An Investigation into Event Decay from Large Personal Media Archives

Aiden R. Doherty and Cathal Gurrin and Alan F. Smeaton
CLARITY: Centre for Sensor Web Technologies
Dublin City University, Ireland.
adoherty@computing.dcu.ie

ABSTRACT

With the growth of digital lifelogging technologies there are challenges in terms of detecting and annotating real world events from this multimedia lifelog data. In this paper we use the SenseCam, a passively capturing wearable camera, worn around the neck, which captures about 3,000 photos per day, thereby creating a personal lifelog or visual recording of the wearer's life, which could be helpful as a human memory aid. For such a large amount of visual information to be of any value, it needs to be structured into semantic events. In this paper we are particularly interested in how a user's perceptions of real world events decays over time. In particular we investigate several questions including whether data owners have different perceptions of event boundaries to non-owners, whether the passage of time changes what we believe to be events and if so then do we forget about the weakly defined original events. We carry out these investigations using real visual lifelog data gathered and annotated, twice, by three users.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval; E.m [**Data**]: Miscellaneous; H.1.m [**Models and Principles**]: Miscellaneous

General Terms

Design, Experimentation, Human Factors

Keywords

lifelogging, SenseCam, event decay, personal life semantics

1. INTRODUCTION

These days many aspects of our lives are some way monitored or logged digitally. From shopping online to browsing the Internet, computers store a log of our actions. We've come to accept, or maybe just ignore, this massive surveillance because it brings us benefits (e.g. better recommendations when shopping online or our browser's history feature for example). We may even have a more secure feeling when we know there is CCTV present, we get itemised billing from phone companies, and we get convenience and even loyalty bonuses with some of our regular purchases.

In recent years we have seen some people take this a step further and start logging many aspects of their own lives, from writing blogs to maintaining their own personal MEMEX [6] of life experiences. Such MEMEX style recording of different aspects of our daily life, in digital form is called LifeLogging. It is a form of reverse surveillance, sometimes termed *sous*veillance, referring to us, the subjects, doing the watching of ourselves. Lifelogging can take many forms, such as the application which runs on a mobile phone to 'log' all your activities and then present all those activities in a calendar format [15] to the MEMEX style recording of Bell & Gemmell [3].

Memory is a hugely important aspect of our lives, so much so that the important events in our lives define our being. Conway argues that autobiographical memory affects us to the very core and defines who we are [7]. In fact shared memories are a part of our social self, and define who we are as a group of people e.g. it is the memories of shared experiences with one's family that create the close ties and bonds, and reminiscing these memories together not only improves one's memory, and the memory of the family, but also creates a tighter social bond between those people. Given that populations are generally getting older and that the number of those with memory impairments is ever increasing [2], information scientists are now exploring using technology to aid memory.

Human memory is recognised as being far from perfect and through the ages we have devised many methods to aid it [2]. In terms of individuals remembering aspects/activities of what they themselves have been doing (autobiographical memory), the written diary has been popular for years. More recently, with the advent of digital technology and the Internet, many individuals have begun to maintain online blogs to detail aspects of their activities. One of the goals of *lifelogging* is to move individuals towards software-aided total memory/experience recall [3]. The lifelogging community attempt to achieve this goal by automatically capturing electronic data from numerous sources (Figure 1), e.g. web pages visited, e-mails sent and received, audio recordings of conversations, etc.

EiMM'09, October 23, 2009, Beijing, China.

