

# Exploring the Technical Challenges of Large-scale Lifelogging

Cathal Gurrin, Alan F. Smeaton, Zhengwei Qiu

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

Aiden Doherty

CLARITY and Dept. of Public Health  
University of Oxford  
Oxford, U.K.  
aiden.doherty@dph.ox.ac.uk

## ABSTRACT

Ambiently and automatically maintaining a lifelog is an activity that may help individuals track their lifestyle, learning, health and productivity. In this paper we motivate and discuss the technical challenges of developing real-world lifelogging solutions, based on seven years of experience. The gathering, organisation, retrieval and presentation challenges of large-scale lifelogging are discussed and we show how this can be achieved and the benefits that may accrue.

## ACM Classification Keywords

H.3.7 Information Storage and Retrieval: Digital Libraries—*systems issues, user issues*

## Author Keywords

Lifelogs, Personal Life Archives, Multidisciplinary, SenseCam

## INTRODUCTION

Recent advances in sensing, search and interaction technologies have helped to bring us to a point where anybody with a cell-phone or a custom off-the-shelf device (such as OMG Life's Autographer or Google Glass) can engage in a process of lifelogging. Lifelogging, as defined by Kitcher and Dodge is "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" [14]. Such lifelogs can contain information about all the activities an individual participates in, such as where they go, what they do there, who they speak with, what they see and what information they access, in short, a complete digital life diary. Given appropriate power-efficient and non-intrusive capture technologies, coupled with a new generation of data organisation techniques and semantic multimedia search technologies, the potential for lifelogging is clear. With the same ease that we execute a Google search now, we should be able to locate any nugget of information from

our past experiences, such as a phone number, a location, an entire event in great multimedia detail, or even perform an analysis of lifestyle trends over many years.

Maintaining rich lifelogs, has the potential to influence our healthcare, working lives, education and social activities [5] and will give new levels of self-awareness to the quantified individual. Specifically, some immediate obvious benefits for the individual include; the easy sharing of natural life experiences, enhanced productivity, increased knowledge for personalised wellness, better understanding of the functioning of human memory and a better understanding of the associations between lifestyle, environmental context and mortality. However, to achieve such benefits, there are a myriad of challenges to be overcome; ranging from semantically rich data capture, to data mining, knowledge extraction and retrieval tools and finally, the provision of appropriate interaction methodologies. With this in mind, we present the aims of this paper:

- Provide a brief overview of past efforts to capture and organise lifelogs,
- Propose how lifelogs may be captured on cell phones, SenseCams or other wearable devices; how they should be organised for subsequent use,
- Discuss the challenges in providing effective user access to data from lifelogs, and
- Motivate the benefits of post-capture analysis and real-time context-aware access to lifelogs.

In the rest of this paper we will briefly describe the overall challenges before discussing the particular requirements for data capture, organisation, retrieval and presentation. Finally we motivate the benefits of lifelogs.

## A BRIEF HISTORY OF LIFELOGGING

As far back as 1945 Vannevar Bush introduced the world to the Memex, a life knowledge organisation hypermedia system operating as a desk-based device [8]. Memex was described as an "*enlarged intimate supplement to one's memory*". In these words, Bush had identified some of the key issues for maintaining lifelogs, that they be enlarged (store as much information as feasible over an extended period of time), intimate (private to the owner) and supplemental (working in synergy with one's memory). These remain some of the key guiding aspects of our research today. Initial research into lifelogging focussed on data gathering devices (e.g. Steve

Mann [29], Aizawa [39], SenseCam [26], Southampton deja-view [13]). While early lifeloggers such as Steve Mann [29] have dedicated decades of research into developing wearable life *capturing* technologies, the most famous attempt to address the challenges of retrieval and value extraction from lifelogs is the seminal MyLifeBits project [22] at Microsoft Research.

The MyLifeBits personal life experience archiving tool was concerned with gathering and making searchable, a long-term personal life archive for one individual. Recent work in the area has produced systems exploiting conventional smartphones as capture devices with real-time analysis of life experience [24], or using custom hardware devices in conjunction with smartphones to enable real-time capture and feedback [13]. It is our conjecture that real-time feedback will support a significant number of additional use-cases, beyond the current generation devices, such as the SenseCam.

The SenseCam (developed by Microsoft Research), is a wearable camera worn via a lanyard around the neck and captures data continually, regulated by onboard sensors. We will not describe the SenseCam in detail here, except to say that images captured using the SenseCam have been shown to operate as powerful autobiographical retrieval cues [6] and that the SenseCam has since been used to support personal reflection, both in those who are cognitively impaired [32], and those who have normal cognitive functioning [18]. In many cases, these efforts have relied on segmenting groups of images into distinct events [17] and then reviewing those images to elicit autobiographical retrieval cues which are triggered by the visual images [6]. Figure 1 shows a first generation interface to a large SenseCam image archive that implements event segmentation. This has served as an inter-disciplinary visual lifelogging platform to support access to personal digital memories [12, 40].

