can be recorded and different reasons to log all this data. Before we look into that, we give a clear definition of what lifelogging is. A commonly used ([1, 9, 13, 14, 18]) definition is the definition by Dodge and Kitchin [5]: "A life-log is conceived as a form of pervasive computing consisting of a unified digital record of the *totality* of an individual's experiences, captured multimodally through digital sensors and stored permanently as a personal multimedia archive." This definition describes multiple facets of lifelogging, which we are going to break down in the following paragraphs.

The first thing important term in the definition is "pervasive", which describes that lifelogging is always happening. If someone lifelogs, he/she (commonly) lifelogs constantly as to gather as much data as possible such that it captures every aspect of the life of the lifelogger.

The second part of the definition we look into is "... a unified digital record of the *totality* of an individual's experiences ..." The unified digital record describes that all data gathered while lifelogging should be present in one digital record. Additionally, this part of the definition again signifies the pervasiveness of lifelogging by italicising the word "totality". Furthermore, the definition tells us that what is recorded is not just data, but rather "an individual's experiences", meaning that the data is highly personal and is grounded in the context of one's life. Capturing an experience is not a trivial task, as the experiences of different people of a simple situation is different based on for example their point-of-view of the situation, their past experiences, their current mood, etc. Lifelogging data should try to overcome these boundaries and capture the experiences of lifeloggers.

"... captured multimodally through digital sensors ..." indicates that lifelogging data is multimodal, meaning it consists of multiple modalities (e.g. sound, images, location, time, current temperature, heart rate etc.). To measure these multiple modalities, various digital sensors are needed.

Finally, "... [lifelogging data is] stored permanently as a personal multimedia archive." This suggests that digital sensors need to be connected to some kind of archive to create a permanent record of the experiences represented by the captured data. Furthermore the archive is personal and therefore it indicates that lifelogging is most commonly used for personal uses rather than using the data for commercial purposes.

At first sight, one might wonder why people lifelog. Collecting data about yourself is in itself not interesting without any grasp of what you can do with this data. Currently, many of the lifelogging technologies are health-related [13]. Examples of health-related lifelogging technologies are the FitBit and the Samsung Gear Fit. These examples are bracelets worn around the arm that for example track the heart rate of the wearer, count the number of steps taken by the wearer, track the sleep pattern/quality of the wearer and if the device is equipped with GPS, the location of the wearer. With this information real-time feedback can be given to the wearers, informing them on their physical activity.

A different reason to lifelog is that lifelogging data can aid the memory of the user [13, 20, 32, 39], which we will discuss more in-depth in Section II-C. Other reasons for lifelogging include logging activities of employees for legal/historical reasons and market research [13].

In summary, lifelogging is constantly capturing all kinds of data about your own experiences, often for personal use. Since constantly capturing a multitude of types of data generates a lot of data, it is necessary to create tools and interfaces that help users browse these kinds of data. In this research, we will focus on one type of lifelogging data, namely photographs. Lifeloggers who capture photographs usually do this by attaching a camera to their body which takes a picture at a regular interval.

B. Image Browsing and Visualisations

Creating, storing and sharing images is cheap and easy and a lot of people have access to (at least) one camera. In 2014, on average 1.8 billion images were uploaded to the web per day Eveleth [8], leading to 657 billion images in 2014. With this extremely large number of images comes the need for good ways of browsing and visualising all these images. In these paragraphs, we will look into the current state of image visualisations. We will focus on images tagged with geographical and or temporal metadata. We do not look into other types of visualisations such a grid views as used in most common online image browsing websites as Google Images [17] and Pinterest[30]. We exclude these visualisations from our analysis as these are not relevant to our research since we focus on using time and place as ways of browsing pictures and feel like lifelogging data can benefit from more advanced visualisations.

Looking into both temporal and spatial visualisations is of interest as each has its advantages as well as its disadvantages Plant and Schaefer [31]. The biggest disadvantage for both temporal and spatial visualisations is that the photographs need to be tagged with temporal and or spatial data, which is not always the case. Different types of visualisations can also be beneficial to different types of use cases. We aim to create a system that fits the most common needs for browsing lifelogging photographs.

1) Geographical Visualisations of Images: In this section, we will discuss the visualisation of images based on the geographical location of where the images were taken. Using geographical information for tagging images is becoming increasingly popular as technology makes tagging images with a location effortless [35]. It is therefore interesting to find useful ways of using this additional information in visualising images.

As will become apparent in the following paragraphs, most of the time these visualisations include some kind of map. However, it has to be noted that map visualisations in itself do not necessarily refer to visualisations that use geographical locations. Reducing the complexity of image data to two dimensions allows for projecting data points on a 2D plane like for example in multi-dimensional scaling, principal component analysis or FastMap [31]. These examples do not use any real-world geographical locations for positioning the images on a map, but do show the images mapped out in a 2D space similar to most geographical visualisations of images. In this research, a map is always representing geographical data.

Torniai et al. [35] reviewed multiple commercial applications that have a geographical visualisation of images. One of these applications is the map view of Flickr [10], a popular photo-sharing website. Images are represented by dots on the map based on the location from the geographical metadata of the images. Nearby images are clustered and clicking on a cluster shows a set of thumbnails that represent the images in the clusters. This way of visualising images is similar to the way Zooomr (a discontinued website) visualised images [35]. Here, images were also placed on the map based on their location but instead of visualising this with a dot, the thumbnail of the images was placed on the map to directly give the user direct information about the image.

Torniai et al. [35] also describe another way of using the spatial information of the images for a browsing method, that unlike most methods, does not show a map to the user. Flickr Photo Compass shows eight images that are related to a selected image by their relative geographical location. In the middle of the visualisation is the selected image shown and on the eight cardinal and intercardinal directions (N, NE, E, SE, S, SW, W, NW) surrounding this image, the images closest to the middle image in said direction are shown. This way, the location of the photos is used for the visualisation, without showing the user a map with all the images. This might not be the most optimal way to browse a large dataset, but might help users find a specific picture near a physical location.

A common way to visualise the spatial data of the images is by placing a simple marker on a map that represents an image [10, 19, 29, 36]. A benefit of this visualisation is that it is simple and intuitive as each marker represents a single image. However, when a lot of images are present at the same place, the scene gets cluttered and interacting with the data can become challenging as markers can occlude each other. Toyama et al. [36] list five ways of visualising images on a map which deal in different ways with this problem. Of those five we methods, we already discussed two: the simple marker and the thumbnail. The third method is an isopleth map, which is a continuous representation of the density of the images per location, similar to a normal density map. The hue, intensity or saturation of the map is related to the density of the images present at each location. The benefit of this method is that the map does not get cluttered with images. However as images are separate entities, presenting them in a continuous way loses the discreteness of an image. Furthermore, Toyama et al. [36] do not describe what interaction with an isopleth map visualisation of images looks like, which looks problematic as selecting an image from a continuous visualisation is not trivial. The fourth visualisation of images on a map that is presented by Toyama et al. [36] is the border-dependent visualisation. In this visualisation, information about the geographical borders of the map is used to subdivide the map into regions. In each region, a marker is placed that scales in size with the number of images present in its region. This technique stimulates useful clustering of images as a single cluster represents a certain predetermined region. This method however also requires additional information about the map as these borders need to be known by the system. Furthermore, scaling and zooming on the map requires different levels of regions as regions might become too big/small to be useful at a certain zoom level. The last method presented by Toyama et al. [36] is the method the authors used in the rest of their paper and is called 'Media Dots'. This visualisation lays a predetermined grid over a map and each cell inside the grid can contain a dot depending on whether there is an image located at the location of the cell. The dot scales logarithmically in size with the number of images present in the cell. Note that some spatial information is lost in the Media Dots representation as all images in one cell are reduced to a single dot with a fixed location inside the cell. However, this representation performs well as it can rescaling the grid and additional precision of the location of an image can be obtained by zooming in on the location of the dot and scaling down the the size of the grid.

Another method for reducing the clutter on the map is by selecting interesting images rather than changing the visualisation type of all images. An example of this can be found in[4]. Here 35 million images from Flickr are analysed based on their spatial distribution and their visual, textual and temporal features. The researchers showed that their visual analysis selects a picture that was most representative of all pictures on a given location and the results showed that in cities this lead to photographs of the top landmarks of that city. What made this task extra hard was that not all pictures at a given location need to capture the landmark and can be completely irrelevant to the task at hand. Having an algorithm that is able to pick out the most interesting locations and images of a large dataset can be useful for automatically generating summaries, which can be used as a method for removing clutter on a map. Jaffe et al. [19] also summarise a large collection of photos. To create the summary, they used the following information from a photo: location, time, photographer, tags, quality (from an externally derived parameter) and relevance to the user. For example, a location is likely interesting if it has a lot of different photographers. An interesting location in its turn should be represented in the summary. They present the summary as thumbnails placed on a map, rather than point markers as they did in their non-summarised visualisation. From their initial evaluation, it showed that summarising the data set was useful and made a good representation of the city they used in their tests. In this research we will not look further into summarising photographs as it is not in the scope of our research, however, we still thought it would be relevant to mention these techniques as they are a valid and useful way of dealing with too much data on a map.

Kalnikaite et al. [20] looked into the effect of including geographical information in Lifelogging data (more about lifelogging data in Section II-A) on the memory of users. The results showed that geographical information supported inferential memory, which is a type of memory that is inferred by some external reference to that memory but was not remembered by the user. This is in contrast to visual cues like photographs which support actual recollection of the memory. This indicates that showing images on a map (in whatever representation) can stimulate the memory of a user to remember certain bits of information that they had forgotten, meaning a map visualisation can be useful and insightful for users to get an overview of their past activities.

2) Temporal Visualisations of Geographically Tagged Images: Image data can contain a timestamp, representing the time and date of creation of the image. This data can be used in the visualisation of a set of images. In most common photo view applications like Windows Explorer and the standard photo applications on Android and iOS, photos are often already organised based on their time stamp, be it either because of all images are sorted on the timestamp or because they are sorted on the generally chronologically determined name of the photos.

Similarly as with spatial data, temporal data can also be clustered. In Gemmell et al. [11, 12] a system called MyLifeBits is presented where 'trips' are made by separating the images based on the temporal distance between images (or their spatial distance). Each trip can be replayed and the images that are taken during each trip are shown chronologically as the replay progresses. Showing the users the images of a single trip in chronological order is a simple way of using time for visualising images. However, when viewing the replay animation, it is hard to keep an overview of the complete data set. MyLifeBits also includes a calendar on which events are marked, related to the data. Browsing this calendar gives insight into the discrete time frame of all pictures.

Another approach to visualising time on a map of images is by colour coding the images markers on the map based on their timestamps as is done in Openpaths (as cited in [33]). It is possible to colour code the images chronologically, meaning the colour of each pin is (slightly) different and the colour indicates the order in which the photos are taken. Additionally, it is also possible to give certain time windows (like 'evening', Tuesday', 'March' or '2016') a specific colour. In this case, multiple markers can have the same colour, but the colour has a real-life grounding, instead of showing the order in which the photographs are taken. Another way of visualising time on a map of images is by connecting markers with lines as in Google Latitude (as cited in [33]). By chronologically connecting markers with lines, the order in which the photographs are taken becomes apparent. A more sophisticated way of using a connector between markers on the map is shown in Eccles et al. [6] where lines are drawn in a 3D space-time cube and where additional annotations can be placed inside the 3D space to further elaborate on the connections and the information shown in the cube. Although the GeoTime visualisation presented in Eccles et al. [6] shows potential, creating the visualisation is not an automated process, making it undesirable to use for large amounts of data like lifelogging data.

Lucero et al. [26] and Ames and Manguy [2] both use a timeline in their visualisation of photographs. In these visualisations, multiple related photos are presented on a timeline. One way of creating a timeline is by grouping the photographs by location and both systems of [2, 26] can do this. By doing so, the timeline shows all photos taken at a single location in chronological order. The timeline gives a quick, organised and playful way of browsing a large data collection of photos [26].

A more advanced visualisation of a timeline can be found in Thudt et al. [33, 34]. Here a visualisation called 'Visits' is presented. This visualisation creates a timeline of maps which each shows a 'visit'. A visit is a part of the map where the user took multiple photographs in a limited time frame. If the number of photos in that certain time frame exceeds a threshold, the visit is added to the timeline. On the timeline, all visits are presented in chronological order. The result is a timeline with parts of a map on it, with each map showing multiple markers which represent photos. The size of each visit/map is related to the number of images on that part of the map and the time spent in that area, where bigger visits/maps are related to more images/more time spent in that area. With this visualisation, a location history is shown where repeated visits of the same location are separated and where the duration of a stay is an important factor. However, all images of actual travel are likely to be excluded from this visualisation as these images are often too far apart from each other on the map to exceed the threshold, which might be undesirable depending on the use case.

C. Browsing Lifelogging Images

Where the previous section discussed ways of browsing 'normal' images, in this section we are going to look at ways of browsing images that were specifically created for lifelogging images. Before we go into ways of visualising lifelogging images, we look at the differences between lifelogging images and normal images.

In Yang et al. [40], 5 characteristics of lifelogging data are discussed: *passive capture* (the user does not need to interfere with the lifelogging device for it to capture images), *continuity* (the lifelogging device captures images continuously), *rich contextual cues* (additional sensor data give extra cues), *large* (lots of data is gathered by the lifelogging device), and *more variety of usage scenarios* (lifelogging supports five different types of memory [32]). These characteristics are all unique to lifelogging images in comparison to more mundane images. These differences lead to other user needs for when users want to interact with the data. Because of the passive capture and continuity of the lifelogging images, a lot of photos that are taken are not that interesting to look at individually, but they might be interesting in the context of other images. The rich contextual cues enable different types of visualisation of as there is for example locational and time-related data to work with (as discussed in Sections II-B.1 and II-B.2). Furthermore having such large amounts of data asks for ways of quickly filtering out data from the data set to get to specific images and asks for new ways of exploring the data set. At last, because of the different types of usage scenarios, different kinds of interfaces can be useful for exploring/browsing lifelogging images. The five types of usage scenarios that Yang and Gurrin [39] refers to, are described by Sellen and Whittaker [32] and are: recollecting, reminiscing, retrieving, reflecting and remembering intentions. Each type has a different purpose and possibly requires another type of interaction with the data. For example, Kalnikaite et al. [20] show that visual cues help users recollect memories while geographical cues have a more inferential nature. It is therefore important to realise what the user needs are when determining what kinds of data are visible in the interface.

1) Visualising Lifelogging Images: Visualising lifelogging data is a challenge in both the literal [15] and figurative sense of the word. As discussed in the previous sections, a good visualisation keeps its users needs in mind and is depending on the availability of certain types of (meta)data. Furthermore, the visualisation should depend on the device the user is using to access the lifelogging data [40]. Additionally, since lifelogging is still a relatively new field, not much research is done into visualising lifelogging images. Given these statements, it comes as no surprise that there are no real standards for the visualisation of lifelogging images. Therefore, we look at specific implementations in this section and see if we can find what does and does not work.

In Kalnikaite et al. [20] two types of visualisations for lifelogging images are shown. First, a simple visualisation focused purely on the images is presented. This visualisation shows pictures ordered by time. The user selects a day and the interface opens a new window, showing each image in chronological order. The second visualisation is centred around a map with map markers on them. When the user hovers over the markers on the map, the picture related to that marker pops up. As discussed early on, their results showed that using spatial cues is beneficial for inferring certain memories whereas images show more details and are better suited for actual recall of memories. The first interface method, focused on the images, might not be that useful for lifelogging data where the user gathers a lot of images each day since browsing the images is done by going over each picture individually. Kalnikaite et al. [20] found that combining location with images was preferred over showing only images as it gave both a good overview of the data as well as a detailed look inside the data.

Gemmell et al. [11, 12] also combine location with lifelogging images and add temporal information to the interface as well. Showing time has the added benefit of creating a logical order in which one can browse the images. Time is visualised in MyLifeBits [11, 12] by adding a calendar where events are shown which are accompanied by pictures and by recreating trips as is discussed in Section II-B.2.

A problem that arises in most of the previously mentioned methods for visualising time and location is that the visualisations focus on transitions overtime/location. With lifelogging data, all moments when the users travel are highly visible on a map as they take up a lot of space since travelling images are hard to cluster at one location. However, the more interesting data is likely to be present at the places the user stayed for a prolonged time. The problem with this is that data at one location is often presented with one marker or a lot of clustered markers that are hard to access individually. Using different kinds of markers that show the density of photographs as discussed in Toyama et al. [36] (see II-B.1), might be a good way to solve this problem. The visualisation by Thudt et al. [33] is also an interesting way of focusing on the specific events rather than on travelling as is more the case in the calendar in MyLifeBits Gemmell et al. [11, 12].

