



Universiteit Utrecht

DEPARTMENT OF INFORMATION COMPUTING SCIENCES

GAME AND MEDIA TECHNOLOGY - MASTER THESIS

Exploring the Suitability of Geospatial Visualization and Temporal Filtering Techniques for Lifelog Data

Author:
Maaïke Konijnendijk

Supervisor:
Dr. Wolfgang Hürst

Second Examiner:
Dr. Michael Behrisch

ICA-6819710

May 18, 2022

Abstract

Lifelogging is the collection of data through sensors carried by a user which record the user's life experiences. Examples of sensors that are used to collect lifelog data are (wearable) cameras, GPS, and sensors collecting biometric data. This activity results in lifelogs, that is, multimodal databases containing large amounts of data that is challenging to navigate. Related research mostly focuses on lifelog retrieval; building a system which can retrieve specific data of the lifelog. These systems are made to enable the user to retrieve certain memories. Yet, they lack in supporting exploration of the data for leisure browsing or situations with a less clear information need, for example, when people only vaguely remember a certain situation. A good visualization of a lifelog's content and the ability to filter it as needed could encourage and support such exploration, but related studies are sparse. Therefore, in this thesis we aim to determine if geospatial visualization and temporal filtering techniques can be used to make lifelogs accessible and easy to search and explore. To achieve this, we implemented and evaluated a system, the Lifelog Browser, which visualizes the LSC'20 test collection, a benchmark lifelog dataset, and makes it accessible. The system features three different geospatial visualizations (Dot map, Cluster map and Heat map) and three different temporal filtering techniques (color-coded Week matrix, color-coded Month matrix and Custom drop-down), which were evaluated for suitability for different search goals in the context of lifelog data. This was done through a user study which included both a lifelog exploration and retrieval challenge, inspired by the Lifelog Search Challenge. Our results show that this type of system is suitable for lifelog access, search and exploration and that out of the chosen approaches, the Cluster geospatial visualization and Custom drop-down temporal filtering technique are considered most useful for lifelog exploration and search.

Acknowledgements

I would like to thank my supervisor dr. Wolfgang Hürst for his guidance and input. His extensive feedback has been essential for this project and with his support, I have achieved more than I had imagined I would before starting. I would also like to thank Maurice Heijligers, for providing his knowledge during a number of technical challenges I faced during this project. Furthermore, I would like to thank all participants for lending me time out of their busy days to partake in my experiment. Finally, I would like to thank my friends, family and partner who were always there to provide encouragement and support when I needed it.

Contents

1	Introduction	3
2	Related Work	5
2.1	Lifelogging Background	5
2.1.1	History	5
2.1.2	Sensors	6
2.1.3	Benefits	6
2.2	Lifelogging Test Collections	7
2.3	Lifelog Systems	8
2.3.1	Geospatial visualization	8
2.3.2	Temporal filtering and visualization	11
2.4	Conclusions	12
3	Lifelog System Design	15
3.1	Motivation	15
3.1.1	Search goals	15
3.1.2	Visualization selection	16
3.2	Design	18
3.2.1	Features and Inspiration	18
3.2.2	Mock-ups and final design	19
3.3	Implementation	24
3.3.1	Database	25
3.3.2	Back-end	26
3.3.3	Front-end	28
3.4	Research questions	30
4	User study	32
4.1	Procedure	32
4.1.1	Introduction	33
4.1.2	Search	34
4.1.3	Evaluation	37
4.2	Materials	37
4.3	Participants	38
5	Results and Discussion	39
5.1	Results	39
5.1.1	Quantitative results	39
5.1.2	Qualitative results	48
5.2	Discussion	51
5.2.1	Usability	51

5.2.2	Suitability per Search goal	52
5.2.3	Search strategies	54
5.2.4	Geospatial visualizations	55
5.2.5	Temporal filtering	56
5.2.6	Overall suitability	57
5.3	Limitations	58
5.4	Future work	58
5.4.1	System extensions	58
5.4.2	Directions for further research	59
6	Conclusion	61
A	Experiment Questionnaire	66
A.1	Consent Form and Explanation	66
A.2	Demographics	66
B	Additional Results	68
B.1	Qualitative Comments	68

Chapter 1

Introduction

Lifelogging is the collection of data through sensors carried by a person which record that person's life experiences. A person who engages in this activity is commonly referred to as a lifelogger. A lifelogger may carry a variety of sensors on their body, such as wearable cameras which take approximately two or three pictures per minute, GPS systems and other sensors collecting biometric data. The collection of data that these sensors generate is called a lifelog, a large multimodal database of sensor output which describes a lifelogger's life experiences. Such a database may help a lifelogger recollect, review and reflect on their life's events.

Current challenge

The problem with such lifelog databases is that they are vast and difficult to navigate manually, both for exploration and search purposes. For this reason, a growing body of research is focused on creating systems for lifelog access. This rise in popularity can be seen in the growing number of participants in challenges like the Lifelog Search Challenge (LSC) [14], a challenge which aims to compare interactive retrieval systems for photo-based lifelog data. These retrieval systems allow users to retrieve specific moments from a lifelog. However, considerably less research has been done into systems with a focus on lifelog exploration, where lifelog data is visualized so that users may leisurely browse it.

The few existing lifelog exploration systems often use geospatial visualization and temporal filtering techniques, such as a point map to visualize where lifelog photos were taken and a calendar interface to allow users to filter the lifelog. Since location and time are important features of lifelog data, using these techniques would appear to be a suitable choice for a lifelog exploration system. However, very few lifelog exploration works provide reasoning or evaluation for the techniques which they have implemented. As such, it is as of yet unclear to what extent geospatial visualization and temporal filtering techniques are suitable for lifelog exploration.

Thesis objective

Currently, there are few formal guidelines available for researchers who wish to create their own lifelog exploration system. To fill the gap in current knowledge, this thesis aims verify whether geospatial visualization and temporal filtering techniques are suitable for displaying lifelog data in such a way that it can be accessed, searched and explored. Since this is a broad aim and there are many techniques which have the potential to be suitable for lifelog exploration, we will evaluate only a selection of geospatial visualization and temporal filtering techniques to determine which of these techniques are most useful for different search goals in the context of lifelog data.

Thesis overview

We start this thesis with an overview of related work, in Chapter 2. In Section 2.1, we provide more context for lifelogging, specifically on its benefits and how geospatial visualization and temporal filtering techniques can support these benefits. In Section 2.2, we provide an overview of available lifelogging test collections and specify our reasoning for using the LSC'20 test collection in our evaluation. In Section 2.3, we present an overview of lifelog systems which employ geospatial visualization and temporal filtering techniques, with a focus on systems which are aimed towards lifelog exploration. In Chapter 3, we present and motivate the design (Section 3.2) and implementation (Section 3.3) of our lifelog exploration system, which will incorporate three different geospatial visualization and temporal filtering techniques each. Our choices for these techniques are motivated by the work we reviewed in Chapter 2. We then introduce our research questions in Section 3.4. Our user study is described next, in Chapter 4. The results from our user study are laid out in Chapter 5 (Section 5.1), where we will also discuss them in relation to the research questions we presented in Chapter 3 and to our general research aim (Section 5.2). Finally, we present our conclusions in Chapter 6.

Contributions

The contributions of this thesis are as follows:

- Our lifelog exploration system, the Lifelog Browser, which is one of the few exploratory lifelog systems with extensive filtering functionalities, which also makes it suitable for lifelog retrieval. It is designed to visualize large amounts of lifelog data and make it accessible and explorable.
- We provide an overview of which geospatial visualization and temporal filtering techniques are most commonly used in lifelog exploration systems. Future work on lifelog exploration systems can use this overview to determine new research directions.
- We provide guidelines for which geospatial visualizations and temporal filtering techniques are most useful for lifelog access, search and exploration.
- Finally, we provide recommendations for which features other than geospatial visualizations and temporal filters should be implemented into a lifelog exploration system with a map interface as its main access point for lifelog photos.

Chapter 2

Related Work

This chapter provides a short background of lifelogging, in which we introduce the concept of lifelogging in more detail and focus on the challenges and benefits of lifelogging. We illustrate how we believe geospatial visualization and temporal filtering techniques can help lifeloggers to profit from the benefits of lifelogging. Next, we discuss lifelogging test collections of visual lifelogging that are currently available, to substantiate our reasoning for visualizing the LSC'20 test collection with our lifelog system. Then we provide an overview of relevant lifelog systems (systems to access lifelogs), specifically systems that employ some type of geospatial visualization or temporal filtering technique, such as a map or calendar visualization. Finally, we draw conclusions from the work that we reviewed, which will serve as inspiration for our study.

2.1 Lifelogging Background

Lifelogging is formally defined by Dodge and Kitchin 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"* [8]. More informally, a person who lifelogs, also referred to as a *lifelogger*, wishes to capture their life's experiences through sensors which they wear on their body. The sensors record details of the lifelogger's day passively, without any interference from the lifelogger themselves. A lifelog is the digital record of these experiences, which has been processed and made accessible through an interface, which we will refer to as a *lifelog system*.

2.1.1 History

The concept of lifelogging was first described, already before any wearable sensors existed, as a system which maintains a collection of all documents that are relevant a person. In 1945, Vannevar Bush introduced his idea of *Memex* [3], which he described as a type of computer device which serves as *"an enlarged intimate supplement to [the user's] memory"*, storing all their books, records, communications and photographs to be presented to them on demand. Of course, this idea could not come into fruition until technology could accommodate the storage of large amounts of data. The first program which brought the idea of *Memex* to life was MyLifeBits [11], created in 2002, which was a system in which all types of personal documents and photographs could be stored and queried.

From then on, lifelogging grew in popularity as both data storage and wearable sensor technologies became more widespread and affordable [17]. Another reason for this increase in popularity that lifelogging experienced is the quantified self movement, where people collect data about themselves and monitor this information as a form of self-improvement or self-reflection [28]. The main

difference between the quantified self movement and lifelogging is that the goal of lifelogging is to record the totality of the user’s life experience rather than a focused aspect of it.

2.1.2 Sensors

Nowadays, a number of different sensors can be used to record life experiences. Many of these sensors, like GPS and accelerometers, are already on the mobile phones that many carry on a day-to-day basis. Smart watches can also record lifelog data, such as steps and heart rate. A less well-known type of lifelogging device is a wearable photo camera, which takes a photo two to three times per minute. These photos describe a lifelogger’s day from their viewpoint. This type of lifelogging is called visual lifelogging and has received a considerable amount of research attention [2]. One of the first visual lifelogging devices was SenseCam [18], which is a photo camera that is worn on a string around the lifelogger’s neck, from where it takes a picture every thirty seconds or when it is triggered by certain changes in the environment, such as a change in light level. The benefit of a device like SenseCam is that it allows users to acquire visual data over longer periods of time, which is more complicated with wearable video cameras due to large storage requirements.

Visual lifelogging data combined with non-visual sensor data provides a detailed record of a person’s life. However, even a single day of (visual) lifelogging can yield up to 2,500 images and a large amount of other raw sensor data. This data must be accessible by the lifelogger to be of value to them, as it is nearly impossible for a lifelogger to find a certain item in their lifelog or explore their raw lifelog data without any kind of interface or query system to assist them, because even a lifelog spanning only a few weeks consists of a vast amount of data. Therefore, research efforts are focused on building systems which allow the lifelogger to more easily access their data, to explore it or retrieve specific items. These systems will be discussed in Section 2.3.

2.1.3 Benefits

For many, especially those participating in the quantified self movement, there are clear benefits to recording certain aspects of their life through sensors. As mentioned, many people already engage in activities such as recording their location or heart rate throughout the day, using devices like smart phones or smart watches. However, only a small number of people also capture a visual lifelog, even though this can also have many benefits. For example, in the medical field, visual lifelogging is used for health and wellness purposes, such as monitoring dietary intake [29] or supporting memory in people with memory impairment [25].

However, for people without medical indications visual lifelogging can also have benefits. Czerwinski et al. [5] suggest that even people without memory impairment can benefit from the memory-supporting features of a lifelog, for example for finding lost objects or reviewing past events and reflecting on them. They also point out the value of a lifelog for sharing personal experiences with others. Similarly, Sellen and Whittaker [31] summarize the benefits of lifelogging in five Rs: recollecting, reminiscing, retrieving, reflecting and remembering intentions. We will briefly touch upon the benefits of recollecting, reminiscing and reflecting, as we believe that these can be attained by providing the user access to their lifelog through geospatial visualization and temporal filtering techniques.

Recollecting is similar to the memory supporting benefits which were suggested by Czerwinski et al., and includes objectives like finding things or remembering specific occurrences. Reminiscing is described as a special type of recollection, where lifeloggers can use their lifelog to re-live experiences for emotional or sentimental reasons. Reflecting includes activities such as examining patterns of past experiences, which enables the lifelogger to learn about themselves. For each of these benefits,

we believe that geospatial visualization and temporal filtering techniques can be helpful, because we believe they may:

- Simplify the querying process, which is useful for recollecting;
- Display the data in a visually appealing and clear way, which is useful for reminiscing;
- Provide an overview of the lifelog data, which simplifies pattern detection and is therefore useful for reflecting.

Additionally, Sellen and Whittaker suggest that *"lifelogs might support a more abstract representation of personal data to facilitate reflection on, and reviewing of, past experience"*. Again, we believe that this abstract representation can be achieved through geospatial visualization techniques specifically, as they provide a lifelogger with an overview of their lifelog, which facilitates said reflection.

In conclusion, lifelogging has a number of benefits even for those without a medical indication, and we hypothesize that a number of these benefits can be more easily attained if lifelog systems employ geospatial visualization and temporal filtering techniques.

2.2 Lifelogging Test Collections

For a long time, there were no standardized test collections for (visual) lifelogs like there are for other multimedia retrieval tasks. This could be for various reasons. An important one is privacy concerns [12], which present themselves in two ways. Lifelogs can contain very personal information about the lifelogger, such as credit card or health details. The lifelog can also contain the faces of possible bystanders, which can be an issue as the lifelogger cannot realistically ask consent from everyone who appears in their lifelog. A test collection must therefore ensure that all personal details of the lifelogger are removed from the lifelog, and that the faces of bystanders who appear in the lifelog are made unrecognizable. This is a huge challenge, as a lifelog of a single day can consist of up to 2,500 pictures. However, having a lifelog test collection opens up a large number of research opportunities, ranging from retrieval to insight generation. In this section, we will discuss the current test collections of visual lifelogging that are available.

The first publicly released test collection for visual lifelog data was the test collection for the NTCIR-12 Lifelog task [13], which was released in 2016. It featured lifelog data collected by three different lifeloggers over the course of a month, recorded by the OMG Autographer wearable camera and a smartphone. The camera gathered about two images per minute, while the smartphone recorded location and physical activities. The data was preprocessed, ensuring temporal alignment of the camera and smartphone data, removal of any data that the lifeloggers were not comfortable with sharing, and anonymisation of the dataset. The metadata of each image contained a semantic location and the physical activity the lifelogger was engaged in at that moment. The output of a visual concept detector was also included in the test collection. Additionally, 48 search topics were included, which could be used to test the retrieval capabilities of a lifelog system.

After the release of the NTCIR-12 Lifelog test collection, a number of other test collections were created. For each of the subsequent NTCIR Lifelog Challenges a new test collection was published, each following the structure of the NTCIR-12 Lifelog test collection. These test collections were used in more lifelog challenges than just NTCIR, such as ImageCLEFlifelog [6]. In 2018, a number of the NTCIR Lifelog test collections were combined into a new test collection, the LSC test collection [16]. This test collection was released for the Lifelog Search Challenge, a new type of lifelog system competition. The main difference between the NTCIR test collections and the LSC test collection is

that the LSC test collection features richer metadata, biometric data, location logs and additional informational data, such as information consumed and created on computer devices. The collection also includes a new type of search topic, which is revealed in parts during the querying process. Just like the NTCIR collections, the LSC test collection is expanded each year with new data.

For this thesis, we have decided to use the LSC’20 test collection [15]. The test collection contains four months of data captured by a single lifelogger, spread over the years 2015, 2016 and 2018. We feel that data captured over a longer time period by a single lifelogger is more representative of lifelog data than data captured by multiple lifeloggers. We also believe that the additional data, such as biometric data, could be an interesting addition to a lifelog exploration system, as this kind of data can be visualized in ways that allow a lifelogger to gain more insight into their activity and overall wellness. Most importantly, the additional metadata that the LSC test collection provides allows us to more easily incorporate retrieval functionalities into our system without having to spend additional time on concept detection. Finally, the large number of search topics that is included in the test collection can be useful for evaluation of our system’s retrieval capabilities.

2.3 Lifelog Systems

In this section, we provide an overview of lifelog systems which make use of geospatial visualization and temporal filtering techniques to visualize lifelog data and make it accessible. We differentiate between two types of lifelog systems, namely lifelog retrieval and lifelog exploration systems. Lifelog retrieval systems focus on enabling the user to search their lifelog data for specific memories and optimizing this search process, while lifelog exploration systems focus on allowing the user to explore their lifelog data. As our research concentrates on the suitability of geospatial visualization and temporal filtering techniques for lifelog exploration, lifelog exploration work is most relevant to us. However, there are also a number of retrieval systems which include geospatial visualizations and temporal filtering techniques which are relevant. As such, these retrieval systems are also included in this overview.

We discuss the relevant systems in two sections, divided by category: the systems that use geospatial visualization (Section 2.3.1) and the systems that use temporal filtering and/or visualization (Section 2.3.2). We emphasize that these categories are not mutually exclusive; if a single system employs both geospatial visualization and temporal filtering techniques, we will discuss it in both sections. We use this subdivision to determine which types of geospatial visualization and temporal filtering techniques are popular in lifelog systems. We also discuss, if applicable, if and how these techniques were evaluated and relevant findings. From this, we can determine gaps in current knowledge and which promising techniques are in need of further evaluation.

2.3.1 Geospatial visualization

An important characteristic of lifelog data is the GPS data that it often contains. This can be used to create geospatial visualizations of the data by displaying the data points on a map. A number of commercial photo applications already show the value of such a visualization, such as Apple’s Photos app (Figure 2.1). Lifelog systems presented in related work seem to have taken inspiration from interfaces like Apple’s Photos and use similar techniques to visualize the GPS characteristic of lifelog data. However, only a few works have actually evaluated the effectiveness of geospatial visualizations for lifelog data; Kalnikaite et al. [23] found that combining GPS data visualized on a map with the lifelog photos helped people to recall more events. Chowdhury et al. [4] performed a comparative study with different geospatial visualizations that each display GPS data in combination with lifelog photos. They found that the version of their system which included a geospatial visualization through which the lifelogger could access their lifelog photos was most preferred by

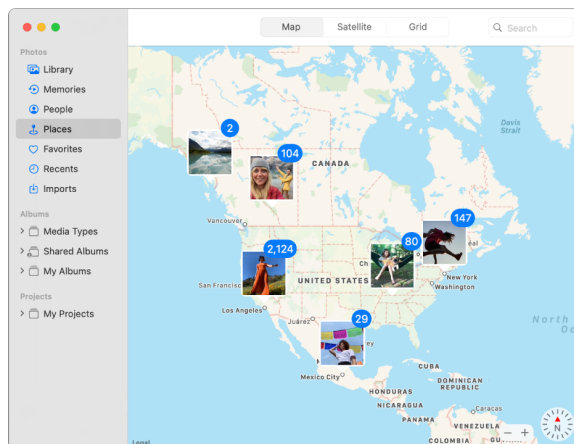


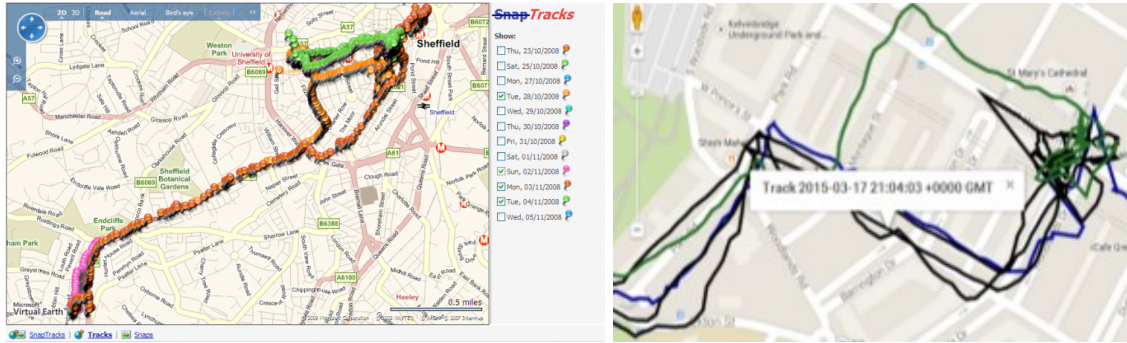
Figure 2.1: Apple’s Photos ‘Places’ interface which shows the user their photos by location.

users. From these studies, we derive that a map-based interface indeed can have a high utility for lifelog data. This is reflected in the number of lifelog exploration works which employ similar techniques.

In this subsection, we provide an overview of these works and their functionalities and we classify which type of geospatial visualization they employ. Figure 2.2 shows examples of the four different types of geospatial visualizations that we distinguished from the works we reviewed in this section, namely *Dot*, *Route*, *Heat* and *Cluster*. However, we will not discuss the relevant systems categorized by which geospatial visualization they employ, as there are a number of systems which use multiple. Instead, we first discuss two systems which visualize smartphone-based lifelog data, then we discuss several desktop-based systems for photo-based lifelog data and finally, we discuss a Virtual Reality system for photo-based lifelog data.