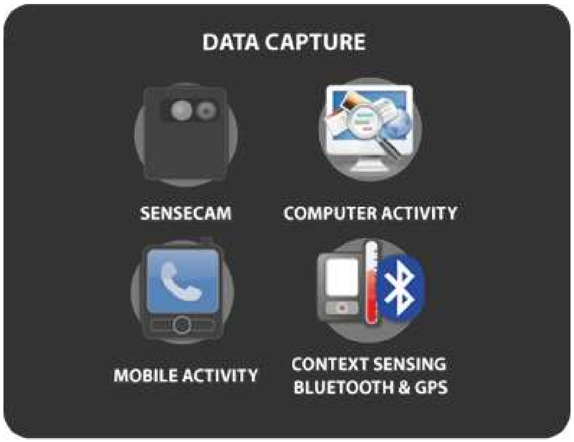

Figure 1: An overview of the many life experiences that can now be captured

The aspect of lifelogging that we are concerned about in this paper is visual lifelogging, i.e. an individual capturing daily activities through the medium of images or video. It is particularly important to do this through passive capture, meaning the use of devices that automatically capture images or video, thus requiring no conscious effort by the user to take images, which leads to him or her acting in a more natural manner. The importance of visual imagery to memory is well established [5]) and we will show later that cue-based recall is very helpful in neural retrieval mechanisms, in particular when those cues are lifelog images taken from one's own perspective.

In this paper we explore a number of questions related to lifelog events, namely:

1. Do owners of lifelog data have a different perception on what constitutes an event on their data than other people do on that owner's data ?
2. How many events are generated in a typical day and what is the typical duration of a lifelog event ?
3. With the passage of time (up to 2 years later) does an owner's view of real world event boundaries on a fixed day of lifelog data, change ?
4. Where the owner's perception does change with the passage of time, are only the weakly identified original events forgotten about ?

We explore these issues through experiments on real visual lifelog data, captured by 3 users over a period of time.

2. BACKGROUND AND RELATED WORK

2.1 Lifelogging

The visual lifelogging community has mainly concentrated on the challenges of miniaturising lifelog devices, and on how to store and manage *vast* quantities of lifelog data. To enable increased non-intrusive capture of visual lifelog material, Microsoft Research in Cambridge, UK, have developed a device known as the SenseCam. The SenseCam is a small wearable device that passively captures a person's day-to-day activities as a series of photographs [12]. It is typically

Figure 2: The Microsoft SenseCam

worn around the neck, and so is oriented towards the majority of activities which the user is engaged in (see Figure 2). Anything within the view of the wearer can be captured by the SenseCam. The device requires no manual intervention by the user as its on-board sensors detect changes in light levels, motion and ambient temperature and then determine when is appropriate to automatically take a photo. For example, when the wearer moves from indoors to outdoors a distinct change in light levels will be registered and photo capture will be triggered. On average the SenseCam captures 3 photos per minute.

Depending on configuration, the SenseCam takes between 2,500 and 5,000 images in a typical morning-to-night day, and as a result a wearer can very quickly build large and rich photo collections. Within just one week, nearly 30,000 images may be captured and over a year the lifelog photoset could grow to well over one million images. The potential benefits of this are numerous and include the ability for a user to easily record events without having to sacrifice their participation, aiding memory and personal recall, and providing insight into a person's life and activities [3]. Notably, preliminary work between Microsoft Research and Addenbrooke's hospital in Cambridge, U.K indicates that a rich photo lifelog can dramatically improve memory and recall for individuals with neurodegenerative memory problems [4].

2.2 Lifelogging events & human memory

Recently, the focus has been placed on how to manage and organise large quantities of lifelogging data in order for it to act as effective memory retrieval cues. To effectively exploit personal lifelog collections, it is necessary to have a basic understanding of how human memory operates.

Atkinson and Shiffrin proposed a multi-store model of the human memory system in 1968 [1]. The model states that there are 3 distinct memory systems present in the human brain, namely one to deal with sensory inputs, one that relates to short-term memories, and finally a distinct system that deals with long-term memories only. Since 1968 more advanced, and complex, memory systems have been proposed, but for the purposes of this paper the Atkinson and Shiffrin model is sufficient. The sensory memory store stores information for fractions of a second, while short-term memories are only retained for between 0.5 and 4 seconds [11]. There are two types of declarative long-term memories: semantic and episodic/autobiographical memories. Semantic memories are knowing about facts e.g. Paris

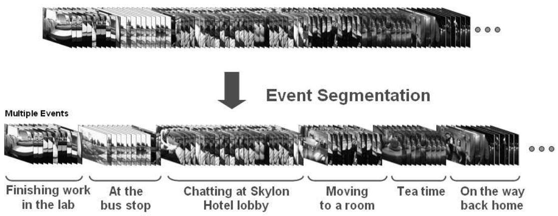

Figure 3: Overview of event segmentation

is the capital of France, England won the soccer world cup in 1966, China is the most populous country in the world, etc. Episodic/autobiographical memories refer to personal experiences e.g. remembering when one's child first walks, recalling a conversation with one's friends from the previous evening, etc. [19]. Given that autobiographical memories are of personal experiences, in this paper we investigate whether owners of SenseCam data have noticeably different views on that data than non-owners.