What becomes clear when looking at interfaces for browsing lifelogging images, is the need for getting to the right information at the right time. An interface should help the user solve problems by showing a good overview of the data that is available. This can be done by using metadata of the images like their temporal or spatial aspects. From this overview, it should be possible to dive further into the data by looking at a smaller subset of images or a single image to get the desired information from the picture(s). When designing an interface, it is important to keep the need for these different levels of exploration in mind.

2) Lifelogging Images in Virtual Reality: Browsing images in virtual reality consists mostly of viewing panorama pictures, even though some of the features of virtual reality are suitable for other types of image browsing as well. Having a wide field-of-view (FOV) can be beneficial as it is possible to show multiple images on the screen without occlusion and without the need to minimise the size of the photographs. Furthermore, because of the immersive nature of VR, it is possible to create a rich environment where the user can actively engage with its data. In this section, we will discuss some ways of visualising (lifelogging) images and their metadata (like location) in virtual reality.

Yang et al. [38] looked into different ways of representing maps/globes in virtual reality. This is relevant for this study as we will use the spatial data of the lifelogging images to pinpoint the location of the images on a map/globe. The study looked at four different ways of visualising earth: an exocentric globe, a flat map, an egocentric globe and a curved map. Users were tasked to compare distances and areas and to estimate directions. The results showed that the exocentric globe was generally the most useful in the tasks of the user study. However, since these tasks differ quite a lot from browsing images, it is questionable if these results translate to our results.

Mengerink [27] looked into how the results of Yang et al. [38] translated into a setting that is more similar to our research. In this research different kind of globes and maps are tested to see what are effective visualisations for map-based image browsing. The results showed that a floor map, a park sign map and a panorama map are viable methods for displaying images on a map in VR.

In the similar vein to [38], Yang et al. [37] looked into visualising flow maps in virtual reality. By creating flow maps and placing them on or over 2D/3D maps or globes, they tested what combination worked best. The results showed that using the third dimension is good for reducing the visual clutter on the screen while using a 2D map with straight link lines was the fastest.

Research that combines virtual reality with lifelogging images is done by Gurrin et al. [15]. The virtual reality part of the research focused on different kinds of interactions with menus for filtering the data. A selection menu with an aim and click method was presented and additionally some slightly more sophisticated ways were presented by having a smaller menu that could be interacted with that was closer to the user. Since no actual experiments were done with this, little insight can be gained about which method works better. At last, they presented a memory wall in the paper, which shows images organised in chronological order which allowed to be scrolled and interacted with. This visualisation made good use of the wide FOV of VR.

Ouwehand [28] continued the work of [27] and created a system for browsing lifelogging images in virtual reality. This system was found to be of high entertainment value and showed that browsing lifelogging images in VR in a map-based visualisation can be fun. This system will also be the base for our research. We describe this system in more detail in Section III-A.

D. Conclusions

In the previous sections, we looked at ways of visualising lifelogging data. We found good research laying the groundwork for our research as well as points that could be improved upon. In this section we will summarise these findings, starting with general image browsing. Geographical visualisations of image collections are often in some form including a map or globe where the images are placed on in some form. The placement of the images on the map or globe represents the spatial aspect of the images. Using the spatial aspect of data by placing the images on a map gives users direct insight into the data, however, it also comes with some problems. For example, a lot of image thumbnails stacked close together might lead to occlusions. Toyama et al. [36] presents methods for dealing with this problem by changing the markers of the images on the maps. Other ways of reducing clutter on the map include grouping images into one cluster of images [7, 36]. However, having very large clusters of data leads to new problems related to visualising and browsing these large clusters of data.

Using temporal information can be of great help to give users more insight into the data. The most common way of using time in the visualisation of images is by ordering the images in chronological order. Another, more complex way of using time in visualising images can be found in MyLifeBits [11, 11], where 'trips' are created by replaying images in chronological order and by creating a calendar which displays events related to the images in the data set. Examples of combining time and space for visualisations purposes can be found in [33]. Here a system is described that colour codes pins on the map based on their time stamp and a system that connect pins on a map by a line when two images were taken successively. The paper furthermore present a more novel way of combing time and space in a system called 'visits' which highlights places where the user took multiple photographs in a limited time and space and creates a timeline which focuses on these places.

Visualising lifelogging images is in essence pretty similar to normal image visualisation. The differences in visualisations come from the differences in the data. Lifelogging images are continuously and passively captured, have rich cues, are generated alongside lots of other data and have a wide range of use cases [32]. For lifelogging image browsing, it is important to be able to use these characteristics to create visualisations that enable the user to get a good overview of the data as well as gain insight into specific parts of the data that are of interest to the user. Virtual reality can be of added benefit to creating such lifelogging image browsing visualisation as VR gives a large field-of-view and is highly immersive. This makes it easy to visualise more images at the same time as well as create an immersive, fun, and natural user experience.

In this research, we create novel lifelogging visualisations that extend a map-based spatial representation in virtual reality with the temporal metadata of lifelogging images. There are previous visualisations that combine both spatial and temporal data into a single visualisation, however, little research is done on the how users experience such visualisations. Furthermore the possibilities for combining spatial and temporal data in a single visualisation has not been fully explored. In this research we create three novel visualisations and analyse the experience of users of these visualisations. The visualisations in this research also use temporal data to decluster

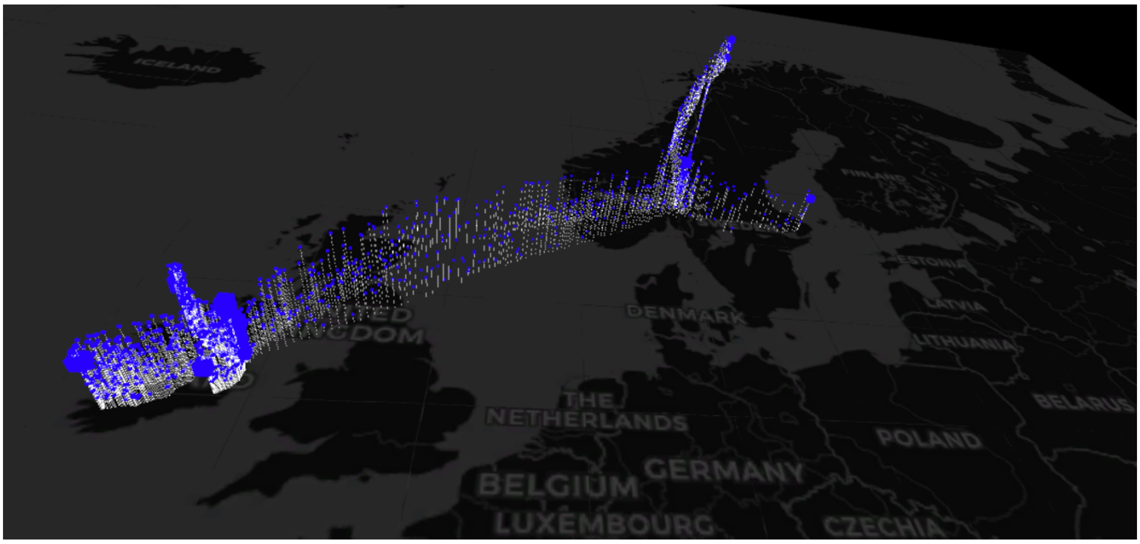

Fig. 1: The spatial visualisation as made by Ouwehand [28] which will be used as the basis of our visualisations. Taken from [28, Figure 7]

clusters of images on a map. Although something similar is done by [33], our visualisations do not select specific parts of the data that are (automatically) marked as interesting but rather show all data, giving the user a complete and unclustered overview of all the data. This is essential in our context as we create visualisations that focus on using time to present users with an overview of the data and to decluster places where a lot of data is clustered. Having places where the data is not clustered, allows us to compare these types of places.

III. Implementation

In this section, we discuss the implementation of the three different interfaces that will be used in this research. First, we start with the spatial visualisation as created by Ouwehand [28] (which is a continuation from Mengerink [27]) as our visualisations will be an extension this. We will not discuss all features of the system from the previous visualisation as Mengerink [27], Ouwehand [28] already provide these descriptions. However, we will highlight the parts that are the most relevant to our research. We continue with discussing the baseline visualisation which, as the name suggests, will be used as a baseline in the experiments. This baseline visualisation is also at the basis for the other two visualisations and most visuals/features that are present in the baseline will also exist in the other two. Furthermore we will discuss our extensions to the baseline visualisation called the continuous visualisation and the discrete visualisation. The application used in this research is made in Unity (Personal Edition, version 2017.3.1f1) and uses C# scripts for the logic of the system. Furthermore, we use the SteamVR plugin to work with the virtual reality headset, which in this case is the HTC Vive. For a more detailed description of the implementation of the back-end functionality of the system like how the system deals with all the data, we refer to Ouwehand [28].

A. *Basic Spatial Visualisation*

We start with discussing the works of Ouwehand [28], as his spatial visualisation will the basis for our visualisations. This interface is shown in figure 1. The main feature of this spatial visualisation is a map with pins placed on it. The map is a tilemap, which means that the image of the map is not a single image, but rather consist of a multitude of smaller images that in conjunction create the larger map. The benefit of using a tilemap over a map with a single image is that a tilemap allows for multiple zoom levels. When the user is far away from the map, the images do not need to show a great level of detail. Seeing a highly detailed map from far away is not useful as the geographical names on the map will be too small to read. Furthermore, having names of individual cities on a map where you can see the whole world is not relevant, as the user is likely to be more interested in the names of the countries at that zoom level. With this tilemap, when users fly towards the map (more on the navigation within the system in the next paragraph), the zoom/detail level of each map tile is calculated by taking the distance of the user to that tile. If the distance is below a certain threshold, the next zoom level of the map is loaded and shown instead of the old zoom level. The opposite process to this is happening when the user moves away from the map. For this research, we used five different zoom levels as having more zoom levels creates problems with computational resources and storage.

Moving around in the visualisation is done using the trackpad of the HTC Vive controller. Pressing on the left or right side of the trackpad lets the user rotate around, enabling users to look around without actually turning their heads. Note, that using the functionality is optional, since turning your head has the same result and might make participants feel less nauseated since their perceived motion is congruent with their actual motion. Translational movement within the system is achieved by clicking the upper and lower part of the trackpad. When the user presses the upper part of the trackpad, the system flies forward in the direction of where the user is looking at. This makes it possible for a user to fly anywhere in 3D space. When pressing the lower part of the trackpad, the user flies away from wherever they are looking at. Mengerink [27] looked at another method of navigation, namely teleportation. In this method, the user aims at an object and teleport directly towards the location of that object by pressing a button. The results of their research show that the flying movement as used in our research and the teleportation method are both viable ways of navigation for our type of map. We choose to go for the flying navigation, as this navigation enables users to navigate towards places in the 3D space where there are no objects to navigate towards. This is particularly relevant for the continuous and discrete visualisations (as described in Sections III-C and III-D) as in these visualisations it can be desirable to move up or down without teleporting directly towards an object.

As shown in figure 1 there are pins scattered around the map. Each pin on the map represents at least one photograph from the data set [25]. The size of the pins scales with the number of images that are represented by the pin. The pins are placed on the map based on the coordinates in the metadata of the pictures. However, placing a pin for all 56450 GPS-tagged images in the LSC data set would lead to an abundance of pins scattered around the map. Therefore, the coordinates of all images are rounded down to two digits, meaning there is an accuracy with a margin of error of 1.1 kilometre. Images that have the same coordinates after rounding down are clustered together. With this method, the number of pins is reduced to 3690.

The pins are placeholders for the images they represent and since the user is likely not interested in these placeholders but rather the actual images, Ouwehand [28] added in a reveal mechanic where when the user got within a certain radius of the pin, the pin would change into a miniature version of the image. This removes the colour of the pins and gives the user direct insight into the data they are working with. The reveal mechanic is shown in 2.

A user can open a pin by aiming with the controller at a pin and pressing the trigger button on the backside of the HTC Vive controller. This will load all images from that pin and present them in an image wall, which is a grid of images with some additional information like the timestamp of the image and their coordinates as shown in 3. In this wall, the data of each day in that pin is represented by a row and each column is a separate photo. This way a user can scroll through the data by day and in each day, the data is ordered chronological.

B. The Baseline Visualisation

The baseline visualisation is similar the visualisation of Ouwehand [28] with some extra features added to it. An overview of the baseline visualisation is shown in figure 5. We reduced the number of buffer tiles on the sides of the map where no data was in comparison to Ouwehand [28] as to give more attention to the relevant parts of the map and to bring aspects of the interface

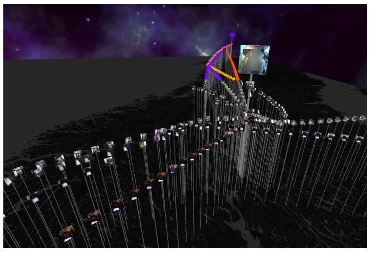

Fig. 2: The reveal mechanic: the representing image of a pin will be shown when a user gets close enough to the pin

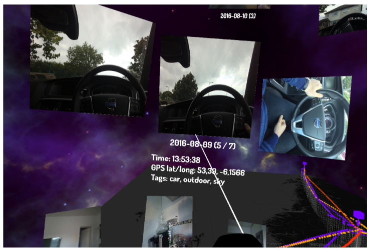

Fig. 3: The image wall, where the images of a single pin are shown. Each row represents a single day

Fig. 4: The date time label which shows the date and time of the image where the controller is aimed at

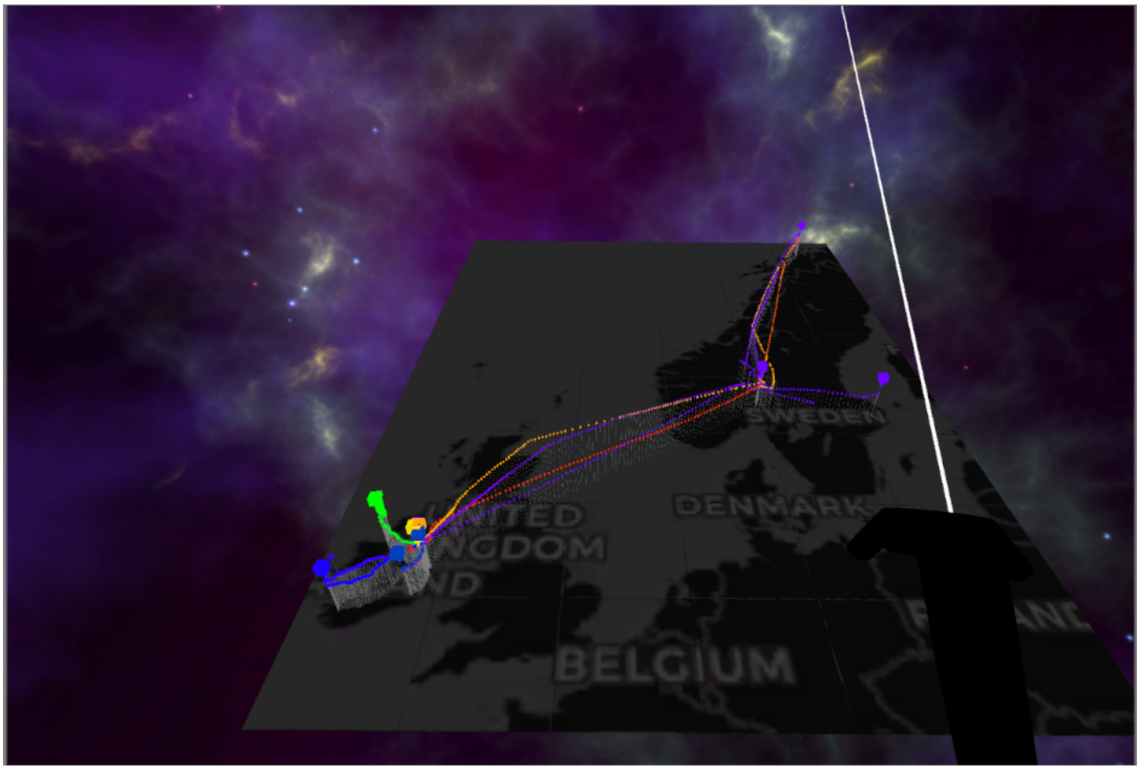

Fig. 5: Overview of the baseline visualisation with the controller

of the continuous and discrete visualisations closer to the user (as described in Sections III-C and III-D).