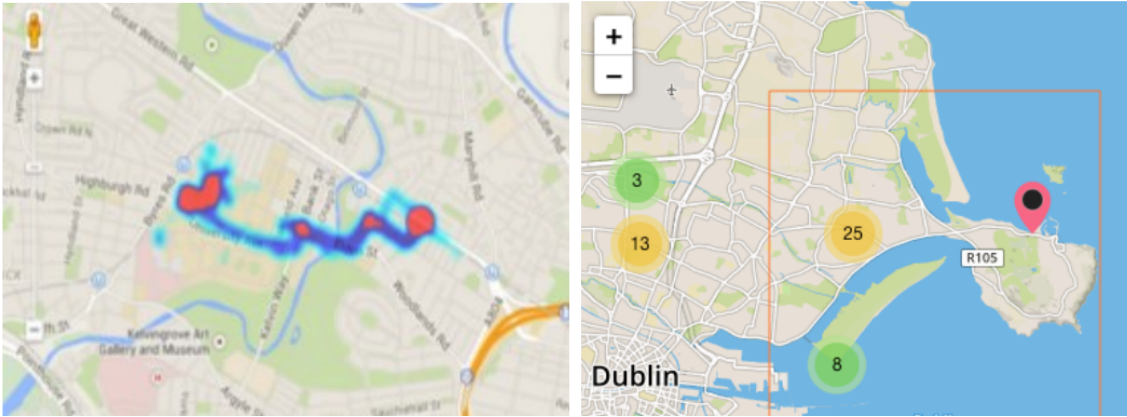
Jeon et al. [22] and Hwang and Cho [21] present systems which use geospatial visualization to display smartphone-based lifelog data. This data does not typically include photos, however, Hwang and Cho still choose to visualize the photos that the lifelogger has taken with their smartphone during the day. In the work of Jeon et al., smartphone usage data is collected in combination with, among other data types, GPS data, motion data and current activity, which participants manually logged via SMS. The main elements of their system’s interface are an information panel which shows a textual summary of the lifelog data, a timeline which shows which activity was carried out when, and a dot map which shows the location of the user at the time of recording data. An interesting feature of this system is its interactivity; when a user hovers over an element in any view, the corresponding elements in the other views light up. Hwang and Cho also visualize lifelogging data collected through smartphones. They focused, among others, on GPS, call and SMS information, photographs taken by the user and weather data. The central part of their interface is the geospatial visualization, however, it also includes a sidebar where photos that the user has taken are placed and where the user can query their lifelog data. The GPS data is displayed using a dot map visualization, which shows the location of the user during the day and includes high-level descriptions of their activities.

Next, we investigate desktop-based lifelog systems which use geospatial visualization to display lifelog data captured with a wearable camera. These are mostly systems which focus on lifelog exploration, but some retrieval systems also employ a map-based visualization to support the querying process. We start with one of the systems that was already mentioned briefly at the start of this



(a) Dot map, as used by Kalnikaite et al. [23].

(b) Route map, as used by Chowdhury et al. [4].



(c) Heat map, as used by Chowdhury et al. [4].

(d) Cluster map, as used by Tran et al. [35].

Figure 2.2: Examples of the different map visualizations that can be found in related work.

section, the system presented by Chowdhury et al. [4]. They compared four different versions of a lifelog exploration system to find the most effective one. Three of the four versions included a type of geospatial visualization, specifically, the dot, route and heat visualization. They found that the version which included the heat map, which showed the user’s most frequently visited locations, was considered most useful for lifelog exploration. The authors believe that this visualization scored so well because it displays a meaningful summary or overview to the lifelogger. They conclude that the design for a lifelog system should follow Shneiderman’s philosophy [32]: *Overview first, zoom and filter, then details-on-demand*. We consider their conclusions to be valuable, however, we feel that the differences between each version of their lifelog system are so significant that from this research we cannot accurately determine which geospatial visualization is most effective for lifelog visualization.

Kalnikaite et al. [23] also performed a comparative study with a desktop-based lifelog system. Their goal was to determine which visualization was most beneficial for memory support: A lifelog system displaying lifelog pictures sequentially (*Snapshots*), a location viewer displaying their routes (*Tracks*) on a dot map, and a combination of these two (*SnapTracks*) compared against unaided memory. They found that *Snapshots* often promoted real recall and *Tracks* promoted inferencing. However, the combination of images and location, *SnapTracks*, was preferred by participants, because it allowed them to see their information in context/overview format before zooming in on details of interest. These results also indicate the relevance of Shneiderman’s philosophy and the role that geospatial visualization can play in displaying lifelogging data to the lifelogger in a useful way.

Because very few lifelog retrieval systems use a map interface in addition to their query tools, we would like to highlight Myscéal, developed by Tran et al. [35]. Myscéal is an interactive lifelog search engine which allows for text querying in order to retrieve relevant lifelog photos. In addition to text and temporal query options, they present the location of the query results on a map. The map uses a cluster visualization (Figure 2.2d) to represent the number of lifelog photos recorded at a certain location. The user can filter results by marking a rectangular area on the map, so the system only returns results that have been captured within that region. Even though Tran et al. do not evaluate the map interface separately, we believe that, similar to heat map interface, the cluster visualization can be suitable for geospatial visualization of lifelog data because of the overview it provides.

Finally, we want to mention a Virtual Reality lifelog system presented by Hürst et al. [20], which visualizes photo-based lifelog data. In this system, the main method to access the lifelog photos is through the map; all photos are placed on the map in a dot map visualization. The photos are displayed to the user as they get close to the dot markers. To inspect a single picture, the user can 'pick it up' and look at it in Virtual Reality it as they would a printed photograph in the real world. Of course, this exact type of interaction does not extend to lifelog visualization on a desktop directly. However, we find that it is still relevant to mention this work, as it illustrates that geospatial visualization of lifelog data can be an effective method for multiple platforms.

2.3.2 Temporal filtering and visualization

Another important feature of lifelog data is time. In related work, we identify a number of ways to filter the lifelog data temporally or to visualize how the data is distributed temporally, namely a drop-down temporal filter, a calendar temporal filter and a timeline visualization, which is often combined with event segmentation techniques. We discuss each technique separately.

Several of the lifelog systems that we reviewed in the previous section employ a very minimalistic way of temporal filtering. Examples are a simple checkbox [23] to select a certain date or drop-down menu [21, 4], where users select a date range by selecting the desired dates in a drop-down.

Surprisingly few systems use a calendar interface to allow lifeloggers to access data at or between certain dates. In fact, from related work the impression could arise that a calendar interface for temporal filtering does not provide a benefit, as very few works seem to employ this, even though it might seem like an obvious choice for temporal filtering. The systems which do incorporate a type of calendar interface are presented by Leibetseder and Schoefmann [26], Jeon et al. [22] and Lee et al. [24]. Leibetseder and Schoefmann present an interesting take on the calendar interface in their LifeXplore retrieval system. Instead of having a calendar widget, the lifelog data is presented inside a calendar, with a summary of each day which consists of a number of lifelog images and some statistics about that day, such as how many unique locations were visited. In the work of Jeon et al. and Lee et al. a calendar interface is combined with other temporal visualizations, such as a timeline [22].

A more frequently occurring temporal technique is the timeline, which is used to visualize how events and activities are distributed temporally. Often, a timeline is combined with a calendar interface, like in the work of Lee et al. [24], as seen in Figure 2.3. In their system, the lifelogger can select which day they wish to view with the calendar widget. The system then shows the lifelog photos of that day for the user to explore further. When mousing over a photo, the time at which it was taken lights up on a timeline. Jeon et al. [22], whose work we discussed before in the previous section, have a similar approach. There are also a number of works which have not implemented a lifelog system themselves but which provide suggestions for the use of a timeline as a valuable visu-

alization for lifelog data. Hopfgartner et al. [19] suggest using a timeline as a way for a lifelogger to relive their day. According to them, each event of a day should be depicted on a timeline through a representative lifelog photo. The lifelogger can then select an event and view the photos of that event one by one. For lifelog exploration on tablets, Yang et al. [38] suggest a similar approach. For desktop PC visualization, they propose a rich timeline where events, level of social interaction and physical activity are all displayed in an abstract way using graphs and colors. They suggest that the lifelog photos are only shown when the user drags across the timeline. However, we have not seen this type of timeline implemented as such in related work.

A temporal technique that is closely related to the timeline is event segmentation. Lifelog systems that work with event segmentation automatically separate lifelog photos into events; e.g. having breakfast, commuting to work or attending a meeting. This way, a lifelog can be displayed in a more comprehensive manner; instead of showing the complete lifelog, only a single representative photo for each event needs to be displayed. This makes the lifelog data easier to explore for the lifelogger. Event segmentation seems to have been researched more thoroughly than other temporal visualization techniques [9, 10]. Lee et al. [24] have built their lifelog system around lifelog data which has been segmented into events, where the events of the day are displayed as a collage of photos and each photo of the collage has a certain size denoting the event’s importance, as seen in Figure 2.3. Yang and Gurrin [37] emphasize the importance of event segmentation, stating that it is necessary in order to reduce complexity of a lifelog. We believe that event segmentation and visualization can aid in lifelog exploration, as it provides the user with a more concise overview over their day, week or month. However, as event segmentation is primarily a computer vision problem, we will not be applying it in our work, even though we recognize its value.

Finally, we would like to note that several retrieval systems do not include any kind of temporal filtering or visualization functionality at all. It appears that temporal filtering either is not something that is considered to be important for lifelog retrieval, or that it is considered sufficient to use a minimalistic interface for this. For the exploration and retrieval systems we reviewed which do incorporate one or multiple temporal filtering and visualization techniques, it is as of yet unclear how effective these techniques are.

2.4 Conclusions

In this chapter, we identified that even though lifelogging is a relatively new area of research, the concept of it has existed for a long time already. Only recently the technology to enable lifelogging has become available, but the large amount of data it produces is challenging to work with. Lack of access to lifelogs is therefore one of the most important challenges of lifeloggers and research focused on providing this access is needed, as lifelogging has many benefits.

One reason why access of lifelog data still remains an open challenge could be because until recently there were no standardized test collections of lifelog data for researchers to work with. This is mostly because making lifelog data publicly available comes with privacy concerns. However, currently there are a number of test collections available, out of which we find that the LSC’20 test collection is the most suitable for our purposes.

In this chapter we provided an overview of lifelog systems that use geospatial visualization and temporal filtering techniques. We identified that common or promising techniques are the dot, route, heat, and cluster maps for geospatial visualization and drop-down menu, calendar interface, timeline and event segmentation for temporal visualization or filtering. We have provided an overview of our findings in Table 2.1, where we marked which systems use which techniques. From Table 2.1, it

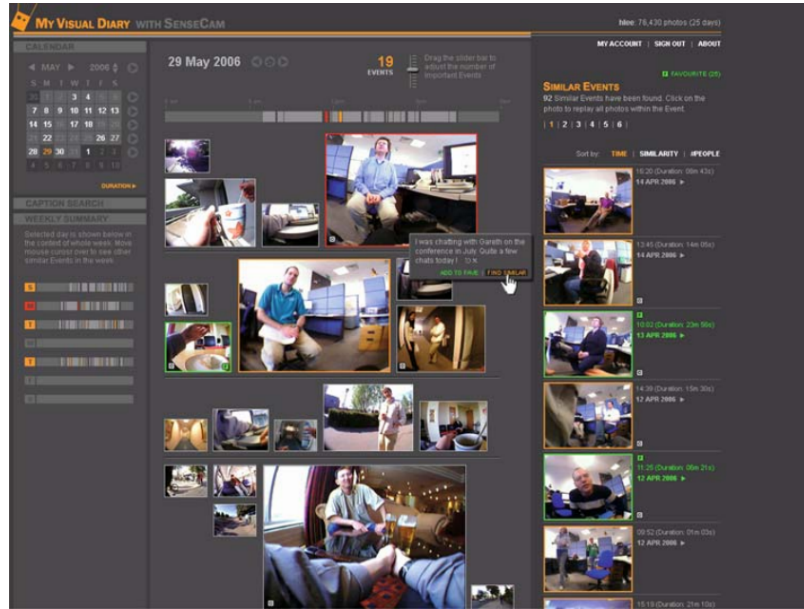


Figure 2.3: The lifelog visualization system from Lee et al. [24], which employs a calendar filter, timeline visualization and event segmentation simultaneously.

becomes clear that visualization systems more frequently employ geospatial visualization techniques than retrieval systems. We note that we inspected more retrieval systems than the ones referenced in this chapter, but these were not included in our discussion as they did not employ any standard visualization techniques.

From Table 2.1 we identify that the dot map is the geospatial visualization technique that is used most frequently. We believe that this is the case because it is easy to implement; placing markers on a map and using these as access points to the corresponding data points is a simpler task than implementing a method which allows for photo access through, for example, a heat map, as with a heat map there is no clear method to access a data point. However, none of these works have investigated whether the dot map was the most suitable geospatial visualization for their system. Only Chowdhury et al. [4] have compared different lifelog visualizations, featuring various geospatial visualizations, on how effective they are for lifelog exploration. However, the visualizations which they compare each contain differing combinations of features, rendering each lifelog visualization method too distinct to confirm which geospatial visualization is most effective for lifelog exploration. Therefore, we conclude that more research is needed into the suitability of geospatial visualization techniques for lifelogging.

From Table 2.1 we also identify that one of the most frequently used temporal visualizations is event segmentation, especially in combination with a timeline. This is a natural combination, as segmenting lifelog data into events and visualizing these on a timeline is an effective way of providing a lifelog overview. This way, lifeloggers need only inspect a small number of photos to remember what happened in a day. However, as mentioned previously, even though we recognize the value of event segmentation for lifelog access and exploration, we will not incorporate it into our research since it is not within our scope. We will also not implement a timeline into our system, since a timeline interface is heavily connected with event segmentation. Without event segmentation, it will be difficult to properly research its effectiveness. Instead, we explore other time-related visualizations,

such as calendar-like visualizations.

Table 2.1 also shows that surprisingly few systems employ a calendar visualization. For systems built for the Lifelog Search Challenge, this can be explained by the fact that other than a year range, date is not used in the search topics, and for that reason querying by date is likely not considered to be a valuable feature. For exploration systems, an explanation could be that participants only collected lifelog data for a short period of time and therefore a calendar was not considered useful enough to implement. However, we note that if a lifelog system intends to display lifelog data from an extended period of time, including temporal filtering and visualization techniques is imperative. More research is needed into the suitability of different temporal filtering and visualization techniques for lifelog data.

Finally, we would like to emphasize the importance of providing the user with an overview of their lifelog data, as indicated in several works reviewed in this chapter [4, 23, 37]. We believe that this overview can be provided by geospatial and temporal visualization techniques, which is a theory that the results of Chowdhury et al. and Kalnikaite et al. support.

In the next chapter we introduce, based on related work presented in this chapter, which visualizations we consider most suitable or representative for lifelog data and thus will incorporate in our system.

	Dot	Route	Heat	Cluster	Drop-down	Calendar	Timeline	Event
Kalnikaite et al.	x				x			
Chowdhury et al.	x	x	x		x			x
Jeon et al.	x					x	x	
Hwang & Cho	x				x			x
Hürst et al.	x							
Hopfgartner et al.							x	x
Yang et al.							x	x
Lee et al.						x	x	x
Tran et al.				x				x
Leibetseder et al.								x

Table 2.1: Overview of which visualizations and/or filtering methods are used by each of the lifelogging systems that we have discussed. *Dot*, *Route*, *Heat* and *Cluster* refer to map visualizations (as seen in Figure 2.2), while *Drop-down*, *Calendar*, *Timeline* and *Event* refer to temporal filters or visualizations. All works listed are focused on lifelog exploration, except for the bottom two from Tran et al. and Leibetseder et al., which are focused on lifelog retrieval.

Chapter 3

Lifelog System Design

In the previous chapter, we reviewed a number of related studies on lifelog systems which incorporated geospatial and temporal filtering and visualization techniques. We determined that a number of techniques, such as a dot map or timeline, are frequently employed in lifelog systems, specifically lifelog exploration systems. However, few of these related studies provide reasoning or evaluation for the techniques they implement. For this reason, it is as of yet unclear if geospatial visualization and temporal filtering techniques are suitable for lifelog access, search and exploration and if so, which techniques are most useful for different search goals in the context of lifelog data.

To fill the gap in current knowledge, we built and evaluated our own lifelog system which incorporates a selection of the techniques we reviewed in Chapter 2. This chapter describes which techniques we have chosen to incorporate in our lifelog system and why, the design process of the system, which includes how we substantiated our design choices through related work and our understanding of lifelogger search goals, a brief overview of the system’s implementation, and finally, the research questions we hope to answer through evaluation of our system.

3.1 Motivation

Before selecting which visualization and filtering techniques to evaluate, we find it important to understand the different search goals that a lifelogger may have when accessing their lifelog data, so that we can select appropriate techniques to support these goals. Section 3.1.1 describes these search goals. Section 3.1.2 introduces the visualization and filtering techniques we have chosen to evaluate, accompanied with our expectations on how they can support these search goals.

3.1.1 Search goals

We previously differentiated between two different types of lifelog systems, namely exploration and retrieval systems. Exploration systems focus on allowing lifeloggers to **explore** their lifelog data (often through lifelog visualization), while retrieval systems focus on allowing lifeloggers to **search** lifelog data. However, these are global descriptions of system goals and do not describe clearly with which goals in mind a lifelogger can access their lifelog data, as exploratory search can take many different forms. To understand better with which search goals lifeloggers may access their lifelog, we look to the work of Rose and Levinson [30]. They have defined a set of search goals that users may have for web search. This does not completely extend to lifelog search, however, we feel that a number of the search goals presented by them are representative of the search goals that a lifelogger could have when accessing their lifelog data. We have selected the goals that we consider relevant

and put them in the context of lifelog search.

In addition to this, we want to refer to the work of Wilson and Elswiler [36], who introduce a specific subcategory of exploratory search that we consider applicable to lifelog search: Casual-leisure searching. This is a search goal where there is no specific information need and the experience of searching is considered more important. This brings us to four different types of search goals that we consider relevant for lifelog data. We also mention how we believe geospatial visualization and temporal filtering techniques can support each search goal:

- **Directed, closed search:** The lifelogger queries with the goal of finding a specific photo. Geospatial visualization and temporal filtering techniques can support this search goal when a lifelogger has a very specific geospatial or temporal query.
- **Directed, open search:** The lifelogger queries for an event or location with the goal of finding an undetermined number of photos that fit their query. For example, when a lifelogger wants to find every time they visited McDonald's, a geospatial visualization could help them identify which McDonald's they like to visit.
- **Undirected search:** The lifelogger queries for an event or location to explore their photos of that event or location. Here, geospatial visualization and temporal filtering techniques are useful when the lifelogger wants to, for example, explore where they went during in a certain week or remember what they did during their holiday abroad.
- **Casual-leisure search:** The lifelogger places random queries to see what they can find. Especially a geospatial visualization or a temporal visualization can support this goal, as visualizing a lifelogger's data in a visually appealing way may support this browsing behavior.

We categorize *Directed, open search*, *Undirected search* and *Casual-leisure search* as exploration search goals that apply mostly to a lifelog exploration system. *Directed, closed search* is the only search goal that directly applies to a lifelog retrieval system. The other *Directed open*, *Undirected* and *Casual-leisure* search goals may also be satisfied by placing a query in a retrieval system, although in such a systems it might be more difficult for a lifelogger to adequately explore the results.

3.1.2 Visualization selection

Next, we present select which geospatial visualizations and temporal visualization and filtering to evaluate, as we believe they can support the search goals we introduced in Section 3.1.1. In the previous chapter, we identified different geospatial visualizations and temporal filtering and visualization techniques which are commonly used in lifelog systems. We concluded that some techniques are more commonly implemented in lifelog systems than others, and that many of these have not been evaluated for how effective they are for access, search and exploration of lifelog data. Rather, they seem to be the result of a more or less arbitrary choice.

The success of the systems that have employed these techniques might imply that they are suitable, but no focused research has been done on this. To fill this knowledge gap, we investigate a number of visualizations to determine if they are suitable for access, exploration and search of lifelog data and to determine which are the most preferred (combinations of) techniques. From the related work that we have examined, we distilled a number of geospatial visualization and temporal filtering techniques to evaluate. Below we have listed the techniques, together with our reasoning and expectations for them:

- **Geospatial Dot map:** This is the most basic way to visualize geospatial data and is used by many lifelog exploration systems, thus, we selected it for evaluation. Because of its detailed

nature, we expect that the Dot map will be most useful during *Directed, closed search*, as it is easy to inspect a small number of query results one by one through the dot map.

- **Geospatial Heat map:** The Heat map is one of the most common ways to visualize how many data points are recorded at a location, which is why we selected it for evaluation. We expect that the Heat map will be the preferred visualization when exploring lifelog data, during *Undirected search* or *Casual-leisure search*, as it provides a clear overview of which locations the lifelogger likes to visit most. However, lifelog access through the Heat map might be less intuitive, as there is not a single clear access point to a photo.
- **Geospatial Cluster map:** The Cluster map visualizes how many data points are located in a certain area by showing a color-coded icon with the number of data points it contains displayed inside it, as seen in Figure 2.2d. Individual data points can be accessed by opening the clusters. We expect that the Cluster map will be a good balance between the accessibility of the Dot map and the overview that the Heat map provides, and can therefore work well for all types of search goals. There are no lifelog exploration systems which use this visualization, but due to the overview it provides it could be a suitable visualization for lifelog exploration.
- **Temporal Drop-down temporal filter:** This is a basic temporal filtering technique, where users input a date range by selecting day, month and year from a drop-down. We expect that this feature will be used most during *Directed, closed search*, as it is one of the quickest ways of selecting a time range. However, there are some usability issues with a drop-down temporal filter, as the user needs to complete a relatively large number of clicks to fill in a complete date range. Nevertheless, we are interested in seeing if a simple temporal filtering interface like this can suffice for lifelog access, search and exploration.
- **Temporal Calendar:** The Calendar is a common temporal filtering technique, where users select a date range by making two date selections on a calendar interface. We expect that this technique will be popular both for search and exploration purposes, as a calendar allows for quick date selection but also has an element of exploration to it.
- **Temporal Color-coded calendar matrix:** We imagine this filter technique as a combination of a temporal filter and a visualization. This interface will look somewhat like a spreadsheet, where a timeframe can be selected through clicking and dragging on cells which represent days or hours depending on which view is active; Week, Month or Year. The cells have a color intensity, which indicates the number of photos captured in that period. This interface is not something that we have encountered often in related work, however, a similar idea was suggested by Yang et al. [37] for activity visualization and we find that it could be an interesting technique to explore, as it also allows us to incorporate information about the temporal distribution of the data into a calendar-like filtering feature. Therefore, we expect this to be the most popular interface for exploration-based search goals, as it offers a visualization in addition to providing filtering functionality.

