Geospatial Access to Lifelogging Photos in Virtual Reality

Wolfgang Hürst
Utrecht University
Utrecht, Netherlands
huerst@uu.nl

Kevin Ouwehand
Utrecht University
Utrecht, Netherlands
k.m.r.ouwehand@students.uu.nl

Marijn Mengerink
Utrecht University
Utrecht, Netherlands
M.H.Mengerink@students.uu.nl

Aaron Duane
Dublin City University
Dublin, Ireland
aaron.duane@insight-centre.org

Cathal Gurrin
Dublin City University
Dublin, Ireland
cgurrin@computing.dcu.ie

ABSTRACT

We present a virtual reality system for accessing geotagged photos taken with a lifelogging camera. Photos are spatially located on a world map that can be explored with a head-mounted display. Using a virtual reality headset allows users to easily and intuitively explore this large information space. Images are initially represented by icons but become visible once a user gets closer to a particular area of interest. While not suitable for all search tasks, this visualisation has benefits in situations where location plays a significant role; be it because the actual content is location-related or because the owner of the lifelog remembers and associates the related event with certain places. Likewise, our spatial representation of the data often implicitly reveals a temporal relationship, which can be helpful in the search process as well.

CCS CONCEPTS

• **Human-centered computing → Virtual reality**; *Human computer interaction (HCI)*; *Visualization systems and tools*; • **Applied computing** → *Education*;

KEYWORDS

Lifelogging, virtual reality, geospatial photo browsing

ACM Reference Format:
Wolfgang Hürst, Kevin Ouwehand, Marijn Mengerink, Aaron Duane, and Cathal Gurrin. 2018. Geospatial Access to Lifelogging Photos in Virtual Reality. In *LSC'18: The Lifelog Search Challenge, June 11, 2018, Yokohama, Japan.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3210539.3210547

1 INTRODUCTION

Your private photos are more than just pictures showing some content. For you, they represent a story, a context, a situation in your life. Looking at them triggers the memory of all those things that you associate with them or the moment they were taken. This is even more true for lifelogging photos, that is, pictures taken automatically by body worn cameras at certain time intervals [5]. Here, photos are not carefully composed but literally present a snapshot of your life. For example, an apparently random shot of the car wing of an old car might seem useless and irrelevant for most people, suggesting it can be thrown away. But if this is your lifelog photo and the car is your grandfather's, it might have a very high value for you. For example, it may remind you of this meeting you had with him right before he sold this car. Thus, inherent characteristics of lifelogging photos are that they are *personal* and that they represent more than the content shown on the image; they have a *meta context* that may be of high relevance to the owner. Looking at them can trigger the memory of this context.

Yet, this also works the other way around when accessing your lifelog data in search for a particular photo. There are of course situations where you exactly remember a picture's content, but forgot about the meta context. Yet, more often, you also remember the context in which the photo was taken. Sometimes, you might not even remember the actual content of the photo at all, but just the things that happened around it (as in the example above, where you might very well remember the meeting with your grandfather and when and where it took place, but have no memory of the actual photos taken at the time).

Considering such contexts and offering features to search for them is therefore an essential aspect that should be considered by search engines for lifelogging data. Yet, we also claim that there are situations where a pure visualisation of such meta contexts can be helpful or even better than, for example, specifying them in queries. For example, imagine a situation like this: "When spending a long weekend in Dublin, I remember that we took a half-day trip to this nice place in the country side about an hour's drive from our hotel in the city center. I want to look at my lifelogs from that place." Such an information need is difficult to specify in a query, especially if you do not remember what the place is called or where exactly it was located. But even without this knowledge, it should be very easy to spot on a map that visualises labels at all places where lifelogging data has been recorded; it must be the only place with many lifelog photos outside of Dublin, and there must be a trace of photos from the drive connecting the photos taken in the city and the ones from the target location.