One of the most important facets related to autobiographical memory, in terms of technology research, is the evidence pointing towards "cued recall" leading to retrieval of memories that can't be accessed via "free recall" [17]. This indicates that users prefer to be prompted with cues (e.g. SenseCam images), rather than being asked to retrieve memories from scratch. However what exactly are we trying to retrieve from the memory system, how are memories stored ?

Zacks, who studies how representation in the brain works, states that humans store memories as events *"... segmenting ongoing activity into events is important for later memory of those activities ... "* [21]. Staying in the field of cognitive science, Newtson, Engquist, & Bois note that *"... breakpoints* [between events] *tend to correspond to points at which the most physical features of the action are changing ... "* [16].

This research suggests that the area of lifelogging should rely heavily on taking an event-centric approach to representing the surrogate human memories that lifelogging (using a device such as a SenseCam) generates. Given that a lifelog is by definition a long-term storage device, understanding the effects of time on automatic identification of event boundaries is an interesting and perhaps crucial question to examine. In this paper we investigate the effects of time on our understanding of what characterises events in a human digital memory.

2.3 EVENT SEGMENTATION

Previous work has motivated the need to automatically divide visual lifelogs into discreet events [9, 14]. This task is somewhat similar to that of scene boundary detection in video as events or activities have an inherent underlying semantic meaning e.g. see Figure 3.

In previous work we carried out an extensive evaluation to optimise event segmentation for lifelog images from the SenseCam [9]. Essentially an event segmentation approach attempts to identify periods of visual or sensory change, and identifies those occasions as most likely to be boundaries between distinct events or activities. Now in this paper we retrospectively investigate the groundtruth created for the aforementioned automated segmentation evaluation. In principle we are most interested in two questions: 1) Can people other than the owner of SenseCam images create similar event boundary groundtruths, and 2) Will the owners' perceptions of event boundaries (on the same data) change quite noticeably with the passage of time?

3. EVENT DECAY

> *... the susceptibility of a long-term memory trace to decay in storage is assumed to depend upon both its strength and its resistence ...*
>
> Wickelgren [20]

One of our main interests in this paper is to investigate the phenomenon of "event decay". Traditionally the notion of event decay is considered when, with the passing of time, the human mind forgets about an event that has happened. This process occurs quite naturally and in the majority of cases is at worst frustrating, but in instances of dementia this poses a major problem, and is an area that is of great interest to the cognitive neuropsychology community. In the field of lifelogging, we can observe another type of "event decay", where people's perception changes (over time) of what exactly constitutes a boundary between two distinct and adjacent semantic events. For example, if one analyses their lifelog a day after capture, and subsequently analyses the same day one (or more) years later, will the person make the same manual segmentation of events ?

Our perception of events changes with the passing of time - events can retrospectively become more important e.g. the stranger you briefly talked to may later become a valued friend (or partner) and this was the first meeting. So initially speaking to this person may not count as an event. Even reviewing the images a few months later this may not be considered a distinct event, but it could happen that many months/years later, people may wish for this conversation to be regarded as a distinct semantic event in its own right.

The converse of this may be a brief meeting with a research colleague to discuss progress on an important paper. On initial review of one's SenseCam images shortly later, it may appear that this event quite likely deserves to be recognised as a distinct event. However perhaps with the passing of time, as one reviews SenseCam images of this day many months later, they may not believe that this brief conversation any longer merits recognition as a distinct event.

In this paper we investigate the issue of "event decay", whereby we look into how the perceptions of three individuals changes after a year between engaging in a manual segmentation process for a full week of lifelog images.