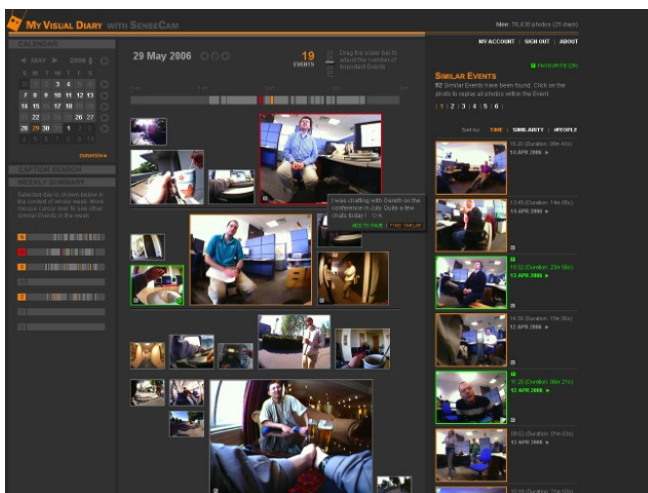


Figure 1. An event-based browsing tool for Visual SenseCam archives (from [17])

Although the SenseCam gathers thousands of photos and tens of thousands of sensor readings per day, it is limited to a fixed set of on-board sensors and off-line processing after data upload. A cell-phone (or other wearable technologies) could generate a richer lifelog which can

support real-time access [24]. Once the data is gathered, a core requirement is to organise the content so that it can be located when needed and consequently, the whole lifelog can become more useful for the lifelogger [4]. We discuss our consideration of the challenges of developing lifelogging solutions next.

## CHALLENGES OF DEVELOPING LIFELOG SOLUTIONS

This far, much research into using lifelogging tools such as the SenseCam has focused on either manually examining captured data or employing an event-based browser tool such as [17]. While there are many domain-specific solutions for focused use cases [16], such as human physical exercise or location sharing, there are still few attempts to fully grasp the potential of lifelogging. As lifelog technology progresses towards a more enhanced person sensing capability (often called total capture), this will pose new challenges for the areas of multimedia contextual sensor capture, multimedia data organisation, multimedia search and retrieval as well as the human factors that define how we can interact with such lifelogs, not to mention various issues such as privacy, security of data and supporting the human need to forget.

Developing lifelogging solutions requires a multidisciplinary approach, from capture of life experience, to presentation and summarisation for multimodal access. For *capture of life experience* the sensors that are utilised need to sample life activities in as much detail as possible, yet be unobtrusive enough to use every day and operate all day without requiring re-charging. Such sensors can be wearable, environmental, or informational (sensing our knowledge activities) and developing useful solutions would require the input of hardware/software engineers as well as personal and environmental sensing experts. Prior work [37] has suggested that focused capture of only the information needed is the best approach for supporting human memory; however we suggest that the idea of ‘total capture’, sampling life experience in high-fidelity will provide for a more useful, future-proof and flexible lifelog. Many of the computational techniques that will mine knowledge and patterns from our lifelogs may not be known yet, so it is more useful to capture as much data as possible now. With total capture, the ‘*wouldn’t it be great if I had a camera*’ moment will become a relic of the past.

We propose that making sense of such high-fidelity streams of raw (lifelog) data requires an understanding of why people will lifelog, and based on this understanding, we propose that it is important to segment data into meaningful life experiences, such as the event segmentation performed in [17], and subsequently to semantically enrich these segmented life experiences to make them easier to find and use later. This requires extensive use of machine learning and multimedia content analysis [15], which, would require the input of cognitive psychologists and healthcare professionals.

Once the lifelog archives are gathered, it becomes necessary to support *permanent archiving*, *real-time search*, *recommendation*, *analytics* and *summarisation* which brings together expertise from search, recommendation, scalable data processing/analytics and cognitive psychology to develop the underlying scalable indexing and

search tools that are effective in allowing a user to exploit their digital archives. Conventional information retrieval models will need to be augmented with cognitive information retrieval models that combine knowledge and best practice techniques from areas such as cognitive science, human-computer interaction, information retrieval, and memory science. It has already been suggested that this requires synergy with, not substitution of, human memory [37].

Finally, the fourth challenge is to support *effective multimodal interaction* with the lifelog, integrating HCI experts and cognitive psychologists [18]. To support effective user interaction, the lifelog needs to operate in real-time, be omnipresent (but not intrusive) and should take into account the cognitive abilities of a person to formulate queries over a lifetime of data, as well as considering the inherent human need to support forgetting.

