

MemoryMesh – Lifelogs as Densely Linked Hypermedia

Lijuan Marissa Zhou

CLARITY Centre for Sensor Web Technologies,
School of Computing, Dublin City University
Dublin, Ireland
mzhou@computing.dcu.ie

Cathal Gurrin

CLARITY Centre for Sensor Web Technologies,
School of Computing, Dublin City University
Dublin, Ireland
cgurrin@computing.dcu.ie

ABSTRACT

Lifelogs are huge archives of multimedia data and consequently, they need to incorporate organization methodologies to fully exploit their potential. Early work in organizing lifelogs based on either video-style playback or event segmentation with browsing or basic search. In this work we propose that lifelogs can be represented as a densely linked hypermedia archive, called a MemoryMesh. We introduce how this can be constructed and the potential to improve retrieval performance.

Categories and Subject Descriptors

H.3.2 [Information Storage and Retrieval]: Information Storage – indexing methods.

General Terms

Algorithms

INTRODUCTION

The SenseCam introduced the research community to the potential of wearable cameras as lifelogging tools to gather media rich lifelogs for an individual. We define lifelogging 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*”, as used by Kitcher and Dodge [1]. The focus of our work is in developing the data structures to support these lifelog archives by exploring the linked-data potential of the multimedia content that exists in the archive.

A lifelog captured using a device such as a SenseCam can consist of more than 4,000 photos captured daily, along with hundreds of times more sensor readings. Very quickly, such a lifelog becomes too large to browse, so it becomes necessary to organize this data and to support search and retrieval. Doherty et. al. identified the ‘event’ as a suitable atomic unit of retrieval and proposed automatic segmentation of lifelog data into events, made accessible through a browsing interface [2]. However it was found that 75% of browsing effort fails to find a known event from a large lifelog [3]. Adding a search facility over automatic (sensor-based) event annotations reduced the failure-to-find error rate to

25% and the search-time by a factor of ten. In the MyLifeBits project [4], a database search mechanism is provided, which was shown to be effective at locating nuggets of information from Bell’s extensive archive.

Both Doherty and Bell & Gemmel’s work show the potential of search interfaces to lifelogs for supporting a user with an information need. In order to understand the potential information needs, Sellen & Whittaker [7], as a guide to future development of lifelogging technologies, have identified the five reasons why people would access their memories, and by association, their lifelogs. The five R’s of memory access are Recollecting, Reminiscing, Retrieving, Reflecting and Remembering Intentions.

While a flat database or text-index based representation of lifelog events can support the five R’s retrieval, it is our conjecture that a better organization structure and data access methodology will support more efficient and effective lifelog retrieval. Hence, we propose the MemoryMesh, which draws knowledge from WWW search, coupled with cognitive psychology, to develop a novel lifelog index structure that models the lifelog as a densely linked hypermedia archive. This allows for the application of new types of information retrieval concepts such as the WWW-inspired PageRank algorithm and supports multi-faceted browsing through the lifelog. We now discuss the MemoryMesh, in terms of its construction and potential for enhanced interaction with lifelogs.

MEMORYMESH – A LIFELOG DATA STRUCTURE

We propose that a densely linked hypermedia is a more suitable lifelog data structure than a flat database-based organisation. In effect we consider the lifelog to be more like a WWW structure, based on documents (events) and links. We know from prior research that exploiting the linkage structure of the WWW allowed the PageRank [6] algorithm to significantly enhance the effectiveness of large-scale information retrieval on the WWW. PageRank was deployed in the Google search engine and was considered an integral part of the ranking process. In the case of lifelogs, we propose that modeling lifelog as a linked data archive (as is done in the human memory system) will bring similar benefits, which will not only support more efficient and effective retrieval, but also better support real-time user interaction with lifelogs and the five R’s of memory access. By applying PageRank-style algorithms, the MemoryMesh will know the importance, novelty of events and will precalculate the links between them, thus making retrieval more efficient and effective.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SenseCam 2013 San Diego, USA

Copyright 2013 ACM 978-1-4503-2247-8 ...\$15.00

MemoryMesh Construction

On the WWW, the links between web pages are pre-existing and created by web page authors as they create new websites and webpages. Consequently, the WWW grows organically and algorithm can mine the latent qualitative judgments inherent in each WWW link. This enhances the effectiveness of search algorithms.