4. EXPERIMENT OVERVIEW

We now introduce the participants and datasets that we used to carry out our experiments. We elaborate on how users were instructed to create event boundary groundtruths, how much annotation was carried out to investigate the number and duration of events, and also we discuss the data gathered by three users to investigate the nature of event decay.

4.1 Creating an event boundary groundtruth

To create a groundtruth, users were asked to review *their own* SenseCam image collections and manually mark the boundary images between all events. As motivated earlier in Section 2.2 it is normally considered important that the

User	Age	Total Num Images	Avg Daily Duration
1	30-35	80,934	13h 08m
2	20-25	76,810	9h 27m
3	25-30	44,447	10h 41m
4	20-25	27,929	7h 45m
5	30-35	41,043	9h 15m

Table 1: Experimental setup: Data gathered by five male computing researchers to analyse the number of lifelog events captured over one month

owner of the images is the person to identify the boundary between semantic events, as the owner is the only person who can truly understand the importance of seemingly minor episodes and encounters in the event segmentation process, as it is *their* images of *their* life. Later in Section 5.1 we will investigate the validity of this consideration. It was stressed to the users that they should judge an event boundary based on what it semantically means to them. Given that one's perception of what constitutes an event boundary may change with the passage of time, we provided our users with instructions asking them to create the groundtruth with the following two levels of granularity:

> *". . . The number 1 bookmarks/boundaries will be those that you feel that are definitely an event. e.g. entering a restaurant, getting into your car, starting a 15 minute talk with an old friend you haven't seen in a while, going to lunch, arriving into the office in the morning, etc.*
>
> *The number 2 bookmarks/boundaries will be those that you think may possibly be an event, i.e. if you were reviewing that day at another time, would you possibly like to see this as an event? Possible situations include:*
>
> - *sitting at desk then going to whiteboard to discuss with work colleague (you decide if you talked to them long enough to merit an event boundary)*
> - *talking to someone in corridor*
> - *at table at home then joined by someone else*
> - *out for walk but then talking to someone in the middle of it*
> - *sitting in bus/plane and then starting to read paper/magazine . . . "*

4.2 Analysing the number and duration of events captured in a lifelog

In May 2007, as part of experiments to train an event segmentation system we organised the collection of one month's worth of SenseCam images from five users [9]. In all this equated to 271,163 images, as detailed in Table 1.

4.3 Event Decay: Event decay observations

In September 2008 and once again in May 2009[1], three of our five users[2] repeated the segmentation groundtruthing task on one week of data from the original dataset. This allows us to investigate the event decay properties related to these images, and Table 2 provides a breakdown of the week's images from this groundtruth, where 45,016 images were captured by our 3 users. Unfortunately two of our users moved abroad in the intervening time, thus highlighting the unique nature of the data and the difficulty of obtaining this for even just three people.

[1]one user was busy until May

[2]Users 1, 2, and 5 from Table 1

To analyse the perception of event decay in the next section, we examine where the event boundaries marked in both groundtruths occur. If any given boundary in the first groundtruth has an associated boundary within ± 5 minutes in the second groundtruth, we consider this a "hit". While this may initially appear a very coarse measurement, especially in relation to the TRECVid shot boundary detection task [18], it must be considered that "sessions" of lifelog data can be captured over 15+ hour periods which then must relate to an individual's semantic interpretation of events. Taking this into account, it is our belief that accuracy to within ± 5 minutes is sufficient.

5. RESULTS AND DISCUSSION

We now discuss the results of our experiments, which allow us to note that only owners of lifelog data truly understand the semantic and worldly significance of *their* data, to find the number and duration of captured lifelog events and finally we look into the nature of perceptions of event decay in lifelogs.