People often remember where events took place. Therefore, we argue that a map-based visualisation of your lifelog photos can be very helpful in many search tasks. And in addition to a targeted search, such as the example above, there are other reasons why

LSC'18, June 11, 2018, Yokohama, Japan

ACM ISBN 978-1-4503-5796-8/18/06…$15.00
https://doi.org/10.1145/3210539.3210547

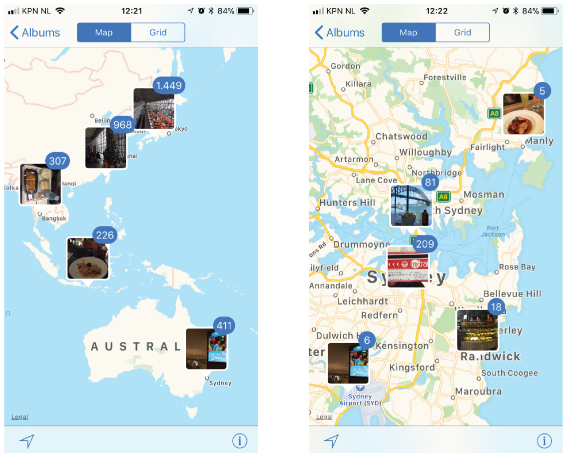

Figure 1: Map-based visualisation of photos taken on a mobile phone (iPhone photo app).

people access their lifelog data. For example, like with regular photo albums, someone might just want to look at lifelog data from a weekend trip to Dublin for leisure purposes. Again, we claim that a good visualisation, and in particular a geospatial arrangement of photos on a map, can be helpful in such cases. And in fact, such representations are already heavily used in common photo organising tools. See for example Figure 1, which shows a map-based visualisation of photos taken with a mobile phone. Not only is such a representation useful for quickly accessing photos from a particular trip, but it also offers a fun and entertaining way to leisurely explore your photo library.

One of the problems with such visualisations is however that they are usually more effective if there is more space available for the visual representations. While browsing photos on your phone using a map might be fun, the small display also limits the experience quite significantly. Similar to how this situation improves when using the large monitor of a laptop or desktop PC, we can argue that a fully immersive head-mounted display, like provided with state-of-the-art virtual reality (VR) headsets, might be beneficial for geospatial visualisation of lifelog data. Such displays could enable us to show a whole world map, enriched with labels of spots where photos have been taken that can then easily and intuitively be accessed.

In this paper, we present a first prototype of this idea, that is, a VR environment to access and explore geotagged lifelogging photos. Our system is intended for the Lifelog Search Challenge (LSC) 2018, an initiative from the Information Retrieval (IR) community with the aim to create comparative benchmarking exercises for interactive lifelog retrieval systems [10]. After summarising related work (Section 2), we describe the implementation of our system (Section 3). Then, we illustrate for what kind of tasks from the LSC 2018 dataset our system might be useful and for which not (Section 4). We conclude with a discussion of our current and future work (Section 5).

2 RELATED WORK

Our idea differs from most existing systems for lifelog data access with respect to the following two aspects: First, we heavily rely on information visualisation as a means of data access, in particular geospatial data representation. Second, we are using VR headsets for data access in contrast to common devices such as desktop PCs and laptops or handheld mobile devices such as phones and tablets. We therefore discuss related work with respect to these two aspects.

When we constantly and automatically record all kinds of data from our life, it is obviously interesting and helpful to somehow visualise the information represented by it. Examples include timeline-based visualisation (e.g., how many calories did I burn today compared to other days this months), map-based visualisations (e.g., where did I go this week), and many more. Kwak et al. [8] present a comparative discussion of common tools for logging sensor data on smartphones. Jeon et al. [6] give an example of visualisations that is explicitly targeted towards supporting users to explore their data. In relation to photos captured with Lifelogging cameras, Yang and Gurrin [11] discuss different options, including a visual diary and calendar summary view. Chowdhury et al. [1] also study different visualisations including a map-based one. Further, more abstract options are discussed by Duane et al. [2]. While map-based visualisations are often mentioned as useful in certain contexts and also used by some (e.g., [1, 6]), most of the time, people see such visualisations of lifelog data as an end in itself. That is, the pure representation of the information is the sole purpose of it. We argue that especially for location-based information, a good visualisation can also be used as means of access to the actual data. We will illustrate this in the use case example discussed in Section 4.