In the interface of Ouwehand [28], pins were randomly placed within a certain height range. This was done such that close pins overlap as little as possible, however, since the height in between pins differed, it also suggested that the height was related to data in some way, which was not the case. In our implementation, we related the height of each pin to the data by increasing the height of a pin with the number of images clustered in the pin. This lead to a more levelled field of pins with some pins being above others, giving extra weight to these pins. This, however, reinforces the problem of having some close pins overlap. The continuous and discrete visualisations (sections III-C and III-D) present different solutions to this problem.

In this research, we also introduced colours to the pins, similar to as in Openpaths (as cited in [33]). In the old interface [28], all pins were a single colour (blue). In our visualisations, the colour of a pin is related to the timestamp of the image of the pin. All pins were sorted on their timestamp in a time sorted pin list. We then created a list of colours where each entry was a gradual change of colour in relation to the previous. Going through the time sorted pin list in order, we assigned the first colour of the gradient list to the first day present in the data. Each pin related to that day was assigned the same colour. When a new date was encountered in the time sorted pin list, the next colour of the gradient was selected. If there were more days then colours in the gradient list, the list was reused for the remaining pins. During the development of this feature, we also looked at other options of colouring the pins to indicate time. Another option was to use a gradient that started at one colour at the first pin in the time sorted pin list and ended at another colour at the last pin in the list. All the other pins were a gradual change of colour from the first to the last colour depending on their chronological order in the time sorted pin list. This resulted in very small gradient changes in between pins. This seemed useful when two completely different dates were close to each other, but it was hard to see changes in the colours between dates close to each other. Furthermore, the actual temporal distance between two consecutive pins in the time sorted pin list had no effect on the colour, meaning that pins that were separated by minutes and pins that were separated by days would have the same gradual difference if they were consecutive pins in the time sorted pin list. Another option that was considered for colouring the pins, was giving each month or each hour a different colour instead of each day as is the case in the final implementation. Although colouring each month was useful for getting an overview of the data, as soon as the user would start to look a bit deeper into the data, there was no information to be gained from the colour of the pin. Using a different colour for each hour lead to a chaotic colouring of the pins that was not useful.

An addition that was made to more quickly get information from the photographs without the need to open a pin was the addition of a date label near the controller. This label showed the time stamp of the first picture of a pin when the user aimed their controller at a pin. This label was rotated such that it was always aimed at the head of the user and remained in place for a short time after the user stopped aiming at the pin. This was done such that it would be possible to use this feature on small pins and pins far away that were a bit harder to consistently aim at without having the label appear and disappear in quick succession.

Another improvement that could help visualise the temporal nature of the data that was considered in this research was connecting pins chronologically to each other with a line similar to Google Latitude (as cited in [33]). The benefit of this is that it is easy to see the order of the pins as it one can follow the line from one pin to the other. We decided against using this additional feature in our final experiments as the added benefit was marginal because of the inclusion of other temporal features, while it did clutter up the data quite a bit which was a problem we were aiming to solve with our visualisations.

C. The Continuous Visualisation

In the baseline visualisation, we added in some ways of displaying the temporal aspects of the data; in the continuous visualisation, we take this a step further by making the visualisation interactable based on the temporal aspects of the data by adding in scroll-able timelines which create a continuous stream of chronological ordered pins as visible in 6. On each of the four corners of the map visualisation, a timeline is placed. If the user clicks the trigger while aiming at one of the timelines, they can drag the time line up or down. The pins will move up and down simultaneously with the timelines, depending on their position in the time sorted pin list. At the start, the timeline is raised a little bit as to show the functionality of the scroller. Most pins will be at map level at the start, however the first couple of pins are raised. When the user scrolls using the timeline, more pins will rise from the map. The distance between each risen pin is fixed. The pins that are still at the map are not influenced by the scrolling until the user scrolls the timeline up far enough such that these pins should rise too. The order in which the pins rise is based on the time sorted pin list as described in Section III-B. This means that the first pin will always be the highest pin and the last pin will always be the lowest pin. The fully risen continuous visualisation is shown in figure 7.

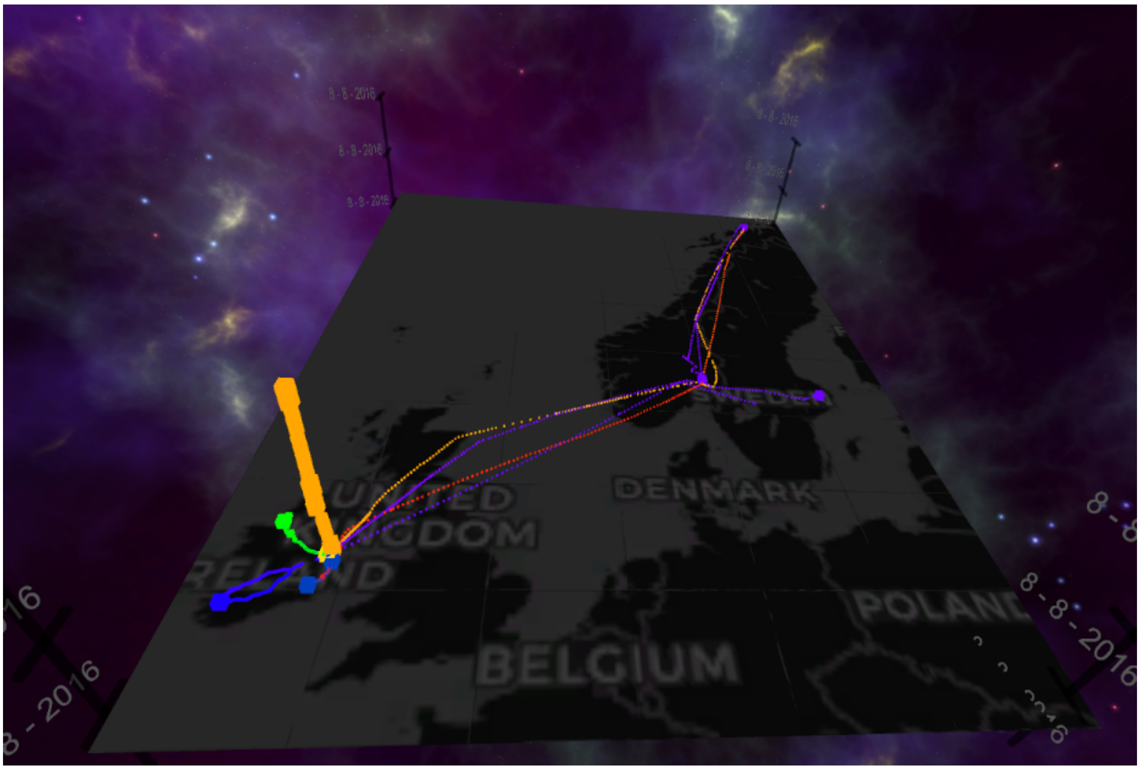

Fig. 6: The start position of the continuous visualisation with the time lines on all four corners. Notice that some pins in Dublin are risen from the map.

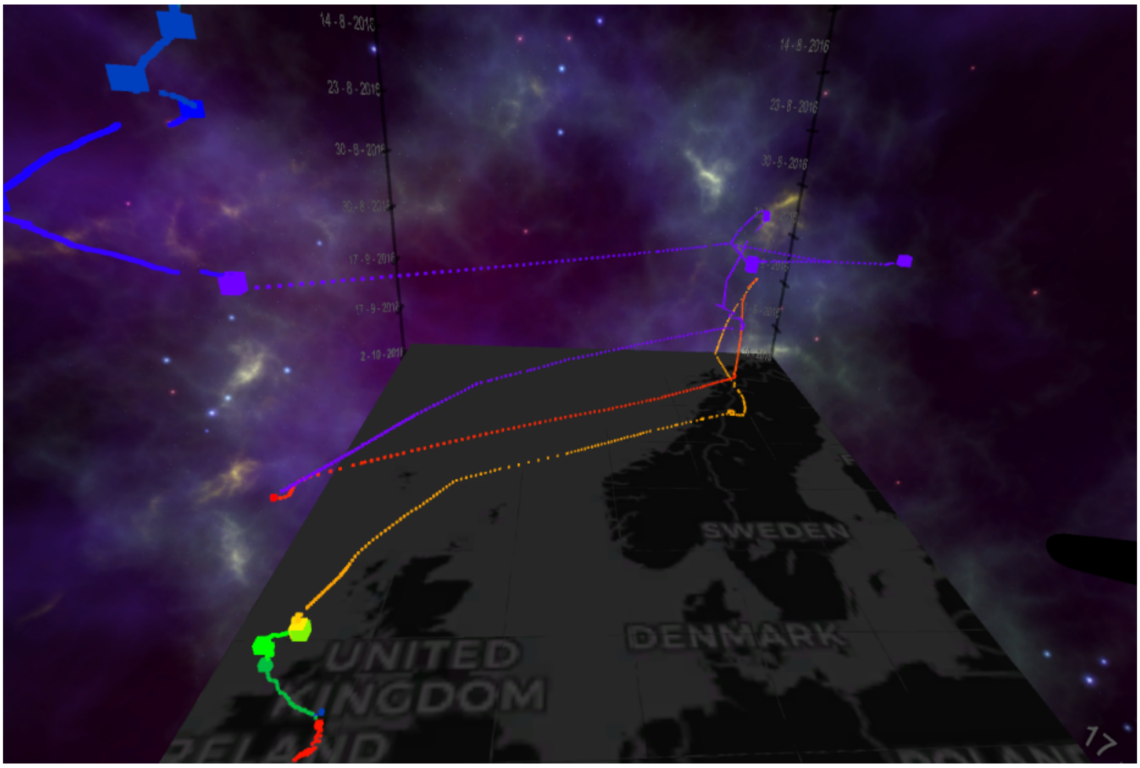

Fig. 7: The continuous visualisation fully expanded. All pins are now risen from the map in chronological order.

It has to be noted that the pin lines that connect the pin blocks to the map that were present in the baseline visualisation are not present in this visualisation any more. Instead, when pin blocks are at their lowest level, they are partially lowered into the map such that it is clear where they are located at on the map. The reason for excluding these pin lines is that having these pin lines present while the pins are raised from the map would result in a lot of visual clutter in the scene. So in this visualisation, some precision of the location of the pins is traded for the temporal data gained from this visualisation.

Each timeline on the corners of the map consists of multiple elements to make up the whole timeline. Each element consists of two parts, the actual timeline object itself and the accompanying date labels. The timeline object itself can be used for scrolling the timeline up and down and the date label indicates what the date is of the pins at the height of that timeline element. When the user scrolls the timeline up, it is possible that the bottom of the currently present timeline elements surpasses the level of the map. In this case, the system will automatically generate a new timeline element. If the user scrolls down, timeline elements that fully drop below the map will be removed from the scene. This way the length of the timeline is always optimal for the current situation. When a user scrolls to the top of the timeline (all pins have risen), the timeline disables scrolling upwards and the visualisation is fully extended. The same is true for the opposite, where the user collapsed all pins on the map, in that case it is not possible to scroll any lower.

We expect that the strength of this visualisation lies in seeing a clear order in which the photographs are taken. Seeing the order of places you visited, helps boost someones inferential memory [20]. Additionally, by enabling users to expand the visualisation in a vertical way, pins are less clustered making it easier to select specific pins. A possible weakness of the visualisation is that it can be hard to see the exact location of a pin when it is far away from the map because it is risen upwards. Another weakness of this visualisation is that every pin has a fixed vertical distance to the next pin and therefore the height tells little about the absolute date of a pin. The discrete visualisation presents a solution for this problem by making the height of a pin connected to a specific time frame.

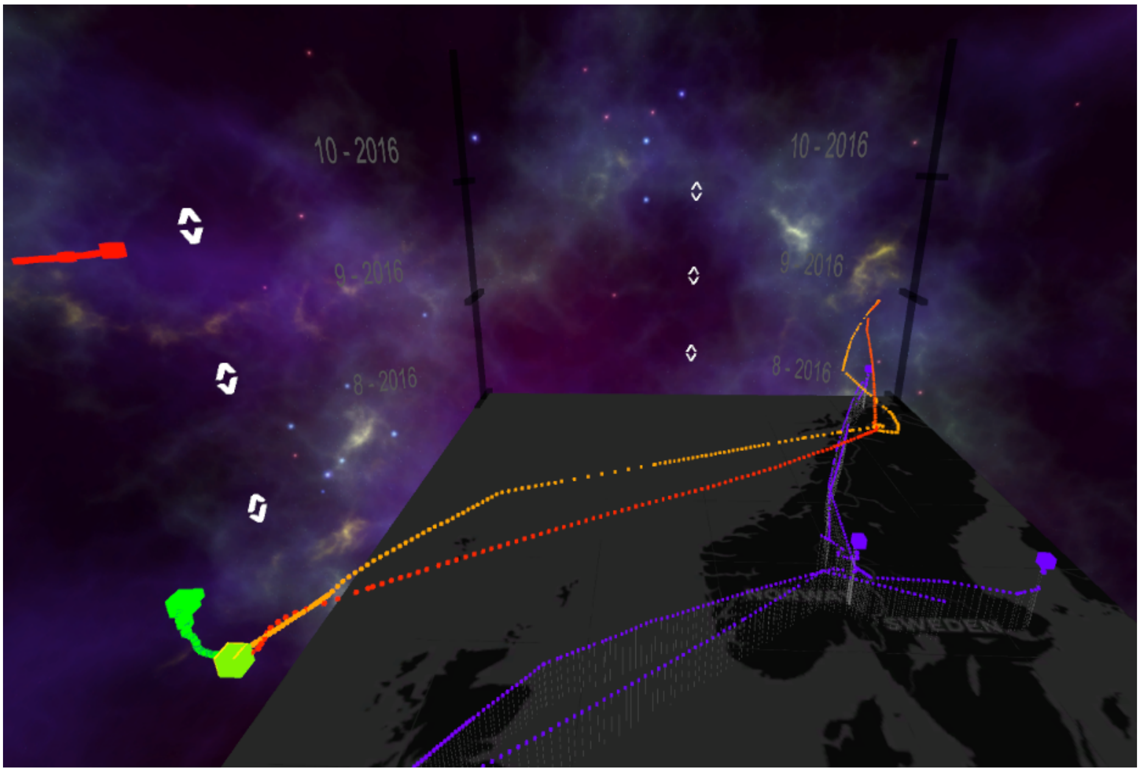

Fig. 8: Overview of the discrete visualisation

D. The Discrete Visualisation

The discrete visualisation borrows concepts from the continuous visualisation but makes these concepts more discrete. In this visualisation, pins are stored in different layers that are vertically stacked on top of each other as shown in image 8. Each layer represents a specific time frame and in this layer are all the pins from which the pictures are within this time frame. When the user loads this visualisation for the first time, each layer represents a month. Users can scroll the layers up and down in a similar fashion to the scrolling of the continuous visualisation. However, each scroll movement is discretised in this visualisation. This means that a user can scroll the interface with a distance that is equal to a multiple of the fixed distance in between layers. The map is also present in this visualisation and doubles as a limit on how far up or down the layers can be scrolled as the uppermost layer can never get below the map and the bottommost layer can never get higher than the level of the map. This limit, in combination with the discretised scrolling, makes sure that there is always one layer positioned at map level.

A key feature in the discrete visualisation is the ability to expand and collapse layers. At the start, each layer represents a month, meaning all data of that month is present in that single layer. When the user clicks on one of the expand buttons placed on each side of the map by aiming the controller at the button and using the trigger, the layer corresponding to that button will expand. This means new layers will be shown which represent the days of that month and those layers and will replace the old month-layer as shown in figure 9. The pins are repositioned from the month-layer to the day-layers. The expansion of a layer leads to an increase in the number of layers, causing the visualisation to grow vertically. This expansion process can be repeated for layers that represent days, splitting up a day-layer in several hour-layers. Collapsing a layer is the opposite of expanding a layer, reverting back to a layer of a greater time unit. For example, when a user triggers the collapse button of a day-layer, all day layers from the corresponding month will be hidden and the month-layer will reappear. To give the user a clear idea of what the current time frame is of the data they are looking at, the scrollable timeline on the corner of the map show a label displaying the current time frame. So for example for a month-layer, it could show '08-2016', for a day-layer '27-08-2016', and for an hour-layer '12h 27-08-2016'.

The layers are fully recursive making it easy to expand the system to also include other time units like years, weeks or minutes. Layers have a parent layer which represents a larger time unit (if the layer is not of the largest time unit) and have a list of child layers which represent a smaller time unit (if the layer is not of the smallest time unit). Layers can show or hide all their child layers by a single function call and the positioning of all layers is automatically set correct. The position of each layer is calculated by recursively calculating its internal height (the number of child layers active). Iterating in order from the bottom to the top of the visualisation over each layer and placing the layers on the sum of the internal height of all layers below the current layer, ensures that each layer is at the correct position. On start-up, all layers are created by iterating over the data and create a layer for each month/day/hour. This is done such that there are no extreme loading times when a user expands or collapses a layer but rather have one moment at the start-up of the visualisation in which the creation of all layers takes up some time.

