

Multiple multimodal mobile devices: Lessons learned from engineering lifelog solutions

Daragh Byrne^{1,2}, Liadh Kelly², Gareth J.F. Jones²

CLARITY: Centre for Sensor Web Technologies¹ &

Centre for Digital Video Processing, Dublin City University, Ireland²

ABSTRACT

For lifelogging, or the recording of one's life history through digital means, to be successful, a range of separate multimodal mobile devices must be employed. These include smartphones such as the N95, the Microsoft SenseCam – a wearable passive photo capture device, or wearable biometric devices. Each collects a facet of the bigger picture, through, for example, personal digital photos, mobile messages and documents access history, but unfortunately, they operate independently and unaware of each other. This creates significant challenges for the practical application of these devices, the use and integration of their data and their operation by a user. In this chapter we discuss the software engineering challenges and their implications for individuals working on integration of data from multiple ubiquitous mobile devices drawing on our experiences working with such technology over the past several years for the development of integrated personal lifelogs. The chapter serves as an engineering guide to those considering working in the domain of lifelogging and more generally to those working with multiple multimodal devices and integration of their data.

INTRODUCTION

This chapter discusses the role and use of multiple mobile devices in life capture or 'lifelogging'. Lifelogging technologies afford us the potential to record a digital account of our personal life histories. A lifelog collection seeks to collect as much digital data on the activities and life of an individual as possible. Through a range of mobile technologies, not only can the digital content encountered in our day-to-day activities be preserved, but also an individual's current contextual factors determined, for example, through environmental- or personal- sensing. The digital artifacts of significance to us are thereby automatically and passively assembled into a multimodal collection. Such a collection might for instance include emails sent and received, text messages, web pages or documents reviewed or created, photos and videos, along with contextual factors such as places visited or people encountered. By bringing these mobile devices and software solutions, and the data they amass, into confluences, we can gain huge insight into the user and empower life capture and the potential for subsequent retrieval, sharing and reminiscence. Lifelog capture, however, poses software engineering and design challenges, not only in the actual collection and recording of life data, but also in the requirements for data management, data processing, integration and consolidation.

For over two years we have been actively working with large scale multimodal lifelogs created from a range of content and context sources by using a diverse range of applications, platforms and devices. From our experience of working with these capture technologies and in developing software solutions to collect, manage and access these collections, we are painfully aware of the challenges raised when attempting to acquire, assemble, aggregate, and use the information from the range of sources required to deliver interesting content, and relevant understanding about the collection owner. From our practical experiences, we have first-hand knowledge of the difficulties in enabling, using, managing and collecting from the multiple sources used to compile a lifelog. From working with these mobile devices we have

also gained invaluable insights into bringing them into confluence. The challenges posed by our activities have implications on engineering the design and creation of software solutions to enable lifelogging and lifelogging applications, and on applications dependent on multiple multimodal devices in general. This chapter is intended to serve as a practical tool for those seeking to gather and use information compiled from diverse mobile devices.

BACKGROUND

In his seminal work Vannevar Bush (Bush, 1945) conceived the notion of lifelogging as a device on which all a person's personal information could be stored and from which it could then later be retrieved. Towards realizing this vision, Microsoft's Gordon Bell (Bell, 2001) has invested both effort and time in the archival and digital capture of all of his personal data. His efforts and the initial focus for lifelogging technology emphasized desktop retrieval, e.g. (Dumais et al., 2003), however in more recent years equal importance has been placed on mobile access and capture (Mase et al., 2006). These technologies, and indeed the wealth of personal information they capture through mobile devices has been explored and exploited not just for personal use but across a range of domains. These have been outlined by Byrne et al (2008b), and include for example therapeutic and medical solutions (Berry et al., 2007; Hodges et al., 2006; Al Mahmud et al., 2008), the obvious use in reminiscence (McCarthy et al., 2007) and more diverse and playful applications (Wood et al 2004).

Recently much attention has been given to lifelogging and research has focused on addressing many of the challenges presented in the management of (Doherty & Smeaton, 2008) and retrieval from (Kelly et al., 2007) such voluminous multimodal collections. The focus of our work within this domain has been to first build up long-term, large-scale multimodal lifelogs containing content generated from multiple mobile devices. Our archives now contain over one and a half years worth of continuously archived personal data, and are much larger and more heterogeneous than typically personal created digital archives. These have then been used to explore key issues pertaining to lifelogs, specifically in context-based retrieval (Kelly et al., 2008), interface development (Chen & Jones, 2009), and narrative generation for reminiscence (Byrne & Jones, 2008).

The lifelogs in use contain data from a variety of complementary sources, including:

Passive Capture Media: Continuous passive image capture is enabled through the use of a small wearable device, the Microsoft Research SenseCam (Hodges et al., 2006; see Fig. 1).

Explicitly Captured Media: Photos and explicitly captured media created on a mobile device can additionally be added to the lifelog to augment passive captured photos.

Mobile Context Data: This data is captured on subjects' Nokia N95 mobile phones by constantly running the *Campaignr* software, provided to us by UCLA (USA) (Joki, Burke & Estrin, 2007). This provides location cues as GPS data; wireless network presence and GSM location data, from which place names, light status and weather conditions can be derived, and co-present Bluetooth devices, from which people present can be uncovered (Lavelle et al. 2007; Kostakos, 2008).

Mobile Activity Data: Mobile phone activity in the form of call logs are recorded using a proprietary piece of software and SMSs are captured using an application developed in-house.

Biometric Information: Physiological and heart rate readings are taken from wearable biometric devicesⁱ, see Figure 2. These allow us to monitor physiological conditions and infer the wearer's emotional state.

Desktop Activity: All laptop and PC activity is also monitored (every item (email, Word document, web page, etc) accessed by the user, with time and duration of access, contents of item, path to item information, etc.) for subjects using a combination of MyLifeBits (Gemmell et al., 2002), Slifeⁱⁱ and in-house scripts.

The use of a variety of multimodal and mobile sources, sampled at high frequency enables a rich picture of a person's life to be assembled. These are then available for use in emerging applications which enable past experience to be assessed, retrieved or re-lived.



Figure 1. The Microsoft Sensecam, inset as worn by a user. (Adapted from Microsoft Research SenseCam homepage).



Figure 2. Biometric devices as worn by a user – the BodyMedia Armband is fitted to the arm.