While most interfaces for more sophisticated access to lifelog data are developed for standard PCs and laptop computers, researchers have started to explore other platforms as well. Yang et al. [12, 13] address the problem of Lifelog access on mobile handheld devices, which given the ubiquity of this platform is a very important aspect. Here, the challenges are mostly coping with the small size of the devices' screens. We argue that head-worn Virtual Reality glasses might offer a significant advantage for the interactive exploration of huge Lifelog archives given their immersive screens which can offer 360 degree surround views. Schaefer et al. [9] and Khanwalkar et al. [7] both present interesting ideas and concepts for exploring large photo archives in VR. Yet, given the different nature of the data, it seems doubtful that such interface designs will be suitable for Lifelog photos. Duane et al. [3, 4] introduce interface designs for accessing large Lifelog photo sets in VR that are based on standard filtering approaches known from desktop PC interfaces. While appealing and certainly promising, it is yet unclear if VR really provides a benefit here or if the added value is mostly due to the larger screen size.

3 IMPLEMENTATION OF MAP-BASED LIFELOG PHOTO BROWSER IN VR

Our system has been implemented in Unity VR for the HTC Vive head-mounted display using the two handheld controllers of the device for interaction. Because of the 360 degree view, there are various possibilities to visualise a world map in this 3D space. Our implementation features three options (see Fig. 2): (a) a flat map

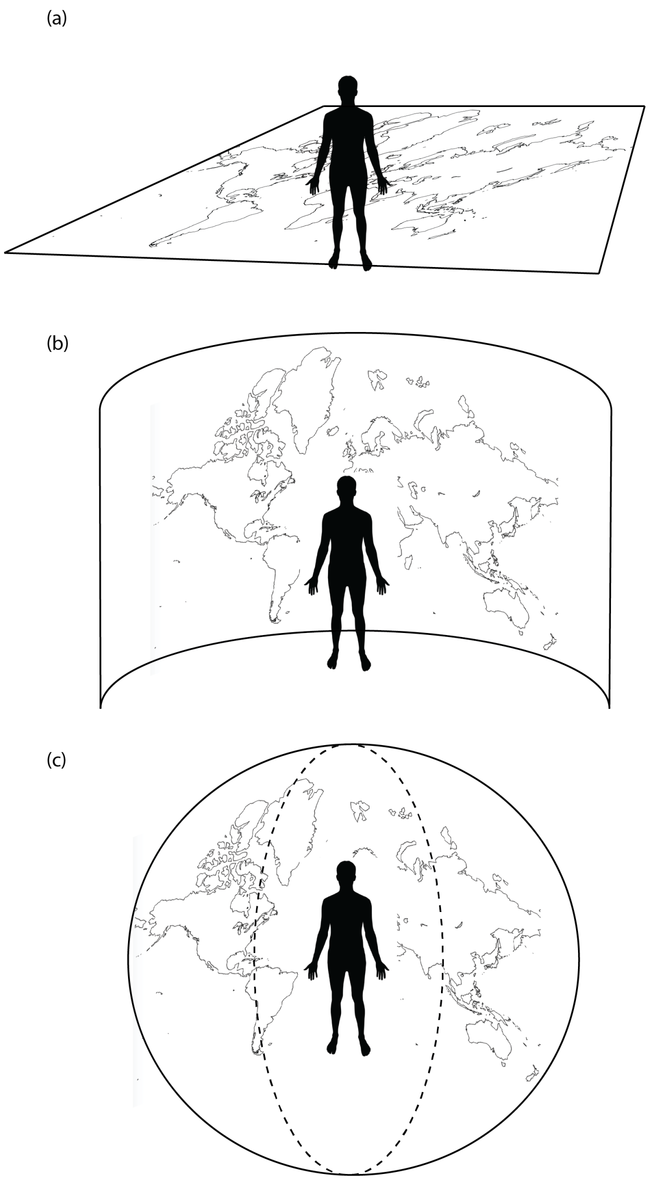

Figure 2: Different map visualisations in VR: (a) flat, (b) cylinder, (c) sphere (with mirrored map).

on the floor that can be explored by "walking" on it (subsequently called *flat map*, (b) a map that is "wrapped" around the user in a cylinder-like shape (subsequently called *cylinder map*), and (c) a map that is projected onto a sphere (subsequently called *sphere map*). Because the viewer is placed inside of the sphere in the latter case, the map is mirrored to create the familiar view that people are used to when looking at a 2D map of the world. Which of those three options is the most suitable for our purposes is the focus of a user study that is currently been performed.

To explore the map and the photos, there are currently two interaction modes implemented. One is *teleportation*, where users select a spot in the 3D space using one of the handheld controllers and are then "teleported" to that particular point. The other is *zooming*, where users get closer or further away from the map again using the handheld controllers. Both options allow users to get a more detailed view of the map by approaching it. Advantages and disadvantages of either of these interaction modes are currently being identified in the user study mentioned above.