### ADDRESSING THE CHALLENGES

In order to make lifelogs more useful for a large variety of use-cases, we need to both address the key challenges that were just described, as well as understand why and how a user would access a lifelog. If we consider how a user would access a lifelog, we can refer to the five R's of memory access by Sellen & Whittaker [37]. The five R's are recollecting, reminiscing, retrieving, reflecting and remembering intentions. Each of the five Rs define a different reason why people want to access their memories, and by inference, their lifelogs.

- *Recollecting* is concerned with reliving past experiences for various reasons. For example, we may want to recall who was at an event, what we spoke about.
- *Reminiscing*, which is a form of recollecting, is about reliving past experiences for emotional or sentimental reasons. It is often concerned with story-telling or sharing of life experiences with others.
- *Retrieving (information)*, is a more specific form of recollecting in which we seek to retrieve specific information from the lifelog, such as an address, a document or a piece of information.
- *Reflecting*, is a form of quantified self-analysis over the life archive data to discover knowledge and insights that may not be immediately obvious.
- *Remembering Intentions*, which is more about prospective (remembering plans) memory than episodic memory (past experiences). This is a form of planning future activities which is a life activity that everyone engages in.

These provide valuable clues as to how to develop the organisation, search and presentation elements of lifelogs, which will be the focus of the remainder of this paper. We will begin by exploring capture and storage, before looking at organisation, search and representation.

### Lifelog Capture & Storage

The starting point in generating lifelogs is gathering the data in a non-intrusive manner. The SenseCam, provided some major advancements; all-day operation on a single battery charge and the use of the on-board sensors as data streams, not simply capture triggers. Gurrin et

al. have performed an analysis on the data gathered by a SenseCam over multi-year archives [25]. However, if we consider the five R's of memory access as a basic requirement, a device such as the SenseCam has a number of very important drawbacks in its acceptance as a lifelogging tool. Principal among these is the off-line nature of the upload of content; a practical implementation of remembering intentions becomes impossible because it assumes real-time contextual analysis, likewise any other real-time recommendation or reminders are also impossible. It is necessary to bulk upload the content at the end of a usage session (typically a day), at which time the SenseCam device is non-operational. This limits the usefulness of the device to be purely a retrospective archive generation tool, with no potential for real-time access. A second shortcoming is that the device is not extensible; it is difficult to include new sensors on the device. The authors note that new technologies such as the Narrative Clip and OMG Life's Autographer will be on the market by the end of 2013, which provide SenseCam like functionality.

The ideal device would include sensors that can capture a rich life-experience archive, not impose an additional user burden by bringing/wearing additional devices, operate all day without requiring additional power sources, have onboard storage and support real-time communication. A smartphone fulfils these requirements [24]. Qiu et al. have developed smartphone software that can sample data constantly from a wide range of on-board physical sensors [34]. Such a range of sensors can easily be extended using on-board communications such as bluetooth, as is the case with Deja-view [13]. Automatic analysis of the user context (directly on the wearable device) allows for power-efficient sampling that can (today), realise a full day's lifelogging from a conventional smartphone and enable real-time recommendation/interventions based on user context. Smart upload of content means that a subset of an event (only some video/photos and sensor data) is streamed to the server, with the remainder uploaded when charging. This helps to conserve battery life; high-bandwidth upload of a lifelog drains the battery.

Once sensed data has been captured, it needs to be stored for access at any point in time. In order to support any of the five R's (especially reflection), it is our conjecture that the lifelog should not be time-limited, i.e. should extend back indefinitely and life experiences (unless by user request) should not be deleted for reasons of storage capacity or processing overhead. A typical SenseCam wearer will generate about 100GB of data per year; for newer devices, this is likely to be significantly higher, up to 1TB per year for newer wearable cameras. For an individual, it is reasonable to assume storage on a single computer and the assumption of Kryder's law, that hard drive densities will continue to increase, should see the storage capacity keep pace with data storage requirements. Notwithstanding the relatively low cost of digital storage, once data services scale to thousands or millions of people, then local storage solutions would tend to be replaced by cloud-hosting. However in the case of the data storage requirements of lifelogs, (at current pricing models) the data storage cost for cloud-hosting is prohibitive. This we see as a com-

mercial, not a technical challenge. We have a belief that the data storage technologies will continue to increase at a near exponential rate into the foreseeable future and that such challenges will be met.

Used alone, such raw readings (images, other sensor data) do not provide much semantic value, but a phase of semantic enrichment (see next section) can enhance the usefulness of the archive. The inclusion of additional sensors via on-board communications, or from external or WWW sources, can also significantly increase the potential for semantic analysis and annotation, though this has not yet received significant research attention.