We expect multiple advantages of this visualisation over the other visualisations. First, by dividing the data into layers, it enables users to focus on specific data and filter out irrelevant data by focusing on specific layers. Furthermore, by splitting the data into layers, you can break up clusters of pins. This separation is especially useful in situations where a lot of pins are close to each other. Another advantage of this visualisation is that finding specific dates within the dataset is relatively easy due to the recursive structure of the interface.

For this interface, we also implemented two additional features that were disabled for the experiments as they needed some adaptation to work more intuitively and due to time constraints had to be excluded from the visualisations. These features are the selection of layers and the suggestion system. Layer selection enables users to select specific layers and display them on

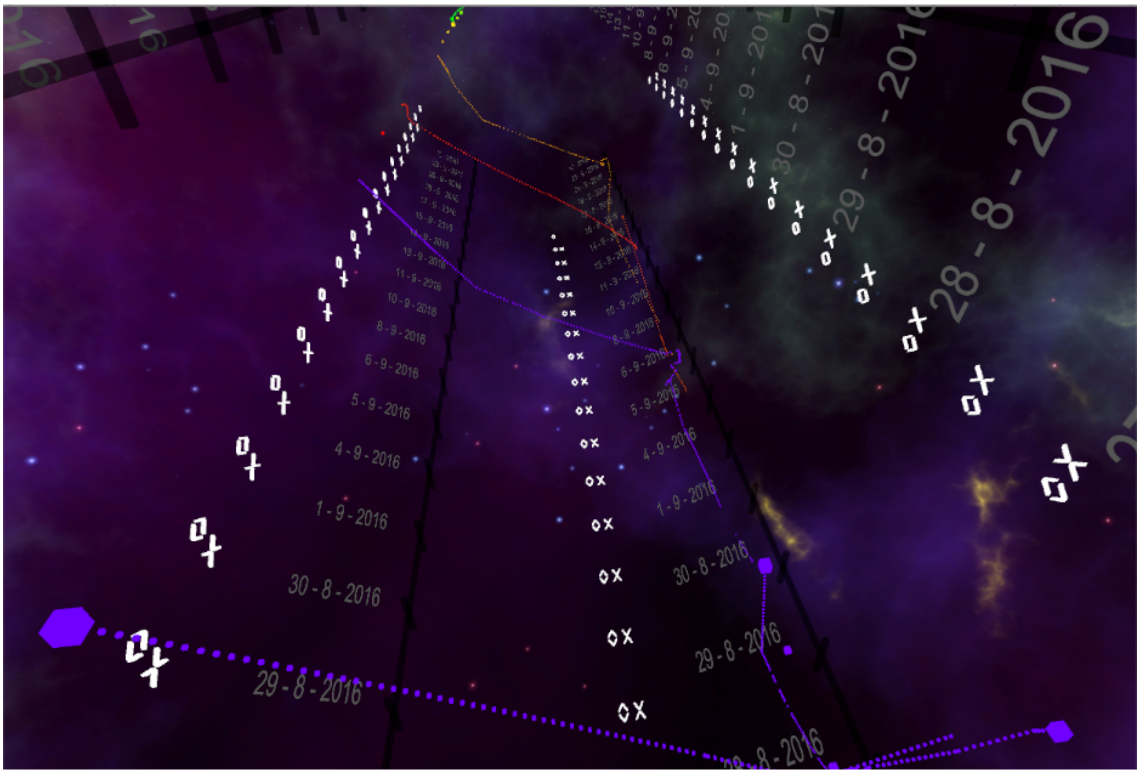

Fig. 9: The discrete visualisation with a month layer expanded into several day layers

the map, no matter if those layers are actually positioned on the map by scrolling them to the map level. The benefit of this feature is that a user can select multiple layers and see multiple specific time frames at a single glance. The suggestion system is related to layer selection. Whenever a user selects a layer, a pop up near the selection button appears. This pop up suggests multiple selection criteria which the user can select. If the user selects one of these selection criteria, the suggestion system iterates over all layers and checks for each layer if it meets the requirements of the selection criteria. If all criteria are met, the layer is selected automatically and thus placed on the map. For example, a user selects a day-layer with the date 06-08-2016. A pop up appears with the options to select 'every Saturday this month', 'every Saturday', and 'every first Saturday of the month'. With this suggestion system, it is possible to quickly select a pattern of layers without the hassle of individually selecting each layer. Doing this might reveal patterns in the data that were otherwise invisible to the user. Although these features are not in the final version, having created these features shows the potential of adding additional features to this system.

IV. Experimental Design

In this section, we will discuss the design of the experiments in this research where participants will be working with our visualisations. We will gather data while the participants are interacting with the system, but we will also ask the participants some questions before and after the experiments to get more insight into their behaviour and experiences. The experiments will be focused on four evaluation points: the enjoyment of the visualisations (how much does the participant enjoy using the visualisations), the usability of the visualisations (how well do the designs of the visualisations work), the effectiveness of the visualisations (how well do the visualisations enable users to complete the tasks), and how well the participants feel while using the visualisation (i.e. our system's nauseousness). The system supports two types of usage: using the system to explore the data and using it to find specific data. Both usage types will be evaluated in the three visualisations.

A. *Focus of the experiments*

The focus of the experiments is the effectiveness, usability, enjoyment and nauseousness of the system. The effectiveness of the system is evaluated because we want to see if the novel visualisations help the user with finding data. A reason why users might use our visualisation is to find specific data and therefore it is important that the visualisations are enabling users to do this. The usability is related to this as a more usable visualisation helps the user to work with the data. The enjoyment is of importance as we want to create visualisations that are fun to use. Virtual reality has multiple benefits and one of the main reasons for people to use VR today is to have fun. If our system causes nausea by its users, we can expect the enjoyment of the system to go down as well as the potential of users wanting to use the system again. Therefore we also explicitly test the nauseousness of the system.

Rather than testing the performance of the participants, we want to test how well the interfaces perform on the aforementioned focus points. The enjoyment of the interfaces by the participants will be tested using a quantitative and a qualitative method as the enjoyment is highly reliant on the individual experience and there needs to be room in the evaluation method for these experiences to be expressed. For the usability we will use a qualitative method (as described in Section IV-B) as a qualitative method can pinpoint the pros and cons of our system in more detail than a quantitative method. The effectiveness of the system will be evaluated using a quantitative method as the effectiveness is a more absolute evaluation point and can be measured automatically, giving us some hard data to work with. The nauseousness effect of our interfaces will be tested with a quantitative method. The combination of using both quantitative and qualitative methods in our experiments is beneficial to the research as these methods enable us to focus on our four evaluation points (usability, effectiveness, enjoyment and nauseousness) in one experiment. Furthermore, this combination gives us also a good idea about what the overall strengths and weaknesses of our interfaces are.

1) Testing the exploratory nature of the interface: We test the ability to explore the data using the different interfaces by letting participants browse the data and give them some general tasks related to the data. These are tasks in which the user is not looking for specific images, but rather explores the data. This is relevant as users might use our system to get a better idea of the data they have. Our interfaces are designed to let participants explore the data using spatio-temporal information and this should be reflected in the tasks given to the participants. The tasks given to the participants should not be too obvious (e.g. in what country does the subject live) nor too specific (e.g. On what street was the subject at 9:27 a.m. on Friday the 14th of May 2007). The following questions/tasks meet these requirements:

- At what day and time did the subject arrive in Oslo for the first time?
- What place did the subject visit first, Oslo or Stockholm?
- What is the first day present in the data?

2) Testing how well users can find specific data: Users will also be tasked to find specific data points from all the data in the large data set. We created tasks that focus on finding specific data to see how well the visualisations suit this style of use. This is relevant as users of our system might be looking for specific information. The tasks for testing this should be specific enough such that users need to interact with specific data points to answer the questions. The following questions will be used to test if the interfaces enable users to find specific data:

- What did the subject drink after he arrived in Tromsø (North Norway) for the first time?
- At what time did the second flight from Oslo to Tromsø (North Norway) depart?
- Where did the subject go to on 23-8-2016?
- What is the first city the subject went to after arriving in Dublin after visiting Oslo for the second time?
- How long did the drive from Collooney (North West Ireland) to Dublin take?
- How many different planes did the subject take in all the data in the dataset?

B. Data Gathering

Our research question will be answered with data gathered before, during, and after the experiments. The questionnaire before the experiment will be used to get an insight into who our participants are. Since the experiments are focused on testing the interfaces and not the participants, some of these questions are related to the interfaces, e.g. 'How much experience do you have with virtual reality?' and 'How often do you use systems for browsing images?'. The data from these questions will be used to see if, for example, people with a lot of virtual reality experience will use our systems differently than people who do not have much/any virtual reality experience. This should tell us that our system is intuitive to use for people who know the standard way of interaction in VR or maybe just the exact opposite. We will also gather data about the demographics of our participants to see if we can find any correlation between those demographics and our evaluation points. We will also ask participants if they have any eye deficiencies as previous research in similar virtual reality experiments have found effects of wearing glasses on the enjoyment of VR [3, 28].

During the experiments, we will gather data for a quantitative analysis of the effectiveness of our visualisations. We will measure the performance of the participants as an indication as to how well the interfaces fit the tasks given to the participants. Most of these data will be automatically gathered by the system. The data which will be gathered for each task are:

- The time it takes for the participant to complete the task
- The time the participant is looking at images on the image wall
- The time the participant is spending on navigation and browsing the interface
- The number of (image) pins that were opened by the participant
- For the discrete visualisation: the number of times the user expanded and/or collapsed a layer

Some data will not be gathered automatically as recording this by hand is more suitable than doing this automatically. These data are:

- A Boolean indicating if the participant found the right answer
- The number of attempts it took for the participant to get the right answer
- The wrong answers that were given by the participant

These data will be used to explain anomalies in the data. Furthermore, we will record a screen capture and a sound recording during the experiment. With these recordings, we can see whatever the participant was seeing and hear whatever the participant said during the experiment. This will be used to explain anomalies in the data.

After each round of tasks, we will ask questions about the experience of the participants. As stated earlier on, the experiments will focus on four aspects: effectiveness, enjoyment, usability and nauseousness. The effectiveness of the interfaces is measured by the quantitative data gathered during the experiments. Testing the enjoyment of the participants is done by asking the participant to fill out a questionnaire when they finished the tasks for an interface. Participants might rate their enjoyment of the interfaces higher by the fact that they enjoyed being in a virtual reality world since this will likely be a novelty for the majority of the participants. This is something we should be wary of as this is not the aim of our research. Therefore we start with four questions about their virtual reality experience. These questions are mostly filler questions and focus specifically on working with the virtual reality headset. After these four questions we shift the attention from the virtual reality headset to the visualisations. Our aim is that as we shift the attention of the questions to the visualisations, participants also shift their attention to the visualisations. We do this to prevent that the participants let the novelty of using a VR headset influence their opinions on the visualisations. The questions for testing the enjoyment of the participants are derived from the Game Experience Questionnaire (GEQ) [16]. This is a questionnaire aimed at testing the experiences of players playing video games. Our research does not have any game elements, however, since our virtual reality environment has similar characteristics to a video game environment we argue that the GEQ is still relevant for our research. We omitted eight items from the GEQ that were not relevant to our research. The

TABLE I: Selected items from the Game Experience Questionnaire [16] with their original item number and category. Table adapted from [16].

#	Item	Category
1	I felt content	Positive affect
2	I felt skillful	Competence
4	I thought it was fun	Positive affect
6	I felt happy	Positive affect
7	It gave me a bad mood	Negative affect
8	I thought about other things	Negative affect
9	I found it tiresome	Negative affect
10	I felt competent	Competence
11	I thought it was hard	Challenge
12	It was aesthetically pleasing	Sensory and Imaginative Immersion
14	I felt good	Positive affect
15	I was good at it	Competence
16	I felt bored	Negative affect
17	I felt successful	Competence
19	I felt that I could explore things	Sensory and Imaginative Immersion
20	I enjoyed it	Positive affect
22	I felt annoyed	Tension/Annoyance
23	I felt pressured	Challenge
24	I felt irritable	Tension/Annoyance
25	I lost track of time	Flow
28	I was deeply concentrated in the game	Flow
29	I felt frustrated	Tension/Annoyance
30	It felt like a rich experience	Sensory and Imaginative Immersion
32	I felt time pressure	Challenge
33	I had to put a lot of effort into it	Challenge

items of the GEQ are subdivided into seven categories, which are: competence, sensory and imaginative immersion, flow, tension/annoyance, challenge, negative affect, positive affect. The items we selected to be used in our research are shown in Table I (table adapted from [16]). Participants will be able to give a score on a scale of 1 to 5 about how much they relate to each item with 1 being 'not at all' and 5 being 'extremely'.

The usability will be measured using a qualitative method. The definition we use for usability is from Lewis [24] who states that usability consist of the systems usefulness, the quality of the interface, and the quality of the information in the interface. At the end of the GEQ questionnaire, the participants are asked to point out the strengths and weaknesses of the interface. Furthermore, they are asked to suggest new features that would improve the system. Before they submit the (digital) questionnaire form, the answers are discussed with the experiment leader to dive deeper into the answers of the participants and to clear up any unclear answers. Furthermore, parts of the Game Experience Questionnaire that will be used to test the enjoyment of the interfaces can give some additional insight into the usability of the interfaces.

Additionally, we ask participants to comment on certain features of the visualisation. These comments can later be used to pinpoint specific parts of the interfaces that worked well and parts that need improvement. After the experiments, the experiment leader will discuss the answers to these questions with the participants to get more in-depth comments about the visualisation/features.

The last thing that we ask the participants about is if they experienced any motion sickness. It is not unlikely that there will be participants that experience motion sickness. However it is hard to do any predictions about this as the occurrence of motion sickness depends on quite some factors, some of which depend on the system and some of which depend on the task or the user [23]. Since the system can be part of the cause of people experience motion sickness, it is good to get an idea if our environment creates experience-disrupting levels of motion sickness. If this would be the case, we need to look into its causes and try to fix the issues for future research of the system. To test for motion sickness we use the virtual reality sickness questionnaire (VRSQ) [22]. This is a questionnaire specifically developed for this purpose and

TABLE II: Virtual reality sickness questionnaire (VRSQ) [22]. The total score can be calculated by this formula: ((Oculomotor Score/12)*100 + (Disorientation Score/15)*100))/2

VRSQ symptom	Oculomotor	Disorientation
1. General discomfort	x	
2. Fatigue	x	
3. Eyestrain	x	
4. Difficulty focusing	x	
5. Headache		x
6. Fullness of head		x
7. Blurred vision		x
8. Dizzy (eyes closed)		x
9. Vertigo		x
Total	[1]	[2]

is based on the simulator sickness questionnaire [21], which is a well known and often used method for testing sickness in visual simulators. The VRSQ tests sickness on two components: oculomotor and disorientation. Participants will fill out the questionnaire by scoring 9 items on a scale from 0 (none) to 6 (severe). Each item is a symptom related to motion sickness and is displayed in Table II [22]. In the end, three scores can be calculated, a score on oculomotor sickness, a score on disorientation related sickness, and a weighted average that indicates a total level of sickness.

C. *The Experimental Setup*

At the start of the experiments, participants will be given a short introduction about the experiments and the goal of the experiments. This includes an introduction into lifelogging so that participants know what kind of data they are presented and to get a better understanding of the experiment as a whole. Furthermore, the participants will be told that they will get to use a virtual reality headset to explore this data. Participants will be informed that they can take a break or even stop the experiments completely if they want to. This is important as virtual reality can make the participants nauseous and we do not want to make our participants feel sick. Since motion sickness is a possibility, participants are obligated to sign a consent form. After signing the consent form, participants are also asked to fill out a questionnaire with some questions about their demographics (age, gender, possible eye deficiencies, and experience with VR and image browsing).