We note that this list is not exhaustive. For example, a route map would also be a valid geospatial visualization. However, we feel like this is not an appropriate visualization for large amounts of geospatial data as we expect it may quickly appear cluttered and it will only be meaningful when the lifelogger is traveling, which is why we have not selected it for evaluation. The timeline is also a popular visualization, but we chose not to evaluate it as it is best combined with some type of event segmentation, which is outside the scope of this work.

3.2 Design

Next, we show how we have incorporated the visualizations selected in Section 3.1.2 into a lifelog exploration system design. The system is built as a desktop web application, named the Lifelog Browser. In this section, we will describe the inspirations from related systems which have led to the design of the Lifelog Browser. Then, we present the first designs we created for the system through mock-ups of the interface and compare this to the final design of the system, which has undergone some changes during implementation.

3.2.1 Features and Inspiration

For the design of our system and its features we have taken inspiration from Flickr Cities [7] (see Figure 3.1) and a number of related works. Flickr Cities is an exploratory interface designed to browse the Yahoo Flickr Creative Commons 100 Million Dataset [34], a large collection of geotagged Flickr photos. The interface provides a number of filtering options, such as date, hour, weekday and tag. The most notable of the features is the map interface, on which the locations where the photos were taken are visualized in a heatmap-like manner. We found this interface a strong inspiration for a lifelog exploration system, as it is designed for photo exploration.

Our main takeaway from Flickr Cities is its layout, which features a large map interface on the left and a sidebar on the right where temporal, tag and location filter options are presented. The size and positioning of the map encourages exploration, as it is the first feature a user sees, while the filtering options remain in the background to be used when the user desires. We have chosen to follow a comparable design to incite a similar reaction in our users.

Another interesting feature of the Flickr Cities interface is how it visualizes the temporal distribution of photos. On the right of the interface, users may select a time range from a timeline or a circular visualization of a 24-hour or week period, where each time range has a grey bar. The length of the bar denotes the number of photos that have been taken during that time period. This inspired us to explore a similar combination of a visualization and temporal filtering interface, which came to be the Temporal color-coded matrix.

Additionally, both the Flickr Cities interface and the numerous lifelog retrieval systems built for the LSC have inspired us to incorporate a location and tag filter into our system. The systems built for the LSC have shown us that incorporating a location and tag filter is crucial to the success of a lifelog retrieval system. Even though few lifelog exploration systems we discussed provide the functionality to filter by location (other than navigating to a location on the map) or tag, we imagine that even when accessing their lifelog for exploration, a lifelogger may want to filter by a location they frequent or activity they like to perform, which a location or tag filter allows them to do. Since the LSC'20 test collection already includes semantic locations and image annotations for each photo, it is still within our scope to include this functionality into our system.

We also wish to include features in our lifelog system that are more geared towards exploration specifically for lifelog data. Hopfgartner et al. [19] emphasize the importance of providing a visual summary of a day's events to the lifelogger, where they may inspect each event photo by photo if they desire. Such a functionality is essential for exploration of lifelog data, as it consists of many photos which in conjunction describe a period of time. Hence, it is imperative that a lifelogger can inspect all these photos in chronological order so they may remember the events of a period of time. A technique like event segmentation would be optimal for this, however, it is not within our scope to incorporate event segmentation into our system, but we still wish to follow the recommendations Hopfgartner et al. provide. Therefore, we designed a feature, additionally inspired by Shneiderman's

philosophy [32](*overview first, zoom and filter, then details-on-demand*), which provides lifeloggers with the possibility to view photos from a specific day and scroll through them chronologically. We call this feature the *Detail view*, as it allows the user to view all details and context of a lifelog photo. Upon clicking on a photo, the *Detail view* will appear as a pop-up where the user may inspect a large version of the photo and the photo’s metadata. Below, they find a scrollable field where the photos of that day are displayed chronologically. Leibetseder and Schoeffmann [26] have a similar view which they call the *Image view*, which shows the metadata of a photo and similar photos. We have used their design as inspiration, although we replaced the similar photos in their design with photos that were taken on the same day.

Finally, in our design we wish to include an overview below the map, where a selection of photos that were found within the bounds of the map is shown. This is inspired by Hwang and Cho [21]. We have called this feature *Photos in this area*, as it shows a chronological subselection of approximately 100 photos of the photos shown on the map. This feature also follows Shneiderman’s philosophy, as it provides the user with an overview of what kind of data they may find on the map. While the geospatial visualization provides the user with an overview of where each photo is located, *Photos in this area* shows the user an impression of what the photos look like. Additionally, with this feature access to photos is simplified, as the user is not required to open all photos from the map to inspect them.

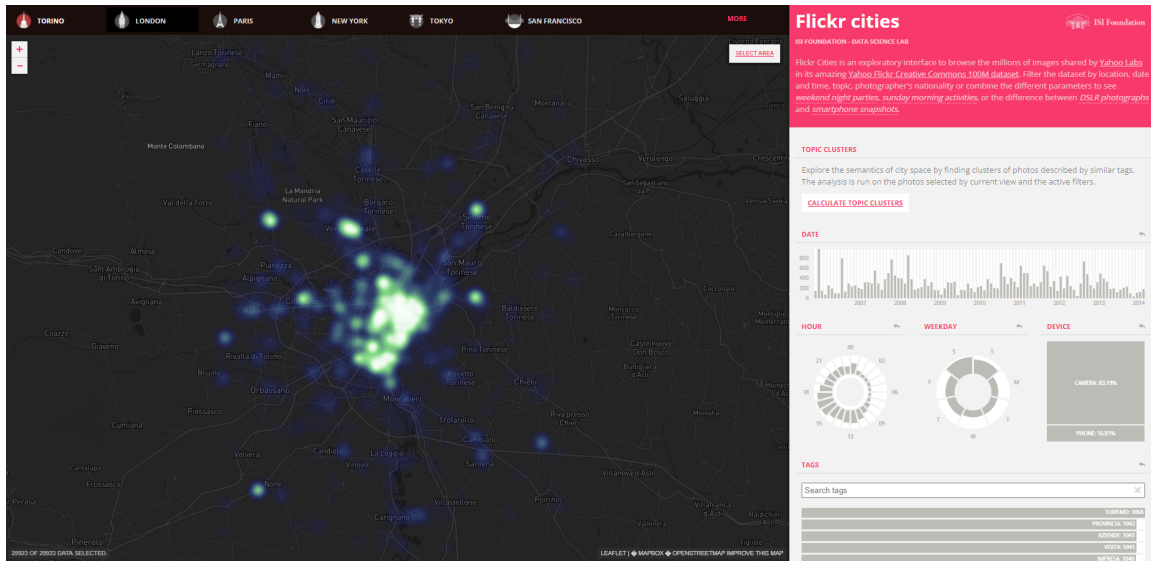
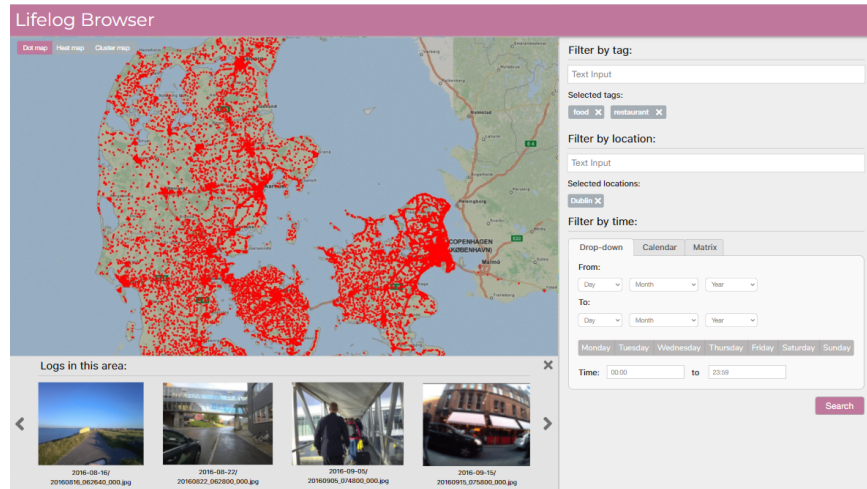


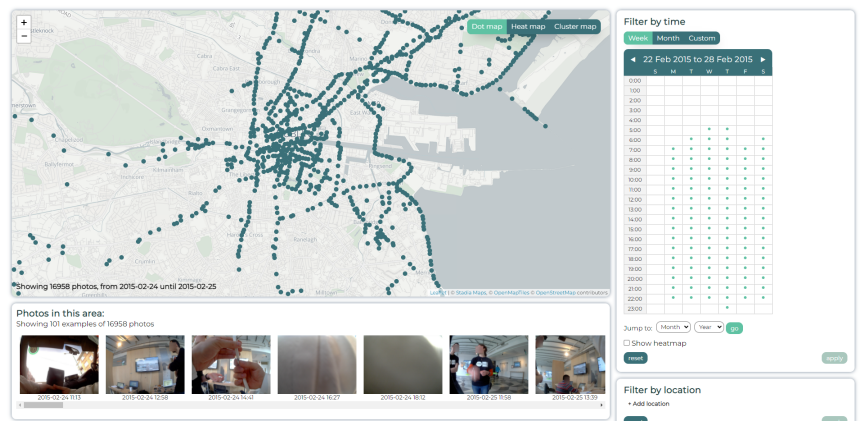
Figure 3.1: The Flickr Cities interface, showing a map interface visualizing the location of the photos on the left and various filtering options on the right.

3.2.2 Mock-ups and final design

After deciding which features to implement into the Lifelog Browser, we created mock-up images of each feature of our lifelog system to serve as a guideline during implementation. However, during the implementation process, a number of features have undergone several design changes. Below, we present each mock-up image that we created before starting implementation, side-by-side to a screenshot of the final implementation in the system. For each image, we explain and motivate our redesign choices.



(a)



(b)

Figure 3.2: A comparison of our original design and the final design for the Lifelog Browser.

Overview

We start with an overview of the system. The images for the overview can be seen in Figure 3.2. As described in Section 3.2.1, we chose a layout with a large map of the left and a smaller search bar on the right to encourage exploration. The map is the main access point for the photos, which is also why it is larger than the other features. The user may switch map visualization with the buttons at the top of the map. We chose to place the *Photos in this area* feature below the map, as it is closely related to the map in the sense that it shows a subselection of the photos that are displayed on the map. This layout has not changed from our original designs.

A small design change has been done to the filter bar. Originally, we intended to have only one *apply* button, which upon being clicked, would apply all filters that the user had chosen, updating the map with the new data. However, during the implementation process we found that it would be more intuitive to place each filter in its own box, with a separate *reset* button, so it is clear to the user that filter use is optional. This also allows users to reset each filter separately, which can be convenient when forming complex search queries.

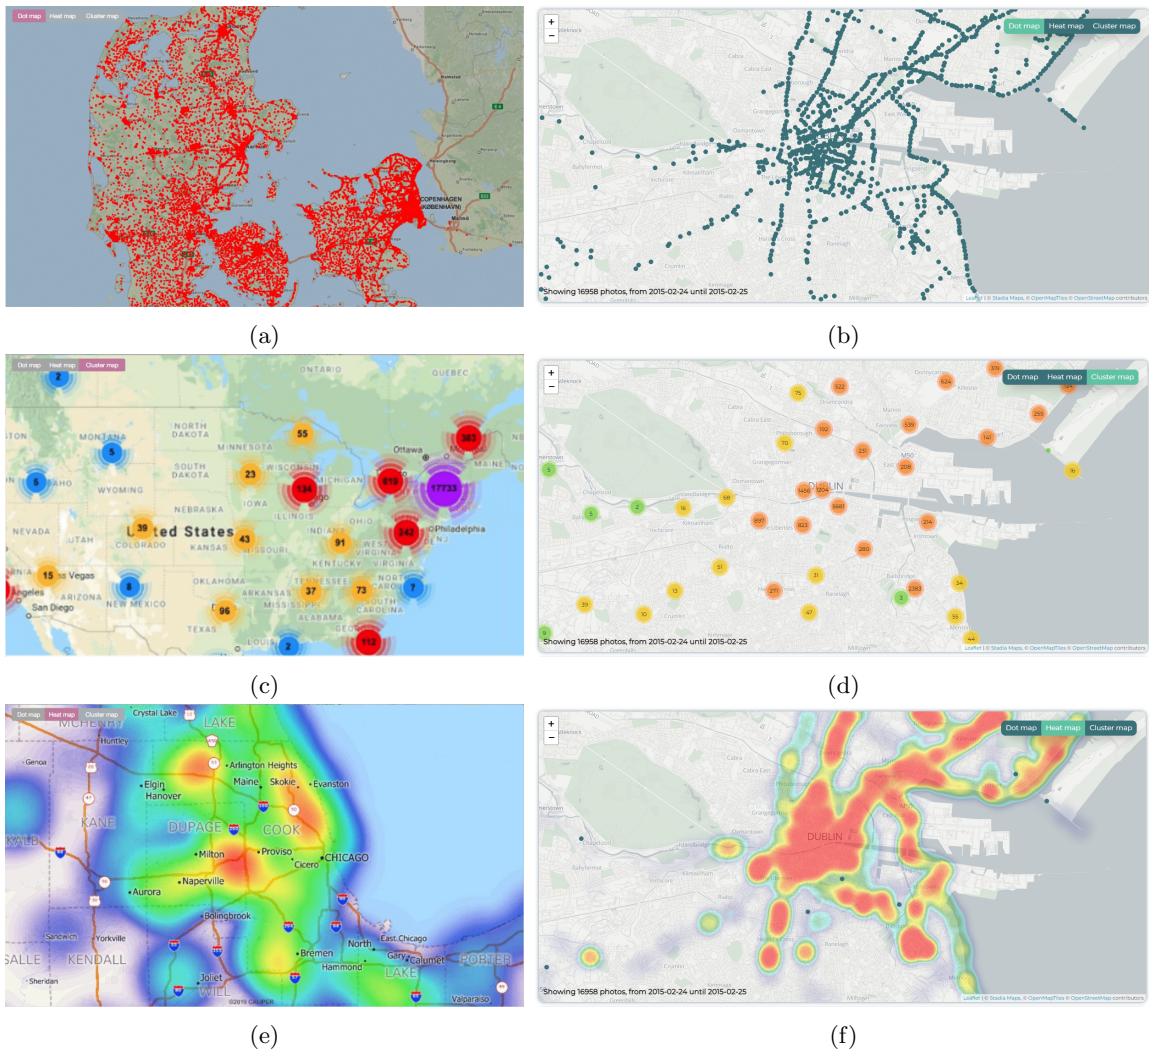


Figure 3.3: A comparison of our original geospatial visualization designs (note that these images were taken as examples from Google Images) and the final geospatial visualizations.

Geospatial visualizations

Next, we compare the original and final designs of the geospatial visualizations, as shown in Figure 3.3. Few design changes were done. The images for our mock-up were taken from Google Images, hence, the styling looks different in the final version of the system, but the functionality is the same. The styling of the geospatial visualizations in the final version of the system is mostly the default styling that is provided by the libraries that we used to implement these visualizations, as we found them suitable and clear for our purposes. Only for the Dot map, we have increased the size of the dot icons to make them easier to click. The color of the dot icons was changed to blue to suit the overall color palette of the system.

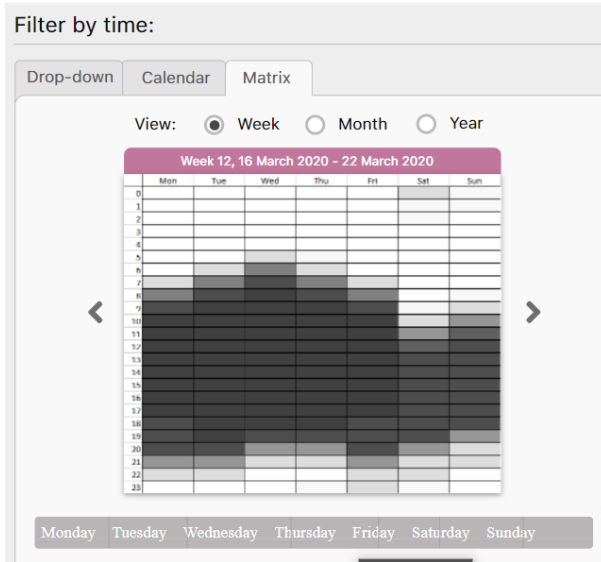
Temporal filtering

Next, we present our designs for the temporal filtering interfaces, as seen in Figure 3.4. Many changes were done to this design. After implementing our original design with the three different temporal filters as presented in Section 3.1.2, we found that it was felt counter-intuitive to include three different temporal filtering interfaces which effectively all provided exactly the same filtering functionality. This was especially the case with the Month view of the Color-coded calendar matrix, which was functionally and visually very similar to the Calendar interface. The number of choices that the different variations of the Color-coded calendar matrix provided (Week, Month and Year), in addition to the Drop-down and Calendar filter functionalities, was also overwhelming. We feared that this would lead to confusion for users and therefore decided to remove the Calendar interface entirely and replace it with the Month view of the Color-coded calendar matrix. We also removed the Year view of the Color-coded calendar matrix, as we felt all its functionality could also be provided by either the Drop down temporal filter of the Month view of the Color-coded calendar matrix. The Week view of the Color-coded calendar matrix was kept as a separate temporal filter. In summary, in the final design, instead of having the Drop-down, Calendar and Color-coded matrix temporal filters, we instead have a Week and Month temporal filter, derived from the original design of the Color-coded calendar matrix, and a Custom temporal filter, which is derived from the original design of the Drop-down temporal filter.

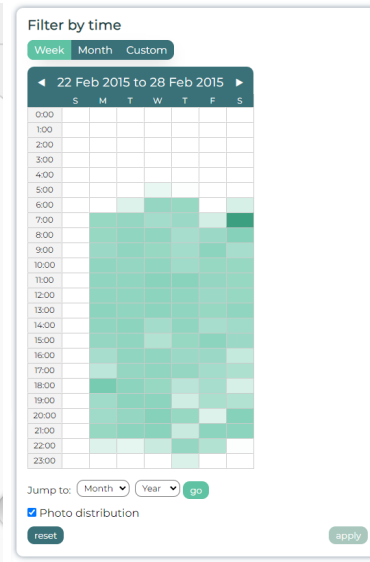
In the final design of the Week temporal filter (as seen in Figure 3.4b), the spreadsheet columns represent days and the rows represent hours. Users may select a time frame by dragging on the spreadsheet. If photos were taken during a certain hour, a dot appears in the corresponding spreadsheet cell. Users have the option to activate the *Photo distribution* feature, which replaces the dots with colors which correspond to the number of photos that were taken that hour. We have refrained from implementing the color-coding as the default functionality, to be able to evaluate its value separately from the Week and Month temporal filters.

The final design of the Month temporal filter (as seen in Figure 3.4d) is very similar to the previous design of the calendar interface (as seen in Figure 3.4c). We have chosen to display two months at once to the user, so they may select dates from two consecutive months. Just like for the Week view, a dot appears in cells which correspond to a time-frame when photos were taken. The user may also view the *Photo distribution* of each day by clicking the checkbox. The Month temporal filter also contains a separate feature where the user may select a time. This was implemented to provide equal functionality the Week temporal filter, as in the Week temporal filter users may select a time range by placing their selection in the spreadsheet.

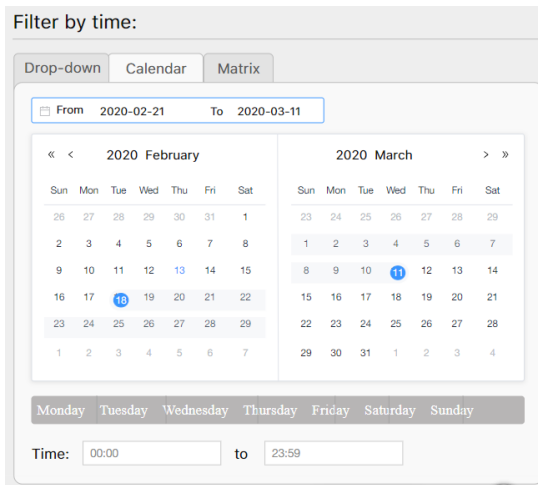
Finally, for the Custom temporal filter (as seen in Figure 3.4f), not much has changed compared to our original design mock-up for the Drop-down temporal filter (as seen in Figure 3.4e). The only change made was to convert the time selection method from an input field to buttons, as this was more aesthetically pleasing and usable for time selection.



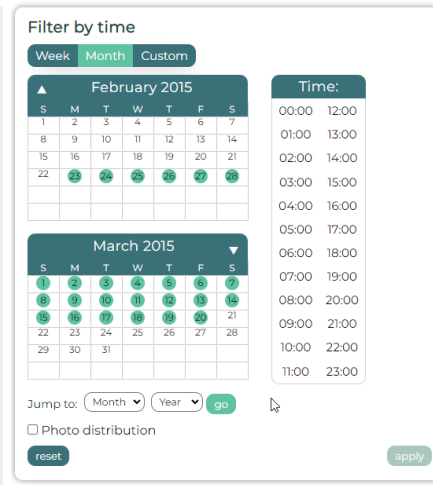
(a)



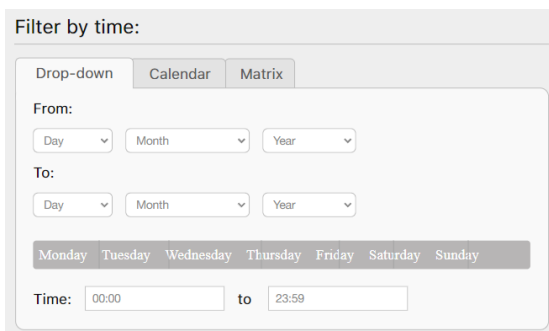
(b)



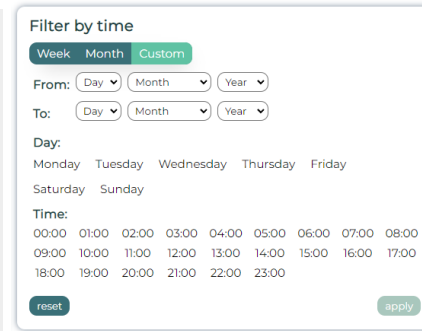
(c)



(d)



(e)



(f)

Figure 3.4: A comparison of our original temporal filter interface designs and the final interfaces.