### Lifelog Organisation

The human memory system has evolved over thousands of years to store autobiographical memories, and we believe that lifelogs need to consider how the human mind operates (synergy not substitution) and support the user accordingly. Literature has motivated that the human mind stores information in distinct events, that similar events are associated with each other, and that more important events are more strongly remembered [18] and easier to recall. Therefore, a starting point for lifelog organisation is proposed as follows:

- *Raw data should be arranged into events:* Typically, in a full day, we know that a person encounters anything upwards of 20 individual *events*, with each lasting about 30 minutes, though there is a lot of variety [19, 27]. Prior work on event segmentation analyses sensor streams from wearable cameras to segment of life-experience into events, post-capture [17], however this poses a problem. The human memory operates in real-time, so we propose that segmentation and processing should occur in real-time and lifelog knowledge made available to the user as it happens. Prior work has defined a generic event segmentation model, in which the event is the atomic unit. Such a situation is not flexible to present query-specific information in response to a user information need. To fully support the five Rs of memory access, we need to consider a sub-event model of retrieval that can retrieve specific nuggets of information as required or provide query specific data retrieval.
- *Events should be semantically described:* To support both post-hoc review and real-time lifelog interactions, a suite of semantic analysis tools are needed. These act as software sensors to enrich the raw sensor streams with semantically meaningful annotations. Such software sensors are multi-layered to allow for additional derivations to be mined from existing sensor outputs. For example, raw accelerometer values on a smartphone can identify the physical activities of a user [3], bluetooth and GPS sensors allow us to determine where and with whom people are with [10], while using automatic detection of concepts is possible from images [15] to 'understand' the image content. This is just a small example, and event ontologies could help to infer higher-level semantics on the lifestyle of individuals. In addition the relative importance or potential memorability of each event can be determined via semantic analysis; a combination of image face detection, bluetooth people recognition and a measure

of the relative uniqueness of the user context (relative to their normal lifestyle) [20] has been employed to automatically determine event importance.

- *A rich narrative summary of each event should be generated:* Concise narratives are an important building block for lifelogs in that they are shown to produce emotional responses to autobiographical memories and help support many of the 5 Rs of memory access. Xu explains why such simple summaries are desired on a cognitive functioning level [42]. A benefit of narrative generation is that a textual narrative can also be used to support keyword text search (Googling for knowledge or events).

### Lifelog Search and Retrieval

In order to retrieve life experiences for search or recommendation, either post-capture or in real-time, some form of search engine would be needed. An initial assumption would be to employ state-of-the-art techniques from database search and information retrieval to scalably index the life-experience events and provide omnipresent access via keyword/database search, ranking, recommending and presenting the multimedia rich life experience archive through multimodal interfaces. However, we contend that to better understand how to develop lifelogging solutions and how to support effective access to these lifelogs, it becomes necessary to understand how people will use and access their life archives. As a starting point, we turn again to the five R's of memory access from Sellen & Whittaker. Each of the five Rs provide valuable clues as to why people want to access their memories, and by inference, their lifelogs.

- *Recollecting, reliving past experiences, in the case of lifelogging is concerned with accessing episodic memories.* Recollecting will require highly accurate search engines that semantically rank events and represent these events in a format required to aid recollection. This will require conventional information retrieval, coupled with query-specific experience segmentation (somewhat similar to the Shot Boundary Detection [38] of digital video. This in itself is a motivation for 'total capture'.
- *Reminiscing, is a form of recollecting and is about reliving past experiences for emotional or sentimental reasons, sometimes alone, often with others.* From information retrieval, it will require new techniques for narrative generation [11], storytelling [9], topic detection and tracking [1] and novelty detection [43] from single (and potentially multiple) individual's archives.
- *Retrieving (information) requires the retrieval of specific nuggets of information from the lifelog.* Retrieval will require highly accurate text, multimedia and sensor-data search engines that retrieve and extract just the nugget of information that is most pertinent to the user. The conventional information retrieval concept of top N ranked lists does not transfer to lifelogs (unless  $N = 1$ ); after all, there is marginal benefit in a system that provides a ranked list of locations for where the car has been left. The query will define the type of knowledge that is required.
- *Reflecting, is about discovering knowledge and insights from the lifelog.* It includes data analytics,

information summarisation from lifelog streams [31], event detection [21], various forms of data analysis to infer and evaluate the importance of new semantic knowledge [15, 39, 10, 33] from the lifelog and optimised presentation methodologies. Typically such data analysis approaches rely on artificial intelligence, machine learning and various forms of statistical analysis and should proactively recommend new knowledge, not solely relying on a human information need as input.

- Remembering intentions, is a form of planning future activities which everyone engages in. This assists people to remind or prompt them on tasks they would like to do (*e.g. post that letter*), or real-time prompts on who they are talking to (*e.g. this is Paul*), or giving prompts on conversation cues (*e.g. last time here together, you had just come away from seeing the new Batman movie*). Past lifelogging efforts were exclusively focused on episodic memory as it was always a post-hoc analysis (i.e. constrained by technology); however with real-time technology, one can consider situational awareness (and past history of user) to provide prospective memory prompts.