CHALLENGES PRESENTED BY LIFELOGGING TECHNOLOGIES

There are many challenges associated with collecting and joining data from diverse mobile devices. From our practical experiences of building, developing and maintaining lifelog collections, along with qualitative and quantitative explorations of the collections which have been amassed, we have discerned a number of core challenges which must be born in mind when undertaking work within the lifelogging domain. These can be broadly demarcated into three categories which are elaborated upon below.

Practicalities of mobile device use

There are clear challenges resulting from the use of mobile devices as tools to enable continuous recording of a person's day-to-day activities. There are a number of practical and operational limitations and constraints inherent in the use of mobile devices in this manner and in confluence.

Battery Life

Since the goal of lifelogging is to offer continuous and non-invasive recording of a person's day-to-day activities, a significant stumbling block to successfully achieving this lies with the current power requirement for uninterrupted recording with mobile devices. Within lifelogging applications, careful consideration for the battery consumption is required and in particular: what modalities and cues are essential to capture and how expensive they are (the

capture of some cues is more expensive than others); in addition to striking a balance between frequency and richness of information captured. If an appropriate balance is not struck then this will have significant affect on the completeness of a lifelog – as the devices will power down unexpectedly before a full day can be observed. In the case of a device such as the SenseCam, if a battery fails it may be several hours before it is charged and ready for use again. In addition to this, power management places a burden on the user to regularly charge and monitor power consumption for their multiple devices. Not only will loss of power have an effect on the completeness of the collection, but it may additionally have impact on the accuracy and correctness of the data within the collection. Often when a device loses power completely, or should a battery be removed (for a protracted period), some of the information held on the device, such as the time, may be lost and manual effort to realign and correct this information may be required.

The SenseCam offers a battery life of between 14 and 18 hours (depending on the frequency of capture) and while this is generally sufficient to capture a typical day, in case of particularly busy or active days the wearer may engage in activities beyond the scope of the battery. Further to this, while most mobile phones and devices will have a standby battery life of several days, their use in actively sensing contextual or environmental cues can significantly and rapidly decrease the lifetime of the battery to hours. Mobile software can poll extremely frequently for contextual cues such as co-present Bluetooth devices or wireless signals. The frequency at which this software polls has obvious implications on power consumption. In our work we employ capture at a rate of every 30 seconds, this provides an extremely rich detailed picture of the environment the owner is located within, and allows us to discern subtle transitions within a narrow window. We found this allows less than 10 hours capture using a Nokia N95 with a standard 950mAh battery. As such, we employ a bulkier but greater capacity battery (2200mAh) which allows for slightly more than one full day's uninterrupted capture. This, however, is often overly detailed for most practical applications and so it may be acceptable to sacrifice data capture frequency for battery life and lose granularity and richness within the captured data. This sentiment is echoed in work of Eagle & Pentland (2006).

Stability, Reliability and Errors of the devices

The stability of the platforms employed within lifelogging is of paramount importance. Continuous frequent capture with mobile devices unfortunately can both degrade performance and reliability resulting in intermittent device or software crashes, hardware failure or forced rebooting. These crashes often happen unexpectedly and the onus is placed on the user to intervene to recover from the failure, e.g. restart the software, particularly on mobile devices. However, these crashes are often neither obvious nor visible to a user. Consider a mobile phone: it is predominantly in standby mode with a dimmed screen, further to this it is often out of sight either in a pocket or on a table or desk until it is needed – so even if visible cues were provided at the time of crash (e.g. screen lighting up or the device rebooting and presenting the welcome screen) they would most likely go unnoticed. A fault on devices with lower fidelity output such as a biometric device or the SenseCam can be even more difficult to discern and discover. This presents a difficulty, as the user may not detect a problem for several hours or longer. Indeed, while the software may appear to be running correctly, inherent limitations in the hardware or technology which underlies this software can affect data capture. From our experience, this is particularly evident in the capture of contextual information at high frequencies and is illustrated in Figure 3. For example in looking at day 4 in this figure, we can see that an unusually low number of polls have been captured. This is the result of a software crash which went unnoticed by the user leaving a large portion of the users activities during this day unrecorded. This has been similarly noted by Eagle and Pentland (2006). Software and hardware crashes may additionally result in the corruption of data stored onboard the device. While in most cases this is a minor annoyance (such as the loss of a single field of information), in exceptional circumstance the loss of data can be catastrophic. In one instance, for example, within our collection activities six-weeks of a participant's data stored locally in a mobile device's database became entirely corrupted and

irrevocably lost. Due to the nature of the software in question, the corruption was not obvious, could not be easily corrected and only presented when the user attempted to upload data upon return from travels.

It is also often the role of the collection owner to monitor the many devices engaged in recording the data to ensure they have not crashed and to discern if they have to restore them to working order. There is as such a burden on the owner first to recognize and diagnose a crash, second to take the appropriate action to recover from that error and finally to do so in a timely manner to ensure good coverage of the recordings. However, as these errors are often not visible or intuitive, one cannot expect a user to ‘check-in’ on all the recording apparatus extremely frequently as this would be highly intrusive into their day-to-day activities. Instead it must be accepted that even with the most diligent lifelogger, it may be several hours before an error or crash is discovered and corrected.

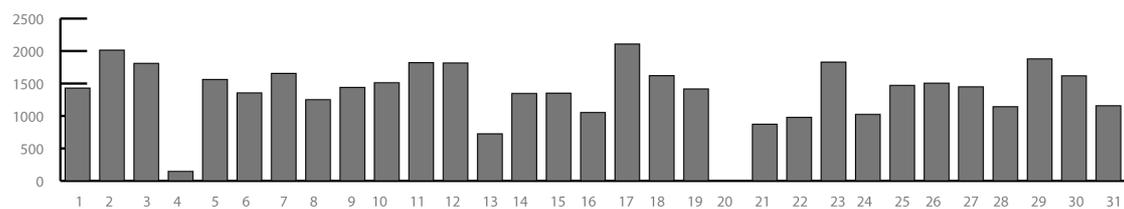


Figure 3. A graph displaying the number of context polls achieved for a representative user during a one month period. A major crash on day 19/20 is clearly visible as is inconsistent performance throughout.

Storage Constraints

It should be born in mind that while mobile devices have increasingly large storage, it is not unlimited. Lifelogging solutions are designed to operate by collecting large volumes of rich information on a person’s day-to-day activities *in the longer term*. In combination with high-frequency of capture, even large capacity devices quickly become filled. Consequently when developing lifelog solutions, careful consideration of and provision for the storage constraints of the devices is required.