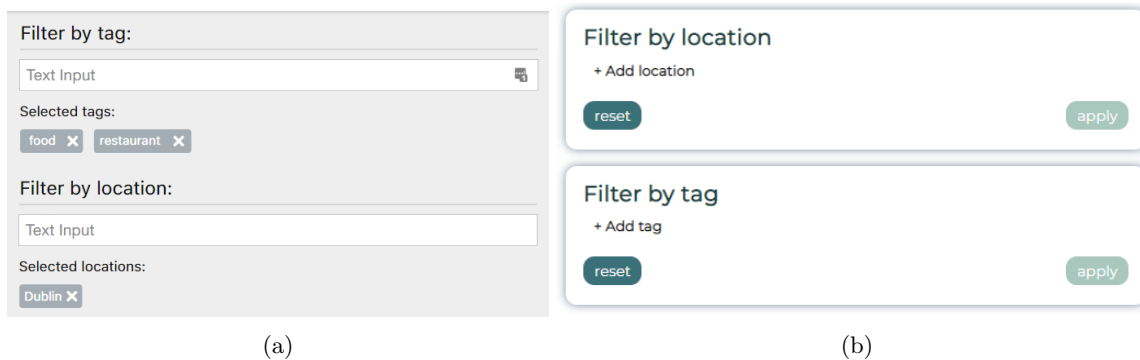


Figure 3.5: The mock-up (a) and final (b) versions of the Location and Tag filters.

Location and Tag filter

Next, we present our designs for the Location and Tag filters, as seen in Figure 3.5. Mostly aesthetic changes have been done to these designs. The only functional change that has occurred is that both the Location and the Tag filter now have their own *reset* and *apply* buttons, to allow both fields to be reset separately.

Detail view

Finally, we present our designs for the Detail view, as seen in Figure 3.6. Mostly aesthetic changes have been done to the design of the Detail view. A non-aesthetic change is that we intended to show the cumulative steps in the metadata table, but this idea was discarded due to time constraints.

3.3 Implementation

Following the designs presented in the previous section, we implemented our lifelog exploration system, the Lifelog Browser, as a web application. From this application, the LSC'20 test collection can be accessed and explored.

The back-end of the Lifelog Browser was built using Django¹. The metadata belonging to the lifelog photos of the LSC'20 test collection is stored in a MariaDB database² which runs in a Docker³ container. The front-end was built using JavaScript, HTML and CSS and the libraries Leaflet⁴ (with plug-ins Markercluster⁵ and Leaflet.heat⁶), jQuery⁷, and Js spreadsheet CE⁸. In this chapter, we will briefly describe the details of the system's implementation, which includes the steps we took to set up the database (Section 3.3.1) and how the back-end and the front-end were built (Sections 3.3.2 and 3.3.3).

¹<https://www.djangoproject.com/>

²<https://mariadb.org/>

³<https://www.docker.com/>

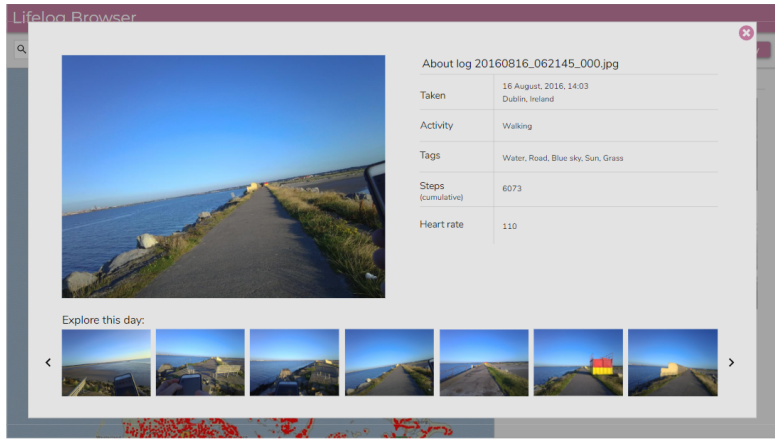
⁴<https://leafletjs.com/>

⁵<http://leaflet.github.io/Leaflet.markercluster/>

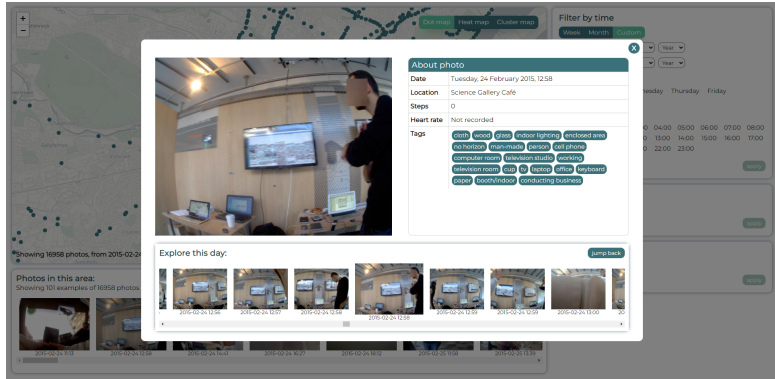
⁶<https://github.com/Leaflet/Leaflet.heat>

⁷<https://jquery.com/>

⁸<https://github.com/jspreadsheet/ce>



(a)



(b)

Figure 3.6: The mock-up (a) and final (b) versions of the detail view.

3.3.1 Database

Metadata preprocessing

The LSC'20 test collection comes with two different CSV metadata files. The first metadata file contains, for **each minute** that is covered in the test collection, a metadata entry which contains a minute ID, the UTC time, the local time, the local timezone, the latitude and longitude of the lifelogger's location, the semantic name of the lifelogger's location, the elevation of that location, the lifelogger's speed, heart rate, calories burned, the type of activity he was engaged in, and the number of steps he took that minute. The second metadata file contains image annotations **per photo** which were automatically generated by Microsoft Cognitive Services. Since one file contains metadata **per minute** and the other file contains the image annotations **per photo**, we first combined the two files into a single file which contained all metadata and image annotations **per photo**. This means that a single minute metadata entry may be duplicated over multiple photos.

The CSV file with metadata per minute included an *activity* field. The activities recorded in this field are limited to movement-based activities like *walking* or *transport*. We felt that it could be useful for users to be able to query by these activities, but did not find them informative enough to be able to rationalize an entirely separate filter feature for them. Therefore, these activities, if available, were included into the tags for each photo.

The CSV file with image annotations for each photo contained a confidence value and bounding box coordinates for each image annotation (tag). To ensure that all information about a photo was shown to a user, we included all image annotations of every photo, regardless of their respective confidence value. We discarded the bounding box coordinates since we felt it was unnecessary to display these to users.

The relevant metadata from both files was combined into a single CSV file with the help of a Python script using the Pandas⁹ library, which simplifies working with CSV files. This process resulted in a single file CSV with contained the metadata per photo for each photo in the LSC'20 test collection. The following metadata was included: *minute ID*, *image path*, *UTC time*, *local time*, *timezone*, *day*, *latitude*, *longitude*, *semantic name*, *elevation*, *speed*, *heart rate*, *calories*, *steps* and *tags*.

Database structure

Our system uses a MariaDB database which runs in a Docker container. Our main database table is the Photo table, which contains each lifelog photo as a Photo object with its corresponding metadata. The Photo database objects were generated with Django from the combined metadata CSV file. The database also contains tables which separately store locations and tags as objects. A Location object has a one-to-many relationship with the *semantic name* field of a Photo object. The Photo and Tag objects have a many-to-many relationship, which is indexed in a separate table with an entry for each combination of photo and tag in the database.

We created two additional database tables specifically for the Week and Month temporal filters. For the Week temporal filter, the table contains an object for each week lifelog photos were taken. That object stores an array, which contains an Integer element for each hour of that week. This element describes how many lifelog photos were taken during that hour. For the Month temporal filter, the table contains an object for each month lifelog photos were taken. That object stores an array which contains an Integer element for each day of that month. This element describes how many lifelog photos were taken during that day. We have chosen to store this information separately in a database table instead of querying it live from the Photo, this process was too time-consuming.

A visual overview of the database structure can be found in Figure 3.7.

3.3.2 Back-end

We use Django to query the database and communicate the results back to the front-end. There are a number of calls from front-end to back-end. The calls are executed from the front-end using jQuery, as this allows for asynchronous communication between the back-end and front-end. The following calls are implemented in our system:

- **Photo search:** This call returns all photos in the database which match the current filters that are applied. Since, for performance reasons, we only return the photos that are found within the bounds of the map view, it is executed whenever the user has finished moving the map or when the *apply* button of any of the filters is pressed. The call sends all the active filters (date, day of the week, tags, locations, hours and current bounds of the map) to the back-end. The back-end returns a JSON response with all photos that match the filters. Specifically, of each photo, its latitude, longitude and image path is returned. We return only these fields to ensure that the response takes minimal time to process. The latitude and longitude are used

⁹<https://pandas.pydata.org/>

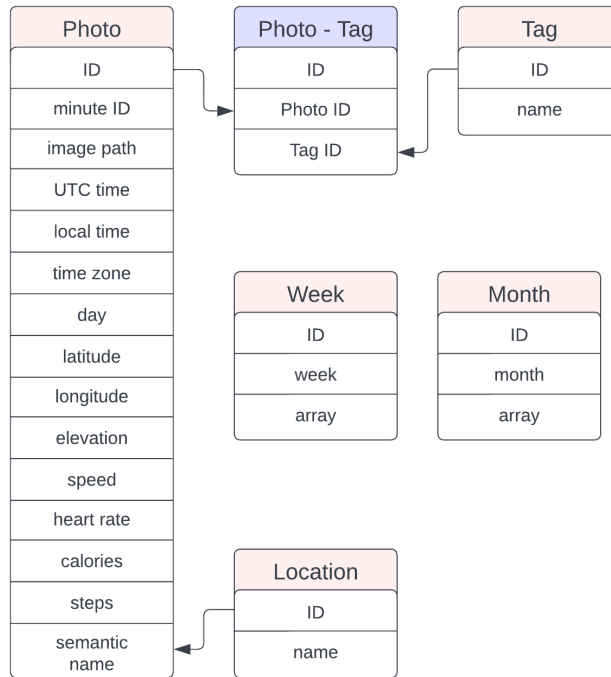


Figure 3.7: A visual overview of the database structure.

to place a marker for the photo on the map, while the image path is used to retrieve the photo file if the user opens it from the map.

- **Location or Tag search:** This call is executed whenever the user types in the Location or Tag filter fields. It returns locations or tags from the Location or Tag table that contain the user's query.
- **Single photo search:** This call is executed whenever the Detail view is opened by clicking on a photo. This call retrieves a photo's complete metadata. The metadata that is returned is its *image path*, *local time*, *UTC time*, *weekday*, *latitude*, *longitude*, *semantic name*, *tags*, *steps* and *heart rate*.
- **Day search:** This call is also executed whenever the Detail view is opened. The photos from day of the photo which is opened in the Detail view are returned. Since a person may go to bed past 00:00, we found it most suitable to return photos between 03:00 that day and 03:00 the next day, if the photo was taken before 00:00. Else, we return photos between 03:00 the previous day and 03:00 the same day. This way, we ensure that we display all photos of the lifelogger's entire day, from when they got up to when they went to bed.
- **Get matrix data:** This call is executed whenever the user skims to the next or previous time period in either the Week or Month temporal filter. For the Week temporal filter, it returns an array which describes how many photos were taken during each hour of that week. For the Month date picker, it returns two arrays which describe how many photos were taken each day of two consecutive months.

3.3.3 Front-end

The front-end implementation consists of a number of major features: the map with the three geospatial visualizations, the temporal filters, the Location and Tag filters and the Detail view. We briefly describe the front-end implementation for each of these features.

Map

The map was implemented using Leaflet, an open-source Javascript library for building web mapping applications. The advantage of Leaflet is that it can be extended with plug-ins that simplify implementation of geospatial visualizations. In addition to these plug-ins, Leaflet offered a number of base functionalities which were useful for our purposes. Leaflet allows placement of customized markers on the map, and it has a built-in popup functionality when clicking on a marker, which we programmed to display the lifelog photo that corresponds to that marker. Leaflet only works in combination with a tileset provider, for which we chose Stadia Maps¹⁰. Stadia Maps provides a number of attractive map themes and does not charge anything when a website runs locally, which made it a suitable provider.

Because under the hood all geospatial visualizations use the Markercluster library for their various functionalities, a plugin for Leaflet which we used to implement the Cluster map, we will explain the implementation of the Cluster map first and then elaborate on the customizations we performed for the Dot map and Heat map. Figures 3.3b, 3.3d and 3.3f show an example image of the styling of the Dot, Cluster and Heat map, respectively.

Cluster: For the Cluster geospatial visualization, we used the default behavior of the Markercluster library. By default, this library clusters all markers on the map within an 80 pixel radius away from each other into a single cluster. Each cluster is color-coded depending on the number of markers that are contained within it. The number of markers within each cluster is also displayed on the cluster. Clicking on a cluster will zoom into the area it covers and the cluster will split into smaller clusters. Once the user reaches a low enough zoom level, clicking on a cluster will 'spiderify' the cluster, unfolding all the markers that the cluster contains into a spiral pattern, so the individual markers can be accessed. We have not changed anything about this default behaviour and styling, as we found it suitable for our purposes.

Even though the Markercluster library allows many markers to be displayed on the map simultaneously, a limit for the number of markers on the map still had to be introduced to ensure that the system always ran smoothly. After testing several different limits, we found that we could place 40,000 markers on the map without sacrificing performance. How many markers are drawn on the map and whether this limit of 40,000 has been reached is communicated to the user through text in the bottom-left corner of the map. This text also contains a description of the date range that the data on the map covers, to allow users to apply a temporal filter should they wish to view the remaining photos currently not displayed on the map. This limit also applies to the Dot and Heat maps.

Dot: For the Dot map implementation, we faced a challenge. Because of the large number of photos to display on the map, it was virtually impossible to place each individual photo as a marker on the map using the default Leaflet functionality, as each marker on the map is a separate Div element. This resulted in the map becoming too slow to navigate when more than 6000 markers were placed on it. Additionally, placing each photo as an individual marker complicated access, because in locations where many photos are taken, such as the lifelogger's home or work, there are many

¹⁰<https://stadiamaps.com/>

markers overlapping each other, which made it impossible to access some photos.

We therefore opted to use the Markercluster library as a way to cluster multiple markers which were close together into a single marker, which reduced the load on the map and allowed users to access markers that would otherwise be overlapping with the 'spiderify' feature. To achieve the look and functionality of a traditional dot map, we modified the default behavior and styling of the Markercluster library. The cluster radius was reduced from the default of 80 pixels to 5 pixels, so as few markers as possible would be clustered together but each individual cluster could still be accessed, even in a crowded area. The styling was adapted so both the clusters and single markers were rendered as a dark blue dot. This implementation has allowed many markers to be displayed on the map while still retaining the look of a dot map and its functionality.

Heat: To generate the Heat map we used the Leaflet.heat plugin. The only modifications we added to its default behavior is ensuring that the heat layer was slightly transparent, so it is still possible to read text on the map below it. However, because Leaflet.heat does not provide any functionality for individual marker placement and interaction, we decided to overlay an invisible layer of clusters on top of the Heat map. For this, we used the Markercluster plugin with customized styling.

To retain the look of the Heat map and not provide any distractions from the heat visualization, we programmed the clusters to only become visible when the user hovers over them. Under the cursor, a faded red dot appears with the cluster bounds shown in red lines. Accessing individual markers works the same as with the Cluster map, where a cluster 'spiderifies' into a spiral of individual blue dot markers. Single markers that are not part of a cluster are also rendered as dark blue dots.

Temporal filters

We will separate our explanation of the temporal filter implementation into the Week and Month temporal filters and the Custom temporal filter, as the same library was used for the Week and Month temporal filters and their functionalities are similar.

Week and Month: The Week and Month temporal filters have functionalities and styling that are very similar to that of a spreadsheet, hence, we used a spreadsheet library to implement these features. We chose the Jspspreadsheet CE plugin, as this was one of the few free options available. Its basic functionality allows the user to select one or multiple adjacent cells on a spreadsheet and to add content to the cells. However, because this is a free plugin, other functionalities are limited and a number of manual modifications with JavaScript needed to be done.

The main modification was to ensure that when the user clicks outside of the spreadsheet, their selection is not erased, as is the default behavior of Jspspreadsheet CE, but it remains active. We accomplished this by storing which cells are selected on the spreadsheet when the user interacts with it, and setting that selection on the spreadsheet once the user clicks on another element of the page.

Converting the selected cells into a date or hour selection was done for the Week temporal filter by calculating from the columns which were selected which date range to query and from the rows which were selected which hours to query. For the Month temporal filter, which date ranges to query was inferred both from which rows and columns were selected.

Depending on whether the Photo distribution feature is activated, the cells of the Week and Month temporal filter either show how **if** there were photos taken during an hour or day, or how **many photos** were taken during an hour or day, respectively. As explained briefly in Section 3.3.2,

this information is stored in arrays in the database and retrieved via a back-end call. Each element in one of these arrays corresponds to a cell in the Week or Month date picker. If an array element is equal to 0, the corresponding cell is left blank. If an array element is higher than 0, a dot is placed in the corresponding cell, or if the Photo distribution feature is activated, the number is converted to a color intensity.

Custom: The custom date picker was implemented using basic HTML and JavaScript, as the drop-down input functionality is a standard functionality in HTML. To ensure that the user provides a valid date range, basic validation is performed while the user selects a date range with the drop-downs. A date selection can only be applied if the date range the user entered is valid.

Location and Tag filter

As explained briefly in Section 3.3.2, while the user is searching for a location or tag, a call is sent to the back-end which returns an array of locations or tags that contain their search query. These results are displayed to the user in a drop-down, so they may select which location or tag to apply. Before displaying the results to the user, the locations or tags are ranked according to relevance. We defined three levels of relevance: the location or tag **starts with** the user's query string, a **second or third word** of the location or tag **starts with** the user's query string, or finally, the location or tag **contains** the user's query string but does not start with it. The results are presented in a drop-down to the user in order of relevance, with the user's query string highlighted in each suggested location or tag. No libraries other than jQuery for the calls to the back-end are used for the Location and Tag search functionalities.

Detail view

When the user opens the Detail view by clicking on a photo on the map or in the *Photos in this area* feature, two calls are made to the back-end. The first call (Single photo search) returns all metadata associated with the photo that the user has clicked on, which is displayed in a table next to the photo. The second call (Day search) returns all photos that were taken on that day. These photos are added to the 'Explore this day' field, which is a scrollable field that is automatically centered on the photo that is currently displayed in the Detail view. The user may re-center the field by clicking a button, which returns the scroll-bar to the position of the active photo. No libraries other than jQuery for the calls to the back-end are used for the Detail view.

3.4 Research questions

After finalizing the design and implementation of our lifelog exploration system, we defined our research questions. As previously mentioned, our global research aim is to identify whether geospatial visualization and temporal filtering techniques are suitable for displaying lifelog data in such a way that it can be accessed, searched and explored. However, since this aim is broad and we cannot test all techniques, we focus on determining which of the geospatial visualization and temporal filtering techniques that we have chosen to incorporate into our lifelog system are the most useful for the different lifelog search goals that we presented in Section 3.1.1. Our research questions are as follows:

- **RQ1:** Out of a Dot, Heat, and Cluster geospatial visualization, which one is considered most useful for different search goals in the context of lifelog data?
- **RQ2:** Out of a color-coded Week matrix, color-coded Month matrix, and Custom drop-down temporal filtering technique, which filtering technique is considered most useful for different search goals in the context of lifelog data?

We believe that 'useful' can be described in a number of different ways. First and foremost, it is how users subjectively rate the usefulness of a feature. In addition to this, we believe a feature is useful if it helps a user achieve their search goal successfully and in an efficient way.

We also want to emphasize that the usefulness of a geospatial visualization or temporal filter is not something that is always standalone. It is likely that users may have a preference to use a certain geospatial visualization and temporal filter in combination with each other, which might enhance their usefulness. As an example, a lifelogger might use the Week or Month temporal filter in combination with the Heat map to identify when they tend to stay up particularly late and where they are if this occurs.

We also expect that features may be more useful to a user when used in a specific order. For example, when a lifelogger is looking for a photo of a specific event that took place in a certain year, they might first apply a temporal filter and a tag filter before interacting with the map to navigate to the location where the event occurred, as it might become easier to find the event on the map when they have other filters applied. Therefore, we will not only be focused on investigating the usefulness of different geospatial visualization and temporal filtering techniques in isolation, we are also interested in determining whether there is a certain interaction order between different parts of the interface that users consider most useful.

Chapter 4

User study

To answer our research questions and general research aim we defined in Chapter 3, we performed a user study with our lifelog system. In this chapter, we define the design of this user study. First, we describe the procedure of the experiment, then the materials we used to perform the experiments and finally, the participants who took part in the experiment.

To evaluate our system and its features, we would ideally let participants collect their own visual lifelog data which could then be made accessible through the system, in a similar process of the user study in the work of Chowdhury et al. [4]. In this scenario, the participants would explore their own lifelog data at their leisure and provide feedback on which system features they found most useful. However, a user study like this is not within our current scope. Instead, to replicate this scenario, we designed a number of questions about the LSC'20 test collection for participants to answer. These questions are intended to evoke search behavior in participants which is similar to lifeloggers exploring and searching their own lifelog data.