Taking the five Rs as a guide, we need to identify how to develop search solutions that can effectively support autobiographical memory access across a range of tasks. This should support both explicit user queries as well as real-time contextual cues. An underlying caveat in all of this work is that a human's ability to generate a query or to find a suitable browsing point is highly dependent on their own ability to remember details correctly from the event containing the information they are seeking.

### Lifelog Interaction

As with any information retrieval system, we need to understand how the lifelog will be accessed. However with lifelogs, there is little prior work into interaction design, though there are some initial guidelines [1]. To begin our analysis, consider how people access their own digital photo or video archives. For sufficiently small archives, a browsing mechanism is usually sufficient (e.g. browsing by date or location). This could be enhanced by manual or automatic clustering/grouping of related content. However, when the archives become larger and less organised, a search or search/browse metaphor is normally chosen to support fast and effective access.

Considering lifelogs, the sheer scale of these multi-year or multi-decade archives suggests that a browsing methodology is not sufficient from the outset. Initial (and the only) experiments into multi-year lifelogs suggests that even a basic search methodology increases the possibility of a user locating desired content by a factor of three in a tenth of the time [18]. There are a number of alternative search methodologies that could be considered. Firstly keyword based search can be processed over textual narratives generated from lifelog data. Another alternative approach is to support the user in generating a new type of multi-axes query in an efficient manner; for example, *I know that my friends Paul and Jack were there, it was a Sunday evening, and we were in Barcelona watching a football game in a bar*. A third option is the real-time context-driven automatic querying that is somewhat of a

holy grail of this research area. The realtime sampling of life experience can trigger contextual queries to support recollection, retrieval of information and remembering intentions, which, if presented to the user in a suitable manner, can provide for truly novel and currently undiscovered applications for lifelogs. Applications that can remind you that the person you have just met is having a birthday today or that the last time you bought this type of soup, you felt ill the following day. We are likely to see such applications being developed for real-time sensing technologies like Google Glass.

In order to more clearly define the interaction challenges and help to articulate our vision, let us again refer to the five R's of memory access. Since each R defines a different way that people are able to interact with their memories, and until such time as we have sufficient numbers of people maintaining lifelogs to get real-world usage data, they serve as a source of different proposed interaction scenarios available:

- Recollecting is concerned with reliving past experiences and we know that visual media, especially captured from the first person viewpoint provide very powerful memory cues and leads to what is referred to as Proustian moments of recall, where the recollected experience is recalled by the user in vivid detail. Any interaction mechanism that supports recollecting from lifelogs will need to be able to support user query formulation (in whichever incomplete form it is provided) and visually present the result events to the user, with a focus on visual media, and a secondary use for other sampled experience data.
- Reminiscing, which is concerned about reliving past experiences for emotional or sentimental reasons will rely heavily on storytelling and narrative generation. Experience suggests that although reminiscing may be an individually activity, it is more frequently a social activity. This poses interesting challenges for interaction design with consideration needed for collaborative search interfaces and a new type of result content that merges various sensed data from lifelogs to compose a story from the archive of one (or more) people. In addition, a social and relaxed environment would be expected for reminiscing, therefore the query generation process needs to be enjoyment oriented with fast response time, minimum overhead querying, most probably employing a gamification methodology to hide the query mechanism altogether.
- Retrieving (information) is concerned with locating information from the lifelog, such as a document, a location, a sound, a recipe, and so on. In this case, the focus of the interaction methodology should be on supporting the user in the query generation process. As described above, this is heavily dependent on the ability of the user to recall query cues to generate an effective query. The actual result presentation is dependent on the information need, and as such could be a video clip, a photo sequence, a face, a location and so on.
- Reflecting will analyse patterns and discover knowledge/insights that may not be obvious to the user. The user should be able to define source data for

analysis, for example, the activity levels correlated with location and time to identify where the user is most active on weekend mornings. The presentation of this data should be highly visual and employ interface metaphors such as timelines, charts and infographics, each of which could support click-through analysis of the underlying data to support drill-down reflection.

- Remembering intentions is concerned with reminders and recall of future activities. The key driver to support this is real-time monitoring and query triggering. Context sensitive reminders could be automatically triggered based on the current context of the user, for example in a certain location at a certain time a reminder is triggered. The key interaction challenge is how to present the reminder to the user.

The five Rs provide a guiding insight into the interaction requirements for lifelogs. As can be seen, elements of assisted query formulation, engaging storytelling, summarisation, visualisation and potentially disruptive recommendation are the key research points that will need to be explored in greater detail.