Following these preliminary steps, it is time for the participants to put on the HTC Vive VR HMD. To get the users familiar with the interface, participants will start with using the interface without any objectives. This exploring will follow a script. First, the participants will get to wear the VR headset and are told to look around. The participants will not get to use the controllers just yet, this stage is just so they get familiar with wearing the headset and being in a virtual reality environment. After this, the experiment leader will give the participant a controller. The controller can be used for moving around and interacting with the interface. The participants will be instructed to use the controller for moving around. This can be done by using the trackpad. Participants will be asked to use the trackpad to turn around and to move forwards and backwards. Participants are tasked to move closer to the map, as to see the different zoom levels of detail of the map as well as the reveal mechanic of the pins. Furthermore, participants are asked to aim the controller at a pin and notice the date label that appears near the controller. Additionally, participants are requested to interact with a pin to get a grasp of how to interact with the image wall. Since the first interface will be the baseline visualisation for everyone, this is where the introduction of the controls will stop for now. When the participant is ready to move to one of the other two interfaces, additional instructions are given to the participants to learn them how to use the other interfaces. The order of the visualisations is a dependent variable. To half of the participants, we present the continuous visualisation after the baseline and before the discrete interface. The other half of the participants will get the visualisations in the order: baseline, discrete, continuous. All participants will start with the baseline interface, since this is a good introduction into the environment as a whole and since this interface does not contain any features that the other two interfaces do not contain. In the continuous visualisation, participants will be instructed to use the timeline to scroll through the interface. This will get them familiar with the main mechanic of the interface. For the discrete visualisation, participants are asked to scroll the interface, but also to expand and collapse layers. After these additional instructions, the participants can ask questions about the interfaces if they want to. When all questions about the visualisations are answered, the participants will start with the tasks for the current interface.

The experiments consist of the participants completing three of the tasks as described in Sections IV-A.1 and IV-A.2 for each interface. For each task in the experiment, participants are requested to complete the task. We explicitly tell participants that they do not need to do this as fast as possible as we do not want them to think it is a race but rather see how the participants would be using the system in a more natural situation. As specified, participants can always take a break or quit the experiment for whatever reason, including motion sickness. If a participant chooses to take a break, the experiment will be paused and no data will be gathered during the pause. After completing the tasks for one of the interfaces, participants will fill in the questionnaires and we will continue with the next interface. A task is completed when the user finds the answer to the question and tells this out loud to the experiment leader who will confirm whether the answer is correct or incorrect. If the answer is incorrect, participants are asked to try again. If the participant is not able to find the correct answer or when the experiment leader notices that the participant is struggling for too long with one task, participants are told the current task is registered as not-complete. Again, participants are reminded that this is not a problem as it is not them who are tested, but rather the system that is. After finishing a task, the experiment leader will follow up with a new one until no tasks are left for the current interface. When a participant completes all tasks for the three interfaces and after answering all questionnaires are finished, the participant is allowed to ask questions about the experiment if they have any. At last, the experiment leader will thank the participant for taking part in the experiment and conclude it.

D. Participants

For the experiments in this research, we will have 12 participants. These participants will be students from the Utrecht University who are friends and acquaintances of the experiment leader. Participants were not compensated in any way for the experiments other than the cookies and lemonade that were present at the experiments. Our participants are likely unfamiliar with lifelogging and we will introduce them to this topic to get a better grasp of the data they will be working with. Furthermore, it has to be noted that in a real-life situation, it is likely that a user of our interface will be familiar with the data as it would be data that the user has gathered themselves. Since it is not feasible to let a user collect enough data themselves such that they could use their own data during our experiments, all participants will be using the same data (data is from [25]). This of course influences the performance of the participants as they are not familiar with the data. However since all users are unfamiliar with the data, we argue that the effect of this on the performance of the participants is similar amongst all participants and none of the interfaces is likely to benefit from this. Additionally, the tasks as described in Sections IV-A.1 and IV-A.2 are specifically designed for participants without preexisting knowledge of the data in mind.

V. RESULTS

In this section, we will discuss the results of the experiments. Before diving into each of the four evaluation points, we start with some general results and discuss demographic data.

First of all, there were twelve participants with an average age of 23.5 (*STD* = 1.8). Eight of the participants were male, three were female and one participant identified as non-binary. The participants were all students of the Utrecht University, studying for various degrees ranging from social sciences to computer science. All participants were right-handed; three participants had eye deficiencies and wore glasses to correct these deficiencies. For three participants the experiments were their first time using a virtual reality headset, the other nine all had limited experience with VR. Seven of the participants used applications with virtual maps more than ten times a month, four participants used such applications regularly and one participant did not use it on a regular basis but was familiar with such applications. Seven participants used applications for browsing images more than ten times a month, the other five used them regularly.

The experiments took about 1.5 hour per participant. For most experiments, all the data collection perform as expected. However, due to some technical errors, four individual tasks did not collect any data. All the results in the following sections take this missing data into account and are correct for it. Furthermore, the order in which the visualisations were presented to the participant was changed around by accident for one participant. In the end, we did not use any of the quantitative data from this participant and redid the experiment with another participant to make sure we had enough participants. The qualitative data from this user was used as the order of the visualisations had little effect on these data.

A. *Enjoyment*

After completing three tasks given to the participant for each interface, they were asked to answer questions of the Game Experience Questionnaire as described in Section IV-B. The data gathered from these questions are presented in figure 10. This graph shows two trends. First of all, the continuous visualisation seems to outperform the other two visualisations on all categories except flow. Note that for the 'Negative Affect', 'Challenge', and 'Tension/Annoyance' categories a lower score is positive for the visualisation. For the other categories, higher scores are better. This indicates that participants enjoyed the continuous visualisation the most of the three visualisations. This can be confirmed by the experiment leaders observations who saw that most participants were the most enthusiastic about the continuous visualisation. Some of the participants who got the visualisation in the order baseline $\rightarrow$ discrete $\rightarrow$ continuous noted that they understood why the experiments were ordered this way as they saw that the visualisations got better and more fun with each new visualisation they used. This was however not intended to be the case as we changed the order around for half of the participant to counter any effects of the order of the visualisations. This example shows that those participants had a clear preference of which visualisation they liked the most.

The second observation that can be made from figure 10 is that the discrete visualisation scores similar to the baseline visualisation with the exception of the 'Competence' category. This shows that all additional features that are added in the discrete visualisation in comparison to the baseline visualisation did not add to the enjoyment of the visualisation. Again, the observations from the experiment leader confirm these results; most participants thought the features were interesting but also difficult to use, sometimes even leading to some frustration. This explains why even though the visualisation showed the potential to be more fun, it did not deliver on that promise.

Four of the twelve participants seemed to better understand and use the discrete visualisation than the other eight based upon observations by the experiment leader. When looking at the enjoyment of these four participants of the discrete visualisation in comparison to the other eight participants, we see this effect in the 'Negative affect' and 'Challenge' scores of the GEQ. We plotted these differences in figure 11. Most scores of the GEQ are pretty similar between the two types of participants, but it is clear that the participants who did not have an intricate understanding of the interface struggled more with the visualisation as it was perceived as more challenge and evoked more negative emotions for those participants than for participants who did seem to have an intricate understanding of the interface.

Ouwehand [28] and Bolwork [3] found that participants who wore glasses enjoyed their virtual reality experience less. We checked to see if a similar pattern arose in our data by comparing

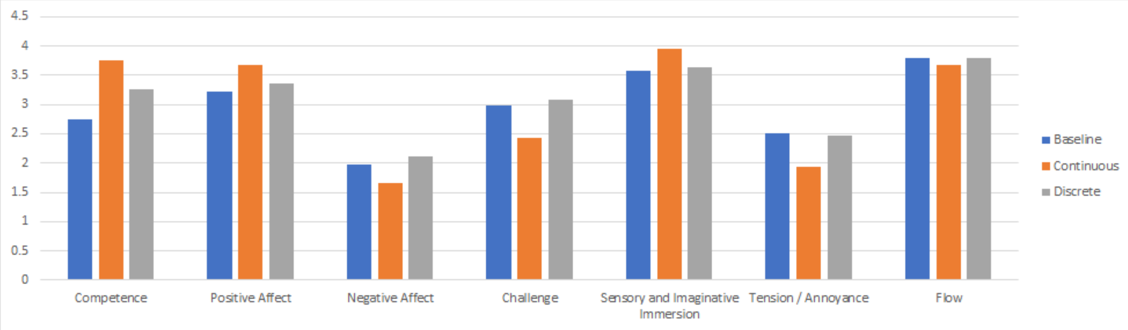

Fig. 10: Game Experience Questionnaire results per visualisation and category. The scores are the averages of the Likert-scale answers for each category. For Competence, Positive Affect, Sensory and Imaginative Immersion, and Flow higher scores are better. For Negative Affect, Challenge and Tension/Annoyance lower scores are better.

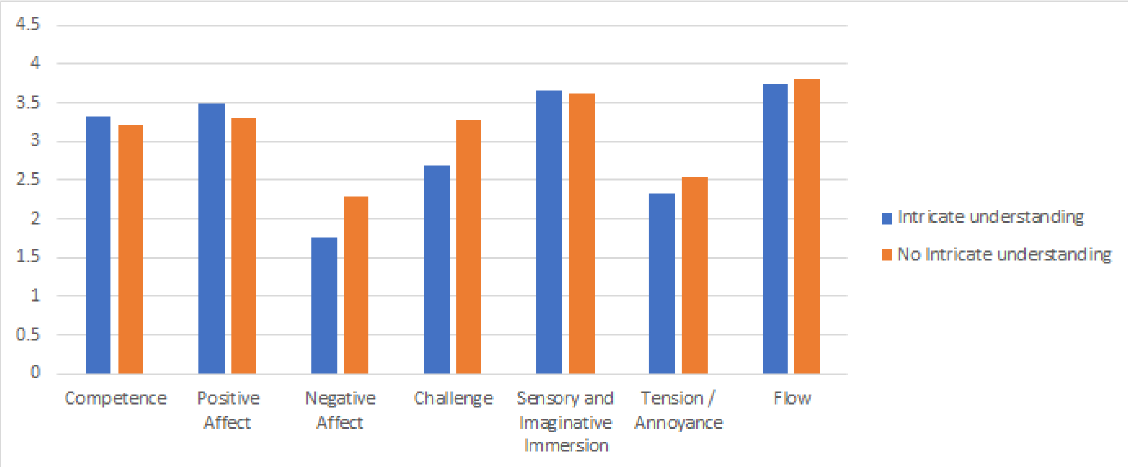

Fig. 11: Game Experience Questionnaire results of the discrete visualisation, subdivided by participants who had an intricate understanding of the visualisation (N = 4) and those who did not (N = 8).

the results from the Game Experience Questionnaire per category between the participants who wore glasses and those who did not. We did not find such a pattern, as the scores between participants with eye deficiencies and those without eye deficiencies did not differ.

B. Usability

The usability related data were gathered by asking the participants open-ended questions about their experiences with interfaces. Some of the categories of the game experience questionnaire (GEQ) used for measuring the enjoyment also show insight into the perceived usability of the visualisations. In this section, we are looking at those data and present the usability results.

We start by discussing the open-ended answers in which the participants told us about the strengths and weaknesses of the interfaces. All the answers of the participants are aggregated and categorised based on their contents. We discuss every visualisation and start with the categories with the most entries. As a reminder, the definition we used for usability is from Lewis [24] which focuses on the system's usefulness, the quality of the interface, and the quality of the information presented by the interface.

First of all, we discuss the baseline visualisation as it was the first one that every participant got to work with. Most of the feedback on the baseline visualisation was related to the different trips in the were shown as nine out of twelve participants made comments about this. Participants seemed to enjoy seeing that the it was easy to see that the man took several trips to Scandinavia and one of the ways participants were able to see this so quickly were the newly introduced pin colours. The following quote from a participant shows this: ”[The interface was a] Good visualisation of the routes the participant took. Using the colours, it was easy to make a distinction between the different dates and routes.” Participants also liked the overview that the baseline visualisation gave as one participant stated: ”Broadly speaking, you can immediately see where someone went and what places he often visits.” Participants also pointed out weaknesses of the interface. Seven of the twelve participants commented on the clustering of pins, mentioning that at locations where there were lots of data (e.g. Dublin and Oslo), it was hard to select a specific pin as it was too crowded with pins. This also made it harder to track the chronological order of the pins. This relates to our second research goal and we expected these comments in the baseline visualisation. Another problem with the visualisation was that the lines connecting the pins to the map, sometimes occluded the geographical names on the map, making it hard for participants to read the exact location of the pins on the map. Also, the resolution of the map was too low at times to be able to read the words on the map. Combining these two problems, participants struggled with reading the geographical names on the map.

The next visualisation we discuss is the continuous one. All twelve participants commented on the positive effects of the time line, expressing how it improved the visualisation. Participants liked that it was easy to see the chronological order of the pins as ”the continuous time line made it pretty easy to follow the journey [of the man who collected the data]”. Furthermore, seeing the time labels on the time line were a good indicator of where in time participants were looking at. One participant told us ”The data on the time line made the visualisation a whole lot clearer [...] as it is way easier to find what you are looking for now”. The problem with too much cluttered pins at places where there was a lot of data was not present in this visualisation accordingly to the participants: ”The time line created a situation where the pins did not overlap too much”. Participants also addressed the weaknesses of this interface, although there was much less consensus about the weaknesses than about the strengths of the interface. Two participants indicated that they did not agree with the order of the pins, which was that the oldest pin is at the top of the visualisation and the newest pin at the bottom. Another problem that a participant mentioned was that ”since there is no constant time interval [in between the pins], it is not easy to tell how long someone stayed at the different places”.

For the discrete visualisation, most participants commented on the overview of this visualisation, as having different time layers made it easier for participants to focus on specific data. One participant worded this sentiment quite clearly: ”It is easy to get a clear overview of all specific moments/time intervals. Furthermore looking into the movement around a specific location is easier by splitting the layers.” Another participant shared this opinion, ”Being able to filter on time, something which reduced the amount of information at one moment, made its easier to see what happened at what time.” Some participants did not agree with these statements at all: ”Adding days to the visualisation [by expanding a month-layer] did not improve the overview. In fact, expanding layers decreased the overview.” Another participant further elaborated on this by commenting ”Having the disconnected layers made the data lose its continuity, while this continuity is exactly what it is all about regarding this [lifelogging] data.” This difference between participants who did find this visualisation usable and those who did not, correspondents with difference between participants who had an intricate understanding of the interface and those who did not, as discussed earlier on in Section V-A.

To compare these strengths and weaknesses in between the visualisations, we counted how many comments on strengths and weaknesses were given per visualisation. We read each comment and counted on how many different aspects they noted. Some of the comments where not specifically related to a single interface (e.g. problems with navigating using the controller) and those comments are not counted here. The baseline visualisation counted 11 strengths and 14 weaknesses. For the continuous visualisation these numbers were 23 strengths and 10 weaknesses and for the discrete visualisation this was respectively 15 and 16.

The Game Experience Questionnaire used for testing the enjoyment of the visualisations can also give some insight into the usability by looking at the categories competence, challenge and flow. In figure 10, the results of these (and the other) categories are shown. The continuous visualisation outperformed the discrete visualisation on competence, which in its turn outperformed the baseline on that measure. This shows that participants felt like they were able to more efficiently use the continuous visualisation than the other visualisations. Also the continuous visualisation was perceived as less of a challenge than the other two visualisations, once again confirming the high usability of this visualisation. For flow, there were no differences found.

We also asked participants to comment on specific features of the visualisations to see what specific parts of the interface worked well. For example we asked them to compare the scrolling mechanics in both visualisations. Most participants seemed to favour the continuous scroller found in the continuous visualisation. The discrete scroller was found to be too slow, even though the scrolling speed was equal to the continuous visualisation, but the resulting scroll movement was rounded to the nearest layer. Observing the participants use the scrollers, we saw that participants were likely to use the scroller on the corner of the map closest to them. Scrolling large amounts with the scroller close to the user is harder to do as grabbing a high point on the scroller is harder to do when up close in comparison to the ones farther away from the user. Furthermore, participants seemed to struggle at times to grab the scroller which led them to focusing on that object too much instead of looking at the data they were scrolling. Additionally, participants indicated they would like to get some visual feedback on grabbing the scroller as they sometimes were unsure whether they actual grabbed it correctly.

Another feature discussed with participants was the navigation. The main issue people reported with the navigation was that people lost the overview of where they were. The main cause for this, as described by the participants, was that the camera moved forwards in the direction of the camera when pressing the forward button. Rather, participants would like to see that the navigation was directed by the rotation of the controller, such you would fly to where ever you are aiming at, instead of where you are looking at. This would also solve another cause of losing overview, namely that participants had troubles with moving to the side. In the current implementation, the user had to look away from the data at times to move to the position they wanted, which led to loosing track of where they were positioned in the space. Also, participants indicated that travelling long distances took too much time. Furthermore, moving too the side without having many objects to look at, was a cause of nauseousness. Therefore changing this navigation method could have a positive effect on multiple evaluation points.

When asked about the date time label near the controller when the controller was aimed at a pin, almost all participants responded positively. Participants often used the label when finding answers to the tasks and stated that it was a useful feature. Sometimes, the label was a bit redundant as in both the continuous and the discrete visualisation there were the time lines that could be used. However, participants still liked the functionality of getting temporal information of a specific pin. The positioning of the date label could improve a bit a some participants stated that it got out of the field of view when aiming the controller at a pin.