When we presented the geospatial visualizations and temporal filtering techniques which we selected for evaluation in Section 3.1.2, we expressed that we believe that each geospatial visualization or temporal filtering technique might be more useful for a different type of search goal. We consider it likely that therefore, a useful lifelog exploration system may incorporate multiple geospatial visualizations and/or temporal filtering techniques, to support every type of search goal a lifelogger might have. For this reason, in our user study, we chose not to compare the chosen geospatial visualizations and temporal filtering techniques to each other in isolation in a comparative study, but rather, to make all these techniques available in our system at the same time and ask participants to answer a number of exploration and retrieval questions about the LSC'20 test collection data, for which they were allowed to use any part of the system they believed would aid them best in their task. This strategy allows us to, in addition to determining which geospatial visualization and temporal filtering technique is the most useful, judge if there are benefits to implementing multiple geospatial visualizations and temporal filtering techniques into a single lifelog exploration system.

4.1 Procedure

During our experiment, participants were first asked for their informed consent and introduced to the system through a tutorial. Next, they were asked to answer a number of questions about the LSC'20 test collection data which they answered by searching the data with our system. Each of these questions was designed to replicate a certain lifelogger search goal and information need. Finally, participants were asked to provide feedback on the system as a whole through the System Usability Scale [1] and on its features separately through a number of Likert questions and open

questions.

The participant was guided through the experiment with a questionnaire, which contained the tutorial, the instructions for the experiment and the evaluation questions. For the tutorial, we opted to have the participants follow written instructions rather than explaining it to them verbally, as the tutorial described each system feature in detail and we wished to ensure each participant received the exact same instructions.

Since we thought it likely that participants might forget or would not see the need to switch geospatial visualizations while searching, we varied the order of which geospatial visualizations was presented as default to the participants. This was done in a Latin Square design, where in one third of the tests the order of the buttons to switch geospatial visualization was Dot-Cluster-Heat, one third was in the order Heat-Dot-Cluster, and one third was in the order Cluster-Heat-Dot. This means that upon starting the Lifelog Browser, one third of participants had the Dot map active by default, one third had the Heat map active by default and one third had the Cluster map active by default. This ensured that the Dot map was not favoured unfairly because participants were always presented with it first. We found it unnecessary to vary the order of the Date pickers in a similar way, as these were ordered from more low-level (Week) to high-level (Custom) and randomizing this order might have caused confusion amongst participants.

Our experiment was designed to be carried out in-person, where the experimenter would be present with the participant. However, due to the COVID-19 pandemic, this was not always possible. Therefore, a number of experiments were carried out remotely, where the experimenter was present with the participant through a digital meeting.

In the following sections, we explain the procedure of our experiment in more detail.

4.1.1 Introduction

In the introduction, the experiment was explained to the participant and they were asked to provide their informed consent. This explanation and consent form can be found in Appendix A.1. Next, participants were asked about their gender, age, digital literacy and previous experience with lifelog data. The exact questions that were asked can be found in Appendix A.2. We chose to collect data on the digital literacy on participants as we believe that the target audience for the Lifelog Browser (lifeloggers) is an audience with a higher digital literacy level, hence, we wanted to ensure that this was reflected in our participants.

After answering the demographic questions, the participants were introduced to the system through a written tutorial explaining all functionalities of the system. The tutorial covered all functionalities of the system and included a number of practice tasks which participants had to complete that were focused on the most important features of the system (map navigation, temporal and location filtering and the Detail view). The tutorial was split into four parts:

- **Overview:** The participants were shown a screenshot of the Lifelog Browser and all features (with the exception of the Detail view) were explained briefly.
- **Geospatial visualizations:** Each geospatial visualization was explained to the participant. Specifically, they were explained what each visualization shows and how to access the lifelog photos from the visualization. How the Detail view worked was also explained in this part of the tutorial. At the end of the explanation, participants were asked to view a photo with each geospatial visualization, each in a different city. They were also asked to open the Detail view and scroll to see what the lifelogger did that day.

- **Temporal filters:** Each temporal filter was explained to the participant. They were explained how to select dates, days or times with each filter and, for the Week and Month temporal filters, the functionality of the Photo distribution checkbox. At the end of the explanation, they were asked to select a certain time-frame with each temporal filter.
- **Location and Tag filters:** The participants were explained how to add and remove locations and tags with the Location and Tag filters. They were also informed about which results to expect when adding multiple locations or tags. At the end of the explanation, they were asked to find photos with certain locations and tags.

Once we confirmed that participants understood the core features of the system, they were introduced to the next part of the experiment, Search.

4.1.2 Search

We wished to evaluate how well suited our system is for different lifelog search goals, specifically Directed closed search, Directed open search, Undirected search and Casual-leisure search. We established previously that these search goals can be subdivided into *exploration* (Directed open search, Undirected search and Casual-leisure search) and *retrieval* (Directed closed search). For this reason, we subdivided the Search phase of our experiment into two different phases, Exploration and Retrieval. In the Exploration phase, participants answered a number of questions about the data which were designed to stimulate them to explore the data. In the Retrieval phase, participants were asked to answer a number of Search topics from the LSC'20 test collection. Participants were instructed to use whichever part of the system that they felt would best aid them to answer the questions they were asked.

During both phases, participants were asked to use the think-aloud protocol, which means they were asked to verbally describe their search process and voice any thoughts that came to their mind while searching. Their comments and search actions were recorded through voice and screen recording, to later review their opinions and search strategies. During search, the experimenter also wrote down which elements the participant interacted with to solve a question and any comments of particular note. To determine which features participants used most during the Search process, we recorded per feature how much time participants spent interacting with it. The system records in the background per feature how much time a user spends hovering over it. Even though this is only an approximation of how much time participants truly interact with a feature, we could use this to identify which features were used most.

The Exploration and Retrieval parts of the Search phase of the experiment are outlined below.

Exploration

As previously defined, a lifelogger accessing their lifelog data for exploration purposes can have three different search goals; Directed open search, Undirected search and Casual-leisure search. Per search goal we designed two questions which we felt represented a realistic lifelogger information need, resulting in six different exploratory questions. Below, we present these exploratory questions. For each question, we describe how we believe they tie in with a lifelogger search goal, how they reflect a realistic information need a lifelogger might have and finally, the search strategy we expect participants to employ to answer this question. Note that the questions below are not grouped by search goal, rather, they are organized from higher level information needs to lower level information needs, in such a way that the participant becomes increasingly familiar with the system while answering the questions. To more easily discern the results for the different questions in following

chapters, we have named the questions after their question type (the **E** denotes exploration) and the order in which they appeared to the participant.

- **E1:** *Name 3 cities outside of Ireland that the lifelogger visited.* This question represents a *Casual leisure* search goal, as one of the most apparent casual browsing queries is to see which locations a dataset covers. We imagine that a lifelogger could access their lifelog data to simply take a look at where they traveled during, for example, the past year. Our expectation for this question is that participants use any geospatial visualization they feel most comfortable with to navigate away from Dublin, where the lifelogger lives, and that they then look for locations where more photos are taken.
- **E2:** *Did the lifelogger travel to China in winter?* This question also represents a *Casual leisure* search goal, as this could be an exploratory query that a lifelogger might place after recalling that, for example, they traveled to China multiple times, but don't remember the details of the trip. We expect that participants will navigate to China using the geospatial visualization of their choice and then use the *Photos in this area* feature to obtain an overview of the dates that the lifelogger stayed in China. Alternatively, we think they could use the Month date picker to see if the dataset contains any data which was recorded during winter months and then, if any data was found, select a time range to confirm whether or not the photos were taken in China.
- **E3:** *If the lifelogger stays up late (past 23:00), what does he typically do?* This question represents an *Undirected* search goal, as the goal is to explore the photos of a certain event (in this case, staying up late). A lifelogger might perform a similar query when, for example, they notice in their data a certain week where they consistently stayed up late and they want remember why they did so. We expect participants to use any of the temporal filters to select a time past 23:00 and use the *Photos in this area* feature and detail view to determine what activities the lifelogger engaged in during that time.
- **E4:** *How many days did the lifelogger spend in Norway?* This question represents a *Directed, open* search goal, as the goal is to find an undetermined number of photos of a certain location. A lifelogger might perform a similar query when they want to remember how long a certain trip lasted. To answer this question, we expect participants to use any geospatial visualization to navigate to Norway and read the text in the bottom-left of the map, which informs the user at which dates the photos that are displayed on the map were taken.
- **E5:** *Name 3 places the lifelogger ate while in Tromsø.* This question represents an *Undirected* search goal, as participants are asked to explore specific photos from a certain location. A lifelogger could place a similar query when they are exploring what they did during a trip. We expect participants to navigate to Tromsø using any geospatial visualization and then use the Tag filter to select photos where the lifelogger is eating. To inspect photos in more detail, we expect them to use the Detail view.
- **E6:** *The data shows that the lifelogger has visited 3 different McDonald's in the area of Dublin. At which McDonald's (north, south or west) has the lifelogger spent the most time?* This question represents a *Directed, open* search goal, as participants are looking to find an undetermined number of photos that fit this query. A lifelogger could place a similar query when they are, for example, looking for the restaurant where they spend the most time. To solve this question, we expect participants to add all McDonald's locations in the Location filter and then navigate to Dublin with a geospatial visualization that allows them to see the number of photos taken at a location, such as the Cluster or Heat map.

Since the goal during the Exploration phase is not for participants to find the answer to questions as quickly as possible, but rather, to instigate them to explore the data, they were reminded that

they should focus more on getting an idea of what the data looks like rather than answering the questions as quickly as possible. To ensure that the experiment would not run unnecessarily long, participants were asked after about three minutes if they wanted to continue to the next question, or to search for a while longer. After another minute, they were asked to continue to the next question.

Retrieval

Lifeloggers may also access their lifelog data with the goal to find a specific photo, in which case they have a Directed, closed search goal. To validate whether our system and its geospatial visualization and temporal filtering techniques can also support this search goal, we chose three search topics from the LSC'20 test collection for participants to answer. We have chosen to incorporate three LSC'20 search topics into the experiment because we wish to test various levels of difficulty. The topics we have selected can be found below. To more easily discern the results for the different questions in following chapters, we have named the questions after their question type (the **R** denotes retrieval) and the order in which they appeared to the participant. These specific topics were chosen because we believe that each of them can be answered using a different search strategy, which can provide us with information on which search strategies our system is suitable for.

- **R1:** *Find one of the photos that matches the following description: It was the best cake I had in years, in an antiques store. I was alone drinking tea and eating cake. I think I finished all the cake in 3 minutes. It was in the UK on a Saturday morning.* We chose this search topic because it is relatively simple to solve, which makes it a good introduction to the Retrieval phase of the experiment. Even though the photo is most easily found using a tag (few photos contain the *cake* tag), participants may still use a temporal filter to narrow down their query.
- **R2:** *Find one of the photos that matches the following description: There was a Chandelier (large glass light) hanging from the ceiling, and red chairs. Actually there were lots of chandeliers. It was just after eating pizza. It was only across the street from the restaurant in Dublin, in the Science Gallery at Trinity College. It was in 2015 on a Tuesday.* We chose this search topic because the system does not contain tags that clearly describe the photo, so it is not easy to find it by tag alone. The location information provided in the description can also only be used as a guide, since the photo was not taken at Trinity College but nearby, so this only serves as an indication of in which area to look. The main strategy we expect participants to use to find this image is to filter by time and tag to find the event that took place **before** the photo was taken (having pizza). We therefore expect participants to use the Detail view in order to solve this search topic.
- **R3:** *Find one of the photos that matches the following description: The Red House. I remember there were lots of cars that day and the weather was very nice with a beautiful blue sky. It was just after eating an ice cream by the sea. Before I went home in someone else's car. It was in 2018. On a Monday evening.* This search topic requires creativity to solve, as alternate tags (ocean instead of sea, sunny instead of sky) must be provided in order to find the image. However, if the participants manage this, the photo is relatively easy to find.

Participants were informed that during the Retrieval phase, their goal was to find the answer as quickly as possible. They were given three minutes per question to find the photo. Unlike how the search topics are presented to the participants during the Lifelog Search Challenge, where a new part of the description is displayed to the participant every thirty seconds, we chose to display the entire description to the participant immediately. This way, the participants have more freedom to choose a search strategy that they feel will be most effective, and we hope to be able to more easily discern patterns in which search strategies participants employ.

4.1.3 Evaluation

After completing the Exploration and Search phases, the participants are asked to rate the system as a whole and specific features of the interface separately. Below, we outline how participants were asked to provide their ratings for our system.

System Usability Scale: First, the participants were asked to fill in the System Usability Scale (SUS) [1]. We chose this tool as it is short and reliable. We believe that it is important to obtain an overall usability score for the system, to ensure the quality of our other results. For example, if we find that our usability score is low, our results on the usefulness of the geospatial visualization and temporal filtering techniques might be too influenced by other problems with our system to draw a clear conclusion.

Interface feature ratings: Next, participants were asked to give ratings for three types of features: the geospatial visualizations (Dot, Cluster and Heat), the temporal filters (Week, Week with Photo distribution, Month, Month with Photo distribution, and Custom) and the Location and Tag filters. For each of the features, participants were asked how **easy to use** they were and how **useful** they found them. For ease of use, participants could choose between the options *very difficult*, *somewhat difficult*, *neither difficult nor easy*, *somewhat easy* and *very easy*. For usefulness, participants could choose between the options *very useless*, *somewhat useless*, *neither useless nor useful*, *somewhat useful* and *very useful*. Participants were also given the option to fill in that they did not use a feature, if they found they had not used it enough to provide their opinion on it. Then, they were asked to provide positive and negative comments on each of the features.

Overall: Finally, the participants were asked if they had any comments on the system as a whole. If needed, after providing these comments, the experimenter asked for clarification on participants' comments. The experimenter also asked for other ideas for improvements, or entirely new features, that participants would have liked to see in the system.

4.2 Materials

For both the in-person and remote execution of the experiment, the Lifelog Browser was run in a Chrome browser on a Windows 10 device with an AMD Ryzen 7 3700X 8-Core Processor (3.60 GHz) and 32 Gigabytes of RAM. We used Open Broadcaster Software (OBS)¹ for screen recording and pen and paper to record which system features the participants interacted with per question and their most notable comments. The questionnaire which was used to guide participants through the experiment and record their answers was created in Qualtrics². However, because Qualtrics experienced a temporary outage while we were performing the experiments, a number of responses were also recorded through Google Forms³, where we temporarily replicated original questionnaire.

For in-person experiments we used the Voice Recorder app⁴ on Android to record participants' comments. Our in-person setup included two screens, a main screen on which the participants could interact with the Lifelog Browser and another screen to the side, on which participants could read the questionnaire.

For experiments that were conducted remotely, we used Microsoft Teams⁵ for video conferencing

¹<https://obsproject.com/>

²<https://www.qualtrics.com/>

³<https://www.google.com/forms/>

⁴<https://play.google.com/store/apps/details?id=com.media.bestrecorder.audiorecorder>

⁵<https://teams.microsoft.com/>

and recording. Since the Lifelog Browser only ran locally, TeamViewer⁶ was used so the participants could interact with the Lifelog Browser as if it were running on their own device. Participants ran TeamViewer in full screen mode on their own laptop or desktop and accessed the questionnaire through another screen, laptop, desktop or on their mobile phone.

4.3 Participants

Partly due to COVID-19 and partly due to the length of the experiment (often exceeding 60 minutes), participants were recruited through convenience sampling. However, we did focus on including participants who were reasonably skilled with technology, to more accurately reflect the average lifelogger. We did not specifically select participants which had previous experience with lifelogging, even though this would be optimal for our study. Participants were not compensated for their participation.

In total, 21 participants took part in the experiment, of which 14 were male and 7 were female. 7 participants were aged between 18-24 and 14 participants were aged between 25-34. 1 participant rated their digital literacy to be poor, 3 participants rated it to be neutral, 8 participants rated it to be good and 9 participants rated it to be very good. 17 participants indicated that they had no prior experience with lifelog data, 3 participants indicated that they had a little experience with lifelog data and 1 participant indicated they had a moderate amount of experience with lifelog data.

⁶<https://www.teamviewer.com/>

Chapter 5

Results and Discussion

In this chapter, we present the results that were obtained from our user study and discuss them in respect to our general research aim and research questions. Finally, we discuss limitations of our study and suggestions of how our findings can be expanded upon with future work.

5.1 Results

5.1.1 Quantitative results

We have obtained five different quantitative data points from our user study that we would like to present here. First, we present our System Usability Scale score, as this is the only quantitative data point we have collected that describes our system as a whole. Then, we present the ratings participants have provided for the different features of our system and how much time each participant has used each system feature during the Search phase of the experiment. Finally, we present per question how many participants answered it correctly, in conjunction with which system features participants used to answer each question.

System Usability Scale

On average, our system received a SUS score of 70.60, which is slightly above the average score of 68. An overview on the responses per question can be seen in Figure 5.1.

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

Table 5.1: SUS statements. Participants were asked to rate each statement on a scale ranging from Strongly disagree to Strongly agree.




Figure 5.1: SUS scores per statement. An overview of the statements can be found in Table 5.1.

Feature ratings

Below, we present an overview of the ratings that participants provided on the geospatial visualizations, temporal filtering techniques, and Location and Tag filter features of our system. The participants were asked to rate each feature on how easy they found it to use (*ease of use*) and how useful they found it (*usefulness*). The ratings were converted to a scale from 1 to 5, where 1 corresponds to *very difficult to use* or *very useless* and 5 corresponds to *very easy to use* or *very useful*.

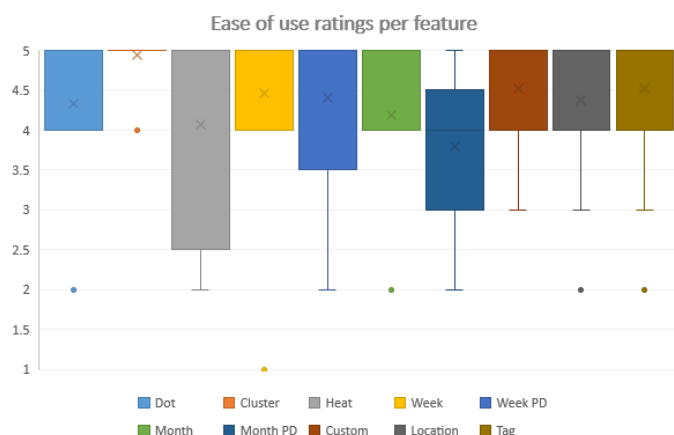
A visual summary of the results is presented in Figure 5.2. Since participants were not forced to rate any feature that they believed they did not interact with enough to provide an opinion on it, some features only have a very limited number of data points. Since our total of 21 participants is already quite limited for statistical testing, we believe there is no added value in performing statistical tests on this data. We will therefore simply report our results.

Dot map: For *ease of use*, 15 participants gave the Dot map an average rating of 4.33 with a SD (standard deviation) of 0.789. 6 participants indicated that they did not use the dot map enough to provide an opinion. For *usefulness*, 17 participants gave the Dot map an average rating of 4.00 with a SD of 1.085. 4 participants indicated that they did not use the Dot map enough to provide an opinion.

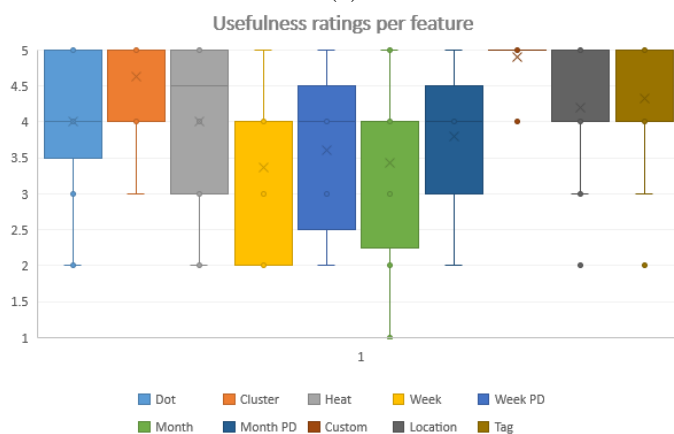
Cluster map: For *ease of use*, 19 participants gave the Cluster map an average rating of 4.95 with a SD of 0.223. 2 participants indicated that they did not use the Cluster map enough to provide an opinion. For *usefulness*, 19 participants gave the Cluster map an average rating of 4.63 with a SD of 0.581. 2 participants indicated that they did not use the Cluster map enough to provide an opinion.

Heat map: For *ease of use*, 14 participants gave the Heat map an average rating of 4.07 with a SD of 1.269. 7 participants indicated that they did not use the Heat map enough to provide an opinion. For *usefulness*, 16 participants gave the Heat map an average rating of 4.00 with a SD of 1.173. 5 participants indicated that they did not use the Heat map enough to provide an opinion.

Week temporal filter: For *ease of use*, 11 participants gave the Week temporal filter an average rating of 4.45 with a SD of 1.157. 10 participants indicated that they did not use the Week temporal filter enough to provide an opinion. For *usefulness*, 11 participants gave the Week temporal filter an average rating of 3.36 with a SD of 0.980. 10 participants indicated that they did not use the Week



(a)



(b)

Figure 5.2: Box plots for ease of use ratings (a) and usefulness ratings (b). PD stands for Photo distribution.

temporal filter enough to provide an opinion.

Week temporal filter with Photo distribution: For *ease of use*, 5 participants gave the Week temporal filter with Photo distribution an average rating of 4.40 with a SD of 1.200. 16 participants indicated that they did not use the Week temporal filter with Photo distribution enough to provide an opinion. For *usefulness*, 5 participants gave the Week temporal filter with Photo distribution an average rating of 3.60 with a SD of 1.020. 16 participants indicated that they did not use the Week temporal filter with Photo distribution enough to provide an opinion.

Month temporal filter: For *ease of use*, 16 participants gave the Month temporal filter an average rating of 4.19 with a SD of 0.726. 5 participants indicated that they did not use the Month temporal filter enough to provide an opinion. For *usefulness*, 16 participants gave the Month temporal filter an average rating of 3.44 with a SD of 1.273. 5 participants indicated that they did not use the Month temporal filter enough to provide an opinion.