5.1 Owner Semantic Knowledge: Only the owner can review his/her own lifelog data

> *". . . It is usually considered important that the owner of the images is the one to identify the boundary between semantic events, as the owner is the only person that can truly understands the importance of seemingly minor episodes and encounters in the event segmentation process, as it is their images of their life . . . "*

To qualify the validity of this statement, user 2 in Table 2 agreed to share 12,522 of his images collected over an 8 day period. This user, and three other participants (marked A, B, & C) then manually identified event boundaries as per the guidlines in Section 4.1. Table 3 outlines the consistency between users' judgements on event boundary images. It is immediately obvious that the consistency of judgements between the 3 non-owner participants B:C (53%), A:C (51%), and A:B (49%) is much closer to each other than to any judgements by the SenseCam owner (43% to C, 37% to B, and just 31% to A). It is also interesting to note the difference in the actual number of marked events too, whereby the non-owners identify more candidate events than the owner (144, 126, 151 vs. 117). These 2 observations support the fact that judgements on lifelog data should be made by the owner of that data given that a level of semantic understanding of the gathered data is required, which is quite natural given that these events normally trigger autobiographical memories (Section 2.2). This also indicates that automated techniques should be evaluated against semantic judgements made by the owner of any given lifelog images.

User	Age	Images Taken	First Groundtruth	Second Groundtruth	Total Num Images	Avg Daily Duration	Avg Num Images/Day
1	30-35	Apr 2007	May 2007	May 2009	18,120	13h 20m	2,589
2	20-25	Oct 2006	May 2007	Sep 2008	12,894	9h 37m	1,611
3	30-35	Nov 2006	May 2007	Sep 2008	14,002	10h 01m	1,750

Table 2: Experimental setup: Data gathered by three male computing scientists to analyse event decay

	Num Events Marked	Owner	A	B	C
Owner	117		31%	37%	43%
A	144	31%		49%	51%
B	126	37%	49%		53%
C	151	43%	51%	53%	

Table 3: Event boundary consistency between different judges on the same data

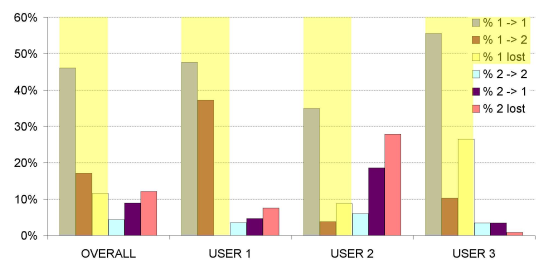

Figure 4: Shift pattern of judgements between two groundtruths. Original type 1 boundaries are shaded.

5.2 Number Events: Number and duration of captured lifelog events

After our 5 users manually segmented their 271,163 images into events, we were left with an average of 19.1 events /activities/episodes per user, per day, with an average event duration of 32 minutes (see Table 4).

It is interesting to compare these findings to Kahneman *et. al.* who detail the activities of 909 women where they carried out a questionnaire, and prompted the women to list down their activities and how they felt during these activities [13] . They discovered that there was an average of 14.1 episodes during the day and that the average episode duration was 61 minutes (vs. 19.1 events and average duration of 32 minutes for us). We believe that more events would have been identified by Kahneman *et. al.* if the users had access to digital lifelog data of their days, as they would have improved recall. The most obvious difference is the duration of the respective episodes, for which there may be two explanations: 1) The differing lifestyles between the groups of subjects, and 2) Our subjects had access to a visual lifelog thus almost every minute detail of their day was recorded and thus minor events would not be forgotten about.

5.3 Event Decay: Overlap of boundaries between 2 groundtruths

Table 5 provides an overview of the number of event boundaries identified in the first and second groundtruthing tasks, and also the number of overlapping events between them (63% overall). An initial indicator of event boundary decay taking place is that 14% less boundaries overall (406 vs. 472) were identified by the users in the 2nd groundtruthing task, although this wasn't the case for user 3 (117 vs. 134). Table 6 graphically illustrates that a significant number of the event boundaries overlap. User 1 has a very significant overlap in the number of event boundaries identified, while user 2 has obviously discarded many of the originally annotated event boundaries, and lastly user 3 has significant overlap but also a number of events that were unique to either the first or second groundtruthing stage. Indeed while we note that there is some degree of event decay, in that less boundaries are identified overall across our 3 users, we can see that this effect is not especially strong.