Figure 3: Image representations for the cylinder map via labels (the top screenshot shows all data from the LSC 2018 dataset) and thumbnails of photos when getting closer (bottom right)

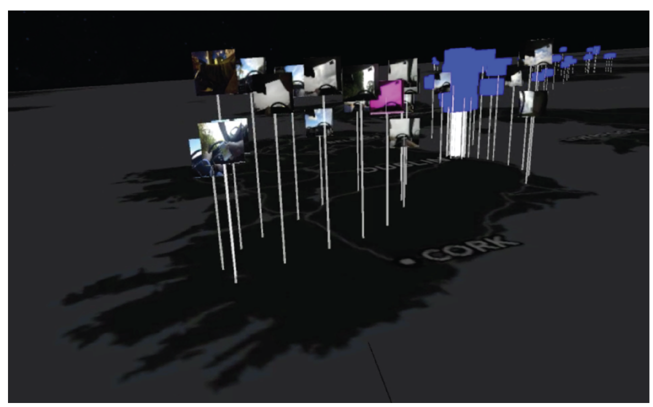

Figure 4: Image representations for the flat map.

Initially, images are represented by labels at the position where they have been taken using the GPS data recorded with them. These labels (blue squares) indicate that a photo has been taken at this location but do not provide any other information. Once the user approaches the map via teleportation or zooming, the labels get replaced by photos (random ones are selected when there are too many pictures with similar GPS location). Figure 3 shows related screenshots with the LSC dataset, which contains images from Northern Europe (mostly Dublin, UK, and Norway). Figure 4 shows a screenshot of the flat map, where labels and photos are put above the map for better visibility. More images become visible once a user gets even closer. Users can "pick up" a single image by selecting it with the controllers. The enlarged version of the photo can then be explored and looked at like a physical photo in the real world (see Fig. 5 for some examples).

Figure 5: Examples for exploration of individual lifelog photos.

4 APPLICATION TO LSC 2018 DATASET

For the Lifelog Search Challenge provides a dataset for the competition along with a set of so-called Known-Item Search (KIS) tasks. The dataset contains images captured by a single lifelogger on 27 days in 45 second intervals with about 1,500 images per day. The KIS tasks resemble a situation where the user has a clear idea of what the target (i.e., the "known item") is. In the competition, this is simulated by providing a textual transcription of the image or its context.

Obviously, our system is not intended to fulfil all possible search tasks, but it is only suited for situations where location plays an eminent role. We illustrate this with the following example topic from the LSC dataset 2018:

I am taking a photo of a white building with a unique blade-like design. The weather is cloudy and it is getting dark, being evening time. There are a number of buildings clearly visible in the image, including a hotel and a Norwegian style house. I had just walked from a sushi restaurant to the hotel where I was staying and I had taken the photo just before entering the hotel. A large yellow pipe is also visible in the image. Just before taking the photo, I had been walking beside the sea. This happened on a Wednesday.

The lifelog photo in question is shown in Fig. 5, bottom right. The first two sentences describe the content of the photo. With this information alone, our map-based visualisation will be of no help in finding it because there are no location-related cues (aside from the comment about the Norwegian style house, which makes it slightly more likely that it was taken in Norway). Yet, the statements that the lifelogger "walked from a sushi restaurant to the hotel", that it has been taken "just before entering the hotel" and that the lifelogger has been "walking beside the sea" just before are comments that do not only indicate a time order but also a spatial relationship. Still, a random user of our system will have a difficult time finding the photo on the map. Here it comes into play what we said in the introduction in Section 1, that is, that lifelogging data is highly personal. Given the description of the topic, it seems quite likely that the lifelogger has at least some vague idea where this event took place. That, together with the context of the sushi restaurant and the walk along the seaside should be enough hints for a successful, location-based access to the right photo. It should also be noted that the implied time order of the events at different locations can be helpful in the search process, as there are likely many labels in the sushi restaurant and the hotel after entering, plus a consecutive sequence of single shots from the walk along the seaside between these two locations. Our current and future work is exploring how such temporal information can be integrated into our map-based visualisation in the most beneficial way.