### REALISING THE POTENTIAL

Heretofore lifelogging has appeared to be an extreme activity carried out only by a small number of pioneering enthusiasts e.g. Steve Mann [29], Gordon Bell [4], and Cathal Gurrin [15]. There are three main reasons why not everyone could have these lifelogs automatically gathered: 1) privacy and ethical concerns [30]; 2) overwhelming amounts of data [37]; and 3) limited device availability. In 2013, we see new devices on the market that will solve the issue of device availability, including smartphones with customisable apps. The search and storage technologies will solve the challenge of overwhelming data quantities. Finally on 1) the privacy and ethical concerns, it is our conjecture that once the personalised experience, wellness, and memory capture and sharing provide a wide range of benefits to end users, that an acceptable usage policy will emerge.

We firmly believe that lifelogging will very soon be a phenomenon available to everyone and maintaining lifelogs will positively impact on individuals, namely:

- **Easy sharing of natural life experiences:** Photo sharing on social networking sites such as Facebook has become popular. Until now the user has had to make a conscious decision to capture any given photo, thus fracturing the experience. Passive capture of lifelogging allows users to enjoy their experiences without interruption and share the media rich experience after the fact. Indeed with Google Glass, one already can see the ideal of real-time life experience streaming.
- **Prospective memory feedback to enhance productivity:** Memory research exploiting visual lifelogging has focused on retrospective episodic memory tasks until now [12]. However as real-time upload becomes commonplace, lifelogs will support the human prospective memory system [2]. Given a user's prior set of experiences and preferences (historical lifelog data), and their current situation (lifelog context), prospective memory prompts can be provided e.g. to remind you that wanted to pick up some more milk

from this grocery store; or providing real-time prompts about whom you are talking to.

- **Personalised wellness feedback:** Poor diet and lack of physical activity are strongly associated with the early death of millions of people [41]. Passive capture lifelogging devices offer the potential to automatically log physical activity [28] and dietary input [35]. This can help to promote healthy lifestyle choices.
- **A greater knowledge of self:** Coupled with the concept of personal wellness, the era of lifelogs will provide information to the individual about their own life activities; information that otherwise would go unnoticed. One can identify trends and pattern in lifestyle and wellness over an extended period of time.
- **A life archive that never forgets:** Neurodegenerative disease affects a large proportion of an ageing population and maintaining a lifelog has been shown in initial small-scale studies to help offset some of the debilitating effects of memory impairment [26]. Even for individuals with fully functional memories, the ability to refer back to the lifelog could allow for disambiguation of faded memories and more accurate recall of the past.
- **An opportunity to better understand the functioning of human memory:** Self-reporting is notoriously error prone [28], visual lifelogs offer an potential to verify contextual details (who/what/when/where) of episodic memories recalled by individuals. When visual lifelogging is available via a widely used medium, experimental psychologists can carry out tests at a large scale.
- **An opportunity to better understand the associations between lifestyle, environmental context and mortality:** Understanding the determinants and barriers to physical activity behaviours is important in designing interventions to positively change behaviours [36]. Accurate measurement of physical activity events is therefore important [7]. Examples of important context attributes of an event of physical activity include: whether it occurs indoors or outdoors; the time of day it occurs; if it is alone or in companionship; and its domain (home, occupational, etc.). Currently, some of these attributes are subjectively measured via self-reporting, but for interventions to be successful, accurate measurement of existing behaviour on what people are doing and when, as well as under what conditions, is critical.
- **A cross-population archive:** Although we have spoken about the personal nature of lifelogging, there is likely to be enormous potential for mankind when the life archives of whole populations can be analysed for trends and statistics. To take just one example, in the spirit of the Framingham study [23], analysis by reputable and trusted organisations of anonymised archives of populations could potentially lead to greater understanding of diseases and illnesses, on a scale heretofore unimaginable.

The potential for lifelogs is enormous. We do acknowledge that there are challenges to be overcome, such as privacy concerns, data storage, security of data, and the

development of a new generation of search and organisation tools, but we believe that these will be overcome and that we are on the cusp of a positive turning point for society; the era of the quantified individual who knows more about the self than ever before, has more knowledge to improve the quality of their own life and can share life events and experiences in rich detail with friends and contacts.

## CONCLUSIONS

In this paper we have presented the first steps into a world of lifelogging, in which we will gather vast archives of data about the individual. This topic is set to affect everybody over the coming years as we learn more about ourselves and have access to technologies that will help us in many aspects of everyday life. We motivate this research area by presenting a number of potential benefits. However, there are numerous research challenges to be met, but experiences from cognitive science, information retrieval, human computer interaction and memory science provide valuable insights into how we can begin to address these challenges.

## ACKNOWLEDGEMENTS

This work is supported by Science Foundation Ireland under grants 07/CE/I1147 and 07/SK/I1186.