Additionally, while many high-end mobile devices now offer storage in excess of several gigabytes, many of the mobile devices employed for lifelogging will not offer such capacities. For example, biometric devices often require extremely small form factors and as such are highly constrained in the amount of storage they can be equipped with. In particular the BodyMedia armband has only sufficient storage for one day of data when sampling at a high rate. Unless the data is removed from the device nightly it will not allow for operation the next day. The Microsoft SenseCam has enough onboard memory to allow continuous capture for between one and two weeks. However, should a user reach or attempt to go beyond the capacity of the device, the device itself will freeze and may begin to overwrite existing images corrupting them in the process. In severe cases, the sensor data for the entire period may become entirely corrupted. This situation is compounded by the fact that no visual feedback on the availability (or lack thereof) of free storage space is given. As such the onus is on the user to ensure data is retrieved in advance of this.

Data Extraction

The data produced from lifelogging technologies offers most utility when it is brought into confluence. In order to achieve this, the various data sources from the mobile devices must be brought together in a central store. However, typically this information is not automatically or easily centralized. Each of the devices stores the data it collects on-board, often in considerably different formats, without direct communication with, knowledge of or awareness of the target ‘store’. As such the collection owner must take charge of numerous, independent, and often manual, data extraction processes: removal of images from the

SenseCam; context data from mobile devices, collection of mobile activity information, etc. This is a wholly cumbersome and time-consuming process.

One might envisage, however, that mobile devices, which can take advantage of built-in wireless internet connections, would automatically upload their data as it is collected. Certainly many lifelogging technologies make provision for this, including the Campaignr Software (Joki, Burke & Estrin, 2007), however, from our experience it is not an ideal approach for several reasons. First, there is not always access to wireless signals in all scenarios where context might be captured. Failing the availability of a wireless signal, a mobile handset's GPRS/3G connection might be utilized, however, in cases of uploading large or verbose packets this could be prohibitively expensive for the end user, given that mobile service providers charge (often a high-premium) for large volume use of such a connection. Given these factors, within our architecture we chose not to employ automatic upload. Also in our lifelogging architecture, the mobile phone used to collect context information uploads the data to a central web-enabled server via a wireless internet connection. Often this data is recorded for a protracted period and then uploaded in bulk. Transmitting a month's worth of gathered context data to the repository could take upwards of several hours. During this period of upload no new information is captured and this is obviously problematic as it places unnecessary holes within the collection. As such it is important that we support the data extraction and exchange process in the most ergonomic and streamlined fashion. One means by which this can be achieved is by adopting less verbose structures for data representation thereby making any communications more terse.

Evolving Platforms and Changing Devices

Mobile and desktop platforms continually evolve through operating system or firmware upgrades. Additionally the software on which lifelogging is contingent will also be upgraded iteratively either in response to such evolution of the platform, to address outstanding issues or to advance the functionality provided. Ideally lifelogging solutions would be built upon robust, stable and unchanging platform, which would deliver reliable and consistent performance for the course of its use. Unfortunately, such systems do not exist, and the continual refinement of the software and hardware employed pose a challenge to the goal of lifelogging. Each iteration of an operating system or software release can potentially disrupt important links in the lifelogging chain.

In using mobile devices in the long term, one can expect natural wear and tear to occur and this will result in either down-time for the devices or total replacement. For example, mobile phones are prone to wear and tear, but may also need to be replaced as a result of loss, theft or damage for example as a result of being dropped. The wear-and-tear is not immediately problematic to lifelogging as a device may be easily swapped out for another. It can however create knock-on-effects particularly in terms of gathering contextual information on the environment of the user. In our lifelogging architecture, we make use of Bluetooth 'sniffing' to identify co-present people by associating a Bluetooth hardware ID to a person via an address book lookup (in order to achieve this mapping, the device ID must be known.) However, such mappings are not always straightforward as a person may upgrade their mobile device (introducing a new and unmapped ID) or devices may change hands and be used by multiple people. We have engineered our systems to provide for such complexities by supporting the addition of new mappings and by allowing those mappings to be time-sensitive and rule-based if required (e.g. only mapping a device to user A between May and November, then mapping that device to user B after it changes hands).

Associated Human Factors

While there are obviously issues relating to the use of these mobile devices within a lifelogging scenario, there are also implications for and burdens placed upon the user who must operate, wear and manage lifelogging technology. These challenges are further compounded by the fact that they must negotiate several of these tools in tandem. This section

explores the human factors relating to the use of these technologies and overviews the implications for designing user managed lifelogging solutions.

Forgetfulness

A major drawback to current lifelogging practices is the responsibility placed upon the collection owner to manage, maintain and monitor capture across multiple devices. Unfortunately individuals are prone to forgetting to perform such actions and this '*forgetfulness*' may present in a number of ways. This may include forgetting to wear or carry a device on one's person, forgetting to turn on the device, forgetting to remove lens caps (Byrne et al., 2008b). Forgetting a mobile phone can be particularly problematic, given that the mobile device is running context-logging software when it is left behind the *actual* contextual environment of the individual is not being sensed. By forgetting the device, the user is falsely attributed as being in the environment and setting where they left the device, rather than their actual contextual setting for that period.

Particular to the SenseCam, the wearer will often engage the '*privacy*' mode when using the bathroom or engaged in an activity they would prefer not to be captured, and while this mode expires after 7 minutes, it is preferable to disengage it when leaving and returning to normal activities. This too is an often-neglected action. Incidentally, this forgetfulness may result in unwanted capture too, for example, wearing the SenseCam becomes so habitual, often its wearer may forget they have it on and may accidentally record inappropriate footage of themselves or others around them - for example getting dressed either in the home or in changing rooms – which requires later removal.

Cumbersome and intrusive nature

Requiring an individual to carry and wear multiple devices about their person for extended periods of time is obviously intrusive, especially so in the case of biometric devices. The BodyMedia device (see Fig 2) for example is an armband, which must be worn on the upper arm so as to make contact with the skin. This is inconvenient, as it must often be worn under clothes in order to facilitate recording. Another biometric device, which can cause discomfort, is the heart-rate monitor band that must be strapped around the torso, again making contact with the skin. This device can cause irritation and can be particularly difficult to wear for protracted periods.

While the SenseCam is not uncomfortable to wear owing to it being reasonably lightweight and possessing a small form factor, it impedes physical activities e.g. when running or performing any action with significant continuous movement it is highly impractical, and may intrude upon social activities. Also, from our experiences and from related studies with the SenseCam, it is clear that people become very aware of the presence of the SenseCam in social settings and the social dynamics may change as result (Byrne et al., 2008b). As a result, the users of such technology are very conscious of wearing the technology (McAtamney & Parker, 2006) and may become uncomfortable in certain situations.