Obviously, there are many topics where our approach with a map-based search is not useful at all; be it because the lifelogger does not remember the correct location or because there are just too many photos taken in a particular area (e.g., all photos in the LSC 2018 dataset labeled with HOME or WORK). Yet, even in such cases, a combination with a more traditional approach, such as querying, could have benefits. Either, one could use the map to pre-filter the data and place the query only on photos within a certain region. Or a query could be used to pre-filter the data that is displayed on the whole map. A good example for the latter case would be the search for a particular person. After querying the person's name, photos of that individual will be shown at every spot in the world where the lifelogger met him or her.

It is noteworthy that the LSC data does contain images without GPS information, which are instead only generally labelled as either "Home" or "Work". This had to be done to protect the privacy of the lifelogger who provided the data. We could address this issue in our system by providing a separate way of accessing these photos, for example, via related labels that pop up at locations independent of the map. Yet, given the large amount of data on these particular locations, an additional filtering approach is needed and part of our current work in progress. Finally, it is important to keep in mind that the LSC 2018 only addresses Known-Item search tasks, where people have a concrete idea of that they are looking for. As stated in the introduction, we believe that our system is particularly well suited for leisure browsing and random exploration of one's personal lifelogging photos. Such a situation is not addressed by the LSC 2018, but a common usage of the data by private lifeloggers.

5 CURRENT AND FUTURE WORK

We presented a system that uses a map-based visualisation in VR for easy and intuitive access to geotagged lifelogging photos. While not suited for all kinds of search needs, it is useful for location-related searches and has high potential for leisure-based exploration of one's lifelog images. We discussed different implementations, for example, for the visualisation and interaction with the data. A detailed comparative study between those is currently in the works. Future work related to the visualisation include a better implementation of the zooming when approaching the map and an improved way to access images in areas that are too cluttered with many photos at close by locations. Maybe the most promising, but from a visualisation point of view also the most challenging extension of the system will be the combination with other approaches; in particular the aforementioned combination with a time-based visualisation and a traditional search approach such as querying.

REFERENCES

[1] Soumyadeb Chowdhury, Md Sadek Ferdous, and Joemon M Jose. 2016. A user-study examining visualization of lifelogs. In *Content-Based Multimedia Indexing (CBMI), 2016 14th International Workshop on.* IEEE, 1–6.

[2] Aaron Duane, Rashmi Gupta, Liting Zhou, and Cathal Gurrin. 2016. Visual insights from personal lifelogs. In *Proceedings of the 12th NTCIR Conference on Evaluation of Information Access Technologies, Tokyo.* 386–389.

[3] Aaron Duane and Cathal Gurrin. 2017. Pilot Study to Investigate Feasibility of Visual Lifelog Exploration in Virtual Reality. In *Proceedings of the 2nd Workshop on Lifelogging Tools and Applications.* ACM, 29–32.

[4] Aaron Duane and Cathal Gurrin. 2018. Lifelog exploration prototype in virtual reality. In *International Conference on Multimedia Modeling.* Springer, 377–380.

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

[6] Jae Ho Jeon, Jongheum Yeon, Sang-goo Lee, and Jinwook Seo. 2014. Exploratory visualization of smarphone-based life-logging data using Smart Reality Testbed. In *Big Data and Smart Computing (BIGCOMP), 2014 International Conference on.* IEEE, 29–33.

[7] Sanket Khanwalkar, Shonali Balakrishna, and Ramesh Jain. 2016. Exploration of large image corpuses in virtual reality. In *Proceedings of the 2016 ACM on Multimedia Conference.* ACM, 596–600.

[8] Sojung Kwak, Joeun Lee, and Jieun Kwon. 2015. An Analysis of Infographic Design for Life-Logging Application. In *Advances in Computer Science and Ubiquitous Computing.* Springer, 185–191.

[9] Gerald Schaefer, Mateusz Budnik, and Bartosz Krawczyk. 2017. Immersive browsing in an image sphere. In *Proceedings of the 11th International Conference on Ubiquitous Information Management and Communication.* ACM, 26.

[10] LSC website. 2018. Lifelog Search Challenge 2018. http://lsc.dcu.ie/

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

[12] Y Yang, H Lee, and C Gurrin. 2013. Lifelogging: New Challenges for Information Visualization on Mobile Platforms. *CHI'11 Extended Abstracts on Human Factors in Computing Systems* (2013).

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