## REFERENCES

1. James Allan. Topic detection and tracking, chapter Introduction to topic detection and tracking, pages 1–16. Kluwer Academic Publishers, Norwell, MA, USA, 2002.
2. Alan Baddeley, editor. *Your Memory: A User's Guide*. Carlton Books, 2004.
3. Ling Bao and Stephen Intille. Activity recognition from user-annotated acceleration data. In Alois Ferscha and Friedemann Mattern, editors, *Pervasive Computing*, volume 3001 of *Lecture Notes in Computer Science*, pages 1–17. Springer Berlin / Heidelberg, 2004. 10.1007/978-3-540-24646-6-1.
4. Gordon Bell and Jim Gemmell. A digital life. *Scientific American*, 2007.
5. Gordon Bell and Jim Gemmell, editors. *Total Recall: How the E-Memory Revolution Will Change Everything*. Penguin Books, 2009.
6. Emma Berry, Adam Hampshire, James Rowe, Steve Hodges, Narinder Kapur, Peter Watson, Georgina Browne Gavin Smyth, Ken Wood, and Adrian 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.
7. Nicola W Burton, Asaduzzaman Khan, and Wendy J. Brown. How, where and with whom? Physical activity context preferences of three adult groups at risk of inactivity. *Br J Sports Med*, Jan 2012.
8. Vannevar Bush. As we may think. *The Atlantic Monthly*, 176(1):101–108, Jul 1945.
9. Daragh Byrne, Aisling Kelliher, and Gareth J. F. Jones. Life editing: third-party perspectives on lifelog content. In Desney S. Tan, Saleema Amershi, Bo Begole, Wendy A. Kellogg, and Manas Tungare, editors, *CHI*, pages 1501–1510. ACM, 2011.
10. Daragh Byrne, Barry Lavelle, Aiden R. Doherty, Gareth J.F. Jones, and Alan F. Smeaton. Using bluetooth and gps metadata to measure event similarity in sensecam images. In *IMAI'07 - 5th International Conference on Intelligent Multimedia and Ambient Intelligence*, pages 1454–1460, 2007.
11. Charles B. Callaway and James C. Lester. Narrative prose generation. *Artif. Intell.*, 139(2):213–252, August 2002.
12. Martin A. Conway and Catherine Loveday. SenseCam: The Future of Everyday Memory Research. *Memory*, 19(7):685–807, Oct 2011.
13. Dirk De Jager, Alex L Wood, Geoff V Merrett, Bashir M Al-Hashimi, Kieron O'Hara, Nigel R Shadbolt, and Wendy Hall. A low-power, distributed, pervasive healthcare system for supporting memory. may 2011.
14. 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):431–445, 2007.
15. Aiden R. Doherty, Niamh Caprani, Ciaran O Conaire, Vaiva Kalnikaite, Cathal Gurrin, Alan F. Smeaton, and Noel E. O Connor. Passively recognising human activities through lifelogging. *Computers in Human Behavior*, 27(5):1948–1958, 2011.
16. Aiden R. Doherty, Steve E Hodges, Abby C. King, Alan F. Smeaton, Emma Berry, Chris J.A. Moulin, Siân Lindley, Paul Kelly, and Charlie Foster. Wearable cameras in health: The state of the art and future possibilities. *Am J Prev Med*, 44(3):320–323, Mar 2013.
17. Aiden R. Doherty, Chris J.A. Moulin, and Alan F. Smeaton. Automatically assisting human memory: A sensecam browser. *Memory*, 7(19):785–795, 2011.
18. Aiden R. Doherty, Katalin Pauly-Takacs, Niamh Caprani, Cathal Gurrin, Chris J.A. Moulin, Noel E. O Connor, and Alan F. Smeaton. Experiences of aiding autobiographical memory using the sensecam. *Human-Computer Interaction*, 27(1-2):151–174, 2012.
19. Aiden R. Doherty and Alan F. Smeaton. Automatically segmenting lifelog data into events. In *WIAMIS 2008. The 9th International Workshop on Image Analysis for Multimedia Interactive Services*, pages 20–23. IEEE, 2008.
20. Aiden R. Doherty and Alan 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.