Personal & Social Privacy and Ethics of Continuous Capture

Lifelogs are by their nature extremely personal collections containing a wealth of rich fine-grained pictures of the activities an individual has been engaged in over extended periods of time. As such there are many places, spaces and situations in which for personal reasons a user will not wish to capture, for example, bathroom breaks, early mornings while getting dressed etc, or in the evening preparing for bed. The individual may also retrospectively wish to temporarily or permanently '*forget*' particularly painful, upsetting or embarrassing moments contained within such collections (Bannon, 2006). There may in addition to the personal considerations for privacy be socially-mediated constraints placed upon recording and capture. For example, friends or family members may explicitly object to being captured (perhaps to any recording or just in specific situations). Further to both the social and personal limitations on recording, there may additionally be ethical, moral and legal implications imposed. In any situation where an individual (known or unknown) has a reasonable right to

privacy, including hospitals, waiting rooms, schools, changing rooms, bathrooms, etc. the lifelogger should not record any material.

Practicalities in the Use of Resulting Data

The use of multiple mobile sources results in a large volume of multimodal sources which must be harmonized in order to offer utility. The data produced by the various devices is often unstructured and in a variety of formats, making unification of the sources a very complicated matter. Significant additional challenges in terms of efficacy of data processing may be created by the need to temporally align, extract meaning, analyze, transform and append information to the collection to enhance its richness and bring it into confluence – particularly given the size of such collections.

Alignment and Synchronization

Synchronization of data from diverse mobile sources is a time-consuming and tedious process. Synchronization of these data types with each other and also with computer, SenseCam, and mobile phone activity is a difficult process, particularly as the various devices capture data at different rates. This process is further complicated if devices have not been initially time aligned or have lost their time settings as a result of battery drain or hardware/software crash. We found that subjects had several minute disparities between time settings on their various mobile devices, and that these disparities were not consistently maintained. Further complicating this is the fact that while traveling between time zones subjects changed the clocks on some devices, but left the clocks on other devices in the time zone of their home country. Faced with these issues, alignment and synchronization of data from disparate mobile devices becomes a seriously challenging process.

Data Processing and Augmentation

Due to the high frequency of capture, huge volumes of data are collected by these devices even during short periods. In order to add utility to the collection there are three main phases for processing. First the raw data from multiple modalities must be extracted, converted to a usable format, deposited in a central repository and aligned with each other. In addition to this some media often require the *extraction* of low level features to facilitate reasoning and retrieval or in order to make them usable within the collection. For example, the extraction of low-level visual features of an image is extremely useful in order to facilitate later processing including: the relation of images to one another and retrieval (Doherty et al., 2008a), to facilitate clustering or aggregation (O’Connaire et al., 2007; Doherty & Smeaton, 2008), to identify keyframes (Blighe et al., 2008), or in the filtering of low-quality redundant imagery (Doherty et al., 2008b). However, this can often be a computationally expensive and time consuming process, and over an extended period of time can be overwhelming if not efficiently and effectively processed – particularly when dealing with modalities which are continuously recording, such as the SenseCam.

The raw or extracted information, for example sensor data or image features, tends not to be usable until it is transformed into usable representations. This forms the second major phase of processing: *transformation*. To provide an illustrative example, cell tower information and wireless signals associated with mobile phones can provide cues to the location of the individual at any given moment. This can be converted to GPS location using geo-coding lookup services, which can be used to determine a precise named location. By doing so, much richness and value can be added to the collection. Additionally, semantic information can be extracted from SenseCam media frames (Byrne et al. 2008a) which can add extremely useful metadata. However, captured at a rate of every 30 seconds, thousands of data points are available per day and as a result processing this information can take weeks for just a single collection. Additionally, aligning this metadata with the many forms of content in a time efficient manner is a sizeable challenge. Some solutions to this will be discussed later.

Finally, the richness of the collection can be further enhanced through *augmentation* of the available data with information from new external sources. This is particularly useful within a

retrieval scenario (Fuller et al., 2008; Kelly et al., 2008). For example, we employ external weather data-sources to enable retrieval based on remembered conditions such as sunny or rainy, as outlined in (O'Hare et al., 2006). Further to this new media data gathered from online or social media sources may be appended into the collections in order to attempt to provide further richness to the lifelog collection.

IMPLICATIONS AND RECOMMENDATIONS

With a variety of challenges now outlined, in this section the implications for designing and developing lifelog solutions are outlined. Further to this attention is drawn to the use of such data within a research context.

Building Lifelogging Systems

In the previous section we outlined the major challenges presented by operating lifelogging solutions across multiple devices. As a result of our work in this area we have identified six fundamental considerations that anyone undertaking to develop technologies or solutions for lifelogs should be mindful of.

Spend Time Upfront

It is particularly important to spend time upfront in the design of the architecture for a lifelogging system. Careful thought should be given to the components of the architecture and how they will interact. In particular attention should be given to: how the mobile devices employed will exchange information with one another; how many devices will be employed, what the implications of this are – specifically for managing and synchronizing time across them; what form will the central repository take and how should it store its data; the data required for processing; and to whether there is available hardware and/or software support for the architecture and functionality you seek to provide. It is extremely important to consider in detail the central repository and provision for the efficient addition, retrieval, and processing of potentially millions of records. The various operations required to be performed using the data it houses should be evaluated in terms of operation time and computational complexity. The goal is to provide a flexible but supportive platform, and it is extremely important to get this right from the start.

Test Early and Often

A lifelogging solution will have complex interactions and interchanges between many devices. Even with the best of intentions and effort, it is extremely difficult to ensure systems will operate as anticipated. It is also important to remember that these systems are designed to operate over protracted and prolonged periods, so it is important that testing be carried out for a sufficiently realistic period. This will ensure that issues occurring from long-term use and as a result of the strain continuous logging places on mobile devices will be uncovered and can be resolved before larger scale deployment.

Design for Robustness and Reliability

For personal, social or legal reasons recording will be disabled at various points, and as such a lifelog collection will never be fully complete. However, completeness should not be sacrificed as a result of issues pertaining to the capture technologies. When designing the software bear in mind the specific modalities that are being captured and their capture rates since this will have significant implications for power consumption. Try to balance the requirements of richness and longevity of recording through experimentation. Remember that continued high frequency polling, particularly for contextual sources, will put extra strain on the hardware and software. Ideally strive to have highly reliable crash-free software, but if in your scenario of use, stable software is particularly difficult to achieve, then design so that the device/software will gracefully recover from any fault.