Month temporal filter with Photo distribution: For *ease of use*, 5 participants gave the

Month temporal filter with Photo distribution an average rating of 3.80 with a SD of 0.980. 16 participants indicated that they did not use the Month temporal filter with Photo distribution enough to provide an opinion. For *usefulness*, 5 participants gave the Month temporal filter with Photo distribution an average rating of 3.80 with a SD of 0.980. 16 participants indicated that they did not use the Month temporal filter with Photo distribution enough to provide an opinion.

Custom temporal filter: For *ease of use*, 21 participants gave the Custom temporal filter an average rating of 4.52 with a SD of 0.587. No participants indicated that they did not use the Custom temporal filter. For *usefulness*, 21 participants gave the Custom temporal filter an average rating of 4.90 with a SD of 0.294. No participants indicated that they did not use the Custom temporal filter.

Location filter: For *ease of use*, 21 participants gave the Location filter an average rating of 4.38 with a SD of 0.844. No participants indicated that they did not use the Location filter enough to provide an opinion. For *usefulness*, 21 participants gave the Location filter an average rating of 4.19 with a SD of 0.794. No participants indicated that they did not use the Location filter enough to provide an opinion.

Tag filter: For *ease of use*, 21 participants gave the Tag filter an average rating of 4.52 with a SD of 0.794. No participants indicated that they did not use the Tag filter enough to provide an opinion. For *usefulness*, 21 participants gave the Tag filter an average rating of 4.33 with a SD of 0.777. No participants indicated that they did not use the Tag filter enough to provide an opinion.

Feature use time

For each feature, we recorded how much time participants spent interacting with it (hovering over it with their mouse). Below, we present our results per feature group (Map, Temporal filter and Other) and, if applicable, our statistical test to determine whether any feature within those feature groups was used significantly more than another of the same feature group. The average use time per feature is pictured in Figure 5.3 and the results per participant are pictured in Figure 5.4.

Map: The Dot map was used for 154.95 seconds on average, with a SD of 178.95. The Cluster map was used for 220.48 seconds on average, with a SD of 159.51. The Heat map was used for 83.67 on average, with a SD of 115.40. A Shapiro-Wilk test revealed that only the data for the Cluster map was normally distributed. Therefore, to test whether any of the map visualizations was used significantly more than the others, we performed a Mann-Whitney U test between each combination of geospatial visualizations. The test between the Dot map and the Cluster map determined that the Cluster map **was not** used significantly more than the Dot map ($p=0.1069$). The test between the Dot map and the Heat map determined that the Dot map **was not** used significantly more than the Heat map ($p=0.1714$). Finally, the test between the Cluster map and the Heat map determined that the Cluster map **was** used significantly more than the Heat map ($p=0.0030$).

Temporal filter: As seen in Figure 5.3, the Week temporal filter with Photo distribution and the Month temporal filter with Photo distribution were hardly used. The Week temporal filter with Photo distribution was used for only 3.19 seconds on average with a SD of 11.86 and the Month temporal filter with Photo distribution was used for only 7.00 seconds on average with a SD of 31.30. We have therefore decided to include these use times with the Week and Month temporal filters, respectively. In total, the Week temporal filter was used for 31.14 seconds on average with a SD of 22.44. The Month temporal filter was used in total for 76.62 seconds on average with a SD of 75.14. The Custom temporal filter was used for 140.57 seconds on average with a SD of 54.06. A Shapiro-Wilk test revealed that the data for the Month date picker was not normally distributed.

Therefore, to test whether any of the temporal filters was used more than the others, we performed a Mann-Whitney U test between each of the combinations of temporal filters. The test between the Week temporal filter and Month temporal filter determined that the Month temporal filter **was** used significantly more than the Week temporal filter ($p=0.0391$). The test between the Week temporal filter and Custom temporal filter determined that the Custom temporal filter **was** used significantly more than the Week temporal filter ($p=1.466e-7$). Finally, the test between the Month temporal filter and Custom temporal filter determined that the Custom temporal filter **was** used significantly more than the Month temporal filter ($p=0.0008$).

Other: Since we are not comparing the Location and Tag filters, Photos in this area feature, and the Detail view amongst each other, we will simply report our results here. The Location filter was used for 61.24 seconds on average with a SD of 21.47. The Tag filter was used for 65.76 seconds on average with a SD of 28.15. The 'Photos in this area'-feature was used for 75.14 seconds on average with a SD of 53.82. Finally, the Detail view was used for 115.71 seconds on average with a SD of 49.47.

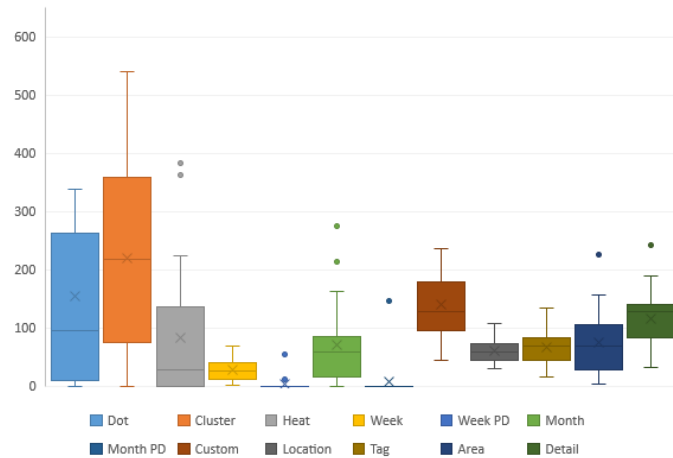


Figure 5.3: Box plot of the average use time (in seconds) per feature. Note that one outlier (of the Dot map) is not pictured in the plot.

Correctly answered questions

Even though for the Exploration questions, the main goal was not for participants to answer them correctly, but rather, to stimulate participants to explore the data, we still wish to report how participants performed per question. Combined with which features participants used to answer questions, we can use this data to determine if there are any search strategies which can be considered more successful.

Below, per question, we present how many participants answered a question correctly, partly correctly, incorrectly or did not provide an answer for a question. We also define how we determined correctness for each of the questions, as for the exploration questions this is not always straightforward. Note that for the retrieval questions, we have repeated only a short description of the question rather than the entire question, which can be found in Section 4.1.2.

E1: *Name 3 cities outside of Ireland that the lifelogger visited.* 16 participants answered this question correctly and 5 participants answered this question partly correctly. We only considered




Figure 5.4: Plot of each participant’s feature use time (in seconds).

places where the lifelogger stayed for an extended amount of time (i.e. did not only fly over with the plane) to be correct.

E2: *Did the lifelogger travel to China in winter?* 14 participants answered this question correctly, 5 answered incorrectly and 2 participants did not answer this question. The participants were informed they had to look for photos in the months December, January and February.

E3: *If the lifelogger stays up late (past 23:00), what does he typically do?* 16 participants answered this question correctly and 5 participants answered this question partly correctly. We did not provide a specific criteria of what ‘typically’ meant in this context. If it is an activity that the lifelogger ever performed after 23:00, we considered it to be correct.

E4: *How many days did the lifelogger spend in Norway?* 8 participants answered this question correctly, 8 answered incorrectly and 5 participants did not answer this question. Even though the correct answer was 13 days, we counted answers ranging from 10 to 13 days as correct, as some participants did use the right approach to answer the question but made a mathematical mistake while calculating the exact number of days.

E5: *Name 3 places the lifelogger ate while in Tromsø.* 16 participants answered this question correctly and 5 participants answered this question partly correctly. Both the type of restaurant and the name of the restaurant were considered correct answers for this question.

E6: *The data shows that the lifelogger has visited 3 different McDonald’s in the area of Dublin. At which McDonald’s (north, south or west) has the lifelogger spent the most time?* 10 participants answered this question correctly, 7 answered incorrectly and 4 participants did not answer this question. A number of participants correctly guessed south, as that was the only McDonald’s location they were able to find.

R1: *Find the photo where the lifelogger is eating cake on a Saturday morning in the UK.* 20 participants answered this question correctly and 1 participant did not answer this question, which means they were not able to find the photo within the given time.

R2: *Find the photo with red chairs and chandeliers, just after having pizza.* 10 participants

	Correct	Partly	Incorrect	Unanswered
E1	16	5	0	0
E2	14	0	5	2
E3	16	5	0	0
E4	8	0	8	5
E5	16	5	0	0
E6	10	0	7	4
R1	20	0	0	1
R2	10	0	0	11
R3	4	0	0	17

Table 5.2: Overview of how many participants answered each question correctly, partly correctly, incorrectly or did not provide an answer.

answered this question correctly and 11 participants did not answer this question, which means they were not able to find the photo within the given time.

R3: *Find the photo of the Red House which was taken after having ice cream by the sea.* 4 participants answered this question correctly and 17 participants did not answer this question, which means they were not able to find the photo within the given time.

Feature use per question

Per question per participant, we wrote down which features the participant used or attempted to use in order to solve the questions. We have also recorded in which order they used these features (repeated use was not recorded). Table 5.3 shows, for each participant, a subtable which describes in which order the participant used the system features per question. The use order is denoted by a number, where the feature that was used **first** is marked with a 1 in the table, the feature that was used second is marked with a 2, and so on, up until 6. We chose not to include features that were used after that, as those were often only used by participants in a final attempt to, often unsuccessfully, obtain the answer. The system features are indicated by an abbreviation, Md for Dot map, Mc for Cluster map, Mh for Heat map, Pw for Week temporal filter (the P denotes *picker*), Pm for Month temporal filter, Pc for Custom temporal filter, L for location filter, T for tag filter, A for *Photos in this area* and D for Detail view. The bottom row of each table shows the total number of times participants have used a type of feature (Map, Temporal filter or Other) **first** in their search process.

The tables are ordered by how well participants performed, which can be derived from the color-coding in the first column of each subtable. Each cell in the first column has a color which denotes how participants performed on that question. If participants answered a question correctly the cell is green, if they answered partly correctly it is yellow, if they answered incorrectly it is orange and if they did not provide an answer it is gray. The ranking is determined first by how many questions participants answered correctly, then questions that were answered partly correctly, then incorrectly and finally, unanswered questions.

Table 5.4 provides a similar overview to Table 5.3, but **per question**. The penultimate row of these tables show how often a specific feature was used first in participants' search process, and the final row of these tables show how often a certain feature type was used first. The order in which the participants are listed in the table is determined by which feature they used first, in order to more easily discern patterns.

5.1.2 Qualitative results

For each feature participants were asked to rate for *ease of use* and *usefulness*, they were also asked to provide positive and negative comments. Finally, they were asked to provide comments on the system as a whole. Here, we present a summary of the positive and negative comments that participants provided per feature and general comments per feature type. Additionally, we present a number of observations we made about how participants interacted with our system. A complete overview of the comments that were given can be found in Appendix B.

Map

Dot map: Comments for the Dot map were approximately equally distributed between positive and negative comments. Participants found that the Dot map was useful for tasks for which they needed a more detailed overview of the data, such as finding travel routes or finding a specific photo on a lower zoom level. On a higher zoom level, however, participants found the Dot map too cluttered and overwhelming. Table B.1 shows all comments participants provided for the Dot map.

Cluster map: The Cluster map received predominantly positive comments, which noted how easy to use participants found it. They found that it was a good combination of the functionalities of the Dot map and the Heat map and that it provides a clear overview of the data. Only a few participants provided negative comments about the Cluster map. These comments stated that too many clicks were needed to open photos and that the areas that the clusters covered were too random. Table B.2 shows all comments participants provided for the Cluster map.

Heat map: Comments on the Heat map were mixed; it received a number of positive comments but a larger number of negative comments. Participants appreciated that the Heat map provided them with an overview, which they considered to be useful when lots of photos were displayed on the map. However, they found that it was sometimes difficult to access the photos and that the colors were not always visible or clear enough, both when zoomed in or zoomed out far. Table B.3 shows all comments participants provided for the Heat map.

General: Comments on the map interface in general were also mixed. A number of participants mentioned that all geospatial visualizations were useful, albeit on different zoom levels. Generally, the map interface was also described as intuitive and useful. However, many participants noted that they found the map to be slow and sometimes unresponsive, which made it feel unreliable to them. They also expressed that they found the cap of 40,000 photos to be bothersome, as they did not feel it was enough to get an adequate overview when viewing a larger area. Table B.4 shows all comments about the map interface in general.

Temporal filters

For the temporal filters, we combined the comments for the Week and Month temporal filters into one overview, as they were very similar. We show the comments on the Photo distribution feature separately, as these can be interpreted separately from the Week and Month temporal filters.

Week and Month: Not many participants commented on the Week and Month temporal filters. A few participants indicated that they liked the styling of these filters and the additional information that was provided about the temporal distribution of the data. However, participants also indicated that they found the Week and Month temporal filters to be too limited and sometimes even confusing or unclear. Many participants commented that they did not consider the Week and Month temporal filters to have added value over the Custom date picker. All comments for the

Week and Month temporal filters together can be found in Table B.5.

Photo distribution: Few participants commented on the Photo distribution feature. Even though the majority of the participants who commented on this feature indicated that they appreciated the feature and thought it had the potential to be very useful, they also stated that they did not see a use for it to answer the questions they were asked in the experiment. However, for the majority of participants who commented, the functionality of the Photo distribution was clear and they imagine it could be useful for other search tasks. All comments on the Photo distribution can be found in Table B.6.

Custom: Many participants commented on the Custom temporal filter and comments were largely positive. Participants found the Custom temporal filter to be by far the most useful, easiest to use and fastest to use. They also found that just having the Custom temporal filter for temporal filtering would provide enough functionality for them. Negative comments were mostly related to the number of clicks that were needed to provide a time range, especially when selecting a year. They also expressed that they would like buttons to filter more quickly by month or year. All comments on the Custom temporal filter can be found in Table B.7.

General: Few positive comments were provided by participants on the temporal filters in general. A number of participants appreciated that all temporal filter options were intuitive and easy to use and that they enjoyed how filtering affected the data that was shown on the map. More negative comments were left by participants, where many found that it was not necessary to have all temporal filters available and that having only one (preferably Custom) would suffice. Many participants also expressed that they struggled with remembering to apply after changing or resetting something in the temporal filters. All comments on the temporal filters in general can be found in Table B.8.

Location and tag

Location: Few participants provided positive comments on the Location filter. The only positive comment we received was that participants appreciated that it quickly allowed them to find what they were looking for, if they knew exactly what to look for. More negative comments were provided on the Location filter, the most occurring one being that participants would like to be able to filter by country and city using the Location filter and that the map should move immediately to a location once it has been added. They also expressed that they disliked that there were duplicate locations and that locations were very limited. Table B.9 shows all comments participants provided on the Location filter.

Tag: More participants provided positive comments on the Tag filter. They appreciated that it was easy to find the tags they were looking for and that there was a large variety of them. However, a large number of participants also complained that they found the tags very limited and especially that the tags did not describe the photos accurately. They also complained that adding a large number of tags often quickly filtered out all the data, with no clear indication of which tag caused the lack of result. Table B.10 shows all comments participants provided on the Tag filter.

General: With both the Location and the Tag filters, participants appreciated that they were easy to use and understand, and that the drop-down with suggestions was very helpful. However, some participants complained that the 'reset' button was more visible than the 'add' button, which confused them. A number of participants also found it confusing that adding multiple locations and tags affected the data differently. Table B.11 shows all comments participants provided on the Location and Tag filters.

Overall

Participants also gave a number of general comments on the system that did not specifically relate to any of the features discussed above. Many participants commented that they enjoyed using the system and that generally, they found it easy to use. They also appreciated the styling of the system. However, common complaints were that the system was slow and that they would have liked a more seamless experience when filtering, with the filters being applied automatically upon adding them rather than having to click the 'apply' button. Some participants also found there were too many options. All comments on the system in general can be found in Table B.12.

Suggestions

Next, we present a selection of the most occurring and most interesting suggestions that participants have provided for the system. We have subdivided these suggestions into either improvements of the features that are currently in the system, or suggestions for entirely new features.

Improvements: Participants suggested a number of improvements for the Detail view. Some wished to see the photo and the *Explore this day* feature in the Detail view bigger. Another participant suggested to move the *Explore this day* feature to the main view of the system, close to the *Photos in this area* feature, so it would be less easy to confuse them. Participants also suggested a number of ideas for filtering photos within an opened cluster, so it would become easier to find the photo they were looking for in a large cluster. Other interesting suggestions were to ensure that the dots or Photo distribution in the Week and Month temporal filters corresponded to the data shown on the map, or to color-code the dots on the Dot map by the date that the photo was taken. Table B.13 shows all suggestions for improvement.

New features: Participants also suggested a number of entirely new features that they would have liked to see in the system. A very valuable suggestion was to transform the Week or Month temporal filter into more of a visualization which appears after filtering the data to inform the user on how the data is distributed temporally. Some participants suggested that they would have liked the Detail view to have a map feature, as they struggled with the fact that they could not see the locations of the photos they viewed in the Detail view. Small suggestions for new features were to have a favorite functionality for photos or locations, or to have a random explore mode. Table B.14 shows all suggestions for new features.

Observations

Finally, we provide a number of observations that we made during the tests. These observations are findings that we believe are not yet reflected clearly enough in the comments above. The most important observations include how participants struggled with remembering to press the 'apply' button after selecting a filter and particularly after resetting. The Week and Month temporal filters were not always clear to participants, as they did not always understand how the dots related to the map and that to filter, they needed to select a date range in the spreadsheet rather than it happening automatically. A number of participants also did not seem to understand the difference between *Photos in this area* and *Explore this day*. A complete overview of observations can be found in Table B.15.

5.2 Discussion

Our study aims to determine whether geospatial visualization and temporal filtering techniques are suitable for displaying lifelog data in such a way that it can be accessed, searched and explored. Specifically, we wished to determine which geospatial visualization, out of a Dot, Cluster and Heat map, and which temporal filter technique, out of a Week, Month and Custom temporal filter is considered most useful for different search goals in the context of lifelog data. In this section, we analyze our results in order to answer these questions.

We start with analyzing our System Usability Scale results, as these results potentially influence our more in-depth results on the different system features. Then, we analyze our system’s suitability in relation to each search goal as presented in Section 3.1.1. Next, we briefly review the most successful search strategies participants used to answer questions. We then attempt to answer our specific research questions on the geospatial visualization and temporal filtering techniques. Finally, using all findings presented in this section, we discuss the overall suitability of our system for lifelog access, search and exploration.

5.2.1 Usability

As reported in Section 5.1.1, our system has received an average SUS-score of 70.60, which is slightly higher than the average of 68. This is a reasonably good score, which likely indicates that our other results were not significantly influenced by any major usability issues. However, there is room for improvement.

Figure 5.1 shows that participants generally responded somewhat in favor of the system, but they rarely indicated that they strongly agreed or disagreed with statements in favor of the system. Mainly, participants indicated that they found the system easy to use (statement 3) and that they imagined others would learn to use the system quickly (statement 7), that they did not imagine they needed any technical support to be able to use the system (statement 4) and that they did not feel like they had to learn too many things to get started (statement 10). These results suggest that our interface is likely well-designed and it is quickly clear to users how to use each functionality. This is supported by our other data, such as comments from the participants, in which we observe that on most features at least a number of participants remarked that they consider them intuitive or easy to use. *Ease of use* ratings for features were also high (above 4.00) , with the exception of the Month date picker with Photo distribution, which received a rating of 3.80.

Even though ratings for the remaining SUS statements were still predominantly positive, we see here that participants voted somewhat less in favor of our system. We see this most strongly in statement 2, which gauges if the participant found the system unnecessarily complex, and statement 6, which gauges if the participant thought there was too much inconsistency in the system. For both of these statements, more participants gravitated towards neutral. We believe this tendency may be caused by the large number of functionalities available in the system. As a number of participants commented, they found the number of filter choices, especially for the temporal filters, too overwhelming. Possibly, this explains their more neutral response to the statement *I thought there was too much inconsistency in this system*.

Finally, we want to note that overall, participants found the SUS statements clear and easy to respond to, but some struggled to respond to the first statement, *I think that I would like to use this system frequently*, as they did not see a situation in which they would ever use the system, since they did not have any lifelog data. The participants therefore were instructed to answer the question as if they did have lifelog data that they wished to access. However, a number of participants expressed

a deep concern with the privacy sensitive nature of lifelog data and therefore answered this question with 'strongly disagree' or 'somewhat disagree', as they could not imagine themselves ever having any kind of lifelog data. This emphasizes the importance of recruiting participants who are lifeloggers for the evaluation of lifelog systems, as this issue would not have come up had the tests been performed with lifeloggers as participants.

5.2.2 Suitability per Search goal

In this section, we discuss the suitability of our system for lifelog access, search and exploration in relation to each lifelog search goal: Directed closed search, Directed open search, Undirected search and Casual leisure search.

Directed, closed search

The questions which corresponded to this search goal were **R1**, **R2** and **R3**. In Table 5.2, we observe that R1 was answered correctly by almost all participants, R2 was answered correctly by approximately half of the participants and R3 was answered correctly by only a fraction of the participants. We believe that participants struggled to answer R2 and R3 due to limitations of the Tag filter. Many participants were focused on applying the tags that they thought best described the photo they were looking for, but found that either the desired tags were not available or that they did not describe the photos well. However, the results of R1 show that when a tag (*cake*) is available in the system and it describes the photo it is attached to well, participants have little trouble finding a specific photo.

The results of questions R2 and R3 also showed that the Detail view allowed participants to find consecutive events. However, not all participants realized from the description that, for example, the *pizza* event in R2 took place **before** the photo they were trying to find was taken. We consider it likely that if more participants had been aware of this, more participants would have found this photo.

We conclude that, even though our system was built mostly with lifelog exploration in mind, it is also suitable for retrieval purposes. The main limiting factor in the system's suitability for Directed, closed search goals seems to be the availability and correctness of tags to describe the data.

Directed, open search

Questions **E4** and **E6** were designed to reflect the *Directed, open* search goal. In Table 5.2, we observe that fewer than half the participants answered E4 correctly, and approximately half the participants answered E6 correctly. This suggests that our system is not highly suitable for queries with this search goal. In the case of E4, the majority of the participants overlooked the text on the map which informed them of which time frame the data on the map covered. In the case of E6, several participants did not realize they needed to apply multiple McDonald's locations in order to see all the relevant data. A few participants also did not realize that they still had tags applied from the previous question. This reveals two of things about our system and its suitability for these types of queries.