21. Aiden R. Doherty and Alan F. Smeaton. Automatically augmenting lifelog events using pervasively generated content from millions of people. *Sensors (Peterborough)*, 10(3):1423–1446, 2010.
22. Jim Gemmell, Gordon Bell, and Roger Lueder. MyLifeBits: a personal database for everything. *Communications of the ACM*, 49(1):88–95, 2006.
23. Tavia Gordon, William P. Castelli, Marthana C. Hjortland, William B. Kannel, and Thomas R. Dawber. Predicting coronary heart disease in middle-aged and older persons. *JAMA: The Journal of the American Medical Association*, 238(6):497–499, 1977.
24. Cathal Gurrin, Zhengwei Qiu, Mark Hughes, Niamh Caprani, Aiden R. Doherty, Steve E. Hodges, and Alan F. Smeaton. The smartphone as a platform for wearable cameras in health research. *American Journal of Preventive Medicine*, 44(3):308–313, March 2013.
25. Cathal Gurrin, Alan F. Smeaton, Daragh Byrne, Neil O Hare, Gareth J.F. Jones, and Noel E. O Connor. An examination of a large visual lifelog. In *AIRS 2008 - Asia Information Retrieval Symposium*, 2008.
26. Steve Hodges, Emma Berry, and Ken Wood. Sensecam: A wearable camera that stimulates and rehabilitates autobiographical memory. *Memory*, 7(19):685–696, 2011.
27. Daniel Kahneman, Alan B. Krueger, David A. Schkade, Norbert Schwarz, and Arthur A. Stone. A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306:1776–1780, 2004.
28. Paul Kelly, Aiden R. Doherty, Emma Berry, Steve Hodges, Alan M. Batterham, and charlie Foster. Can we use digital life-log images to investigate active and sedentary travel behaviour? results from a pilot study. *International Journal Behavioural Nutrition Physical Activity*, 8:44, 2011.
29. Steve Mann. Wearable computing: A first step toward personal imaging. *Computer*, 30(2):25–32, 1997.
30. David H. Nguyen, Gabriela Marcu, Gillian R. Hayes, Khai N. Truong, James Scott, Marc Langheinrich, and Christof Roduner. Encountering sensecam: personal recording technologies in everyday life. In *UbiComp '09: Proceedings of the 11th international conference on Ubiquitous computing*, pages 165–174, New York, NY, USA, 2009. ACM.
31. Y. Ogras and Hakan Ferhatosmanoglu. Online summarization of dynamic time series data. *The VLDB Journal*, 15(1):84–98, January 2006.
32. Katalin Pauly-Takacs, Chris J. Moulin, and Edward J. Estlin. SenseCam as a rehabilitation tool in a child with anterograde amnesia. *Memory*, 19(7):705–712, Oct 2011.
33. Zhengwei Qiu, A.R. Doherty, C. Gurrin, and A.F. Smeaton. Mining user activity as a context source for search and retrieval. In *Semantic Technology and Information Retrieval (STAIR), 2011 International Conference on*, pages 162–166, June 2011.
34. Zhengwei Qiu, Cathal Gurrin, Aiden R. Doherty, and Alan F. Smeaton. A Real-Time life experience logging tool. *Advances in Multimedia Modeling*, pages 636–638, 2012.
35. Sasank Reddy, Andrew Parker, Josh Hyman, Jeff Burke, Deborah Estrin, and Mark Hansen. Image browsing, processing, and clustering for participatory sensing: lessons from a dietsense prototype. In *EmNets '07: Proceedings of the 4th workshop on Embedded networked sensors*, pages 13–17, New York, NY, USA, 2007. ACM.
36. James F Sallis, Neville Owen, and Michael J. Fotheringham. Behavioral epidemiology: a systematic framework to classify phases of research on health promotion and disease prevention. *Ann Behav Med*, 22:294–298, 2000.
37. Abigail J. Sellen and Steve Whittaker. Beyond total capture: a constructive critique of lifelogging. *Comm. ACM*, 53(5):70–77, 2010.
38. Alan F. Smeaton, Paul Over, and Aiden R. Doherty. Video shot boundary detection: Seven years of trevid activity. *Computer Vision and Image Understanding*, 114(4):411–418, 2010. [jce:title;Special issue on Image and Video Retrieval Evaluation;ce:title;](#)
39. Datchakorn Tancharoen, Toshihiko Yamasaki, and Kiyoharu Aizawa. Practical life log video indexing based on content and context. In *Multimedia Content Analysis, Management, and Retrieval, In Proceedings of SPIE-IST Electronic Imaging*, January 2006.
40. Elise van den Hoven, Corina Sas, and Steve Whittaker. Designing for personal memories: Past, present, and future. *Human-Computer Interaction*, 27(1-2):1–12, 2012.
41. WHO. *A guide for population based approaches to increasing levels of physical activity: Implementation of the WHO global strategy on diet, physical activity and health*. World Health Organisation, 2007.
42. Jiang Xua, Stefan Kemenya, Grace Parka, Carol Frattalib, and Allen Brauna. Language in context: emergent features of word, sentence, and narrative comprehension. *Neuroimage*, 25(3):1002–1015, 2005.
43. Le Zhao, Min Zhang, and Shaoping Ma. The nature of novelty detection. *Inf. Retr.*, 9(5):521–541, November 2006.