Mitigate against Human Factors

Thought should be spared for the individual who will manage and operate these devices – they will be providing you with a detailed insight into their lives and collecting intimate recordings of their lives. It is important then that the benefits to them for conducting such activities are high in comparison to the costs associated with these activities and in particular the burden that the recording solutions place on them. The solutions developed should attempt to minimize the need for intervention by the users whenever and wherever possible. While the temptation is to offload the management entirely to the users, this inclination should be avoided – for example, through automated streamlined solutions to data centralization, graceful recovery from error, scheduled reminders or prompts to overcome forgetfulness, etc. However, realistically automated and graceful recovery from failure is a lofty goal, and there will be times where the users need to take corrective action. In such situations, the devices should make it obvious that intervention is required and the nature of the problem – this may be as simple as playing a custom alert tone or the flashing of a particular indicator. Once the user has been guided to identifying the error, they should be supported by the system in correcting it through helpful instructions, etc. In fact, the solutions should be generally supportive offering well-labeled prompts and generally conforming to good usability conventions. Such issues can be diagnosed through frameworks such as Nielsen’s Heuristics (1994A, 1994B, 2002.)

Collect Data regularly and Check it regularly

With the limited storage capacities of the mobile devices and the large volumes of data that lifelogging generates, it is crucial that the data is collected centrally and removed from these finite resources. This can either be managed manually by offering simple low-tech solutions, for example the provision of regular reminders to those collecting by automated text messages or calendar based reminders, or by automating upload on software launch or at regular and appropriate intervals throughout its use. The solution employed will be dependent on the architecture employed – do the devices predominately require a wired cable based connection to transfer information or are wireless communication protocols available? Additionally, by seeking to collect data regularly issues resulting from data corruption and hardware or software failure can be minimized.

Build in support for privacy

Most lifelogging devices and software provide the opportunity to enable a temporary ‘privacy’ mode where no information is captured for a short period or until user intervention is taken. It is important to provide such functionality in the solutions you build, given the intimate and personal nature of the content being collected, however, it should be remembered the user does not always remember to use this feature. Over the course of the data capture users will inevitably capture data of a personal nature that they would rather not have committed to a lifelong archive. This might include SenseCam images captured in a bathroom, text messages containing sensitive communications or web content with personal access information such as online banking details. Consequently, opportunity to remove such content from the archive through review must be provided.

Due to the personal nature of data capture there is a strong need for maintaining data privacy and providing facilities for users to remove content from their archives. There is also a huge burden on the user to do this with regularity due to the voluminous and ever-increasing nature of such collections. So we should additionally seek to support the user in this process through automatic means, e.g. through the automatic identification of images which may have been accidentally captured in privacy-conscious or sensitive locations, such as in changing rooms or bathrooms, as described by Byrne et al. (2008a).

Using data for research

Beyond the challenges of engineering, developing and employing lifelog solutions which utilise multimodal devices, there are also considerations to be made for qualitative and quantitative research purposes.

Participants: Lifelogs are highly personal collections. As such, it can be difficult to recruit participants to amass such collections over the longer term. Added to this the burden imposed on participants to manage the many recording devices, and also to manage and review their data over a long period of capture, may reduce the willingness or likeliness of potential candidates to participate. As such, and from our own experiences, we found it is best to recruit participants with an vested interest in amassing these collections, i.e. those directly or peripherally involved with the project - Microsoft's Gordon Bell is an excellent example of such an individual (Bell, 2001). Given that these people have a stake in the project they are also more likely to exert additional effort to ensure that their collections are correctly, properly and reliably captured. The practical constraints and overheads of managing any more than a small number of participants within such a collection activity must also be considered. Irrespective of the participants affiliation to the research project, plain language statements and informed consent agreements should be sought upfront. Additionally, given the personal nature of these collections, the individual's willingness to participate may change over time. As such provision to allow an individual to withdraw their participation and data without penalty should be made.

Collection Storage: Ethical consideration around the storage and aggregation of an individual's personal artefacts is required. We have taken great efforts to ensure that all possible data protection steps have been ensured in our work. First and most importantly each participant's lifelog collections should be independent of all others and placed on separate secured server volumes. The only individual with access to create, maintain or review this content should be the collection owner themselves. Only in exceptional circumstances and with prior consent of the collection owner should any of the principal investigators have access to this volume – this might include objection from an external individual to the recording and a request to remove any capture of them from a collection.

Instructions to Participants: All participants should be clearly informed of the situations in which they should not engage in recording practices. This should include any circumstance where an individual has a reasonable right to privacy, for example within a doctor's surgery, hospital, school, etc., and where a given individual has specifically objected to recording taking place either generally or during a specific activity.

Allow for Teething Issues: Given the number of complementary devices and sources in use in a lifelog solution, there are likely to be complications particularly in the early days. Within our studies, we provisioned a one-month period at the beginning of the trial, where many of these issues and complications could be resolved. This also allows participants sufficient time for the recording practices to become habitual. Further to this, it allows those encountered by participants to become accustomed to these novel and perhaps initially socially intrusive devices.

Experimental Design - Developing Test Collections & Mitigating Bias: Experimental work will require experiments to be conducted directly upon any amassed lifelogs. We have found it particularly advantageous to utilise early portions of the collection as development sets for initial testing, e.g. the first month of unencumbered recording. When evaluating lifelog approaches over the longer term it is important to mitigate against collection owner exposure issues, in particular it is often essential to limit exposure to the collection (beyond review and removal of personal sensitive information), as this can lead to bias in any experimental output – especially in any investigation pertaining to recall or remembering. For example, if a user has reviewed a particular instance regularly and in detail, they are more likely to remember details about this instance vividly, thus biasing them towards easily finding this instance in retrieval experiments. It is highly advisable that any supplementary studies e.g. diary studies, observational work, groundtruthing, etc., which may affect the naturalness or ecological validity of a user's day-to-day interactions be conducted during the

development set period. Alternatively, such studies should be delayed until the last period of capture, i.e. final month, as conducting at this point will not impede upon long-term, recall-based investigation, since this content is likely to be too recent to offer utility in such research. There is obvious utility in considering the evaluations and methodologies early and also in the consolidation of potential studies, so as not to intrude on the naturalness of interaction.