Firstly, it was not clear enough to participants how they could find which time frame the data on the map covered, even though this was illustrated in the tutorial. It is possible that after familiarizing themselves further with the interface, this would no longer be an issue. Alternatively, the functionality of the system could be extended to, after a query, show an overview of how the data is distributed temporally with a visualization similar to the Week or Month date picker, as suggested

by a participant.

Secondly, it emphasizes the importance of intuitive location (and tag) filtering. As a number of participants suggested, it would be more intuitive if the Location filter functioned more like Google Maps, where all instances of, for example, McDonald’s are shown on the map when querying for McDonald’s. We believe that with this functionality, more participants would have answered question E4 correctly. However, implementing such a functionality was beyond the scope of our research.

Undirected search

Questions **E3** and **E5** were designed to reflect the *Undirected* search goal. Participants performed well on both questions, with only a few participants providing partly correct answers. These partly correct answers can be attributed to participants not inspecting the photos closely enough and misinterpreting them.

Generally, participants did not have trouble finding the answer to these questions. To inspect the photos they had found after applying the appropriate filters, they viewed the photos in the Detail view, which they accessed through the map or the *Photos in this area* feature. They often examined events in more detail by scrolling through the *Explore this day* feature.

We therefore conclude that our system is very suitable for Undirected search goals. However, it must be noted that this suitability also depends on the quality of the locations and tags in the system. For example, if it is impossible to find photos of a certain event because the tags available in the system do not describe it well, then users will not be able to explore the photos of this event properly.

Casual-leisure search

Finally, questions **E1** and **E2** were designed to reflect the *Casual-leisure* search goal. Question E1 was answered correctly by almost all participants, with only a few participants answering it partly correctly because they did not inspect the data close enough to determine whether the lifelogger visited a city or merely flew over it. Question E2 was answered correctly by the majority of the participants, although there were also a number of participants who struggled to answer this question and either provided an incorrect answer or indicated that they did not know how to answer the question.

Since the main feature of our system is the map interface, it is fairly straightforward that participants would be able to easily answer question E1. For question E2 however, we identify the same problem as we previously discussed for question E4, which is that there is no immediately apparent method for users to determine on which dates the photos that are displayed on the map have been taken, or that the available methods for this were not clear enough to participants. In addition to the solutions we presented for this previously, we want to present an additional solution to this problem that was suggested by a participant, which is to color-code the Dot map depending on when photos were taken. This could allow users to get a quick impression of when photos were taken.

We believe that, with the addition of a feature which allows participants to more easily obtain an overview of how the data is temporally distributed, our system can be extremely suitable for Casual-leisure search. The map interface supports this goal, as before the Search phase of the experiment started, a number of participants out of curiosity navigated to places they had visited themselves to explore the lifelog photos taken there. However, successful exploration could be more easily attained if the system could display all photos which were found in a query on the map instead

of being limited to 40,000. This way, participants could in all cases immediately be provided with an overview on the map of the entirety of the data, which could elicit a stronger exploration response and allow for more intuitive exploration.

5.2.3 Search strategies

From Tables 5.3 and 5.4, we can deduce if there are any search strategies that participants have employed which have led to participants more successfully answering the search questions. We will briefly discuss search strategies in respect to the Map and Temporal filter features and finally, we try to discern if any participants used search strategies that seem particularly successful.

Map

Overall, we observe that the map interface is often the first thing that participants use for exploration questions, but not for retrieval questions. However, this pattern is influenced by the type of question participants are asked, as a number of exploration questions could be solved by exclusively using the map. We observe that if an exploration question has a temporal aspect, participants often first use the temporal filter rather than the map interface. Likewise, we note that whether or not participants use the map first may also depend on the preceding question they were asked, as they are more likely to interact with the map first if it a question requires them to navigate to a new location.

A more clear, explainable pattern is seen in which of geospatial visualizations participants used to answer different types of questions. We observe that the Heat map was hardly used to answer retrieval questions. This may be due to the abstract nature of the Heat map; when participants are looking for a specific photo in a limited amount of time, they likely do not want to use a geospatial visualization that provides an overview of how many photos were taken on a location. Rather, participants may want to inspect all results of their query one by one. For this, the Dot map or Cluster map are a more suitable visualization, as participants note that with these geospatial visualizations it is easier to access photos. A few participants commented that they found the Dot map most useful for retrieval, however, from our current results we cannot yet conclude if the Dot map is more suitable for retrieval tasks than the Cluster map. To determine this, we suggest a comparative study is performed.

Temporal filter

No clear conclusions can be drawn about whether participants who used the Week or Month temporal filter were generally more successful in answering the questions, as there is not enough data to conclude this. However, it is apparent that participants more often used the Custom temporal filter to answer retrieval questions, most likely because this provided them with the functionality to select an entire year range. The Month temporal filter is used occasionally for exploratory questions, likely because in these questions participants were not asked to query a year range.

Per participant

In general, we have not found an overall search strategy in Table 5.3 that can be considered more effective or successful. Participants who performed well (6 questions or more answered correctly) used a range of different features in a variety of orders, but a similar pattern can be seen with participants who performed less well. We do observe that participants who started more often with the Location or Tag filter generally seemed to perform slightly worse, however, this can also be attributed to the quality of the locations and tags included in the LSC'20 test collection.

Something that shows more strongly from Table 5.3 is that participants often have a strong preference for a certain geospatial visualization, with several participants almost exclusively using one single geospatial visualization. It seems that participants who switched more often between geospatial visualizations generally performed slightly worse, however, participants who struggled to answer a question often switched geospatial visualization in the hopes of finding a strategy to solve said question. Therefore, this observation does not directly prove anything about the success of a search strategy where only a single geospatial visualization is used.

5.2.4 Geospatial visualizations

In this section, we will attempt to answer our research question in relation to the geospatial visualizations: *Out of a Dot, Heat and Cluster geospatial visualization, which one is considered most useful for different search goals in the context of lifelog data?* To do this, we will look in unison at the ratings for *ease of use* and *usefulness* per geospatial visualization as presented in Figure 5.2, at how much time each map visualization was used as presented in Section 5.1.1, and participants' search strategies in Table 5.4. If applicable, we also analyze comments presented in Section 5.1.2.

The first thing that stands out from the ratings per feature presented in Figure 5.2 is that the Cluster map scores very high for both *ease of use* and *usefulness* ratings. The Cluster map was also used significantly more than the Heat map, but not than the Dot map. We note that use time is not necessarily an indication of suitability, since it is also possible that a longer use time indicates that participants did not understand how to use a feature effectively. However, since the *ease of use* and *usefulness* ratings for the Cluster map are very high, we conclude that participants simply used the Cluster map more because they found it easy to use and useful. The comments that were provided on the Cluster map also support this. Participants mention that they gravitated towards using the Cluster map since they found it to be a good combination of the functionalities of the Dot map and the Heat map. However, it must be noted that, looking at Figure 5.4 and 5.3, participants who used the Cluster map almost exclusively did not necessarily perform better than participants who used a mix of geospatial visualizations.

From Figure 5.2, we observe that the Dot map and the Heat map receive comparable ratings. The Dot map is rated slightly higher than the Heat map for *ease of use*, but equally for *usefulness*. We do see that the Heat map sees slightly less use time than the Dot map, although this is not statistically significant. This data suggests that the Dot map and the Heat map are generally both suitable for lifelog access, but not as suitable as the Cluster map.

Using Table 5.4, we gained more insight into which geospatial visualizations were used for certain search goals. Overall, we see that the Cluster map is used consistently over all types of search goals, but slightly more for retrieval questions compared to the other geospatial visualizations. For the exploration questions, participants seem to slightly prefer to use the Cluster map over the Dot map, even though the differences are small, and they highly prefer to use the Cluster map over the Heat map. We also observe that the Dot map is used consistently for all types of search goals, albeit slightly less than the Cluster map. Finally, we see an interesting pattern in the use of the Heat map. It is hardly used for Directed, closed search questions, likely because it provides too abstract of an overview for this type of search goal. It is used more for exploratory questions, but still not as much as the Dot or Cluster map. There were a number of questions for which participants used the Heat map slightly more, which were the Casual-leisure search questions (E1 and E2) and specifically the Directed, open search question where participants were asked to identify at which location more photos were taken (E6). This implies that, even though the Heat map was used less frequently and seems to be less popular amongst participants, it is still suitable for search tasks where users require an overview. In addition to this, we note that participants provided many suggestions for the

improvement of the Heat map and that perhaps if these suggestions were implemented, the results for the Heat map might see improvement.

In conclusion, all geospatial visualizations that were implemented into our system are suitable for access, search and exploration of lifelog data, although results suggest that some visualizations are most suitable for specific search goals. The Dot map seems to be more suitable for search tasks where only a few photos are shown on the map at a time. The Heat map seems more suitable to provide an overview, so we consider it more useful for exploration tasks where users do not need to access the photos often. The Cluster map appears to be suitable for both search goals, and we therefore consider it to be the most useful for different search goals in the context of lifelog data. However, wish to emphasize the value of incorporating multiple map visualizations into a single lifelog exploration system. Comments from participants indicate that each visualization may be useful for different search goals and that they appreciated that multiple visualizations were available to them. If it is not possible to implement multiple geospatial visualizations in a lifelog exploration system, we suggest that a visualization similar to the Cluster map is used, as it provides both the overview functionality of a Heat map and the more detailed view of a Dot map.

5.2.5 Temporal filtering

In this section, we will attempt to answer our research question in relation to the temporal filtering techniques: *Out of a color-coded Week matrix, color-coded Month matrix and Custom drop-down temporal filtering technique, which technique is considered most useful for different search goals in the context of lifelog data?* To do this, we will look in unison at the ratings for *ease of use* and *usefulness* per map visualization as presented in Figure 5.2, at how much each temporal filtering technique was used during the Search phase of the experiment as presented in Section 5.1.1, and participants' search strategies in Table 5.4. Where applicable, we also analyze relevant comments presented in Section 5.1.2.

From the ratings in Figure 5.2 and use time in Section 5.1.1, we can conclude that participants appreciated and used the Custom temporal filter the most out of all temporal filters. For *ease of use*, all temporal filters received comparable ratings, which implies that the Week and Month temporal filters did not receive lower *usefulness* ratings because participants did not understand how to use them. Instead, we believe that participants rated them lower for *usefulness* because they found them too limited. Where the Custom temporal filter allowed them to select a wide date range, the Week temporal filter allowed them to choose only a date range of a week and the Month temporal filter of up to two months. The flexibility of the Custom temporal filter was beneficial for questions like E2, where most participants wished to applied a date range of 3 months, and the retrieval questions, where a date range of a year was given.

These results might suggest that temporal filtering techniques like the Week and Month temporal filters are simply not suited for lifelog exploration and retrieval, however, comments from participants suggest that these features might still have potential for additional exploratory use. The same applies to the Photo distribution feature, which only five participants supplied ratings for, comments show that participants imagine it could be valuable and interesting in different use cases.

The fact that the Week and Month temporal filters were hardly used reveals a limitation in our study, which is that we did not include questions in our user study for which the Week or Month temporal filters could be useful. However, in this case, there is a risk to design questions which are biased towards a certain interface feature. Therefore, we again emphasize the importance of performing lifelog system evaluations with lifeloggers or participants who record their own lifelog data, as then there is no need to design questions to replicate lifelogger search goals and information

needs, which removes any possible bias from the user study.

In conclusion, the Custom temporal filter seems to be most suitable for access, search and exploration of lifelog data. The most important reason for this is its flexibility. Hence, our advice is to implement a temporal filtering technique is very flexible, but also allows users to apply a date range quickly. An example of such a technique is a calendar widget extended with functionalities which allow for quick hour, day, month and year selection, although the suitability of such a temporal feature would need to be confirmed in future work. Finally, we note that we do not discard that an interface like the Week and Month temporal filter can be useful for different, more low-level exploratory search goals. They also have potential to be used as an additional temporal visualization, to inform the user on how the data is temporally distributed.

5.2.6 Overall suitability

Our global research aim is to determine whether geospatial visualization and temporal filtering techniques are suitable for access, search and exploration of lifelog data. From participants' performance in our user study, comments they provided on the system and the SUS-score our system received, we can conclude that these techniques are indeed generally suitable. Even though participants were not very successful in answering questions which focused on the temporal distribution of the data, most other question types were answered successfully. Participants occasionally struggled to answer some retrieval questions, which can mostly be attributed to the quality of the tags which were included in the LSC'20 test collection rather than the lack of suitability of the geospatial visualization and temporal filtering techniques.

However, even though our research shows that lifelog access, search and exploration can be achieved through a map interface and simplified with temporal filtering techniques, we believe that a number of features must be included into such a system to make it feasible. There must be a way for users to view events or photos in a chronological order, a functionality which we provided in our system through the Detail view. We believe that providing a preview of the photos on the map, as we did through the *Photos in this area* feature, may help participants in their search and exploration process. It must also be possible to filter by tag and location. This is shown through the high ratings for *usefulness* for the Location and Tag filters. The Detail view and *Photos in this area* feature were not rated, but we can see from Figures 5.3 and 5.4 that nearly all participants used the Detail view and the *Photos in this area* feature. From Table 5.4, we observe that all four additional features are used quite intensively. However, we emphasize that it is still unclear whether our implementation of these features is optimal. This is an interesting direction for further research.

Finally, we want to highlight a number of limitations of lifelog access through a system like ours, for which we are currently unable provide solutions and which we suggest are further investigated in future work. At present, our system is unable to deal with lifelog photos which do not include a latitude and longitude. When querying, these photos are simply displayed to the user, as by default queries to the database include the bounds of the map. A method must be designed to allow users to access these photos in a way that does not distract from the results on the map, but still allows for intuitive access. Another problem that might occur in large lifelog data sets which are accessed through a map is that it is difficult to access areas where a substantial number of photos were taken, such as home or work. We already identified this problem in our system, where unfiltered clusters of these areas grew so large that opening them could lead to a significant slow down, if not crash, in the browser. If access through a map is to remain a viable solution for growing lifelog databases, a solution must be found for this issue.

5.3 Limitations

In the previous section we already identified a number of limitations of our study. Here, we will repeat any limitations which were previously mentioned and present a number of other limitations which we have not discussed before.

The main limitation in our study is one that was already mentioned a number of times, namely that the participants of our study were not lifeloggers interacting with their own lifelog data. To solve this issue, we designed questions intended to replicate a lifelog search goal. However, it is difficult, if not impossible, to guarantee that these questions are representative of all different lifelogger search goals and information needs. This ties in with a limitation we presented in Section 5.2.5, where we mentioned that it is likely that the questions we designed were not focused enough towards the use of the Week and Month temporal filters. However, incorporating questions for a specific interface feature can introduce a bias towards those features. Having lifeloggers interact with their own data would eliminate the need for representative questions for different search goals entirely, which would provide more realistic and unbiased results.

There are a number of other limitations in our study design. The order in which participants interacted with the features of our system per question was recorded by writing it down. However, it would be more accurate if this was automatically recorded. Automating this process could also have allowed us to record how much time participants interacted with each feature **per question**, rather than over the duration of the entire experiment. This could also have shown us more clearly how participants interacted with the system, for example, whether they spend a lot of time on one feature before moving on to the next or if they jump between features instead. The most accurate way to record this would be to use eye tracking, which would be the most appropriate and accurate method to determine which features and search strategies participants use most [33].

The final limitation we want to discuss is related to the additional system features, outside of the geospatial visualizations and temporal filters. Even though we asked participants to provide ratings for the Location and Tag filters, for the *Photos in this area* and Detail view features it is currently still unclear how they contribute to the success of our system and how their design could be improved.

5.4 Future work

In this section, we discuss possible extensions of our current system and directions for future research.

5.4.1 System extensions

The current implementation of the Lifelog Browser only contains a fundamental array of features, which are essential to enable basic lifelog search and exploration. In this Section 5.2, we already discussed a number of extensions which could be done to improve our current design. In addition to these, we have several other extensions which we believe could be valuable. During our study of related work and during the implementation process of the system, we compiled a list of features that could be valuable additions to future versions of our system. These extensions and improvements were not implemented for our current study because of time constraints. The features that were inspired by related work are listed below:

Slideshow: Lee et al. [24] implemented a feature in their lifelog system where if the user mouses over a lifelog photo, a number of preceding and succeeding photos are shown in a slideshow. We believe that such a feature can be valuable for both retrieval and exploration tasks, as this allows

the user to inspect the context in which a photo was taken without having to open the Detail view.

Scalable insets: Lekschas et al. [27] present the concept of *scalable insets*, a technique where a summary of noteworthy visual data is presented in one or multiple thumbnails next a map or highly detailed photo. In our case, these scalable insets could be presented next to clusters with many photos, to provide users with an overview of photos inside that cluster.

Extended geospatial filtering: In their lifelog retrieval system, Tran et al. [35] provide a useful additional map filtering option, where users may filter by location by drawing a rectangle on the map. Implicitly, this functionality is already provided in our system by allowing the user to focus the map on a certain area that they wish to see, however, we believe drawing on the map provides a more flexible filtering method as the size of the rectangle can be adapted. As a concrete example of the added value, in the current design of our system, when trying to find photos in Norway, participants had difficulty to position the map so that only Norway could be viewed on it. This feature would simplify such a search task.

In addition to the features presented above, there are also a number of improvements that can be done to the current features of the system which would increase the system’s usability and suitability for lifelog search and exploration. First and foremost, many participants remarked that the map interface was quite slow, which made the system feel unreliable at times. For a more seamless interaction, the photo querying process should be optimized so that map loading times are reduced. In addition to this, we believe it would be valuable if the map did not have a marker limit, so all photos from the database could be displayed on the map simultaneously. This would improve the suitability of our system as a lifelog data visualization tool, as it would allow users to more easily obtain an overview of the geospatial distribution of lifelog data.

Finally, we believe improvements could be done to the tags. As described in Section 3.3.1, during implementation of the Tag filter functionality, we chose to ignore the confidence value that came with each tag in the LSC’20 test collection, as we did not want tags to be erroneously excluded. However, this caused the issue that photos often came with tags which did not describe the photo well. A solution would be to discard tags below a certain confidence level, or to design a method which uses this confidence value to rank photos by relevance.

5.4.2 Directions for further research

In Section 5.2, we already indicated a number of directions for further research. In this section, we will repeat the suggestions that were previously made and introduce a number of additional suggestions. Our suggestions cover four distinct directions:

Geospatial visualizations: In our study, we evaluated three geospatial visualizations which we believed would be suitable for lifelog access, search and exploration. However, there are more geospatial visualizations which could be suitable for this. A geospatial visualization which we purposefully excluded from our evaluation but which we believe could still be valuable is the route map. We foresee a number of possible complications in the usability of a route map to display lifelog data, as we explained in Section 3.1.2, therefore, further research is needed to determine whether it is indeed suitable for lifelog access. We also encourage research into the issues which we discussed in Section 5.2.6 that arise from the use of a map as the main access point for lifelog photos, specifically access to photos which do not have a location and how to handle large clusters of photos.

Temporal filtering techniques: We originally intended to evaluate three different types of temporal filtering techniques which provided almost identical functionality, but had different input

methods. However, because we felt it could be confusing to users to provide them with these three very similar temporal filtering options, we adapted our design. Even though we feel this was an appropriate decision for our study, we believe it could still be valuable to perform a user study with different types of temporal filtering techniques that provide the same functionality, as our current study only provides suggestions on what functionality a temporal filter must supply, not on which input method participants prefer. We recommend a comparative user study for this.

Temporal visualizations: Our research shows that the Week and Month temporal filters were not preferred by participants because they did not supply participants with enough flexibility to select the desired date ranges. However, they do show potential for supplemental temporal visualization rather than temporal filtering. We suggest an additional feature where the temporal distribution of the data is visualized in a similar manner as in the Week and Month temporal filters with Photo distribution. We also want to highlight the timeline as a possible temporal visualization technique. Future research must be done into the suitability and usefulness of these types of features for lifelog access, search and exploration.

Supportive functions: Finally, we believe it may be valuable to perform further research into features which are supportive of the map and temporal filter interfaces, such as the Detail view and the *Photos in this area* feature. As mentioned before, our results suggest that these features are useful for lifelog access, search and exploration, but it is unclear to what extent. Future research is needed to determine the optimal design of these types of features.

Chapter 6

Conclusion

This thesis aimed to identify whether geospatial visualization and temporal filtering techniques are suitable for access, search and exploration of lifelog data. Specifically, we intended to determine which geospatial visualization (out of a Dot map, Cluster map and Heat map) and which temporal filtering method (out of a color-coded Week matrix, color-coded Month matrix and Custom drop-down technique) is considered most useful for different search goals in the context of lifelog data.

From our user study, we observe that the geospatial visualization and temporal filtering techniques we chose to implement are indeed generally suitable for lifelog access. The current implementation of our system is especially suitable for Undirected and Casual-leisure search. The system also seems suitable for Directed, open search (retrieval tasks), but its success in this area is largely dependent on the quality of a Tag filter feature. Participants had difficulty to answer questions of the Directed, open search type, which indicated that a number of improvements should be done to the system to better allow for these types of queries. We suggest that a feature is implemented which visualizes how the data is temporally distributed. A visualization similar to our Week or Month temporal filter or a timeline could be suitable for this.

From the geospatial visualizations that we evaluated, we conclude that the Cluster map is overall the most useful for different search goals in the context of lifelog data. The Cluster map received the highest score for our metrics *ease of use* and *usefulness*, and participants mostly provided positive comments on it. However, our results show that even though the Cluster map appears to be the most suitable all-round geospatial visualization, it is not entirely clear whether it is also most suitable for every search goal in isolation, since the Dot map and the Heat map also received generally favorable ratings and our results suggest that they might be useful for specific search goals. We therefore recommend that lifelog systems that aim to provide exploration functionalities implement multiple geospatial visualizations to allow users the freedom to choose the most appropriate visualization for their goal. If it is not feasible to implement multiple geospatial visualizations in a lifelog exploration system, we recommend implementing a geospatial visualization similar to the Cluster map. Finally, for systems that are focused on lifelog retrieval, we do not recommend implementing an abstract visualization like the Heat map, but rather, a visualization that allows for quick photo access like the Dot or the Cluster map.