Since we propose that a lifelog is a densely linked hypermedia archive, we can model it as a graph. From mathematics, we know that a graph is an ordered pair $G = (V, E)$ comprising a set V nodes together with a set E of edges, which are 2-element subsets of V (i.e. connect two nodes). In the MemoryMesh, the graph is a representation of a set V of events where some pairs of events are symmetrically connected by edges (the set E). An edge e is a link that is created between events that are considered sufficiently similar. Each event is represented by a semantic annotation (typically from wearable sensors as in [3]), which forms the content for both query/retrieval and for linkage generation within the MemoryMesh. Given a set of events V , the strength of edges between the events are calculated and the most appropriate N edges are inserted into the graph. There are many methods for selecting N edges. Here we present three example methods:

- Visual similarity between SenseCam images from one event and SenseCam images in other events. Either low-level (e.g. SIFT or regionalised colour) or high-level (visual objects co-occurrence) visual similarity can be calculated and can inform the strength of links, with links above a threshold bring selected,
- Multi-axes similarity sources from lifelog metadata, such as events at the same location, time, noise level, actors involved, and many other sources of linkage evidence,
- External sources of similarity, by looking to semantic-web-style external sources of semantic data to identify real-world links between events.

In addition, the linkage model could be single or multi-layer. A single-layer mesh would allow for a single link between events whereas a multi-layer mesh would allow for multi-faceted browsing, which would provide additional flexibility to the MemoryMesh. Regardless of the layering, it is likely that a proposed linkage distribution would need to be adhered to when calculating a meaningful value for N , across single or multiple layers. As new events are continually inserted into the lifelog, the MemoryMesh would need to re-calculate linkages either dynamically or periodically.

RETRIEVAL FROM THE MEMORYMESH

Having the links between events in the MemoryMesh provides a number of benefits. Firstly, it supports real-time browsing for reminiscence and reflection, without the need to continually, dynamically generate links to related events. Secondly, the rich linked hypermedia allows for a user browsing session to be targeted into the best region of the lifelog to begin a linked reminiscence session or provided guided reminiscence tours. In addition, the pre-calculation of multi-layer linkages would make the MemoryMesh more flexible to new and novel use-cases. Finally, links between events in the MemoryMesh provides an additional source of evidence when retrieving information from the lifelog, which would allow for the application of common WWW algorithms such as HITS and PageRank, which have proven so useful on the WWW (another large-linked graph).

Applying a HITS [5] style algorithm to the MemoryMesh could identify the most important user-context/query related events from the MemoryMesh. Indeed exploring the top non-principal eigenvectors from the MemoryMesh (applying HITS techniques) could help to identify clusters of similar experience from within the lifelog. Employing PageRank could support the efficient selection of both mundane and novel events in the lifelog. As lifelogs increase in size, identifying novel events becomes increasingly important.

Enhancing the MemoryMesh to incorporate a multi-level event linkage model, it becomes possible to retrieve from, and browse through, the lifelog using various criteria, such as exploring the lifelog via user activity links, behavioral similarity, co-occurrence of similar objects, similar environmental context, or even the colour or textures of the SenseCam images grouped into events. It is our conjecture that moving from the flat collection of annotated events into a linked hypermedia will produce a more useful lifelog.

CONCLUSION

We have proposed and presented the MemoryMesh, which organises a lifelog as a densely linked hypermedia. We suggest that applying a MemoryMesh organization model to a lifelog will increase the flexibility and usefulness of the lifelog and support new types of user interaction and better support real-time user search and retrieval. Our next step in this research is building and evaluating a real-world implementation of the MemoryMesh.

ACKNOWLEDGEMENTS

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under grant no. "SFI 11/RFP.I/CMS/3283".

REFERENCES

- [1] Martin Dodge and Rob Kitchin. "Outlines of a world coming into existence": Pervasive computing and the ethics of forgetting." *Environment and Planning B* 34.3 (2007): 431-445.
- [2] Aiden R. Doherty and Alan F. Smeaton. "Automatically segmenting lifelog data into events." *WIAMIS'08, the Ninth International Workshop on Image Analysis for Multimedia Interactive Services*. Klagenfurt, Austria, 2008.
- [3] Aiden R. Doherty, Katlin Pauly-Takacs, Niamh Caprani, Cathal Gurrin, Chris J.A. Moulin, Noel E. O'Connor and A. F. Smeaton. Experiences of aiding autobiographical memory using the SenseCam. In: *Human-Computer Interaction*, 27(1-2):151-174, 2012.
- [4] Jim Gemmell, Gordon Bell and Roger Lueder. "MyLifeBits: a personal database for everything." *Communications of the ACM* 49.1 (2006): 88-95.
- [5] Jon Kleinberg, 1999. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 5 (1999), 604-63
- [6] Larry Page, Sergey Brin, Rajeev Motwani and Terry Winograd. (1999) *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report. Stanford InfoLab.
- [7] Abigail J. Sellen, and Steve Whittaker. "Beyond total capture: a constructive critique of lifelogging." *Communications of the ACM* 53.5 (2010): 70-77.