Experimental Design - Conducting Groundtruthing & Evaluations: Within all of our experimental evaluations and groundtruthing annotations, no access to the individual collection was available to ensure sound ethical practice. As such all annotation, groundtruthing and experimental investigations were conducted on the individual's computer in isolation, independently. Any information extracted could not contain specific identifiable details to ensure that those examining the groundtruths could not discern any private information. Unique identifiers, timestamps and/or judgements should be the only items visible in any output. Any other required information should be obfuscated by enumerating or anonymizing. This allows usable output to be available to those developing models, algorithms and implementations without requiring the participant to provide personal or sensitive information.

Quantitative Output: While the low number of collections garnered in such studies may create a barrier to quantitative output, it should be born in mind that these collections survey extremely broad periods in rich detail and contain many thousands if not millions of artefacts. While it may be difficult to achieve experimental results that are widely generalizable, the inferences which can be gained are extremely valuable. We would recommend that all quantitative output should be supplemented with qualitative feedback to support the findings making them in many ways more applicable and generalisable. The utility of this approach is highlighted not only in our work but also by others in the domain, for example Elswailer et al. (2007).

Bias & Subjectivity: It is worth providing a cautionary note on the low numbers of users and the use of qualitative measures for evaluation. In all qualitative work there is the potential for subjective feedback and while in other experimental approaches the use of sufficiently large numbers can weed out 'wild card' participants and spurious feedback, it may be extremely difficult to do so given the limited number of individuals likely to amass a lifelog. It is also important to consider when using members who have a direct stake in the success of a research project that they may be inherently subjective and/or biased towards specific outcomes of the evaluation (particularly if privy to the particulars of the experiment). Where possible we would recommend that experimental details be developed in isolation from the participants, thereby removing the potential for such issues. Any use of an investigator's collections should be treated as a special case and acknowledged as such in reporting.

FUTURE RESEARCH DIRECTIONS

Lifelogging is a nascent domain and as such much of the ongoing work is exploratory: designed to tease out the possibilities of such technology while uncovering its limitations in meeting the vision's demands. We have outlined several important considerations for the design of lifelogging solutions however there is one in particular which was unaddressed. As lifelogs are designed to be long-term stores of personal life histories, the question of how we design technology not only to meet the needs of an archive which spans 20, 50 or even 100 years, but also in how we manage the technology which is responsible for its continual collection across such periods. While within our work we have tried to maintain consistent platforms through which capture is facilitated, going forward we must consider how existing lifelogs, and the software facilitating their capture, can be migrated to new hardware platforms and file-systems as the current ones become obsolete or as new state-of-the-art is introduced. Issues relating to defunct file-formats or access to early portions of the collections

may become particularly acute for all involved, especially if operating across the scale of an entire lifetime. Careful thinking about the future perspective is required and such challenges are open questions within the domain.

We additionally need to give thought to how lifelogging technologies behave and respond when we move beyond research in academic contexts into real-world deployments. Given the additional strain continuous capture places upon mobile devices, the need to ensure robustness and reliability and requirements for the streamlined, smooth operation and management of multiple mobile devices in confluence, immense challenges are created.

Given that the technology is likely to change, grow, evolve and adapt as new technology and case scenarios become available the challenges and strategies for managing lifelogs will also have to evolve. There are a variety of ways that lifelogs can be achieved from varying the multiple devices employed, the modal channels collected or the intended purpose or end goals, this too has bearing on the strategies employed to manage such collections, as such no perfect all-encompassing solution is likely to be developed, however we can strive for a solution best tailored to the intended application and domain.

CONCLUSION

This chapter has served to outline the broad challenges and implications of developing lifelog solutions. It is hoped by outlining these it will serve as a practical guide to anyone considering developing such solutions in the future and will give them a platform to fast-track their design, development and engineering of such solutions. It is hoped that by bearing these considerations and challenges in mind, better lifelog solutions can be developed, not only from a software perspective but also in catering for the various human factors involved when creating and managing such personal collections. It should be noted that the challenges and implications addressed here do not provide an exhaustive list of all those encountered in our work, however they should provide an overview of the salient and most problematic.

We would encourage others working in the lifelog domain, to not only consider our perspective on lifelogging, but also advance, outline, address and share their own perspectives so that the community may benefit.

ACKNOWLEDGEMENTS

We would like to thank the Irish Research Council for Science, Engineering and Technology. This work is also supported by Science Foundation Ireland Research Frontiers Programme 2006 and under grant 07/CE/I1147.

REFERENCES

Al Mahmud, A., Aliakseyeu, D., & Martens, J.B. (2008). Enabling Storytelling by Aphasics in an Augmented Home Environment. In D. England (Ed.), *22nd BCS HCI Group Conference* (pp. 3-6). British Computer Society (BCS).

Bannon, L.J. (2006). Forgetting as a feature, not a bug: the duality of memory and implications for ubiquitous computing. *CoDesign*, 2(1), 3-15.

Bell, G. (2001). A Personal Digital Store. *Communications of the ACM*, 44 (1), 86-92.

Berry, E., Kapur, N., Williams, L., Hodges, S., Watson, P., Smyth, G., Srinivasan, J., Smith, R., Wilson, B., and Wood, K. (2007). The use of a wearable camera, SenseCam, as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis. *Special issue of Neuropsychological Rehab.* 17 (4/5), 582-681.

Blighe, M., Doherty, A., Smeaton, A.F., & O'Connor, N. (2008). Keyframe Detection in Visual Lifelogs. In K. Valavanis & I. Maglogiannis (Eds.), *1st International Conference on Pervasive Technologies Related to Assistive Environments*, (Article No. 55). ACM.

Bush, V. (1945). As We May Think. *The Atlantic Monthly*, 176(1), 101-108.

Byrne, D., Doherty, A., Snoek, C.G.M., Jones, G.J.F., & Smeaton, A.F. (2008a). Validating the Detection of Everyday Concepts in Visual Lifelogs. In L. Hardman, A. Hauptmann & D. Duke (Eds.), *3rd International Conference on Semantic and Digital Media Technologies*, (pp. 15-30). Springer-Verlag.

Byrne, D., Doherty, A.R., Smeaton, A.F., Jones, G.J.F., Kumpulainen, S., & Järvelin, K. (2008b). The SenseCam as a Tool for Task Observation. In D. England (Ed.), *22nd BCS HCI Group Conference*, (pp. 19-22). British Computer Society.

Byrne, D., & Jones, G.J.F. (2008). Towards Computational Autobiographical Narratives through Human Digital Memories. In the 2nd ACM Workshop on Story Representation, Mechanism and Context, (pp. 9-12). ACM.