5.4 Event Decay: Detailed analysis on strong and weakly identified boundaries

As described in the previous section, users were asked to identify event boundaries and judge their significance by a graded relevance scale, i.e. '1' where users are certain that a semantic event boundary has taken place, or '2' where users are not convinced that a semantic event boundary has taken place but they could see arguments on why it could be considered so. Figure 4 provides a detailed analysis of how users' perceptions of event boundaries changed. It can be seen that there are many more type '1' event boundaries identified (71% in total), and also we can see from the left-hand column that the majority of *all* events that were originally identified as strong '1' event boundaries, remained as strong '1' boundaries the second time around too. Also, it can be noted that the 2nd largest number of events are those that were originally identified as strong '1' event boundaries but "downgraded" to weaker type '2' event boundaries when groundtruthed again. It is quite interesting to note that the original type '2' boundaries produce an unexpected result in Figure 4 in that more of the boundaries migrate to type '1' rather than staying at type '2', although more in line with expectations even more again original type '2' events are quite simply not identified second time around. It is interesting to note however that as many original type '1' event boundaries are not identified 2nd time around as type '2' boundaries (approximately 10% each of *all* the events). From this we can take that the concept of binary division of event types is not a concept we consider to be persistent over time.

In Figure 5 we consider event boundaries that were originally annotated as strong type '1' boundaries. In looking for strong indicators of event boundary decay we would expect that a large percentage of the original type '1' boundaries would remain, with a reasonable amount "downgraded" to type '2' boundaries, and very few completely lost. Indeed Figure 5 somewhat follows this trend, although users 2 and 3 appear to have "lost" over 20% of events.

User	Age	Avg Daily Duration	Total Num Images	Groundtruthed Events	Avg Events Per Day	Avg Images Per Day	Images Per Event
1	30-35	13h 08m	80,934	995	28	2,312	81
2	20-25	9h 27m	76,810	875	18	1,600	88
3	25-30	10h 41m	44,447	348	17	2,116	128
4	20-25	7h 45m	27,929	329	13	1,117	85
5	30-35	9h 15m	41,043	439	19	1,783	93

Table 4: Results: data segmented by our five male computing scientists into semantic events.

User	Age	Num Events 1st GT (images/event)	Num Events 2nd GT (images/event)	Num Overlapping Events	1st GT recall	2nd GT Precision	Percentage Overlap
1	30-35	172 (105 i/e)	155 (116 i/e)	144	0.84	0.93	79%
2	20-25	183 (70 i/e)	117 (110 i/e)	113	0.62	0.97	60%
3	30-35	117 (119 i/e)	134 (104 i/e)	83	0.71	0.62	49%
Total		472	406	340	0.72	0.84	63%

Table 5: Results: Event decay analysis on event boundary groundtruth (GT) statistics

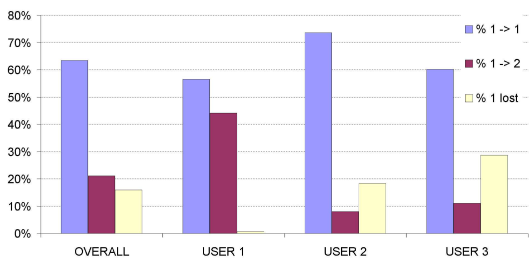

Figure 5: Shift pattern of judgements between two groundtruths where original boundaries were type 1

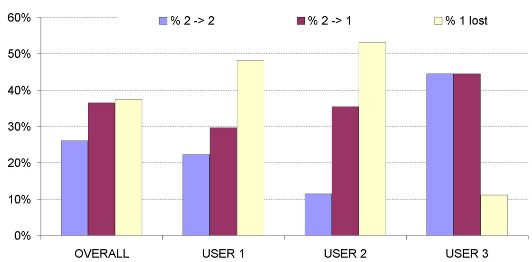

Figure 6: Shift pattern of judgements between two groundtruths where original boundaries were type 2

Then, in Figure 6 we consider event boundaries that were originally annotated as weak type '2' boundaries. In looking for strong indicators of event boundary decay we would expect that a large percentage of the original type '2' boundaries would simply decay and become lost second time around. Interestingly though a large percentage of events were "upgraded" to type '1' event boundaries, and even many remained as type '2' boundaries. Unsurprisingly a large number of events were "lost" second time around, especially for users 1 and 2, but overall we believe that there isn't sufficient evidence to point toward a strong value placed on such event boundary type segmentation decay in the minds of our users.