Finally, participants were asked to comment on the colouring of the pins. Most participants seemed to like the colours, especially the different colours of the trips/flights as those colours made it easy to distinguish between different trips. Although most participants seemed to be using these colours, they did not understand the logic of the colouring. Most guessed that each individual colour indicated an single trip, while in reality it indicated an individual day. One participant told that the colours were hard to use for him as he was colour blind and complementary colours were close to each other, making it hard to distinguish between the colours.

C. *Effectiveness*

Analysing the effectiveness of the system can be done with the quantitative data that was gathered during the experiments. This data gives us insight into how well the participants were

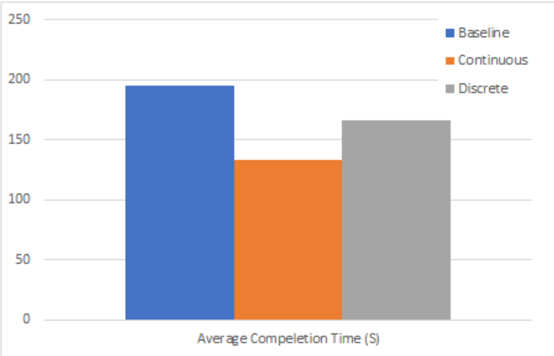

Fig. 12: Average completion time in seconds of all tasks per interface.

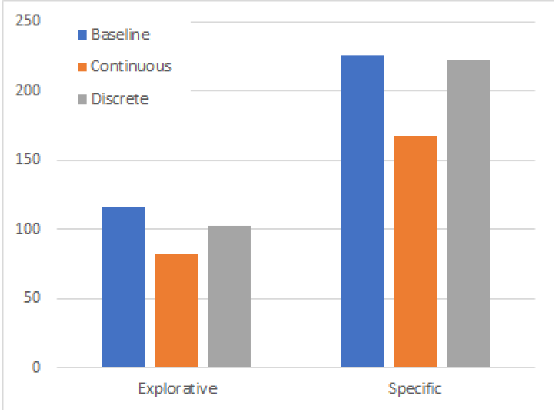

Fig. 13: Average completion time in seconds per interface per use type.

able to perform the tasks. Comparing the different visualisations gives us insight into what visualisation enabled the user to complete the tasks as good and quick as possible. The first data we look at is the average completion time per task grouped per visualisation, which is shown in figure 12. The tasks in the continuous visualisation were completed the fastest on average while in the baseline the tasks took the longest.

We can also look at the average completion time of all tasks per interface on a per task basis. This data is shown in figure 14. This figure shows that except for task 6, the continuous visualisation has the lowest or is close to having the lowest average completion time per task, indicating that the continuous visualisation is on average almost always the quickest visualisation to use for completing these tasks. The discrete visualisation and the baseline visualisation trade turns for which has the highest average completion time. Looking at the data for task 6 (What is the first city the subject went to after arriving in Dublin after visiting Oslo for the second time?), we see that one of the participants who used the continuous visualisation for this task had a remarkably high completion time, skewing the results. When excluding this outlier from the calculations, the continuous visualisation has a lower average completion time than the other interfaces on task 6.

In Sections IV-A.1 and IV-A.2 we described two types of tasks that the participants had to perform, namely task of a explorative nature and tasks testing the ability to find specific data. In figure 13 we split the average completion time on these types of tasks. As expected, the exploratory tasks were simpler and resulted in a quicker completion time than the tasks in which participants were tasked to find specific data. Furthermore, participants using the continuous visualisation were the quickest finishing their tasks in both usage types. For the exploratory tasks, the discrete visualisation seemed to be faster than the baseline visualisation; for the specific data finding tasks this difference is negligible.

Another way of looking at the effectiveness of the interfaces is looking at the percentage of tasks completed per interface. Remember, participants were allowed to stop a task if the answer was too hard to find. The completion percentage per interface is shown in figure 15. The continuous visualisation had a completion percentage of 100%, showing that no participant stopped a task unsuccessful when using the continuous visualisation. The percentages of the other two visualisations are close to each other with 89% and 91% for the baseline and discrete visualisation respectively. One participant completed two tasks unsuccessfully, the other uncompleted tasks are distributed over the other participants.

Similar to the completion percentage is the average number of wrong answers given by the participant before the task was ended (either successful or unsuccessful). This data is shown in figure 16. As visible, the results of the baseline visualisation were once again the worst, with an average of 0.75 wrong answers per task. The average for the continuous one was 0.6 and

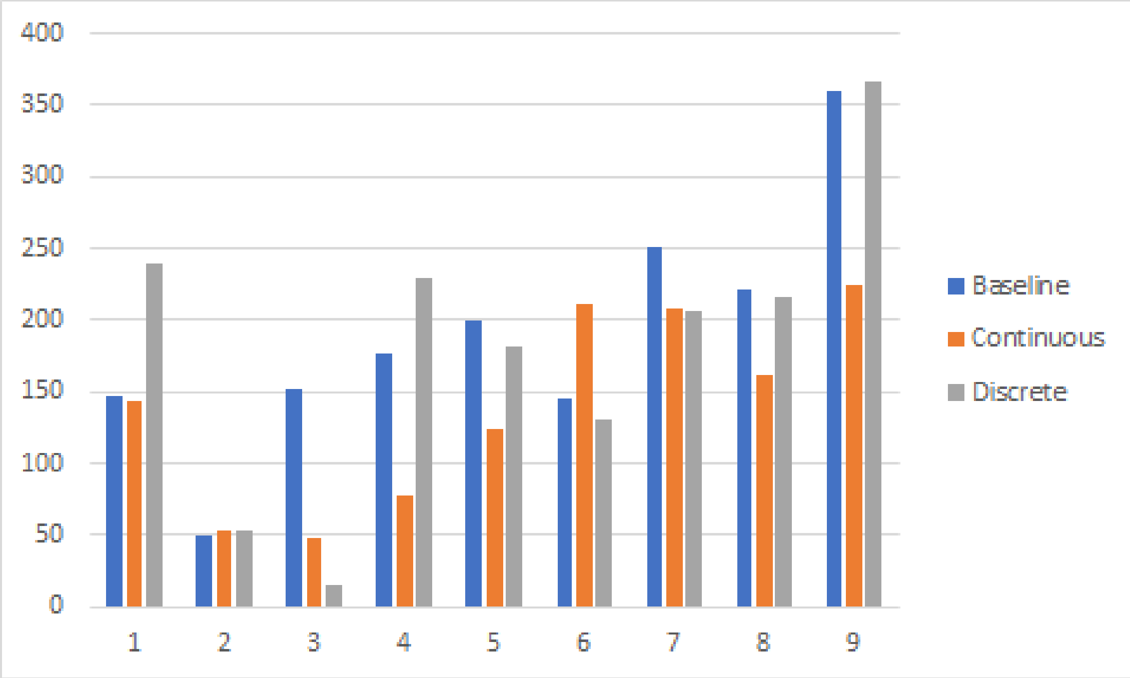

Fig. 14: Average completion time in seconds per task per interface.

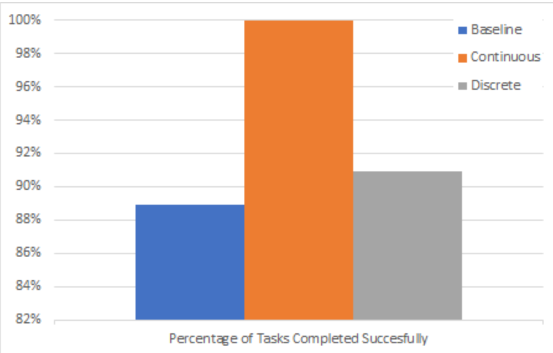

Fig. 15: Successful completion percentage of the tasks per interface.

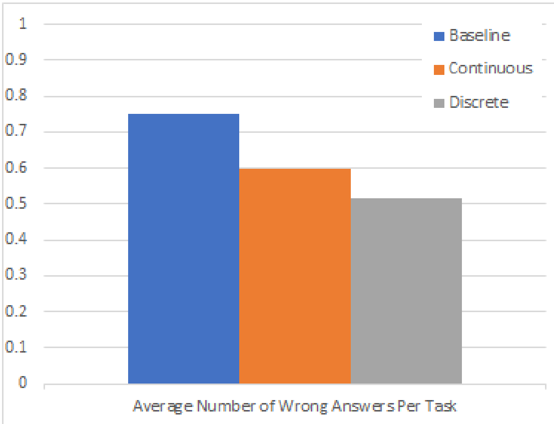

Fig. 16: Average number of wrong answers given by the participant before completion of the task.

for the discrete interface, this number was 0.52. Related to this data is the percentage of tasks completed without a wrong answer. For the baseline visualisation, 52.8% of the tasks were completed without a wrong answer. For the continuous and discrete visualisation, this number was 71.4% and 69.7% respectively.

The last measure of the effectiveness of the interfaces is related to the pins of the visualisation, namely the number of unique images visited per task and the number of pins opened per task which are visualised in figure 17. The most amount of images and pins were opened in the baseline visualisation; in the continuous interface the least amount of images and pins were opened. This indicates that participants were able to find the right answer in less tries in the continuous visualisation than in the other visualisations.

As stated above, there were two types of users for the discrete visualisation: those with and those without an intricate understanding of the visualisation, as described earlier. We checked if we could find any patterns in the quantitative data that could distinguish between the types

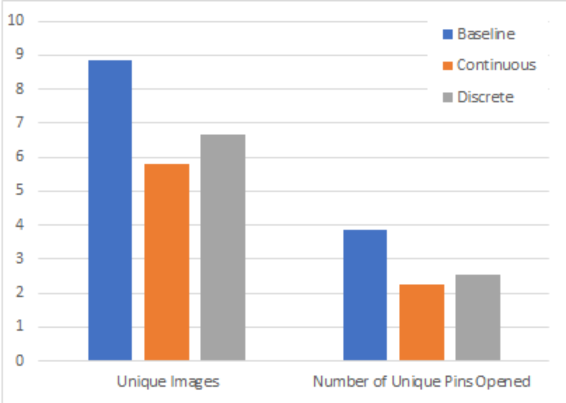

Fig. 17: Pin related data per visualisation: the average number of unique images opened per task and the number of pins opened per task.

of users. We did not find such patterns as there was too little data to work with to say anything about this possible difference between user types in the effectiveness of the interface.

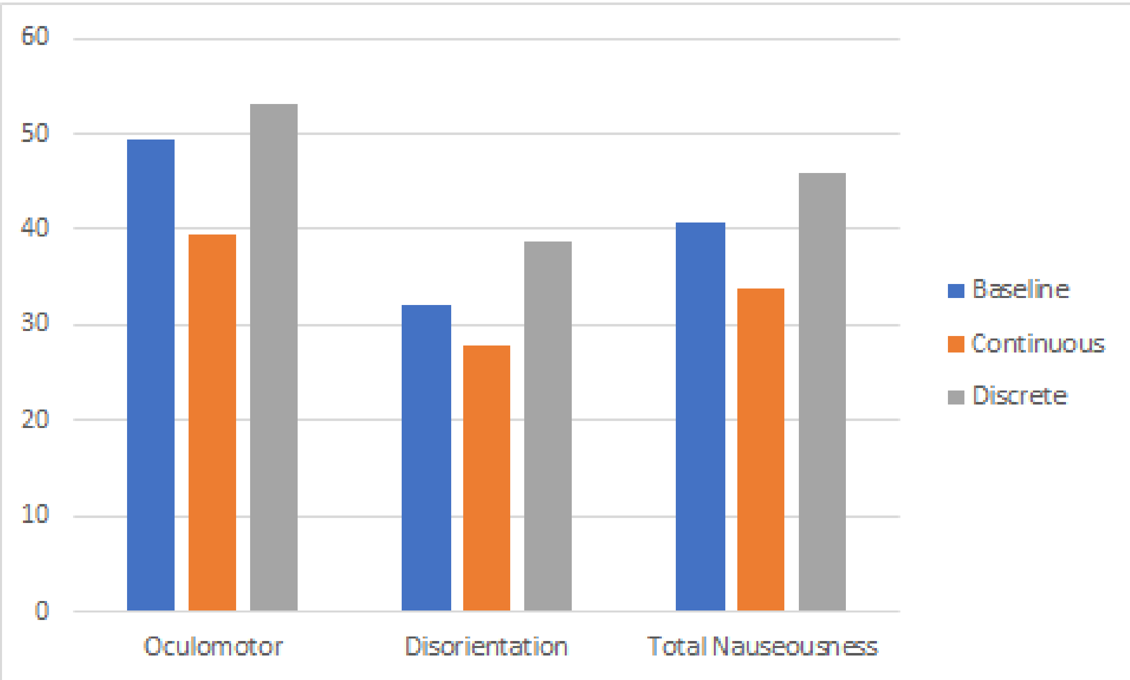

Fig. 18: Results from the Virtual Reality Sickness Questionnaire per visualisation split on oculomotor, disorientation, and total nauseousness.

D. Nauseousness

To test if our visualisations make participants feel sick, we used the Virtual Reality Sickness Questionnaire (VRSQ) by Kim et al. [22]. This questionnaire uses nine questions that can be divided over two types of sickness: oculomotor and disorientation. Furthermore, a total can be calculated by taking a weighted average of the two types. The results of the VRSQ in our research are shown in figure 18. From this figure we can see that our visualisations had some problems as it mainly induced oculomotor related sickness. This is in line with the qualitative data about the nauseousness, where participants indicated that the navigation within the system was mostly the cause for them feeling a bit nauseous. Especially rotating by using the controller was an often mentioned cause of generating illness when using our visualisations. Furthermore, figure 18 also shows that the continuous visualisation performed the best on the VRSQ and the discrete visualisation performed the worst. This is again in line with the feedback we received from the participants. The discrete visualisation made participants feel nauseous as the visualisation could get tall (by expanding layers) and looking up and down between the layers made participants feel nauseous, especially in combination with navigation by using the controller. The continuous visualisation performed better than the baseline visualisation. Participants indicated that this was the case because having some objects flying above them map (but not too high) like the pins and the timelines enabled them to use those objects as mental anchors to navigate around, without loosing the feeling of where they were in the virtual space.

Participants also indicated that the resolution was a cause of illness for them. Especially when participants wanted to look at the city names on the map. Participants had to get very close to the map and the low resolution of both the map (even at its highest level of detail) and the head set caused some illness. Other participants indicated trouble with visually focusing on objects as a cause of illness. The duration of the experiments also led to some discomfort for two participants. For one participant, the induced illness caused by the experiment was too much. He stopped the experiment early before finishing the last of the nine tasks. This was taking into account for all effectiveness measures by taking averages, filtering out any effect of missing this data. Another participant said he felt some discomfort caused by the frame rate, which he stated was too low. It is possible that other participants had trouble with this as well, but could not pinpoint the frame rate as the cause as they were less familiar with virtual reality than this particular participant.

VI. Discussion

In this research we looked at three different visualisations of spatio-temporal lifelogging data in virtual reality and evaluated them on four points, namely enjoyment, usability, effectiveness, and nauseousness, in order to answer the research question introduced earlier: “What is the effect of including different kinds of time-related features on a map-based lifelogging visualisation in virtual reality on the enjoyment, usability, effectiveness, and nauseousness of the visualisation?” Here we will discuss the results presented in Section V, evaluating the effects of the different interfaces on the four points posed in the research question.

A. Effects on Enjoyment

The results show that the continuous visualisation is the one that was most liked by the participants. As shown in figure 10, on six of the seven categories of the Game Experience Questionnaire, the continuous visualisation outperforms the other two. We believe that there are multiple reasons for this difference between the levels of enjoyment of the different visualisations. First, the visualisation was simple to use. There was only one additional interaction in this visualisation in comparison to the baseline visualisation, namely dragging up a time line to create a continuous and chronological stream of pins. To do this, users had to aim their controller at a time line, press the trigger button, and drag the time line upwards. Because of this simplicity, participants easily understood the visualisation, which helped them feel comfortable. This is the opposite of the discrete visualisation, where most participants got a bit stressed out by the visualisation, as it was not intuitive enough for them. This had a negative effect on the GEQ-score for the discrete visualisation.