Chen, Y., & Jones, G.J.F. (2009). An Event-Based Interface to Support Personal Lifelog Search. In C. Stephanidis, et al. (Eds.), *13th International Conference on Human-Computer Interaction*. Springer.

Doherty, A., O'Conaire, C., Blighe, M., Smeaton, A.F., & O'Connor, N. (2008a). Combining Image Descriptors to Effectively Retrieve Events from Visual Lifelogs. In A. del Bimbo & E. Bakker (Eds.), *ACM International Conference on Multimedia Information Retrieval*, (pp. 10-17). ACM.

Doherty, A., Byrne, D., Smeaton, A.F., Jones, G.J.F., & Hughes, M. (2008b). Investigating Keyframe Selection Methods in the Novel Domain of Passively Captured Visual Lifelogs. In J. Luo, L. Guan, A. Hanjalic, M.S. Kankanhalli & I. Lee (Eds.), *7th ACM International Conference on Image and Video Retrieval*, (pp. 259-268). ACM.

Doherty, A., & Smeaton, A.F. (2008). Automatically Segmenting Lifelog Data Into Events. In H. Hellwagner, F. Pereira & S. Kollias (Eds.), *9th International Workshop on Image Analysis for Multimedia Interactive Services*, (pp. 20-23). IEEE.

Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., & Robbins, D. (2003). Stuff I've seen: a system for personal information retrieval and re-use. In J. Callan, D. Hawking & A. Smeaton (Ed.), *26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, (pp. 72-79). ACM.

Eagle, N., & Pentland, A. (2006). Reality mining: sensing complex social systems. *Personal Ubiquitous Computing*, 10(4), 255-268.

Elsweiler, D., Ruthven, I., & Jones, C. (2007). Towards Memory Supporting Personal Information Management Tools, *Journal of the American Society for Information Science and Technology*, 58(7), 924-946.

Fuller, M., Kelly, L., & Jones, G.J.F. (2008). Applying Contextual Memory Cues for Retrieval from Personal Information Archives. In J. Teevan & W. Jones (Eds.), *Personal Information Management Workshop, CHI*.

Gemmell, J., Bell, G., Lueder, R., Drucker, S., & Wong, C. (2002). MyLifeBits: Fulfilling the Memex Vision. In L. Rowe, M. Muhlhauser, K. Ross & N. Dimitrova (Eds.), *ACM Multimedia* (pp. 235-238). ACM.

Gurrin, C., Byrne, D., O'Connor, N.E., Jones, G.J.F., & Smeaton, A.F. (2008). Architecture and challenges of maintaining a large-scale, context-aware human digital memory. In E. Izquierdo & G. Liu (Eds.), *5th IET Visual Information Engineering Conference*. Institution of Engineering and Technology.

Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., & Wood, K. (2006). SenseCam: a Retrospective Memory Aid. In P. Dourish & A. Friday (Eds.), *Ubicomp*, (pp. 177-193). Springer-Verlag.

Joki, A., Burke, J., & Estrin, D. (2007). *Campaignr - a framework for participatory data collection on mobile phones*. Technical Report 770, Centre for Embedded Network Sensing, University of California, Los Angeles, October 2007.

Jones, G.J.F., Gurrin, C., Kelly, L., Byrne, D., & Chen, Y. (2008). Information Access Tasks and Evaluation from Personal Lifelogs. In M. Sanderson & T. Sakai (Eds.), *2nd International Workshop on Evaluating Information Access, A Satellite Workshop of NTCIR-7*, (pp. 75-86).

Kostakos, V., O'Neill, E., & Jones, S. (2008). Social networking 2.0. In D. Tan, M. Czerwinski & A. Lund (Eds.), *extended abstracts of the conference on Human factors in computing systems*, (pp. 3555-3560). ACM.

Kelly, L. (2007). The Information Retrieval Challenge of Human Digital Memories. In A. MacFarlane, L. Azzopardi & I. Ounis (Eds.), *BCS IRSG Symposium: Future Directions in Information Access*, (pp. 114-119). British Computer Society.

Kelly, L., Chen, Y., Fuller, M., & Jones, G.J.F. (2008). A Study of Remembered Context for Information Access from Personal Digital Archives. In Pia Borlund et al. (Eds.), *2nd International Symposium on Information Interaction in Context*, (pp. 44-50). British Computer Society.

Lavelle, B., Byrne, D., Gurrin, C., Smeaton, A.F., & Jones, G.J.F. (2007). Bluetooth Familiarity: Methods of Calculation, Applications and Limitations. In G. Broll et al. (Eds.), *Mobile Interaction with the Real World, Workshop at the 9th International Conference on Human Computer Interaction with Mobile Devices and Services*, (pp. 55-58).

Mase, K., Sumi, Y., Toriyama, T., Tsuchikawa, M., Ito, S., Iwasawa, S., Kogure K., & Hagita, N. (2006). Ubiquitous Experience Media. *IEEE Multimedia*, 13(4), 20-29.

McAtamney, G., & Parker, G. (2006). An Examination of the Effects of a Wearable Display on Informal Face-to-Face Communication. In R. Grinter et al. (Eds.), *SIGCHI conference on Human Factors in computing systems*, (pp. 45-54). ACM.

McCarthy, S., Sayers, H., & McKeivitt, P. (2007). Investigating the Usability of PDAs with Ageing Users. In T. Ormerod & C. Sas (Eds.), *21st British Human Computer Interaction Conference*, (pp. 67-70). British Computer Society.

Nielsen, J. (1994). Usability inspection methods. In C. Plaisant (Ed.), *Conference Companion on Human Factors in Computing Systems*. ACM.

Nielsen, J. (1994). Enhancing the explanatory power of usability heuristics. B. Adelson, S. Dumais, & J. Olson (Ed.), *Conference on Human Factors in Computing Systems: Celebrating interdependence*. ACM.

Nielsen, J. (2002). Ten usability heuristics. Retrieved from http://www.useit.com/papers/heuristic/heuristic_list.html, accessed on August 20, 2002.

O'Conaire, C., O'Connor, N., Smeaton, A.F., & Jones, G.J.F. (2007). Organising a Daily Visual Diary Using Multi-Feature Clustering. In *SPIE Electronic Imaging - Multimedia Content Access: Algorithms and Systems*.