6. CONCLUSIONS AND FUTURE WORK

In this work we have provided a motivation on why lifelog data should be automatically segmented into manageable segments by identifying the boundaries between different daily events. This takes advantage of the fact that *"... segmenting ongoing activity into events is important for later memory of those activities ... "* [21]. Here we investigated a number of questions related to lifelog events:

1. **Do owners of lifelog data have a different perception on what constitutes an event on their data than other people do on that owner's data ?** The answer we found is that owners of lifelog data are significantly better judges (than non-owners) on *their* data as *they* have the best knowledge of the semantic meaning, and worldly context, of that data (Section 5.1). We believe this is an indication that any automated techniques applied to lifelog data should be measured against semantic groundtruths e.g. [9].

2. **How many events are generated in a typical day and what is their typical duration ?** On data gathered by male computing research scientists, our answer is 19.1 events per user, per average day, (Section 5.2) and these are 32 minutes per average event (Section 5.2).

3. **With the passage of time (up to 2 years later) does an owner's perception of event boundaries change on a fixed day of lifelog data ?** Our answer is that while we observed evidence of event boundary decay (only 63% overlap), we do not believe that this poses as large a challenge for maintaining a long-term lifelog as we had expected, e.g. for determining event boundaries automatically [9], determining the interestingness of events [10], or indeed in retrieving events from a lifelog [8](Section 5.3)

4. **Where the owner's perception does change with the passage of time, are only weakly identified original events forgotten about ?** The answer is that 36% of original "weak" boundaries are lost, whereas just 20% of original "strong" boundaries are lost (Section 5.4).

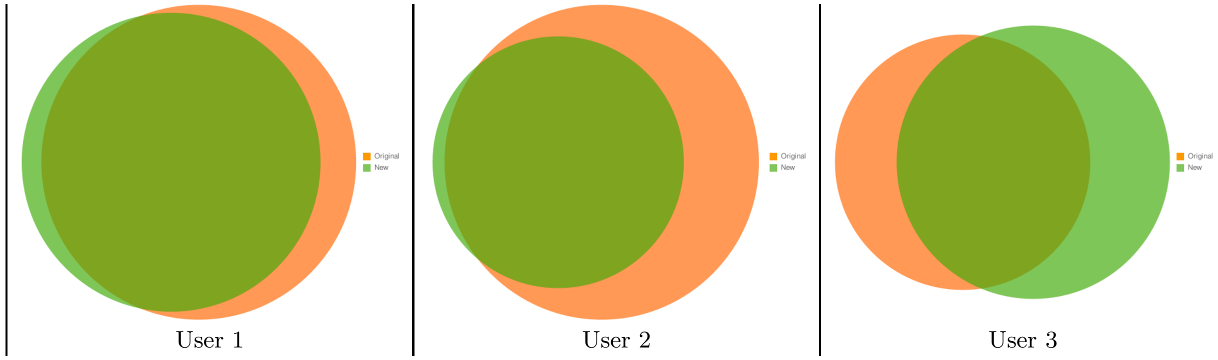

Table 6: Results: Graphical illustration of boundary overlap between groundtruths (original judgements in orange, new judgements in green)

The work we reported here points to work we should now undertake to use these insights into event decay as part of lifelog search and browsing systems. In particular, when presenting visual lifelog data from some time in the past, these insights will inform us on how such older should be presented.

Acknowledgements

Science Foundation Ireland under grant 07/CE/I1147, and Microsoft Research (equipment grant number 2007-056). We also would like thank Hyowon Lee, Shane McGrath, Mohiuddin (Lincoln) Patoary, and Pragathi Shetty for assisting in our annotation effort.

7. REFERENCES

[1] R. Atkinson and R. Shiffrin. Human memory: A proposed system and its control processes. *The psychology of learning and motivation*, 2:89–195, 1968.