The baseline visualisation scored mostly similar to the discrete visualisation with the exception of ’Competence’. The reason why the discrete visualisation was scored similarly is likely to be that participants did not really enjoy the baseline visualisation, but did not have anything to complain about it either. However, participants did seem to like the discrete visualisation, but it was also difficult to use leading to frustration and decreasing the enjoyment of it. The difference in score on competence shows that the discrete visualisation made participants feel like they were not capable of working with the visualisation. One might suspect that this would also affect the scores on 'Negative Affect', 'Challenge' and 'Tension\Annoyance', which, at first look, does not seem to be the case. However, when looking further into the data and splitting the participants into groups based on whether or not they had an intricate understanding of the data, we see this effect for the participants who did not have an intricate understanding of the visualisation on 'Negative Affect' and 'Challenge', see fig 11. This confirms that participants who did not understand the discrete visualisation increased the scores of the GEQ, which lead to it matching the scores of the baseline visualisation. It also shows that the GEQ scores are better than the baseline visualisation if participants did understood the visualisation, showing that a visualisation would likely have outscored the baseline visualisation on the GEQ, if it was more intuitive.

The second reason why the continuous visualisation scored best on the GEQ was that, since it was easy to use, participants were also able to consistently complete all the tasks. The continuous visualisation had a completion percentage of 100%, whilst the other two visualisations scored lower, see figure 15. Whenever a participant was unable to complete a task, it was noted by the experiment leader that their attitude towards the visualisation deteriorated.

Another reason why the continuous visualisation outscored the other visualisations on the GEQ was that it did not cause as much nausea as the other visualisations (this will be elaborated in Section VI-D). Feeling nauseous is a reason to enjoy yourself less and thus a partial explanation of the difference in the GEQ scores.

It is harder for us to tell the exact effects of the improvements of the baseline visualisation over the visualisation made by Ouwehand [28] on the enjoyment of the visualisations, as we did not compare the baseline visualisation to his visualisation. However, the results still show that the addition of colours to the pins and the date time label were beneficial to the enjoyment of the visualisation. When asked to pin point the strengths of the baseline visualisation, the participants often referred to the pin colours. Furthermore, the date time label was a useful addition to the system (especially in the baseline visualisation), as participants often used that feature. Excluding that feature would have made some of the tasks a lot harder, probably leading to more frustration by the participants. Thus we think that the inclusion of the date time label improved the score of the GEQ, concluding that it had a positive effect on the enjoyment of the participants.

B. *Effects on Usability*

The continuous visualisation outperformed the other two visualisations on usability, as participants pointed out more strengths than weaknesses, as opposed to the contrasting results from the other two visualisations. In the previous section, we already touched upon the usability of the visualisations, as it also relates to the enjoyment of the visualisation. In this section we take a closer look at the usability, and pinpoint specific parts of the visualisations that lead to the continuous visualisation being the most usable of the three.

First of all, the chronological nature of the visualisation was found to be helpful to the participants as participants indicated they used this often in solving their tasks. As the strengths and weaknesses analysis showed in the results section, this visualisation performed well because it was easy to get an overview of all the data and to get temporal data of the pins, since they were placed in chronological order. When the time line was fully expanded, the order of all the pins became clear at a single glance. This chronological order also enabled users to quickly search through the data in search for a specific date, which was helpful in some tasks, according to the participants.

The visualisation also fits the type of data quite well. Lifelogging data is a continuous stream of data that is captured non stop [5]; visualising the continuous stream of data in a continuous line of pins scattered across a map and a time line fits this data type well. Also, by making the map vertically dynamic, the user could drag the pins upwards at a busy cluster of pins and separate and focus on the specific data that was of interest to the user. Furthermore, the problem of not being able to see the geographical names on the map because the pin(lines) were occluding those names was not so much of a problem in this visualisation, as participants could easily drag the pins upwards and read the locations on the map. Some participants mistook the order of the pins, where the newer pins where closer to the map than the older pins, as opposed to what the participants expected. We chose to go for this order of the pins, as we argued that, in practice, it is likely relevant to look at the newest pins and having them closer to the map when the visualisation is expanded, makes it easier for users to see the exact location of the newer pins.

The discrete visualisation showed potential to be usable. However, it was not usable to most participants in its current state, as the visualisation was found to be confusing at times. The time line and the expanding and collapsing of the layers made the data more approachable, as grouping data on a time frame helped focus participants on specific data. However, because of its implementation, where participants were unsure about what data they were looking at and were struggling with time line, the visualisation confused participants. A quote of a participant who just used the discrete visualisation shows this problem: ”It was not really clear for me at what moment in time I was. The slider (time line) was not always of the same interval, which was confusing. Also the amount of buttons on the side of the map for expanding and collapsing was too much.” The problem of the time line having different intervals, was a problem that was often observed by the experiment leader. A number of participants did not realise that the other layers who were not expanded were still visible if another layer was expanded. This lead to situations where, for example, a participant thought he was only looking at layers of a specific month, was actually also looking at other months as well. The reason that the visualisation functioned this way was to never fully remove data from the visualisation, as this could remove the overview of the data as participants might not be aware they made some data invisible. However, it is debatable if this was beneficiary to the visualisation as it appeared to confuse participants. Similar to this problem, it was possible to scroll data to a position below the map, making the data invisible for the user depending on their location in the virtual world. During the implementation we looked at options to overcome this problem, like having the data collapse at the map in a similar fashion to the continuous visualisation. Though this solution did work for the continuous visualisation, it lead to problems in the discrete version where it became unclear at what data your were looking at. Although the layers caused some confusing for some participants, other participants seemed to really understand the visualisation and being able to focus on a single layer helped them with completing the task. As a participant stated, ”Being able to see the data per month and expand on it by going to a specific day or time, it became clear what data belonged to what month and you could still really dive into the data.”

C. Effects on Effectiveness

The effectiveness of the visualisations was measured with the quantitative data, which was automatically gathered by the system during the experiments. Similar as with the usability and enjoyment, the continuous visualisation scored the best of the three on this evaluation point. We believe that the main reason for this is the usability of the visualisation. Having a more usable interface leads to being able to focus more on the task than on the interface itself. Where in the discrete visualisation participants struggled with getting the right information at the right place, participants did not struggle with this in the continuous visualisation. Additionally, as there were less interactions possible in the continuous visualisation, it was easier to get the right information in the right place. In the discrete visualisation some participants wanted to get the right hour-layer at the right place in the visualisation before they started looking at the answer to the task, while in the continuous it did not matter as much to the participants as to where the relevant pins were exactly. This might be explained by the nature of the interfaces. In the discrete visualisation you look at specific time frames and the more specific the frame, the less clutter is visible and therefore looking at the data as zoomed in as possible might be desirable to some participants. In the continuous visualisation the absolute vertical position of a pin was not relevant at all for most tasks, but the relevant position to the other pins was. Therefore, in the discrete visualisation, participants seemed to spend more time preparing the visualisation to answer the tasks than actually answering the task. In the continuous visualisation, this did not happen and this is what lead to a higher effectiveness of the visualisation.

Comparing the effectiveness of the discrete visualisation to the baseline visualisation, we see that additional features of the discrete visualisation helped the effectiveness even though the additions were not as useful as the additions of the continuous visualisation. The possibility to split the data in layers and therefore focus on specific data parts helped the discrete visualisation as it reduced the cluttered places where a large number of pins were stacked. Furthermore having some additional temporal data available to the user helped find specific dates seemed helpful, which also contributed to a better score on the effectiveness of the discrete visualisation.

D. Effects on Nauseousness

The level of nauseousness of the participant while using the visualisations differs between the visualisations. The continuous visualisation scores once again the best, with the lowest score on both the oculomotor and disorientation related sickness and consequently also on the total nauseousness. This shows that the continuous visualisation improved the level of nauseousness of the participants in comparison to the baseline. The continuous interface had some vertical elements, which was not the case for the baseline. Participants were able to use to higher placed objects as mental anchor points when moving through the space. Having an anchor point to relate your movement to is helpful with decreasing motion sickness. One might argue that because of this, the discrete visualisation would also perform well. The opposite was true however as participants found it nauseating to look up and down a lot. This caused the discrete visualisation to score even lower than the baseline. It is hard to draw any general conclusion here about the effect of temporal aspects on the nauseousness of participants as it was not the temporal aspects of the objects that made them nauseating or not, but rather the position of the elements that had an effect on this.

E. Comparing the continuous and discrete visualisation directly

This research gives an answer the question stated in Section I: ”What is the effect of different temporal visualisations based on a spatial data representation of lifelogging visualisation in virtual reality on the enjoyment, usability, effectiveness, and nauseousness of the visualisation?”. The baseline visualisation introduces some improvements over the system of Ouwehand [28]. These improvements (amongst other additional features), are also present in the two other visualisations, making them more interesting and promising for real life applications and as such, we compare the evaluations of the continuous and discrete visualisation in this paragraph. Furthermore, these two visualisations outperformed the baseline on almost all fronts, so it is more interesting to look at those two. We are not looking at the differences of the actual visualisation but rather the implications and uses of the visualisations.

The results of the experiments show that the continuous visualisation is more liked and more useful than the discrete one. The overview and chronological order of the pins is what participants seemed to liked the most about this visualisation. However, the discrete visualisation is also able to offer both of these things. Before expanding any layers, the user is presented with an overview of the data in just a couple of layers as visible in figure 8. Although the layers are vertically separated, having the data spread apart vertically should not be a problem as this is also often the case in the continuous visualisation, see figure 7. Participants do seem to lose the overview of the discrete visualisation when one or more month layers are expanded to day layers even though it is visually pleasing. One participant called this experience ”... like being in a cathedral or some kind of big building where you like to look around” (see figure 9 for this effect). So although it might make for an visually pleasing experience, it is questionable if having this much verticality in the visualisation is desirable, especially since this also caused nausea.

Another difference between both visualisations is in how they handle the chronological nature of the data. Whereas in the continuous visualisation, the chronological order of each data point becomes immediately clear when you expand the visualisation by dragging up the time line (albeit some participants mistaken the start for the end of the data set and vice versa), the discrete interface never fully organises all pins in the right order as the maximum level of detail is provided by the hour-layers. Furthermore, it takes more effort to fully expand all the layers in the discrete visualisation than it takes to expand the whole continuous visualisation.

Because of the way the two visualisations provide different kinds of overview and how they show the chronological order of the pins, we would suspect that there are two different usage types that each fit one of the visualisations. These usage types correspondent to the ones discussed in Sections IV-A.1 and IV-A.2. From the qualitative results we would expect that the continuous visualisation would work better for exploring the data on a surface level and the discrete visualisation would work better for finding specific data. However, as figure 13 shows, this is not the case. Yet, it is possible that this might be the case for the participants who got a good grasp of how the discrete visualisation worked. As stated earlier on, some participants struggled with the discrete visualisation while others used the visualisation in smart ways to complete the tasks. It might be the case that for participants with intricate understanding of the visualisation, the average completion time of tasks per usage type might be very different. However, since this research focused on acquiring and analysing qualitative data, we cannot say anything about this as we have too little data.

Another aspect that is relevant when comparing these two visualisation is the scalability, something we have not yet discussed anywhere in this research. Scalability, in this context, relates to how well the visualisations perform when the number of images that might vary between a very small and a very large number of images. In terms of computational performance, the continuous visualisation would scale the best of the two visualisations. Having more pins and dragging them up is not demanding heavy resources that a PC that can handle virtual reality does not have. For the discrete visualisation however, having a large number of pins definitely affects performance. On start-up, the system needs to loop over all pins, organise them in a chronological order (up till this point this is similar to the continuous visualisation), create month- day- and hour-layers for each month, day and hour that are present in the pins, assign each pin to the right layer, hide the layers that are not visible and position the layers that are visible correctly. With the current amount of 56,450 images, the discrete visualisation takes about 32 seconds to set everything up. It is possible that this might be optimised further, however since it calls some core Unity features like create objects and setting them inactive, the performance that can be gained might be marginal. Luckily, this all takes place at the start-up of the system (as doing this in real time would cause too much loading times) and is therefore not a problem while the system runs.

The performance of the system however is not the only aspect where scalability comes into play as the amount of data also changes the visualisations themselves. Currently we have about three months of data, but lifeloggers might want to visualise all the data they gathered over years of lifelogging in our visualisation. For the continuous visualisation, this could mean two things. Either the visualisation would need to get decrease the vertical distance between pins so that more pins can fit in the same vertical space or the visualisation would expand in vertical size, something that might lead to a lot of scrolling for the user when it wants to find specific data. Both situations are not ideal and call for additional features like scaling the time line to deal with these problems. For the discrete visualisation, not much changes as having large amounts data in the discrete layers does not necessarily lead to situations that are very different for the current situations. It might be helpful to add layers for years and possibly decades instead of only having layers for months, days and hours, but the interaction with the visualisation itself does not need to be changed to suit large data sets. Actually implementing these extra layers is very easy as the system is set up completely recursively and adding extra layers is as easy as creating a new time unit for 'years' and possibly 'decades' and change the starting conditions of the recursive loop to start at these newly created time units.

Both the continuous and discrete visualisations have their own advantages and disadvantages. The continuous visualisation performs better for the tasks in this research, is more liked by the participants accordingly to the GEQ, and causes less nauseousness. However, to properly display a larger amount of data, the visualisation needs some additional features implemented to still perform well. The discrete visualisation can performs decent, however, this depends on the user as some users understand the visualisation very well while others do not. This visualisation is easily scalable to larger data sets which translates well to other lifelogging data sets.

In the end, participants can put more focus on the data in the continuous visualisation as it is intuitive to everyone, leading to a low average completion time and high likeability scores. The discrete visualisation demands users to pay more attention to the interface itself rather than the data, although when users understand the discrete visualisation well, this problem disappear.

VII. Future work

This research focused on the qualitative analysis of three different spatio-temporal visualisations of lifelogging images and their effects on enjoyment, usability, effectiveness, and nauseousness. Future research could extend on our findings. We will explore some of these possibilities here and explain why they might be interesting.

As shown in previous sections, both the continuous and discrete visualisations have their own (dis)advantages. It would be interesting to take the strengths of both visualisations and combine them into a new visualisation and see how it would perform. A way of doing this is, for example, by starting out the new visualisation with the discrete interface and when a user decides to expand a month, that month is turned into a continuous like visualisation. This way, users get the advantage of the overview of seeing the data in month layers (or year layers, depending on the data set) from the discrete visualisation and the advantage of having a continuous stream of data points from the continuous visualisation. Some other changes to the visualisations are required when merging these interfaces, like having buttons to switch from the discrete visualisation to the continuous visualisation and back and making sure this transition goes well and runs smooth. This interface might get rather complicated to use by combining the two types of interfaces in visualisation, however, it also enables researchers to get feedback on both systems while presenting the participants only one visualisation. Furthermore, the problem of scaling disappears for the continuous visualisation as it now never show more than a months worth of data. The scalability of the discrete visualisation is exploited as dealing with data of multiple years is still not a problem in such a hybrid visualisation.

Another direction to take this research in, is giving users some kind of summary of (parts of) the data set in the environments created for this research. As discussed in Section II-B.2, there exist systems that summarise geographical tagged images like in Gemmell et al. [12], Thudt et al. [33, 34]. These systems present images on some chronological time line to show users a summary of (parts of) the dataset. In our continuous visualisation, we also present a time line of all the images to the user when the user fully expands the visualisation as shown in 7, although this time line is in 3D. A system that automatically flies users over this time line, guiding them alongside all images might be a way to give users a summary of parts of the data. As a participant of this research pointed out "It would be nice to have sometime like a 'play'-function where some photos are presented and you can go back and forth between those images". A way to implement this is to create an invisible track which is close to all the images and that acts like rails among which the user can move back and forth. This way, users can fly over this predetermined route and get a good overview of the data. Users might select a part of the data they want to get an overview of and only that part is shown during the fly-by and as such can be viewed in more detail. Some participants of this research did something similar to the behaviour that would arise when creating the system described in this paragraph as they manually followed a line of pins by flying over it. The participants that did so indicated that they liked doing that as ”flying over the pins was fun” and ”the stars on the side of the visual and going up or down as you follow the route made it feel like you were time travelling.”

Another possible direction for future research that directly follows from this research is to implement some of the suggestions given by participants and do quantitative research to see how those improvements might change the experience of the users. These changes might be minor, like changing the order around of the pins in the continuous visualisation and AB-test what order works the best or could be bigger changes like overhauling the navigation of the camera and see how this has an effect on the nauseousness of participants. An example of an overhaul for the navigation might be to change the direction of the movement to the direction of the controller as opposed to the camera as in current implementation. Also multiple participants reported that they would like to see a button to return to the start position as this gave them a good overview of the data and that they would like to see an increase of the movement speed over time when the user is moving continuously as currently the navigation felt too slow for longer distances. Other suggestions that might improve the visualisations are listed in appendix I.