O'Hare, N., Lee, H., Cooray, S., Gurrin, C., Jones, G., Malobabic, J., O'Connor, N., Smeaton, A.F., & Uscilowski, B. (2006). MediAssist: Using Content-Based Analysis and Context to Manage Personal Photo Collections. In M.R. Naphade et al. (Ed.), *5th International Conference on Image and Video Retrieval* (pp. 529-532). Springer-Verlag.

Wood, K., Fleck, R., & Williams, L. (2004). Playing with SenseCam. In S. Benford et al. (Eds.), *Playing with Sensors: Exploring the boundaries of sensing for playful ubiquitous computing (W3), UbiComp*.

ADDITIONAL READING SECTION

Abowd, G.D., Dey, A.K., Brown, P. J., Davies, N., Smith, M., & Steggle, P. (1999). Towards a Better Understanding of Context and Context-Awareness. In H. Gellersen, (Ed.), *1st international Symposium on Handheld and Ubiquitous Computing*, (pp. 304-307). Springer-Verlag.

Abowd, G.D., Mynatt, E.D., & Rodden, T. (2002). The human experience [of ubiquitous computing]. *Pervasive Computing*, 1(1), 48-57.

Aizawa, K., Tancharoen, D., Kawasaki, S., & Yamasaki, T. (2004). Efficient Retrieval of Life Log Based on Context and Content. In *the 1st ACM Workshop on Capture, Archival and Retrieval of Personal Experiences*, (pp. 22-31). ACM.

Banavar, G., & Bernstein, A. (2002). Software infrastructure and design challenges for ubiquitous computing applications. *Communications of the ACM*, 45(12), 92-96.

Bellotti, V., & Sellen, A. (1993). Design for privacy in ubiquitous computing environments. In G. de Michelis, C. Simone & K. Schmidt (Eds.), *3rd European Conference on Computer-Supported Cooperative Work ECSCW*. Kluwer Academic Publishers.

Blum, M., Pentland A., & Troster, G. (2006). InSense: Interest-Based Life Logging. *IEEE Multimedia*, 14(4), 40-48.

Brown, P.J., & Jones, G.J.F. (2002). Exploiting contextual change in context-aware retrieval. In *2002 ACM symposium on Applied computing*, (pp. 650-656).

Cherry, S. (2005). Total Recall Life Recording Software. *IEEE Spectrum*, 42, 24-30.

Davies, N., & Gellersen, H. (2002). Beyond Prototypes: Challenges in Deploying Ubiquitous Systems. *IEEE Pervasive Computing*, 1(1), 26-35.

Ellis, D.P.W., & Lee, K. (2006). Accessing Minimal-Impact Personal Audio Archives. *IEEE Multimedia*, 13(4), 30-38.

Gemmell, J., Williams, L., Wood, K., Lueder, R., & Bell, G. (2004). Passive Capture and Ensuing Issues for a Personal Lifetime Store. In *1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences*. ACM.

Gurrin, C., Smeaton, A.F., Byrne, D., O'Hare, N., Jones, G.J.F. & O'Connor, N. (2008). An Examination of a Large Visual Lifelog. In W. Ma, K. Wong, H. Li & J. Nie (Eds.), *4th Asia Information Retrieval Symposium*, (pp. 537-542). Springer-Verlag.

Hori, T., & Aizawa, K. (2003). Context-based Video Retrieval System for the Life-Log Applications. In *the 5th ACM SIGMM International Workshop on Multimedia Information Retrieval*, (pp. 31-38). ACM.

Jones, J., & Marsden, G. (2005). *Mobile Interaction Design*. London, UK: Wiley & Sons.

Jones, W. (2007). *Keeping Found Things Found: The Study and Practice of Personal Information Management*. Burlington, MA, USA: Morgan Kaufmann.

Lahlou, S., Langheinrich, M., & Röcker, C. (2005). Privacy and trust issues with invisible computers. *Communications of the ACM*, 48(3), 59-60.

Langheinrich, M. (2001). Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. In G.D. Abowd, B. Brumitt, & S.A. Shafer (Eds.), *3rd international Conference on Ubiquitous Computing*. Springer-Verlag.

Lee, H., Smeaton, A.F., O'Connor, N., Jones, G.J.F., Blighe, M., Byrne, D., Doherty, A.R., & Gurrin, C. (2008). Constructing a SenseCam Visual Diary as a Media Process. *Multimedia Systems Journal, Special Issue on Canonical Processes of Media Production*, 14(6), 341-349.

Mann, S. (2004). "Sousveillance": inverse surveillance in multimedia imaging. In *12th Annual ACM international Conference on Multimedia*, (pp. 620-627). ACM.

O'Hara, K., Tuffield, M., & Shadbolt, N. (2009). Lifelogging: Privacy and Empowerment with Memories for Life. *Identity in the Information Society*, 1(2).

Satyanarayanan, M. (1996). Fundamental challenges in mobile computing. In *Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, (pp. 1-7). ACM.

Scott, J. (2008). Challenges for Privacy with Ubiquitous Sensor Logging. In *Workshop on Security and Privacy Issues in Mobile Phone Use, Pervasive*.

Tan, D., Sellen, A., Czerwinski, M., Bell, G., Gemmell, J., Wood, K., Robertson, G., Oliver, N., Berry, E., Hodges, S., & Kapur, N. (2007). Save Everything: Supporting human memory with a personal digital lifetime store. In J. Teevan & W. Jones (Eds.), *Personal Information Management: Challenges and Opportunities*. University of Washington Press.

Tancharoen, D., Yamasaki, T., & Aizawa, K. (2005). Practical Experience Recording and Indexing of Life Log Video. In *the 2nd ACM Workshop on Capture, Archival and Retrieval of Personal Experiences*. ACM.

KEY TERMS & DEFINITIONS

Biometric sensors: Devices placed in direct contact with individual's skin, to record skin conductivity which can be used to determine individual's emotional state or arousal levels.

Desktop activity: Items accessed on personal computer, for example word documents, emails, web pages etc.

Explicit photo capture: Photographs explicitly taken by an individual.

Lifelog: Digital archive of personal information, containing for example desktop activity, mobile activity, passively captured photos, explicitly captured photos.

Mobile activity: Activity performed on mobile phone, e.g. SMSs sent and received, phone calls made, received and missed, web browsing.

Mobile context: Individual's current environment captured by mobile devices, e.g. geo-location, motion, light status, weather, people present.

Passive photo capture: Image capture triggered by sensors (e.g. motion, infra red) or a preset time interval.

ⁱ BodyMedia ArmBand: <http://www.bodymedia.com>

ⁱⁱ <http://www.slifelabs.com>