From the temporal filtering techniques that we evaluated, we conclude that the Custom temporal filter is the most useful for different search goals in the context of lifelog data. This was clearly shown in the ratings for *ease of use* and *usefulness*, as well in how much time participants used this feature and the comments they provided on it. The most important reason for the Custom temporal filter's popularity is the flexibility that it provides over the Week and Month temporal filters, as with those,

participants were not able to select, for example, an entire year. Our advice for both lifelog retrieval and exploration systems is therefore to choose a temporal filtering technique which provides as much flexibility as possible. However, we do not necessarily recommend to use our exact implementation of the Custom temporal filter, as we believe there are more optimal ways to implement such a temporal filtering feature. We suggest further research is done in the form of a comparative study between different temporal filtering techniques which provide the same functionality.

Bibliography

- [1] A. S. f. P. Affairs. System usability scale (sus), Sep 2013.
- [2] M. Bolanos, M. Dimiccoli, and P. Radeva. Toward storytelling from visual lifelogging: An overview. *IEEE Transactions on Human-Machine Systems*, 47(1):77–90, 2016.
- [3] V. Bush et al. As we may think. *The atlantic monthly*, 176(1):101–108, 1945.
- [4] S. Chowdhury, M. S. Ferdous, and J. M. Jose. A user-study examining visualization of lifelogs. In *2016 14th International Workshop on Content-Based Multimedia Indexing (CBMI)*, pages 1–6. IEEE, 2016.
- [5] M. Czerwinski, D. W. Gage, J. Gemmell, C. C. Marshall, M. A. Pérez-Quinones, M. M. Skeels, and T. Catarci. Digital memories in an era of ubiquitous computing and abundant storage. *Communications of the ACM*, 49(1):44–50, 2006.
- [6] D. T. Dang Nguyen, L. Piras, M. Riegler, G. Boato, L. Zhou, and C. Gurrin. Overview of imageclef lifelog 2017: lifelog retrieval and summarization. 2017.
- [7] DataInterfaces. Flickr cities. <http://www.datainterfaces.org/projects/flickr/#torino>.
- [8] M. Dodge and R. 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.
- [9] A. R. Doherty and A. F. Smeaton. Automatically segmenting lifelog data into events. In *2008 Ninth International Workshop on Image Analysis for Multimedia Interactive Services*, pages 20–23. IEEE, 2008.
- [10] A. Garcia del Molino, J.-H. Lim, and A.-H. Tan. Predicting visual context for unsupervised event segmentation in continuous photo-streams. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 10–17, 2018.
- [11] J. Gemmell, G. Bell, R. Lueder, S. Drucker, and C. Wong. Mylifebits: fulfilling the memex vision. In *Proceedings of the tenth ACM international conference on Multimedia*, pages 235–238, 2002.
- [12] C. Gurrin, R. Albatal, H. Joho, and K. Ishii. A privacy by design approach to lifelogging. In *Digital enlightenment yearbook 2014*, pages 49–73. IOS Press, 2014.
- [13] C. Gurrin, H. Joho, F. Hopfgartner, L. Zhou, and R. 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, 2016.

- [14] C. Gurrin, B. T. Jónsson, K. Schöffmann, D.-T. Dang-Nguyen, J. Lokoč, M.-T. Tran, W. Hürst, L. Rossetto, and G. Healy. Introduction to the fourth annual lifelog search challenge, lsc'21. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*, pages 690–691, 2021.
- [15] C. Gurrin, T.-K. Le, V.-T. Ninh, D.-T. Dang-Nguyen, B. T. Jónsson, J. Lokoč, W. Hürst, M.-T. Tran, and K. Schoeffmann. Introduction to the third annual lifelog search challenge (lsc'20). In *Proceedings of the 2020 International Conference on Multimedia Retrieval*, pages 584–585, 2020.
- [16] C. Gurrin, K. Schoeffmann, H. Joho, B. Munzer, R. Albatat, F. Hopfgartner, L. Zhou, and D.-T. Dang-Nguyen. A test collection for interactive lifelog retrieval. In *International Conference on Multimedia Modeling*, pages 312–324. Springer, 2019.
- [17] C. Gurrin, A. F. Smeaton, and A. R. Doherty. Lifelogging: Personal big data. *Foundations and trends in information retrieval*, 8(1):1–125, 2014.
- [18] S. Hodges, L. Williams, E. Berry, S. Izadi, J. Srinivasan, A. Butler, G. Smyth, N. Kapur, and K. Wood. Sensecam: A retrospective memory aid. In *International Conference on Ubiquitous Computing*, pages 177–193. Springer, 2006.
- [19] F. Hopfgartner, Y. Yang, L. M. Zhou, and C. Gurrin. User interaction templates for the design of lifelogging systems. In *Semantic Models for Adaptive Interactive Systems*, pages 187–204. Springer, 2013.
- [20] W. Hürst, K. Ouwehand, M. Mengerink, A. Duane, and C. Gurrin. Geospatial access to lifelogging photos in virtual reality. In *Proceedings of the 2018 ACM Workshop on The Lifelog Search Challenge*, pages 33–37, 2018.
- [21] K.-S. Hwang and S.-B. Cho. A lifelog browser for visualization and search of mobile everyday-life. *Mobile Information Systems*, 10(3):243–258, 2014.
- [22] J. H. Jeon, J. Yeon, S.-g. Lee, and J. Seo. Exploratory visualization of smartphone-based lifelogging data using smart reality testbed. In *2014 International conference on big data and smart computing (BIGCOMP)*, pages 29–33. IEEE, 2014.
- [23] V. Kalnikaite, A. Sellen, S. Whittaker, and D. 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, 2010.
- [24] H. Lee, A. F. Smeaton, N. E. O'connor, G. Jones, M. Blighe, D. Byrne, A. Doherty, and C. Gurrin. Constructing a sensecam visual diary as a media process. *Multimedia Systems*, 14(6):341–349, 2008.
- [25] M. L. Lee and A. K. Dey. Lifelogging memory appliance for people with episodic memory impairment. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 44–53, 2008.
- [26] A. Leibetseder and K. Schoeffmann. lifexplore at the lifelog search challenge 2021. In *Proceedings of the 4th Annual on Lifelog Search Challenge*, pages 23–28. 2021.
- [27] F. Lekschas, M. Behrisch, B. Bach, P. Kerpedjiev, N. Gehlenborg, and H. Pfister. Pattern-driven navigation in 2d multiscale visualizations with scalable insets. *IEEE transactions on visualization and computer graphics*, 26(1):611–621, 2019.
- [28] D. Lupton. *The quantified self*. John Wiley & Sons, 2016.

- [29] S. Reddy, A. Parker, J. Hyman, J. Burke, D. Estrin, and M. Hansen. Image browsing, processing, and clustering for participatory sensing: lessons from a dietsense prototype. In *Proceedings of the 4th workshop on Embedded networked sensors*, pages 13–17, 2007.
- [30] D. E. Rose and D. Levinson. Understanding user goals in web search. In *Proceedings of the 13th international conference on World Wide Web*, pages 13–19, 2004.
- [31] A. J. Sellen and S. Whittaker. Beyond total capture: a constructive critique of lifelogging. *Communications of the ACM*, 53(5):70–77, 2010.
- [32] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of information visualization*, pages 364–371. Elsevier, 2003.
- [33] T. Strandvall. Eye tracking in human-computer interaction and usability research. In *IFIP Conference on Human-Computer Interaction*, pages 936–937. Springer, 2009.
- [34] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- [35] L.-D. Tran, M.-D. Nguyen, N. T. Binh, H. Lee, and C. Gurrin. Myscéal: an experimental interactive lifelog retrieval system for lsc’20. In *Proceedings of the Third Annual Workshop on Lifelog Search Challenge*, pages 23–28, 2020.
- [36] M. L. Wilson and D. Elswiler. Casual-leisure searching: the exploratory search scenarios that break our current models. 2010.
- [37] Y. Yang and C. Gurrin. Personal lifelog visualization. In *Proceedings of the 4th International SenseCam & Pervasive Imaging Conference*, pages 82–83, 2013.
- [38] Y. Yang, H. Lee, and C. Gurrin. Visualizing lifelog data for different interaction platforms. In *CHI’13 Extended Abstracts on Human Factors in Computing Systems*, pages 1785–1790. 2013.

Appendix A

Experiment Questionnaire

A.1 Consent Form and Explanation

You are invited to participate in evaluating our lifelog exploration system, the Lifelog Browser. The Lifelog Browser was designed to make lifelog data accessible and explorable. Lifelog data is data that is recorded by wearable sensors that a person (a lifelogger) carries around during their day. The sensors record details of the lifelogger’s day passively, without any interference from the lifelogger themselves. By doing this, the lifelogger can capture their life’s experiences in detail. Specifically, the lifelog data you will look at with the Lifelog Browser is a collection of images recorded by someone using a wearable camera which takes a photo approximately twice every minute.

During the study, you will be asked to perform a set of tasks. These tasks consist of answering questions about the lifelog data and finding a number of specific photos. This will help us to evaluate the design of the Lifelog Browser. We are not evaluating you or your performance in any way.

As you perform various tasks with the system, your actions and comments will be recorded, through screen recording and voice recording. You will be asked to verbally describe what you are doing and voice any thoughts you may have about the system. You may be asked questions during and after the evaluation, in order to clarify our understanding of your actions and view of the system. The evaluation session takes around 60 minutes. After you have completed all the tasks, you will fill in a short questionnaire and you may give additional comments.

The information you provide will be processed anonymously. Your name will be removed and only a subject number will be used to identify you during analyses and any written report of the research. The evaluation will be recorded, and all data stored securely, viewed only by experimenters.

Your participation is voluntary and unpaid. You are free to withdraw from this study at any time without providing a reason.

A.2 Demographics

Participants were asked to answer the following questions before starting the tutorial:

- What is your gender?
- What is your age?

- How would you rate your own ability to navigate and evaluate information online or in a digital format?
 - Very poor
 - Poor
 - Neutral
 - Good
 - Very good

- How much experience do you have with lifelog data?
 - None at all
 - A little
 - A moderate amount
 - A lot
 - A great deal

Appendix B

Additional Results

B.1 Qualitative Comments

Positive

- Easy to open different photos (not too much clicking).
- Useful when there are not too many pictures on the map.
- Useful to find paths or routes.
- Useful to find specific photos
- Useful on a low zoom level to find exactly where a photo was taken.

Negative

- Can be cluttered and overwhelming when there are too many dots on the map.
- Not the best way to show this kind of data.
- It's difficult to see what is going on in crowded areas.
- It's redundant if you also have the cluster map.

Table B.1: **Dot map** comments.

Positive

- It's the most useful and easy to use.
- It merges the dot and heat map into one.
- It provides a good overview, so the amount of data that is presented is less overwhelming and is shown clearly.
- Useful when there are many photos on the map.
- You can easily see how many photos were taken where.
- It really excels in showing you places that were visited more often or where the lifelogger traveled.

Negative

- Many clicks are needed to get to the level that you can open the photos, this can be frustrating also because then the system keeps on loading.
- It would be nice to have more logical zones for clusters, like a certain neighborhood rather than random divisions.
- It wasn't clear how to get through the Cluster map to the specific pictures.

Table B.2: **Cluster map** comments.

Positive

- Useful to get a quick sense of things, to let you recognize places that were visited more often or where the lifelogger traveled.
- Useful when there are many photos on the map.
- It's nice that you can see where the lifelogger traveled with the blue 'trace'.
- Useful when there are many photos on the map.
- Gives a general overview.
- It's nicer to see the overview in colors rather than in numbers, that's why heat map is nicer than the cluster map.

Negative

- It is sometimes difficult to translate the color to how many photos are in that spot. When you zoom in, the color value changes because there are fewer photos in total, resulting in just bright red spots on the map.
- It does not add much over the cluster map, it feels like a less informative version of it. It is also not as intuitive.
- Could not really be used to answer the questions that were asked.
- Becomes unclear when zoomed in.
- If there are too few pictures the blue can be hard to notice, especially when you're moving around the map.
- It generalizes too much, so it's hard to distinguish areas.
- It was difficult to see where to click because the mouse over areas are small.
- There are too many gradations of colors. It would have been nicer if they had fit in with the color scheme of the design.
- It causes too much clutter.
- It is not useful for finding specific photos, because it's difficult to associate the areas with single pictures.

Table B.3: **Heat map** comments.

Positive

- It's fast.
- It's intuitive, all the options make sense without much explanation.
- All map options are useful. Some participants commented this too, but with the note that they were all useful on different zoom levels.
- The 'spiderification' works well to see photos from the same location.
- The zooming animations look really nice.
- It's easy to see how much data there is at a specific point.
- In combination with the filters, it's very easy to find certain pictures.

Negative

- It can be slow, unresponsive and doesn't always load consistently, which makes the system feel unreliable.
- You can zoom out too much.
- It should be easier to jump to cities or countries.
- Quite a few clicks are needed to access a specific photo.
- Sometimes 40.000 pictures is not enough when you're viewing a large area.
- After 'spiderification' of a cluster, the dots are still the same color which makes it difficult to see which ones are part of the spider in a cluttered area.

Table B.4: **General map** comments.

Positive

- The visuals make it clear and it's nice to look at.
- Might have been used if there had been questions about a specific week or month.
- It is really nice to see where the data is so you know there should be results.
- It is great for exploration.
- For the **week picker**, it's efficient that you can select both the day and hour at the same time.

Negative

- It's confusing or unclear and too limited.
- You should be able to select disjoint areas.
- It's unclear whether the dots are connected in some way to the map or not.
- It has no added value over the Custom date picker.
- For the **week picker**, it's unclear what the date is as you only see the day.
- For the **week picker**, no use was found for it with the questions that were asked.
- For the **month picker**, it wasn't clear that the arrows could be used to navigate through the months.

Table B.5: **Week and Month date picker** comments.

Positive

- It gives a clear overview of how many pictures are available.
- It can be really helpful, even though with the questions asked there was no reason to use it.
- It looks nice.
- It could be the default.
- It could have its own feature.

Negative

- It is unclear what the use is for the questions that were asked.
- The checkbox could be more visible.
- It was unclear what it did.

Table B.6: **Photo distribution** comments.

Positive

- It is by far the most useful and fastest to use.
- It is nice to be able to select the different days of the week and times by clicking.
- It is easy to understand and concise.
- It does everything that the others do, except showing the photo distribution.
- It has very extensive filtering possibilities.
- It gets the data you want to see loaded pretty quickly.
- It is best for searching for a specific thing.
- Just having the Custom date picker would be enough.

Negative

- Filling it in takes lots of clicking, which is tedious. It takes too much time to fill in.
- When clicking the dropdown, it would be better to have a calendar appear to select the start/end day, as selecting dates with a calendar is less time-consuming.
- It would be nice if it had an option to select the time of day (morning, afternoon or evening), month (or season) and year without filling in a date range.
- It did not display an error when an invalid date range was selected.
- It would be nice to be able to leave certain drop-down items empty so you can only filter on the ones you have filled in.
- Would be nice if it had a 'relative' option, so you can select a date and say you want to find photos, for example, a year from then.

Table B.7: **Custom date picker** comments.

Positive

- The filter option worked seamlessly, it's really nice to see that when you change something in the filter it changes on the map.
- Intuitive and quick and easy to use.
- Well-implemented.

Negative

- Having all of the different date pickers is too many options, it would be fine to only have one, or combine the existing one so you can drill down to different levels (week, month, year).
- For the week and month, it is not clear in what situation you would use them over the custom date filter. It would need to be a very specific question in order to use those.
- The different selection methods were counter-intuitive, as you could drag in the month and week pickers but not when selecting hours or days.
- It's confusing when you forget to press apply or reset.
- It is unclear how the different interfaces differ.

Table B.8: **General date picker** comments.

Positive

- Lets you quickly find photos if you know what to look for.

Negative

- It should also contain countries and cities, not just specific locations.
- The map should immediately jump to the location once you add it.
- It is annoying that there are duplicate locations for places like McDonald's, it would be nicer if one location shows all McDonald's instead of having to add it.
- It is kind of finicky, but mostly because locations are limited.
- If it only allows you to filter by specific location, maybe it's better to change the name 'location' to something like 'places', so it's more obvious that filtering by country/city is not possible.
- If you don't know what you're looking for exactly, it's hard to use/useless.

Table B.9: **Location filter** comments.

Positive

-
- It was easy to find the tags you were looking for.
 - There is a large variety of tags.

Negative

-
- The tags don't really describe the data well, so it's difficult to rely on them as you won't know what results will show up.
 - There are tags missing, there should be more.
 - Combinations of multiple tags would often quickly filter out all the data so no results were shown.
 - It is unclear which tags affect how much data is being shown.
 - It is hard to know which tags to use. It's a guessing game to find the ones you want.
 - It's not nice that sometimes you need to scroll down to see the tag filter.
 - Photos have too many tags.

Table B.10: **Tag filter** comments.**Positive**

-
- They are easy to use and understand.
 - The drop-downs with suggestions were helpful.
 - The design is great.
 - The high number of tags and locations made it easy to find the data I wanted.

Negative

-
- The reset button attracts more attention than the 'add' button, which is counter-intuitive.
 - It would work more nicely if the tag and location fields were combined, so it would have more of a Google Maps functionality, where you can look for specific locations, countries but also types of places, such as supermarkets.
 - It's confusing that adding multiple tags or locations affect the data differently.
 - Both filters sometimes give suggestions which are not exactly what you were looking for.

Table B.11: **General location/tag filter** comments.

Positive

- The system is useful to find photos.
- The layout makes sense, with everything in its own box. It's clear and you always know where everything is that way.
- It could get very intuitive if you use it more, like if you know which tags to pick to find what.
- The system is nice, it would be cool if it could be integrated with your phone album.
- The Photos in this area preview is useful, so you can quickly scan the photos. It provides a nice overview.
- It was a bit difficult to use at the beginning, but I quickly got the hang of it and then it was easy.
- There are not too many options, which makes it easy to narrow down to what you are looking for.
- The system is exhaustive, there are a lot of options. This makes it user friendly, because people might prefer to look for different things.
- The interface is clear and simple, not too much to look at. The colors are not too in your face.
- It's nice that you can see where the lifelogger traveled.
- It is simple and easy to grasp, super clear how to use.

Negative

- Having to constantly click apply, especially after resetting, is tiresome. It would be better if the filters could be applied immediately after doing a single change, to make the experience more seamless.
- There are a lot of options, maybe some simplification can be done?
- I might need more practice to really get used to using it.
- It is slow, should run more smoothly.
- Having the Photos in this area and Explore this day look so similar is a little confusing. Would be nice to get more feedback on what the filters are doing to the data that is shown.
- Very data sensitive, so I would hesitate to use it myself.

Table B.12: **General** comments.

Improvements

- It would be nice to have the photos of Explore this day and Photos in this area in the same overview, so it's less easy to confuse them. Then the 'Explore this day' in the detail view would be obsolete. The Detail view would then only have information about the image that was clicked.
- For the Detail view, make the photo itself bigger, details about the photo a little smaller and make the Explore this day bigger and scrollable vertically, as the photo and the day are the most important.
- Clicking a tag in the detail view should show all the photos with that tag on the map.
- Photos in this area could be sorted by activity rather than chronologically.
- Photos in this area could be used to show what is in a 'spiderified' cluster.
- It would be nice to be able to order the photos in a cluster when 'spiderified' by date.
- When opening a cluster, it would be nice to be able to filter those photos separately by for example tag.
- It would be nice if it was possible to collapse the Photos in this area and Time filters, to only be able to focus on the map.
- It would be better to constrain the photo distributions on Week and Month to the map only, so you can see how the data is spread out across the map.
- It might be nice to color code the dots on a color scale which communicates at which time they were taken.
- The first 40.000 photos could be more meaningful, so a subselection of events could be shown rather than just the first 40.000 photos. This would make it easier to see at a glance where the lifelogger went.

Table B.13: Suggestions for **improvement**.

New features

- It would be nice to have a histogram of tags, so you can see how often certain tags were used in the area that you are looking at.
- It would be nice to be able to see the map while being inside the Detail view. The photo you're viewing will then be smaller, but can be enlarged further when clicking on it.
- It would be nice to have a random explore mode.
- Having a favorite feature for either photos or locations.
- Showing a list of most visited cities.
- If the system is for personal use only, it would be nice to be able to tag persons.

Table B.14: Suggestions for **new features**.

Observations

- Participants really struggle with the fact that they have to click apply after adding something. Specifically, they find it annoying that after resetting, they still have to click apply. Some participants even suggested that it would be better if all changes (not only reset) were immediately applied. This highlights the importance of having a system that is very responsive. As users are searching, they don't want to pay attention to which buttons exactly they have to press. Possibly, the 'apply' highlighting is also not clear enough and highlighting this more would make it more intuitive.
- For a number of participants, it seemed to be unclear that pressing reset and applying one of the date picker options resets the whole date filter. They were going through the week, month and custom and pressing reset on all of them.
- For the month filter, it seemed that for some participants it would be more intuitive if it automatically filtered whichever month they saw on the spreadsheet. They did not understand that they had to select dates on the spreadsheet before it would show.
- It also seemed like users thought that the dots on the week and month filter also somehow corresponded with what was shown on the map, in the sense that if no data was shown on the map, the dots on the calendar also should not have appeared. This could be an interesting functionality to actually quickly view how the data on the map is distributed date-wise.
- It was not always clear to participants that if they wanted to see all the data, they did not need to apply any time filters.
- Some participants did not understand that the Detail view showed all photos of the day and that this view was not influenced by the filters they had applied.
- Even though the Photo distribution feature was hardly used, a number of participants commented that they really like the feature, even though they could not find a use for it to answer the questions they were given.

Table B.15: **Observations** made during the experimenting process.