[2] A. Baddeley, editor. *Your Memory: A User's Guide.* Carlton Books, 2004.

[3] G. Bell and J. Gemmell. A digital life. *Scientific American*, 2007.

[4] E. Berry, A. Hampshire, J. Rowe, S. Hodges, N. Kapur, P. Watson, G. B. G. Smyth, K. Wood, and A. M. Owen. The neural basis of effective memory therapy in a patient with limbic encephalitis. *Neurology, Neurosurgery, and Psychiatry with Practical Neurology*, 80(3):582–601, 2009.

[5] W. F. Brewer. *Practical Aspects of Memory: Current Research and Issues*, chapter Qualitative analysis of the recalls of randomly samples autobiographical events, pages 263–268. Wiley, 1988.

[6] V. Bush. As we may think. *The Atlantic Monthly*, 176(1):101–108, Jul 1945.

[7] M. A. Conway. Memory and the self. *Journal of Memory and Language*, 53(4):594–628, 2005.

[8] A. R. Doherty, C. O. Conaire, M. Blighe, A. F. Smeaton, and N. E. O'Connor. Combining image descriptors to effectively retrieve events from visual lifelogs. In *MIR '08: ACM international conference on Multimedia information retrieval*, pages 10–17, New York, NY, USA, 2008. ACM.

[9] A. R. Doherty and A. F. Smeaton. Automatically segmenting lifelog data into events. In *WIAMIS: 9th International Workshop on Image Analysis for Multimedia Interactive Services*, pages 20–23, Washington, DC, USA, 2008. IEEE Computer Society.

[10] A. R. Doherty and A. F. Smeaton. Combining face detection and novelty to identify important events in a visual lifelog. In *CIT: 8th International Conference on Computer and Information Technology, Workshop on Image- and Video-based Pattern Analysis and Applications.*, Washington, DC, USA, 2008. IEEE Computer Society.

[11] D. Elsweiler. *Supporting Human Memory in Personal Information Management*. PhD thesis, Department of Computer and Information Sciences, University of Strathclyde, 2007.

[12] S. Hodges, L. Williams, E. Berry, S. Izadi, J. Srinivasan, A. Butler, G. Smyth, N. Kapur, and K. Wood. Sensecamă: A retrospective memory aid. In *UbiComp: 8th International Conference on Ubiquitous Computing*, volume 4602 of *LNCS*, pages 177–193, Berlin, Heidelberg, 2006. Springer.

[13] D. Kahneman, A. B. Krueger, D. A. Schkade, N. Schwarz, and A. A. Stone. A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306:1776–1780, 2004.

[14] W.-H. Lin and A. Hauptmann. Structuring continuous video recordings of everyday life using time-constrained clustering. In *Multimedia Content Analysis, Management, and Retieval Ű SPIE-IST Electronic Imaging*, volume 6073, pages 111–119, 2006.

[15] A. Myka. Nokia lifeblog - towards a truly personal multimedia information system. In *MDBIS - Workshop des GI-Arbeitkreises "Mobile Datenbanken and Informationsysteme"*, pages 21–30, 2005.

[16] D. Newtson, G. Engquist, and J. Bois. The objective basis of behavior units. *Personality and Social Psychology*, 35(12):847–862, 1977.

[17] J. E. Purdy, M. R. Markham, B. L. Schwartz, and W. C. G. (2nd edition), editors. *Learning and Memory*. Thomson Learning, 2001.

[18] A. F. Smeaton, P. Over, and A. R. Doherty. Video shot boundary detection: Seven years of trecvid activity. *Computer Vision and Image Understanding (in press)*, 2009.

[19] E. Tulving. *Basic mechanisms in cognition and language*, chapter Neurocognitive processes of human

memory, pages 261–281. Elsevier, 1998.
[20] W. A. Wickelgren. Trace resistance and the decay of long-term memory. *Mathematical Psychology*, 9:418–455, 1972.
[21] J. M. Zacks, N. K. Speer, J. M. Vettel, and L. L. Jacoby. Event understanding and memory in healthy aging and dementia of the alzheimer type. *Psychology and Aging*, 21(3):466–482, 2006.