VIII. Conclusion

In this research we looked at ways of visualising spatio-temporal lifelogging data in a virtual reality environment. We did this by creating three different visualisations and we looked at what aspects of these visualisations worked well and what aspects have room for improvement. We tasked ourselves with two research goals as stated in Section I. In this section we discuss if we reached those goals and what we can learn from our experience.

A. Including temporal data in a map-based visualisation of lifelogging images

The first research goals was to use temporal data to improve a map-based visualisation of lifelogging images to help users explore the data. To reach this goal, we created the three visualisations described in this paper and let participants comment on the features and the strengths and weaknesses of the visualisation. We found that colouring the data based on the time stamps of the data, was a feature that helped participants quite a lot with exploring the data. In the previous implementation [28], all data points had the same colour which made it impossible to see the order of the data points in one glance, as shown in 1. With the addition of colours to the pins, it became immediately clear to participants what pins followed each other. A clear example of this can be seen in figure 5 where there are distinct differences between, for example, the different trips from Dublin to Oslo. In one glance it is clear that there are four different flights between these two locations. Another addition that helped participants explore the data was the date time label that was present near the controller when the participant aimed the controller at a pin. This label showed the date and time of the pin that was aimed at. This feature made it possible to get an idea of the absolute time of each pin, without needing to open the pin and its images. The combination of the relative order of the pins, as provide by the colouring of the pins, and the absolute time stamp of the date time label made it possible for participants to quickly find specific data points based on their time stamp.

The continuous and discrete visualisations added additional time related features to the baseline visualisation which helped participants exploring the data and find specific data easier. First, we will discuss the discrete visualisation. The discrete visualisation included the addition of time layers in which each pin that was related to that specific time layer was placed on as described in Section III-D. This visualisation outperformed the baseline visualisation on usability and effectiveness, but also caused more nauseousness than the baseline and scored similar to the baseline on its enjoyment. This visualisation has a lot of potential, but need some improvements to actually work for all types of users. In this research we found that there were two types of users for this visualisation. First, there were people who used the system quite logically and intuitive and therefore performed well. On the other hand, there were people who did not get this visualisation at all and got lost in the interface when using the visualisation. Even though there was this difference in user types, the average completion time of the this interface was lower than that of the baseline visualisation, indicating that the discrete visualisation was still an improvement over the baseline. The different layers helped the participants to find specific data. This was done by having discreticed time units making the time of each layer and its pins clear. Though the chronological order of the pins is more clear in the continuous visualisation, the discrete visualisation still benefited from ordering the pins in a semi-chronological way as each layer was chronologically ordered but the pins inside each layer were not.

The continuous visualisation had the addition of a scrollable time line which lifted the pins upwards from the map in the order they were taken to create a continuous and chronological stream of pins as described in Section III-C. Participants liked the addition of the time line as they rated this visualisation higher on usability and enjoyment over the baseline and continuous visualisations. Furthermore, this visualisation caused less nausea and participants were able to find the right answers to the tasks quicker in this visualisation than in the other two visualisation. This indicates that the continuous visualisation is the overall better visualisation of the three. The main reason for this is that it was easy for participants to see the chronological order of the pins, as seeing the continuous stream of pins made the temporal relation of each pin to every other pin clear. Comparing the time stamp of each pin is as simple as comparing height in the visualisation. If a participant was looking for a specific point in time, it could narrow it down pretty quickly where this point was. Furthermore the addition of the time line on the side of the visualisation made it possible to see absolute time data as well, making it even more easy to find specific data. Also in comparison to the discrete visualisation, all these benefits were gained without adding a lot of visuals clutter to the interface, making it a cleaner visualisation which helped participants to focus on the data and not on the visualisation itself.

In summary, the three visualisations show different ways of using time to help users explore lifelogging data in a map based virtual reality environment. There are additions that are more focused on the relative relationship between the time stamps of the images like the colouring and the continuous visualisation and additions that are more focused on the absolute time stamp of the images like the date time label and the discrete visualisation. Both have their own strengths and weaknesses. Focusing on the relative time is interesting when comparing different pins, but is less useful when looking for specific dates. The opposite is true for the absolute time additions where finding specific dates is more easy but getting an idea of the order of the pins is not as trivial. From this research it looks like having relative cues about the time of an image is more important than having absolute cues as participants seemed to perform the best in the continuous visualisation and indicated they liked that visualisation more than the others.

B. Declustering clusters of lifelogging images in a map-based visualisation by using temporal data

The second research goal in this research was to use temporal data to spread out clusters of lifelogging images in a map based visualisation to make the data more accessible and help user find specific data. The continuous and discrete visualisations are our methods to reach this goal.

To decluster pins that are close together on a plane (the map), we decided to use the third dimension as that dimension was unused in the previous visualisation and offered opportunities to creatively decluster pins. The continuous and discrete visualisation both use the third dimension to make the data more insightful. The discrete visualisation does this by splitting the data into discrete layers of time. If a user splits the data by, for example, expanding a month layer to multiple day layers, the data automatically declusters. Clusters of pins often form when there is a lot of data close to each other, but not close enough to be grouped into one pin. This means that the lifelogger needs to be at multiple locations that are somewhat close to each other to form these clusters, which often does not happen in a single day but over multiple days. Therefore splitting the data into day-layers automatically declusters the clusters of pins.

The continuous visualisation uses a similar method to decluster pins, however, without the use of discrete layers, but by creating the continuous stream of pins. The stream of pins helps to decluster the pins as it places all pins on a different vertical position (if the visualisation is fully expanded). It is hard to say if this method of declustering pins is better than the method of the discrete visualisation as we did not explicitly test to see what is the better method for declustering pins. We can conclude that participants seem to prefer the continuous visualisation over the discrete visualisation and the declustering method might have contributed to this.

Note, that to reach this research goal, we decided to create visualisations in which all data was still present. Methods including completely filtering out parts of data were also considered for this research. However, we decided to focus on other ways of declustering the clusters of pins, as filtering out data inherently makes some parts of the data invisible which might confuse the user and makes getting an overview of the complete dataset harder. Furthermore, developing an interface to filter out data asks for menus and a lot of button to select specific data. This however, decreases the immersion of the user as a menu takes the attention of the user and directs it away from the data and might fit better in a desktop application than in a virtual reality application.

C. Final Conclusion

In this research, we created two novel visualisation of lifelogging data in a spatio-temporal virtual reality environment and looked at the effects these visualisations had on the enjoyment, usability, effectiveness, and nauseousness of the user. We found that a continuous representation of time in a map-based visualisation was perceived as most enjoyable and usable. The continuous visualisation was also the most effective visualisation and caused the less nauseousness amongst the users. The new visualisations also helped to decluster clusters of pins on the map. Furthermore, we introduced some features that give users extra temporal information (the pin colours and the date time label) and these features helped the user find specific data in the dataset.

IX. Acknowledgements

We would like to thank Kevin Ouwehand and Marijn Mengerink for creating the initial implementation of the system that we used to modify and extend on for our own research. Kevin Ouwehand also supported us during the initial phase of this research to get everything up and running, which was of great assistance to this research. Furthermore we like to thank friends and family for their continued support whom without this paper would have never been completed. Special thanks goes to Jarno Verhagen for helping conceptualise some of the programming of the visualisations and Ingeborg Lip for helping us improve the text of this thesis. At last, we like to thank all the participants for participating in the experiments.

References

[1] Anita L Allen. Dredging up the past: Lifelogging, memory, and surveillance. *The University of Chicago Law Review*, 75(1):47–74, 2008.

[2] Morgan G Ames and Lilia Manguy. Photoarcs: ludic tools for sharing photographs. In *Proceedings of the 14th ACM international conference on Multimedia*, pages 615–618. ACM, 2006.

[3] Joris Bolwork. Viewpoint-controlled navigation in virtual reality for exploratory image browsing. Master's thesis, Utrecht University, 2017.

[4] David J Crandall, Lars Backstrom, Daniel Huttenlocher, and Jon Kleinberg. Mapping the world's photos. In *Proceedings of the 18th international conference on World wide web*, pages 761–770. ACM, 2009.

[5] Martin Dodge and Rob Kitchin. 'outlines of a world coming into existence': pervasive computing and the ethics of forgetting. *Environment and planning B: planning and design*, 34(3):431–445, 2007.

[6] Ryan Eccles, Thomas Kapler, Robert Harper, and William Wright. Stories in geotime. *Information Visualization*, 7(1):3–17, 2008.

[7] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96 - 34, pages 226–231, 1996.

[8] Rose Eveleth. How many photographs of you are out there in the world?, Nov 2015. URL `https://www.theatlantic.com/technology/archive/2015/11/how-many-photographs-of-you-are-out-there-in-the-world/413389/`.

[9] Md Sadek Ferdous, Soumyadeb Chowdhury, and Joemon M Jose. Analysing privacy in visual lifelogging. *Pervasive and Mobile Computing*, 40:430–449, 2017.

[10] Flickr, 2018. URL `https://www.flickr.com/map`.

[11] Jim Gemmell, Aleks Aris, and Roger Lueder. Telling stories with mylifebits. In *2005 IEEE International Conference on Multimedia and Expo*, pages 1536–1539. IEEE, 2005.

[12] Jim Gemmell, Gordon Bell, and Roger Lueder. Mylifebits: a personal database for everything. *Communications of the ACM*, 49(1):88–95, 2006.

[13] Cathal Gurrin, Alan F Smeaton, Aiden R Doherty, et al. Lifelogging: Personal big data. *Foundations and Trends® in Information Retrieval*, 8(1):1–125, 2014.

[14] Cathal Gurrin, Hideo Joho, Frank Hopfgartner, Liting Zhou, and Rami Albatal. Ntcir lifelog: The first test collection for lifelog research. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 705–708. ACM, 2016.

[15] Cathal Gurrin, Hideo Joho, Frank Hopfgartner, Liting Zhou, Rashmi Gupta, Rami Albatal, Dang Nguyen, and Duc Tien. Overview of ntcir-13 lifelog-2 task. NTCIR, 2017.

[16] Wijnand IJsselsteijn, Yvonne De Kort, Karolien Poels, Audrius Jurgelionis, and Francesco Bellotti. Characterising and measuring user experiences in digital games. In *International conference on advances in computer entertainment technology*, volume 2, page 27, 2007.

[17] Google Images, 2018. URL `https://images.google.com/`.

[18] Tim Jacquemard, Peter Novitzky, Fiachra O'Brolcháin, Alan F Smeaton, and Bert Gordijn. Challenges and opportunities of lifelog technologies: A literature review and critical analysis. *Science and engineering ethics*, 20(2):379–409, 2014.

[19] Alexandar Jaffe, Mor Naaman, Tamir Tassa, and Marc Davis. Generating summaries and visualization for large collections of geo-referenced photographs. In *Proceedings of the 8th ACM international workshop on Multimedia information retrieval*, pages 89–98. ACM, 2006.

[20] Vaiva Kalnikaite, Abigail Sellen, Steve Whittaker, and David Kirk. Now let me see where i was: understanding how lifelogs mediate memory. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2045–2054. ACM, 2010.

[21] Robert S Kennedy, Norman E Lane, Kevin S Berbaum, and Michael G Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993.

[22] Hyun K Kim, Jaehyun Park, Yeongcheol Choi, and Mungyeong Choe. Virtual reality sickness questionnaire (vrsq): Motion sickness measurement index in a virtual reality environment. *Applied ergonomics*, 69:66–73, 2018.

[23] Eugenia M Kolasinski. Simulator sickness in virtual environments. Technical report, Army research Inst for the behavioral and social sciences Alexandria VA, 1995.

[24] James R Lewis. Ibm computer usability satisfaction questionnaires: psychometric evaluation and instructions for use. *International Journal of Human-Computer Interaction*, 7 (1):57–78, 1995.

[25] LSC. Lsc website 2019. lifelog search challenge 2019. `http://lsc.dcu.ie/`, 2019. Accessed: 2019-06-28.

[26] Andrés Lucero, Marion Boberg, Sanna Malinen, and Kaisa Väänänen-Vainio-Mattila. Image space: an empirical study of geotagged mobile media content capture and sharing. In *Proceeding of the 16th International Academic MindTrek Conference*, pages 187–194. ACM, 2012.

[27] Marijn Mengerink. Geospatial image browsing in virtual reality. Master's thesis, Utrecht University, 2019.

[28] Kevin Ouwehand. Geospatial access to lifelogging images in vr. Master's thesis, Utrecht University, 2019.

[29] Antoine Pigeau and Marc Gelgon. Organizing a personal image collection with statistical model-based icl clustering on spatio-temporal camera phone meta-data. *Journal of Visual Communication and Image Representation*, 15(3):425–445, 2004.

[30] Pinterest, 2018. URL `https://pinterest.com/`.

[31] William Plant and Gerald Schaefer. Visualisation and browsing of image databases. In *Multimedia Analysis, Processing and Communications*, pages 3–57. Springer, 2011.

[32] Abigail J Sellen and Steve Whittaker. Beyond total capture: a constructive critique of lifelogging. *Communications of the ACM*, 53(5):70–77, 2010.

[33] Alice Thudt, Dominikus Baur, and Sheelagh Carpendale. Visits: A spatiotemporal visualization of location histories. In *Proceedings of the eurographics conference on visualization*, 2013.

[34] Alice Thudt, Dominikus Baur, Samuel Huron, and Sheelagh Carpendale. Visual mementos: Reflecting memories with personal data. *IEEE transactions on visualization and computer graphics*, 22(1):369–378, 2016.

[35] Carlo Torniai, Steve Battle, and Steve Cayzer. Sharing, discovering and browsing geotagged pictures on the world wide web. In *The Geospatial Web*, pages 159–170. Springer, 2009.

[36] Kentaro Toyama, Ron Logan, and Asta Roseway. Geographic location tags on digital images. In *Proceedings of the eleventh ACM international conference on Multimedia*, pages 156–166. ACM, 2003.

[37] Yalong Yang, Tim Dwyer, Bernhard Jenny, Kim Marriott, Maxime Cordeil, and Haohui Chen. Origin-destination flow maps in immersive environments. *IEEE transactions on visualization and computer graphics*, 2018.

[38] Yalong Yang, Bernhard Jenny, Tim Dwyer, Kim Marriott, Haohui Chen, and Maxime Cordeil. Maps and globes in virtual reality. In *Computer Graphics Forum*, volume 37 - 3, pages 427–438. Wiley Online Library, 2018.

[39] Yang Yang and Cathal Gurrin. Personal lifelog visualization. In *Proceedings of the 4th International SenseCam & Pervasive Imaging Conference*, pages 82–83. ACM, 2013.

[40] Yang Yang, Hyowon Lee, and Cathal Gurrin. Visualizing lifelog data for different interaction platforms. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pages 1785–1790. ACM, 2013.

APPENDIX I
SUGGESTIONS FOR IMPROVEMENT OF THE VISUALISATIONS

Improvements as suggested by the participants. The irrelevant suggestions are removed from these list. First the baseline visualisation as these improvements can also be applied to the other two visualisations:

- In the current implementation you cannot see where you touch the touch pad on the controller, while some participants indicated that they would like to see this as they felt like this would help move more correctly.
- Show some kind of legend where people can look up what the colours represent.
- Create a navigation method for quicker navigation.
- Filter out some images based on a selection by the user to let them focus on specific parts of the data.
- Show additional geographical information like a compass and have a clearer/more detailed map.
- Decluster the pins by grouping them together.
- Let pins shrink in size when you zoom in so you can identify individual pins better when you zoom in on a cluster of pins.
- Show the colours of the pins, even when the images of the pins are loaded in by having a coloured border around the images/pins as to keep the temporal data present. Also let pins change back to the colours when zooming out.
- Use the colours differently. Instead of reusing colours when all of the selected colours are used, use more colours so that each colour really represents a single day.
- Add music for more immersion

Improvements for the continuous visualisation:

- Add a play function to fly over all images.
- Add a function to let users choose the accuracy of the scroller to make more precise scrolling possible.
- Make it possible to change the vertical distance between pins to separate them further if wanted.
- Switch around the order of the pins on the time line such that new pins are at the top and old pins at the bottom.
- Let the y-axis scroll up a little bit further that even the bottom pins can be lifted from the map.
- Create a plane that users can drag up to see what pins came before and what pins came after a certain date.

Improvements for the discrete visualisation:

- Do something to make sure participants see where the data is. Some participants suggested things to remove layers where there was no data, but every layer had data present, otherwise the layer would not have been created.
- Use the buttons on the controller to scroll through the layers.
- Have bigger buttons on the side to more easily select them and have them a little bit more spaced apart from each other.
- Use a filter to only show certain layers or locations in the visualisation.
- Let the buttons and the time line light up when you click/grab them as to give some feedback on this.
- Make it possible to flick the time line instead of only being able to grab and drag it.